Advances in Group Theory and Applications

15 (2023), pp. 63–97
ISSN: 2499-1287
DOI: 10.32037/agta-2023-003

Introduction to Hyperbolic Groups

Davide Spriano

Abstract

This is an extended version of the lecture notes of the minicourse "Introduction to hyperbolic groups and generalizations" held at GABY 2022.

Mathematics Subject Classification (*2020*): 20F05, 20F65, 20F67, 20F69

Keywords: geometric group theory; hyperbolic group; quasi-isometry; Milnor–Schwarz Lemma; algebraic entropy; Morse geodesics

[*Geometric group theory*] *is about using geometry*
(*i.e. drawing pictures*) *to help us understand groups,*
which can otherwise be fairly dry algebraic objects
(*i.e. a bunch of letters on a piece of paper*)
Ric Wade

1 Geometric group theory

This minicourse is thought to be a first course in geometric group theory. As such, we will assume little knowledge and spend significant time with the basics of geometric group theory, with a particular focus on pictures and examples to build intuition, up to the definition of hyperbolic groups and quasi-isometric invariance of hyperbolicity. After that, we will provide a very incomplete survey on some notable properties of hyperbolic groups. Most of the material is classical, and can be found, for instance, in the excellent books [5, 8] or in the original survey by Gromov [10].

1.1 Preliminaries: Cayley graphs and quasi-isometries

The first question we need to answer is the following: *What is Geometric Group Theory?* The keyword is *group*. Geometric Group Theory is concerned in understanding groups, but using "geometric" techniques. More precisely, given a group G we want to study G via its action on a metric space (X, d). To achieve that, we need a way to cook up a metric space associated to a given group. There are several such constructions in the literature, but one is by far the most important and used one, the *Cayley graph*. We say that a subset X of a group G is *symmetric* if for all $x \in X$ we also have $x^{-1} \in X$.

For us, a graph will be a simplex complex where each edge has length one and is identified with the real interval $[0, 1]$. In particular, a graph is naturally equipped with a metric that turns it into a *geodesic metric space* i.e. a metric space X where for any two points x, y there is an isometric embedding $\gamma\colon [0, b] \to X$, where $b = d_X(x, y)$ such that $\gamma(0) = x$ and $\gamma(b) = y$. When considering geodesics or, in general, maps $p\colon [a, b] \to X$, we will often identify the map with its image.

Definition 1.1 Let G be a group and S a symmetric set of generators for G. The *Cayley graph* of G with respect to S is the simplicial graph $\mathrm{Cay}(G, S)$ with vertex set G and an edge connecting g, h if and only if there is $s \in S$ such that $h = gs$.

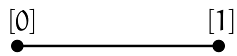

Figure 1: The Cayley graph of $\mathbb{Z}/2\mathbb{Z}$ with respect to the generating set $\{[1]\}$.

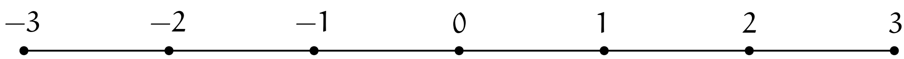

Figure 2: The Cayley graph of $\mathbb{Z}$ with respect to the generating set $\{\pm 1\}$.

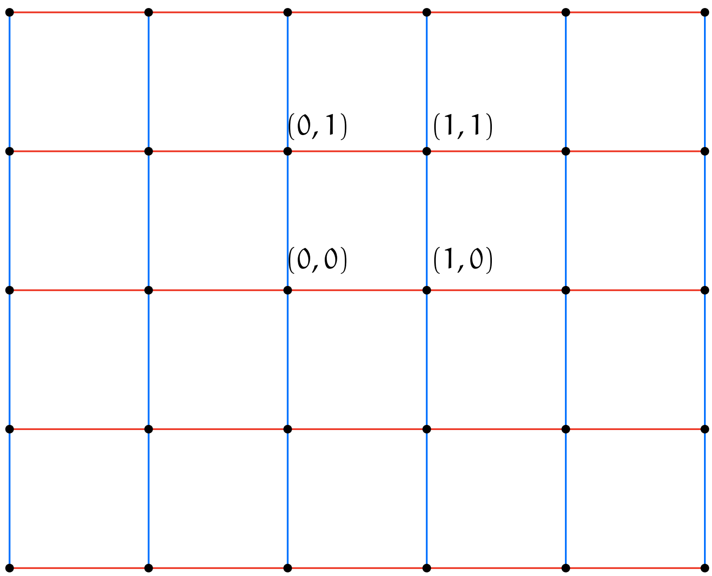

Figure 3: The Cayley graph of $\mathbb{Z} \times \mathbb{Z}$ with respect to the generating set $\{\pm(1,0), \pm(0,1)\}$.

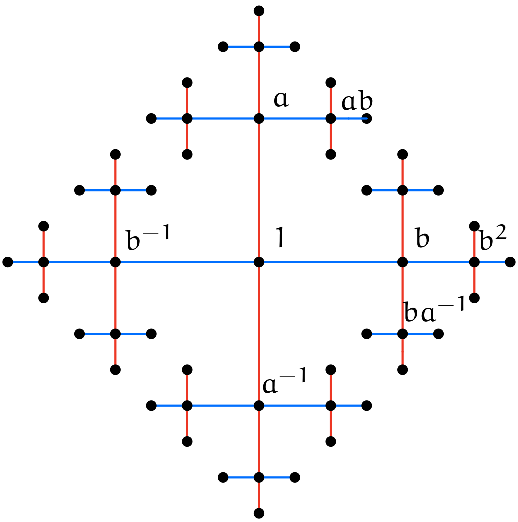

Figure 4: The Cayley graph of the free group on $\{a, b\}$ with respect to the generating set $\{a, b\}$.

Every non-symmetric generating set S can be symmetrized by considering the union $S \cup S^{-1}$, and for our purposes the sets S and $S \cup S^{-1}$ carry the same amount of information. Thus, for a non-symmetric generating set S, we abuse notation and use the shorter $\mathrm{Cay}(G, S)$ to denote the Cayley graph $\mathrm{Cay}(G, S \cup S^{-1})$. Note that if S is finite, then $S \cup S^{-1}$ is also finite.

Observe that G naturally acts on the Cayley graph: given $g, h, k \in G$, the group element g acts on the vertices h, k by $g \cdot h = gh$ and $g \cdot k = gk$. This action is simplicial: if $k = hs$, i.e. there is an edge between h and k, then $gk = ghs$, i.e. there is an edge between $g \cdot h$ and $g \cdot k$.

Note that we can use the Cayley graph to turn G into a metric space. Indeed, define $d_S \colon G \times G \to \mathbb{R}$ as the the distance in $\mathrm{Cay}(G, S)$, i.e. the function that assigns to g, h the length of the shortest path between g and h in $\mathrm{Cay}(G, S)$. Note that d_S depends on the Cayley graph and in particular on the choice of the generating set for G.

We will now establish some basic properties that relate the geometry of the Cayley graph to the algebra of the group G. The first tool that we need is the *label* associated to an oriented edge. If e is an oriented edge from g to h, we define $\mathsf{label}(e) \in S$ to be the unique generator in S such that $h = gs$. Note that $h = gs$ is equivalent to $g = hs^{-1}$. Thus, denoting by $\overline{e}$ the edge with the opposite orientation of e, we have $\mathsf{label}(\overline{e}) = \mathsf{label}(e)^{-1}$. In particular, given a path, we can associate to it a sequence of generators in S using the function label. Considering the opposite direction, observe that the degree of each vertex of $\mathrm{Cay}(G, S)$ is going to be precisely $|S|$, as every vertex g is connected precisely to $gs, s \in S$. In particular, if we fix a vertex g and we consider a sequence of generators $s_1, s_2, \ldots, s_n$, this defines a path in the Cayley graph starting from g and ending to the group element $gs_1s_2 \ldots s_n$. For example, if we start from g and consider the sequence s, s^{-1}, we will associate to it the path that starts at g, crosses the edge labeled by s, and then crosses it back to g. This provides a one to one correspondence

$$\{\text{paths starting at a fixed vertex}\} \xleftrightarrow{\mathsf{label}} \{\text{sequences of generators}\}.$$

There are two important consequences of this fact. The first is that given a path p starting at g, the path ends at h if and only if $h = g \cdot \mathsf{label}(p)$. The second is that the group G acts by isometries on $\mathrm{Cay}(G, S)$. Indeed, consider two vertices h, k and let p be the shortest path between them. Then $d_S(k, h) = \mathrm{length}(p)$ and $h = k \cdot \mathsf{label}(p)$. Since $gh = gk \cdot \mathsf{label}(p)$ it follows that there is a path p' with the same label of p connecting gh and gk. Thus, $d_S(gk, gh) \leqslant d_S(h, k)$. By the symmetric argument we obtain $d_S(h, k) = d_S(gh, gk)$ for all g, h, k, i.e. that G acts by isometries.

Note that the definition of Cayley graph heavily depends on the choice of the generating set. In particular, one can easily see that topological and local properties such as vertex degree and connectivity, are not preserved when we change generating sets. However, large scale properties such as "the graph goes into two opposite directions" are preserved. The notions required to make all of this precise are the ones of *quasi-isometric embedding* and *quasi-isometry*.

