Groups with Context-Free Conjugacy Problems

Derek F. Holt, Sarah Rees, Claas E. Röver

September 30th 2009, Newcastle

Abstract

The conjugacy problem and the inverse conjugacy problem of a finitely generated group are defined, from a language theoretic point of view, as sets of pairs of words. An automaton might be obliged to read the two input words synchronously, or could have the option to read asynchronously. Hence each class of languages gives rise to four classes of groups; groups whose (inverse) conjugacy problem is an (a)synchronous language in the given class. For regular languages all these classes are identical with the class of finite groups. We show that finitely generated groups with asynchronously context-free inverse conjugacy problem are precisely the virtually free groups. Moreover, the other three classes arising from context-free languages are shown all to coincide with the class of virtually cyclic groups, which is precisely the class of groups with synchronously one-counter (inverse) conjugacy problem. It is also proved that for a δ-hyperbolic group the intersection of the inverse conjugacy problem with the set of pairs of quasi-geodesics is context-free.

1 Introduction

The relationship between the complexity of the word problem of a group as a formal language and the algebraic characteristics of the group has been well studied; it is well known that the word problem of a group is regular precisely when the group is finite [3] and context-free precisely when the group is virtually-free [9]; and in fact when the word problem is context-free it is even deterministic context-free. The purpose of this article is to investigate the relationship between the complexity of the conjugacy problem of a group as a formal language and the algebraic characteristics of the group.

Throughout this paper we assume that all group generating sets are finite and symmetric; that is, they are closed under taking inverses. For a group G generated by a set X, we define the *conjugacy problem* of G to be the set of ordered pairs of words (u, v) such that u and v are conjugate in G.

We will consider machines that read pairs of words synchronously as strings over an alphabet $X \times X$ as well as machines that read inputs asynchronously from two strings. For synchronous reading, in any reading move of the machine (as opposed to a non-reading or ε-move), a single symbol from each of the two input words is read, where the shorter of the two words is padded at the end with a padding symbol that maps onto the identity element of the group. For asynchronous reading, in each move of the machine, zero or one symbols from each of the two input words may be read, so there are four types of moves altogether. Observe that synchronous reading is a special case of asynchronous reading.

Using additional states as necessary, we can clearly assume that a 2-variable asynchronous machine $\mathcal{M}$ never reads from both input strings in a single move. . This allows us to define a corresponding one variable machine $\mathcal{M}_1$ with alphabet the union $X^L \cup X^R$ of two disjoint copies of X. Define $\rho_L : X^L \cup X^R \to X^*$ by $\rho_L(x^L) = x$ where x^L is the element of X^L corresponding to $x \in X$ and $\rho_L(x^R) = \varepsilon$ for all $x^R \in X^R$, and extend ρ_L to $(X^L \cup X^R)^*$. Define ρ_R similarly. Then, the automaton $\mathcal{M}_1$ that simulates $\mathcal{M}$ should accept w precisely when $\mathcal{M}$ accepts $(\rho_L(w), \rho_R(w))$ by reading the inputs in the order encoded in w. Such an $\mathcal{M}_1$ can obviously be constructed.

We will call a language of pairs of strings *synchronously* or *asynchronously regular* if it is accepted by a finite state automaton reading the two input words synchronously or asynchronously. We define *synchronously* and *asynchronously context-free* similarly, with finite state automaton replaced by pushdown automaton. If the pushdown automaton has only one work symbol besides the bottom of stack marker, then it and the accepted language are called *synchronously* or *asynchronously one-counter*.

Note that the word problem for G is the restriction of the conjugacy problem to pairs (u, ε), where ε is the empty string, and hence the conjugacy problem is never easier than the word problem.

It is straightforward to prove that the conjugacy problem for G is regular if and only if G is finite, and we include a proof in Section 2. Our first main result classifies groups whose conjugacy problem is context-free.

Theorem A *Let G be a finitely generated group. The following are equivalent.*

(i) The conjugacy problem of G is synchronously one-counter.

(ii) The conjugacy problem of G is asynchronously context-free.

(iii) G is virtually cyclic.

Theorem A is proved in Section 3, where it is broken down into Theorems 5, 8 and 10.

By reversing the order of one of the input words we can deal with a wider class of groups, but only with asynchronous automata. We call the set of all pairs of words (u, v) for which u and v^{-1} are conjugate in G, the *inverse conjugacy problem* for G. Our second main result classifies groups whose inverse conjugacy problem is context-free.

Theorem B *Let G be a finitely generated group. Then the inverse conjugacy problem of G is*

(i) synchronously context-free if and only if G is virtually cyclic;

(ii) asynchronously context-free if and only if G is virtually free.

Theorem B is broken down into Theorems 5, 8, 18 and 19, and is the subject of Section 4. Theorem 19 is of independent interest, showing that, for any $\lambda \geq 1$ and $\epsilon \geq 0$ the restriction of the inverse conjugacy problem of a word hyperbolic group to pairs of (λ, ϵ)-quasi-geodesics is asynchronously context-free.

In Section 5 we sketch the following result which implies that with respect to conjugacy problems indexed languages are more powerful than context-free languages.

Theorem C *The conjugacy problem of a finitely generated virtually free group is asynchronously indexed.*

This should be compared to the still open problems of whether indexed languages lead to larger classes of groups than context-free languages in the realms of word and coword problems; see [6].

2 Preliminaries

Let us begin by showing that the choice of generating set makes no difference and hence that we are dealing with group theoretic properties. In a way this justifies our terminology introduced and used before.

Theorem 1 *If the conjugacy problem or inverse conjugacy problem for a group is synchronously or asynchronously context-free with respect to one generating set then the same is true with respect to any other generating set.*

PROOF: Suppose first that the conjugacy or inverse conjugacy problem for the group G is asynchronously context-free with respect to X. Given another generating set Y, for each $y \in Y$, we choose a word w_y in X^* that represents the element y, and hence define a homomorphism from Y^* to X^*. Given any pair of input words (u, v) over Y, the automaton to solve the (inverse) conjugacy problem first rewrites over X using that homomorphism and then operates as the machine solving the (inverse) conjugacy problem over X.

This part of the argument works for any class of languages closed under inverse homomorphisms.

Now suppose that the (inverse) conjugacy problem for G is synchronously context-free with respect to X. Given another generating set Y, define the words w_y just as above. If the words w_y all had the same length, we could proceed as above. Otherwise by padding as necessary, as we rewrite, with words of the form xx^{-1} we can ensure that the rewrites of the prefixes of u, v which we have read so far differ in length by at most 1 at any stage. Hence using some additional memory to deal with that small shift we can adapt our solution to the problem over X to give a solution over Y.

This part of the argument has made some small assumptions about the power of the machine, and hence it is not immediately clear that it works for any class of language closed under inverse homomorphism. $\square$

We include the following elementary result for completeness.

