Conjugacy of finite subsets in hyperbolic groups

Martin R. Bridson* and James Howie†

Abstract

There is a quadratic-time algorithm that determines conjugacy between finite subsets in any torsion-free hyperbolic group. Moreover, in any k-generator, δ-hyperbolic group Γ, if two finite subsets A and B are conjugate, then $x^{-1}Ax = B$ for some $x \in \Gamma$ with $\|x\|$ less than a linear function of $\max\{\|\gamma\| : \gamma \in A \cup B\}$. (The coefficients of this linear function depend only on k and δ.) These results have implications for group-based cryptography and the geometry of homotopies in negatively curved spaces.

In an appendix, we give examples of finitely presented groups in which the conjugacy problem for elements is soluble but the conjugacy problem for finite lists is not.

1 Introduction

Many of the central ideas in modern geometric group theory flow from Gromov's theory of hyperbolic groups [9], which encapsulates a remarkable correspondence between the geometry of negatively curved manifolds and the complexity of the basic decision problems in group theory. The class of hyperbolic groups is much more extensive and diverse than the class of fundamental groups of closed negatively curved manifolds. Nevertheless, when such fundamental groups enjoy a property as a consequence of the convexity of the metric on the manifold, one expects that a suitable distillation of the geometry should allow one to establish a similar property for all hyperbolic groups.

For example, closed geodesics in compact negatively curved spaces provide canonical representatives for the free homotopy classes of loops in the space, and a canonical curve shortening process provides an efficient homotopy from an arbitrary rectifiable loop to its geodesic; correspondingly, conjugacy classes in hyperbolic groups contain a small number of shortest representatives, and there is an algorithm that quickly reduces an arbitrary word in the generators to such a representative, hence solving the conjugacy problem.

More precisely, if Γ is a δ-hyperbolic group with generating set S, then there is a subquadratic-time algorithm that decides if a given pair of words in the letters $S^{\pm 1}$ represent conjugate elements of Γ. Moreover, if $a, b \in \Gamma$ are conjugate, then there is an element $x \in \Gamma$ such that $x^{-1}ax = b$ and

$$\|x\| \leq \|a\| + \|b\| + c_0,$$

*Department of Mathematics, Imperial College London, London SW7 2AZ; m.bridson@imperial.ac.uk
†Department of Mathematics, Heriot-Watt University, Edinburgh EH14 4AS; J.Howie@ma.hw.ac.uk

where $\|\gamma\| = d(1, \gamma)$ is measured in the word metric corresponding to S, and c_0 is a constant depending only on δ and the cardinality of S (see Proposition 2.3.)

But not all decision problems in hyperbolic groups are soluble: the generalised word problem is insoluble in general, as is the generation problem (wherein one must decide which finite subsets of a given cardinality generate the group); see [2]. As a special case of this last fact one sees that in general there does not exist an algorithm to decide which finitely generated subgroups of a hyperbolic group are conjugate.

Our goal in this article is to prove that conjugacy for finite subsets in hyperbolic groups follows the paradigm set by single elements rather than that set by finitely generated subgroups. (This corresponds to the fact that in a negatively curved space X homotopies between graphs mapped into X behave much like homotopies between loops in X; see §8.)

Theorem A *Let Γ be a group that is δ-hyperbolic with respect to a finite generating set of cardinality k. Then there exist constants α and β (depending only on δ and k) with the following property: if two finite lists $A = [a_1, \ldots, a_m]$, $B = [b_1, \ldots, b_m]$ of elements are conjugate in Γ, then there exists $x \in \Gamma$ such that $x^{-1} a_i x = b_i$ for $i = 1, \ldots, m$ and*

$$\|x\| \leq \alpha\mu + \beta,$$

where $\mu = \max\{\|c\| \colon c \in A \cup B\}$.

This theorem leads immediately to an algorithm that determines conjugacy between finite lists of elements in Γ, and therefore finite subsets — i.e. it solves the Whitehead problem for inner automorphisms of arbitrary word-hyperbolic groups. (In fact, the existence of a solution to this problem, but not the linear bound on the length of the conjugating element, is implicit in the work of Gersten and Short [8]; cf. [14, 16].)

A naïve implementation of the algorithm provided by Theorem A requires exponential time but a careful refinement of it yields a *quadratic time algorithm* in the torsion-free case, which is important from the point of view of group based crytosystems [1]. In the general case, we prove the following result. Here, $\|A\|$ and $\|B\|$ denote the sums of the lengths of the words in the given lists, while $\mu_{A,B}$ denotes the maximum length of a word in $A \cup B$.

Theorem B *Assume Γ is δ-hyperbolic with respect to a generating set of cardinality k. Then there is a constant $C = C(\delta, k)$ and an algorithm that, given two finite lists of elements $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_m]$ of Γ (as words in the generators), will either*

1. *terminate after at most $Cm\mu_{A,B}^2$ steps having determined whether the lists are conjugate, outputting a conjugating element if it exists; or else,*
2. *terminate after at most $C(\|A\| + \|B\|)$ steps with the conclusion that all of the elements listed have finite order.*

The rather limp conclusion of case 2 reflects the difficulties that large centralisers pose when one is investigating conjugacy: in a hyperbolic group, the centraliser of any

element of infinite order is virtually cyclic, but the centralisers of torsion elements can be large. Related difficulties concerning torsion underlie most of the technical problems that arise in the proof of Theorem A; indeed the theorem admits a rather short proof in the torsion-free case.

Some aspects of the linear bound on $\|x\|$ in Theorem A are worthy of comment. For example, the leading coefficient α turns out to be rather large, whereas if we restrict attention to the case where A and B are single elements then, as indicated above, we may take $\alpha = 2$.

On the other hand, it is worth noting that our bound on $\|x\|$ is independent of the cardinality of A and B.

In proving Theorem A we have made little attempt to minimize the value of α and β; indeed we have repeatedly traded sharpness for economy of expression. Nevertheless, we always give explicit bounds so as to make it clear that the constants we obtain are computable functions of δ and k (see Section 6, in particular).

In the context of decision problems, it is worth noting that there exist finitely presented groups in which the conjugacy problem for elements is soluble, but the conjugacy problem for finite lists is not. Explicit examples of such groups are constructed in the Appendix to this paper.

Finally, a few comments are in order concerning the import of our results with regard to *group-based cryptosystems*. In the past few years there has been considerable interest in the idea of basing public-key encryption systems on suitably chosen algebraic analogues of the discrete logarithm problem; a seminal paper in this regard is [1]. An idea that has attracted particular interest is that of basing such a system on the (generalised) conjugacy search problem in a suitable group. Ideally, one would like a group in which there is a rapid solution to the word problem, but in which there does not exist a sub-exponential time algorithm that, given the information that two lists of elements in the group are conjugate, will find a conjugating element (cf. [1] and [18]). Since hyperbolic groups can be characterised as those finitely presented groups that admit a particularly rapid and practical solution to the word problem (Dehn's algorithm), it would be of profound interest if one could construct a hyperbolic group in which the conjugacy search problem for lists was algorithmically complex. Our results show that, in fact, such groups do not exist.

This paper is organised as follows. In §2 we gather the basic facts about hyperbolic groups that we require for the proof of Theorem A. The mechanics of the proof are divided into three separate cases and hence three sections, §§3-5. (One anticipates that this division should be made according to whether the subgroups generated by A and B are finite, virtually cyclic, or non-elementary, but in practice the demarcation is more subtle.) We complete the proof of Theorem A in §6. In §7 we prove a number of algorithmic results. In particular we prove Theorem B by explicitly describing the promised algorithm. As mentioned above, much of the technical difficulty in our proofs arises from the need to allow for torsion in Γ. If one assumes from the outset that Γ is torsion-free, then one can achieve significant improvements in the speed of the algorithm. A streamlined algorithm for the torsion-free case is also described in §7. Finally, in §8, we describe an application of Theorem A to the width of homotopies between maps of metric graphs into negatively curved spaces. The estimate that we

obtain is closely related to an inequality of T. Kappeler, S. Kuksin and V. Schroeder [12, Theorem 0.1] (see also [13, Theorem 5.1]). It was these results that inspired us to consider the problems solved in this paper. In this context, the second author is grateful to Sergei Kuksin for a number of helpful discussions.

2 Preliminaries

Let δ be a positive real number. We say that a metric space X is *δ-hyperbolic* if it is *geodesic* (any two points are joined by a geodesic arc) and any geodesic triangle in X is *δ-slim* (each side is contained in the closed δ-neighbourhood of the union of the other two sides).

This definition agrees with that in [3] and in [10]. One should be aware that other books and articles use different definitions, but all are equivalent to ours modulo a calculable scaling of the parameter δ (see [3, Chapter III.H], for example).

One immediate consequence of our definition is the following.

Lemma 2.1 *Any geodesic quadrilateral in a δ-hyperbolic metric space is 2δ-slim (i.e., each side is contained in the closed 2δ-neighbourhood of the union of the other three sides).*

A group Γ is said to be *δ-hyperbolic* with respect to a finite generating set S if the corresponding Cayley graph $X = X(\Gamma, S)$ is δ-hyperbolic as a metric space. Here the metric on X is the length metric that makes each edge isometric to the real unit interval; its restriction to the vertex set Γ is the word metric, i.e., $d(g, h)$ is the length of the shortest word in the letters $S^{\pm 1} := S \sqcup S^{-1}$ representing $g^{-1}h$.

For the remainder of the paper, we shall regard Γ and S as being fixed. We denote by k the cardinality of the finite set S, and fix a positive integer δ such that $X(\Gamma, S)$ is δ-hyperbolic.

The *boundary* of Γ is the set $\partial\Gamma$ of equivalence classes of geodesic rays $\rho : [0, \infty) \to X$, where two rays ρ, ρ' are defined to be equivalent if $d(\rho(t), \rho'(t))$ is bounded for $t \in [0, \infty)$.

Fix a linear ordering on the (monoid) generating set $S^{\pm 1}$ of Γ. This induces a lexicographical ordering on words in the letters $S^{\pm 1}$ that have a fixed length n, and hence also on geodesics with fixed endpoints.

Following Delzant [5] we say that a (directed) geodesic from $g \in \Gamma$ to $h \in \Gamma$ is *special* if it is the (unique) least geodesic from g to h with respect to this lexicographical ordering. An infinite geodesic is *special* if each of its finite segments is special.

The following summarises the basic facts about special geodesics which we will require in the sequel. Details may be found in [3, page 466] or [5, pages 678-679].

Proposition 2.2 *Let $a \in \Gamma$ be an element of infinite order.*

1. *There exists an infinite special geodesic γ joining two (distinct) points $a^{-\infty}$ and a^{∞} of $\partial\Gamma$, such that the powers of a are all a bounded distance from γ.*
2. *The pair $a^{\pm\infty} \in \partial\Gamma$ satisfying condition 1 is uniquely determined by a.*

3. *The number of special geodesics joining $a^{-\infty}$ to a^{∞} is bounded above by a constant R, depending only on δ and k;*
4. *a permutes this finite set and $a^{R!}$ acts on each special geodesic by a translation in the direction of a^{∞}.*
5. *There is a constant K depending only on δ and k such that, for any special geodesic σ joining $a^{-\infty}$ to a^{∞}, the centraliser $C_\Gamma(a)$ of a in Γ is contained in the closed $(1+\|a\|)K$-neighbourhood of σ.*

The constant R in Proposition 2.2 may be taken to be the volume of the ball of radius 2δ about $1 \in \Gamma$. In particular, $R \leq 1 + 2k\sum_{i=1}^{2\delta-1}(2k-1)^i \leq (2k+1)^{2\delta}$.

Proposition 2.3 *There is a constant c_0, depending only on δ and $k = |S|$, such that if a and b are conjugate in Γ, then there exists $x \in \Gamma$ with $x^{-1}ax = b$ and*

$$\|x\| \leq \|a\| + \|b\| + c_0.$$

(It suffices to let $c_0 = (2k+1)^{4\delta} + 4\delta$.)