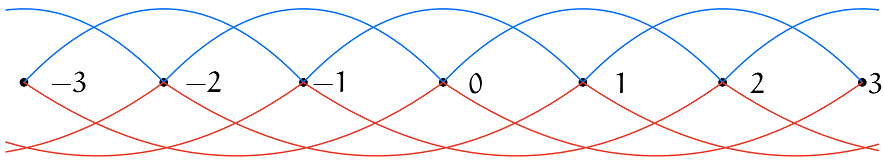

Figure 5: The Cayley grah $\mathrm{Cay}(\mathbb{Z}, \{\pm 2, \pm 3\})$. The fine and topological properties (connectivity, fundamental group, degree...) are very different from $\mathrm{Cay}(\mathbb{Z}, \{\pm 1\})$, but they both "keep going in two opposite directions".

Definition 1.2 Let (X, d_X) and (Y, d_Y) be metric spaces. For constants $\lambda \geqslant 1, \epsilon \geqslant 0$ we say that a map $f\colon X \to Y$ is a (λ, ϵ)–*quasi-isometric embedding* if for any $x, y \in X$ it holds:

$$\frac{d_Y\big(f(x), f(y)\big) - \epsilon}{\lambda} \leqslant d_X(x, y) \leqslant \lambda d_Y\big(f(x), f(y)\big) + \epsilon.$$

We say that a map f is a *quasi-isometric embedding* if it is a (λ, ϵ)–quasi-isometric embedding for some (λ, ϵ).

In the literature it's often presented a slightly different definition of (λ, ϵ)–quasi-isometric embedding compared to the one above, namely requiring that the additive error on the left-hand side is ϵ instead of $\frac{\epsilon}{\lambda}$. The advantage with choice of constants as above is that the equivalent inequalities with the roles of $d_X(x, y)$ and $d_Y\big(f(x), f(y)\big)$ exchanged will have the same constants. This is, ultimately, only a matter of taste, as each condition implies the other up to a small change in the constants.

Example 1.3 In $\mathbb{R}^2$, equipped with the standard Euclidean metric, consider the path connecting $(0,0)$, $(0,n)$, (n,n) and $(n,0)$, parameterized by unit speed. This is a $(3,0)$–quasi-isometric embedding $\gamma\colon [0,3n] \to \mathbb{R}^2$.

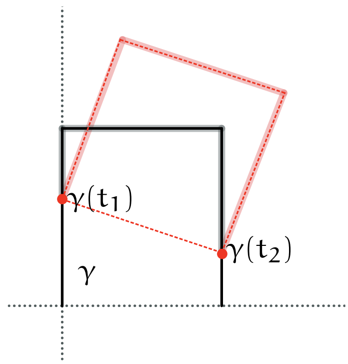

Figure 6: In black, the image of γ in $\mathbb{R}^2$. Basic Euclidean geometry shows that the length of γ between the red points is less than three times the distance between the endpoints, i.e. the length of the red dashed path.

Indeed, since γ is parameterized by arclength we have that for all $t_1 < t_2 \in [0,3n]$ it holds $d_{\mathbb{R}^2}\big(\gamma(t_1), \gamma(t_2)\big) \leqslant t_2 - t_1$. For the other direction, there a couple of cases to consider, the more significant is perhaps when $\gamma(t_1) \in \big[(0,0),(0,n)\big]$ and $\gamma(t_2) \in \big[(n,n),(n,0)\big]$, as in the picture. Some basic Euclidean geometry shows

$$t_2 - t_1 = \ell\big(\gamma|_{[t_1,t_2]}\big) \leqslant 3d\big(\gamma(t_1), \gamma(t_2)\big).$$

Inspired by the above example, we introduce the following definition.

Definition 1.4 A (λ, ϵ)–*quasi-geodesic* of a metric space X is a (λ, ϵ)–quasi-isometric embedding $\gamma\colon I \to X$, where $I \subset \mathbb{R}$ is a (possibly unbounded) closed interval. If I is bounded, we say that γ is a (quasi-geodesic) *segment*, if I has the form $[a, \infty)$ or $(-\infty, a]$, we say that γ is a *ray* and if $I = (-\infty, \infty)$ we say that γ is *bi-infinite*.

Note that quasi-geodesics do not need to be continuous. For this reason, some authors prefer to define them with a discrete domain, i.e. $\gamma\colon I \cap \mathbb{Z} \to X$. There is a clear dictionary between the two definition: given $\gamma\colon I \cap \mathbb{Z} \to X$ define $\gamma'\colon I \to X$ by setting

$$\gamma'\big([n, n+1)\big) = \gamma(n),$$

for $n \in \mathbb{Z}$.

However, quasi-geodesics are, in a sense, almost continuous.

Lemma 1.5 *If $\gamma\colon I \to X$ be a (λ, ϵ)–quasi-geodesic and let x be in the image of γ. Then either x has distance at most $\epsilon + 1$ from one of the endpoints of γ, or, if $x = \gamma(t_0)$ there exist t_1, t_2 with $t_1 < t_0 < t_2$ such that $d\big(\gamma(t_i), x\big) \leqslant \epsilon + 1$ and $\gamma(t_i) \neq \gamma(t)$.*

Proof — Suppose that $x = \gamma(t_0)$ has distance strictly more than $\epsilon+1$ from the endpoints of γ (this is an empty condition if I is unbounded), and define

$$T = \big\{ t \in I \mid \gamma(t) = \gamma(t_0) \big\}.$$

Note that for all $t \in T$ the $\mathbb{R}$–distance between t and the endpoints of I is at least $\frac{(\epsilon+1)-\epsilon}{\lambda} = \frac{1}{\lambda}$. Let $t_1 < t$ be such that $t - t_1 < \frac{1}{\lambda}$. Then

$$d\big(\gamma(t_1), \gamma(t)\big) < \frac{\lambda}{\lambda} + \epsilon = \epsilon + 1.$$

Since T is bounded away from the endpoints of I, there exists a point $t_1 \in I \setminus T$ as above, yielding $\gamma(t_1) \neq \gamma(t_0)$. The same argument works for t_2. $\square$

When dealing with quasi-geodesics, the lack of continuity prevents expressions like "the first $t \in I$ such that ...", essentially because infimi might not be realized. Thus it is often useful to consider the notion of ρ–*first point*. We say that $t \in I$ is a ρ–first point satisfying P if for any other $t' \in I$ satisfying P we have that $t' > t - \rho$.

To obtain a notion of "isometry up to coarse geometry", we need to add some form of surjectivity to quasi-isometric embeddings. Indeed, we saw that it is possible to quasi-isometric embed a line into a plane, and it is straightforward to see that a point is quasi-isometrically embedded in any metric space.

Definition 1.6 Let (X, d_X) and (Y, d_Y) be metric spaces. For constants $\lambda \geqslant 1, \epsilon \geqslant 0$ we say that a map $f\colon X \to Y$ is a (λ, ϵ)–*quasi-*

isometry if it is a (λ, ϵ)–quasi-isometric embedding and $N_\epsilon\big(f(X)\big) = Y$. The last condition is called being *ϵ–quasi-surjective*.

We say that a map is a quasi-isometry if it is a (λ, ϵ)–quasi-isometry for some (λ, ϵ). We say that two metric spaces are quasi-isometric if there is a quasi-isometry between them.

Note that we are using the notion of being quasi-isometric as if it was an equivalence relation. This is not obvious from the definition, but it turns out to be true. If $f\colon X \to Y$ is a (λ, ϵ)–quasi-isometry, we can associate to it a map $\bar{f}\colon Y \to X$ as follows: for each $y \in Y$ choose a closest point $f(x) \in f(X)$, and then define $\bar{f}(y) = x$. Note that there are many choices involved, for instance the choice of a closest point in $f(X)$ and the preimage x. However, quasi-surjectivity ensures that $d\big(y, f(x)\big) \leqslant \epsilon$, and the fact that f is a quasi-isometry ensures $\operatorname{diam} f^{-1}f(x) \leqslant \epsilon$.

Exercise 1.7 *Show that for every (λ, ϵ) there are (λ', ϵ') such that any quasi-inverse of a (λ, ϵ)–quasi-isometry is a (λ', ϵ')–quasi-isometry.*

In particular, this shows that being quasi-isometric is an equivalence relation.

Note that composing a quasi-isometry with its quasi-inverse do not yield the identity map but only, as pointed before, a map

$$\mathrm{QId}\colon X \to X$$

and a constant ϵ satisfying $d\big(x, \mathrm{QId}(x)\big) \leqslant \epsilon$ for all $x \in X$.

Exercise 1.8 *Show that the image of a geodesic under a quasi-isometry is a quasi-geodesic.*

As an example, we show that many infinite valence trees are quasi-isometric to each other. It is true that regular trees of degree at least 3 are quasi-isometric to each other, but for the sake of simplicity we present a proof of a special case.

Lemma 1.9 *Let T_n be the infinite valence tree of degree n. For all $n, m \geqslant 2$, the trees T_{2n} and T_{2m} are quasi-isometric.*

Proof — Our goal is to show that for each $n > 2$ there is a quasi-isometry f from T_{2n} to T_4. Select a vertex $o \in T_{2n}$. Then the result follows from the fact that being quasi-isometric is an equivalence relation. We will inductively define the image of vertices and edges of larger and larger balls around o. We start with the 1–neighbourhood of o as in Figure 7. The idea is that we substitute the 1–neigbourhood of o with a 4–valence with as many leaves as the degree of o. To achieve this, we split the edges incident into o into two groups of 3 and as many groups of 2 as necessary. Then we proceed inductively as in Figure 8.