Theorem 2 *A group has asynchronously regular conjugacy problem if and only if it is finite, in which case the problem is synchronously regular.*

Proof: If the conjugacy problem is (synchronously or asynchronously) regular, then so is the word problem as the intersection of the conjugacy problem with the regular set $X^* \times \{\varepsilon\}$, and hence a group with regular conjugacy problem must be finite.

Conversely suppose that G is finite, in which case we may store the conjugacy classes of its elements. Using a finite state automaton based on the Cayley graph for $G \times G$, we can identify the elements of G represented by two input words u, v and hence decide their conjugacy. $\square$

Of course the same result holds for the inverse conjugacy problem.

3 Conjugacy Problems

In this section we begin by solving the conjugacy and inverse conjugacy problem for virtually cyclic groups on a one-counter automaton Then we show that these are the only groups with synchronously context-free conjugacy or inverse conjugacy problem. This is an application of the classical pumping lemma for context-free languages.

Lemma 3 (Pumping Lemma [5, Lemma 6.1]) Add exact lemma reference *Let L be a context-free language. Then there exists a constant $N > 0$ such that if $w \in L$ with $|w| > N$, then $w = \upsilon\phi\chi\psi\omega$ where $|\phi\chi\psi| \leq N$, $|\phi\psi| \geq 1$, and $\upsilon\phi^n\chi\psi^n\omega \in L$ for all $n \geq 0$.*

Finally we prove that groups with asynchronous context-free conjugacy problem must also be virtually cyclic, making use of Ogden's generalisation of the pumping lemma.

Lemma 4 (Ogden's Lemma [5, Lemma 6.2]) *Let L be a context-free language. Then there exists a constant $N > 0$ such that if $w \in L$, with at least N letters in w marked, then $w = \upsilon\phi\chi\psi\omega$ where at most N letters in $\phi\chi\psi$ are marked, and at least one in $\phi\psi$, and $\upsilon\phi^n\chi\psi^n\omega \in L$ for all $n \geq 0$.*

3.1 Virtually Infinite Cyclic Groups

Theorem 5 *Every virtually cyclic group has synchronously one-counter conjugacy problem, and synchronously one-counter inverse conjugacy problem.*

PROOF: Since regular languages are also one-counter, we may and will assume that $G = \langle X \rangle$ is virtually infinite cyclic. Let z be a generator of a normal cyclic subgroup of finite index in G, and let T be a set of coset representatives for $\langle z \rangle$ in G. Then every element of G is uniquely representable by a word $z^k t$ for $k \in \mathbb{Z}, t \in T$, and we refer to this word as its normal form.

Since T and X are finite we may assume that the machine knows the normal forms for all r^t, r^{-1} and rx, for $r, t \in T$ and $x \in X$.

First we consider the conjugacy problem. The conjugate of $z^i r$ by $z^k t$, $(z^k t)^{-1} z^i r z^k t$, can take one of four forms:

(a)	$z^i r^t$,	when r, t both centralise z
(b)	$z^{-i} r^t$,	when r centralises but t inverts z
(c)	$z^{i-2k} r^t$,	when t centralises but r inverts z
(d)	$z^{2k-i} r^t$,	when r, t both invert z.

We are now ready to describe the operation of our machine.

Suppose words u, v over X are input, and that u and v have normal forms $z^i r$ and $z^j s$ with $r, s \in T$. Before starting to read, the machine will guess one of the four cases above as well as a specific choice of t, and attempt to verify that v is a conjugate of u of that type. On a one-counter automaton, one can rewrite a single input word into normal form using the states to store the coset representative and the stack to store the exponent of z. But for a pair of words we cannot compute both exponents using a single stack. In fact, we compute both coset representatives r, s using the states. In addition we compute either $i-j$ in case (a), $i+j$ in case (b), or the parity of $i-j$ in cases (c) or (d); we only need to use the stack in the first two cases.

After reading u and v, the machine knows r and s. It checks that r^t and s are in the same coset of $\langle z\rangle$, and also whether each of r, t centralise or invert z, as specified by the selected case (a), (b), (c) or (d). If either of these checks fails then the computation is abandoned.

Otherwise, $r^t = z^l s$, for some $l \in \mathbb{Z}$. It remains to check the value of l against the information we have computed for i and j. To verify that u, v are conjugate by an element of the form $z^k t$, we check in case (a) that $i-j=-l$, in case (b) that $i+j=l$ and in cases (c) and (d) that the parity of $i-j$ matches that of l.

The machine recognising the inverse conjugacy problem operates very similarly. We assume as above that u, v are on the two tapes, and that u and v have normal forms $z^i r$ and $z^j s$. We aim to verify that u is conjugate to v^{-1}. We have $v^{-1} = z^{-\epsilon j}s^{-1}$, where ϵ is -1 or 1 depending on whether or not conjugation by s inverts z. Let $s^{-1} = z^m s'$, where $s' \in T$. Note that if u and v^{-1} are indeed conjugate, then r, s, s' either all centralise or all invert z. So ϵ is equal to 1 in cases (a) and (b), and -1 in cases (c) and (d).

In all cases r^t and s^{-1} must be in the same coset of $\langle z\rangle$; otherwise we abandon the computation. This time we check in case (a) that $i+j=m-l$, in case (b) that $j-i=m-l$ and in cases (c) and (d) that the parity of $i-j$ matches that of $m-l$. □

3.2 Cancellation Lemmas and Transversals

We recall that a reduced word $x_1x_2\cdots x_r$ is said to be *cyclically reduced* if $x_1 \neq x_r^{-1}$. So the free reduction of any word has the form $\alpha^{-1}\beta\alpha$ with β cyclically reduced, and we shall call β the *cyclic reduction* of the original word. By [8, Theorem 1.3] two words over a free generating set of a free group are conjugate in the group if and only if their cyclic reductions are cyclic conjugates of each other.

For the remainder of the article, $u \equiv v$ will mean that u and v are equal as words, and $u =_G v$ means that u and v are G-equivalent; that is, they represent the same element of the group G. We denote the length of a word w by $|w|$.

Lemma 6 *Let u and v be cyclically reduced words over a free generating set of a free group, such that $u \neq v^{-1}$ and neither u nor v is a proper power. Let $m = \max(|u|, |v|)$. Then, for $k, l \geq 1$, the cancelled suffix of u^k and prefix of v^l in the free reduction of the product u^kv^l each have length less than $2m$, and hence the freely reduced length of u^kv^l is greater than $k|u| + l|v| - 4m$.*

PROOF: Suppose without loss of generality that $m = |v| \geq |u|$. The result is clear if $l = 1$, so suppose that $l > 1$ and assume for a contradiction that the prefix v^2 of v^l is cancelled in the product $u^k v^l$. Then $v \equiv u^{-r}w$ for some $r \geq 1$ with $|w| < |u|$ and, since $v \neq u^{-1}$ and v is not a proper power, we have $|w| > 0$. But now, $v^2 \equiv u^{-r}wu^{-r}w$, and since its subwords wu^{-r} and v have equal length and each is completely cancelled when multiplied on the left by u^{k+r}, we must have $wu^{-r} \equiv v$. So v is a nontrivial cyclic conjugate of itself, and hence must be a proper power, which is a contradiction. $\square$

We use the following technical lemma in the proofs of both Theorem 8 and Theorem 10.