Proof. Choose x so that $ax = xb$ and $n := \|x\|$ is minimal. Let $x(i)$, $i = 0, \ldots, n$ be the vertices of a geodesic $\hat{x}$ from 1 to x in the Cayley graph of Γ, and consider a geodesic quadrilateral with vertices $1 = x(0)$, $x = x(n)$, a and $ax = xb$, where $\hat{x}$ and $a\hat{x}$ are two of the sides.

Each vertex of $\hat{x}$ is within a distance 2δ of a vertex on one of the other sides. If

$$2\delta + \|a\| < i < n - 2\delta - \|b\|,$$

then this vertex must be on the side $a\hat{x}$; denote it $ax(j(i))$. Moreover, we must have $|j(i) - i| \leq 2\delta$. (For example, if $j(i) - i > 2\delta$ then we would have

$$\|ax\| \leq i + d(\hat{x}(i), a\hat{x}(j(i))) + n - j(i) < n = \|x\|,$$

which would contradict the minimality of $\|x\|$, since ax conjugates a to b.) Thus

$$d(\hat{x}(i), a\hat{x}(i)) \leq d(\hat{x}(i), a\hat{x}(j(i))) + d(a\hat{x}(j(i)), a\hat{x}(i)) \leq 4\delta.$$

For different values of i in the displayed range, the elements $\hat{x}(i)^{-1}a\hat{x}(i)$ must be distinct: if $\hat{x}(i)^{-1}a\hat{x}(i) = \hat{x}(i')^{-1}a\hat{x}(i')$ with $i < i'$, then $x' = \hat{x}(i)\hat{x}(i')^{-1}x \in \Gamma$ satisfies $ax' = x'b$ and $\|x'\| < \|x\|$.

It follows that $(n - 2\delta - \|b\|) - 2\delta - \|a\|$ is less than the number of elements in the ball of radius 4δ about $1 \in \Gamma$; call this V. Then

$$\|x\| = n \leq \|a\| + \|b\| + V + 4\delta,$$

as required. □

3 Conjugate lists of torsion elements

In this section we study conjugacies between finite lists of torsion elements in our fixed δ-hyperbolic group Γ.

Lemma 3.1 *Suppose that $a \in \Gamma$ is an element of finite order N. Suppose that $x_1, x_2, b \in \Gamma$ are such that $ax_1x_2 = x_1x_2b$, $\|x_1\| \geq \|a\| + 2\delta$, $\|x_2\| \geq \|b\| + 2\delta$, and*

$$\|x_1x_2\| = \|x_1\| + \|x_2\| \geq (N+3)(\|a\| + 2\delta) + \|b\| + 2\delta.$$

Then

$$\|x_1^{-1}ax_1\| \leq 8\delta.$$

Proof. We assume that $\|x_1^{-1}ax_1\| > 8\delta$ and derive a contradiction. Write

$$x = x_1x_2 = g_1g_2\cdots g_{\|x\|},$$

where $g_i \in S^{\pm 1}$ for each i, and $x_1 = g_1\cdots g_{\|x_1\|}$. This decomposition determines a geodesic $\hat{x}$ joining 1 to x in the Cayley graph $X(\Gamma, S)$; at integer times $t = 0, \ldots, \|x\|$ this geodesic visits the vertices $\hat{x}(t) = g_1\cdots g_t$. Since quadrilaterals in Γ are 2δ-slim, there exists, for each integer $t \in I := \{\|a\| + 2\delta, \ldots, \|x\| - \|b\| - 2\delta\}$ at least one integer $s(t)$ such that $d(\hat{x}(t), a\hat{x}(s(t))) \leq 2\delta$. Note that

$$\begin{aligned} |t - s(t)| &= |d(1, \hat{x}(t)) - d(a, a\hat{x}(s(t)))| \\ &\leq d(1, a) + d(\hat{x}(t), a\hat{x}(s(t))) \\ &\leq \|a\| + 2\delta, \end{aligned} \tag{3.1}$$

by the triangle inequality. On the other hand, again by the triangle inequality,

$$\begin{aligned} |\ \|x_1\| - s(\|x_1\|)\ | &= d(a\hat{x}(\|x_1\|), a\hat{x}(s(\|x_1\|))) \\ &\geq d(a\hat{x}(\|x_1\|), \hat{x}(\|x_1\|)) - d(\hat{x}(\|x_1\|), a\hat{x}(s(\|x_1\|))) \\ &> 8\delta - 2\delta = 6\delta, \end{aligned} \tag{3.2}$$

where the assumption $\|x_1^{-1}ax_1\| > 8\delta$ has been used to bound the first term on the second line, and the definition of $s(\|x_1\|)$ used to bound the second.

Let us assume that $s(\|x_1\|) > \|x_1\|$. The proof in the other case is entirely analogous. Note that, whenever $t \in I$ is such that $(t - \|x_1\|)(s(t) - s(\|x_1\|)) \geq 0$, yet another application of the triangle inequality gives

$$\begin{aligned} |(s(\|x_1\|) - \|x_1\|) - (s(t) - t)| &= |(t - \|x_1\|) - (s(t) - s(\|x_1\|))| \\ &= |d(\hat{x}(t), \hat{x}(\|x_1\|)) - d(a\hat{x}(s(t)), a\hat{x}(s(\|x_1\|)))| \\ &\leq d(\hat{x}(t), a\hat{x}(s(t))) + d(\hat{x}(\|x_1\|), a\hat{x}(s(\|x_1\|))) \\ &\leq 4\delta. \end{aligned}$$

From inequality (3.2) we deduce that $s(t) - t > 2\delta$.

This last conclusion begins to seem absurd when one considers its implications in terms of the geometry of the cycle of N quadrilaterals formed by successive pairs

$(a^{i-1}\hat{x}, a^i\hat{x})$, with indices modulo N. To tease out this absurdity, we define a sequence of integers $t_1, \ldots, t_N$ in I satisfying the condition $(t - \|x_1\|)(s(t) - s(\|x_1\|)) \geq 0$. Note that in this construction, the estimate (3.1) is used to bound $t_{i+1} - t_i$, thus ensuring that $t_j \leq \|x\| - \|b\| - 2\delta$ (in the light of our hypothesised lower bound on $\|x\| = \|x_1 x_2\|$).

In the light of (3.1) and (3.2), we may begin with

$$t_1 := \begin{cases} \|x_1\| & \text{if } \|x_1\| < 2\|a\| + 4\delta, \\ \|a\| + 2\delta & \text{otherwise.} \end{cases}$$

If $1 \leq k < N$, and t_k has been defined, then

$$t_{k+1} := \begin{cases} s(\|x_1\|) + \|a\| + 2\delta & \text{if } t_k \leq \|x_1\| - \|a\| - 2\delta < s(t_k), \\ s(t_k) & \text{otherwise.} \end{cases}$$

With these choices, for each $i = 1, \ldots, N-1$ we have

$$t_i + 2\delta < s(t_i) \leq t_{i+1}. \tag{3.3}$$

Noting that $a^N = 1$, we apply the triangle inequality to

$$\begin{aligned} s(t_N) - t_1 &= d(\hat{x}(t_1), a^N \hat{x}(s(t_N))) \\ &\leq \sum_{i=1}^{N} d(a^{i-1}\hat{x}(t_i), a^i \hat{x}(s(t_i))) + \sum_{i=1}^{N-1} d(a^i \hat{x}(s(t_i)), a^i \hat{x}(t_{i+1})). \end{aligned}$$

Because $\hat{x}$ is geodesic and multiplication by a is an isometry, the summands of the second sum can be written as $t_{i+1} - s(t_i)$. The summands in the first sum can be written as $d(\hat{x}(t_i), a\hat{x}(s(t_i)))$, which is at most 2δ by definition. These observations explain the first line in the following inequalities. The second line comes from the left-hand inequality in (3.3).

$$\begin{aligned} s(t_N) - t_1 &\leq 2N\delta \ + \ \sum_{i=1}^{N-1} (t_{i+1} - s(t_i)) \\ &< \sum_{i=1}^{N} (s(t_i) - t_i) + \sum_{i=1}^{N-1} (t_{i+1} - s(t_i)) \\ &= s(t_N) - t_1. \end{aligned}$$

This is the desired contradiction. □

Corollary 3.2 *Let $[a_1, \ldots, a_m]$ be a list of m pairwise distinct, nontrivial elements of finite order in Γ. Suppose that $x \in \Gamma$ satisfies*

$$\|x\| \geq (2k+5)^{4\delta+2}(\mu + 2\delta),$$

where $\mu = \max\{\|a_1\|, \|x^{-1}a_1 x\|, \ldots, \|a_m\|, \|x^{-1}a_m x\|\}$. Then

$$m \leq (2k)^{8\delta}.$$

Proof. Since every finite subgroup of Γ is conjugate to one contained in the ball B of radius $4\delta+2$ centred at the identity element (see [3, page 460, proof of Theorem 3.2]), the order of each a_i is bounded above by the number of vertices in this ball, which is in turn bounded above by $(2k+1)^{4\delta+2}$, since Γ is k-generated.

Hence Lemma 3.1 applies, for each $i = 1, \ldots, m$, with $a = a_i$ and $b = x^{-1}a_i x$. Choose a subdivision $x = x_1 x_2$ as in Lemma 3.1, with $\|x\| = \|x_1\| + \|x_2\|$, $\|x_1\| \geq \mu + 2\delta$ and $\|x_2\| \geq \mu + 2\delta$. Then, by Lemma 3.1, $\|x_1^{-1} a_i x_1\| \leq 8\delta$ for all $i = 1, \ldots, m$. Since the a_i are pairwise distinct and nontrivial, the same is true for the $x_1^{-1} a_i x_1$, hence m is bounded above by the number of nontrivial elements in the ball of radius 8δ around the identity element, which in turn is bounded above by $(2k)^{8\delta}$. □

Theorem 3.3 *Let $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_m]$ be sets of torsion elements in a k-generator, δ-hyperbolic group Γ. Let $\mu = \max\{\|a_1\|, \|b_1\|, \ldots, \|a_m\|, \|b_m\|\}$. If A and B are conjugate, then there exists $x \in \Gamma$ such that $x^{-1}Ax = B$ and*

$$\|x\| \leq (2k+5)^{4\delta+2}(\mu + 2\delta) + (2k)^{8\delta(2k)^{8\delta}}.$$

Proof. There is no loss of generality in assuming that the a_i are pairwise distinct and nontrivial, and (renumbering the b_i if necessary) that there exists an element $x \in \Gamma$ such that $x^{-1}a_i x = b_i$ for all $i = 1, \ldots, m$.

We choose such an x of shortest possible length ℓ, say

$$x = g_1 g_2 \cdots g_\ell,$$

(with each $g_i \in S^{\pm 1}$). If $\|x\| = \ell$ satisfies the inequality in the statement, then we are done. Suppose then that

$$\ell > (2k+5)^{4\delta+2}(\mu + 2\delta) + (2k)^{8\delta(2k)^{8\delta}}.$$

For $t = 0, \ldots, \ell$, define $x_t = g_1 g_2 \cdots g_t$, and for each $i = 1, \ldots, m$ define $w(t, i) = x_t^{-1} a_i x_t$. Then it follows from Lemma 3.1 that $\|w(t,i)\| < 8\delta$ for all $i = 1, \ldots, m$ and for all $t = \mu + 2\delta, \ldots, \ell - \mu - 2\delta$. It also follows from Corollary 3.2 that $m \leq (2k)^{8\delta}$. Since the number of nontrivial elements of Γ of length $\leq 8\delta$ is at most $(2k)^{8\delta}$, the number of distinct m-tuples of such elements is at most

$$(2k)^{8\delta m} \leq (2k)^{8\delta(2k)^{8\delta}}.$$

Hence there exist values s, t with $\mu + 2\delta \leq s < t \leq \ell - \mu - 2\delta$ such that $w(s,i) = w(t,i)$ for all $i = 1, \ldots, m$. Now define

$$y = x_s x_t^{-1} x = g_1 \cdots g_s g_{t+1} \cdots g_\ell,$$

so that $\|y\| \leq \ell + s - t < \|x\|$, and for each $i = 1, \ldots, m$ we have $y^{-1} a_i y = x^{-1} a_i x = b_i$ (since $x_s^{-1} a_i x_s = w(s,i) = w(t,i) = x_t^{-1} a_i x_t$).

