

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/161702>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

© 2022 Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

The compressed word problem in relatively hyperbolic groups

Derek Holt and Sarah Rees

Dedicated to Patrick Dehornoy, our collaborator and friend, from whom we learned a lot

Abstract

We prove that the compressed word problem in a group that is hyperbolic relative to a collection of free abelian subgroups is solvable in polynomial time.

1. Introduction

The main result of [8] is that the compressed word problem in a hyperbolic group is solvable in polynomial time. Here we generalise this result to a group hyperbolic relative to a set of free abelian subgroups.

Theorem A. *The compressed word problem for a group that is hyperbolic relative to a collection of free abelian subgroups is solvable in polynomial time.*

We prove the theorem by extending the arguments of [8] from hyperbolic to relatively hyperbolic groups. Our principal source for results about this class of groups is [2]. In particular, we use the automaticity of groups of that type, proved in [2, Theorem 7.7]; however we need to construct a new asynchronously automatic structure, with particular properties that we need.

We believe that our result remains true if we substitute arbitrary (finitely generated) abelian groups for free abelian groups in the theorem statement but (as we shall explain after the definition of the *components* of a word in Section 3) extending the proof in that way would result in further technical difficulties in a proof that is already highly technical, so we decided not to attempt it here. It might be more interesting to try to extend the result to groups hyperbolic with respect to a collection of virtually abelian subgroups, but we are not currently able to extend our methods to cover that case in general.

We introduce basic concepts and notation, including the definitions of straight line programs (SLP) and the compressed word problem, in Section 2, which follows this introduction. Section 3 contains the definition of relatively hyperbolic

groups (following Osin [12]) and properties of those that we shall need in the article. Section 4 provides basic background material on SLPs, Section 5 gives some results for SLPs associated with finitely generated abelian groups, Section 6 examines the geometry of the compressed Cayley graph $\widehat{\Gamma}$ of a relatively hyperbolic group G , and Section 7 establishes some results for SLPs associated with such a group. The main result of the article, Theorem A, is proved across the final two sections, Sections 8 and 9, as the concatenations of two theorems, Theorem 9.1 and Theorem 8.1.

The authors would like to thank Saul Schleimer, for introducing us to the study of the compressed word problem, especially for hyperbolic groups [8], and Yago Antolin, for some very helpful discussions on the properties of relatively hyperbolic groups and their automatic structures. We are also grateful to an anonymous referee for a careful reading of the paper and several helpful comments and suggestions.

2. Definitions and notation

To a large extent (but not entirely) our notation and definitions follow [8].

2.1. Words

For a finite set Σ (which we call an *alphabet*), we define a *word* w over Σ to be a string $x_0 \cdots x_{n-1}$, where each x_i is in Σ . We denote by $|w|$ the length n of w , by ϵ the empty word (of length 0), and for $0 \leq i < j \leq n$ we denote by $w[i : j]$ the substring $x_i \cdots x_{j-1}$, which we call a *subword* of w (in [8], and elsewhere, such a substring is called a *factor*). We abbreviate the prefix $w[0 : j]$ of w as $w[: j]$, its suffix $w[i : n]$ as $w[i :)$, and $w[i : i + 1) = x_i$ as $w[i]$, and also consider ϵ to be a subword of w .

For words $v, w \in \Sigma^*$, we write $v = w$ if v and w are equal as strings and, when Σ is a generating set of a group G , we write $v =_G w$ if v and w represent the same element of G . All group generating sets Σ in this article will be assumed to be inverse closed (that is, $x \in \Sigma \Rightarrow x^{-1} \in \Sigma$).

Suppose that Σ is an ordered finite generating set for a group G and that $w \in \Sigma^*$ represents an element $g \in G$. Then we define $\text{slex}(w)$ (or $\text{slex}_\Sigma(w)$) to be the shortlex minimal word representing g ; that is, $\text{slex}(w)$ is the lexicographically least among the shortest words that represent the same group element as w .

2.2. Straight-line programs

Let Σ be a finite alphabet and V a finite set with $V \cap \Sigma = \emptyset$. Let $\rho : V \rightarrow (V \cup \Sigma)^*$ be a map and extend the definition of ρ to $(V \cup \Sigma)^*$ by defining $\rho(a) = a$ for

all $a \in \Sigma \cup \{\epsilon\}$ and $\rho(uv) = \rho(u)\rho(v)$ for all $u, v \in (V \cup \Sigma)^*$. We define the associated binary relation \succeq on V by $A \succeq B$ whenever the symbol B occurs within the string $\rho^k(A)$, for some $k \geq 0$.

We define a *straight-line program* (SLP for short) over an alphabet Σ to be a triple $\mathcal{G} = (V, S, \rho)$, with $S \in V$ and $\rho : V \rightarrow (V \cup \Sigma)^*$ a map such that the associated binary relation \succeq on V is acyclic, that is the corresponding directed graph contains no directed cycles. The set V is called the set of *variables* of \mathcal{G} , and S is called the *start variable*. Where necessary, we write $V_{\mathcal{G}}, S_{\mathcal{G}}, \rho_{\mathcal{G}}$, rather than simply V, S, ρ .

An SLP \mathcal{G} is naturally associated with a context-free grammar (V, Σ, S, P) , where P is the set of all productions $A \rightarrow \rho(A)$ with $A \in V$, and we will often use the name \mathcal{G} also for this grammar. It follows from the definition of an SLP that this associated grammar derives exactly one terminal word, which we call the *value* of \mathcal{G} and denote by $\text{val}(\mathcal{G})$.

SLPs are used to provide succinct representations of words that contain many repeated substrings. For instance, the word $(ab)^{2^n}$ is the value of the SLP $\mathcal{G} = (\{A_0, \dots, A_n\}, \rho, A_0)$ with $\rho(A_n) = ab$ and $\rho(A_{i-1}) = A_i A_i$ for $0 < i \leq n$. We provide more background on SLPs in Section 4.

2.3. The compressed word problem

The *compressed word problem* for a finitely generated group G with the finite symmetric generating set Σ is the following decision problem:

Input: an SLP \mathcal{G} over the alphabet Σ .

Question: does $\text{val}(\mathcal{G})$ represent the group identity of G ?

It is an easy observation that the computational complexity of the compressed word problem for G is independent of the choice of generating set Σ ; more precisely, if Σ' is another finite symmetric generating set for G , then the compressed word problem for G with respect to Σ is log-space reducible to the compressed word problem for G with respect to Σ' [11, Lemma 4.2]. So, when proving that the compressed word problem for G is solvable in polynomial time, we are free to choose whichever finite symmetric generating set of G is most convenient for the purpose.

2.4. Fellow travelling and automatic groups

Suppose that $\Gamma = \Gamma(G, \Sigma)$ is the Cayley graph for a group G with finite symmetric generating set Σ . Let v, w be words over Σ , and let γ_v, γ_w be the paths traced out in Γ by v, w from the identity vertex 1 of Γ . For $0 \leq i < |v|$, we

denote the vertex of Γ labelled $v[i]$ by $\gamma_v[i]$, and similarly for $w[i]$ and $\gamma_w[i]$. (So, in particular, $\gamma_v[0] = \gamma_w[0]$ is the identity vertex.) For the following definition, we define $\gamma_v[i] := \gamma_v[|v| - 1]$ for integers $i \geq |v|$.

We say that the words v and w *fellow travel* at distance k (or, more briefly, *k-fellow travel*) if $d_\Gamma(\gamma_v[i], \gamma_w[i]) \leq k$ for all $i \geq 0$. In other words, the distance in Γ between the vertices on γ_v and γ_w at the ends of subpaths traced out by the subwords $v[:i]$ and $w[:i]$ of v and w is at most k . In this situation we also say that the paths γ_v and γ_w *k-fellow travel*.

We can extend this terminology to paths γ_v, γ_w that do not start at the same vertex in Γ . In particular, if two such paths *k-fellow travel*, then their start points and also their end points are at distance at most k from each other.

We say that the words v and w *asynchronously fellow travel* at distance k if we can choose sequences $(i_0 = 0, i_1, \dots, i_n = |v| - 1)$ and $(j_0 = 0, j_1, \dots, j_n = |w| - 1)$, with $i_{t+1} \in \{i_t, i_t + 1\}$, $j_{t+1} \in \{j_t, j_t + 1\}$ for each t , such that $d_\Gamma(\gamma_v[i_t], \gamma_w[j_t]) \leq k$ for each t with $0 \leq t \leq n$, and again we may also apply this terminology to the paths that they trace out in Γ .

We call the pairs of vertices $\gamma_v[i_t], \gamma_w[j_t]$ *corresponding vertices* in the fellow travelling. Note that a vertex on one of the paths can have more than one corresponding vertex on the other path, but in that case the corresponding vertices are the vertices lying on a contiguous subpath of the other path.

We say that G is *automatic* if there exists a finite state automaton A over Σ for which every element of G has at least one representative word in $L(A)$, as well as an integer k such that, whenever $v, w \in L(A)$ and either $v =_G w$ or $v =_G wx$ with $x \in \Sigma$, then v, w fellow travel at distance k . We say that G is *asynchronously automatic* if the same is true but with an asynchronous fellow traveller property, and *(asynchronously) biautomatic* if, in addition, for words u, w that satisfy $xu =_G w$ with $x \in \Sigma$, that paths traced out by u, w from the vertices x and 1 , respectively, (asynchronously) fellow travel at distance k .

We call the pair (A, k) an *automatic structure* (or *asynchronously automatic structure*) for G . An automatic structure (A, k) is called *geodesic* or *shortlex* if each word in $L(A)$ is a representative of minimal length, or minimal within the shortlex word ordering, of the element that it represents, respectively. We say that (A, k) is a *structure with uniqueness* if it contains a unique representative of each element of G . We refer to [5] for basic properties of automatic structures and automatic groups.

3. Relatively hyperbolic groups

The purpose of this section is to give the definition of a relatively hyperbolic group and to list the properties that we shall need in this article. The properties

we need are proved in the article [2], and build on results of [12]. We have used Osin's definition of relatively hyperbolicity; it is proved in [12, Theorem 1.5] that (for finitely generated groups, as in our case) this is equivalent to the definition of [3], also to the definition of [6] combined with the Coset Penetration Property (see below), called strong relative hyperbolicity in [6]. Below we have (essentially) used notation and statements from [2].

We suppose that Σ is a finite generating set for a group G , and that $\{H_i : i \in \Omega\}$ is a finite collection of subgroups of G , which we call the collection of *parabolic subgroups* of G . Define $\mathcal{H} := \bigcup_{i \in \Omega} (H_i \setminus \{1\})$, and $\widehat{\Sigma} := \Sigma \cup \mathcal{H}$. We let $\Gamma := \Gamma(G, \Sigma)$ and $\widehat{\Gamma} = \widehat{\Gamma}(G, \widehat{\Sigma})$ be the Cayley graphs for G over Σ and $\widehat{\Sigma}$, respectively. (So $\widehat{\Gamma}$ has the same vertices as Γ but more edges than Γ .) We call a word over Σ (or $\widehat{\Sigma}$) *geodesic* if it labels a geodesic path in Γ (or $\widehat{\Gamma}$).

Following [2, Definition 2.5] and [12, Section 1.2], we define F to be the free product of groups

$$F := (*_{i \in \Omega} H_i) * F(\Sigma)$$

and suppose that a finite subset R of F exists whose normal closure in F is the kernel of the natural map from F to G ; in that case we say that G has the *finite presentation*

$$\left\langle X \cup \bigcup_{i \in \Omega} H_i \mid R \right\rangle$$

relative to the collection of subgroups $\{H_i : i \in \Omega\}$. Now if u is a word over $\widehat{\Sigma}$ that represents the identity in G , then u is equal within F to a product of the form

$$\prod_{j=1}^n f_j r_j^{\eta_j} f_j^{-1},$$

with $r_j \in R$, $f_j \in F$ and $\eta_j = \pm 1$ for each j . The smallest possible value of n in any such expression of this type for u is called the *relative area* of u , denoted by $\text{Area}_{\text{rel}}(u)$.

We say that G is *hyperbolic relative to* the collection of subgroups $\{H_i\}$ if it has a finite relative presentation as above and a constant $C \geq 0$ such that

$$\text{Area}_{\text{rel}}(u) \leq C|u|$$

for all words u over $\widehat{\Sigma}$ that represent the identity in G .

We note that if G is relatively hyperbolic then the graph $\widehat{\Gamma}$ is Gromov-hyperbolic. Note also that, by [12, Proposition 2.36], the intersection $H_i \cap H_j$ for $i \neq j$ is finite.

The notation and results that follow are all taken from [2]. Given a path p in $\widehat{\Gamma}$, we say that the path p *penetrates* the left coset gH_i if p contains an edge labelled by an element of H_i that connects two vertices of gH_i . An H_i -*component* of

such a path is defined to be a non-empty maximal subpath of p that is labelled by a word in H_i^* . Two components s and r (not necessarily of the same path) are *connected* if both are H_i -components for some H_i , and if the start points of both paths lie in the same left coset gH_i of H_i .

A path p is said to *backtrack* if $p = p'sr's'p''$ where s, s' are H_i -components, and the word labelling r represents an element of H_i ; if no such decomposition of p exists, then p is *without backtracking*. A path p is said to *vertex backtrack* if it contains a subpath of length greater than 1 labelled by a word that represents an element of some H_i ; otherwise p is said to be *without vertex backtracking*. We note that if a path does not vertex backtrack then it does not backtrack and all of its components have length 1.

We denote the start and end points of a path p in $\widehat{\Gamma}$ by p_- and p_+ , respectively, and say that paths p, q in $\widehat{\Gamma}$ are *k-similar* if $\max\{d_\Gamma(p_-, q_-), d_\Gamma(p_+, q_+)\} \leq k$. The following fundamental result about k -similar paths in $\widehat{\Gamma}$, proved as [12, Theorem 3.23], is also stated as [2, Theorem 2.8].

Proposition 3.1. [12, Theorem 3.23]. (*Bounded Coset Penetration Property*) *Let G be relatively hyperbolic, as above. For any $\lambda \geq 1$, $c \geq 0$, $k \geq 0$, there exists a constant $\mathbf{e} = \mathbf{e}(\lambda, c, k)$ such that, for any two k -similar paths p and q in $\widehat{\Gamma}$ that are (λ, c) -quasigeodesics and do not backtrack, the following conditions hold.*

- (1) *The sets of vertices of p and q are contained in the closed \mathbf{e} -neighbourhoods of each other in Γ .*
- (2) *Suppose that, for some i , s is an H_i -component of p with $d_\Gamma(s_-, s_+) > \mathbf{e}$; then there exists an H_i -component of q that is connected to s .*
- (3) *Suppose that s and t are connected H_i -components of p and q , respectively. Then s and t are \mathbf{e} -similar.*

The next three results are derived in [2] from the Bounded Coset Penetration Property.

We define the *components* of a word $w \in \Sigma^*$ to be the nonempty subwords of w of maximal length that lie in $(\Sigma \cap H_i)^*$ for some parabolic subgroup H_i . In general, since $H_i \cap H_j$ is finite for $i \neq j$, it is possible for the end of one component to overlap the beginning of the next, where the overlapping generators lie in a finite intersection. In this paper, we shall generally be assuming that $H_i \cap H_j$ is trivial for $i \neq j$, in which case the components are necessarily disjoint. This holds in particular when the H_i are free abelian groups, and it is the main reason why we have not attempted to generalise our main theorem to groups that are hyperbolic relative to arbitrary finitely generated abelian groups. We strongly believe that such a generalisation would be possible, but it might involve significant additional technicalities, of a similar nature to those involved in [2].

Let $w := \alpha_0 u_1 \alpha_1 u_2 \cdots u_n \alpha_n$, where the u_j are its components. Then, following [2, Construction 4.1], we define the *derived word* $\hat{w} := \alpha_0 h_1 \alpha_1 h_2 \cdots h_n \alpha_n \in \widehat{\Sigma}^*$, where each h_j is the element of a parabolic subgroup represented by u_j . So the components of paths in Γ and $\widehat{\Gamma}$ labelled by w and \hat{w} are labelled by the subwords u_i and h_i of w and \hat{w} , respectively.

A word over Σ is said to have a *parabolic shortening* if, for some i , it has a component over $\Sigma \cap H_i$ that is non-geodesic; otherwise it has *no parabolic shortenings*.

Proposition 3.2. [2, Lemma 5.3, Theorems 7.6, 7.7] *Let G be a finitely generated group, hyperbolic with respect to a family of subgroups $\{H_i\}_{i \in \Omega}$, and let Σ' be a finite generating set of G .*

Then there exist $\lambda \geq 1$, $c \geq 0$ and a finite subset \mathcal{H}' of \mathcal{H} such that, for every finite, ordered generating set Σ of G with $\Sigma' \cup \mathcal{H}' \subseteq \Sigma \subseteq \Sigma' \cup \mathcal{H}$ for which each H_i has a geodesic biautomatic structure over $\Sigma \cap H_i$, we have:

- (i) *G has a geodesic biautomatic structure over Σ which is a shortlex structure if the structures on H_i are shortlex;*
- (ii) *there exists a finite set Φ of non-geodesic words over Σ such that, for each word $w \in \Sigma^*$ with no parabolic shortenings and no subwords in Φ , the word $\hat{w} \in (\Sigma \cup \mathcal{H})^*$ is a (λ, c) -quasigeodesic without vertex backtracking.*

For the remainder of this paper, given a relatively hyperbolic group G and its finite generating set Σ , we shall say that (G, Σ) is *suitable for parabolic geodesic biautomaticity* if

- (i) $H_i \cap H_j = \{1\}$ for all $i \neq j$;
- (ii) each parabolic subgroup H_i has a geodesic biautomatic structure over $\Sigma \cap H_i$, and there exist λ and c such that the conclusions (i) and (ii) of Proposition 3.2 hold for G and Σ .

So this applies in particular when the parabolic subgroups H_i are free abelian, as in the hypothesis of our main result.

We are now in a position to state and prove a corollary to the above proposition.

Corollary 3.3. *Let (G, Σ) be relatively hyperbolic and suitable for parabolic geodesic biautomaticity. If $w \in \Sigma^*$ is a geodesic word that represents an element of H_i for some i , then $w \in (\Sigma \cap H_i)^*$. In particular, we have $H_i = \langle \Sigma \cap H_i \rangle$ for each i .*

Proof. Since w is geodesic, it cannot contain subwords in Φ or have parabolic shortenings. The fact that \hat{w} is without vertex backtracking implies that $|\hat{w}| = 1$ and hence that w is a word over $\Sigma \cap H_i$. Since every element of H_i can be represented by some geodesic word w , H_i is generated by $\Sigma \cap H_i$. \square

Lemma 3.4. *Let (G, Σ) be relatively hyperbolic and suitable for parabolic geodesic biautomaticity, and let $w \in \Sigma^*$ be a geodesic word. Then there exists a constant $\lambda \geq 1$ such that \hat{w} labels a $(\lambda, 0)$ -quasigeodesic path in $\hat{\Gamma}$ that does not vertex backtrack.*

Proof. It follows from Proposition 3.2 that \hat{w} labels a (λ, c) -quasigeodesic path for some $\lambda \geq 1$ and $c \geq 0$ and, since w is geodesic, it cannot represent 1_G unless $w = \epsilon$, and so by increasing λ if necessary we may assume that $c = 0$. \square

Unfortunately, the biautomatic structure for G that is given by Proposition 3.2 does not appear to have all of the properties that we need in the proofs of our main results. For that we need a structure in which \hat{w} is a geodesic for words w in the language. Since no such structure appears in the literature, we need to establish its existence here. It turns out that this structure could be asynchronous, but that will be adequate for our purposes.