Lemma 7 *Suppose that G has a non-abelian free normal subgroup F of finite index, and that a, b are part of a free basis Y for F. Then there exists a transversal T for F in G such that for each $t \in T$*

(i) *the word over Y for $a^t := t^{-1}at$ is a cyclically reduced word α,*

(ii) *the word over Y for b^t has the form $\rho^{-1}\beta\rho$ where β is cyclically reduced, and $|\rho|$ is as small as possible.*

For all such t there is no cancellation in the products $\alpha\rho^{-1}$ or $\rho\alpha$.

PROOF: Each coset Ft contains representatives for which (i) is satisfied. For if a^t is not cyclically reduced, then it has the form $\sigma^{-1}\alpha\sigma$ with α cyclically reduced, $\sigma \in F$, and $t\sigma^{-1}$ satisfies (i). From amongst all those representatives of Ft satisfying (i) we choose one satisfying (ii).

Now if there were cancellation in forming either of the products $\alpha\rho^{-1}$ or $\rho\alpha$, then we could replace α by a cyclic conjugate and reduce the length of ρ. $\square$

3.3 Synchronous (Inverse) Conjugacy Problems

Theorem 8 *A finitely generated group with synchronously context-free conjugacy problem or inverse conjugacy problem is virtually cyclic.*

PROOF: First suppose that the group G has synchronously context-free conjugacy problem. Then the word problem for G is context-free, and hence G is virtually free [9]. Suppose that G is not virtually cyclic. Then G has a finite index free normal subgroup F of rank at least 2. Let a, b be part of a free generating set Y for F, and let $A = a^{-1}, B = b^{-1}$. We choose a generating set X for G that contains Y. So, by Theorem 1, the conjugacy problem of G is synchronously context-free with this generating set.

We choose a transversal T for F in G satisfying the conditions of Lemma 7.

Now, for some suitably large k, let $u = a^k b^k a^k B^k$ and $v = b^k a^k B^k a^k$. Then u and v are conjugate in F, and so $w = (a,b)^k(b,a)^k(a,B)^k(B,a)^k$ is in the conjugacy problem. According to the pumping lemma (Lemma 3), we can write w in the form $\upsilon\phi\chi\psi\omega$ where $|\phi\chi\psi| \leq N$, $|\phi\psi| \geq 1$, and for all $n \geq 0$, $(u_n, v_n) := \upsilon\phi^n\chi\psi^n\omega$ is in the language, and hence must also represent a conjugate pair in G. In particular this is true of $(u_0, v_0) := \upsilon\chi\omega$. Our aim is to establish a contradiction by showing that u_0 and v_0 are not conjugate.

Provided that $k > N$, the subword $\phi\chi\psi$ must be contained within one of the three subwords $(a,b)^k(b,a)^k$, $(b,a)^k(a,B)^k$ or $(a,B)^k(B,a)^k$ of w. We shall only deal with the first of these three possibilities; the arguments dealing with the other two cases are similar.

In this first case we have $u_0 \equiv a^{k-r}b^{k-s}a^kB^k$, $v_0 \equiv b^{k-r}a^{k-s}B^ka^k$, where at least one of r, s is greater than 0, and $r + s \leq N < k$. It is clear that u_0 and v_0 are not cyclically conjugate, and hence are not conjugate within F; we still need to verify that they are not conjugate within the larger group G.

Suppose $g \in G$ with $g^{-1}u_0g = v_0$. Then $g = tx$ for some $t \in T$ and $x \in F$, so the cyclically reduced length of $t^{-1}u_0t$ is equal to $|v_0| = |u_0|$. By choice of the transversal T, we have $t^{-1}at =_G \alpha$ and $t^{-1}bt =_G \rho^{-1}\beta\rho$, as described in Lemma 7. If ρ is nonempty then, since there is no cancellation in the products $\alpha\rho^{-1}$ and $\rho\alpha$, $t^{-1}u_0t =_G \alpha^{k-r}\rho^{-1}\beta^{k-s}\rho\alpha^k\rho^{-1}\beta^{-k}\rho$ is clearly cyclically reduced, with length greater than $|u_0|$. Hence ρ is empty and $\alpha =_G t^{-1}at$ and $\beta =_G t^{-1}bt$ are both cyclically reduced words. Furthermore, since α and β form part of a free basis of F, we cannot have $\alpha = \beta^{\pm 1}$ and neither α nor β is a proper power.

Now by Lemma 6 at most $4\max\{|\alpha|, |\beta|\}$ generators cancel in any product $\alpha^i\beta^j$ or $\beta^j\alpha^i$, and hence the length of the cyclic reduction of $t^{-1}u_0t =_G \alpha^{k-r}\beta^{k-s}\alpha^k\beta^{-k}$ is at least $(2k-r)|\alpha| + (2k-s)|\beta| - 16\max\{|\alpha|, |\beta|\}$. Provided that k is sufficiently large, this can only be equal to $|v_0| = 4k - r - s$ if $|\alpha| = |\beta| = 1$. In that case, $t^{-1}u_0t$ is a cyclic conjugate in F of v_0, so we must have $\alpha, \beta \in \{a, b, A, B\}$. (We can calculate how large k needs to be for this purpose after the choice of the transversal T, which determines all the words α and β that can arise.)

But it is easily seen that no word formed by substituting any of a, b, A, B for a and b in u_0 is a cyclic conjugate of v_0. So u_0 is not conjugate to v_0 and we have our contradiction. Once the other two cases have been similarly dealt with we conclude that G must be virtually cyclic.

The proof of the fact that groups with synchronously context-free inverse conjugacy problem are virtually free is very similar. We suppose that G has context-free inverse conjugacy problem, and that F is a free finite index normal subgroup of rank at least two with generators $a, b, \ldots$. But now we choose $u = a^kb^ka^kB^k$ and $v = A^kb^kA^kB^k$. Then u and v^{-1} are conjugate in

F, so (u, v) is in the inverse conjugacy problem. From then on the argument is the same. □

3.4 Asynchronous Conjugacy Problems

We now proceed to generalise Theorem 8 to groups with asynchronously context-free conjugacy problem. The proof is similar to that of Theorem 8, but considerably more difficult technically. In particular, we need a further lemma along the lines of Lemma 6.

Lemma 9 *Let u and v be cyclically reduced words over a free generating set Y of a free group F, such that neither u nor v is a proper power. Let $m = \max(|u|, |v|)$, and let w be any freely reduced word over Y. Let $k, l \geq 1$. Suppose that either*

(a) u is not a cyclic conjugate of v^{-1}, or

(b) $u =_F v^{-1}$ and w is not a power of u.