This contradicts the choice of x, completing the proof. □

4 The virtually cyclic case

Again, Γ is a δ-hyperbolic group with respect to the word-metric defined by a finite generating set S of cardinality k. In this section we consider conjugacies between finite subsets of Γ that generate infinite, virtually cyclic subgroups.

The following lemma allows us to say more about the structure of an infinite, virtually cyclic subgroup of Γ, in terms of its action on $\partial\Gamma$.

Lemma 4.1 *Suppose that a subgroup H of the hyperbolic group Γ contains an element a of infinite order. Let $a^{\pm\infty}$ denote the corresponding elements of $\partial\Gamma$ (cf. Proposition 2.2). Then*

(1) *H is virtually cyclic if and only if it fixes the pair $\{a^{\pm\infty}\}$ setwise.*

(2) *H has infinite centre if and only if it fixes $\{a^{\pm\infty}\}$ pointwise.*

(3) *Suppose that H is virtually cyclic and let K denote the pointwise stabiliser of $\{a^{\pm\infty}\}$ in H. Then every element of $H \smallsetminus K$ has finite order.*

Proof. If H fixes $\{a^{\pm\infty}\}$ setwise, then it permutes the finite set of special geodesics between $a^{\pm\infty}$. The kernel of this permutation representation has finite index in H and fixes each special geodesic setwise. Since the action of Γ on its Cayley graph is free, this kernel must act by translation on each special geodesic, and so is infinite cyclic.

Conversely, if H is virtually cyclic then it contains a normal subgroup of finite index which is infinite cyclic, generated by some power a^t of a. Now $a^{\pm\infty}$ are the endpoints of the quasi-geodesic $\Lambda = \{a^{nt} \colon n \in \mathbb{Z}\}$. Given $h \in H$, $h(a^{\infty})$ and $h(a^{-\infty})$ are the endpoints of the quasigeodesic $h.\Lambda = \{ha^{nt} \colon n \in \mathbb{Z}\}$, and hence of the asymptotic quasigeodesic $\Lambda^h = \{ha^{nt}h^{-1} \colon n \in \mathbb{Z}\}$. But this last set is simply $\{a^{nt} \colon n \in \mathbb{Z}\}$, since $N := \langle a^t \rangle$ is normal in H. Thus H fixes $\{a^{\pm\infty}\}$ setwise. This proves (1).

If we were in the setting of (1) and N were not central in H, then there would exist $h \in H$ such that $ha^th^{-1} = a^{-t}$ and hence $h(a^{\infty}) = a^{-\infty}$. Thus if H is contained in the pointwise stabiliser of $\{a^{\pm\infty}\}$ then N must be central.

It remains to establish the "if" implication in (2). The centraliser in Γ of any element γ of infinite order fixes $\gamma^{\pm\infty}$ pointwise and hence, in the light of (1), contains $\langle\gamma\rangle$ as a subgroup of finite index. It follows that if H has an infinite centre Z, then Z contains $\langle a^p \rangle$ as a subgroup of finite index for some $p > 0$. But this means that $a^{\pm\infty}$ are the only fixed points of Z in $\partial\Gamma$ (Proposition 2.2(2)). Since the fixed point set of Z is H-invariant, it follows from (1) that H is virtually cyclic. And H must fix $a^{\pm\infty}$ pointwise because $ha^{pn}h^{-1} = a^{pn}$ for all $h \in H$ and $n \in \mathbb{Z}$.

Finally we prove (3). Note that $h^2 \in K$ for every $h \in H$. If h has infinite order, then $(h^2)^{\pm\infty} = h^{\pm\infty}$ are the only fixed points of h^2, so $\{a^{\pm\infty}\} = \{h^{\pm\infty}\}$ and $h \in K$. □

In Section 6 we shall explain how part (3) of the preceding lemma can be used to prove Theorem A in the special case where the lists A and B generate virtually cyclic groups with finite centre; this is done by a reduction to the case considered in the previous section, where the lists consisted only of torsion elements. For the rest of this section we concentrate on the infinite centre case.

Recall from Proposition 2.2 that there exists a global bound $R \leq (2k)^{2\delta}$ on the number of special geodesics joining the two points $a^{\pm\infty} \in \partial\Gamma$ determined by any element $a \in \Gamma$ of infinite order.

Theorem 4.2 *Suppose that $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_m]$ are conjugate lists of elements in Γ. Suppose that the subgroup generated by A is virtually cyclic with infinite centre, and suppose that a_1 has infinite order. Then $\exists x \in \Gamma$ such that $a_i x = x b_i$ for all i, and*

$$\|x\| \leq (R! + 8\delta R + 2)\mu_1 + 8\delta + (2k+1)^{4\delta},$$

where $\mu_1 = \max\{\|a_1\|, \|b_1\|\}$.

Proof. Let H be the subgroup of Γ generated by A, and fix $y \in \Gamma$ such that $a_i y = y b_i$ for all i. The idea of the proof is to identify an element z in the centre of H so that the length of $x = z^{-1}y$ is bounded as in the statement of the theorem. Thus we begin with an exploration of the centre of H.

Let $\mathcal{S} = \{\sigma_1, \ldots, \sigma_r\}$ be the finite set of special geodesics from $a_1^{-\infty}$ to a_1^{∞}. As in Lemma 4.1, the action of H permutes $\mathcal{S}$, with the kernel K of the action $H \to \Sigma(\mathcal{S})$ acting freely on each σ_j by translations. In particular K is cyclic. Arguing as in the proof of Lemma 4.1(2), we see that K is central in H. Note that $1 \neq a_1^{r!} \in K$.

By Proposition 2.3, we know there exists $c \in \Gamma$ such that $a_1 c = c b_1$ and

$$\|c\| \leq \|a_1\| + \|b_1\| + 4\delta + (2k+1)^{4\delta}.$$

Noting that $yc^{-1} \in C_\Gamma(a_1)$, we investigate the centralizer of $C_\Gamma(a_1)$.

For some $p \leq r$, the action of a_1^p on $\mathcal{S}$ must leave some σ_j invariant, acting on it as a translation by an integer distance. Thus, taking $g \in \sigma_j$ we see that the translation number $p\tau(a_1) = \tau(a_1^p) := \lim_n d(1, a_1^{pn})/n = \lim_n d(g, a_1^{pn} g)/n$ is a positive integer. Translation numbers are conjugacy invariants, so it follows that a_1^{pm} is not conjugate to an element of length less than 4δ if $m \geq 4\delta$. The proof of [3, Corollary 3.10 (2), page 462] shows that if $\gamma \in \Gamma$ is not conjugate to such a short element, the centralizer $C_\Gamma(\gamma)$ is contained in the $(2\|\gamma\| + 4\delta)$-neighbourhood of $\langle\gamma\rangle$. Hence $yc^{-1} \in C_\Gamma(a_1) \subset C_\Gamma(a_1^{4\delta p})$ lies in the $(2\|a_1^{4\delta p}\| + 4\delta)$-neighbourhood of $\langle a_1^{4\delta p}\rangle$.

Thus we may write $yc^{-1} = a_1^N \eta$, where $\|\eta\| \leq 4\delta(2p\|a_1\| + 1)$. Write $N = r!q + \rho$, where $0 \leq \rho < r!$. Recall that $p \leq R$ and $r \leq R$. Define $z = a_1^{r!q} \in Z(H)$ and $x = z^{-1}y = a_1^\rho \eta c$. Then $a_i x = x b_i$ for all i, and

$$\begin{aligned}
\|x\| &\leq \rho\|a_1\| + \|\eta\| + \|c\| \\
&\leq R!\|a_1\| + 4\delta(2R\|a_1\| + 1) + \|a_1\| + \|b_1\| + 4\delta + (2k+1)^{4\delta} \\
&\leq (R! + 8\delta R + 2)\mu_1 + 8\delta + (2k+1)^{4\delta}.
\end{aligned}$$

□

5 Non-elementary lists

We continue to assume that Γ is a δ-hyperbolic group with respect to the word-metric defined by a finite generating set S of cardinality k. In this section we consider conjugate lists of elements of Γ that generate non-elementary subgroups. It turns out that for the bound we seek in Theorem A, it is sufficient to consider (sub)lists of length 2.

We begin with a quantification of the fact that if periodic geodesics in hyperbolic spaces do not diverge quickly, then they are forced to remain uniformly close.

Lemma 5.1 *Let ℓ be a positive integer. Suppose that σ_1 and σ_2 are geodesics in the Cayley graph of Γ, and $a_1, a_2 \in \Gamma$ are such that a_i fixes σ_i setwise, acting on it as a translation of length ℓ. Let $p_1 \in \sigma_1$, $p_2 \in \sigma_2$ and $n \in \mathbb{Z}$. If*

$$n > (2k+1)^{6\delta} + (d(p_1,p_2) + d(a_1^n(p_1), a_2^n(p_2)) + 4\delta)/\ell,$$

then σ_1 is contained in the 2δ-neighbourhood of σ_2, and $\langle a_1, a_2\rangle$ is virtually cyclic.

Proof. Let p_0 be the vertex on σ_1 between p_1 and $a_1^n(p_1)$ with $d(p_1,p_0) = d(p_1,p_2)+2\delta$. The lower bound we have imposed on n implies that for $j = 0, \dots, (2k+1)^{6\delta}$, the elements $a_1^j p_0 \in \Gamma$ lie on the arc of σ_1 joining p_1 to $a_1^n(p_1)$. Moreover none of these elements is within a distance $d(p_1,p_2) + 2\delta$ of p_1, nor $d(a_1^n(p_1), a_2^n(p_2)) + 2\delta$ of $a_1^n(p_1)$.

For any point x on $[p_1, a_1^n(p_1)]\sigma_1$ bounded away from the endpoints by the above inequalities, the 2δ-slimness of a quadrilateral with vertices $[p_1, a_1^n(p_1), a_2^n(p_2), p_2]$ ensures that there exists a point $q(x)$ on σ_2, between p_2 and $a_2^n(p_2)$, with $d(x, q(x)) \le 2\delta$. Define $q_0 = q(p_0)$.

For those x between p_0 and $a_1^n(p_1)$ with $d(p_0,x) > 6\delta$, we compare $q(x)$ with the point $q'(x)$ between q_0 and $a_2^n(p_2)$ such that $d(p_0,x) = d(q_0,q'(x))$. Note that $q(x)$ cannot lie between p_2 and q_0 on σ_2, because if it were 2δ-close to any of the sides $[p_1,p_0]$, $[p_0,q_0]$ or $[p_2,p_1]$ of the geodesic quadrilateral $[p_1,p_0,q_0,p_2]$, this would contradict the inequality

$d(y,x) \ge d(p_0,x) \ge 6\delta$ for $y \in [p_1,p_0] \subset \sigma_1$;
$d(y,x) \ge d(p_1,x) - d(p_1,p_2) \ge 8\delta$ for $y \in [p_1,p_2]$; or
$d(y,x) \ge d(p_0,x) - d(p_0,q_0) \ge 4\delta$ for $y \in [p_0,q_0]$.