Proposition 3.5. *Let (G, Σ) be relatively hyperbolic and suitable for parabolic geodesic biautomaticity, and let $\Gamma := \Gamma(G, \Sigma)$, $\hat{\Gamma} := \hat{\Gamma}(G, \hat{\Sigma})$. Let u, v, w_1, w_2 be words over Σ satisfying $w_1 u =_G v w_2$, with $|w_1|, |w_2| \leq k$ for some $k \geq 0$ and, in quadrilaterals in Γ and $\hat{\Gamma}$ whose sides are labelled by the words in that equation, let p and \hat{p} be paths (in Γ and $\hat{\Gamma}$) labelled by u and \hat{u} , and let q and \hat{q} be the paths labelled by v and \hat{v} .*

Suppose that

- (a) for each parabolic subgroup H_i , all H_i -components of both u and v lie in the specified geodesic biautomatic structure;
- (b) for some $\lambda \geq 1$ and $c \geq 0$, the paths \hat{p} and \hat{q} are (λ, c) -quasigeodesics that do not backtrack.

Then

- (i) there is a constant $e' = e'(\lambda, c, k)$ such that the paths p and q e' -fellow travel in Γ , in such a way that those vertices of p that are also vertices of \hat{p} have at least one corresponding vertex on q that is also a vertex of \hat{q} , and vice versa;
- (ii) there is a constant $k' = k'(\lambda, c, k)$ such that, whenever two vertices b_1 and b_2 on q (or p) are both at Γ -distance at most e' from the same vertex on p (or q), then the distance in the path q (or p) between b_1 and b_2 is at most k' ;
- (iii) we have $|u| \leq (k' + 1)|v|$ and $|v| \leq (k' + 1)|u|$.

Proof. Applying Proposition 3.1 to the paths \hat{p} and \hat{q} , we choose $e_1 \geq e(\lambda, c, k)$ such that also $e_1 \geq k$, and then choose $e_2 \geq e(\lambda, c, e_1)$ such that $e_2 \geq e_1$. For a subword $u[i : j]$ of u , we denote by $p[i : j]$ the subpath of p labelled by $u[i : j]$. So, if $u[i : j]$ is a component of u , then $p[i : j]$ is a component of p .




Figure 1: Fellow travelling quasigeodesics

Let $u[i_1 : j_1], u[i_2 : j_2], \dots, u[i_t : j_t]$ with $i_1 < i_2 < \dots < i_t$ be the components of u of length greater than e_2 , and let the corresponding e_1 -similar components of v to which they are connected be $v[k_1 : l_1], v[k_2 : l_2], \dots, v[k_t : l_t]$. (Since \hat{v} does not backtrack, these components are uniquely defined.)

We claim that $k_1 < k_2 < \dots < k_t$; that is, the connected components on v occur in the same order as their corresponding components on u . We shall prove by induction on r that $k_1 < k_2 < \dots < k_r$ for all $1 \leq r \leq t$. It is vacuously true for $i = 1$, so suppose that it is true for some value of r . Since $e_1 \geq k$, the subpaths $\widehat{p[j_r : \cdot]}$ and $\widehat{q[l_r : \cdot]}$ of \hat{p} and \hat{q} are e_1 -similar (λ, c) -quasigeodesics that do not backtrack and, by choice of e_2 , the component $u[i_{r+1} : j_{r+1}]$ of u is connected to a component of $v[l_r : \cdot]$. But since \hat{v} does not backtrack, this component must be $v[k_{r+1} : l_{r+1}]$, so $k_r < l_r \leq k_{r+1}$, which completes the induction and establishes the claim. See Figure 1.

Since the component subwords of u and v lie in a geodesic biautomatic structure, the paths $\widehat{p[i_r : j_r]}$ and $\widehat{q[i_r : j_r]}$ must $e_1 L_1$ -fellow travel for $1 \leq r \leq t$, where L_1 is the maximum of the fellow travelling constants for the biautomatic structures of the parabolic subgroups H_i .

On the other hand, the subpaths $\widehat{p'}$ and $\widehat{q'}$ of \hat{p} and \hat{q} , where $p' := p[j_{r-1} : i_r]$ and $q' := q[l_{r-1} : k_r]$, are e_1 -similar for $1 \leq r \leq t + 1$ (where for convenience we define $j_0 = k_0 = 0$, $i_{t+1} = |u|$ and $k_{t+1} = |v|$), so they are contained in closed e_2 -neighborhoods of each other. Now the components of $u[j_{r-1} : i_r]$ have length at most e_2 , and those of $v[l_{r-1} : k_r]$ have length at most $e'_2 := e_2 + 2e_1$ (since otherwise they would be connected to a component of $u[j_{r-1} : i_r]$, which would have length greater than e_2). So the paths $\widehat{p[j_{r-1} : i_r]}$ and $\widehat{q[l_{r-1} : k_r]}$ are $(\lambda e'_2, ce'_2)$ -quasigeodesics that lie within $(e_2 + e'_2)$ -neighborhoods of each other. By the argument used in the proof of [9, Proposition 3.1], it follows that these paths L_2 -fellow travel for some constant L_2 that depends only on λ , k and c . Furthermore, by increasing L_2 by at most e_2 we can ensure that all vertices on either of these paths that are also vertices of \hat{p} or \hat{q} have at least one corresponding vertex with the same property. So we have proved (i) with $e' = \max(e_1 L_1, L_2)$.

Now let b_1 and b_2 be vertices on q as in (ii), and assume that b_2 is further along

q than b_1 . Then $d_\Gamma(b_1, b_2) \leq 2e'$. Let $v' \in \Sigma^*$ be a geodesic word with $v' =_G v$, and let q' and \widehat{q}' be the paths in Γ and $\widehat{\Gamma}$ labelled by v' and \widehat{v}' , respectively. Then \widehat{q}' is quasigeodesic by Lemma 3.4, and by applying (i) to v and v' with suitable constants, we find that v and v' L_3 -fellow travel for some constant L_3 . Let b'_1 and b'_2 be vertices of q that correspond to b_1 and b_2 in the fellow-travelling. Then $d_\Gamma(b'_1, b'_2) = d_{q'}(b'_1, b'_2) \leq 2(e' + L_3)$. Now any two vertices c_1 and c_2 of q that lie between b_1 and b_2 are each at distance at most L_3 from some vertex of q' that lies between b'_1 and b'_2 , so we have $d_\Gamma(c_1, c_2) \leq 2e' + 4L_3$. So, in particular, since the components of q are geodesics in Γ , the subpath of q between b_1 and b_2 cannot have any components of length greater than $2e' + 4L_3$. So this subpath of q is a quasigeodesic for suitable constants, and (ii) now follows.

We get a corresponding result with p and q interchanged, and if this results in a larger constant k' , then we replace k' by this larger value. Now (iii) follows directly from (ii). \square

Corollary 3.6. *Let (G, Σ) be relatively hyperbolic and suitable for parabolic geodesic biautomaticity. Let $g_1, g_2 \in G$ and let $\lambda \geq 1$ and $c \geq 0$. Then there is a deterministic 2-tape finite state automaton $\mathcal{A} = \mathcal{A}(\lambda, c, g_1, g_2)$ such that:*

- (i) *for all u, v with $(u, v) \in L(\mathcal{A})$, we have $g_1 u =_G v g_2$;*
- (ii) *for all $u, v \in \Sigma^*$ with $g_1 u =_G v g_2$ that satisfy hypotheses (a) and (b) in the statement of Proposition 3.5, we have $(u, v) \in L(\mathcal{A})$.*

Proof. Let $w_1, w_2 \in \Sigma^*$ be geodesic words defining g_1 and g_2 , and let $k = \max(|w_1|, |w_2|)$, and let $e' = e'(\lambda, c, k)$ and $k' = k'(\lambda, c, k)$ be the constants defined in Proposition 3.5 (i) and (ii). We define a 2-tape asynchronous *word-difference* automaton \mathcal{A}' , following the recipe of [5, Definition 2.3.3]. The states \mathcal{A}' are elements of G of Σ -length at most e' , with start state g_1 and single accepting state g_2 , and transitions $g^{(x,y)} = x^{-1}gy$ for all $x, y \in \Sigma \cup \{\epsilon\}$ such that g and $x^{-1}gy$ are both states of \mathcal{A}' . Then \mathcal{A}' clearly satisfies (i) and it satisfies (ii) by Proposition 3.5 (i). But since \mathcal{A}' contains ϵ -transitions, it is non-deterministic.

Proposition 3.5 (ii) enables us to define a deterministic 2-tape automaton \mathcal{A} with $L(\mathcal{A}) = L(\mathcal{A}')$. After reading a prefix u_1 of u , we know that the lengths of the prefixes v_1 of v for which $v_1^{-1}g_1 u_1$ has Σ -length at most e' can differ by at most k' . So \mathcal{A} can read the longest such prefix v_1 and remember all shorter such prefixes. After reading each letter of u_1 , \mathcal{A} need never read forward more than k' letters of v_2 to maintain this information. If at any time there are no eligible prefixes v_1 of v then \mathcal{A} stops and rejects (u, v) . \square

Proposition 3.7. *Let (G, Σ) be relatively hyperbolic, and suitable for parabolic geodesic biautomaticity. Define \mathcal{L} to be the set of all words $u \in \Sigma^*$ such that the derived word $\widehat{u} \in \widehat{\Sigma}^*$ is geodesic and such that, for each parabolic subgroup H_i , all H_i -components of u lie in the specified geodesic biautomatic structure. Then \mathcal{L} is the language of an asynchronous biautomatic structure for G .*

Furthermore, for any ordering of Σ with associated lexicographical ordering \leq_{lex} of Σ^* , the language $\mathcal{L}_0 = \{u \in \mathcal{L} : u =_G v, v \in \mathcal{L} \Rightarrow u \leq_{\text{lex}} v\}$ is an asynchronous biautomatic structure for G with uniqueness.

Proof. If we can prove that \mathcal{L} is a regular language, then the first claim will follow immediately from Corollary 3.6. (The hypothesis that \hat{u} is geodesic for $u \in \mathcal{L}$ implies that \hat{u} does not backtrack.) The required property that all H_i -components of words in \mathcal{L} lie in the specified geodesic biautomatic structure is certainly testable by a finite state automaton, so we may restrict our attention to words that satisfy it.

For all $u \in \mathcal{L}$, we have $u_1 \in \mathcal{L}$ for all prefixes \widehat{u}_1 of \hat{u} . So if $u \in \Sigma^*$ with $u \notin \mathcal{L}$, then \hat{u} has a shortest prefix u_1 such that \widehat{u}_1 is a prefix of \hat{u} (i.e. such that u_1 is a union of complete components of u and subwords in $(\Sigma \setminus \mathcal{H})^*$) and $u_1 \notin \mathcal{L}$. Since the maximal proper prefix of \widehat{u}_1 is geodesic, \widehat{u}_1 labels a $(1, 2)$ -quasigeodesic in $\widehat{\Gamma}$, and it is not hard to see that \widehat{u}_1 cannot backtrack. Let $v_1 \in \mathcal{L}$ with $v_1 =_G u_1$. Then $(u_1, v_1) \in L(\mathcal{A})$, where $\mathcal{A} := \mathcal{A}(1, 2, 1_G, 1_G)$ as defined in Corollary 3.6. Furthermore, if (u_2, v_2) is a prefix of (u_1, v_1) in an accepting path of (u_1, v_1) through \mathcal{A} , then the Σ -length and hence also the $\widehat{\Sigma}$ -length of the element of G defined by $u_2^{-1}v_2$ is bounded. So \mathcal{A} can be modified to keep track of the difference between the $\widehat{\Sigma}$ -lengths of these pairs of prefixes (u_2, v_2) , and hence it can detect that $|\widehat{u}_1| > |\widehat{v}_1|$.

So a word $u \in \Sigma^*$ lies in \mathcal{L} if and only if its components satisfy the required condition, and if it has no prefix u_1 consisting of a union of complete components of u and subwords in $(\Sigma \setminus \mathcal{H})^*$ for which there exists v_1 with $(u_1, v_1) \in L(\mathcal{A})$ and $|\widehat{u}_1| > |\widehat{v}_1|$. So \mathcal{L} is regular by [5, Lemma 7.1.5].

The final statement follows immediately from [5, Theorem 7.3.2] (together with its proof, in which \mathcal{L}_0 is defined as we have done so here). \square

In Subsection 6.1, for words $w \in \Sigma^*$, we shall define $\text{nf}(w)$ to be the unique word in \mathcal{L}_0 with $w =_G \text{nf}(w)$. Anticipating that definition, we have:

Proposition 3.8. *There is a constant D (depending on G and Σ) such that $|\text{nf}(w)| \leq D|w|$ for any $w \in \Sigma^*$.*

Proof. It is sufficient to prove this when w is a geodesic word. Then by Lemma 3.4 \hat{w} is a $(\lambda, 0)$ -quasigeodesic for some $\lambda \geq 1$, and the result follows by applying Proposition 3.5 (iii) with $u = w$ and $v = \text{nf}(w)$ and $w_1 = w_2 = \epsilon$. \square

Remark 3.9. *The languages \mathcal{L} and \mathcal{L}_0 are not necessarily closed under subwords, but if $u \in \mathcal{L}$ or $u \in \mathcal{L}_0$, and u_1 is a subword that contains only complete components of u together with subwords in $(\Sigma \setminus \mathcal{H})^*$, then $u_1 \in \mathcal{L}$ or \mathcal{L}_0 .*

Remark 3.10. *It is not difficult to show that the paths in Γ labelled by words in \mathcal{L} and \mathcal{L}_0 are quasigeodesics, but we shall not need that property.*

4. Background on straight line programs

Let $\mathcal{G} = (V, S, \rho)$ be a straight line program over an alphabet Σ . For a variable $A \in V$, the word $\rho(A)$ is called the *right-hand side* of A . We define the *size* of \mathcal{G} to be the total length of all right-hand sides: $|\mathcal{G}| := \sum_{A \in V} |\rho(A)|$.

It follows from the acyclicity condition that each variable A in V has a well defined *height*, $h(A)$, namely the smallest positive integer r for which $\rho^r(A) \in \Sigma^*$. By removing from \mathcal{G} all variables that do not occur in any image $\rho^k(S)$ for $k \in \mathbb{N}$, we obtain (in time that is a polynomial function of $|\mathcal{G}|$) an SLP in which S is the only variable of maximal height. We call such an SLP *trimmed*; we will often want to assume this property for an SLP.

Suppose that A is a variable of height r in \mathcal{G} . We define $\text{val}_{\mathcal{G}}(A) := \rho^r(A) \in \Sigma^*$, and observe that $\text{val}_{\mathcal{G}}(S) = \text{val}(\mathcal{G})$. We also define a (trimmed) SLP $\mathcal{G}_A := (V_A, A, \rho_A)$ over Σ , the *restriction* of \mathcal{G} to A , with start variable A , set of variables V_A consisting of all $B \in V$ that appear within $\rho^k(A)$ for some $k \geq 0$, and map ρ_A defined to be the restriction of ρ to V_A . We note that for any $B \in V_A$, $\text{val}_{\mathcal{G}}(B) = \text{val}_{\mathcal{G}_A}(B)$, and in particular $\text{val}(\mathcal{G}_A) = \text{val}_{\mathcal{G}}(A)$.

If every right-hand side (except that of $\rho(S) = \epsilon$, if it occurs) has the form $a \in \Sigma$ or BC with $B, C \in V$ (that is, the grammar associated with \mathcal{G} is in Chomsky normal form), then we shall say that \mathcal{G} itself is in *Chomsky normal form*. We shall often want to assume this property.

Given SLPs \mathcal{G}_1 and \mathcal{G}_2 over a single alphabet Σ , we denote by $\mathcal{G}_1\mathcal{G}_2$ the SLP with $\text{val}(\mathcal{G}_1\mathcal{G}_2)$ equal to the concatenation $\text{val}(\mathcal{G}_1)\text{val}(\mathcal{G}_2)$, which we derive from \mathcal{G}_1 and \mathcal{G}_2 by adding to their disjoint union a single variable $S_{\mathcal{G}_1\mathcal{G}_2}$ and the single production $S_{\mathcal{G}_1\mathcal{G}_2} \rightarrow S_{\mathcal{G}_1}S_{\mathcal{G}_2}$. We extend this definition to a concatenation of any number of SLPs, and also to a concatenation of SLPs and words over Σ , so that a concatenation of the form $\mathcal{G} = u_0\mathcal{G}_1u_1 \cdots \mathcal{G}_ku_k$, where each \mathcal{G}_i is a SLP and each u_i a possibly empty word over Σ is formed by the addition of a single production

$$S_{\mathcal{G}} \rightarrow u_0S_{\mathcal{G}_1}u_1S_{\mathcal{G}_2}u_2 \cdots S_{\mathcal{G}_k}u_k$$

to the disjoint union of the productions of the CSLPs \mathcal{G}_i .

We will make use of a number of results from the literature, which we collect together here as a single proposition:

Proposition 4.1. *Let $\mathcal{G} = (V, S, \rho)$ be a straight line program over a finite alphabet Σ .*

- (i) *An algorithm exists that transforms \mathcal{G} into an SLP \mathcal{G}' in Chomsky normal form with $|\mathcal{G}'| \leq 2|\mathcal{G}|$ and $\text{val}(\mathcal{G}) = \text{val}(\mathcal{G}')$, in time that is a linear function of $|\mathcal{G}|$; see for example [11, Proposition 3.8].*
- (ii) *We have $|\text{val}(\mathcal{G})| \leq 3^{|\mathcal{G}|/3}$ [4, proof of Lemma 1].*

- (iii) The length $|\text{val}(\mathcal{G})|$ can be computed in time that is a polynomial function of $|\mathcal{G}|$ [11, Proposition 3.9].
- (iv) For $0 \leq i \leq j \leq |\text{val}(\mathcal{G})|$, an SLP with value $\text{val}(\mathcal{G})[i : j]$ can be computed in time that is a polynomial function of $|\mathcal{G}|$ [11, Proposition 3.9].
- (v) Given a deterministic finite state automaton M over the alphabet Σ and an SLP \mathcal{G} over the alphabet Σ , it can be determined in time that is a polynomial function of $|\mathcal{G}|$ whether $\text{val}(\mathcal{G}) \in L(M)$ [11, Theorem 3.11].
- (vi) Given two SLPs \mathcal{G} and \mathcal{H} , it can be checked in time that is a polynomial function of $|\mathcal{G}| + |\mathcal{H}|$ whether $\text{val}(\mathcal{G}) = \text{val}(\mathcal{H})$ [13].