Then the freely reduced length of $u^k w v^l$ is greater than $k|u| + l|v| - |w| - 4m$.

Proof: If the whole of the word w is not cancelled in the product $u^k w v^l$, then the reduced length of the product is at least $k|u| + l|v| - |w|$.

Otherwise, after cancelling just w, the resulting word has length $k|u| + l|v| - |w|$ and is of the form $u^{k'} u_1 v_2 v^{l'}$ where $0 \leq k' \leq k$, $0 \leq l' \leq l$, $u = u_1 u_2$, $v = v_1 v_2$. But this word can also be written as $u_1 (u')^{k'} (v')^{l'} v_2$, where $u' = u_2 u_1$ and $v' = v_2 v_1$ are cyclic conjugates of u and v. In this case, we observe that $w =_F u^{-k+k'+1} u_2^{-1} v_1^{-1} v^{-l+l'+1}$.

Now the result clearly follows from Lemma 6 if $u'v' \neq_F 1$.

In case (a) we certainly have $u'v' \neq_F 1$.

In case (b), we have $uv =_F 1$. Suppose that we also have $u'v' =_F 1$, with u', v' as above. Then $u_2^{-1} v_1^{-1}$ centralises u and so is a power of u, and it follows that w a power of u, contradicting our assumption. Hence also in this case $u'v' \neq_F 1$, and the result follows. □

Theorem 10 *If a group has asynchronously context-free conjugacy problem then it is virtually cyclic.*

Proof: Let G have asynchronously context-free conjugacy problem. As in Theorem 8, G is virtually free. Suppose that G is not virtually cyclic. Then G has a finite index free normal subgroup F of rank at least three. Let a, b, c

be part of a free generating set Y for F, and let $A = a^{-1}, B = b^{-1}, C = c^{-1}$. Again, we choose a generating set X for G that contains Y.

This time, we use Lemma 7 to choose two transversals T_b and T_c of F in G, such that:

(i) for all $t \in T_b \cup T_c$, the reduced word α over Y for $t^{-1}at$ is cyclically reduced;

(ii) for all $t \in T_b$, the reduced word over Y for $t^{-1}bt$ has the form $\rho^{-1}\beta\rho$ with β cyclically reduced and $|\rho|$ as small as possible; and

(iii) for all $t \in T_c$, the reduced word over Y for $t^{-1}ct$ has the form $\tau^{-1}\gamma\tau$ with τ cyclically reduced and $|\tau|$ as small as possible.

(iv) there is no cancellation in the words $\rho\alpha$ or $\alpha\rho^{-1}$ for $t \in T_b$ or in the words $\tau\alpha$ or $\alpha\tau^{-1}$ for $t \in T_c$.

From our earlier discussion on page 2, there is a context-free language L_1 with alphabet $X^L \cup X^R$ of two disjoint copies of X such that two words u and v are conjugate in G if and only if there is a word $w \in L_1$ with $\rho_L(w) = u$ and $\rho_R(w) = v$.

We choose some suitably large k and l with $k \neq l$ and put

$$\zeta = b^k a^l b^k a^l B^k, \quad \eta = c^l a^k c^l a^k C^l, \quad u = \zeta\eta, \quad v = \eta\zeta.$$

Then u, v are conjugate in F (hence in G), and we let $w \in L_1$ be as above. In other words, w is a shuffle of u^L and v^R, copies of u and v over X^L respectively X^R. It is clear that either all of ζ^L occurs before ζ^R or all of η^R occurs entirely before η^L.

We shall suppose that the first of these possibilities occurs. We omit the other case, for which the argument is similar, the main difference being that where we use the transversal T_b in the arguments below, we use T_c in the other case.

So $w \equiv w_1 w_2$, where ζ^L occurs in w_1 and ζ^R within w_2.

We now apply Ogden's Lemma, stated as Lemma 4, to w. We shall assume that k and l have been chosen to be larger than the constant N of the lemma. We mark all the symbols of ζ^L in w_1 and all the symbols of ζ^R in w_2, and no other symbols. (In the case we are omitting, η plays the role of ζ here.)

We conclude that we can write $w \equiv \upsilon\phi\chi\psi\omega$, where $\phi\chi\psi$ has at most N marked symbols, $\phi\psi$ has at least 1 marked symbol, and $\upsilon\phi^n\chi\psi^n\omega \in L_1$ for all $n \geq 0$. In particular, putting $n = 0$, we have $\upsilon\chi\omega \in L_1$, which means that $u_0 = \rho_L(\upsilon\chi\omega)$ and $v_0 = \rho_R(\upsilon\chi\omega)$ must be conjugate in G. Our aim is to establish a contradiction by showing that u_0 and v_0 are not conjugate in G.

The situation is more complicated than in Theorem 8, because we have no information about how many of the unmarked symbols are involved in ϕ and ψ. We distinguish between two cases, the first of which splits into two (similar) subcases.

Case 1a. The substring $\phi\chi\psi$ ends before the final marked symbol in w_1. Then, since $k, l > N$ and $\phi\chi\psi$ has at most N marked symbols, $\rho_L(\phi\chi\psi)$ can only involve letters from within 2 neighbouring blocks of the prefix ζ of u, and hence

$$u_0 \equiv \begin{cases} b^{k-r}a^{l-s}b^k a^l B^k \eta, & \text{or} \\ b^k a^{l-r} b^{k-s} a^l B^k \eta, & \text{or} \\ b^k a^l b^{k-r} a^{l-s} B^k \eta, & \text{or} \\ b^k a^l b^k a^{l-r} B^{k-s} \eta, \end{cases}$$

and

$$v_0 \equiv c^{l_1} a^{k_1} c^{l_2} a^{k_2} C^{l_3} \zeta,$$

where $0 < r + s \le N$, and $0 \le k_i \le k$, $0 \le l_i \le l$ for the relevant i.

Case 1b. The substring $\phi\chi\psi$ begins after the first marked symbol in w_2. This is similar to Case 1a, after inverting u_0 and v_0 and interchanging their rôles, which we are allowed to do, as we only need to show that they are not conjugate.

Case 2. The substring $\phi\chi\psi$ includes the final marked symbol in w_1 or the first marked symbol in w_2 or both. Then

$$u_0 \equiv b^k a^l b^k a^l B^{k-r} c^{l_4} a^{k_3} c^{l_5} a^{k_4} C^{l_6} \text{ and } v_0 \equiv c^{l_1} a^{k_1} c^{l_2} a^{k_2} C^{l_3} b^{k-s} a^l b^k a^l B^k$$

where again $0 < r + s \le N$ and $0 \le k_i \le k$, $0 \le l_i \le l$ for the relevant i.