Since both $q'(x)$ and $q(x)$ lie on the geodesic arc $[q_0, a_2^n(p_2)] \subset \sigma_2$, we have

$$\begin{aligned} d(q(x),q'(x)) &= |d(q_0,q(x)) - d(q_0,q'(x))| \\ &= |d(q_0,q(x)) - d(p_0,x)| \\ &\le d(p_0,q_0) + d(x,q(x)) \le 4\delta, \end{aligned}$$

where the last line is the triangle inequality applied to the path $(p_0,q_0,q(x),x)$, recalling that $d(p_0,q_0) \le 2\delta$ and $d(x,q(x)) \le 2\delta$.

Hence

$$d(x,q'(x)) \le d(x,q(x)) + d(q(x),q'(x)) \le 6\delta.$$

Note that if $x = a_1^j p_0$, then $q'(x) = a_2^j q_0$, because a_2 translates σ_2 by the same distance ℓ as a_1 translates σ_1. Also, $d(p_0, a_1^j p_0) = j\ell > 6\delta$ if $j > 6\delta$. Thus for

$j = 6\delta + 1, \ldots, (2k+1)^{6\delta}$, the element $(a_1^j p_0)^{-1} \cdot (a_2^j q_0)$ belongs to the ball of radius 6δ about the identity in Γ. This ball contains fewer than $1 + (2k)^{6\delta}$ elements, so there must be two distinct values of j (say s and t) for which these elements coincide. In other words

$$(a_1^s p_0)^{-1}(a_2^s q_0) = (a_1^t p_0)^{-1}(a_2^t q_0),$$

whence

$$a_1^{t-s} = a_2^{t-s}.$$

It follows that the endpoints in $\partial\Gamma$ of both σ_1 and σ_2 are $a_1^{\pm\infty} = a_2^{\pm\infty}$, and by the slimness of geodesic quadrilaterals σ_1 is contained in the 2δ-neighbourhood of σ_2. Indeed, given $y \in \sigma_1$, if one chooses $m \in \mathbb{Z}$ so that $a_1^{m(t-s)}(y) \in [a_1^s(p_0), a_1^t(p_0)]$, then

$$d(y, \sigma_2) \leq d(y, a_2^{m(s-t)}(q'(a_1^{m(t-s)}(y)))) = d(a_1^{m(t-s)}(p)), q'(a_1^{m(t-s)}(y))) \leq 2\delta.$$

Moreover, by Lemma 4.1(1), the subgroup $H = \langle a_1, a_2 \rangle$ of Γ is virtually cyclic, since it fixes $\{a^{\pm\infty}\}$ setwise. □

The geometric intuition behind the non-elementary case is as follows. We are told that there is an element x conjugating a_1 to b_1 and a_2 to b_2. By Proposition 2.3 there are short words c_1 and c_2 conjugating a_1 to b_1 and $a_2 a_1 a_2^{-1}$ to $b_2 b_1 b_2^{-1}$, respectively. Any other element conjugating a_1 to b_1 must differ from c_1 by an element of the centraliser of a_1. This centraliser is contained in a tubular neighbourhood of the union of the geodesic lines with endpoints $a_1^{\pm\infty}$. It follows that the distance from x to each of these geodesics is bounded by a linear function of $\|a_1\| + \|c_1\|$. The same argument shows that x must also lie in a certain tubular neighbourhood of the geodesic lines connecting $a_2 \cdot a_1^{\pm\infty}$, where the width of the neighbourhood is a linear function of $\|a_2\| + \|c_2\|$. If $\langle a_1, a_2 \rangle$ is non-elementary, then $a_2 \cdot a_1^{\pm\infty} \neq a_1^{\pm\infty}$, the two sets of geodesics we have been discussing are distinct, and we may use Lemma 5.1 to control the size of the intersection of the two tubular neighbourhoods in which we have argued x must lie.

The constants K and R in the following statement are as in Proposition 2.2.

Theorem 5.2 *Let $a_1, a_2, b_1, b_2, x \in \Gamma$ be such that $xb_1 = a_1 x$, $xb_2 = a_2 x$, a_1 has infinite order, and a_2 does not fix the set $\{a_1^{\pm\infty}\} \subset \partial\Gamma$. Then*

$$\|x\| \leq [(2k+1)^{6\delta} R + 18K + 9R + 18]\, \mu_2 \; + \; 5R + 10K + 24\delta + 5(2k+1)^{4\delta},$$

where $\mu_2 = \max\{\|a_1\|, \|a_2\|, \|b_1\|, \|b_2\|\}$.

Proof. By Proposition 2.2 there is a special geodesic σ_1 joining $a_1^{\pm\infty}$ and an integer t with $0 < t \leq R$ such that a_1^t leaves invariant σ_1, acting as a translation of length ℓ, say. Recall that $C_\Gamma(a_1)$ is contained in the $(1 + \|a_1\|)K$-neighbourhood of σ_1.

Note that $a_2 a_1^t a_2^{-1}$ leaves invariant the special geodesic $\sigma_2 = a_2(\sigma_1)$ and acts on it as a translation of length ℓ. And $C_\Gamma(a_2 a_1 a_2^{-1}) = a_2 C_\Gamma(a_1) a_2^{-1}$ is contained in the $(1 + \|a_2 a_1 a_2^{-1}\|)K$-neighbourhood of σ_2. By hypothesis, the endpoints of σ_2 in $\partial\Gamma$ are distinct from those of σ_1.

By Proposition 2.3, we know there exist $c_1, c_2 \in \Gamma$ so that $a_1 c_1 = c_1 b_1$ and $a_2^{-1} a_1 a_2 c_2 = c_2 b_2^{-1} b_1 b_2$, with

$$\begin{aligned} \|c_1\| &\leq \|a_1\| + \|b_1\| + 4\delta + (2k+1)^{4\delta} \quad \text{and} \\ \|c_2\| &\leq \|a_2^{-1} a_1 a_2\| + \|b_2^{-1} b_1 b_2\| + 4\delta + (2k+1)^{4\delta}. \end{aligned}$$

Let p_1 be the point of σ_1 closest to x. Since $x c_1^{-1} \in C_\Gamma(a_1)$ lies in the $(1 + \|a_1\|)K$-neighbourhood of σ_1,

$$\begin{aligned} d(x, p_1) &\leq \|c_1\| + (1 + \|a_1\|)K \\ &\leq (1+K)\|a_1\| + \|b_1\| + 4\delta + (2k+1)^{4\delta} + K. \end{aligned} \tag{5.1}$$

Similarly, if p_2 is the closest point of σ_2 to x, then

$$\begin{aligned} d(x, p_2) &\leq \|c_2\| + (1 + \|a_2^{-1} a_1 a_2\|)K \\ &\leq (2K+2)\|a_2\| + (K+1)\|a_1\| + 2\|b_2\| + \|b_1\| + 4\delta + (2k+1)^{4\delta} + K. \end{aligned}$$

So by the triangle inequality,

$$d(p_1, p_2) \leq 4(K+2)\mu_2 + 8\delta + 2(2k+1)^{4\delta} + 2K. \tag{5.2}$$

The identity element 1 also lies in the $(1 + \|a_1\|)K$-neighbourhood of σ_1, so there is an integer n such that

$$d(1, a_1^{tn}(p_1)) \leq (1 + \|a_1\|)K + t \cdot \|a_1\| < (R+K)(\|a_1\| + 1). \tag{5.3}$$

Similarly, there exists an integer m such that

$$d(1, a_2^{-1} a_1^{tm} a_2(p_2)) \leq (1 + \|a_2^{-1} a_1 a_2\|)K + t\|a_1\| < (R+K)(\|a_1\| + 1 + 2\|a_2\|).$$

Hence

$$d(a_1^{tn}(p_1), a_2^{-1} a_1^{tm} a_2(p_2)) \leq 2(R+K)(\|a_1\| + \|a_2\| + 1). \tag{5.4}$$

By combining (5.2) and (5.4) we obtain

$$\begin{aligned} |n-m|\ell &= |d(p_1, a_1^{tn}(p_1)) - d(p_2, a_2^{-1} a_1^{tm} a_2(p_2))| \\ &\leq d(p_1, p_2) + d(a_1^{tn}(p_1), a_2^{-1} a_1^{tm} a_2(p_2)) \\ &\leq (4R + 8K + 8)\mu_2 + 2R + 8\delta + 2(2k+1)^{4\delta} + 4K. \end{aligned}$$

Hence

$$\begin{aligned} d(a_1^{tn}(p_1), a_2^{-1} a_1^{tn} a_2(p_2)) &\leq d(a_1^{tn}(p_1), a_2^{-1} a_1^{tm} a_2(p_2)) + |n-m|\ell \\ &\leq (8R + 12K + 8)\mu_2 + 4R + 6K + 8\delta + 2(2k+1)^{4\delta}. \end{aligned}$$

Since σ_1 and σ_2 have different endpoints in $\partial\Gamma$, neither is contained in the 2δ-neighbourhood of the other. Hence, by Lemma 5.1, we must have

$$\begin{aligned} n &\leq (2k+1)^{6\delta} + \frac{1}{\ell}[d(p_1, p_2) + d(a_1^{tn}(p_1), a_2^{-1} a_1^{tn} a_2(p_2)) + 4\delta] \\ &\leq (2k+1)^{6\delta} + \frac{1}{\ell}[(8R + 16K + 16)\mu_2 + 8K + 4R + 20\delta + 4(2k+1)^{4\delta})]. \end{aligned} \tag{5.5}$$

By using the triangle inequality with estimates (5.3), (5.5) and (5.1), we see that

$$\begin{aligned}\|x\| &= d(1,x)\\ &\le d(1,a_1^{tn}(p_1)) + d(a_1^{tn}(p_1),p_1) + d(p_1,x)\\ &= d(1,a_1^{tn}(p_1)) + n\ell + d(p_1,x)\\ &\le (R+K)(\|a_1\|+1)\\ &\qquad + \Big(\ell(2k+1)^{6\delta} + [(8R+16K+16)\mu_2 + 8K + 4R + 20\delta + 4(2k+1)^{4\delta})]\Big)\\ &\qquad + \Big((1+K)\|a_1\| + \|b_1\| + 4\delta + (2k+1)^{4\delta} + K\Big).\end{aligned}$$

Thus, gathering terms and noting that $\ell = \tau(a_1^t) = t\tau(a_1) \le R\|a_1\| \le R\mu_2$, we finally obtain

$$\|x\| \le \mu_2[(2k+1)^{6\delta}R + 18K + 9R + 18)] \; + \; 5R + 10K + 24\delta + 5(2k+1)^{4\delta},$$

as claimed. □

6 Proof of Theorem A

In this section we complete the proof of Theorem A (as expanded below to include specific bounds), and note the obvious application to the conjugacy problem for finite subsets of Γ.

Theorem A′ *Let Γ be a group that is δ-hyperbolic with respect to a finite generating set of cardinality k. Then there exist constants α and β (depending only on δ and k) with the following property: if two finite lists $A = [a_1, \dots, a_m]$, $B = [b_1, \dots, b_m]$ of elements are conjugate in Γ, then there exists $x \in \Gamma$ such that $x^{-1}a_i x = b_i$ for $i = 1, \dots, m$ and*

$$\|x\| \le \alpha\mu + \beta,$$

where $\mu = \max\{\|c\| \colon c \in A \cup B\}$.

(It suffices to take $\alpha = ((2k+1)^{2\delta+1})!$ and $\beta = (2k+1)^{8\delta(2k)^{8\delta}}$.)

Proof of Theorem A′. There are four cases to consider.

The torsion case: Suppose first that A (and hence B) consists only of torsion elements. Then by Theorem 3.3 we can find a conjugating element x of length bounded by

$$\|x\| \le \alpha_1\mu + \beta_1,$$

where $\alpha_1 = (2k+5)^{4\delta+2}$ and $\beta_1 = 2\delta(2k+5)^{4\delta+2} + (2k)^{8\delta(2k)^{8\delta}}$.