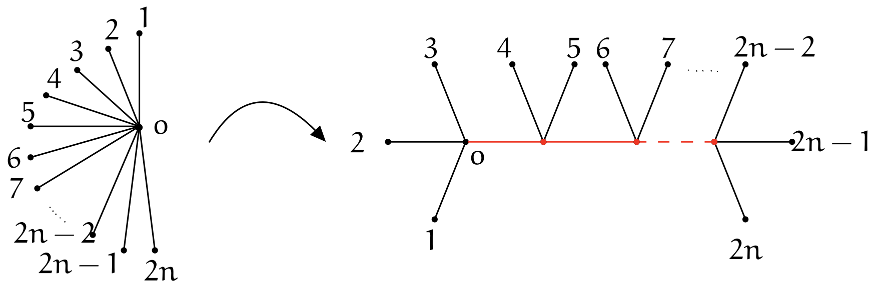

Figure 7: The 1–neighbourhood of o. The vertex o is "blown-up" to a path of length $n-2$, where the degree of o is $2n$.

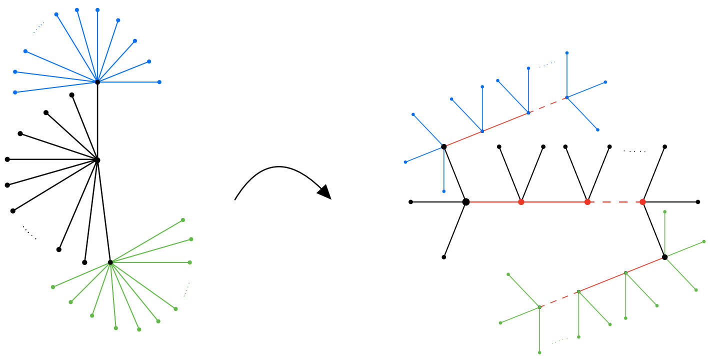

Figure 8: The 2–neighbourhood of o. We proceed inductively to blow-up first all the vertices at distance 1 from o, then 2 and so on.

Note that the map f is injective on vertices and edges, but does not preserve adjacency: in particular, if e_1 and e_2 are adjacent, then $f(e_1)$ and $f(e_2)$ have distance at most $n-2$. Thus, arguing as in Figure 9, we get that f is a $(n-1, n-2)$–quasi-isometry.

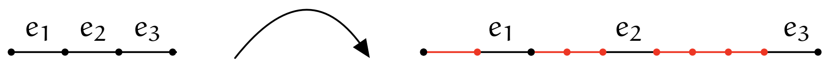

Figure 9: The image of a path. Every edge gets mapped to an edge, and every vertex gets mapped to a path of length at most $n-2$.

The statement is proved. $\square$

We are now ready to show some quasi-isometric invariance for Cayley graphs. Given two sets X, Y the *symmetric difference* between them, denoted $X \Delta Y$ is the set $(X \cup Y) \setminus (X \cap Y)$.

Proposition 1.10 *Let* G *be a group,* S_1 *and* S_2 *be two generating sets such that* $S_1 \Delta S_2$ *is finite. Then the identity* $\mathrm{Id}\colon \mathrm{Cay}(G, S_1) \to \mathrm{Cay}(G, S_2)$ *is a quasi-isometry.*

Proof — Let $T_1 = S_1 \setminus S_2$. By assumption, T_1 is a finite set. Thus, we can define $C_1 = \max\left\{d_{S_2}(1,t) \mid t \in T_1\right\} < \infty$. Then, given vertices x, y of $\mathrm{Cay}(G, S_1)$, let $p = s_1, \ldots, s_n$ be a shortest path between them. Now construct a path in $\mathrm{Cay}(G, S_2)$ between x and y as follows: starting from x, if $s_1 \in S_2$ then connect x and xs_1 by the edge labelled by s_1. Otherwise, choose a shortest path between x and xs_1. Since $d_{S_2}(x, xs_1) = d_{S_2}(1, s_1) \leqslant C_1$, we can proceed this way and obtain a path between x, y of length at most $C_1 d_{S_1}(x,y)$. Thus, $d_{S_2}(x,y) \leqslant C_1 d_{S_1}(x,y)$. By reversing the roles of S_1 and S_2 we obtain the other inequality. $\square$

Finite symmetric difference is a necessary hypothesis. The two Cayley graphs of the group $\mathbb{Z}$ given by $\mathrm{Cay}\big(\mathbb{Z}, \{\pm 1\}\big)$ and $\mathrm{Cay}\big(\mathbb{Z}, \mathbb{Z}\big)$ are not quasi-isometric. The first is a line, as in Figure 2, and the second is a complete graph over a countable set.

The above proposition shows that, up to consider generating sets with finite symmetric difference, the Cayley graph of a group is well defined up to quasi-isometry. However, the expression "up to consider generating sets with finite symmetric difference" is a bit cumbersome. Luckily, there is a very large, extremely natural class of groups that has a best generating set, which is well-defined up to finite symmetric difference, namely *finitely generated groups*. Indeed, Proposition 1.10 can be translated as the following.

Proposition 1.11 *Let* G *be a finitely generated group. Then for any two finite sets of generators* S_1, S_2*, the identity*

$$\mathrm{Id}\colon \mathrm{Cay}(G, S_1) \to \mathrm{Cay}(G, S_2)$$

is a quasi-isometry.

For this reason, in the vast majority of cases geometric group theorists consider Cayley graph of finitely generated groups with respect to *finite* generating sets. Soon, we will also restrict our attention to finite generating sets, but before doing that we need infinite generating sets for a little more, namely to discuss the *Milnor–Schwarz Lemma*, also known as *the fundamental theorem of geometric group theory*.

1.2 The Milnor–Schwarz Lemma

Nomen Omen?

The Milnor–Schwarz Lemma essentially constitutes the base onto which most of geometric group theory is built, as it allows to relate the geometry of a group with an action on a metric space. Thus, it is quite an exceptional coincidence the fact that the second one of its authors is the (unfortunate) example on how composing a quasi-isometry with a quasi-inverse is not the identity map. Indeed, Schwarz is a German-Jewish surname that was transliterated to Russian. When the work of Albert Schwarz was translated to English, the translators didn't realize that his surname had a natural Latin-characters form, and transliterated it to Švarc, showing that two transliterations (quasi-isometries) do not compose to the identity. Albert Schwarz is currently professor at UC Davis, and a mathematical autobiography of his life can be found at `www.math.ucdavis.edu/~schwarz/bion.pdf`.

The Milnor–Schwarz Lemma states that if a group acts "nicely" on a metric space X, then it is quasi-isometric to the space. Before stating the Lemma, we need to recall some definitions.

Definition 1.12 An action of a group G on a metric space X is *cobounded* if there exists a bounded set $B \subseteq X$ such that $G \cdot B = X$.

The above definition can be alternatively formulate as: there exists $R < \infty$ such that for any basepoint o and point $x \in X$ there exists some $g \in G$ such that $d(g \cdot o, x) \leqslant R$.

Definition 1.13 An action of a group G on a metric space X is *metrically proper* (or simply *proper*) if for any ball B, the set

$$\{g \in G \mid g \cdot B \cap B \neq \emptyset\}$$

is finite.

The word "proper" unfortunately means different, but very related things in the literature. For this reason, we add the specification "metrically" to avoid confusion. Indeed, often authors mean the topological definition of proper, i.e. a map $f\colon G \times X \to X$ such that the preimage of every compact set is compact. To make things more complicated, a space is said to be proper if closed balls are compact, and it is true that an action on a proper space is topologically proper if and only if it is metrically proper. The main advantage of the metric perspective is that quasi-isometries send bounded sets to bounded sets, but do not send proper metric spaces to proper metric spaces, arguably making it the correct notion for working in the quasi-isometry setting.

It is often convenient to bundle the notions of the previous definitions as follows.

Definition 1.14 An action of a group G on a metric space X is *geometric* if it is by isometries, metrically proper and cobounded.

The last piece of terminology we need is the following.

Definition 1.15 A metric space X is (λ, ϵ)–*quasi-geodesic* if for any pair $x, y \in X$ there exists a (λ, ϵ)–quasi-geodesic connecting them. We simply say that X is quasi-geodesic if it is (λ, ϵ)–quasi-geodesic for some (λ, ϵ).

Example 1.16 There are several examples of metric space that are not quasi-geodesic. The easiest of these examples is made by the subset $\{2^n \mid n \in \mathbb{Z}\} \subseteq \mathbb{R}$ equipped with the restriction of the metric on $\mathbb{R}$. It is clear that this is a metric space, but it is not quasi-geodesic.

A non-disconnected example is the following: In the hyperbolic plane $\mathbb{H}^2$, let C_n be a circle of radius 2^n and let C be the union $\bigcup_{n=1}^{\infty}$ as in the picture. If we equip C with the restriction of the $\mathbb{H}^2$–metric, we cannot connect the wedging point with any of the opposite points. The rough idea is that those points get too far away to simply "jump" between them, and using the rest of C is also not possible, as in the hyperbolic plane the circumference is exponential in the radius.

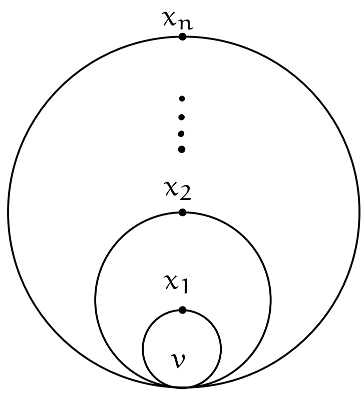

Figure 10: The wedge point v and the opposite points x_n cannot be connected by *uniform* quasi-geodesics.

Often, the Milnor–Schwarz Lemma is formulated as: "if G acts geometrically on a quasi-geodesic space X, then G is finitely generated and quasi-isometric to X". However, to the best of my knowledge, all of the proofs of this fact actually show the following, significantly stronger, result.