Definition 4.2. Let $\mathcal{G} = (V, S, \rho)$ be an SLP over an alphabet Σ , with value $w_1 u w_2$. We say that u has a *root*, A in V , if for some k and $\ell = h(S) - k$, we have $\rho^k(S) = \alpha A \beta$, where $\alpha, \beta \in (V \cup \Sigma)^*$, $\rho^\ell(\alpha) = w_1$, $\rho^\ell(\beta) = w_2$, and $\rho^\ell(A) = u$.

4.1. Extensions of SLPs

It will sometimes be convenient in our proofs to make use of particular extensions of SLPs, namely cut-SLPs and tethered-SLPs.

Cut-SLPs (which we shall abbreviate as CSLPs) are defined analogously to SLPs, but in addition a cut-SLP may contain right-hand sides that are written in the form $B[: i]$ or $B[i :]$ for a variable B and an integer $i \geq 0$ [11]; CSLPs are used in the construction of a polynomial time algorithm for the compressed word problem of a free group in [10], where they are called *composition systems*. We call $[: i]$ and $[i :]$ cut operators. It is convenient to allow cut operators of the form $B[i : j]$ for $0 \leq i < j$, which we can achieve as the combination of two cut operators $B[: j][i :]$. If $\rho(A) = B[i : j]$ in a CSLP \mathcal{G} , then we define $\text{val}_{\mathcal{G}}(A)$ to be the string $\text{val}_{\mathcal{G}}(B)[i : j]$.

Given a CSLP \mathcal{G} , we denote by $\mathcal{G}[i : j]$ the CSLP with value $\text{val}(\mathcal{G})[i : j]$ that we derive from \mathcal{G} by adding to \mathcal{G} a single variable $S_{\mathcal{G}[i:j]}$ (the start variable of $\mathcal{G}[i : j]$) and the single production $S_{\mathcal{G}[i:j]} \rightarrow S_{\mathcal{G}}$ to the CSLP \mathcal{G} .

In a CSLP, as in a SLP, we require that the associated binary relation is acyclic. We define concatenations of CSLPs analogously to concatenations of SLPs. The following results can be found in the literature; the second follows from the first together with Proposition 4.1 (vi).

- Proposition 4.3.**
- (i) From a given CSLP \mathcal{G} one can compute, in time that is a polynomial function of $|\mathcal{G}|$, an SLP \mathcal{G}' such that $\text{val}(\mathcal{G}) = \text{val}(\mathcal{G}')$ [7]; see also [11, Theorem 3.14].
 - (ii) Given CSLPs \mathcal{G}_1 and \mathcal{G}_2 , it can be checked in polynomial time whether $\text{val}(\mathcal{G}_1) = \text{val}(\mathcal{G}_2)$.

As in [8], the proof of the main theorem also involves extensions to SLPs and CSLPs that we call TSLPs and TCSLPs, respectively (T stands for “tethered”).

These extensions only make sense when the alphabet Σ is the (inverse closed) generating set of a group G , and when G is equipped with a normal form $\text{nf}(w)$ for words $w \in \Sigma^*$; that is, $\text{nf}(w) \in \Sigma^*$, $\text{nf}(w) =_G w$, and, for $v, w \in \Sigma^*$, $v =_G w \Rightarrow \text{nf}(v) = \text{nf}(w)$.

We extend the definition of a SLP or CSLP to that of a TSLP or TCSLP over such an alphabet Σ by allowing additional right-hand sides that are expressions of the form $B\langle\alpha, \beta\rangle$ with a variable B , and strings α, β over Σ , each of length at most J , where $J = J_{\mathcal{T}} \in \mathbb{N}$ is a parameter of the TSLP or TCSLP \mathcal{T} . Given a right-hand side $\rho(A) = B\langle\alpha, \beta\rangle$ in a TSLP or TCSLP \mathcal{T} , we define

$$\text{val}_{\mathcal{T}}(A) := \text{nf}(\alpha \text{val}_{\mathcal{T}}(B) \beta^{-1}).$$

In TSLPs and TCSLPs, as with SLPs and CSLPs, we require that the associated binary relations are acyclic, and we define concatenations of TSLPs and TCSLPs analogously to concatenations of SLPs.

We define the *size* of a variable A in a TCSLP \mathcal{T} to be the total number of occurrences of symbols from $\Sigma \cup V$ in $\rho(A)$, and the size of \mathcal{T} is obtained by taking the sum over all variables.

As we did for SLPs, for a TCSLP \mathcal{U} over Σ , we define the height, $\text{h}(A)$, of a variable A of \mathcal{U} to be the smallest integer k such that $\rho_{\mathcal{U}}^k(A) \in \Sigma^*$. Proposition 4.1 (i) extends to TCSLPs: for a given TCSLP \mathcal{U} over Σ with $\text{val}(\mathcal{U}) \neq \epsilon$, we can in linear time construct a TCSLP \mathcal{U}' with the same value in which $|\mathcal{U}'| \leq 2|\mathcal{U}|$ and all right hand sides $\rho_{\mathcal{U}'}(A)$ that lie in $(V_{\mathcal{U}'} \cup \Sigma)^*$ have the form $a \in \Sigma$ or BC with $B, C \in V_{\mathcal{U}'}$.

We say that the TCSLP (or TSLP) \mathcal{T} is *nf-reduced* if for every variable A the word $\text{val}_{\mathcal{T}}(A)$ is nf-reduced; that is $\text{nf}(\text{val}_{\mathcal{T}}(A)) = \text{val}_{\mathcal{T}}(A)$.

In Section 8 we shall prove Theorem 8.1 that, given a TCSLP \mathcal{T} over a suitably chosen generating set Σ of a group that is hyperbolic relative to a collection of free abelian subgroups (and which satisfies suitable conditions relating to the ‘splitting of components’, explained in Definition 7.2), we can, in time that is a polynomial function of $|\mathcal{T}|$ (depending on $J_{\mathcal{T}}$), compute an SLP \mathcal{G} with $\text{val}(\mathcal{T}) = \text{val}(\mathcal{G})$. This result will be a vital component of our main result. Its proof will require the following generalisation of Proposition 4.1 (ii) to nf-reduced TCSLPs.

Proposition 4.4. *Let $\mathcal{T} = (V, S, \rho)$ be an nf-reduced TCSLP. Then there is a constant J' (depending on Σ and $J_{\mathcal{T}}$) such that $|\text{val}(\mathcal{T})| \leq (J')^{|\mathcal{T}|}$*

Proof. As we saw above, there is a TCSLP $\mathcal{T}' = (V', S, \rho')$ with $\text{val}(\mathcal{T}') = \text{val}(\mathcal{T})$, $|\mathcal{T}'| \leq 2|\mathcal{T}|$ in which all right hand sides $\rho'(A)$ that lie in $(V' \cup \Sigma)^*$ have the form $a \in \Sigma$ or BC with $B, C \in V'$. If $\rho(A) = B\langle\alpha, \beta\rangle$ with $|\alpha|, |\beta| \leq J_{\mathcal{T}}$ then $\rho'(A) = \rho(A)$ and, by 3.5 (iii), we have $|\text{val}(A)| \leq (k' + 1)\text{val}(B)$ with $k' = k'(1, 0, J_{\mathcal{T}})$.

We claim that $\text{val}_{\mathcal{T}'}(A) \leq \max(2, k' + 1)^{h(A)}$ for all $A \in V'$. Since $h(S) \leq |\mathcal{T}'| \leq 2|\mathcal{T}|$ this will prove the proposition with $J' = \max(2, k' + 1)^2$.

The proof of the claim is by induction on $h(A)$, and the base case $h(A) = 0$ is clear. Otherwise, if $\rho'(A) = BC$ then $|\text{val}_{\mathcal{T}'}(A)| = |\text{val}_{\mathcal{T}'}(B)| + |\text{val}_{\mathcal{T}'}(C)|$ with $h(B), h(C) \leq h(A)$; if $\rho'(A) = B[i :]$ or $B[: j]$, then $|\text{val}_{\mathcal{T}'}(A)| \leq |\text{val}_{\mathcal{T}'}(B)|$; and if $\rho'(A) = B\langle \alpha, \beta \rangle$ with $|\alpha|, |\beta| \leq J_{\mathcal{T}}$ then, as saw above, $|\text{val}_{\mathcal{T}'}(A)| \leq (k' + 1)|\text{val}_{\mathcal{T}'}(B)|$. In all three cases the claim follows immediately from the inductive hypothesis. \square

5. Some constructions of SLPs for abelian groups

We need a few results about SLPs for finitely generated abelian groups, which will be the parabolic subgroups of our relatively hyperbolic groups.

Lemma 5.1. *Let G be a finitely generated abelian group with generating set $X = \{x_1^{\pm 1}, \dots, x_k^{\pm 1}\}$, and let \mathcal{G} be an SLP over X . Then, in time that is polynomial in $|\mathcal{G}|$, we can*

- (i) *compute the vector (n_1, n_2, \dots, n_k) , defined by $\text{val}(\mathcal{G}) =_G \prod_{i=1}^t x_i^{n_i}$ (where the integers n_i are output as binary numbers),*
- (ii) *construct an SLP \mathcal{G}' with*

$$\text{val}(\mathcal{G}') = \prod_{i=1}^t x_i^{n_i}, \quad |\mathcal{G}'| \leq \max(4k \log_2(|\text{val}(\mathcal{G}')|), 1).$$

Proof. The fact that we can compute the integers n_i in polynomial time follows from Proposition 4.1 (ii) together with the fact that we can perform addition and subtraction on integers of absolute value at most N in time $O(\log N)$.

For part (ii), note first that for any $x \in X$ and integer $n > 0$, we can define an SLP with value x^n and size at most $\max(3 \log_2(n), 1)$ by expressing n in binary and introducing a variable A_i with size 2 and value x^{2^i} for each i with $2^i < n$. For example, with $n = 14$, we write $14 = 2^1 + 2^2 + 2^3$, and define the SLP $(\{S, A_1, A_2, A_3\}, S, \rho)$ with $\rho(S) = A_1 A_2 A_3$, $\rho(A_1) = xx$, $\rho(A_2) = A_1 A_1$, $\rho(A_3) = A_2 A_2$, which has value x^{14} and size 9. So we can construct an SLP \mathcal{G}' with value $\text{val}(\mathcal{G})$ and size at most $(3 \log_2(|\text{val}(\mathcal{G}')|) + 1)k \leq \max(4k \log_2(|\text{val}(\mathcal{G}')|), 1)$. \square

Proposition 5.2. *Let $G = \langle X \rangle$ be a finitely generated abelian group. Then X is contained in a finite subset Y of G for which, given an SLP \mathcal{G} over Y , we can construct an SLP \mathcal{G}' over Y with $\text{val}(\mathcal{G}') = \text{slex}_Y(\text{val}(\mathcal{G}))$ and $|\mathcal{G}'| \leq \max(4|Y| \log_2(|\text{val}(\mathcal{G}')|), 1)$ in time that is a polynomial function of $|\mathcal{G}|$.*

Proof. Choose generators $z_1, \dots, z_r, z_{r+1}, \dots, z_s$ of the cyclic direct factors of G , where z_i has infinite order if and only if $i \leq r$. We can write each element x of X as $\prod_{i=1}^s z_i^{n_i(x)}$ with $n_i(x) \in \mathbb{Z}$; note that the set of all $n_i(x)$ is finite, and so its elements can be regarded as constants. Let e be the exponent of the torsion subgroup of G , and let $M := e(\max_{x \in X} \sum_{i=1}^r |n_i(x)| + 1)$.

Now we define Y to be the union of the three sets $\{z_i^{\pm M} : 1 \leq i \leq r\}$, $\{z_i^{\pm 1} : 1 \leq i \leq s\}$, and X ; we order Y so that

- pairs of mutually inverse generators are adjacent in the ordering; and
- the generators $z_i^{\pm M}$ come first in the ordering, with $z_1^M < z_1^{-M} < \dots < z_r^M < z_r^{-M}$.

We denote by y_i and y_i^{-1} the elements in positions $2i - 1$ and $2i$ of this ordered set of generators. Suppose that there are $2t$ such generators in total.

We claim that, in any geodesic word $\prod_{i=1}^t y_i^{n_i}$ with $n_i \in \mathbb{Z}$, we have $|n_k| < M$ for all $k > r$.

To see this, given such a geodesic word, suppose that $r < k \leq t$. If $y_k = z_j$ for some j , then we could replace y_k^M by the shorter word y_j if $j \leq r$ and by the empty word if $j > r$; so we must have $|n_k| < M$. Otherwise, we have $y_k \in X$, and then $x := y_k =_G \prod_{i=1}^s z_i^{n_i(x)}$ and (because $e|M$), $x^M =_G \prod_{i=1}^r z_i^{Mn_i(x)} =_G \prod_{i=1}^r y_i^{n_i(x)}$. This last word is shorter than x^M , since $M > \sum_{i=1}^r |n_i(x)|$. So replacing x^M by $\prod_{i=1}^r y_i^{n_i(x)}$ would be a reduction in length, and hence again we must have $|n_k| < M$.

Now it follows from Lemma 5.1 that, in time polynomial in $|\mathcal{G}|$, we can write $g := \text{val}(\mathcal{G})$ as an integer vector over the generators in Y , and hence as a product $\prod_{i=1}^t y_i^{n_i} \in G$. Now we want to compute $\text{slex}(g)$. If any $|n_k| \geq M$ for $k > r$, then we write $n_k = qM + n'_k$ with $|n'_k| < M$ and $\text{sgn}(n_k) = \text{sgn}(n'_k)$ and replace the expression for g by an equivalent expression in which y_k^{qM} is replaced by a shorter word, as described in the preceding paragraph. This involves integer arithmetic and can be done in time polynomial in the sizes of the integers. So we may assume that $g = \prod_{i=1}^t y_i^{n_i} \in G$ with $|n_k| < M$ for all $k > r$.

Now suppose that $\text{slex}(g) = \prod_{i=1}^t y_i^{n'_i}$. Then $\prod_{i=1}^t y_i^{(n_i - n'_i)} = 1$. and since $|n'_k|, |n_k| < M$, we have $|n_k - n'_k| < 2M$ for all $k > r$. Since y_1, \dots, y_r are free generators, there are no nontrivial relations in G that involve only y_1, \dots, y_r , and so the number of lists of integers m_1, \dots, m_t for which $\prod_{i=1}^t y_i^{m_i} = 1$ and $|m_i| < M$ is at most $(2M)^{t-r}$. So we can consider each of these equations in turn, and thereby find all possible values of n'_i . From these, we select the shortlex least representative of g . We can now use Lemma 5.1 to construct the required SLP \mathcal{G}' . \square

6. Examining the geometry of $\widehat{\Gamma}$

For the remainder of this article, G will be a group that is hyperbolic relative to a collection of free abelian subgroups H_i , and \mathcal{H} the set of non-identity elements of those subgroups, as in Section 3.

By [5, Theorem 4.3.1], abelian groups are shortlex automatic with respect to any finite ordered generating set, and a little thought shows that any automatic structure for an abelian group must be biautomatic. So the hypotheses of shortlex biautomaticity of the subgroups H_i in Proposition 3.2 are satisfied for all choices of finite, ordered generating sets of the H_i .

6.1. Fixing the generating set Σ and constant δ

Given a generating set Σ' for G , we define, for each i , $X_i := (\Sigma' \cup \mathcal{H}') \cap H_i$, where \mathcal{H}' is the finite subset of \mathcal{H} whose existence is guaranteed by Proposition 3.2. It follows from Corollary 3.3 that, for each i , X_i generates H_i . Now we select finite subsets $Y_i \supset X_i$ as in Proposition 5.2, and define $\Sigma := \Sigma' \cup \mathcal{H}' \cup \bigcup_{i \in \Omega} Y_i$. Our construction ensures that $\Sigma \cap H_i = Y_i$. For the remainder of this article, we use this generating set Σ for G , and denote $\Sigma \cap H_i$ by Σ_i ; we note that (G, Σ) is suitable for parabolic geodesic biautomaticity. For convenience we assume that the elements of Σ represent distinct elements of G and that no element of Σ represents the identity element.

Furthermore, we define \mathcal{L}_0 to be the asynchronous automatic structure of G with uniqueness that is defined in Proposition 3.7, where we use the shortlex biautomatic structure on H_i over Σ_i . We shall call words in \mathcal{L}_0 \mathcal{L}_0 -reduced and, for $v \in \Sigma^*$, we denote the unique element $u \in \mathcal{L}_0$ with $u =_G v$ by $\text{nf}(v)$. We shall use this normal form in all of our TCSLPs over Σ . Our assumptions on Σ ensure that $\text{nf}(a) = a$ for all $a \in \Sigma$.

Now that we have fixed Σ , we can also fix the associated Cayley graphs Γ and $\widehat{\Gamma}$. We know from the definition of relative hyperbolicity that $\widehat{\Gamma}$ is Gromov hyperbolic, and we fix the constant $\delta > 0$ such that $\widehat{\Gamma}$ is a δ -hyperbolic space; that is, all geodesic triangles in $\widehat{\Gamma}$ are δ -thin (and hence also δ -slim) as defined in [1, Chapter 1]. We assume these choices for Σ , Γ , $\widehat{\Gamma}$ and δ for the remainder of this article.

6.2. Properties of some geodesic triangles and quadrilaterals in $\widehat{\Gamma}$

Proposition 6.1. *Let $u, v, w \in \Sigma^*$ be words over our selected generating set for G with $v =_G uw$, where \hat{u} and \hat{v} are geodesic words over $\widehat{\Sigma}$, and $|\hat{w}| \leq \kappa$ for some κ . Then constants $K_1(\kappa)$ and $L_1(\kappa)$ exist such that, for any vertex d on the path in $\widehat{\Gamma}$ labelled by \hat{u} and at distance at least $K_1(\kappa)$ from the end of that path, there is a vertex e of $\widehat{\Gamma}$ on the path labelled by \hat{v} , with $d_\Gamma(d, e) \leq L_1(\kappa)$.*

Proof. Since \hat{u} labels a geodesic path in $\hat{\Gamma}$ and $|\hat{w}| \leq \kappa$, we know that $\hat{u}\hat{w}$ labels a $(1, 2\kappa)$ -quasigeodesic path in $\hat{\Gamma}$. Our aim is to replace $\hat{u}\hat{w}$ by a G -equivalent word $t \in \hat{\Sigma}^*$ that also labels a $(1, 2\kappa)$ -quasigeodesic path, and does not backtrack, and whose prefix of length $|\hat{u}| - K_1(\kappa)$ matches the corresponding length prefix of \hat{u} for some constant $K_1(\kappa)$. The result will then follow directly from Proposition 3.1 (1), applied to t and \hat{v} .