Henceforth we may assume that A contains at least one element of infinite order. Renumbering A and B if necessary, suppose that a_1 has infinite order. Let $a^{-\infty}, a^{+\infty} \in \partial\Gamma$ denote the corresponding boundary points. Lemma 4.1 tells us that the subgroup H of Γ generated by A is virtually cyclic if and only if H fixes $\{a^{\pm\infty}\}$ setwise. Moreover, in this case H has an infinite centre if and only if H fixes $\{a^{\pm\infty}\}$ pointwise.

Virtually cyclic with finite centre: Suppose now that H is virtually cyclic with finite centre. Then the pointwise stabiliser K of $\{a^{\pm\infty}\}$ in H has index 2, and by Lemma 4.1(3) every element of $K \smallsetminus H$ has finite order. In particular $a_1 \in K$. Renumbering A and B again if necessary, we may assume $a_2 \notin K$. We now apply Nielsen transformations to replace A and B by lists of torsion elements. Specifically, define $A' = [a'_1, \dots, a'_m]$, where

$$a'_i = \begin{cases} a_i a_2 & \text{if } a_i \in K \\ a_i & \text{otherwise.} \end{cases}$$

Similarly, define $B' = [b'_1, \dots, b'_m]$, where

$$b'_i = \begin{cases} b_i b_2 & \text{if } a_i \in K \\ b_i & \text{otherwise.} \end{cases}$$

Note that

$$\mu' = \max\{\|c\| \colon c \in A' \cup B'\} \leq 2\mu.$$

By Theorem 3.3 there is an element $x \in \Gamma$ such that $x^{-1}a'_i x = b'_i$ for all i, and

$$\|x\| \leq \alpha_1 \mu' + \beta_1 \leq (2\alpha_1)\mu + \beta_1.$$

Finally, for each $i = 1, \dots, m$ we have

$$x^{-1}a_i x = x^{-1}a'_i x = b'_i = b_i$$

if $a_i \notin K$, while

$$x^{-1}a_i x = x^{-1}a'_i a_2 x = b'_i b_2 = b_i$$

if $a_i \in K$.

Virtually cyclic with infinite centre: Suppose now that H is virtually cyclic with infinite centre. Then by Theorem 4.2 there exists $x \in \Gamma$ such that $x^{-1}a_i x = b_i$ for $i = 1, \dots, m$ and

$$\|x\| \leq \alpha_2 \mu + \beta_2,$$

where $\alpha_2 = R! + 8\delta R + 2$ and $\beta_2 = 8\delta + (2k+1)^{4\delta}$.

The non-elementary case: Finally, suppose that H is not virtually cyclic. Then H does not fix $\{a^{\pm\infty}\}$ setwise. Renumbering A and B if required, we may assume that a_2 does not fix $\{a^{\pm\infty}\}$. If $x \in \Gamma$ is such that $x^{-1}a_i x = b_i$ for $i = 1, \dots, m$, then, by Theorem 5.2,

$$\|x\| \leq \alpha_3 \mu + \beta_3,$$

where $\alpha_3 = ((2k+1)^{6\delta} R + 18K + 9R + 18)$ and $\beta_3 = 5R + 10K + 24\delta + 5(2k+1)^{4\delta}$.

These four cases cover all possibilities, so there exists $x \in \Gamma$ such that $x^{-1}a_i x = b_i$ for all $i = 1, \dots, m$ and

$$\|x\| \leq \alpha\mu + \beta,$$

where $\alpha = \max\{2\alpha_1, \alpha_2, \alpha_3\}$ and $\beta = \max\{\beta_1, \beta_2, \beta_3\}$ depend only on δ and k. (The constants in the statement of Theorem A′ are crude estimates on these numbers.) □

Theorem A leads immediately to an exponential-time algorithm to solve the conjugacy problem for finite lists and finite subsets in hyperbolic groups. More efficient algorithms will be discussed in the next section.

Corollary 6.1 *Let Γ be a hyperbolic group. Then there is an algorithm which, given two finite lists $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_m]$ of words in the generators of Γ, will decide whether or not the lists of elements of Γ represented by A and B are conjugate in Γ.*

Proof. Let α and β be the constants given in Theorem A, and let $\mu = \max\{|w| : w \in A \cup B\}$, where $|\cdot|$ denotes word length. For each of the finitely many words x of length $\leq \alpha\mu + \beta$, one can use the solution to the word problem of Γ to check for $i = 1, \ldots, m$ whether or not $x^{-1}a_i x = b_i$ in Γ. If one finds an x for which these equalities hold, then clearly A is conjugate to B. Otherwise, by Theorem A, A is not conjugate to B. □

Corollary 6.2 *Let Γ be a hyperbolic group. Then there is an algorithm which, given two finite lists $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_n]$ of words in the generators of Γ, will decide whether or not the subsets of Γ represented by A and B are conjugate in Γ.*

Proof. The solution to the word problem for Γ enables us to check algorithmically whether the words in A represent distinct elements of Γ. If we find $a_i = a_j$ for some $i < j$, we delete a_j. Similarly, we delete repetitions from B.

Having eliminated repetitions (after finitely many steps), we can detect whether or not the cardinalities of the subsets of Γ represented by A and B agree. If they do not, then A and B are non-conjugate. Thus we may assume that $m = n$ and that no element of Γ occurs twice in either A or B. Now, the subset represented by A is conjugate to the subset represented by B if and only if the list represented by A is conjugate to some permutation of the list represented by B. Since there are only finitely many permutations to check, this problem is soluble by Corollary 6.1. □

7 Algorithms

In this section we describe a quadratic time algorithm that decides whether finite lists of elements in torsion-free hyperbolic groups are conjugate. We shall also prove Theorem B, as stated in the introduction. Along the way we shall describe efficient algorithms for solving several related problems in hyperbolic groups. In particular we describe an algorithm that enumerates, in time bounded by a quadratic function of $r + |w|$, the elements of length at most r in the centraliser of any element w of infinite order.

Once again, Γ will be a fixed group that is δ-hyperbolic with respect to a finite generating set S of cardinality k. When we speak of *words*, we mean elements of the free monoid on $S^{\pm 1}$.

Recall that there is a partial algorithm that, given an arbitrary finite presentation, will halt if the group is hyperbolic, and calculate a real number δ such that it is δ-hyperbolic; see [15].

7.1 Solving the word problem in linear time

We begin by reminding the reader how Dehn's algorithm can be used to solve the word problem of any hyperbolic group in linear time. The algorithm we describe is

essentially the same as that on pages 33-35 of [17] (see also [11]). The first thing to recall is that if X is a δ-hyperbolic geodesic space and $M > 8\delta$ is an integer, then a non-constant M-local geodesic in X cannot be a loop (Corollary 1.14 on page 407 of [3]). Dehn's algorithm is an application of this fact to words in the generators S, regarding them as labels on edge-paths in the Cayley graph $X(\Gamma, S)$:

Algorithm DA: Fix $M \geq 8\delta + 1$. Given a word w, one reads it from the left searching for subwords of length at most M that are not geodesic; if such a subword u is found, it is replaced by a geodesic word u' such that $u' = u$ in Γ; one then begins reading the (edited) word again, starting $M - 1$ letters to the left of u' (or the beginning of w, if the prefix before u' has length less than $M - 1$). Since the length of w is reduced by an integer amount by each substitution, this algorithm terminates after at most $|w|$ reductions. And since the amount of "backing-up" that one does after each edit is bounded by a constant, the number of subwords of length $\leq M$ that are read in the course of this algorithm is also bounded by a linear function of $|w|$. The algorithm terminates when one reaches the end of the (edited form of) w, at which point one concludes that $w = 1$ in Γ if and only if w has been reduced to the empty word. In general, at the end of the algorithm one has transformed w into a M-local geodesic (and this output will be used in the following subsections).

Remark. Arguing as above, one obtains a finite presentation $\Gamma = \langle S \mid R \rangle$ by taking R to be the set of words of length at most $16\delta + 1$ that equal the identity in Γ. (Indeed, it is such a presentation that is being used implicitly in the algorithm.)

7.2 Conjugacy for individual elements

We recall a strategy that solves the conjugacy problem efficiently in any hyperbolic group. This discussion follows that on pages 452–453 of [3].

First think of how one solves the conjugacy problem in a finitely generated free group: given two words, one performs free reductions and cyclic permutations in order to make the words cyclically reduced; the given pair of words are conjugate if and only if their cyclically reduced forms are the same (up to cyclic permutation).

One can follow the outline of this proof in order to solve the conjugacy problem in an arbitrary hyperbolic group. To this end, one defines a word in the generators of such a group to be *M-cyclically reduced* if it and all of its cyclic permutations are M-local geodesics, and one defines a word to be *fully reduced* if it and all of its cyclic permutations are actually geodesic. A now-standard, diagram-surgery argument is used to prove ([3, Lemma 2.9, page 452]):

Lemma 7.1 *If two fully reduced words u and v represent conjugate elements of Γ, then*

1. $\max\{|u|, |v|\} \leq 8\delta + 1$, *or else*
2. *there exist cyclic permutations u' and v' of u and v and a word w of length at most $2\delta + 1$ such that $wu'w^{-1} = v'$ in Γ.*

If applied naïvely, this lemma does not yield a polynomial time solution to the conjugacy problem, because the process of passing from an arbitrary word in the generators

to the corresponding fully reduced form involves a large number of checks to verify its geodesic nature. One can circumvent this difficulty by exploiting the fact that local geodesics in hyperbolic spaces are good approximations to geodesics. More specifically, one has the following result, (Theorem 1.13 on page 405 of [3]), other aspects of which we'll need shortly.

Proposition 7.2 *Let X be a δ-hyperbolic geodesic space and let $c : [a, b] \to X$ be an M-local geodesic with $M > 8\delta$. Then*

(1) *$im(c)$ is contained in the 2δ-neighbourhood of any geodesic segment $[c(a), c(b)]$ connecting its endpoints,*

(2) *$[c(a), c(b)]$ is contained in the 3δ-neighbourhood of $im(c)$, and*

(3) *c is a $(\lambda, 2\delta)$-quasi-geodesic, where $\lambda = (M + 4\delta)/(M - 4\delta)$.*

This leads us to the following modification of Lemma 7.1.

Lemma 7.3 *Fix $M \geq 8\delta + 1$. There is a constant θ, depending only on δ, k and M, such that if the words u and v are cyclically M-reduced, then*

(1) *$\max\{|u|, |v|\} \leq \theta$, or else*

(2) *there exists a cyclic permutation u' of u and a word w of length at most θ such that $wu'w^{-1} = v$ in Γ.*

Remark. On page 453 of [3] the above lemma is stated with a cyclic conjugate v' of v in the second item, but a simple diagram-surgery shows that this is unnecessary.

Algorithm 7.4 *Fix $M \geq 8\delta + 1$. Given two words u and v, one considers them as cyclic words and proceeds as in* **DA**, *replacing non-geodesic subwords of length $\leq M$. At the end of this process u and v have been replaced by cyclically M-reduced words u_0 and v_0. Lemma 7.3 provides a* finite *set of words[1] T such that u is conjugate to v in Γ if and only if $wu_0'w^{-1} = v_0$ in Γ for some $w \in T$ and some cyclic permutation u_0' of u_0. One uses* **DA** *to decide the validity of each of the putative relations $wu_0'w^{-1} = v_0$.*

Remark. (Sub-quadratic Time) If one allows inputs to be treated as cyclic words, then it is clear that the above algorithm will run in linear time. However, the obvious implementation on a Turing machine will not run in linear time because in order to treat u and v as cyclic words one has to scan back to the beginning of the tape when one reaches the end of the word. (Such a naïve implementation therefore becomes quadratic-time.)

Epstein and Holt [7] avoid this difficulty by returning to the beginning of the word only once. They then exploit the fact that this is sufficient to transform u and v into words all of whose cyclic permutations are quasi-geodesics with fixed constants. Their idea leads to an algorithm of Turing complexity $O(n \log n)$ that will run in linear time if

[1] By definition, this finite set includes one conjugating element for each conjugate pair (u, v) as in Lemma 7.3(1).

one allows a RAM model of computation in which arithmetic operations are performed in constant time.