Case 2 is the more difficult, and we shall discuss it first. Let v_0' be the cyclic reduction of the word v_0. Then

$$v_0' \equiv \begin{cases} v_0 & \text{provided} \quad k_2 \ne 0 \\ c^{l_1} a^{k_1} c^{l_2 - l_3} b^{k-s} a^l b^k a^l B^k & \text{if} \quad k_2 = 0,\ k_1 \ne 0 \\ c^{l_1 + l_2 - l_3} b^{k-s} a^l b^k a^l B^k & \text{if} \quad k_1 = k_2 = 0,\ l_1 + l_2 \ne l_3 \\ a^l b^k a^l B^s & \text{if} \quad k_1 = k_2 = 0,\ l_1 + l_2 = l_3. \end{cases}$$

Let $g \in G$ with $g^{-1} u_0 g =_G v_0$. Then $g =_G tx$ for some $t \in T_b$ and $x \in F$, and $t^{-1} u_0 t$ is conjugate to v_0' within F. So the cyclic reduction u_0' of $t^{-1} u_0 t$ must be a cyclic conjugate of v_0'. We proceed to show that this is impossible.

Let the reduced word in F for $t^{-1} ct$ be $\sigma^{-1} \gamma \sigma$ with γ cyclically reduced.

Then

$$t^{-1} u_0 t \equiv \underline{\rho^{-1} \beta_1^k \rho \alpha_1^l \rho^{-1} \beta_2^k \rho \alpha_2^l \rho^{-1} \beta_3^{-(k-r)} \rho} \sigma^{-1} \gamma^{l_4} \sigma \alpha^{k_3} \sigma^{-1} \gamma^{l_5} \sigma \alpha^{k_4} \sigma^{-1} \gamma^{-l_6} \sigma,$$

with $\alpha_1 \equiv \alpha_2 \equiv \alpha$ and $\beta_1 \equiv \beta_2 \equiv \beta_3 \equiv \beta$; The subscripts are there to enable us to refer unambiguously to the subwords. Note that property (iv) of the transversal T_b ensures that the underlined prefix is freely reduced. We call the non-underlined part θ.

First note that, since α, β and γ are conjugates of the images of the free generators a, b and c, none of them is a cyclic conjugate of any of the others or their inverses. And, for the same reason, none of them is a proper power. Hence any two of them satisfy the conditions of Lemma 9. We shall also use the fact that β cannot be in the normal closure of $\langle \alpha, \gamma \rangle$.

The free reduction of θ either involves some high powers of some or all of α, γ or γ^{-1}, or can be considered to be short.

In the first case, we apply case (a) of Lemma 9 twice to show that there can be only limited cancellation between $\beta_3^{-(k-r)}$ and whatever is the first high power surviving in the free reduction of θ, and between whatever is the last high power surviving in the free reduction of θ and β_1^k.

In the second case, we first note that the fact that β cannot be in the normal closure of $\langle \alpha, \gamma \rangle$ excludes the possibility that $\rho\theta\rho^{-1}$ is a non-trivial power of β. If $\rho\theta\rho^{-1}$ is non-trivial, then we apply case (b) of Lemma 9 to see that there can be only limited cancellation involving $\beta_3^{-(k-r)}$ and β_1^k.

Therefore, precisely one of the following holds.

(A) The free reduction of $\rho\theta\rho^{-1}$ is empty; or

(B) Only a small fraction of β_1^k and $\beta_3^{-(k-r)}$ are cancelled.

Since the underlined prefix of $t^{-1}u_0t$ is freely reduced, in either situation the powers α_1^l, β_2^k and α_2^l survive in their entirety in u_0'.

Now we can deduce that $|\alpha| = |\beta| = 1$ as follows. First, as part of a free basis of F neither α nor β can be a proper power. Then if either has length greater than 1, the number of syllables in u_0' must be at least $\min\{2l, 2k\}$, which is larger than 10, the number of syllables in v_0', so long as k and l are sufficiently large. This contradicts the fact that u_0' is a cyclic conjugate of v_0'.

In Situation (A), $u_0' \equiv \beta^r\rho\alpha^l\rho^{-1}\beta^k\rho\alpha^l\rho^{-1}$ (or $\alpha^l\rho^{-1}\beta^k\rho\alpha^l$ if $r = 0$) and so u_0' contains only three high powers of generators. Clearly, this means that $v_0' \equiv a^lb^ka^lB^s$, in which case ρ must be empty and, since $r + s > 0$, we find that there are no possibilities for α and β that make u_0' and v_0' cyclic conjugates.

So it remains to consider Situation (B). Since the underlined subword is freely reduced, in this case ρ and ρ^{-1} each appear twice in u_0'. But the only negative powers of generators that occur in v_0' occur as a unique power of C and a unique power of B, so ρ must be empty. Hence

$$t^{-1}u_0t =_G \beta^k\alpha^l\beta^k\alpha^l\beta^{-(k-r)}\theta$$

and u_0' must contain a subword of the form $\beta^{k'}\alpha^l\beta^k\alpha^l\beta^{-k''}$, where $k', k'' > 0$.

We now consider the possibilities for α and β, assuming that u_0' is a cyclic conjugate of v_0'. Since both β and β^{-1} occur in u_0' and A does not occur in v_0, β cannot be a or A. Since there are unique powers of B and of C in v_0' but two disjoint powers of β in u_0', we cannot have $\beta = C$ or $\beta = B$. If $\beta = c$, then we would necessarily have $\alpha = a$ and $\alpha^l\beta^k\alpha^l \equiv a^{k_1}c^{l_2}a^{k_3}$ and hence $k_1 = k_3 = l$, $l_2 = k$, which is impossible, since $k_1 \le k$, $l_2 \le l$ and $k \ne l$. So the only possibility is that $\beta = b$ which implies immediately that $\alpha = a$.

Now the subwords $a^lb^ka^l$ of u_0' and v_0' must correspond, so we have

$$a^lb^ka^lB^kc^{l_1}a^{k_1}c^{l_2}a^{k_2}C^{l_3}b^{k-s} =_F a^lb^ka^lB^{(k-r)}\theta b^k \tag{1}$$

and hence

$$c^{l_1}a^{k_1}c^{l_2}a^{k_2}C^{l_3} =_F b^r\theta b^s =_F b^r\sigma^{-1}\gamma^{l_4}\sigma a^{k_3}\sigma^{-1}\gamma^{l_5}\sigma a^{k_4}\sigma^{-1}\gamma^{-l_6}\sigma b^s.$$

Now, working modulo $[F, F]$, we have

$$a^{k_1+k_2}c^{l_1+l_2-l_3} =_{F/[F,F]} a^{k_3+k_4}b^{r+s}\gamma^{l_4+l_5-l_6}.$$

Since the exponent of b on the right hand side must be zero, but $r+s > 0$, we must have $l_4 + l_5 - l_6 \ne 0$. So modulo $[F, F]$, γ must be a product of powers of a, b and c. But the automorphism of $F/[F, F]$ induced by conjugation by t has finite order and fixes a and b, and this is only possible if $\gamma =_{F/[F,F]} c^{\pm 1}$, and hence the exponent of b on the right hand side is non-zero. So Situation (A) is also impossible and hence Case 2 cannot occur.