Theorem 1.17 (Milnor–Schwarz) *Let* G *be a group acting by isometries and cobundedly on a quasi-geodesic metric space* X*. Then there exists a (possibly infinite) generating set* S *for* G *such that for any point* $o \in X$ *the orbit map*

$$\begin{array}{rcl} \mathrm{Cay}(G, S) & \to & X \\ g & \mapsto & g \cdot o \end{array}$$

is a quasi-isometry.

Moreover, if the action is metrically proper, the set S *is finite.*

Proof — Let (λ, ϵ) be such that X is (λ, ϵ)–quasi-geodesic. Since the action of G is proper, there exists a radius R such that for any ball B_R of radius R we have $X = G \cdot B_R$. We claim that the orbit map induces a uniform–quasi-isometry.

Fix a basepoint o, let B_R be the ball of radius R and center o and B

be the ball of radius $2(R+\epsilon+1)$ with center o. Define

$$S = S_{o,2(R+\epsilon+1)} = \{g \in G \mid B \cap g \cdot B \neq \emptyset\}.$$

Let $h, g \in G$ be any elements. We need to show that there are constants (λ, ϵ) that do not depend on the choice of o such that

$$\frac{d_{\mathrm{Cay}(G,S)}(h,g) - \epsilon}{\lambda} \leqslant d_X(h \cdot o, g \cdot o) \leqslant \lambda d_{\mathrm{Cay}(G,S)}(h,g) + \epsilon.$$

Since G acts by isometries both on $\mathrm{Cay}(G,S)$ and X, the first by definition and the second by hypothesis, we can assume $h = 1$. Indeed,

$$d_{\mathrm{Cay}}(h,g) = d_{\mathrm{Cay}}(1, h^{-1}g),$$

so showing the result for all g suffices, and the analogous argument applies to d_X.

Firstly, let

$$\gamma\colon [0,a] \to X$$

be a (λ,ϵ)–quasi-geodesic between o and $g \cdot o$. Let t_1 be the ρ–first point of $[0,a]$ such that $\gamma(t_1) := x_1$ is outside B, for some arbitrarily small ρ. Since the action of G is cobounded, there exists some $g_1 \in G$ such that $x_1 \in g_1 B_R$. Note that

$$d(o, x_1) > 2(R+\epsilon+1),$$

thus

$$t_1 \geqslant 2\frac{R}{\lambda}.$$

Since quasi-geodesics are coarsely continuous (Lemma 1.5),

$$d(B, x_1) \leqslant \epsilon + 1.$$

Thus, $B \cap g_1 B \neq \emptyset$ and hence $g_1 \in S$. We proceed this way to find $x_2 = \gamma(t_2)$, the ρ–first value of t after t_1 that is outside $g_1 B$, and observe again that

$$t_2 - t_1 \geqslant 2\frac{R}{\lambda}.$$

Iterating, we obtain a chain of at most $\frac{a\lambda}{R}$ pairwise intersecting translates of B such that the first has center o and the last one contains $g \cdot o$.

Up to add one final ball, the last one can be assumed to have center $g \cdot o$. Thus, S generates G. Moreover, we have

$$d_{\mathrm{Cay}(G,S)}(1, g) \leqslant \frac{a\lambda}{R} \leqslant \frac{\lambda(\lambda d_X(o, g \cdot o) + \epsilon)}{R}.$$

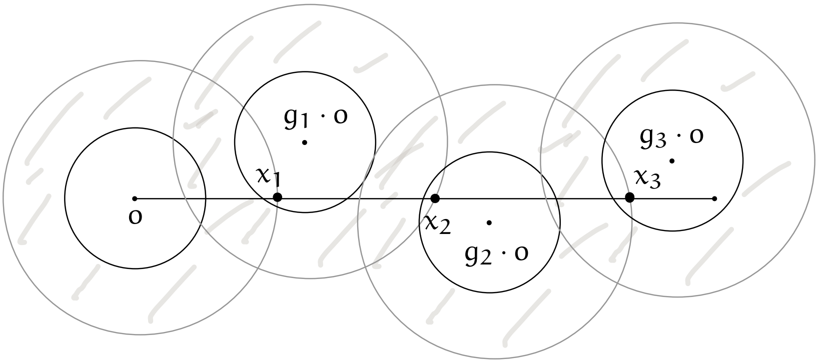

Figure 11: The translates of the ball are chosen as follows: the point x_1 is outside the first *outer* ball, but inside the second *inner* ball. This allows us to guarantee that the sequence x_i is making progress along the quasi-geodesic γ.

For the reverse inequality, let

$$p = s_1, \ldots, s_n$$

be a shortest path in $\mathrm{Cay}(G, S)$ between 1 and g, in particular

$$g = s_1 \ldots s_n.$$

Consider the sequence of balls

$$B, s_1 B, s_1 s_2 B, \ldots, s_1 \ldots s_n B = gB.$$

We claim that any two consecutive ones intersect. Indeed consider

$$s_1 \ldots s_k B \quad \text{and} \quad s_1 \ldots s_{k+1} B.$$

By acting on both by $(s_1 \ldots s_k)^{-1}$ we have that they intersect if and only if B and $s_{k+1}B$ do, which happens by definition of S. As B has radius $2(R+\epsilon+1)$, the triangular inequality yields

$$d_X(o, g \cdot o) \leqslant 4(R+\epsilon+1)n = 4(R+\epsilon+1)d_{Cay(G,S)}(1, g).$$

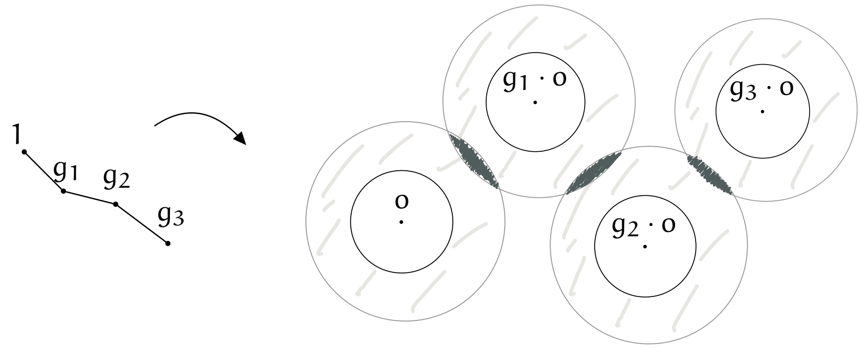

Figure 12: A path in the Cayley graph is naturally associated to an ordered sequence of balls, such that two consecutive balls intersect.

Note that the values we obtained in both estimates do not depend on the choice of the basepoint o, but the definition of $S = S_{o,2(R+\epsilon+1)}$ does. However, it is easily seen that $S_{o,r} \subseteq S_{o',2r}$, so up to enlarge the set S we get a generating set that works for all basepoints.

Finally, observe that the properness of the action implies that all sets S considered are finite. □

As mentioned before, considering cobounded but not proper actions can be very useful. However, for the majority of applications, it is sufficient to restrict to geometric actions of finitely generated groups. Before restricting our attentions to geometric actions only, we provide an incomplete list of examples of natural non proper actions. This is meant to be a zoo for readers with familiarity in geometric group theory, and we will not provide background on the objects.

Example 1.18 The following are examples of cobounded but not-proper action:

1. The action of a group G on its Bass-Serre tree is cobounded, but

proper only in the very special case of splittings of virtually free groups with respect to finite vertex groups.

2. The action of a product $G_1 \times G_2$ on one of the factors G_i by projection is cobounded, and not proper as long as G_{i+1} is infinite. This is the case, for instance of $\mathbb{Z}^2$ acting on $\mathbb{Z}$ by $a^n b^m \cdot k = n + k$. As a matter of fact, no action of $\mathbb{Z}^2$ on $\mathbb{Z}$ can be proper.

3. The action of the Mapping Class Group on the curve graph. It is cobounded by the change of coordinate principle, but the stabilizer of a curve contains the Mapping Class Group of a lower complexity surface, and thus is infinite.

4. The action of a non-hyperbolic, acylindrically hyperbolic group on a hyperbolic space. This includes, for instance, the action of a right-angled Artin group on its contact graph, or a relatively hyperbolic group acting on the coned-off Cayley graph.

From now on, we will consider only finitely generated groups and Cayley graphs of them with respect to finitely generating sets. Thus, when saying that a group G is quasi-isometric to a group H, we mean that one, and hence every, Cayley graph of G with respect to a finite generating set is quasi-isometric to one, and hence every, Cayley graph of H with respect to a finite generating set. In particular, when the specific generating set is not important, we will simply write $\mathrm{Cay}(G)$.

Exercise 1.19 *Show that a (finitely generated) group acts geometrically on its Cayley graph.*

Exercise 1.20 *Let X be a connected, locally simply connected metric space. Then $\pi_1(X)$ acts geometrically on the universal cover $\widetilde{X}$.*

As a corollary, we get some very interesting fact about well-known groups.