For our purposes, the more naïve Algorithm 7.4 will suffice. We shall use it in the following form.

Lemma 7.5 *There is a constant K_c, depending only on δ, k and M, with the following property. Given two words u and v that represent conjugate elements of Γ, one can employ Algorithm 7.4 to find, in time $O(|u|+|v|)^2$, a word w with $|w| \le K_c(|u|+|v|)$ such that $w^{-1}uw = v$ in Γ.*

Proof. First we consider the cyclic permutations of u and v that are made in Algorithm 7.4; our aim is to find short words w_1 and w_2 such that $w_1^{-1}uw_1 = u_0$ and $w_2^{-1}vw_2 = v_0$. The key point to observe is that the only times when one needs to make a cyclic permutation of (the edited form of) u are when the word itself is an M-local geodesic but some cyclic permutation of it is not; and when it is made, the permutation only involves conjugating by a word of length less than M. Immediately following the permutation, the length of the edited form of u is reduced. Thus the sum of the lengths of the conjugating elements over all permutations made is less than $M|u|$. We require the machine implementing our algorithm to record each of the conjugating elements for the permutations, thus finding w_1 with $w_1^{-1}uw_1 = u_0$ and $|w_1| \le M|u|$. Similarly, the algorithm finds w_2 with $w_2^{-1}vw_2 = v_0$ and $|w_2| \le M|v|$.

The final stage of the algorithm finds a word w_0 in the finite set T such that $w_0^{-1}\pi^{-1}u_0\pi w_0 = v_0$, where π is a suffix of u_0.

Define $w = w_1\pi w_0 w_2^{-1}$. Then $w^{-1}uw = v$ and

$$|w| \le (M+1)(|u|+|v|) + \max\{|\omega| \colon \omega \in T\}.$$

Lemma 7.3(2) together with Proposition 2.3 provides a bound on the length of words in T that depends only on δ, k and M. □

If one is unconcerned with the computational complexity of finding the word w in the preceding lemma, then Proposition 2.3 shows that one can sharpen considerably the bound on its length.

7.3 Calculation of Centralizers

Given a word w in the generators, we write $C_\Gamma(w)$ to denote the centralizer of w in Γ. Balls about $1 \in \Gamma$ in the word metric will be denoted $\mathrm{Ball}_r = \{g \in \Gamma \colon \|g\| \le r\}$.

For any hyperbolic group Γ (indeed any biautomatic group) there is an algorithm that, given a word w in the generators, will calculate a set of generators for $C_\Gamma(w)$ and the quasi-convexity constant for $C_\Gamma(w) \hookrightarrow \Gamma$; see [4]. It follows that there is an algorithm that, given w and an integer $r > 0$, will calculate the intersection $C_\Gamma(w) \cap \mathrm{Ball}_r$. Unfortunately, the time function of this algorithm is exponential in both $|w|$ and r. In this subsection we adopt an altogether different approach to prove:

Theorem 7.6 *Let Γ be a group that is δ-hyperbolic with respect to a generating set of cardinality k. There are constants*[2] *$L_1, \ldots, L_5$, depending only on k and δ, and an algorithm that, given a word w in generators of Γ and an integer $r > 0$ will, in time $\leq L_1(|w| + r)^2$, calculate*

$$C_\Gamma(w) \cap \mathrm{Ball}_r$$

and output the elements as a list of at most $L_2(|w| + r) + L_3$ words, each of length at most $L_4(|w| + r) + L_5$, provided that w has infinite order in Γ.

The idea of the proof is to identify a set $Q(w, r)$ that contains $C_\Gamma(w) \cap \mathrm{Ball}_r$ and has a membership that is easily calculable and not too big. We then exploit the efficient solution to the word problem **DA** to check for each $q \in Q$ whether $[q, w] = 1$ in Γ.

The following argument can be applied with any $M \geq 8\delta + 1$, but it is convenient to take $M = 12\delta$.

We remind the reader once more that if $\gamma \in \Gamma$ is not conjugate to an element of length less than 4δ, then the argument on page 463 of [3] shows that $C_\Gamma(\gamma)$ is contained in the $(2\|\gamma\| + 4\delta)$-neighbourhood of $\langle\gamma\rangle$.

Let w be a word that represents an element of infinite order in Γ. As in the proof of Theorem 4.2, we know that w_ϵ^P has integral translation number for some $P \leq R$, and hence $w^{R!M}$ is not conjugate to an element of length less than M. Apply **DA** to $w^{R!M}$ cyclically, as in Algorithm 7.4, to find a cyclically M-reduced word w_0 and a word x with $|x| \leq R!M|w|$ such that $xw^{R!M}x^{-1} = w_0$.

Let $\mathrm{Pre}(w_0)$ denote the set of prefixes of positive powers of the word w_0. Given an integer $r > 0$, define

$$Q(w, r) = \{x^{-1}p\eta x \mid \ |\eta| \leq 7\delta;\ p \in \mathrm{Pre}(w_0);\ |p| \leq 2r + 4R!M\,|w| + 18\delta.\}.$$

Lemma 7.7 *Let $\overline{Q}(w, r)$ denote the image of $Q(w, r)$ in Γ. Then*

$$C_\Gamma(w) \cap \mathrm{Ball}_r \subset \overline{Q}(w, r).$$

Proof. We have fixed $M = 12\delta$.

Write $w^{R!M} = x^{-1}w_0x$ in Γ, where $|x| \leq R!M\,|w|$ and w_0 is cyclically M-reduced.

Let $\gamma_0 \in \Gamma$ be the group element represented by the word w_0. Then w_0 is the label of a geodesic ℓ_0 in $X(\Gamma, S)$ from 1 to γ_0.

Consider the bi-infinite path $\ell = \bigcup_{n\in\mathbb{Z}} \gamma_0^n(\ell_0)$ in $X(\Gamma, S)$.

Since w_0 is cyclically M-reduced, ℓ is an M-local geodesic with endpoints $\gamma_0^{\pm\infty} \in \partial\Gamma$. Hence, for any $z \in C_\Gamma(w_0)$, the path $z(\ell)$ is also an M-local geodesic with endpoints $\gamma_0^{\pm\infty} \in \partial\Gamma$. In other words, ℓ and $z(\ell)$ are a bounded Hausdorff distance (D, say) apart. Proposition 7.2 allows us to specify a more precise bound 7δ for the Hausdorff distance, as follows.

Given a point p_0 on ℓ, choose two points p_1, p_2 on ℓ, on either side of p_0, with $d(p_0, p_i) > D + 4\delta$ for $i = 1, 2$.

[2] one could of course streamline this to a single constant: increase L_2 and L_4 to allow $L_3 = L_5 = 0$, then take the maximum of L_1, L_2, L_4

Let q_1, q_2 be points on $z(\ell)$ with $d(p_i, q_i) \leq D$ for $i = 1, 2$. Choose geodesic segments $[p_1, p_2]$ and $[q_1, q_2]$.

By Proposition 7.2 and the 2δ-slimness of geodesic quadrilaterals, there are points $p' \in [p_1, p_2]$, $q' \in [q_1, q_2]$ and $q_0 \in z(\ell)$ with $d(p_0, p') \leq 2\delta$, $d(p', q') \leq 2\delta$, and $d(q', q_0) \leq 3\delta$. Hence $d(p_0, q_0) \leq 7\delta$, giving the desired bound.

In particular, $z = p\eta$ where p is a prefix of a power of w_0 and $|\eta| \leq 7\delta$. Thus every element of $C_\Gamma(w)$ can be written as $\zeta = x^{-1}p\eta x$. We are interested only in those ζ with $d(1, \zeta) \leq r$. The triangle inequality tells us that

$$d(1, \zeta) \geq d(x^{-1}, x^{-1}p) - d(1, x^{-1}) - d(\zeta, x^{-1}p) \geq d(1, p) - 2\,|x| \, - \, |\eta|.$$

Since $|x| \leq R!M\,|w|$ and $|\eta| \leq 7\delta$, it follows that $d(1, \zeta) > r$ if $d(1, p) > r + 2R!M\,|w| + 7\delta$. Now ℓ is a $(2, 2\delta)$-quasi-geodesic (Proposition 7.2), so $d(1, p) \geq \frac{1}{2}|p| - 2\delta$, and therefore $d(1, \zeta) > r$ if

$$|p| > 2r + 4R!M\,|w| + 18\delta.$$

□

Proof of Theorem 7.6.

Let $\kappa = (2r + 4R!M\,|w| + 18\delta)$.

By construction, the number of words in $Q(w, r)$ is bounded by $((\kappa+1)(2k+1)^{7\delta})$, and the length of each word is bounded by $(\kappa + 4R!M\,|w| + 7\delta)$.

There is an obvious algorithm for constructing $Q(w, r)$: enumerate $\mathrm{Ball}_{7\delta}$ as $\eta_1, \eta_2, \ldots, \eta_V$ (a finite process); apply **DA** cyclically to obtain the decomposition $w = x^{-1}w_0x$; then, taking the first κ prefixes $p \in \mathrm{Pre}(w_0)$ in order of increasing length, form the sublists $x^{-1}p\eta_1 x, \ldots, x^{-1}p\eta_V x$.

It is easy to pass from the resulting list of elements of $Q(w, r)$ to an irredundant list of elements of $\overline{Q}(w, r)$ because all repetitions arise from equalities in Γ of the form $s\eta_i = s'\eta_j$, where s and s' are suffixes of length at most 14δ in some p and p'. Therefore repetitions can be avoided by simply including into the algorithm a search for the finite list of forbidden possibilities for $s\eta_j$.

The processes of generating $Q(w, r)$ and of deleting redundant elements each involve $O(|w| + r)$ implementations of algorithms running in time $O(|w| + r)$. Thus, in time $O(|w|+r)^2$, we can generate a list of words, each of length at most $(\kappa+4R!M\,|w|+7\delta)$, with one word representing each element of $\overline{Q}(w, r)$. For each word q in the list, we can use **DA** to check in time $O(|w| + r)$ whether $[q, w] = 1$ in Γ. We delete q from the list if and only if $[q, w] \neq 1$. Applying this step to all the $O(|w| + r)$ words in the list therefore also takes time $O(|w| + r)^2$. □

7.4 Conjugacy for finite lists

We are now in a position to prove Theorem B, as stated in the introduction. First we shall describe the algorithm, then we shall argue that it does indeed determine if the given lists of elements are conjugate.

Algorithm 7.8 *The algorithm takes as input a pair of finite lists of words* $A = [a_1, \ldots, a_m]$ *and* $B = [b_1, \ldots, b_m]$.

Step 1(i): *Apply* **DA** *to check if* $a_1 = b_1 = 1$ *in* Γ. *If both equal* 1, *and* $m = 1$, *then stop and output YES. If both equal* 1, *and* $m > 1$, *then delete* a_1 *and* b_1 *and start again. If only one of them equals* 1, *then stop and output NO. If neither equals* 1, *then proceed to Step 1(ii).*

Step 1(ii): *Apply* **DA** *to search for the least* j *such that*[3] $a_j^N \neq 1$ *in* Γ. *If* j *is found, simultaneously reorder the indices of both lists so that* $j = 1$, *and pass to Step 2. If no such* j *is found, apply the same procedure to the* b_i*; if* $b_j^N \neq 1$ *for some* j, *output NO; if there is no such* j, *stop and output* TORSION.