If $\hat{u}\hat{w}$ backtracks, then it contains a subword of length greater than 1 that represents an element of H_i for some $i \in \Omega$. Since such elements have length at most 1 over $\hat{\Sigma}^*$, and $\hat{u}\hat{w}$ labels a $(1, 2\kappa)$ -quasigeodesic path, any such subword has length at most $K_1(\kappa) := 1 + 2\kappa$. Furthermore, since \hat{u} does not vertex backtrack, such a subword must intersect the suffix \hat{w} nontrivially. So it does not intersect the prefix $\hat{u}[:|\hat{u}| - K_1(\kappa)]$ of u . So, after replacing any such subwords by G -equivalent words of length 1 over Γ^* , the resulting word t does not backtrack, and has the desired property. \square

Our next result can be proved by two applications of Proposition 6.1, and we omit the details.

Lemma 6.2. *Let $u, v, w_1, w_2 \in \Sigma^*$ be words over Σ with $w_1u =_G vw_2$, and suppose that \hat{u} and \hat{v} are geodesic words, and that $|\widehat{w}_1|, |\widehat{w}_2| \leq \kappa$ for some κ . Then constants $K_2(\kappa)$ and $L_2(\kappa)$ exist, such that, for any vertex d on the path in $\hat{\Gamma}$ labelled by \hat{u} and at distance at least $K_2(\kappa)$ from the beginning and the end of that path, there is a vertex e of $\hat{\Gamma}$ on the path labelled by \hat{v} , with $d_\Gamma(d, e) \leq L_2(\kappa)$.*

Proposition 6.3. *There is a linear function $f : \mathbb{N} \rightarrow \mathbb{N}$ and a constant L' , with the following property. Suppose that u, v, w_1, w_2 are words over our selected generating set Σ for G , with $w_1u =_G vw_2$, and suppose that \hat{u} and \hat{v} are geodesic words. Consider a quadrilateral in $\hat{\Gamma}$ with sides labelled $\widehat{w}_1, \hat{u}, \widehat{w}_2, \hat{v}$.*

Then, for any vertex d on the path labelled by \hat{u} in that quadrilateral, and at distance at least $f(\max\{|\widehat{w}_1|, |\widehat{w}_2|\})$ from the beginning and the end points of that path, there is a vertex e of the path labelled by \hat{v} in the quadrilateral, with $d_\Gamma(d, e) \leq L'$.

Proof. We recall the constant δ defined in Section 6.1. In the quadrilateral defined in the statement, replace the sides labelled by \widehat{w}_1 and \widehat{w}_2 by geodesic paths p_1 and p_2 between their endpoints to give a geodesic quadrilateral in $\hat{\Gamma}$. So $|p_i| \leq |\widehat{w}_i|$ for $i = 1, 2$.

Since $\hat{\Gamma}$ is δ -hyperbolic, any vertex on any side of this quadrilateral is at distance at most 2δ in $\hat{\Gamma}$ from a vertex on one of the other three sides.

Now for a vertex d on the path labelled \hat{u} that is at distance ℓ along \hat{u} from the start point, all vertices on the path p_1 are at distance at least $\ell - |p_1| \geq \ell - |\widehat{w}_1|$ from d in $\hat{\Gamma}$. So, if $\ell \geq |\widehat{w}_1| + 2\delta + 1$, then d cannot be at distance at most 2δ

from a vertex on p_1 . Similarly, if the distance of d on \hat{u} from the end point of \hat{u} is greater than $|\widehat{w}_2| + 2\delta + 1$, then it cannot be at distance at most 2δ from a vertex on p_2 . So if both of those conditions on d hold, then d must be at distance at most 2δ in $\widehat{\Gamma}$ from a vertex e on the side of the quadrilateral labelled \hat{v} .

Now, by Lemma 6.2, there are constants $K_2 := K_2(2\delta)$ and $L_2 := L_2(2\delta)$ depending only on G and Σ , such that any vertex d on the path labelled \hat{u} that is at distance at least $|\widehat{w}_1| + 2\delta + 1 + K_2$ and $|\widehat{w}_2| + 2\delta + 1 + K_2$ from the start and end points of that path, respectively, is at distance at most L_2 in Γ from a vertex on the side of the quadrilateral labelled \hat{v} . This proves the proposition with $L' = L_2$. \square

6.3. Fixing the constant L

We define L to be the larger of the constants $L_1(\delta)$ and L' that were defined in Propositions 6.1 and 6.3, respectively. We will refer to L repeatedly in the final two sections of the paper.

7. Some constructions of SLPs for relatively hyperbolic groups

The proof of our main result (which is split across Theorems 9.1 and 8.1) will need some technical results relating to SLPs for our relatively hyperbolic group G over our selected generating set Σ . In general these results will be applied to sub-SLPs of the SLP that is input and the SLPs that are derived from it within the above two theorems.

Proposition 7.1. *Let \mathcal{G} be an SLP over our selected generating set Σ for G , and let $w := \text{val}(\mathcal{G})$. Then, in time polynomial in $|\mathcal{G}|$, we can construct an SLP \mathcal{G}' in Chomsky normal form with value w such that, for all variables A of \mathcal{G}' , all components of $\text{val}(A)$ have roots in \mathcal{G}'_A .*

Furthermore, for each parabolic subgroup H_i , we can compute a list of those variables A of \mathcal{G}' for which $\text{val}_{\mathcal{G}'}(A) \in \Sigma_i^$.*

Proof. We may assume that \mathcal{G} is trimmed and in Chomsky normal form. We construct $\mathcal{G}' = (V', S, \rho')$ from $\mathcal{G} = (V, S, \rho)$ by modifying the map ρ on some of the variables in V , while also introducing some new variables that are needed within those new images for ρ . For each $A \in V$, we will have $\text{val}_{\mathcal{G}'}(A) = \text{val}_{\mathcal{G}}(A)$. In addition, for each $A \in V$, every component of $\text{val}_{\mathcal{G}'}(A)$ will have a root in \mathcal{G}'_A .

We put the variables of V into increasing order of height, and consider them in that order. The SLP \mathcal{G}' will be the last of a sequence of SLPs $\mathcal{G}_0 = \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n$,

where $n = |V|$, and $\mathcal{G}_i = (V_i, S, \rho_i)$ will be made from \mathcal{G}_{i-1} by considering the i -th variable, A_i , in the list. We have $V_i \supset V_{i-1}$ and form ρ_i as a modification of ρ_{i-1} . The SLP \mathcal{G}_i might not be in Chomsky normal form, since, for some variables A' , $\rho_i(A')$ could be a string of any number of variables ≥ 0 , possibly the empty string, But it will be clear from our construction that this is the only obstruction to \mathcal{G}_i being in Chomsky normal form.

We describe the construction of the sequence $\mathcal{G}_1, \dots, \mathcal{G}_n$, and prove by induction on i that, for every variable A' of \mathcal{G}_i not in the sequence A_{i+1}, \dots, A_n , every component of $\text{val}_{\mathcal{G}_i}(A')$ has a root in $(\mathcal{G}_i)_{A'}$. To prove the i -th inductive step we need to verify the existence of appropriate roots both for the variable A_i and for any new variables that are defined in the construction of \mathcal{G}_i from \mathcal{G}_{i-1} . For each such A_i and all new variables, we can record whether their values lies in Σ_i^* for some i .

If A_i has height one, then $\rho(A_i)$ has length at most 1, with A_i as its root. So no modification is necessary, and $\mathcal{G}_i = \mathcal{G}_{i-1}$. So now suppose that A_i has height greater than one. For notational convenience we define $A := A_i$. The construction of \mathcal{G}_i from \mathcal{G}_{i-1} will involve the definition of some new variables, and of the images of these and of A under ρ_i ; the images under ρ_i and ρ_{i-1} of all other variables will match.

Since \mathcal{G} is in Chomsky normal form, $\rho_{i-1}(A) = \rho(A)$, and $\text{h}(A) > 1$, we have variables $B, C \in V$ with $\rho_{i-1}(A) = BC$. By the inductive hypothesis, every component of $\text{val}_{\mathcal{G}}(A)$ that lies entirely within $\text{val}_{\mathcal{G}}(B)$ or $\text{val}_{\mathcal{G}}(C)$ has a root in $(\mathcal{G}_{i-1})_B$ or $(\mathcal{G}_{i-1})_C$ (respectively), and hence in \mathcal{G}_{i-1} . So no modification is necessary unless $\text{val}_{\mathcal{G}}(A)$ contains a component $u = u_1u_2$ for which $\text{val}_{\mathcal{G}}(B) = v_1u_1$ and $\text{val}_{\mathcal{G}}(C) = u_2v_2$ with $u_1, u_2 \neq \epsilon$. We suppose that u is such a component. We note that (since, as discussed earlier when we defined components of words, any two components of $\text{val}_{\mathcal{G}}(A)$ are disjoint) any other component of $\text{val}_{\mathcal{G}}(A)$ is a subword of either v_1 or v_2 , and hence of $\text{val}_{\mathcal{G}}(B)$ or $\text{val}_{\mathcal{G}}(C)$. By the induction hypothesis, u_1, u_2 have roots D_1, D_2 in $(\mathcal{G}_{i-1})_B$ and $(\mathcal{G}_{i-1})_C$ respectively. By finding the rightmost variable in $\rho^k(B)$ for $k = 1, 2, \dots$ and using our records of which variables have values in Σ_i^* for some i , we can locate the variable D_1 in polynomial time, and similarly D_2 . (In fact we do that in any case to establish whether we are in the situation where modification is required.)

Now we introduce a new variable D , and define $\rho_i(D) := D_1D_2$. We also introduce new variables B', C' and set $\rho_i(A) = B'DC'$. We want to define $\rho_i(B')$ and $\rho_i(C')$ to ensure that $\text{val}_{\mathcal{G}_i}(B') = v_1$, and $\text{val}_{\mathcal{G}_i}(C') = v_2$, and this motivates the definitions that now follow.

First we note the existence of a sequence of variables $B_0 = B, B_1, \dots, B_r = D_1$ in V_{i-1} , for which each B_{j+1} is the last letter of $\rho_{i-1}(B_j)$; that is $\rho_{i-1}(B_j) = \beta_j B_{j+1}$, for some $\beta_j \in V_{i-1}^*$.

Similarly, we find a sequence of variables $C_0 = C, C_1, \dots, C_s = D_2$ in V_{i-1} , such




Figure 2: Adjusting the SLP below A

that, for each j , $\rho_{i-1}(C_j) = C_{j+1}\gamma_j$, for some $\gamma_j \in V_{i-1}^*$. We define $B'_0 := B'$, $C'_0 := C'$ and introduce new variables B'_j for each $j = 1, \dots, r-1$, and C'_j for each $j = 1, \dots, s-1$. We define

$$V_i := V_{i-1} \cup \{B'_0, B'_1, \dots, B'_{r-1}, C'_0, C'_1, \dots, C'_{s-1}, D\}.$$

Then we define

$$\rho_i(B'_0) := \beta_0 B'_1, \quad \rho_i(B'_j) := \beta_j B'_{j+1} \text{ for } j = 1, \dots, r-2, \quad \rho_i(B'_{r-1}) := \beta_{r-1},$$

and similarly

$$\rho_i(C'_0) := \gamma_0 C'_1, \quad \rho_i(C'_j) := \gamma_j C'_{j+1} \text{ for } j = 1, \dots, s-2, \quad \rho_i(C'_{s-1}) := \gamma_{s-1}.$$

and $\rho_i(E) := \rho_{i-1}(E)$ for all variables $E \in V_{i-1} \setminus \{A\}$. Figure 2 illustrates this adjustment of \mathcal{G}_{i-1} to \mathcal{G}_i .

This completes our construction of \mathcal{G}_i , which is certainly acyclic. Our definition of new variables and of their images under ρ_i was designed to ensure that $\text{val}_{\mathcal{G}_i}(A_i) = \text{val}_{\mathcal{G}}(A_i)$. We note that $\text{h}_{\mathcal{G}_i}(B') = \text{h}_{\mathcal{G}'_{i-1}}(B)$, $\text{h}_{\mathcal{G}_i}(C') = \text{h}_{\mathcal{G}_{i-1}}(C)$, and $\text{h}_{\mathcal{G}_i}(D) \leq \text{h}_{\mathcal{G}_{i-1}}(A_i)$ (equality can occur here if $D_1 = B$ or $D_2 = C$). So $\text{h}_{\mathcal{G}_i}(A_i) \leq \text{h}_{\mathcal{G}_{i-1}}(A_i) + 1$, for each i , and hence $\text{h}_{\mathcal{G}_i}(S) \leq \text{h}_{\mathcal{G}_{i-1}}(S) + 1$.

We need to verify our claim that, for all variables $A' \in \mathcal{G}_i$ that are not in the sequence A_{i+1}, \dots, A_n , every component of $\text{val}_{\mathcal{G}_i}(A')$ has a root in $(\mathcal{G}_i)_{A'}$. This is true for all $A' \in V_{i-1} \setminus \{A_i, A_{i+1}, \dots, A_n\}$, because $\rho_i(A') = \rho_{i-1}(A')$ and hence $\text{val}_{\mathcal{G}_i}(A') = \text{val}_{\mathcal{G}}(A')$ for all such A' . Our construction of \mathcal{G}_i was designed to ensure that our claim is true for $A' = A_i$, and it is clearly true for $A' = D, D_1$

and D_2 . Finally, for the new variables B'_j and C'_j , the components of $\text{val}_{\mathcal{G}_i}(B'_j)$ and $\text{val}_{\mathcal{G}_i}(C'_j)$ are components of $\text{val}_{\mathcal{G}_{i-1}}(B_j)$ and $\text{val}_{\mathcal{G}_{i-1}}(C_j)$ respectively and, by the inductive hypothesis, they have roots in $(\mathcal{G}_{i-1})_{B_j}$ and $(\mathcal{G}_{i-1})_{C_j}$. It follows immediately from the definitions of $\rho_i(B'_j)$ and $\rho_i(C'_j)$ that these components have roots in $(\mathcal{G}_i)_{B'_j}$ and $(\mathcal{G}_i)_{C'_j}$.

After the process is complete, we have $\text{h}_{\mathcal{G}'}(S) \leq n + \text{h}_{\mathcal{G}}(S) \leq 2n$. So, during each of the n steps of this process the number of variables we have added has been at most $4n$, and hence the process is bounded by a polynomial in $|V|$, so certainly by a polynomial in $|\mathcal{G}|$. We complete the proof by putting the final SLP into Chomsky normal form. Since this involves only the addition of new variables, it does not affect the property that all components of w have roots. \square

Definition 7.2. Let w be a word over Σ and suppose that $u := w[i : j]$ is a subword of w . We say that u *splits a component* of w if u starts or ends part way through a component of w , but is not a subword of a single component; otherwise we say that u *splits no components* of w .

So if u splits no component of w , then either u is a proper subword of a component of w , or u is a concatenation of components of w and of subwords in $(\Sigma \setminus \mathcal{H})^*$. In that second case, there exist integers k, l with $k < l$ such that $\hat{u} = \hat{w}[k : l]$; we shall sometimes choose to write $w[[k : l]]$ rather than $w[i : j]$ as a notation for this subword u of w .

Now let $\mathcal{G} = (V, S, \rho)$ be an SLP, CSLP, TSLP or TCSP for the group G , and let $w := \text{val}(\mathcal{G})$. We say that \mathcal{G} *splits no components* (or is *non-splitting*) if for any variable A of \mathcal{G} , whenever $\text{val}(A)$ occurs as a subword of w with root A , then that that subword splits no components of w .

Now suppose that A, B are variables in a CSLP \mathcal{G} and that $\rho(A) = B[i : j]$ (so that $\text{val}_{\mathcal{G}}(A) = \text{val}_{\mathcal{G}}(B)[i : j]$). We say that the cut operator $B[i : j]$ *splits a component* if the subword $w_B[i : j] = w_A$ splits a component of w , or if it is a proper subword of a component of w ; otherwise $B[i : j]$ is called *non-splitting*. Notice that for cut operators to be non-splitting, we are also excluding the possibility of w_A being a proper subword of a component of w_B . If $B[i : j]$ is a non-splitting cut-operator and k, l are the integers defined by $\text{val}(B)[i : j] = \text{val}(B)[[k : l]]$, it is often convenient to specify the cut operator in terms of k and l rather than i and j , that is, as $B[[k : l]]$. In this case, we say that the cut operator is *specified relative to compression*.

In general the SLPs (and CSLPs, TSLPs and TCSPs) that we shall construct in this article, as well as cut operators within the TCSPs, will not split components, and the cut operators will be specified relative to compression.

Remark 7.3. *Suppose that \mathcal{G} is an SLP as in the conclusion of Proposition 7.1 (where it is called \mathcal{G}'); that is, for each variable A of \mathcal{G} , every component of*

$\text{val}(A)$ has a root in \mathcal{G}_A . Then \mathcal{G} splits no components. Moreover, if \mathcal{G} is trimmed and $\text{val}(\mathcal{G})$ is nf-reduced, then so is $\text{val}(A)$, for every variable A of \mathcal{G} .

Proof. Suppose that the subword $u = \text{val}(A)$ of $w := \text{val}(\mathcal{G})$ splits a component of w ; so we have $u = u_1u_2$ with u_1 and u_2 nonempty, where either u_1 is a proper suffix or u_2 is a proper prefix of a component v of w . But, by assumption, v has a root B in \mathcal{G} , and these occurrences of $\text{val}(A)$ and $\text{val}(B)$ in w overlap without one being a subword of the other, which is not possible. The final assertion now follows from Remark 3.9. \square

In order to build CSLPs that do not split components, we shall need the following result.

Corollary 7.4. *Let \mathcal{G} be an SLP over our selected generating set Σ for G , let $w := \text{val}(\mathcal{G})$, and suppose that every component of w has a root in \mathcal{G} (and hence \mathcal{G} splits no components). Then given k, l with $0 \leq k < l \leq |\hat{w}|$, in time that is polynomial in $|\mathcal{G}|$, we can compute the integers i, j with $w[i : j] = w[[k : l]]$. Conversely, given i and j such that the subword $w[i : j]$ of w is a union of complete components of w and subwords in $(\Sigma \setminus \mathcal{H})^*$, we can compute k and l with $w[[k : l]] = w[i : j]$, in polynomial time.*

Proof. By regarding the roots of components of w as new terminals, we can regard \mathcal{G} as an SLP $\hat{\mathcal{G}}$ over some finite subset of the infinite alphabet $\hat{\Sigma}$ with $\text{val}(\hat{\mathcal{G}}) = \hat{w}$. Our lists of variables A of \mathcal{G} for which $\text{val}(A) \in \Sigma_i^*$ for some i enable us to identify such variables. We can then use Proposition 4.1 (iv) to compute SLPs over $\hat{\Sigma}$ with values $\hat{w}[: l]$ and $\hat{w}[k : l]$. Then by reinterpreting them as SLPs over Σ , we can (by Proposition 4.1 (iii)) compute their lengths and thereby compute i and j .