Corollary 1.21 *All fundamental groups of closed orientable surfaces are quasi-isometric.*

Proof — If $G = \pi_1(\Sigma)$ and $H = \pi_1(\Sigma')$, we have that G acts geometrically on $\widetilde{\Sigma}$ and H acts geometrically on $\widetilde{\Sigma'}$. By the uniformization theorem of Riemann surfaces, $\widetilde{\Sigma} \cong \widetilde{\Sigma'} \cong \mathbb{H}^2$, where $\cong$ denotes isometries. Hence G and H are quasi-isometric to $\mathbb{H}^2$ and to each other. $\square$

Corollary 1.22 *Any two finitely generated, non-abelian free groups are quasi-isometric.*

Proof — A free group F_n of rank n is the fundamental group of the bouquet of n circles B_n. The group F_n is abelian only if $n = 1, 0$ and hence the group is either trivial or $\mathbb{Z}$. It is finitely generated when $n < \infty$. So, proceeding as in the previous corollary, we want to show that for $1 < n, m < \infty$, the universal covers $\widetilde{B_n}$ and $\widetilde{B_m}$ are quasi-isometric. It is an easy exercise in covering theory to show that $\widetilde{B_n}$ is the infinite regular tree T_{2n} of valence $2n$. Thus, the result follows from Lemma 1.9. □

Corollary 1.23 *Let* H *be a finite-index subgroup of* G*. Then* H *is quasi-isometric to* G*.*

Proof — We need to show that H acts geometrically on $\mathrm{Cay}(G)$. Since G acts properly, by isometries on $\mathrm{Cay}(G)$ and $H < G$, we have that H acts properly and by isometries on $\mathrm{Cay}(G)$. We need to show that H acts coboundedly. If G is finite, this follows trivially, so assume G is infinite. For each coset of H, choose a representative $\{h_i\}$ closest to the identity, and let M be the maximal length of such representatives. Since H is of finite index, there are only finitely many $\{h_i\}$ and thus M is finite. Assume, by contradiction, that H did not act coboundedly, i.e. that there was an infinite sequence of points g_i with $d(g_i, H) \geqslant i$. Then $d(g_{M+1}, H) = d(1, g_{M+1}^{-1}H) > M$, a contradiction. Thus H acts geometrically on $\mathrm{Cay}(G)$ and it is quasi-isometric to G. □

2 Hyperbolicity

Given a finitely generated group G, all of its Cayley graphs are quasi-isometric. In order to make use of this fact, we need some property that is preserved by quasi-isometry and that, hopefully, can be used to draw consequences about the group. The most important notion ticking those boxes is the notion of *hyperbolicity*.

2.1 Definition

If Y is a subset of a metric space X, we denote by $N_r(Y)$ the closed r–neighbourhood of Y, i.e. the set

$$N_r(y) = \{x \in X \mid d(x, y) \leqslant r\}.$$

Definition 2.1 Let $\alpha_0, \alpha_1, \alpha_2$ be the sides of a geodesic triangle T. We say that T is δ–*thin* if $\alpha_i \subseteq N_\delta(\alpha_{i-1} \cup \alpha_{i+1})$, where the indices are taken mod (3).

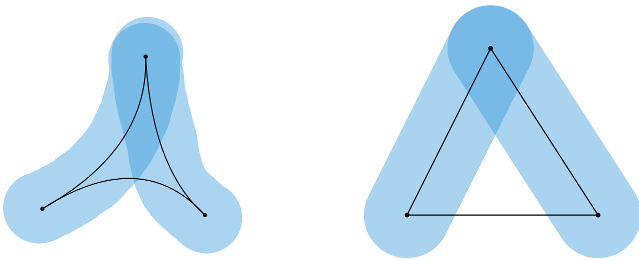

Figure 13: Thin and not thin triangles

Definition 2.2 A geodesic metric space is δ–*hyperbolic* if there exists $\delta \geqslant 0$ such that any geodesic triangle of X is δ–thin. We say that a metric space is hyperbolic if it is δ–hyperbolic for some δ.

Hyperbolic metric spaces have very surprising properties. For instance, there are several chain of properties $P_1 \Rightarrow P_2 \Rightarrow \ldots \Rightarrow P_n$ such that all the opposite implications are false for general metric spaces, but they turn out to hold in hyperbolic spaces. We will see an example of this later, but the theme is that hyperbolic spaces are rigid: as soon as the weaker property P_n is satisfied, the stronger P_1 follows.

We start with a list of elementary examples and non-examples.

Example 2.3 *The following spaces are hyperbolic.*

1. *Any bounded space (by taking δ bigger than the diameter).*
2. *Any tree is 0–hyperbolic.*
3. *The hyperbolic plane $\mathbb{H}^2$ and, in general, the hyperbolic space $\mathbb{H}^n$.*

The following spaces are not hyperbolic.

1. *The Euclidean space $\mathbb{R}^n$ for $n \geqslant 2$, as for any δ there exists a non δ–thin triangle;*
2. *In general, any space that contains a quasi-isometrically embedded copy of $\mathbb{R}^n$ for $n \geqslant 2$, for the same reason.*

Another very notable counterexample are the Cayley graphs of the Baumslag–Solitar groups $BS(n, m)$. This is not hard to see, but requires knowing that every cyclic subgroup of a hyperbolic group is quasi-isometrically embedded (Theorem 3.8).

2.2 The Morse lemma

Our final goal is to show that being hyperbolic is a quasi-isometric invariant property, namely that if X and Y are quasi-isometric geodesic metric spaces and X is hyperbolic, so it is Y. Since all graphs, and in particular Cayley graphs, are geodesic metric spaces, if a Cayley graph is hyperbolic, so any Cayley graph is. So, the notion of *hyperbolic group* would be well-defined. To reach that goal, we need a series of foundational facts on the geometry of metric and hyperbolic spaces, culminating in what is often called *Morse lemma*.

We start with a very useful lemma that allows us to ignore some pathological behaviour of quasi-geodesics.

Definition 2.4 Let X be a metric space and $p\colon [a, b] \to X$. The *length* of p, denoted by $\ell(p)$, is defined by

$$\sup_{a=t_0 \leqslant t_1 \leqslant \ldots \leqslant t_n = b} \left\{ \sum_{i+0}^{n-1} d(\gamma(t_i)\gamma(t_{i+1})) \right\}.$$

If $\ell(p) < \infty$ we say that p is *rectifiable*.

Definition 2.5 Given subsets Y, Z of a metric space X, the *Hausdorff distance* between Y, Z is

$$d_H(Y, Z) = \inf\left\{\epsilon \geqslant 0 \mid Y \subseteq N_\epsilon(Z),\ \ Z \subseteq N_\epsilon(Y)\right\}.$$

Lemma 2.6 (Taming quasi-geodesics) *Let X be a geodesic metric space and let $\gamma\colon [a, b] \to X$ be a (λ, ϵ)–quasi-geodesic of X. There there exists a $\big(\lambda, 2(\lambda+\epsilon)\big)$–quasi-geodesic $\gamma'\colon [a, b] \to X$ such that:*

1. *γ' is continuous;*

2. *γ is rectifiable and for all $t, t' \in [a, b]$ it holds*

$$\ell\big(\gamma'|_{[t,t']}\big) \leqslant k_1 \, d\big(\gamma(t), \gamma(t')\big) + k_2,$$

where k_1, k_2 depend only on λ, ϵ;

3. $\gamma(a) = \gamma'(a)$ *and* $\gamma(b) = \gamma'(b)$;

4. *the Hausdorff distance between (the images of)* γ *and* γ' *is at most* $\lambda + \epsilon$.

PROOF — We will provide a sketch of the proof. For more details, we refer the reader to [5, Chapter III.H, Lemma 1.11].

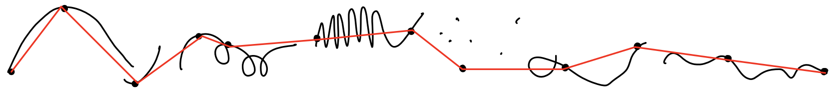

Figure 14: The black quasi-geodesic γ is replaced by the red piecewise-geodesic γ'. The latter is obtained by sampling points of γ at regular intervals.

Let
$$\Sigma = \{a, b\} \cup \big([a, b] \cap \mathbb{Z}\big).$$
For each $s \in \Sigma$, define $\gamma'(s) = \gamma(s)$. Since X is geodetic, given consecutive $s_1, s_2 \in \Sigma$ choose a geodesic segment connecting them, and extend γ' by linear interpolation, so that the image of γ' is the concatenation of such segments. Note that the length of each such segment is at most $\lambda + \epsilon$, and, since γ is a (λ, ϵ)–quasi-geodesic, every point of (the images of) $\gamma \cup \gamma'$ lies in the $\lambda + \epsilon$–neighbourhood of $\gamma(\Sigma)$. In particular, we verified items 1,3,4. It remains to prove that γ' is a quasi-geodesic, and the estimates on the length, that we leave as exercise. $\square$

The main result of this section is the Morse lemma, that in turn will allow us to prove quasi-isometric invariance of hyperbolicity. Before venturing in the proof, let us state it and present some examples and counterexamples.

Definition 2.7 A quasi-geodesic γ of a metric space X is said to be *Morse* (or to have the *Morse property*) if for any λ, ϵ there exists an N such that for any (λ, ϵ)–quasi-geodesic segment η such with endpoints $x_1, x_2 \in \gamma$, we have
$$d_H\big(\eta, \gamma|_{[t_1, t_2]}\big) \leqslant N,$$
where t_1, t_2 are any values such that $\gamma(t_i) = x_i$.

Since the definition of Morse quasi-geodesic includes an assignment $(\lambda, \epsilon) \to N$, it is sometimes useful to bundle this information in a map $M\colon \mathbb{R}_{\geqslant 1} \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}$, and say that a quasi-geodesic is M*–Morse* if the assignment in Definition 2.7 is given by M. The Morse lemma states the following.