Step 2: *Apply Algorithm 7.4 to determine if there exists* y *such that* $y^{-1}a_1y = b_1$, *and (following Lemma 7.5) find it if it exists. If there is no such* y, *output NO; if such* y *exists proceed to Step 3.*

Step 3: *Let* α *and* β *be as in Theorem A. Let* μ *be the length of the longest word in either of the lists* A *and* B, *and let* $r_{A,B} = \alpha\mu + \beta + |y|$. *Proceeding as in the proof of Theorem 7.6, enumerate the words* $q \in Q(a_1, r_{A,B})$. *For each* q, *apply* **DA** *to check whether* $(qy)^{-1}a_i(qy) = b_i$ *in* Γ *for* $i = 1, \ldots, m$. *If one of these equalities fails, proceed to the next element in the enumeration of* Q*; output NO if the enumeration of* Q *has been exhausted. If for some* q *all of the equalities* $(qy)^{-1}a_i(qy) = b_i$ *hold, stop, output YES, and declare* qy *to be a conjugating element.*

Proof of Theorem B. It is clear that Algorithm 7.8 will terminate and when it outputs YES a conjugating element has indeed been found. Likewise, when it outputs TORSION, it does so correctly. In the light of the earlier results in this section, it is also clear that the running time of Step 1 is bounded above by a linear function of $\|A\| + \|B\|$.

Similarly, the running time of Step 2 is bounded by a linear function of μ^2.

The first non-trivial point to check is that if Step 3 outputs NO, then the input lists are not conjugate. Since $y^{-1}a_1y = b_1$, the elements of Γ conjugating a_1 to b_1 are precisely those of the form zy with $z \in C_\Gamma(a_1)$. According to Theorem A, we may restrict our attention to those z such that $\|zy\| \leq \alpha\mu + \beta$. Thus we need only consider those z with $\|z\| \leq \alpha\mu + \beta + |y| =: r_{A,B}$, and Lemma 7.7 assures us that each such z is equal in Γ to at least one the words in the set $Q(a_1, r_{A,B})$. Thus if Step 3 outputs NO, then A is not conjugate to B.

It remains to prove that the running time of Step 3 is bounded by a linear function of $m\mu^2$. Lemma 7.5 tells us that $|y| \leq K_c(|a_1| + |b_1|)$. Therefore

$$r_{A,B} \leq \alpha\mu + \beta + K_c(|a_1| + |b_1|) \leq (\alpha + 2K_c)\mu + \beta.$$

As in Theorem 7.6, it follows that the number of words $z \in Q(a_1, r_{A,B})$is bounded by a linear function of μ, as is the length of these words. Moreover, a list of these words can be generated in time $O(\mu^2)$.

The time it takes **DA** to check whether $(zy)^{-1}a_k(zy) = b_k$ is bounded by a linear function of $(|zy| + |a_k| + |b_k|)$, which in the light of our bounds on $|z|$ and $|y|$, can in turn be bounded by a linear function of μ. As k and z vary, the number of checks

[3] $N = t!$ where $t \leq (2k+1)^{4\delta+1}$ is a bound on the orders of torsion elements

that we must make is bounded by $m\,|Q(a_1, r_{A,B})|$. We have bounded $|Q(a_1, r_{A,B})|$ by a linear function of μ. Hence the running time of Step 3 is $O(m\mu^2)$. □

7.5 A streamlined algorithm for torsion-free groups

Consider the algorithm of the previous section. If one knows that the group is torsion-free, then Step 1(ii) is obviously unnecessary. More importantly, in the torsion-free case one can exploit the fact that virtually cyclic subgroups are actually cyclic, and the fact that the centraliser of a non-elementary subgroup is trivial.

Algorithm 7.9 *The algorithm begins as in Step 1(i) of Algorithm 7.8. It then implements Step 2 of that algorithm before proceeding as follows:*

Step 3′: *Use* **DA** *to check for which i one has $a_1a_i = a_ia_1$ in Γ; stop if some $a_1a_i \neq a_ia_1$ and proceed to Step 5′. If $a_1a_i = a_ia_1$ for all i, proceed to Step 4′.*

Step 4′: *Take y from Step 2 and for each i apply* **DA** *to decide for $i = 2, \ldots, m$ whether $y^{-1}a_iyb_i^{-1} = 1$ in Γ. If there is an i for which the equality is not valid, stop and output NO; otherwise output "YES with conjugator y".*

Step 5′: *Let i be the integer found in Step 3′. Let $r = |y| + \alpha_3\mu + \beta_3$, where $\alpha_3\mu + \beta_3$ is the bound found in Theorem 5.2 on the length of elements conjugating $[a_1, a_i]$ to $[b_1, b_i]$. List $Q(a_1, r)$. For each $z \in Q(a_1, r)$ apply* **DA** *to check if $(zy)^{-1}a_i(zy) = b_i$. If for some z one has $(zy)^{-1}a_i(zy) = b_i$, go to Step 6′. If none of these equalities is valid, output NO.*

Step 6′: *Take zy from Step 5′ and check for $k = 2, \ldots, m$ whether or not $(zy)^{-1}a_k(zy) = b_k$ in Γ. If this equality fails for some k, stop and output NO. If equality holds for all k, output "YES with conjugator zy".*

Theorem 7.10 *Let Γ be a torsion-free group that is δ-hyperbolic with respect to a generating set of cardinality k. Given two finite lists of words in the generators, $A = [a_1, \ldots, a_m]$ and $B = [b_1, \ldots, b_m]$, the above algorithm will determine whether or not A is conjugate to B in Γ.*

The running time of the algorithm is bounded above by a quadratic function of $\|A\| + \|B\|$.

Proof. In a torsion-free hyperbolic group, $a_1a_2 = a_2a_1$ if and only if the subgroup $\langle a_1, a_2\rangle$ is cyclic. Moreover roots are unique, so if $a_1 \neq 1$ and $a_1a_i = a_ia_1$ for $i = 2, \ldots, n$, then $\langle a_1 \ldots, a_n\rangle$ is cyclic. Thus Step 3′ is deciding whether $\langle A\rangle$ is cyclic. If it is cyclic, its centralizer is also cyclic: indeed $C_\Gamma(A) = C_\Gamma(a_1)$. It follows that the lists A and B are conjugate only if y acts as a conjugator — hence the structure of Step 4′.

If $a_1a_i \neq a_ia_1$ then the centralizers of a_1 and a_i intersect trivially, so there is at most one element conjugating a_1 to b_1 and a_i to b_i. Step 5′ will find this conjugating element zy if it exists. Step 6′ then determines whether or not zy conjugates a_k to b_k for the other values of the index k.

The running time of Step 3′ is $O(\|A\|^2)$, while that of Step 4′ is bounded by a constant multiple of $|y|(\|A\| + \|B\|)$, which is $O(\|A\| + \|B\|)^2$. In the light of Theorem

7.6, in Step 5′ one can list $Q(a_1, r)$ in time $O(\|A\|+\|B\|)^2$, because $r \leq |y|+\alpha_2\mu_2+\beta_2 = O(\|A\| + \|B\|)$. Moreover, the time it takes **DA** to check if $(zy)a_i(zy)^{-1} = b_i$ is bounded by a constant multiple of $\sum_{z\in Q}(|z| + |a_i| + |b_i|)$. In the light of Theorem 7.6, each summand is bounded by a linear function of $\|A\| + \|B\|$, as is the number of summands.

Finally, the time it takes **DA** to check the validity of all of the putative equalities in Step 6′ is bounded by a constant multiple of $|zy|(\|A\| + \|B\|)$, which is $O(\|A\| + \|B\|)^2$. □

Arguing as in Corollary 6.2, we have:

Corollary 7.11 *In any torsion-free hyperbolic group Γ, for each integer m, there exists a quadratic-time algorithm to determine conjugacy among the subsets of Γ that have cardinality $\leq m$. (The subsets are given as lists of words in the generators.)*

The appearance of the integer m in this corollary comes from the fact that one must apply the preceding algorithm to all possible descriptions of the subsets as ordered lists, and this introduces a factor of $m!$ to the running time.

8 Application to width of homotopies

Free homotopy classes of maps from a circle into any arc-connected and locally-simply connected space X are in bijection with conjugacy classes of elements in the fundamental group of that space. Similarly, homotopy classes of maps from a fixed compact graph $\mathcal{G}$ are in bijection with conjugacy classes of finite lists of elements in the fundamental group, and homotopies between such maps correspond to conjugacies between the finite lists. In this section we shall examine this correspondence more closely and use it to relate our earlier results to the geometry of homotopies in negatively curved spaces. As we said in the introduction, our motivation comes from [12].

To avoid notational complications, we focus on the case where the graph $\mathcal{G}$ is a rose, i.e. a connected compact graph with one vertex v and finitely many edges $e_1, \ldots, e_n$ (which we regard as maps $e_i : [0, 1] \to \mathcal{G}$). The adaptation to the general case is entirely straightforward.

Fix a basepoint $x_0 \in X$. A continuous map $g : \mathcal{G} \to X$ and a choice of path p from x_0 to $g(v)$ determines a list of elements in $\pi_1(X, x_0)$, namely $[g] = [g_*(e_1), \ldots, g_*(e_n)]$, where $g_*(e_i)$ is the homotopy class of the concatenation $p \cdot (g \circ e_i) \cdot \overline{p}$ (where $\overline{p}$ means p traversed in the opposite direction). A different choice of p leads to a conjugate list. Conversely, one can construct a map homotopic to g from any conjugate of $[g]$ by simply choosing a loop at x_0 representing each element on the list.

Assume that X is a compact geodesic space, e.g. a closed Riemannian manifold. We are interested in *rectifiable* maps $g : \mathcal{G} \to X$, that is maps such that for each edge e_i the path $g \circ e_i$ has finite length, written $l_g(e_i)$. Define $\|g\|_\infty = \max\{l_g(e_1), \ldots, l_g(e_n)\}$.

We are concerned with the lengths of the tracks $h_s : t \mapsto H(s, t)$ of homotopies $H : \mathcal{G} \times [0, 1] \to X$ between rectifiable maps. Thus we define $W(H) = \sup\{l(h_s) \mid s \in \mathcal{G}\}$.

In the following statement, negative curvature is in the sense of A.D. Alexandrov (see [3]). The constant K depends only on the least upper bound on the curvature of

X, the diameter D of X, and the number of elements of $\pi_1 X$ that fail to move an open ball of radius D in $\widetilde{X}$ off itself. K does *not* depend on the number of edges in $\mathcal{G}$.

Theorem 8.1 *If X is a compact, negatively curved, geodesic space, then there is a constant K such that each homotopic pair of rectifiable maps $g_1, g_2 : \mathcal{G} \to X$ is connected by a homotopy $H : \mathcal{G} \times [0,1] \to X$ with $W(H) \leq K(\|g_1\|_\infty + \|g_2\|_\infty + 1)$.*

Proof. We choose a shortest path p_x from the basepoint $x_0 \in X$ to each $x \in X$. This associates to $g : \mathcal{G} \to X$ the homotopic map $\underline{g} : \Gamma \to X$ that sends v to x_0 and sends each e_i to the geodesic loop at x_0 in the (based) homotopy class of $p_{g(v)} g(e_i) \overline{p}_{g(v)}$. Note that $\|\underline{g}\|_\infty \leq \|g\|_\infty + 2D$. Also, because geodesics in X vary continuously and uniquely, $\underline{g} \circ e_i$ and $g \circ e_i$ are homotopic via the homotopy with locally-geodesic tracks whose restriction to $\{v\} \times [0,1]$ is $p_{g(v)}$. This homotopy has width $W(H) \leq D + \frac{1}{2} \max_i \{l_g(e_i) + l_{\underline{g}}(e_i)\}$.

In the light of these observations, given a pair of homotopic rectifiable maps $g_1, g_2 : \mathcal{G} \to X$, we are free to replace them with $\underline{g}_1, \underline{g}_2 : \mathcal{G} \to X$. Our next goal is to relate $\|\underline{g}_i\|_\infty$ to the length of the elements in the lists $[\underline{g}_i] = [g_i]$ (we'll then use Theorem A).