For the converse, we compute SLPs for $w[: j]$ and $w[i : j]$, regard then as SLPs with values $\hat{w}[: l]$ and $\hat{w}[k : l]$ over a finite subset of $\hat{\Sigma}$, and then compute their lengths. \square

Lemma 7.5. *Let \mathcal{G} be an SLP over our selected generating set Σ for G , and let $w := \text{val}(\mathcal{G})$. Suppose that the components of w are all shortlex reduced and that \hat{w} is a (λ, c) -quasigeodesic that does not backtrack for some constants $\lambda \geq 1$ and $c \geq 0$. Then, in time polynomial in $|\mathcal{G}|$, we can construct an SLP with value $\text{nf}(w)$.*

Proof. We assume that \mathcal{G} is trimmed and in Chomsky normal form and that all of its components have roots. So, in particular, for all variables A of \mathcal{G} , $\text{val}_{\mathcal{G}}(A)$ arises (possibly multiple times) as a subword of w . The hypotheses allow us to apply Proposition 3.5 with $v = w$, $u = \text{nf}(w)$, and $w_1 = w_2 = 1$. Let p and q be the paths in Γ that are labelled by $\text{nf}(w)$ and w , and \hat{p} and \hat{q} the corresponding paths in $\hat{\Gamma}$. Then p and q e' -fellow travel, where e' is a constant that depends

only on G , Σ , λ and c , and all vertices on \hat{q} have at least one corresponding vertex on \hat{p} . As in the proof of Proposition 3.5, we define $\mathbf{e}_1 := \mathbf{e}(\lambda, c, 0)$ and $\mathbf{e}_2 := \mathbf{e}(\lambda, c, \mathbf{e}_1)$, where $\mathbf{e}()$ is as defined in Proposition 3.1

Now consider some instance of $\text{val}(A)$ as a subword w_1 of w that is derived from A and labels a subpath q_1 of q . As we observed in Remark 7.3, either (this instance of) w_1 is a concatenation of complete components of w and subwords in $(\Sigma \setminus \mathcal{H})^*$, or else $w_1 \in \Sigma_i^*$ for some i and w_1 is a proper subword of a component of w .

In the first of these cases (Case 1), the start and end vertices of subpath q_1 of q are also vertices of \hat{q} , and the fact that vertices on \hat{q} have corresponding vertices on \hat{p} implies that there exist $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq \mathbf{e}'$ such that $\text{nf}(\alpha \text{val}_{\mathcal{G}}(A) \beta^{-1})$ is a subword of $\text{nf}(w)$ that labels a subpath p_1 of p whose start and end vertices are also vertices of \hat{p} . (Note that we are using Remark 3.9 here, which ensures that the subword of $\text{nf}(w)$ labelled by this subpath of p is nf -reduced.)

In the second case we see from the proof of Proposition 3.5 that, if (Case 2a) the component of w of which w_1 is a proper subword has length greater than \mathbf{e}_2 , then there exist $\alpha, \beta \in \Sigma_i^*$ with $|\alpha|, |\beta| \leq \mathbf{e}'$ such that $\text{nf}(\alpha \text{val}_{\mathcal{G}}(A) \beta^{-1})$ labels a subword of a connected component of $\text{nf}(w)$. (Here we use the fact that the connected components are shortlex reduced words and so their subwords are also shortlex reduced and hence nf -reduced.)

In Cases 1 and 2a, we say that this instance of the subword $w_1 = \text{val}_{\mathcal{G}}(A)$ of w *corresponds to* the subpath of p that is labelled by $\text{nf}(\alpha \text{val}_{\mathcal{G}}(A) \beta^{-1})$. If w_1 is a proper subword of a component of w of length at most \mathbf{e}_2 (Case 2b), then we do not attempt to define a corresponding subpath of p .

We shall construct an SLP \mathcal{S} with value $\text{nf}(w)$. As usual, we do this by processing the variables of \mathcal{G} in order of increasing height. For each variable A of \mathcal{G} and each pair of words $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq \mathbf{e}'$, we attempt to define a variable $A_{\alpha, \beta}$ of \mathcal{S} with $\text{val}_{\mathcal{S}}(A_{\alpha, \beta}) = \text{nf}(\alpha \text{val}_{\mathcal{G}}(A) \beta^{-1})$. We shall not necessarily succeed in doing this for every such α, β , but we shall do so at least for each α, β for which $\text{val}(A)$ corresponds to $\text{nf}(\alpha \text{val}_{\mathcal{G}}(A) \beta^{-1})$ as described above. After carrying out this process for all variables A of \mathcal{G} , we complete the proof by letting the start variable of \mathcal{S} be $S_{\epsilon, \epsilon}$ for $S = S_{\mathcal{G}}$.

Suppose first that either A has height 0, or that $\text{val}(A) \in \Sigma_i^*$ for some i and $|\text{val}(A)| \leq \mathbf{e}_2$. Then we compute $\text{nf}(\alpha \text{val}(A) \beta^{-1})$ for all $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq \mathbf{e}'$. Since $|\alpha \text{val}(A) \beta^{-1}|$ is bounded above by the constant $\mathbf{e}_2 + 2\mathbf{e}'$, we can do this in time bounded by a constant, by using the word acceptor in the asynchronous automatic structure of G to enumerate the words in the accepted language in order of increasing length until one is found that is equal in G to $\alpha \text{val}(A) \beta^{-1}$, and then define $A_{\alpha, \beta}$ with $\rho_{\mathcal{S}}(A_{\alpha, \beta}) = \text{nf}(\alpha \text{val}(A) \beta^{-1})$.




Figure 3: Subwords of fellow travelling words w and $\text{nf}(w)$

Otherwise, we have $\rho_{\mathcal{G}}(A) = BC$ for variables B and C of \mathcal{G} that we have processed already. We proceed as follows for each pair $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq e'$. For all words $\gamma \in \Sigma^*$ with $|\gamma| \leq e'$ for which we have defined variables $B_{\alpha, \gamma}$ and $C_{\gamma, \beta}$ for \mathcal{S} , we check whether the word $\text{val}_{\mathcal{S}}(B_{\alpha, \gamma})\text{val}_{\mathcal{S}}(C_{\gamma, \beta})$ lies in the language of the word acceptor, which we can do in polynomial time by Proposition 4.1 (v). If so, then we introduce the new variable $A_{\alpha, \beta}$ for \mathcal{S} , define $\rho_{\mathcal{S}}(A_{\alpha, \beta}) = B_{\alpha, \gamma}C_{\gamma, \beta}$, and move on to the next pair of words α, β .

We need to show that we succeed in defining $A_{\alpha, \beta}$ whenever A corresponds to $\text{nf}(\alpha \text{val}_{\mathcal{G}}(A)\beta^{-1})$ as described above. Let $w_1 := \text{val}_{\mathcal{G}}(A)$, $w_2 := \text{val}_{\mathcal{G}}(B)$ and $w_3 := \text{val}_{\mathcal{G}}(C)$, so $w_1 = w_2w_3$. Since we are assuming that $|\text{val}_{\mathcal{G}}(A)| > e_2$, all instances of w_1 as subwords of w must lie in Case 1 or Case 2a, and (since all components of \mathcal{G} have roots) we see then that the same applies to the instances of w_2 and w_3 that arise from the derivation $\rho(A) = BC$. So the vertex of Γ corresponding to the end of the subword $\text{val}_{\mathcal{G}}(B)$ is at distance at most e' from a vertex in the corresponding subword $u_1 := \text{nf}(\alpha \text{val}_{\mathcal{G}}(A)\beta^{-1})$ of $\text{nf}(w)$ such that the associated subwords u_2 and u_3 of u_1 correspond to w_2 and w_3 , respectively. See Figure 3.

In other words, there exists $\gamma \in \Sigma^*$ with $|\gamma| \leq e'$ such that $\text{val}_{\mathcal{G}}(B)$ corresponds to $\text{nf}(\alpha \text{val}_{\mathcal{G}}(B)\gamma^{-1})$ and $\text{val}_{\mathcal{G}}(C)$ corresponds to $\text{nf}(\gamma \text{val}_{\mathcal{G}}(C)\beta^{-1})$, in which case the variables $B_{\alpha, \gamma}$ and $C_{\gamma, \beta}$ have been defined, and so we successfully define $A_{\alpha, \beta}$ with $\rho_{\mathcal{S}}(A_{\alpha, \beta}) = B_{\alpha, \gamma}C_{\gamma, \beta}$. \square

Note that the SLP \mathcal{S} that is constructed in the above lemma can have size up to $(e')^2$ times the size of \mathcal{G} . The following result, guarantees a size that is bounded as a function of the length of the derived word \hat{w} .

Proposition 7.6. *Let \mathcal{G} be an SLP over our selected generating set Σ for G , and let $w := \text{val}(\mathcal{G})$. Then we can construct an SLP \mathcal{S} with value equal to $\text{nf}(w)$ and with size at most $\max(C|\hat{w}| \log(|w|), 1)$ in time that is bounded by $j(|\hat{w}|, |\mathcal{G}|)$, for some constant C and some (increasing) polynomial function $j()$.*

Proof. The proof is by induction on $|\hat{w}|$. We shall not attempt to specify the constant C in the statement precisely, but it would be possible to do so by

following the calculation at the end of the proof. This constant will depend on $|\Sigma|$ and on the constant D from Proposition 3.8. Similarly we shall not attempt to specify the polynomial $j(\cdot)$ in the statement, but merely observe that all steps in the proof can be done in time polynomial in $|\hat{w}|$ and $\log(|w|)$, and that there are $|\hat{w}|$ steps in the inductive proof. By Proposition 7.1 we may assume that all components of \mathcal{G} have roots.

There is nothing to prove if $|\hat{w}| = 0$. Otherwise, we can write $\hat{w} = \hat{v}\hat{z}$ with $|\hat{z}| = 1$, where also $w = vz$, and so $w =_G vz$, and we can assume that we have found an SLP \mathcal{S}_v with $\text{val}(\mathcal{S}_v) = \text{nf}(v)$, where the size of \mathcal{S}_v is bounded as required. We denote by v_{nf} the word $\text{nf}(v) = \text{val}(\mathcal{S}_v)$ and by w_{nf} the word $\text{nf}(w)$. Then $w_{\text{nf}} =_G v_{\text{nf}}z$.

Suppose first that z represents an element of H_i for some component subgroup H_i (so $z \in \Sigma_i^*$) and that v_{nf} ends in a letter from Σ_i . Then let y be the H_i -component of v_{nf} that contains this letter; so we have $v_{\text{nf}} = uy$ and $\widehat{v_{\text{nf}}} = \hat{u}\hat{y}$, and it follows from the description of \mathcal{L} in Proposition 3.7 that $\text{nf}(u) = u$. The word u might not be a prefix of w , but $w =_G u(yz)$, and \hat{u} does not end in a letter from Σ_i . In this situation, we replace v_{nf} by u , derive \mathcal{S}_u with value $u = \text{nf}(u)$ from \mathcal{S}_v , and replace z by yz .

So in any case, we may now assume that, if $z \in \Sigma_i^*$ for some i , then v_{nf} does not end in a letter from Σ_i . Let $z_{\text{nf}} := \text{nf}(z)$. So $z_{\text{nf}} = \text{slex}(z)$ when $z \in \Sigma_i^*$; otherwise z consists of a generator that does not lie in any parabolic subgroup and, since we are assuming that the elements of Σ represent distinct elements of G , we have $|z| = 1$ and $z_{\text{nf}} = z$. In the former case we can compute an SLP \mathcal{S}_z with $\text{val}(\mathcal{S}_z) = z_{\text{nf}}$ and $|\mathcal{S}_z| \leq \max(4|\Sigma_i| \log_2(z_{\text{nf}}), 1)$ by Proposition 5.2. So we can in any case compute an SLP \mathcal{G} with value $v_{\text{nf}}z_{\text{nf}}$.

Now the path \hat{p} in $\widehat{\Gamma}$ labelled by $\widehat{v_{\text{nf}}z_{\text{nf}}} = \widehat{v_{\text{nf}}}\widehat{z_{\text{nf}}}$, as an extension of a geodesic by an element of $\widehat{\Sigma}$, is a $(1, 2)$ -quasigeodesic. We claim that \hat{p} does not backtrack. For suppose that it does. Then the final edge of \hat{p} must penetrate the same H_i -coset as one of the other edges of \hat{p} , for some index i , and then we must have $z \in H_i$. So some suffix of $\widehat{v_{\text{nf}}}$ represents an element of H_i and, since the word $\widehat{v_{\text{nf}}}$ is geodesic, this suffix must have length 1. But then $\widehat{v_{\text{nf}}}$ ends in a letter in Σ_i , contrary to assumption.

We can now apply Lemma 7.5 to the SLP \mathcal{G} with value $v_{\text{nf}}z_{\text{nf}}$ and compute an SLP \mathcal{S}_0 with value $w_{\text{nf}} := \text{nf}(v_{\text{nf}}z)$. By Proposition 7.1, we can modify \mathcal{S}_0 as necessary (in polynomial time), to ensure that all of its components have roots.

We now explain how to modify \mathcal{S}_0 such that its size is bounded by $C|\hat{w}| \log(|w|)$ for some constant C . Let $w_{\text{nf}}[i_r : j_r]$ for $1 \leq r \leq t$ be the components of w_{nf} , and let A_r be the root of $w_{\text{nf}}[i_r : j_r]$ in \mathcal{S}_0 . Then by Lemma 5.1 we can construct an SLP \mathcal{S}_r with $\text{val}(\mathcal{S}_r) = \text{val}(A_r)$ and $|\mathcal{S}_r| \leq \max(4|\Sigma| \log_2(|\text{val}(\mathcal{S}_r)|), 1)$, and define $B_r := \mathcal{S}_r$. Note that $|\mathcal{S}_r|$ is bounded by a multiple of $\log(|w_{\text{nf}}|)$, which is itself bounded by a multiple of $\log(|w|)$ by Proposition 3.8. We now define \mathcal{S} ,

whose variables consist of S_S together with all of the variables of each S_r with $r \geq 1$. For each $A \in S_r$, we define $\rho_S(A) := \rho_{S_r}(A)$, and further

$$\rho(S_S) = w_{\text{nf}}[j_0 : i_1]B_1w_{\text{nf}}[j_1 : i_2]B_2 \cdots w_{\text{nf}}[j_{t-1} : i_t]B_tw_{\text{nf}}[j_t, i_{t+1}),$$

where $j_0 := 0$ and $i_{t+1} := |w_{\text{nf}}|$. Then the SLP \mathcal{S} satisfies the required bound on its size. \square

8. Converting TCSLPs and TSLPs to SLPs

The proof of our main result falls into two parts, the first part constructing, from the input SLP \mathcal{G} , a TCSLP that defines the nf-reduced representative $\text{nf}(\text{val}(\mathcal{G}))$ of its value, and then the second part converting that TCSLP into a SLP with the same value; this is the same strategy as was applied to prove the corresponding result [8, Theorem 6.7] for hyperbolic groups, and our proofs of the component results are based on the proofs in [8].

The construction of a TCSLP accepting $\text{nf}(\text{val}(\mathcal{G}))$ is described in the final section of the paper, Section 9. This section is devoted to the proof of the following conversion theorem:

Theorem 8.1. *Let G be a group hyperbolic relative to a collection of free abelian subgroups, and suppose that a generating set Σ for G , and integer L are selected as in Sections 6.1, 6.3. Let \mathcal{T} be an nf-reduced non-splitting TCSLP for G over Σ , with $J_{\mathcal{T}} \leq L$. Suppose further that each cut operator of \mathcal{T} is non-splitting and specified relative to compression. Then we can construct, in time polynomial in $|\mathcal{T}|$, an nf-reduced SLP \mathcal{S} over Σ with $\text{val}(\mathcal{S}) = \text{val}(\mathcal{T})$, whose size is bounded by a polynomial in $|\mathcal{T}|$.*

We choose to stress the polynomial bound on the size of \mathcal{S} , although it follows from the polynomial bound on time.

The proof of Theorem 8.1 is split into the two results that follow, Propositions 8.2 and 8.3. The first of these computes from a given nf-reduced TSLP \mathcal{U} an SLP \mathcal{S} with the same value, and the second computes an nf-reduced TSLP \mathcal{U} from an nf-reduced TCSLP \mathcal{T} . This follows the strategy of the proof of the corresponding result in [8], and our proofs adapt those of the components of that proof.

There are complications in our proofs that do not arise in [8], resulting partly from the fact that we are using Proposition 6.3 rather than [8, Lemma 4.4], and partly because many of the upper bounds on lengths of words are bounds on their lengths as words over $\widehat{\Sigma}$ rather than over Σ .

Proposition 8.2. *Let G, Σ and L be as in Theorem 8.1. Then, given an nf-reduced non-splitting TSLP \mathcal{U} for G over Σ with $J_{\mathcal{U}} \leq L$, we can construct, in*

time polynomial in $|\mathcal{U}|$, an SLP \mathcal{S} over Σ with $\text{val}(\mathcal{S}) = \text{val}(\mathcal{U})$, whose size is bounded by a polynomial function of $|\mathcal{U}|$.

Proof. Let f be the linear function in the conclusion of Proposition 6.3. Modifying f as necessary, we can assume that f is an increasing function with $f(n) \geq n$ for all n .

The result is trivial if $\text{val}(\mathcal{U}) = \epsilon$, so suppose not. We defined the height of a variable in \mathcal{U} in Section 4.1. We also pointed out that the right hand sides $\rho(A)$ that lie in $(V_{\mathcal{U}} \cup \Sigma)^*$ may be assumed to have the form $a \in \Sigma$ or BC with $B, C \in V_{\mathcal{U}}$, and we shall assume that here.

We define the *tether-height* $\mathfrak{t}(A)$ and *tether-depth* $\mathfrak{d}(A)$ of a variable A , via

$$\begin{aligned} \mathfrak{t}(A) &:= \begin{cases} 0 & \rho_{\mathcal{U}}(A) \in \Sigma^* \\ \max\{\mathfrak{t}(B), \mathfrak{t}(C)\} & \text{if } \rho_{\mathcal{U}}(A) = BC \\ \mathfrak{t}(B) + 1 & \rho_{\mathcal{U}}(A) = B\langle\alpha, \beta\rangle, \end{cases} \\ \mathfrak{d}(A) &:= \mathfrak{t}(S) - \mathfrak{t}(A) + 1, \end{aligned}$$

where $S = S_{\mathcal{U}}$. By removing unused variables, we may assume that S has maximal height and maximal tether-height, so that every variable has positive tether-depth.