Theorem 2.8 (Morse lemma) *Let* X *be a hyperbolic space. Then there exists* $M\colon \mathbb{R}_{\geqslant 1} \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}$ *such that all the geodesics of* X *are* M*–Morse.*

Let us introduce a useful notation.

Definition 2.9 Given a quasi-geodesics $\gamma\colon I \to X$ and $\eta\colon [a, b] \to X$ with endpoints on γ, the *restriction of* γ *to* η, denoted as $\gamma|_\eta$, is defined as follows. Let

$$J_1 \subseteq I = \{t \in I \mid \alpha(t) = \eta(a)\} \quad \text{and} \quad J_2 = \{t \in I \mid \alpha(t) = \eta(b)\}.$$

For

$$J = \bigcup_{t_1 \in J_1, t_2 \in J_2} [t_1, \infty) \cap (-\infty, t_2],$$

define $\alpha|_\eta$ as $\alpha|_J$.

Example 2.10 *Every geodesic in a tree is Morse.*

Proof — Indeed, let γ be a geodesic and η be a (λ, ϵ)–quasi-geodesic with endpoints on γ. By Lemma 2.6, we can assume that η is continuous, as we can substitute η with some continuous η', prove the result for η' and then extend it to η, since $d_H(\eta, \eta')$ is uniformly bounded.

For simplicity, simply denote by γ the restriction $\gamma|_\eta$. If the image of η, as a set, coincides with image of γ, we are done. Otherwise, we can assume that there is some $\eta(t)$ at maximal distance from γ. Since η is continuous and the endpoints of η are on γ, there exists a maximal $t_1 < t$ such that $\eta(t) \in \gamma$ and a minimal $t_2 > t$ with $\eta(t_2) \in \gamma$. Since we are in a tree, $\eta(t_1) = \eta(t_2)$, and since η is a (λ, ϵ)–quasi-geodesic we need to have $t_2 - t_1 \leqslant \epsilon$. Thus,

$$d\big(\gamma, \eta(t)\big) = d\big(\eta(t_i), \eta(t)\big) \leqslant \lambda|t - t_i| + \epsilon \leqslant \frac{\lambda\epsilon}{2} + \epsilon,$$

where we are using the easily verifiable fact that $\eta(t_1) = \eta(t_2)$ realize the distance $d\big(\gamma, \eta(t)\big)$.

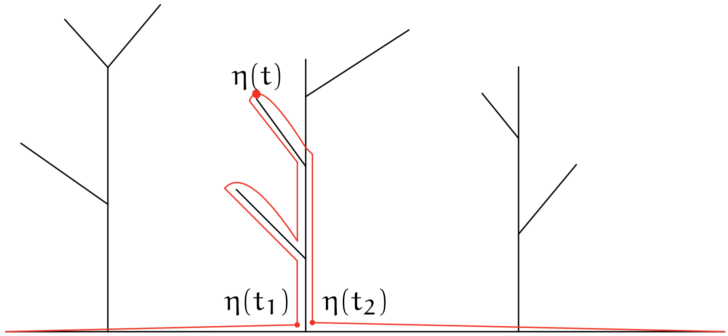

Figure 15: Geodesics of a tree have the Morse property

We are done. □

Example 2.11 *On the opposite side of the spectrum, no infinite geodesic in $\mathbb{R}^2$ has the Morse property. This shows that having the Morse property is not a generic property, and it makes sense to have a name for the concept.*

Proof — Let γ in $\mathbb{R}^2$ be the geodesic corresponding to the x–axis. For $n > 0$, define η_n to be the map $\eta_n \colon [0, 3n] \to \mathbb{R}^2$ joining the sequence $(0,0), (0,n), (n,n), (n,0)$. By Example 1.3, each of the η_n is a $(3,0)$–quasi-geodesic of $\mathbb{R}^2$, and each η_n escapes the $(n-1)$–neighbourhood of $\gamma|_{\eta_n}$. Thus, there is not a uniform neighbourhood containing all of them and γ does not have the Morse property.

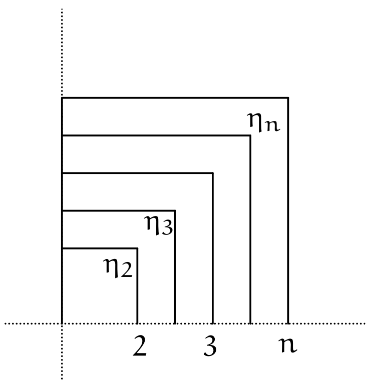

Figure 16: The family η_n.

We are done. □

The next proposition essentially constitutes the first half of the proof of the Morse lemma.

Proposition 2.12 *Let X be a hyperbolic space, γ a geodesic and p a continuous rectifiable path with endpoints on γ. Then the restriction $\gamma|_p$ is contained in the $\left(\delta\lceil\log_2(\ell(p))\rceil + 1\right)$–neighbourhood of p.*

Proof — Up to substitute γ with $\gamma|_p$ we can assume that the endpoints of γ and p coincide (note that we are using that γ is a geodesic to guarantee that the interval J of Definition 2.9 is closed). Let a, b be the endpoints of γ and $x_0 \in \gamma$ be any point. Our goal is to uniformly bound $d(x_0, p)$. If $d(x_0, \{a, b\}) \leqslant \delta$, we are done. Otherwise, let m be the midpoint of p, i.e. the point realizing

$$\ell\left(p|_{[a,m]}\right) = \ell\left(p|_{[m,b]}\right).$$

Then x_0 is on one side of a geodesic triangle with vertices a, b, m. Note that all the vertices of the triangle belong to p. Thus, there exists a point x_0 either on a geodesic from a to m or from m to b satisfying $d(x_0, x_1) \leqslant \delta$.

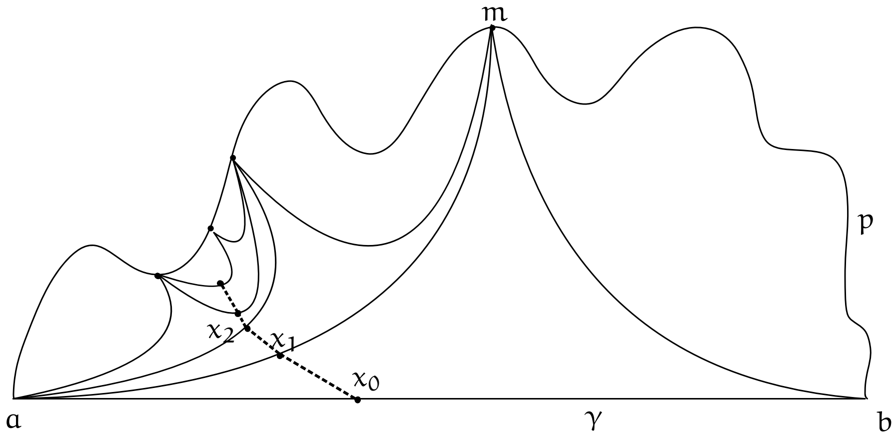

Figure 17: Iterating the x_i.

Now, observe that we are exactly in the same situation as before, but the length of p has halved. Each iteration reduces the length of p by half, and produces a new point x_i at distance at most δ from x_{i-1}. Hence, after at most $\lceil \log_2(\ell(p)) \rceil$ steps, we find x_n at distance at most δ from on the of vertices of the triangle obtained in the n–th iteration, and hence from p.

We conclude

$$d(x,p) \leqslant \delta \lceil \log_2(\ell(p)) \rceil + \delta.$$

As x was generic, the proof is concluded. □

We are now ready to complete the proof of Theorem 2.8.

Proof of Theorem 2.8 — Let γ be a geodesic and η a (λ,ϵ)–quasi-geodesic with endpoints on γ. Up to substitute γ with $\gamma|_\eta$, we can assume that they have the same endpoints, and using Lemma 2.6 we can assume that η is continuous and rectifiable. Let $x \in \gamma$ be a point such that $d(y,\eta) \leqslant d(x,\eta)$ for all $y \in \gamma$. Such a point exists by compactness of γ and η. Let $\xi = d(x,\eta)$. Let A, B be the endpoints of γ and $a, b \in \gamma$ be such that

$$d(a,x) = d(b,x) = 2\xi \quad \text{and} \quad d(a,A) < d(b,A).$$

We will assume that such points exist. The proof in the case in which $d\big(x, \{A,B\}\big) \leqslant \xi$ will follow the same blueprint, up so some small tweaks.

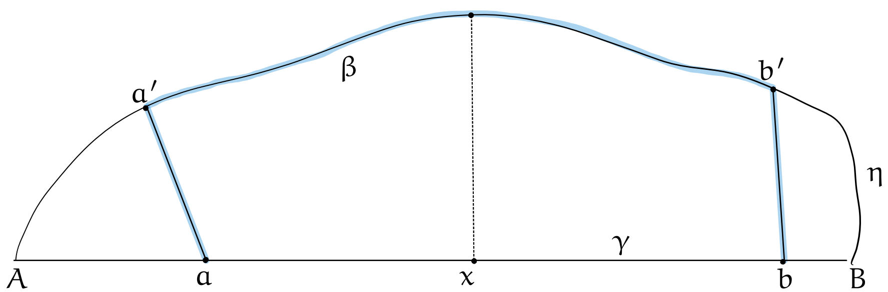

Figure 18: The path β.