Case 1 can be ruled out by similar reasoning, which we shall now describe briefly. In Case 1a, we have

$$u_0 \equiv b^{k-r_1}a^{l-s_1}b^{k-r_2}a^{l-s_2}B^{k-r_3}\eta$$

where $0 < r_1 + s_1 + r_2 + s_2 + r_3 \le N$, and either just one term, or two adjacent terms of the sequence r_1, s_1, r_2, s_2, r_3 are nonzero. As in Case 2, we let v_0' be the cyclic reduction of v_0, and we can find $t \in T_b$ such that the cyclic reduction u_0' of $t^{-1}u_0t$ is a cyclic conjugate of v_0'. Again as in Case 2, we can use a syllable count to show that $|\alpha| = |\beta| = 1$ and then that ρ must be empty.

It follows as before that β cannot be a, A, C or B. If $\beta = c$, then we get $\alpha = a$ and $\alpha^{l-s_1}\beta^{k-r_2}\alpha^{l-s_2} \equiv a^{k_1}c^{l_2}a^{k_3}$, so $l-s_1 = k_1$, $k-r_2 = l_2$, $l-s_2 = k_3$, with $k_i \le k$ and $l_i \le l$. This is impossible provided that $|k - l| > N$, and we can assume that k and l have been chosen with this property. So, as in Case 2, we must have $\alpha = a$, $\beta = b$.

Still proceeding as before, Equation (1) above becomes

$$a^l b^k a^l B^k c^{l_1} a^{k_1} c^{l_2} a^{k_2} C^{l_3} b^k =_F$$
$$a^{l-s_1} b^{k-r_2} a^{l-s_2} B^{k-r_3} \sigma^{-1} \gamma^l \sigma a^k \sigma^{-1} \gamma^l \sigma a^k \sigma^{-1} \gamma^{-l} \sigma b^{k-r_1}$$

This implies immediately that $s_1 = r_2 = s_2 = 0$, so exactly one of r_1 and r_3 must be nonzero, in which case we get the same contradiction as in Case 2 by working modulo $[F, F]$. □

4 Inverse Conjugacy Problems

We turn now to the inverse conjugacy problem. The following theorem is a special case of the later result Theorem 18, but we include it here because its proof is more straightforward, and gives an indication of the proof of Theorem 18.

Theorem 11 *A finitely generated free group has asynchronously context-free inverse conjugacy problem.*

PROOF: We use a free generating set for the free group. Then, by [8, Theorem 1.3], the inverse conjugacy problem consists of all pairs of words (u, v) for which the cyclic reduction of u has the form $\alpha\beta$ while the cyclic reduction of v has the form $\alpha^{-1}\beta^{-1}$. That is, the free reduction of u has the form $\gamma\alpha\beta\gamma^{-1}$, while the free reduction of v has the form $\delta\alpha^{-1}\beta^{-1}\delta^{-1}$. In other words u itself can be written as a concatenation of (not necessarily freely reduced) subwords $u_1u_2u_3u_4$, and v as a concatenation of subwords $v_1v_2v_3v_4$, where each of u_1u_4, v_1v_4, u_2v_2 and u_3v_3 freely reduces to the empty word.

We can recognise a pair of words of this form non-deterministically on a machine with a single stack which guesses the endpoints of each of the subwords. We read the subwords in the order $u_1, v_1, u_2, v_2, u_3, v_3, v_4, u_4$. We use the stack to freely reduce and so put onto it the free reduction of u_1, followed by a marker, then the free reduction of v_1, followed by a marker. Then we use the stack above this to verify that u_2v_2 freely reduces to the empty word (and otherwise abort the computation), and then that u_3v_3 freely reduces to the empty word (and otherwise abort). At this point, if the calculation has not been aborted, then the stack contains the free reduction of u_1, followed by a marker, and then the free reduction of v_1 followed by a marker. Deleting those markers as we read on we can check first whether or not v_1v_4 reduces freely to the empty word, and second whether or not u_1u_4 reduces to the empty word. □

Our main aim in this section is to prove that the inverse conjugacy problem of a virtually free group is asynchronously context-free. The method employed will be basically similar to that used in the proof of Theorem 11, but will involve several additional complications.

We need to start by reviewing some basic material from geometric group theory, particularly from the theory of (word-)hyperbolic groups. The following definition and discussion is from [4, Definition I.8.14].

A (λ, ϵ)-*quasi-isometric embedding* with $\lambda \geq 1$ and $\epsilon \geq 0$ is a map $\phi : (\mathcal{X}_1, d_1) \to (\mathcal{X}_2, d_2)$ between two metric spaces such that, for all $x, y \in \mathcal{X}_1$, we have

$$\frac{d_1(x,y)}{\lambda} - \epsilon \leq d_2(f(x), f(y)) \leq \lambda d_1(x,y) + \epsilon.$$

The map ϕ is called a *quasi-isometry* if, in addition, there is a constant $C \geq 0$ such that for all $y \in \mathcal{X}_2$ there exists $x \in \mathcal{X}_1$ with $d_2(y, f(x)) \leq C$. The two spaces are called *quasi-isometric* if there is a quasi-isometry between them; being quasi-isometric is an equivalence relation on metric spaces.

The Cayley graphs of a finitely generated group G with respect to different finite generating sets are quasi-isometric. So we can unambiguously define two finitely generated groups to be quasi-isometric if their Cayley graphs are quasi-isometric. Furthermore, it is easily seen that if H has finite index in G then the embedding of the Cayley graph of H into that of G, using a finite generating set of G that contains one of H, is a quasi-isometry, and hence G and H are quasi-isometric.

For a word w over X with $G = \langle X \rangle$, $|w|_G$ will denote the geodesic length of w as an element of G; that is, the length of a shortest word v over X with $v =_G w$. We say that w is a *geodesic word* in G if $|w| = |w|_G$.

A (λ, ϵ)-*quasigeodesic* in a metric space $(\mathcal{X}, d)$ is a (λ, ϵ)-quasi-isometric embedding $c : I \to \mathcal{X}$ from a (bounded or unbounded) interval I of the real line to X. In particular, a path of length n in the Cayley graph Γ_X of a group $G = \langle X \rangle$ is a (λ, ϵ)-quasigeodesic if the standard mapping from the interval $[0, n]$ to the path is, and a word $x_1 x_2 \cdots x_n$ over X is a (λ, ϵ)-quasigeodesic word in G if the path that it labels is. Assuming that $\lambda \geq 1$ and $\epsilon \geq 0$, this is equivalent to the property that any subword of length k has geodesic length at least $k/\lambda - \epsilon$.