In the proof, it is convenient to assume that A , B and C have the same tether-depths in all productions of type $A \rightarrow BC$. To achieve this, we can increase the tether-depth of any variable A , if necessary, by introducing a new redundant variable X together with a production $X \rightarrow A\langle\epsilon, \epsilon\rangle$. This will not affect the maximality of the height and tether-height of S .

We process the variables of \mathcal{U} in order of increasing height. Since the number of variables of \mathcal{U} is certainly bounded by $|\mathcal{U}|$, in order to get the bounds we need on time and space, it will be sufficient to bound the time spent processing each variable A , together with the length of right-hand sides added during each such step.

As we process each variable A of \mathcal{U} we shall either define a copy of A within \mathcal{S} or a set of bounded size of new variables for \mathcal{S} that are associated with A , and for each new variable of \mathcal{S} that we introduce we shall define its right-hand side in \mathcal{S} .

Our construction will ensure the following, for any variable A of \mathcal{U} , where $w := \text{val}_{\mathcal{U}}(A)$. (Recall that, since \mathcal{U} is nf-reduced, $w = \text{nf}(w)$.)

- (i) If $|\hat{w}| \leq (8\mathfrak{d}(A) + 1)f(2L)$, then \mathcal{S} contains a copy of A , and the SLP \mathcal{S}_A is computed from \mathcal{U}_A in time bounded by $j(|\hat{w}|, |\mathcal{U}|)$, and has size at most $C|\hat{w}|\log(|w|)$, where C and j are the constant and polynomial of Proposition 7.6. Note that Proposition 4.4 implies that the size of \mathcal{S}_A is bounded by a polynomial function of the input size $|\mathcal{U}|$ in this case.

- (ii) If $|\hat{w}| > (8d(A)+1)f(2L)$, then w decomposes as a concatenation $\ell_A w' r_A$ with $|\widehat{w'}| \geq f(2L)$, and

$$4f(2L)d(A) \leq |\widehat{\ell_A}|, |\widehat{r_A}| \leq (4d(A) + h(A))f(2L).$$

As A is processed, new variables A_ℓ and A_r are adjoined to \mathcal{S} as the roots of SLPs with values ℓ_A, r_A (with size at most $C|\widehat{\ell_A}| \log(|\ell_A|)$ and $C|\widehat{r_A}| \log(|r_A|)$), as well as variables $A_{\alpha,\beta}$ as the roots of SLPs with values $\text{nf}(\alpha w' \beta^{-1})$, for each $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$.

Each of the subwords ℓ_A, w' and r_A of w is a union of complete components of w and subwords in $(\Sigma \setminus \mathcal{H})^*$, and is in normal form.

Suppose that, while processing the variables of \mathcal{U} in increasing order of height, we have reached the variable A of \mathcal{U} . Let $w := \text{val}_{\mathcal{U}}(A)$.

We consider three different possible cases.

Case 1. $\rho_{\mathcal{U}}(A) \in \Sigma$. We define a variable A within \mathcal{S} , and define $\rho_{\mathcal{S}}(A) := a$.

Case 2. $\rho_{\mathcal{U}}(A) = BC$ for variables B, C of \mathcal{U} . Recall that we are assuming that $d(A), d(B), d(C)$ are all equal; let d be their common value. Let $u := \text{val}_{\mathcal{U}}(B)$, $v := \text{val}_{\mathcal{U}}(C)$, so that $w = uv$. Since we are assuming that \mathcal{U} is non-splitting, either $uv \in \Sigma_i^*$ for some i , in which case $|\hat{u}|, |\hat{v}|, |\hat{w}| \leq 1$, or $\hat{w} = \hat{u}\hat{v}$.

Suppose first (Case 2.1) that $|\hat{u}|, |\hat{v}| > (8d+1)f(2L)$ (so we do not have $uv \in \Sigma_i^*$). Then, when we processed the variables B, C , we computed SLPs with values ℓ_B, r_B, ℓ_C, r_C such that:

- (i) $4df(2L) \leq |\widehat{\ell_B}|, |\widehat{r_B}| \leq (4d + h(B))f(2L)$,
- (ii) $4df(2L) \leq |\widehat{\ell_C}|, |\widehat{r_C}| \leq (4d + h(C))f(2L)$,
- (iii) $u = \ell_B u' r_B$ and $v = \ell_C v' r_C$ where $|\widehat{u'}|, |\widehat{v'}| \geq f(2L)$.

Moreover, for all words $\eta, \theta \in \Sigma^*$ with $|\eta|, |\theta| \leq L$, we defined variables $B_{\eta,\theta}$ and $C_{\eta,\theta}$ in \mathcal{S} whose values in \mathcal{S} are $\text{nf}(\eta u' \theta^{-1})$ and $\text{nf}(\eta v' \theta^{-1})$, respectively.

We now define $\ell_A := \ell_B, r_A := r_C$, and $w' := u' r_B \ell_C v'$; so that $w = \ell_A w' r_A$. Note that, since $d(A) = d(B) = d(C) = d$ and $h(B), h(C) \leq h(A)$, we have (from (i) and (ii)) the length constraints $4f(2L)d(A) \leq |\widehat{\ell_A}|, |\widehat{r_A}| \leq (4d(A) + h(A))f(2L)$. And from the lower bounds on $|\widehat{u'}|, |\widehat{v'}|$ in (iii) we can deduce the required bound $|\widehat{w'}| \geq f(2L)$.

We already have SLPs with values ℓ_A and r_A . It remains to define the right-hand sides for the variables $A_{\alpha,\beta}$ for all $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$.

For each such α, β , and for all $\eta, \theta \in \Sigma^*$ with $|\eta|, |\theta| \leq L$, we compute (using Proposition 7.6) an SLP \mathcal{Z} with value $z := \text{nf}(\eta r_B \ell_C \theta^{-1})$; since \hat{z} has length bounded by a constant multiple of $|\widehat{V_{\mathcal{U}}}|$, this computation takes time bounded by




Figure 4: Case 2.1

a polynomial in $|\mathcal{U}|$, and $|\mathcal{Z}|$ is similarly bounded. Then we check (in polynomial time) using Proposition 4.1(v) whether the word

$$\text{val}_{\mathcal{S}}(B_{\alpha,\eta}) z \text{val}_{\mathcal{S}}(C_{\theta,\beta}) = \text{nf}(\alpha u' \eta^{-1}) \text{nf}(\eta r_B \ell_C \theta^{-1}) \text{nf}(\theta v' \beta^{-1})$$

is nf-reduced, in which case it is the word $\text{nf}(\alpha u' r_B \ell_C v' \beta^{-1}) = \text{nf}(\alpha w' \beta^{-1})$; see Figure 4.

We claim that there must be at least one pair η, θ for which this holds. To see this, we apply Proposition 6.3 twice. First we apply it to the quadrilateral with sides labelled $\alpha, w', \beta, \text{nf}(\alpha w' \beta^{-1})$, using $|\widehat{u'}|, |\widehat{v'}| \geq f(2L) \geq f(L)$, to define η , and then to the quadrilateral with sides labelled $\eta, r_B \ell_C v', \beta, \text{nf}(\eta r_B \ell_C v' \beta^{-1})$ using $|\widehat{r_B \ell_C}|, |\widehat{v'}| \geq f(L)$ to define θ . Proposition 6.3 ensures that $|\theta|, |\eta| \leq L'$, and since in Section 6.3 we chose $L \geq L'$, we certainly have $|\theta|, |\eta| \leq L$. We then include the variables of the SLP \mathcal{Z} within \mathcal{S} and define $\rho_{\mathcal{S}}(A_{\alpha,\beta}) := B_{\alpha,\eta} S_{\mathcal{Z}} C_{\theta,\beta}$.

Suppose next (Case 2.2) that $|\widehat{u}| > (8d+1)f(2L)$ and $|\widehat{v}| \leq (8d+1)f(2L)$ (so again we do not have $uv \in \Sigma_i^*$). (Case 2.3 where $|\widehat{u}| \leq (8d+1)f(2L)$ and $|\widehat{v}| > (8d+1)f(2L)$ is similar, and we shall omit the details.) Then we have already computed an SLP with value v , and SLPs for words ℓ_B and r_B where:

- (i) $4df(2L) \leq |\widehat{\ell_B}|, |\widehat{r_B}| \leq (4d+h(B))f(2L)$,
- (ii) $u = \ell_B u' r_B$ for a word u' with $|\widehat{u'}| \geq f(2L)$.

Moreover, for all words $\eta, \theta \in \Sigma^*$ with $|\eta|, |\theta| \leq L$, we have defined variables $B_{\eta,\theta}$ with value $\text{nf}(\eta u' \theta^{-1})$.

If $|\widehat{v}| \leq f(2L)$, then we set $\ell_A := \ell_B$, $w' := u'$, and $r_A := r_B v$, so that $A_{\ell} = B_{\ell}$ and A_r is the root of an SLP with value $r_B v$, and we define $\rho_{\mathcal{S}}(A_{\alpha,\beta}) := B_{\alpha,\beta}$ for all $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$; so the step takes polynomial time and adds right-hand sides of bounded total length. In this case, we have $4df(2L) \leq |\widehat{\ell_A}| \leq (4d+h(B))f(2L) \leq (4d+h(A))f(2L)$ and $4df(2L) \leq |\widehat{r_A}| \leq (4d+h(B)+1)f(2L) \leq (4d+h(A))f(2L)$, as required.

So now assume that $|\widehat{v}| > f(2L)$. Again, we set $\ell_A := \ell_B$. Since we are assuming that \mathcal{U} is non-splitting and $\rho(A) = BC$, an occurrence of $\text{val}_{\mathcal{U}}(B)$ as a prefix of




Figure 5: Case 2.2

an occurrence of $\text{val}_{\mathcal{U}}(A)$ cannot split a component, and so we have $\widehat{r_B v} = \widehat{r_B} \widehat{v}$. Since $|\widehat{r_B v}| \geq (4d+1)\mathfrak{f}(2L)$ we can define r_A as the suffix of $r_B v$ with $|\widehat{r_A}| = 4\mathfrak{d}\mathfrak{f}(2L)$; that is, $r_B v = y r_A$ for some word y .

Since $\widehat{r_B v} = \widehat{r_B} \widehat{v}$ and $|\widehat{r_B}| \geq 4\mathfrak{d}\mathfrak{f}(2L) = |\widehat{r_A}|$, we have $|\widehat{y}| = |\widehat{r_B}| + |\widehat{v}| - |\widehat{r_A}| \geq |\widehat{v}| > \mathfrak{f}(2L)$. (Recall that r_B and v are in normal form, and hence $\widehat{r_B}$ and \widehat{v} are geodesic.) Then $w = \ell_A w' r_A$ with $w' = u' y$. This satisfies the required bounds on the lengths of $\widehat{\ell_A}$, $\widehat{r_A}$ and $\widehat{w'}$. We can use Proposition 7.1 and Corollary 7.4 to define SLPs with values y and r_A , in time polynomial in $|\mathcal{U}|$; we set A_r to be the root of the second of these.

It remains to define the right-hand sides for the variables $A_{\alpha, \beta}$ for all words $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$. Now, for all $\eta \in \Sigma^*$ with $|\eta| \leq L$, we compute (using Proposition 7.6) an SLP \mathcal{Z} with value $z := \mathfrak{nf}(\eta y \beta^{-1})$; this is done in polynomial time and \mathcal{Z} has polynomially bounded size. Then we check whether the word

$$\text{val}_{\mathcal{S}}(B_{\alpha, \eta}) z = \mathfrak{nf}(\alpha u' \eta^{-1}) \mathfrak{nf}(\eta y \beta^{-1})$$

is \mathfrak{nf} -reduced, in which case it is $\mathfrak{nf}(\alpha u' y \beta^{-1}) = \mathfrak{nf}(\alpha w' \beta^{-1})$; see Figure 5. Since $|\widehat{u}| > (8d+1)\mathfrak{f}(2L) > \mathfrak{f}(L)$ and $|\widehat{y}| > \mathfrak{f}(2L) \geq \mathfrak{f}(L)$, Proposition 6.3 implies that there must be at least one such η with $|\eta| \leq L' \leq L$. We include the variables of the SLP \mathcal{Z} within \mathcal{S} and define $\rho_{\mathcal{S}}(A_{\alpha, \beta}) := B_{\alpha, \eta} \mathcal{S}_{\mathcal{Z}}$.

Finally (Case 2.4), suppose that $|\widehat{u}|, |\widehat{v}| \leq (8d+1)\mathfrak{f}(2L)$ (and hence $|\widehat{w}| \leq 2(8d+1)\mathfrak{f}(2L)$). Note that we could now have $w \in \Sigma_i^*$ for some i . In this case, we have already computed SLPs with values u and v when we processed B, C . If $|\widehat{w}| \leq (8d+1)\mathfrak{f}(2L)$ (which is certainly true when $w \in \Sigma_i^*$) then we can define an SLP for w as the concatenation of those for u and v and attach it to \mathcal{S} with A as its root. We can then use Proposition 7.6 to ensure that this SLP satisfies the required bound on its size.

Otherwise we factorise w as $w = \ell_A w' r_A$ with $|\widehat{\ell_A}| = |\widehat{r_A}| = 4\mathfrak{d}\mathfrak{f}(2L)$, and thus $(8d+2)\mathfrak{f}(2L) \geq |\widehat{w'}| \geq \mathfrak{f}(2L)$. We use Proposition 7.1 and Corollary 7.4 to define SLPs with values ℓ_A and r_A , and attach those to \mathcal{S} . For $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$ we compute (using Proposition 7.6), an SLP \mathcal{Z} with value $\mathfrak{nf}(\alpha w' \beta^{-1})$; Then we attach \mathcal{Z} within \mathcal{S} with $A_{\alpha, \beta}$ as its root, so that $\text{val}_{\mathcal{S}}(A_{\alpha, \beta}) = \mathfrak{nf}(\alpha w' \beta^{-1})$. This computation takes time bounded by a polynomial function of its input

size which, as a result of our size restriction on the SLPs \mathcal{S}_B and \mathcal{S}_C is in turn bounded by a polynomial function of $|\mathcal{U}|$; so the same applies to the size of \mathcal{Z} .

Case 3. $\rho_{\mathcal{U}}(A) = B\langle\sigma, \tau\rangle$ for a variable B of \mathcal{U} and words $\sigma, \tau \in \Sigma^*$ with $|\sigma|, |\tau| \leq L$. Let $u := \text{val}_{\mathcal{U}}(B)$ and $w := \text{val}_{\mathcal{U}}(A) = \text{nf}(\sigma u \tau^{-1})$. Let $\mathbf{d} := \mathbf{d}(B)$. Then $\mathbf{d}(A) = \mathbf{d} - 1 \geq 1$, and so $\mathbf{d} \geq 2$.

If $|\hat{u}| \leq (8\mathbf{d} + 1)\mathbf{f}(2L)$ (Case 3.1), then we have already computed an SLP with value u . In that case, we proceed much as we did in Case 2.4, as follows. Using Proposition 7.6, we construct an SLP for $\sigma u \tau^{-1}$ as the concatenation of three SLPs (the word σ , our SLP for u , and the word τ), and then construct from this an SLP \mathcal{Z} for w . This computation takes time bounded by $j((8\mathbf{d} + 1)\mathbf{f}(2L) + 2L)$ and \mathcal{Z} has size bounded by $C|\hat{w}| \log(|w|)$. If $|\hat{w}| \leq (8\mathbf{d} + 1)\mathbf{f}(2L)$, then we just attach \mathcal{Z} to \mathcal{S} with A as its root.

Otherwise, we factorise w as $w = \ell_A w' r_A$ and, just as in the second part of Case 2.4, we define SLPs with values ℓ_A and r_A , and attach those to \mathcal{S} and, for each $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$, we compute an SLP with value $\text{nf}(\alpha w' \beta^{-1})$, and attach it to \mathcal{S} with the new variable $A_{\alpha, \beta}$ as its root, so that $\text{val}_{\mathcal{S}}(A_{\alpha, \beta}) = \text{nf}(\alpha w' \beta^{-1})$.

Otherwise (Case 3.2), we have $|\hat{u}| > (8\mathbf{d} + 1)\mathbf{f}(2L)$. We have computed SLPs for words ℓ_B, r_B such that $4\mathbf{d}\mathbf{f}(2L) \leq |\widehat{\ell}_B|, |\widehat{r}_B| \leq (4\mathbf{d} + \mathbf{h}(B))\mathbf{f}(2L)$ and $u = \ell_B u' r_B$ for a word u' with $|\hat{u}'| \geq \mathbf{f}(2L)$. Moreover, for all words $\eta, \theta \in \Sigma^*$ with $|\eta|, |\theta| \leq L$, we have defined variables $B_{\eta, \theta}$ with value $\text{nf}(\eta u' \theta^{-1})$.

We check for all $\eta, \theta \in \Sigma^*$ with $|\eta|, |\theta| \leq L$ whether

$$\text{nf}(\sigma \ell_B \eta^{-1}) \text{val}_{\mathcal{S}}(B_{\eta, \theta}) \text{nf}(\theta r_B \tau^{-1}) = \text{nf}(\sigma \ell_B \eta^{-1}) \text{nf}(\eta u' \theta^{-1}) \text{nf}(\theta r_B \tau^{-1})$$

is nf -reduced, in which case it is $\text{nf}(\sigma \ell_B u' r_B \tau^{-1}) = \text{nf}(\sigma u \tau^{-1}) = w$. Since $|\widehat{\ell}_B|, |\widehat{r}_B|, |\hat{u}'| \geq \mathbf{f}(L)$, we can show that such words η, θ exist by two applications of Proposition 6.3 as we did in Case 2.1.

Let $s := \text{nf}(\sigma \ell_B \eta^{-1})$ and $t := \text{nf}(\theta r_B \tau^{-1})$, so that $w = s \text{nf}(\eta u' \theta^{-1}) t$. We deduce from the triangle inequality that $|\widehat{\ell}_B| \leq |\hat{\sigma}| + |\hat{s}| + |\hat{\eta}| \leq |\hat{s}| + 2L$ and similarly that $|\widehat{r}_B| \leq |\hat{t}| + 2L$, and hence we have $|\hat{s}|, |\hat{t}| \geq 4\mathbf{d}\mathbf{f}(2L) - 2L \geq (4\mathbf{d} - 1)\mathbf{f}(2L) \geq 7\mathbf{f}(2L)$. (Recall that $\mathbf{d} \geq 2$.) Hence we can factorise these words as $s = vx$ and $t = yz$ with $|\hat{v}| = |\hat{z}| = 4(\mathbf{d} - 1)\mathbf{f}(2L) = 4\mathbf{d}(A)\mathbf{f}(2L) \geq 4\mathbf{f}(2L)$, and $|\hat{x}|, |\hat{y}| \geq 3\mathbf{f}(2L)$.