Let $a' \in \eta$ be a point closest to a, and define similarly b' for b. Since $d(x,\eta)$ is maximal, we have $d(a,a') \leqslant \xi$ and $d(b,b') \leqslant \xi$. By the triangular inequality, we have

$$d(a', b') \leqslant d(a', a) + d(a, b) + d(b, b') \leqslant 6\xi,$$

and hence, by item 2 of Lemma 2.6, we have that there are constants k_1, k_2 depending only on λ, ϵ such that

$$\ell\big(\eta|_{[a',b']}\big) \leqslant k_1(6\xi) + k_2.$$

Choose geodesics segments $[a, a']$ and $[b, b']$ between a, a' and b, b' respectively, and let β be the path between a, b obtained by concatenating $[a, a']$, $\eta|_{[a',b']}$ and $[b', b]$. If is easily seen that

$$\ell(\beta) \leqslant k_1(6\xi) + k_2 + 2\xi = O(\xi).$$

But Proposition 2.12 yields

$$\xi \leqslant O\big(\log_2\left(\ell(\beta)\right)\big) \leqslant O\big(\log_2(\xi)\big).$$

Thus, there is a uniform bound on ξ. □

Corollary 2.13 *Let X and Y be quasi-isometric geodesic metric spaces. Then X is hyperbolic if and only if Y is, possibly with different hyperbolicity constants.*

Proof — Let $f\colon X \to Y$ be a quasi-isometry, and assume that Y is hyperbolic. The case where X is hyperbolic is completely analogous, using a quasi-inverse of f. Let T be a geodesic triangle of X with vertices x_0, x_1, x_2, and let α_i be the side of T not containing x_i. Put $y_i = f(x_i)$. Since f is a quasi-isometry, $f(\alpha_i)$ is a quasi-geodesic of Y with endpoints y_{i+1}, y_{i-1} (indices mod 3), where the quasi-geodesic constants of $f(\alpha_i)$ depend only on f. Choose geodesics $[y_i, y_j]$ of Y between y_i, y_j. Since Y is hyperbolic, $f(\alpha_i)$ is at uniformly bounded Hausdorff distance from $[y_{i+1}, y_{i-1}]$. Since the triangle with vertices $\{y_i\}$ is uniformly thin, we get that $f(\alpha_i)$ is contained in a uniform neighbourhood of $[y_{i-1}, y_i] \cup [y_i, y_{i+1}]$, which in turn is contained in a uniform neighbourhood of $f(\alpha_{i+1}) \cup f(\alpha_{i-1})$. Thus, there is some uniform constant M such that for each $p \in f(\alpha_i)$ there is

$$q \in f(\alpha_{i+1}) \cup f(\alpha_{i-1})$$

with $d_Y(p, q) \leqslant M$. Since f is a quasi-isometry and p is generic, this yields that triangles in X are uniformly thin.

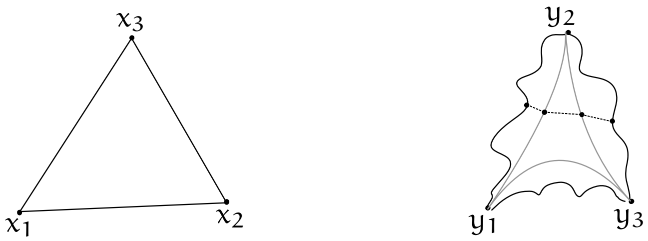

Figure 19: Hyperbolicity of Y forces triangles in X to be thin.

The statement is proved. □

In particular, this allows us to finally define hyperbolic groups.

Definition 2.14 A finitely generated group G is *hyperbolic* if $\mathrm{Cay}(G)$ is hyperbolic.

Using the Milnor–Schwarz Lemma, we get that a group is hyperbolic precisely when it acts geometrically on a hyperbolic metric space. This allows to upgrade Example 2.3 to statements about groups.

Example 2.15 *The following groups are hyperbolic.*

1. *Any finite group.*

2. *The fundamental groups* $\pi_1(\Sigma)$, *where* Σ *is a closed (not necessarily orientable) surface, as it acts geometrically on the hyperbolic plane* $\mathbb{H}^2$.

3. *Free groups, as they act geometrically on trees.*

3 Survey on hyperbolicity

The goal of this section is to provide a, very incomplete, survey on remarkable known results about hyperbolic groups.

By definition, hyperbolic groups are finitely generated. It is not very hard to show that they are finitely presented. If we consider a presentation, we obtain a 2–dimensional complex by gluing 2–cells that correspond to the relation, called the *Cayley complex*. Given a Cayley complex of a finitely presented group, there is a natural notion of area, thus for a closed loop in the 1–skeleton, one can consider the map

$$\mathrm{fill}(p) = \min\{\mathrm{Area}(D) \mid D \text{ disk}, \ \partial D = p\}.$$

Interestingly, determining how hard is to fill a loop in a Cayely complex gives us important information about the group. Namely, let $D \colon \mathbb{N} \to \mathbb{N}$ be the function

$$D(n) = \max\{\mathrm{fill}(p) \mid \ell(p) = n\},$$

called the *Dehn* function. Although the Dehn function depends on the choice of presentation, its asymptotic does not, so it makes sense to talk about linear, quadratic, etc. Dehn function. For the precise definition and more on the topic, we refer to [5, Chapter III.H]. A very important property of hyperbolic groups is the following. The proof idea for the forward direction is essentially due to Dehn himself (see [7]).

Theorem 3.1 (see [10]) *A group is hyperbolic if and only if it has linear Dehn function.*

This highlights an important property of hyperbolic groups. They admit several, apparently completely unrelated, definitions. We will elaborate on this later, and recall several other important facts of hyperbolic groups. We start with the fact that they are not simple in a very strong sense. Recall that a group is *virtually cyclic* if it does contain a cyclic subgroup of finite index. Note that the trivial group is cyclic and hence finite groups are virtually cyclic. Non-virtually cyclic groups are also called *non-elementary*.

Theorem 3.2 (see [10]) *Let G be a non-elementary hyperbolic group. Then there exists an infinite normal subgroup H such that G/H is infinite.*

Following the thread of subgroups of hyperbolic groups, we have the following result, often called the Tits alternative.

Theorem 3.3 (Tits alternative, [10]) *Let G be a hyperbolic group and $H < G$ a subgroup. Then either H is virtually cyclic, or H contains the free group F_2.*

An important consequence of the Tits alternative is that balls in hyperbolic groups grow exponentially. More precisely, let S be a generating set for G, and define the *growth function* as

$$\mathrm{growth}_S(n) = |B(n)|,$$

where $|B(n)|$ represents the number of vertices in the ball of radius n of $\mathrm{Cay}(G, S)$. Clearly, the growth function depends on the choice of generating set, as revealed by a quick count on the ball of radius 1 in Figures 2 and 5. However, as for the Dehn function, the asymptotic of the growth do not depend on a generating set.

Exercise 3.4 *The free group F_n, $n \geqslant 2$ has exponential growth. Moreover, if a group G satisfies $F_n < G$ then G has exponential growth as well.*

As a corollary, we obtain the following result.

Corollary 3.5 *Let G be a non-elementary hyperbolic group. Then G has exponential growth.*

Considering the asymptotic behaviour of the growth function can feel a bit unsatisfying, and it is natural to wonder whether we can have some better estimates on the growth rate of hyperbolic groups, and as before we want those results to not depend on the choice of a generating set.

Definition 3.6 A group G has *uniform exponential growth* if there exists $\lambda > 0$ such that for any finite generating set S of G we have:

$$e^{\lambda n} \leqslant \mathrm{growth}_S(n).$$

It is tempting to assume that groups with exponential growth should, in fact, have uniform exponential growth. However, there are examples where this is not the case, namely of groups of exponential, but not uniform exponential growth. The first such example was constructed in [15], and additional counterexamples can also be found in [2, 12, 14]. It is worth noting that all such examples are infinitely presented and, to the best of my knowledge, it is still open whether finitely presented groups of exponential growth have uniform exponential growth.

In particular, we expect that hyperbolic groups have uniform exponential growth, which is confirmed by a classical result of Koubi.

Theorem 3.7 (see [11]) *Let G be a non-elementary hyperbolic group. Then G has uniform exponential growth.*

Consider an infinite order element $g \in G$, where G is a non-elementary hyperbolic group. Since balls are finite, for any radius r there must be some n such that g^n does not belong to the ball B_r centered at the identity. However, since G has exponential growth, the size of B_r is (asymptotically) exponential in r, and we could in principle need to consider an exponential amount of powers $\{g, g^2, \dots, g^n\}$ to achieve $d(1, g^n) \geqslant r$. This turns out to not be the case, as the next theorem shows.

Theorem 3.8 *Let g be an infinite order element of a hyperbolic group G. Then the map*

$$\begin{array}{ccc} \mathbb{Z} & \to & \mathrm{Cay}(G, S) \\ n & \mapsto & g^n \end{array}$$

is a quasi-isometry.

In other words, the cyclic subgroup $\langle g \rangle$ is quasi-isometrically embedded in G. Given a finitely generated subgroup H of a finitely generated subgroup G, authors in the literature often say that H is *undistorted* in G if the inclusion is a quasi-isometric embedding. Thus the previous theorem states that all infinite cyclic subgroups of a hyperbolic group are undistorted. We remark that by considering a finitely generated subgroup we are secretly fixing a metric on H and a map $H \to G$ given by the inclusion. If we were to consider general metric spaces $Y \subseteq X$, one would need to equip Y with a meaningful metric to talk about the fact that Y is undistorted.

Being an undistorted subgroup is, a priori, not an excessively strong condition. It turns out that in hyperbolic groups this is equivalent to much stronger notions. The first that we will consider is *quasiconvexity*.

Definition 3.9 A subset Y of the geodesic metric space X is K–*quasiconvex* if for any pair of points $x, y \in Y$, any geodesic connecting them is contained in the K–neighbourhood of X.