To simplify some of the technicalities, we shall assume when convenient that λ and ϵ are both integers.

Lemma 12 *Let $\phi \colon G \to H$ be a (λ', ϵ')-quasi-isometry (with respect to some word metrics on G and H) and $C, \kappa \geq 1$ constants. Let $g_0, g_1, \ldots, g_n \in G$ and, for $0 \leq i \leq n$, let $h_i \in H$ with $|h_i^{-1}\phi(g_i)|_H \leq C$. Suppose that $h_i \neq h_{i+1}$ for $0 \leq i < n$ and that there exists a (λ'', ϵ'')-quasigeodesic in H through*

$h_0, h_1, \ldots, h_n$ *(in that order). Then every path* $u = \xi_1 \cdots \xi_n$ *in* G *through* $g_0, g_1, \ldots, g_n$ *in which each segment* ξ_i *from* g_{i-1} *to* g_i *satisfies* $|\xi_i| \leq \kappa$ *is a* (λ, ϵ)*-quasigeodesic in* G*, where* $\lambda = \kappa\lambda'\lambda''$ *and* $\epsilon = (\epsilon' + \epsilon'' + 2C + 2/\lambda'')/\lambda' + 2\kappa$*.*

Proof: There is nothing to show if there is no such path u. Otherwise, let v be a subword of u and $g_j, g_{j+1}, \ldots, g_{j+k}$ be the g_i on v. Since $|\xi_i| \leq \kappa$, we have

$$k \geq |v|/\kappa - 2 \quad \text{and} \quad |v|_G \geq |\xi_{j+1} \cdots \xi_{j+k}|_G - 2\kappa.$$

The hypotheses $h_i \neq h_{i+1}$ and $|h_i^{-1}\phi(g_i)|_H \leq C$ together with the existence of a (λ'', ϵ'')-quasigeodesic through $h_j, \ldots, h_{j+k}$ (which must have length at least k) imply that

$$|\phi(g_j)^{-1}\phi(g_{j+k})|_H + 2C \geq |h_j^{-1}h_{j+k}|_H \geq k/\lambda'' - \epsilon''.$$

Finally, the fact that ϕ is a (λ', ϵ')-quasi-isometry guarantees that

$$|\phi(g_j)^{-1}\phi(g_{j+k})|_H/\lambda' - \epsilon'/\lambda' \leq |g_j^{-1}g_{j+k}|_G = |\xi_j \cdots \xi_{j+k}|_G.$$

These four inequalities immediately imply $|v|/\lambda - \epsilon \leq |v|_G$, whence u is a (λ, ϵ)-quasigeodesic by the remark above. □

Theorem 13 *Let* $G = \langle X \rangle$ *be a finitely generated virtually free group. Then there exist constants* $\lambda \geq 1$*,* $\epsilon, \kappa \geq 0$ *with the following property: for any word* $w = x_1 x_2 \cdots x_n$ *over* X*, there is a* G*-equivalent* (λ, ϵ)*-quasigeodesic word* w'*, obtained from* w *by replacing some of its subwords by* G*-equivalent words of length at most* κ*.*

Proof: By assumption, G has a subgroup H of finite index that is free on some generating set Y and, by the remarks above, for some $\lambda' \geq 1$, $\epsilon' \geq 0$, there is a (λ', ϵ')-quasi-isometry ϕ from G to H with respect to the word metrics $|.|_G$ and $|.|_H$ induced by X and Y respectively. Let $w = x_1 x_2 \cdots x_n$ be a word over X, and let $1 = g_0, g_1, \ldots, g_n =_G w$ be the vertices on the path from the identity vertex in Γ_X that is labelled by w, with x_i labelling the edge joining g_{i-1} to g_i.

Since ϕ is a (λ', ϵ')-quasi-isometry, $|\phi(g_{i-1})^{-1}\phi(g_i)|_H \leq \lambda' + \epsilon'$ for each i. Let ξ_i be the (unique) freely reduced word over Y that labels the geodesic from $\phi(g_{i-1})$ to $\phi(g_i)$; so $|\xi_i| \leq \lambda' + \epsilon'$. The free reduction $\bar{\xi}$ of $\xi := \xi_1 \xi_2 \cdots \xi_n$ is the geodesic from $\phi(g_0)$ to $\phi(g_n)$ in H. There may be more than one way of freely reducing ξ by cancelling inverse pairs of generators, but we choose one such method arbitrarily. Then we define $\bar{\xi_i}$ to be the subsequence of those letters of ξ_i which remain in $\bar{\xi}$; this is well defined since a free reduction is simply deletion of a subword.

Let $I = \{i \mid \bar{\xi}_i \not\equiv \varepsilon\} = \{i_1, i_2, \ldots, i_k\}$, and put $i_0 = 0$. Then $\bar{\xi} \equiv \bar{\xi}_{i_1}\bar{\xi}_{i_2}\cdots\bar{\xi}_{i_k}$ and $\xi_i \equiv \xi_i'\bar{\xi}_i\xi_i''$, as ξ_i is freely reduced. Put $h_0 = \phi(g_0)$ and $h_j = h_0\bar{\xi}_{i_1}\bar{\xi}_{i_2}\cdots\bar{\xi}_{i_j}$ and observe that $h_k =_H \phi(g_n)$. For ease of notation, let us redefine i_k to be n without changing h_k.

Since the ξ_i have length at most $C := \lambda' + \epsilon'$ and $\phi(g_{i_j}) =_H h_0\bar{\xi}_{i_1}\bar{\xi}_{i_2}\cdots\bar{\xi}_{i_j}\xi''_{i_j}$ for $1 \le j \le k-1$, we have $|h_j^{-1}h_{j+1}|_H \le C$ and $|\phi(g_{i_j})^{-1}\phi(g_{i_{j+1}})|_H \le 2C$ for $0 \le j < k$ and $|h_j^{-1}\phi(g_{i_j})|_H \le C$ for $0 \le j \le k$. Since ϕ is a (λ', ϵ')-quasi-isometry, it now follows that $|g_{i_j}^{-1}g_{i_{j+1}}|_G \le \kappa := 2C\lambda' + \lambda'\epsilon'$.

Consequently, it is possible to replace each subword $x_{i_j+1}\cdots x_{i_{j+1}}$ of w (which labels the segment between g_{i_j} and $g_{i_{j+1}}$) by a G-equivalent word of length at most κ and the result is a (λ, ϵ)-quasigeodesic, by Lemma 12; $h_{j-1} \ne h_j$ because $\bar{\xi}_j \ne \varepsilon$, and $\lambda'' = 1$, $\epsilon'' = 0$ as $\bar{\xi}$ is a geodesic. □

A geodesic metric space is *δ-hyperbolic* if all of its geodesic triangles are δ-thin; see [4, Definition III.H.1.16]. We recall that a triangle with sides A, B, C is δ-thin if there are three points, one on each of the three sides, called the *meeting points*, with the properties that the two meeting points on the sides adjacent to any of the three vertices of the triangle are equidistant from that vertex, and two points moving synchronously between a vertex and the two meeting points on the sides through it remain a distance at most δ apart. The group $G = \langle X\rangle$ is defined to be δ-hyperbolic if its Cayley graph Γ_X is.