We define $\ell_A := v$ and $r_A := z$; these words satisfy the required bounds on their lengths. Note that $\text{val}_{\mathcal{U}}(A) = \text{nf}(\sigma u \tau^{-1}) = \ell_A w' r_A$ with $w' := x \text{nf}(\eta u' \theta^{-1}) y$, and $|\widehat{w}'| \geq 6\mathbf{f}(2L) \geq \mathbf{f}(2L)$. As in earlier cases we can apply Proposition 7.1 and Corollary 7.4 to compute SLPs with values $v = \ell_A$, x , y , $z = r_A$, and w' .

It remains to define the right-hand sides of the variables $A_{\alpha, \beta}$ with values $\text{nf}(\alpha w' \beta^{-1})$ for all words $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq L$. The lower bounds on




Figure 6: Case 3

the lengths of $\hat{v}, \hat{x}, \hat{y}, \hat{z}$ allow us to apply Proposition 6.3 to the quadrilaterals with sides labelled σ, ℓ_B, η, vx and θ, r_B, τ, yz , respectively. We can compute in polynomial time words $\mu, \nu \in \Sigma^*$ with $|\mu|, |\nu| \leq L$ and factorisations $\ell_B = v'\mu$, $r_B = y'z'$ such that $\sigma v' =_G v\mu$, $x\eta =_G \mu x'$, $\theta y' =_G y\nu$, and $\nu z' =_G z\tau$. By the triangle inequality, the words x' and y' satisfy $|\hat{x}'|, |\hat{y}'| \geq f(2L)$.

Now consider the quadrilateral with sides labelled $x'u'y'$, $\text{nf}(\alpha\mu)$, $\text{nf}(\beta\nu)$, and $\text{nf}(\alpha\mu x'u'y'\nu^{-1}\beta^{-1})$. Since $|\hat{x}'|, |\hat{y}'|, |\hat{u}'| \geq f(2L)$ and $|\text{nf}(\alpha\mu)|, |\text{nf}(\beta\nu)| \leq 2L$, we can again make two applications of Proposition 6.3 to show that there exist words $\chi, \psi \in \Sigma^*$ with $|\chi|, |\psi| \leq L' \leq L$ such that the word

$$\begin{aligned} \text{nf}(\alpha\mu x'\chi^{-1}) \text{val}_{\mathcal{S}}(B_{\chi,\psi}) \text{nf}(\psi y'\nu^{-1}\beta^{-1}) = \\ \text{nf}(\alpha\mu x'\chi^{-1}) \text{nf}(\chi u'\psi^{-1}) \text{nf}(\psi y'\nu^{-1}\beta^{-1}) \end{aligned}$$

is nf -reduced, in which case the above word is $\text{nf}(\alpha\mu x'u'y'\nu^{-1}\beta^{-1}) = \text{nf}(\alpha w'\beta^{-1})$; see Figure 6. As before, we can find these words χ, ψ in polynomial time. Finally (using Proposition 7.6), we define SLPS \mathcal{Y} and \mathcal{Z} with values $\text{nf}(\alpha\mu x'\chi^{-1})$ and $\text{nf}(\psi y'\nu^{-1}\beta^{-1})$, include their variables within \mathcal{U} , and then define $\rho_{\mathcal{S}}(A_{\alpha,\beta}) := S_{\mathcal{Y}} B_{\chi,\psi} S_{\mathcal{Z}}$. This concludes the definition of the right-hand sides for the variables $A_{\alpha,\beta}$.

We complete the definition of the SLP \mathcal{S} by adding a start variable $S_{\mathcal{S}}$ to \mathcal{S} and setting $\rho_{\mathcal{S}}(S_{\mathcal{S}}) := S_{\ell} S_{\epsilon,\epsilon} S_r$, where $S := S_{\mathcal{U}}$ is the start variable of \mathcal{U} and S_{ℓ} and S_r are the variables with values $\ell_{\mathcal{S}}$ and $r_{\mathcal{S}}$. This ensures $\text{val}(\mathcal{S}) = \ell_{\mathcal{S}} \text{nf}(s') r_{\mathcal{S}}$, where s' is such that $\ell_{\mathcal{S}} s' r_{\mathcal{S}} = \text{val}(\mathcal{U})$. But we are assuming that \mathcal{U} is nf -reduced, so s' is also nf -reduced and we get $\text{val}(\mathcal{S}) = \ell_{\mathcal{S}} \text{nf}(s') r_{\mathcal{S}} = \ell_{\mathcal{S}} s' r_{\mathcal{S}} = \text{val}(\mathcal{U})$. \square

Proposition 8.3. *Let G, Σ and L be as in Theorem 8.1. Then, given an nf -reduced non-splitting TCSLP \mathcal{T} for G over Σ with $J_{\mathcal{T}} \leq L$, such that each of its cut operators is non-splitting and specified relative to compression, we can compute in polynomial time an nf -reduced non-splitting TSLP \mathcal{U} with $J_{\mathcal{U}} \leq L$ and $\text{val}(\mathcal{U}) = \text{val}(\mathcal{T})$, whose size is bounded by a polynomial function of $|\mathcal{T}|$.*

Proof. We follow the proof of [8, Lemma 6.5]. The idea of the proof is taken

from [7], where it is shown that a CSLP can be transformed in polynomial time into an SLP with the same value. Let $\mathcal{T} = (V_{\mathcal{T}}, S_{\mathcal{T}}, \rho_{\mathcal{T}})$ be the input TCSLP. As discussed in Section 4.1, we can assume that all right-hand sides from $(V_{\mathcal{T}} \cup \Sigma)^*$ are of the form $a \in \Sigma$ or BC with $B, C \in V_{\mathcal{T}}$.

Consider a variable A such that $\rho_{\mathcal{T}}(A) = B[[: i]]$; the case in which $\rho_{\mathcal{T}}(A) = B[[i :)]$ can be handled analogously. By considering the variables in order of increasing height, we can assume that no cut operator occurs in the right-hand side of any variable A' with $h(A') < h(A)$. Using Proposition 8.2 we can compute an SLP with value $\text{val}_{\mathcal{T}}(B)$ and then use Corollary 7.4 to compute $n_B := |\widehat{\text{val}_{\mathcal{T}}(B)}|$ in polynomial time.

Now we show how to eliminate the cut operator in $\rho_{\mathcal{T}}(A)$. This involves adding at most $h(\mathcal{T})$ new variables to the TCSLP. Moreover the height of the TCSLP after the cut elimination will still be bounded by $h(\mathcal{T})$. Hence, the final TSLP has at most $h(\mathcal{T}) \cdot |V|$ variables, and its size is polynomially bounded.

The idea of the cut elimination is to push the cut operator towards variables of lesser height. For this, we need to consider the various possibilities for the right-hand side of B .

Case 1. $\rho_{\mathcal{T}}(B) = a \in \Sigma$. If $i = 1$ we define $\rho_{\mathcal{U}}(A) := a$, and if $i = 0$ we define $\rho_{\mathcal{U}}(A) := \epsilon$.

Case 2. $\rho_{\mathcal{T}}(B) = CD$ with $C, D \in V$. Since we are assuming that the cut operator in $\rho_{\mathcal{T}}(A) = B[[: i]]$ is non-splitting, $\text{val}_{\mathcal{U}}(B)$ cannot consist of a single component, and so $\text{val}_{\mathcal{U}}(C)$ and $\text{val}_{\mathcal{U}}(D)$ must consist of complete components of $\text{val}_{\mathcal{U}}(B)$ together with subwords in $(\Sigma \setminus \mathcal{H})^*$. Define $n_C := |\widehat{\text{val}_{\mathcal{T}}(C)}|$. If $i \leq n_C$ then we define $\rho_{\mathcal{U}}(A) := C[[: i]]$. If $i > n_C$ then we define $\rho_{\mathcal{U}}(A) := CX$ for a new variable X and set $\rho_{\mathcal{U}}(X) := D[[: i - n_C]]$. We then continue with the elimination of the remaining cut operator in $C[[: i]]$ or in $D[[: i - n_C]]$.

Case 3. $\rho_{\mathcal{T}}(B) = C\langle \alpha, \beta \rangle$ with $C \in V$ and $\alpha, \beta \in \Sigma^*$ with $|\alpha|, |\beta| \leq J_{\mathcal{T}} \leq L$. Let $u := \text{val}_{\mathcal{T}}(C)$, $v := \text{val}_{\mathcal{T}}(B) =_G \alpha u \beta^{-1}$, and $v_1 = \text{val}_{\mathcal{T}}(A)$. So $|\widehat{v}_1| = i$ and, where we write $v = v_1 v_2$, our condition on the non-splitting of cut operators ensures that not only v_1 but also v_2 must consist of complete components of v together with subwords in $(\Sigma \setminus \mathcal{H})^*$; so $|\widehat{v}_2| = n_B - i$. Thus, we have $\text{val}_{\mathcal{T}}(A) = v_1$ and $v = \text{nf}(\alpha u \beta^{-1})$. By Proposition 8.2 we can assume that we have computed SLPs with values u and v and we may assume by Proposition 7.1 that all components of u and v have roots in these SLPs.

Suppose first that $|\widehat{v}_1|, |\widehat{v}_2| \geq f(L)$. (This corresponds to Case 3.3 of the proof of [8, Lemma 6.5].) Then by Proposition 6.3 there exists a factorisation $u = u_1 u_2$ and $\eta \in \Sigma^*$ with $|\eta| \leq L$ such that $v_1 =_G \alpha u_1 \eta^{-1}$ and $v_2 =_G \eta u_2 \beta^{-1}$. We note that $|\widehat{\alpha}|, |\widehat{\beta}|, |\widehat{\eta}| \leq L$, and applying the triangle inequality in a quadrilateral within $\widehat{\Gamma}$ with geodesic sides labelled by $\widehat{\alpha}$, \widehat{v}_1 , $\widehat{\eta}$, \widehat{u}_1 , we see that $i - 2L \leq |\widehat{u}_1| \leq i + 2L$, and so we can find such a factorisation of u in polynomial time

by computing SLPs for the words $u[[:j])$ and $u[[j:]$) for every integer j with $|i-j| \leq 2L$. Then we apply Proposition 8.2 and compute SLPs for the words $w_1 := \text{nf}(\alpha u[[:j])\eta^{-1})$ and $w_2 := \text{nf}(\eta u[[j:])\beta^{-1})$ for every $\eta \in \Sigma^*$ with $|\eta| \leq L$. Finally we can, by using Proposition 4.1 (vi), check whether $v_1 = w_1$ and $v_2 = w_2$ (we are guaranteed to find at least one such j and η), and we define $\rho_{\mathcal{U}}(A) := X\langle\alpha, \eta\rangle$ for a new variable X and set $\rho_{\mathcal{U}}(X) := C[[:j])$. We then continue with the elimination of the cut operator in $C[[:j])$.

The other two cases are a little more complicated than in [8]. Suppose first that $|\widehat{v}_1| < f(L)$. Then by Proposition 7.6, in polynomial time we can compute an SLP \mathcal{T}_A with $\text{val}(\mathcal{T}_A) = \text{val}(A) = v_1$. We include the variables and productions of \mathcal{T}_A as part of the TSLP \mathcal{U} that we are constructing, and define $\rho_{\mathcal{U}}(A) := S_{\mathcal{T}_A}$.

Otherwise we have $|\widehat{v}_1| \geq f(L)$ and $|\widehat{v}_2| < f(L)$. Let $k := |\widehat{v}| - f(L)$. Then, as above, we can use Proposition 8.2 together with Proposition 6.3 to find l and a word $\eta \in \Sigma^*$ with $|\eta| \leq L$ such that $v[[:k]) =_G \alpha u[[:l])\eta^{-1}$. Now we introduce a new variable X with $\rho_{\mathcal{U}}(X) = C[[:l])$. But in addition, using Proposition 7.6 we compute, in polynomial time, an SLP \mathcal{T}_A with $\text{val}(\mathcal{T}_A) = v[[k:i])$, include the variables and productions of \mathcal{T}_A in \mathcal{U} , and define $\rho_{\mathcal{U}}(A) := X\langle\alpha, \eta\rangle S_{\mathcal{T}_A}$. Then $\text{val}(A) = v[[:k])v[[k:i]) = v[[:i]) = v_1$ as required. As before, we continue with the elimination of the cut operators below $C[[:l])$.

Since each elimination of a cut operator in $\rho(A)$ can lead to further such eliminations in the variables below A , the total number of such eliminations in the processing of \mathcal{T} is bounded by $h(\mathcal{T})^2$. \square

With the completion of the proofs of Propositions 8.2 and 8.3 we have now completed the proof of Theorem 8.1. The following corollary of that theorem will be used in the final section.

Corollary 8.4. *Let \mathcal{G} be an SLP over Σ for which $w := \text{val}(\mathcal{G})$ is nf-reduced, and let $v \in \Sigma^*$ have length at most L . Then we can, in polynomial time, compute an SLP \mathcal{S} over Σ with $\text{val}(\mathcal{S}) = \text{nf}(wv)$.*

Proof. From \mathcal{G} , we can immediately define a TSLP \mathcal{T} with $\text{val}(\mathcal{T}) = \text{nf}(wv)$ and $J_{\mathcal{T}} = L$ by adjoining a new start variable $S_{\mathcal{T}}$ together with the production $\rho(S_{\mathcal{T}}) = S_{\mathcal{G}}\langle\epsilon, v^{-1}\rangle$. We can then construct \mathcal{S} with $\text{val}(\mathcal{S}) = \text{nf}(wv)$ in polynomial time, by Theorem 8.1. \square

9. The final step.

Since a word represents the identity element if and only if its nf-reduction is the empty word, the main theorem, Theorem A, follows immediately from the combination of the following Theorem 9.1 with Theorem 8.1. Hence the proof of Theorem 9.1 is our final step.

Theorem 9.1. *Let G be a group hyperbolic relative to a collection of free abelian subgroups, and suppose that a generating set Σ for G , and integer L are selected as in Sections 6.1, 6.3. Let \mathcal{G} be an SLP for G over Σ . Then we can construct, in polynomial time, a non-splitting nf-reduced TCSP \mathcal{T} with $\text{val}(\mathcal{T}) = \text{nf}(\text{val}(\mathcal{G}))$ and $J_{\mathcal{T}} \leq L$, where each cut operator of \mathcal{T} is non-splitting and specified relative to compression.*

The following lemma, which is a special case of Theorem 9.1, will be used within its proof, and applied to sub-SLPs of the SLP \mathcal{G} within the statement of the theorem. Since the proof of the lemma involves similar but simpler arguments to that of the theorem, it is convenient to defer its proof.

Lemma 9.2. *Let G , Σ and \mathcal{G} be as in Theorem 9.1, and assume also that $\widehat{\text{val}(\mathcal{G})}$ is a geodesic word. Then in polynomial time we can construct a non-splitting nf-reduced TCSP \mathcal{T} with $\text{val}(\mathcal{T}) = \text{nf}(\text{val}(\mathcal{G}))$ and $J_{\mathcal{T}} \leq L$, where each cut operator of \mathcal{T} is non-splitting and specified relative to compression.*

Proof of Theorem 9.1. We know from Proposition 3.2 that G is asynchronously automatic over Σ so, by Proposition 4.1 (v), we can test in polynomial time whether the words defined by SLPs over Σ are nf-reduced. We may assume by Proposition 4.1 that the given SLP $\mathcal{G} = (V_{\mathcal{G}}, S_{\mathcal{G}}, \rho_{\mathcal{G}})$ is trimmed and in Chomsky normal form, and by Proposition 7.1 that all components of $\text{val}(\mathcal{G})$ have roots.

The proof follows the same strategy as that of [8, Theorem 6.7], but the presence of the parabolic subgroups gives rise to complications. Our aim is to construct a TCSP $\mathcal{T} = (V_{\mathcal{T}}, S_{\mathcal{T}}, \rho_{\mathcal{T}})$ with value $\text{nf}(\text{val}(\mathcal{G}))$ that satisfies $J_{\mathcal{T}} \leq L$, where L is the constant defined in Section 6.3.

We consider the variables of \mathcal{G} in order of increasing height; $V_{\mathcal{T}}$ will contain a copy A of each variable A of \mathcal{G} , together with some auxiliary variables. We build \mathcal{T} piece by piece, starting with $V_{\mathcal{T}}$ empty. At each stage, as we consider the variable A , we add a copy of A and possibly some other new variables to $V_{\mathcal{T}}$, and define the image of $\rho_{\mathcal{T}}$ on each of those so as to make $\text{val}_{\mathcal{T}}(A) = \text{nf}(\text{val}_{\mathcal{G}}(A))$. Our construction of TCSPs rather than CSLPs will ensure a polynomial bound on the number of new variables we add at each stage and on the size of \mathcal{T} . That the condition on cut-operators holds will be clear from the construction.

If the variable A under consideration has height one, then $\rho_{\mathcal{G}}(A) = a$, for some $a \in \Sigma$; in that case, we simply define $\rho_{\mathcal{T}}(A) = \text{nf}(a) = a$. So from now on we suppose that $\text{h}(A) > 1$, in which case $\rho_{\mathcal{G}}(A) = BC$, for variables B, C of height less than $\text{h}(A)$. Since we have already processed the variables B and C , we know that \mathcal{T} already contains sub-TCSPs \mathcal{T}_B and \mathcal{T}_C , with start variables $B_{\mathcal{T}}$ and $C_{\mathcal{T}}$, and with $\text{val}(\mathcal{T}_B) = v_1 := \text{nf}(\text{val}_{\mathcal{G}}(B))$, $\text{val}(\mathcal{T}_C) = v_2 := \text{nf}(\text{val}_{\mathcal{G}}(C))$, and $J_{\mathcal{T}_B}, J_{\mathcal{T}_C} \leq L$.

By Theorem 8.1, we can construct, in polynomial time, SLPs \mathcal{S}_B and \mathcal{S}_C , with the same values as \mathcal{T}_B and \mathcal{T}_C . As was the case for \mathcal{G} , we can ensure that \mathcal{S}_B

and \mathcal{S}_C are in Chomsky normal form, and contain roots for all components of v_1 and v_2 , respectively.

We consider a triangle in the Cayley graph $\Gamma = \Gamma(G, \Sigma)$ whose sides are paths labelled by v_1, v_2 and the nf-reduced representative v_3 of the element $v_1 v_2$. Our aim is to construct \mathcal{T}_A as a TCSP with value v_3 . Let a, b and c be the vertices of this triangle with sides from a to b , b to c and a to c labelled by v_1, v_2, v_3 . So the sides labelled $\widehat{v}_1, \widehat{v}_2, \widehat{v}_3$ in the corresponding triangle in $\widehat{\Gamma}$ are geodesic.