The notion of quasiconvexity relies on geodesics, and thus it is not preserved by quasi-isometries, as the following exercise shows.

Exercise 3.10 *Show that being quasiconvex is not preserved by quasi-isometry.* (Hint: the x–axis in $\mathbb{R}^2$)

The natural way to address this is to consider all quasi-geodesics at the same time. In order to deal with neighbourhoods depending on the quasi-geodesic constants, we borrow the formalism and name from the case of Morse geodesics.

Definition 3.11 Let $M\colon \mathbb{R}_{\geqslant 1} \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}$ be a map. A subset Y of a metric space X has the M–*Morse property* if for any (λ, ϵ)–quasi-geodesic η with endpoints on Y, we have that η is contained in the $M(\lambda, \epsilon)$–neighbourhood of Y.

The definition of the Morse property is more involved than the one of quasiconvexity, but it has a coarse geometric meaning in non-hyperbolic spaces that the first notion lacks. For more on the Morse property, we refer the reader to [13], where the author extensively studies the Morse property in metric spaces, and he calls it *strong quasiconvexity*.

Given a set Z of a metric space X and a point $x \in X$ we define the closest point projection of x to Z to be the set

$$\pi_Z(x) = \{z \in Z \mid d(z, x) = d(Z, x)\}.$$

Note that $\pi_Z(x)$ might be empty. This can be fixed by allowing an error of an arbitrarily small ϵ in the defining equation. This is a standard trick in coarse geometry. So, let us assume that $\pi_Z(x)$ is always non-empty.

Definition 3.12 Let $D > 0$ be a constant. We say that a subset Z of a metric space X is D–*strongly contracting* if for any metric ball B disjoint from Z the set $\pi_Z(B)$ has diameter at most D.

For the relations between various notion of contraction and the Morse property we refer the reader to [1].

We can now relate all of those notions together. In general, they are all distinct and form a chain of implications. For hyperbolic groups, this is not the case and all the notions are equivalent.

Theorem 3.13 *Let G be a finitely generated group and H a finitely generated subgroup. Then the following chain of implications holds for H.*

$$\textit{Strongly contracting} \Rightarrow \textit{Morse} \Rightarrow \textit{quasiconvex}^{*} \Rightarrow \textit{undistorted}.$$

If G is hyperbolic, then undistorted implies strongly contracting.

Definition 3.14 A subgroup of a hyperbolic group is *quasiconvex* if it satisfies the conditions of Theorem 3.13.

[*] Being quasiconvex is not preserved by quasi-isometry. Hence, we are assuming that a generating set is fixed. The other notions are preserved by quasi-isometry with the same implications.

Quasiconvex subgroups play a special role in the theory of hyperbolic groups. The first consequence, that can be proven as exercise, is the following.

Proposition 3.15 *Let* H *be a quasiconvex subgroup of a hyperbolic group. Then* H *is hyperbolic.*

Quasiconvex subgroups of hyperbolic groups fail to be normal in a very strong sense. Let H be a subgroup of a group G. Following [9], we say that two conjugates $g_1Hg_1^{-1}$ and $g_2Hg_2^{-1}$ are *essentially distinct* if $g_1H \neq g_2H$. The *width* of H in G is the maximal size of a collection of pairwise essentially distinct conjugates

$$\{g_1Hg_1^{-1}, \dots, g_nHg_n^{-1}\}$$

such that any two conjugates have infinite intersection. Note that if H is normal and of infinite index, then H has infinite width. If H has width 1 we say that H is *malnormal*. Note that those notions can be extended to a family $\{H_i\}$ of subgroups.

Theorem 3.16 (see [9]) *Let* H *be a quasiconvex subgroup of a hyperbolic group* G. *Then* H *has finite width in* G.

There are several generalizations of hyperbolic groups. A notable one, already introduced by Gromov in its original paper, is the one of *relatively hyperbolic group*. For a full definition, we refer the reader to [8]. Intuitively, a group is hyperbolic relative to a collection $\mathcal{P}$ of subgroups if the non-hyperbolicity is relegated to the elements of $\mathcal{P}$ and the elements of $\mathcal{P}$ are *geometrically separated*, meaning that they don't have "parallel" areas. For our purposes, what matters is that, in general, one needs a good understanding of a group to determine a set of peripheral subgroups, but for hyperbolic groups there is an excellent characterization due to Bowditch.

Theorem 3.17 (see [4]) *Let* G *be a hyperbolic group and* $\{H_i\}$ *a family of quasiconvex subgroup. Then* G *is hyperbolic relatively to the family* $\{H_i\}$ *if and only if the family* $\{H_i\}$ *is malnormal.*

3.1 Definitions of hyperbolicity

A truly remarkable fact about hyperbolic groups, which deserves its own section, is the fact that they admit several, apparently completely unrelated, definitions. We will survey some of the more remarkable ones.

We start with an interesting local-to-global principle proven by Gromov.

Theorem 3.18 (see [10]) *A geodesic metric space is hyperbolic if and only if for each pair of quasi-geodesic constants (λ, ϵ), there exist constants L, λ', ϵ' such that every path that L–locally is a (λ, ϵ)–quasi-geodesic, is globally a (λ', ϵ')–quasi-geodesic.*

A key step in the proof is the characterization of Theorem 3.1. We stated it for groups only, as the definition of Dehn function for metric spaces is more involved, but this should be thought as a problem in formalism and not in substance. The proof carries over to general metric spaces.

We saw that a key property of hyperbolic spaces is that every geodesic satisfies the Morse property. Cordes proved that this is, in fact, a characterization.

Theorem 3.19 (see [6]) *A geodesic metric space X is hyperbolic if and only if there exists $M\colon \mathbb{R}_{\geqslant 1} \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}$ such that all geodesics of X are M–Morse.*

Moving to a more dynamical perspective, to every hyperbolic group we can associate a notion of boundary, which is a compact, perfect, metrizable topological space and the group acts on the boundary by homeomorphisms, with north-south dynamics and with a *discrete converge action*. For the latter, given a set X denote by $\Theta(X)$ the triple of X minus the diagonal, that is to say, the set

$$\{(x, y, z) \in X \times X \times X \mid |\{x, y, z\}| = 3\}.$$

Definition 3.20 Let G be a group acting by homeomorphisms on a compact metrizable space X with at least two points. We say that the action is a *discrete convergence action* if the induced action of G on $\Theta(X)$ is properly discontinuous and cocompact.

Theorem 3.21 (see [3]) *Let G be non-elementary. Then G is hyperbolic if and only if G admits a discrete convergence action on a compact metrizable space X with no isolated points and at least two points.*

The boundary is not the only “space at infinity” that can be associated to a hyperbolic group. Other very well studied objects are the *asymptotic cones* of a group. The definition of asymptotic cone is rather involved and requires a choice of a ultrafilter. A single asymptotic cone can tell a lot about a space, but it is often useful to consider all of them at once.

Theorem 3.22 (see [10]) *A geodesic metric space X is hyperbolic if and only if all of its asymptotic cones are $\mathbb{R}$–trees.*

REFERENCES

[1] G.N. Arzhantseva – C.H. Cashen – D. Gruber – D. Hume: "Characterizations of Morse quasi-geodesics via superlinear divergence and sublinear contraction", *Doc. Math.* 22 (2017) 1193–1224.

[2] L. Bartholdi: "A Wilson group of non-uniformly exponential growth", *C.R. Math. Acad. Sci. Paris* 336 (2003), 549–554.

[3] B.H. Bowditch: "A topological characterisation of hyperbolic groups", *J. Amer. Math. Soc.* 11 (1998), 643–667.

[4] B.H. Bowditch: "Relatively hyperbolic groups", *Internat. J. Algebra Comput.* 22 (2012), 1250016, 66pp.

[5] M.R. Bridson – A. Haefliger: "Metric Spaces of Non-Positive Curvature", *Springer*, Berlin (1999).

[6] M. Cordes: "Morse boundaries of proper geodesic metric spaces", *Groups Geom. Dyn.* 11 (2017), 1281–1306.

[7] M.W. Dehn: "Transformationen der Kurven auf zweiseitigen Flaechen", *Math. Ann.* 72 (1912), 413–421.

[8] C. Druţu – M. Kapovich: "Geometric Group Theory", *American Mathematical Society*, Providence, RI (2018).

[9] R. Gitik – M. Mitra – E. Rips – M. Sageev: "Widths of subgroups", *Trans. Amer. Math. Soc.* 350 (1998), 321–329.

[10] M. Gromov: "Hyperbolic groups", in Essays in Group Theory, *Springer*, New York (1987), 75–263.

[11] M. Koubi: "Croissance uniforme dans les groupes hyperboliques", *Ann. Inst. Fourier (Grenoble)* 48 (1998), 1441–1453.

[12] V. Nekrashevych: "A group of non-uniform exponential growth locally isomorphic to IMG($z^2 + i$)", *Trans. Amer. Math. Soc.* 362 (2010), 389–398.

[13] H.C. Tran: "On strongly quasiconvex subgroups", *Geom. Topol.* 23 (2019), 1173–1235.

[14] J.S. Wilson: "Further groups that do not have uniformly exponential growth", *J. Algebra* 279 (2004), 292–301.

[15] J.S. Wilson: "On exponential growth and uniformly exponential growth for groups", *Invent. Math.* 155 (2004), 287–303.

Davide Spriano
Department of Mathematics
University of Oxford
Andrew Wiles Building
OX2 6GG Woodstock Rd (Oxford)
e-mail: spriano@maths.ox.ac.uk