The set S of elements $\gamma \in \pi_1(X, x_0)$ represented by geodesics c_γ of length less than $2D$ based at x_0 is finite and generates the group. Consider the function $\gamma \mapsto \|\gamma\|$ measuring distance from the identity in the word metric defined by S. According to the Švarc-Milnor Lemma, the functions $\gamma \mapsto \|\gamma\|$ and $\gamma \mapsto l(c_\gamma)$ are quasi-Lipschitz, i.e. there exist constants $\lambda \geq 1$ and $\varepsilon \geq 0$, depending only on D, such that

$$\frac{1}{\lambda} l(c_\gamma) - \varepsilon \leq \|\gamma\| \leq \lambda l(c_\gamma) + \varepsilon.$$

(This is almost the same as saying that $\pi_1(X, x_0)$ with the word metric is quasi-isometric to the universal covering $\widetilde{X}$; see [3], page 140.)

Recall that, given a list of group elements $A = [a_1, \ldots, a_m]$ and a word metric, we write $\|A\| = \max \|a_i\|$. The above inequality provides a quasi-Lipschitz relation between $\|\underline{g}\|_\infty$ and $\|[g]\|$ for every $g : \mathcal{G} \to X$. Since X is negatively curved, $\pi_1(X, x_0)$ is δ-hyperbolic, where δ depends only on the upper curvature bound and the quasi-isometry constants λ and ε. Thus Theorem A assures us that if g_1 and g_2 are homotopic, then there will be an element $\gamma \in \pi_1(X, x_0)$, with $\gamma^{-1}[g_1]\gamma = [g_2]$, whose length is bounded by a linear function of $\max\{\|\underline{g}_1\|_\infty, \|\underline{g}_2\|_\infty\}$, and the coefficients of this linear function will depend only on $\delta, |S|$ and D (in the guise of λ and ε). It follows that the geodesic c_γ based at x_0 also has length bounded by such a linear function. And, as at the end of the first paragraph, it follows $\underline{g}_1$ and $\underline{g}_2$ are homotopic via a homotopy whose tracks are also bounded by such a linear function. □

Remark. The above argument made rather mild use of the curvature hypothesis and as a result can be adapted to more general spaces with δ-hyperbolic fundamental group.

The above theorem should be compared with [12, Theorem 0.1], which, in the Riemannian setting, gives a linear bound for the width of H in terms of the *sum* of the lengths of the images under f and g of all the edges of the graph. It is important to note, however, that the constants involved in the linear bound in [12] are small, whereas the constant in our theorem will be very large. Moreover, the result in [12]

can be pushed beyond the hyperbolic world [13, Theorem 5.1]. Here one relaxes the curvature condition on the manifold to 'nowhere positive' – in this case the bound obtained depends on the homotopy class of g_1 and g_2.

Deborah Ruoss of Zürich (private communication) has generalised the main result of [12] to the case of closed Riemannian manifolds with δ-hyperbolic fundamental group.

A Conjugacy for lists in non-hyperbolic groups

There is a general acceptance amongst group theorists that any conceivable variation on the theme of decidability can be realised by a suitably cunning construction of a finitely presented group. The purpose of this appendix is to record such a construction.

Theorem A.1 *There exist finitely presented groups in which the conjugacy problem for elements is soluble, but the conjugacy problem for finite lists is not.*

Our earlier results show that the conjugacy problem for finite subsets of hyperbolic groups differs little in complexity from the conjugacy problem for individual elements. Theorem A.1 shows that this property is truly attributable to the hyperbolicity of the group: in a non-hyperbolic group these two conjugacy problems can have utterly different levels of complexity.

Our construction is based on a remarkable phenomenon discovered by Collins and Miller [6]:

Theorem A.2 (Collins-Miller) *There exists a finitely presented group G and a subgroup $H \subset G$ of index two such that G has a soluble conjugacy problem but H does not.*

A.1 A Fibre-Product Construction

We'll use a fibre-product construction to meld $G \to G/H$ with a group F that admits an epimorphism $\phi : F \to \mathbb{Z}_2$ with the following properties:

(1) there is a 2-element subset U of $K := \ker \phi$ such that the centraliser $C_F(U)$ of U in F is trivial;

(2) for every $f \in F$, there exists $z_f \in C_F(f)$ such that $\phi(z_f) \neq 0$.

Example A.3 *Let $F = \mathbb{Z}^3 \rtimes \mathbb{Z}_2^2$, where the action of a basis $\{a_1, a_2\}$ for $\mathbb{Z}_2^2$ on a basis $\{u_0, u_1, u_2\}$ for $\mathbb{Z}^3$ is given by $a_i u_i a_i = u_i$ and $a_i u_j a_i = -u_j$ if $i \neq j$. Define $\phi : F \to \mathbb{Z}_2$ by $\phi(u_j) = \phi(a_i) \neq 0$ for $i = 1, 2$ and $j = 1, 2, 3$.*

Lemma A.4 *F and ϕ satisfy conditions* (1) *and* (2)*.*

Proof. An easy calculation shows that an element $u_0^\alpha u_1^\beta u_2^\gamma a_1^\delta a_2^\varepsilon \in F$ commutes with $u_1 a_1$ if and only if $\alpha = \gamma = \varepsilon = 0$. Thus $C_F(u_1 a_1) \cong \mathbb{Z}_2 \times \mathbb{Z}$, generated by a_1 and u_1. Similarly, $C_F(u_2 a_2)$ is generated by a_2 and u_2. Hence $C_F(\{u_1 a_1, u_2 a_2\}) = C_F(u_1 a_1) \cap C_F(u_2 a_2)$ is trivial. This proves (1).

To prove (2), note that, for $v \in \mathbb{Z}^3$ and $b \in \mathbb{Z}_2^2$, vb commutes with u_i if $b = a_i$ $(i = 1, 2)$, and with u_0 otherwise. □

Example A.5 *Let G, H and F be as above and define Γ to be the fibre product of the maps $\pi : G \to G/H$ and $\phi : F \to \mathbb{Z}_2$. Thus $\Gamma \subset G \times F$ is the subgroup generated by*

$$H \times \{1\} \bigcup \{(\gamma, a_1)\} \bigcup \{1\} \times K.$$

for a choice of $\gamma \notin H$.

We may now apply our condition (2) to the conjugacy problem for Γ.

Lemma A.6 *If $(g_1, f_1), (g_2, f_2) \in \Gamma$ are conjugate in $G \times F$, then they are conjugate in Γ.*

Proof. Suppose (s, t) conjugates (g_1, f_1) to (g_2, f_2) in $G \times F$. Now, $(s, t) \in \Gamma$ if and only if $\pi(s) = \phi(t)$. Should this not be the case, then condition (2) allows us to replace t by τ so that $\phi(\tau) = \pi(s)$ and $\tau^{-1} f_1 \tau = t^{-1} f_1 t = f_2$. □

Corollary A.7 *Γ has a soluble conjugacy problem.*

Conversely, condition (1) provides an obstruction to the solubility of the conjugacy problem for lists.

Lemma A.8 *Let $h_1, h_2 \in H$. Then the lists*

$$\underline{\gamma}_1 := [(h_1, u_1 a_1), (h_1, u_2 a_2), (h_1, u_0^2), \ldots, (h_1, u_0^{2m-4})],$$

$$\underline{\gamma}_2 := [(h_2, u_1 a_1), (h_2, u_2 a_2), (h_2, u_0^2), \ldots, (h_2, u_0^{2m-4})]$$

of elements of Γ are conjugate by an element of Γ if and only if h_1 and h_2 are conjugate in H.

Proof. Certainly, the elements in $\underline{\gamma}_1$ and $\underline{\gamma}_2$ belong to Γ since $h_i \in H$ and $u_i a_i, u_0^2 \in K$. Moreover, if $h \in H$ with $h^{-1} h_1 h = h_2$, then $(h, 1)^{-1} \underline{\gamma}_1 (h, 1) = \underline{\gamma}_2$.

Conversely, if $(g, f) \in \Gamma$ with $(g, f)^{-1} \underline{\gamma}_1 (g, f) = \underline{\gamma}_2$, then $f = 1$, so $g \in H$ and h_1, h_2 are conjugate in H. □

Corollary A.9 *The conjugacy problems for lists of length $m \geq 2$ and for sets of size $m \geq 2$ in Γ are not soluble.*

Proof. By Lemma A.8, a solution to the conjugacy problem for lists of length $m \geq 2$ would yield a solution to the conjugacy problem for H, contrary to hypothesis.

Moreover, the lists $\underline{\gamma}_1$ and $\underline{\gamma}_2$ of Lemma A.8 are conjugate if and only if the corresponding sets are, since the F-components of the elements in $\underline{\gamma}_1$ are pairwise non-conjugate in F. Hence the conjugacy problem for finite subsets (of cardinality $m \geq 2$) in Γ is also not soluble. □

References

[1] I. Anshel, M. Anshel and D. Goldfeld, *An algebraic method for public-key cryptography*, Math. Res. Lett. **6** (1999), 287–291.

[2] G. Baumslag, C. F. Miller III and H. Short, *Unsolvable problems about small cancellation and word hyperbolic groups*, Bull. London Math. Soc. **26** (1994), 97–101.

[3] M. R. Bridson and A. Haefliger, "Metric Spaces of Non-Positive Curvature", Grund. Math. Wiss. 319, Springer-Verlag, Berlin-Heidelberg-New York, 1999.

[4] M. R. Bridson and L. Reeves, *On the isomorphism problem for biautomatic groups*, preprint, 2003.

[5] T. Delzant, *Sous-groupes distingués et quotients des groups hyperboliques*, Duke Math. J. **83** (1996), 661-682.

[6] D. J. Collins and C. F. Miller III, *The conjugacy problem and subgroups of finite index*, Proc. London Math. Soc. **34** (1977), 535–556.

[7] D. B. A. Epstein and D. F. Holt, *The linearity of the conjugacy problem in word hyperbolic groups*, preprint.

[8] S. M. Gersten and H. B. Short, *Rational subgroups of biautomatic groups*, Ann. of Math. **134** (1991), 125–158.

[9] M. Gromov, *Hyperbolic groups*, in "Essays on Group Theory" (S.M. Gersten ed.), MSRI Publ. 8, Springer-Verlag, 1987, pp. 75–263.

[10] P. de la Harpe, "Topics in Geometric Group Theory", University of Chicago Press, 2000.

[11] D. F. Holt, *Word-hyperbolic groups have real-time word problem.* Internat. J. Algebra Comput. **10** (2000), 221–227.

[12] T. Kappeler, S. Kuksin and V. Schroeder, *On the Poincaré inequality for maps into closed manifolds of negative sectional curvature*, Preprint 21-2002, University of Zürich (revised version, 2003).
`http://www.math.unizh.ch/fileadmin/math/preprints/21-02rev.pdf`

[13] T. Kappeler, S. Kuksin and V. Schroeder, *Perturbations of the harmonic map equation*, Communications in Contemporary Mathematics **5** (2003), 629–669.

[14] G. Levitt and K. Vogtmann, *A Whitehead algorithm for surface groups*, Topology **39** (2000), 1239–1251.

[15] P. Papasoglu, *An algorithm detecting hyperbolicity*, Geometric and computational perspectives on infinite groups (Minneapolis, MN and New Brunswick, NJ, 1994), Amer. Math. Soc., Providence, RI, 1996, pp. 193–200.

[16] E. Rips and Z. Sela, *Canonical representatives and equations in hyperbolic groups*, Invent. Math. **120** (1995) 489–512.

[17] H. Short (ed.), *Notes on word hyperbolic groups* in "Group Theory from a Geometrical Viewpoint", edited by E. Ghys, A. Haefliger and A. Verjovsky, World Scientific 1991, pp. 3–63.

[18] V. Shpilrain, *Assessing security of some group based cryptosystems*, arXiv:math.GR/0311047, 4 Nov 2003.