So, since $\widehat{\Gamma}$ is a δ -hyperbolic space, there are meeting vertices d_1, d_2, d_3 , with d_i on $\gamma_{\widehat{v}_i}$ ($i = 1, 2, 3$) and $d_{\widehat{\Gamma}}(d_i, d_j) \leq \delta$ for $i \neq j$; see Figure 7. Now let $K := K_1(\delta)$ as defined in Proposition 6.1, and recall that the constant L defined in Section 6.3 satisfies $L \geq L_1(\delta)$. Now we apply Proposition 6.1 to the sections of $\gamma_{\widehat{v}_1}$ and $\gamma_{\widehat{v}_2}$ that join b to d_1 and d_2 , and so deduce that, for any vertex b_1 of $\widehat{\Gamma}$ on $\gamma_{\widehat{v}_1}$ between d_1 and b and distance on $\gamma_{\widehat{v}_1}$ at least K from d_1 , there exists a vertex b_2 , on $\gamma_{\widehat{v}_2}$, with $d_{\widehat{\Gamma}}(b_1, b_2) \leq L$. We shall call vertices b_1, b_2 of $\widehat{\Gamma}$ with $d_{\widehat{\Gamma}}(b_1, b_2) \leq L$ that lie on two different sides of the triangle *corresponding vertices*. Note that although the particular corresponding vertices b_1, b_2 whose existence we have just shown are found within the sections of $\gamma_{\widehat{v}_1}$ and $\gamma_{\widehat{v}_2}$ that join b to d_1 and d_2 , in general, corresponding vertices might be found past either or both of d_1 and d_2 on those paths.

We claim that it is possible to decide in polynomial time whether a vertex b_1 on $\gamma_{\widehat{v}_1}$ has a corresponding vertex b_2 on $\gamma_{\widehat{v}_2}$ and, if so, find b_2 together with the label in Σ^* of a path of length at most L in Γ joining b_1 to b_2 .

The justification for this claim is as follows. Let l be the distance in $\widehat{\Gamma}$ from b to b_1 . It is straightforward to define an SLP $\overline{\mathcal{S}}_B$ with value $v_1^{-1} = \text{val}(\mathcal{S}_B)^{-1}$. The word v_1^{-1} might not be nf-reduced but its derived word \widehat{v}_1^{-1} is geodesic, so we can use Lemma 9.2 followed by Theorem 8.1 to construct an SLP with value equal to $(\text{nf}(v_1^{-1}))[:, l)$. Since the number of words over Σ of length at most L is bounded above by the constant $|\Sigma|^{L+1}$, we can in polynomial time, by Corollary 8.4, compute SLPs with values $\text{nf}(\overline{\mathcal{S}}_B[:, l])\eta$ for all words $\eta \in \Sigma^*$ of length at most L . For each of these, we compute the length l' of its derived word, and then check whether its value is equal to $\text{val}(\mathcal{S}_C[:, l'])$. If so, then the vertex b_2 at distance l' from b in Γ corresponds to b_1 , and η is the required path label from b_1 to b_2 . Furthermore all such corresponding vertices b_2 are found by this procedure.

We now consider two possible situations, as follows. In Case 1, there is either a vertex a' on $\gamma_{\widehat{v}_2}$ that corresponds to the vertex a of $\gamma_{\widehat{v}_1}$, or there is a vertex c' on $\gamma_{\widehat{v}_1}$ that corresponds to the vertex c of $\gamma_{\widehat{v}_2}$. In Case 2, no such vertices a' or c' exist. By the claim above, we can check which case we are in.

Case 2 is more difficult, so we shall deal with that first and provide a brief description of the argument for Case 1 at the end. So suppose that we are in Case 2. Now the vertex b of $\gamma_{\widehat{v}_1}$ has a corresponding vertex (b itself) on

$\gamma_{\widehat{v}_2}$, but the vertex a has no corresponding vertex on $\gamma_{\widehat{v}_2}$. We need to find corresponding vertices b_1 and b_2 on $\gamma_{\widehat{v}_1}$ and $\gamma_{\widehat{v}_2}$ with the additional property that the vertex b'_1 that is at distance 1 from b_1 in $\widehat{\Gamma}$, on $\gamma_{\widehat{v}_1}$, between b_1 and a , has no corresponding vertex on $\gamma_{\widehat{v}_2}$. We do this, as in the proof of [8, Theorem 6.7], using the technique of *binary search* to find b_1 by testing whether various vertices that we call b_t have corresponding vertices; that is, where l_0 is the length of $\gamma_{\widehat{v}_1}$, we first test if the vertex b_t at distance $l_0/2$ from b along $\gamma_{\widehat{v}_1}$ has a corresponding vertex on $\gamma_{\widehat{v}_2}$. If b_t has a corresponding vertex, and the next vertex b'_t on $\gamma_{\widehat{v}_1}$ has no corresponding vertex on $\gamma_{\widehat{v}_2}$, then we set $b_1 := b_t$ and set $b'_1 := b'_t$,

Otherwise, our next choice for b_t is either the vertex at distance $l_0/4$ from b along $\gamma_{\widehat{v}_1}$ (when our first choice for b_t has no corresponding vertex) or at distance $3l_0/4$ from b (when our first choices for both b_t and b'_t have corresponding vertices). We continue searching in this way, each time in one half of the previous interval of $\gamma_{\widehat{v}_1}$, until we find the vertices b_1, b_2, b'_1 that satisfy the required conditions. Note that the time taken for this search is logarithmic in the length of \widehat{v}_1 and so polynomial in the size of its defining SLP \mathcal{S}_B .

If b'_1 were distance greater than K from d_1 along $\gamma_{\widehat{v}_1}$ within the section joining d_1 to b , then Proposition 6.1 would imply the existence of a corresponding vertex for b'_1 on $\gamma_{\widehat{v}_2}$; but we know there is no such vertex. It follows that

$$d_{\gamma_{\widehat{v}_1}}(a, b'_1) \leq d_{\gamma_{\widehat{v}_1}}(a, d_1) + K - 1, \text{ and } d_{\gamma_{\widehat{v}_1}}(a, b_1) \leq d_{\gamma_{\widehat{v}_1}}(a, d_1) + K.$$

We claim that

$$d_{\gamma_{\widehat{v}_2}}(c, b_2) \leq d_{\gamma_{\widehat{v}_2}}(c, d_2) + 3L + \delta.$$

Since $d_{\widehat{\Gamma}}(d_1, d_2) \leq \delta$ and $d_{\widehat{\Gamma}}(b_1, b_2) \leq L$, this follows from the triangle inequality if $d_{\widehat{\Gamma}}(d_1, b_1) \leq 2L$. So suppose not. Then, since $L \geq K$, b_1 must be between a and d_1 on \widehat{v}_1 ; note that this position for b_1 is *not* as suggested in Figure 7. But, by Proposition 6.1, if $d_{\gamma_{\widehat{v}_2}}(c, b_2) > d_{\gamma_{\widehat{v}_2}}(c, d_2) + K$, then there is a point b''_1 on \widehat{v}_1 between b and d_1 with $d_{\Gamma}(b_2, b''_1) \leq L$, and hence also $d_{\widehat{\Gamma}}(b_2, b''_1) \leq L$. But then we have $2L \geq d_{\widehat{v}_1}(b_1, b''_1) \geq d_{\widehat{v}_1}(b_1, d_1) > 2L$, a contradiction. So $d_{\gamma_{\widehat{v}_2}}(c, b_2) \leq d_{\gamma_{\widehat{v}_2}}(c, d_2) + K$ and, since $K < 3L + \delta$, the claim also holds in this case.

So we have

$$\begin{aligned} d_{\Gamma}(b_1, b_2) &\leq L, & d_{\gamma_{\widehat{v}_1}}(a, b_1) &\leq d_{\gamma_{\widehat{v}_1}}(a, d_1) + K, \\ & & d_{\gamma_{\widehat{v}_2}}(c, b_2) &\leq d_{\gamma_{\widehat{v}_2}}(c, d_2) + 3L + \delta. \end{aligned}$$

Now it follows from Proposition 6.1 that any vertex on $\gamma_{\widehat{v}_1}$ between a and d_1 that is at distance at least K along the curve from d_1 , must have a corresponding vertex on $\gamma_{\widehat{v}_3}$ that lies between a and d_3 . So now define a_1 to be the vertex on $\gamma_{\widehat{v}_1}$ between a and d_1 that is distance $2K$ along the curve from b_1 (if there is no such vertex, then define a_1 to be a). Then there is a vertex a_3 on $\gamma_{\widehat{v}_3}$ between a and d_3 that corresponds to a_1 .




Figure 7: The hyperbolic triangle

Similarly let c_2 be the vertex on $\gamma_{\hat{v}_2}$ between c and d_2 that is distance $K + 3L + \delta$ along the curve from b_2 (if there is no such vertex, then define c_2 to be c). Then there is a vertex c_3 on $\gamma_{\hat{v}_3}$ that corresponds to c_2 .

There exist words ζ , η and θ over Σ each of length at most L , that label paths in Γ from a_3 to a_1 , b_1 to b_2 , and c_3 to c_2 . We know η already, because (by the claim above) we found it when we defined b_1 and b_2 ; to progress further, we need to find ζ and θ . We find these through an exhaustive search process, as we shall now describe. We generate all possible word pairs (ζ, θ) (that is, word pairs of lengths at most L), and check their validity, until we find a solution.

So suppose that (ζ, θ) is a candidate pair. Define k_1, l_1, k_2, l_2 to be the integers for which $v_1[[k_1 : l_1]]$ labels the path in Γ from a_1 to b_1 and $v_2[[k_2 : l_2]]$ the path from b_2 to c_2 ; then $v_1[[k_1 : l_1]] = \text{val}(\mathcal{S}_B[[k_1 : l_1]])$ and $v_2[[k_2 : l_2]] = \text{val}(\mathcal{S}_C[[k_2 : l_2]])$. The nf-reduced representative of $v_1[[k_1 : l_1]]\zeta^{-1}$ is the value of the TCSP $\mathcal{S}_B[[k_1 : l_1]](\epsilon, \zeta)$, and we can find an SLP \mathcal{S}_1 with the same value in polynomial time, by Theorem 8.1.

The word $\zeta v_1[[k_1 : l_1]]\eta v_2[[k_2 : l_2]]\theta^{-1}$ has length at most $3K + 6L + \delta$ over $\widehat{\Sigma}$, and is the value of the CSP $\zeta \mathcal{S}_B[[k_1 : l_1]]\eta \mathcal{S}_C[[k_2 : l_2]]\theta^{-1}$. By Proposition 7.6 applied with $\kappa = 3K + 6L + \delta$, we can (in polynomial time) find an SLP \mathcal{S}_2 with value $\text{nf}(\zeta v_1[[k_1 : l_1]]\eta v_2[[k_2 : l_2]]\theta^{-1})$. The nf-reduced representative of $\theta v_2[[k_2 : l_2]]$ is the value of the TCSP $\mathcal{S}_C[[k_2 : l_2]](\theta, \epsilon)$, and, again, we can find an SLP \mathcal{S}_3 with the same value in polynomial time. Now the word $\text{val}(\mathcal{S}_1 \mathcal{S}_2 \mathcal{S}_3)$ represents the same element of G as each of the words $v_1 v_2$ and v_3 , and will be equal as a word to v_3 if and only if it is nf-reduced.

By Proposition 4.1 (v), we can test in polynomial time whether $\text{val}(\mathcal{S}_1\mathcal{S}_2\mathcal{S}_3)$ is nf-reduced; if it is, then $\mathcal{S}_1\mathcal{S}_2\mathcal{S}_3$ is an SLP for v_3 . We define a TCSLP $\mathcal{T}_1 := \mathcal{T}_B[[: k_1]](\epsilon, \zeta)$ as an extension of \mathcal{T}_B by a single variable $S_{\mathcal{T}_1}$, and similarly $\mathcal{T}_3 := \mathcal{T}_C[[l_2 :]](\theta, \epsilon)$ as an extension of \mathcal{T}_C . The TCSLPs \mathcal{T}_1 and \mathcal{T}_3 have the same values as $\mathcal{S}_1, \mathcal{S}_3$, respectively, and the concatenation $\mathcal{T}_1\mathcal{S}_2\mathcal{T}_3$ is a TCSLP with value v_3 . We set our copy of A within \mathcal{T} to be the start variable of that TCSLP, and define $\rho_{\mathcal{T}}(A)$ accordingly.

(The reason that we do not simply adjoin the SLPs \mathcal{S}_1 and \mathcal{S}_3 to \mathcal{T} is that, if we did that, then we would be unable to prove that the $|\mathcal{T}|$ remains bounded throughout the complete process by a polynomial function of $|\mathcal{G}|$. This explains why we needed to introduce the concepts of TSLP and TCSLP.)

We shall now briefly describe the corresponding argument in Case 1, which is similar to that for Case 2, but simpler. Suppose that there is a vertex c' on $\gamma_{\widehat{v}_1}$ that corresponds to the vertex c of $\gamma_{\widehat{v}_2}$ - the other case is similar. Then, as explained earlier, we can locate c' , we can find $\eta \in \Sigma^*$ that labels a path of length at most L in Γ from c to c' , and we can calculate the integer k_1 such that $v_1[[: k_1]]$ is the prefix of v_1 from a to c' .

Let a_1 be the vertex on $\gamma_{\widehat{v}_1}$ that is between a and c' and is at distance K in $\gamma_{\widehat{v}_1}$ from c' (or $a_1 = a$ if there is no such vertex), and let k_2 be its distance from a along v_1 . Then by Proposition 6.1, a_1 has a corresponding vertex a_3 on $\gamma_{\widehat{v}_3}$. Let $\zeta \in \Sigma^*$ be the label of a path of length at most L in Γ from a_3 to a_1 . Then

$$v_3 = \text{nf}(v_1[[: k_2]]\zeta^{-1})\text{nf}(\zeta v_1[[k_2 : k_1]]\eta^{-1})$$

and, much as in Case 2, we can find ζ by exhaustive search, then define the TCSLP \mathcal{T} as the concatenation $\mathcal{T}_1\mathcal{S}_2$, where $\text{val}(\mathcal{T}_1) = \text{nf}(v_1[[: k_2]]\zeta^{-1})$, and $\text{val}(\mathcal{S}_2) = \text{nf}(\zeta v_1[[k_2 : k_1]]\eta^{-1})$.

In both Cases 1 and 2, we observe that all of the variables of \mathcal{T}_B and \mathcal{T}_C were defined during the processing of other variables, but a copy of A , the variables $S_{\mathcal{T}_1}$ and $S_{\mathcal{T}_2}$ and also the variables of \mathcal{S}_2 are added to $V_{\mathcal{T}}$ during the processing of the variable A of \mathcal{G} , and their images under $\rho_{\mathcal{T}}$ are correspondingly defined. Since the length of \mathcal{S}_2 as a word over $\widehat{\Sigma}$ is bounded by the constant $3K + 6L + \delta$ we know from Proposition 7.6 that its size is bounded by a constant multiple of $\log(|\text{val}(\mathcal{S}_2)|)$ and hence, by Proposition 4.1 (ii), of $|\mathcal{G}|$.

So after all variables of \mathcal{G} have been processed, the size of the final TCSLP \mathcal{T} is bounded by a quadratic function of \mathcal{G} . This completes the proof. \square

Proof of Lemma 9.2. As we observed in Remark 7.3, since all components of $w = \text{val}(\mathcal{G})$ already have roots, it follows from Proposition 7.1 that for each variable A of \mathcal{G} , all occurrences of $\text{val}(A)$ as subwords of w that are derived from A consist either of complete components of w together with subwords in $(\Sigma \setminus \mathcal{H})^*$, or of subwords of a component. So the assumption that \widehat{w} is a geodesic word implies that $\widehat{\text{val}_{\mathcal{G}}(A)}$ is also geodesic for all variables A of \mathcal{G} .

We follow the proof of Theorem 9.1. When we consider the variable A with $\rho(A) = BC$, we know that $\widehat{\text{val}}_G(A) = \widehat{\text{val}}_G(B)\widehat{\text{val}}_G(C)$ is geodesic, and so the concatenation of the consecutive sides ab and bc of the associated hyperbolic triangle in $\widehat{\Gamma}$, which is labelled by the word $\widehat{uv} = \widehat{u}\widehat{v}$, is a geodesic path in Γ .

Proposition 6.1 now implies that the vertices a_1 and c_2 on the two paths at distance K from b in $\widehat{\Gamma}$ have corresponding vertex on $\gamma_{\widehat{v}_3}$.

After defining a_1 and c_2 in this way, the rest of the proof is the same as the proof of Theorem 9.1, with $\eta = \epsilon$. But it is of course important to point out that we have not used the conclusion of the lemma in its proof! \square

References

- [1] J. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro and H. Short, “Notes on word-hyperbolic groups”, in E. Ghys, A. Haefliger and A. Verjovsky, eds., *Proceedings of the Conference “Group Theory from a Geometric Viewpoint” held in I.C.T.P., Trieste, March 1990*, World Scientific, Singapore, 1991.
- [2] Y. Antolín and L. Ciobanu, Finite generating sets of relatively hyperbolic groups and applications to geodesic languages. *Trans. Amer. Math. Soc.*, 368(11):7965–8010, 2016.
- [3] B.H. Bowditch, Relatively hyperbolic groups, *Internat. J. Algebra Comput.*, 22(03):66pp, 2012.
- [4] Moses Charikar, Eric Lehman, April Lehman, Ding Liu, Rina Panigrahy, Manoj Prabhakaran, Amit Sahai, and Abhi Shelat. The smallest grammar problem. *IEEE Transactions on Information Theory*, 51(7):2554–2576, 2005.
- [5] D.B.A. Epstein, J. Cannon, D. Holt, D., S. Levy, M. Paterson, and W. Thurston, W., *Word Processing in Groups*, Jones and Bartlett, Boston, 1992.
- [6] B. Farb, Relatively hyperbolic groups, *Geom. Funct. Anal.* 8(5):810–840, 1998.
- [7] Christian Hagenah. *Gleichungen mit regulären Randbedingungen über freien Gruppen*. PhD thesis, University of Stuttgart, 2000.
- [8] D. Holt, M. Lohrey, S. Schleimer, Compressed decision problems in hyperbolic groups, *arXiv:1808.06886*, Preprint, 2018.
- [9] D. Holt, S. Rees, Regularity of quasigeodesics in a hyperbolic group, *Internat. J. Algebra Comput.* 13(5):585–596, 2003.

- [10] Markus Lohrey. Word problems and membership problems on compressed words. *SIAM Journal on Computing*, 35(5):1210 – 1240, 2006.
- [11] Markus Lohrey. *The Compressed Word Problem for Groups*. SpringerBriefs in Mathematics. Springer, 2014.
- [12] Denis V. Osin. Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems. *Mem. Amer. Math. Soc.*, 179(843):vi+100, 2006.
- [13] Wojciech Plandowski. Testing equivalence of morphisms on context-free languages. In *Proceedings of ESA 1994*, volume 855 of *Lecture Notes in Computer Science*, pages 460–470. Springer, 1994.