Lemma 14 (Proposition 3.1 in [7]) *If u, v are words in a δ-hyperbolic group $G = \langle X\rangle$ with $u =_G v$, u a geodesic, and v a (λ, ϵ)-quasigeodesic for some λ, ϵ, then u and v boundedly asynchronously K-fellow-travel for some constant K and some asynchronicity bound M, where K and M depend only on λ, ϵ and δ.*

To simplify some of the technicalities, we shall assume that δ, M, K are all integers with $\delta \ge 1$, and that each vertex on u or v is at distance at most K from a vertex on v or u, respectively.

A word w over X with $G = \langle X\rangle$ is said to be *fully reduced* if w and all of its cyclic conjugates are geodesic words in G. Lemma III.Γ.2.9 of [4] says that if the group G is δ-hyperbolic and the fully reduced words u and v represent conjugate elements of G, then either $\max(|u|, |v|) \le 8\delta + 1$ or there exist cyclic conjugates u' and v' of u and v and a word α over X with $\alpha u'\alpha^{-1} =_G v'$ and $|\alpha| \le 2\delta + 1$.

Define w to be (λ, ϵ)-*quasi fully reduced* if w and all of its cyclic conjugates are (λ, ϵ)-quasigeodesic words in G. We can use a similar argument to that employed in the proof of Lemma III.Γ.2.9 of [4] to show (assuming for simplicity that λ and ϵ are integers):

Lemma 15 *If u and v are (λ,ϵ)-quasi fully reduced words representing conjugate elements of a δ-hyperbolic group, then either* $\max(|u|,|v|) \leq \lambda(8\delta + 2K+\epsilon+1)$ *or there exist cyclic conjugates u' and v' of u and v and a word α over X with $\alpha u'\alpha^{-1} =_G v'$ and $|\alpha| \leq 2(\delta+K)$, where K is the constant of Lemma 14.*

PROOF: Suppose that $\max(|u|,|v|) > \lambda(8\delta+2K+\epsilon+1)$, and without loss of generality assume that $|u| > \lambda(8\delta+2K+\epsilon+1)$. Choose a word α that is as short as possible such that there exist cyclic conjugates u' and v' of u and v with $\alpha u'\alpha^{-1} =_G v'$. Let u'' and v'' denote geodesic words with $u'' =_G u'$ and $v'' =_G v'$. Then, since u' and v' are by assumption (λ,ϵ)-quasigeodesics, by Lemma 14 they are at a uniform distance at most K from u'' and v'', and $|u''| > 8\delta+2K+1$.

Consider a geodesic quadrilateral Q in the Cayley graph of G whose sides, read in order from a vertex, trace out edge paths labelled α, u'', α^{-1}, $(v'')^{-1}$. We shall refer to the sides of Q that are labelled $\alpha^{\pm1}$ as the vertical sides and to the sides labelled u'' and $(v'')^{-1}$ as the top and the bottom sides, respectively. We also adjoin paths labelled u' and $(v')^{-1}$ to Q having the same first and last vertices as the top and bottom sides of Q. Since $|u''| \geq 8\delta+2K+2$, we can choose a vertex p on the top side of Q (labelled u'') that is at distance at least $4\delta+K+1$ from both ends of this side. Then p is at a distance at most 2δ from a vertex on one of the other three sides of Q. If it were within 2δ of a point p' on the bottom side (labelled $(v'')^{-1}$), then there would be vertices q,q' on the paths labelled u' and v' with $d(q,q') \leq 2\delta+2K$. But then, if α' is a word labelling a geodesic path from q' to q, α' conjugates a cyclic conjugate of u to a cyclic conjugate of v and hence, by choice of α, we have $|\alpha| \leq |\alpha'| \leq 2(\delta+K)$, as required.

So suppose that p is within a distance 2δ from one of the vertical sides of Q. Let x and y be respectively the top and bottom vertices of the side containing q. Then the choice of α implies that $|\alpha| \leq d(p,y)+K$, and clearly $d(p,y) \leq 2\delta+d(q,y)$ with $|\alpha| = d(q,x)+d(q,y)$, so we have $d(q,x) \leq 2\delta+K$ and therefore $d(p,x) \leq d(p,x)+d(q,x) \leq 4\delta+K$, contrary to the choice of p. □

In the second case of this lemma ($|\alpha| \leq 2(\delta+K)$), by considering the quadrilateral with sides labelled α, u', α^{-1}, $(v')^{-1}$, we see that every vertex p' on u' is either at distance at most $2(\delta+K)$ from a vertex on the side labelled v'^{-1}, or it is at distance at most $2\delta+K$ from a vertex on a side labelled α or α^{-1}. So it is in any case at distance at most $4\delta+3K$ from a vertex on the side labelled v'^{-1}. By choosing the vertex on u' to correspond to the beginning of the original word u, we conclude that there exists a word β with $|\beta| \leq 4\delta+3K$ and $\beta^{-1}u\beta = v''$ for some cyclic conjugate v'' of v.

By a similar argument applied to the quadrilateral with sides labelled β, u, β^{-1} and $(v'')^{-1}$, we find that the vertex on the side labelled $(v'')^{-1}$ that corresponds to the beginning of the word v^{-1} is at distance at most $6\delta + 4K$ from the side labelled u. Hence (see diagram below) we have the following result.

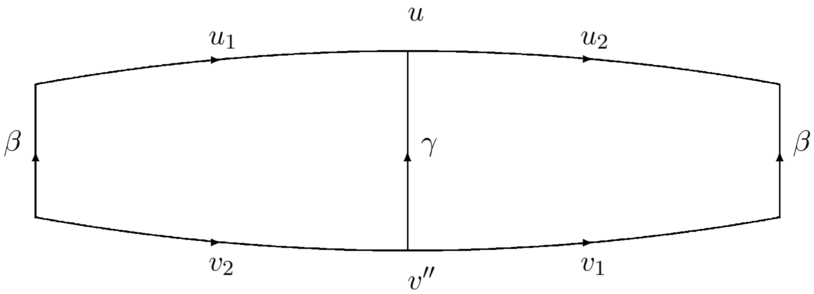

Proposition 16 *Suppose that u and v are (λ, ϵ)-quasi fully reduced words representing conjugate elements of a δ-hyperbolic group G. Then either $\max(|u|, |v|) \leq \lambda(8\delta + 2K + \epsilon + 1)$, or we have $u \equiv u_1u_2$ and $v \equiv v_1v_2$ where there exist $\beta, \gamma \in X^*$ with $|\beta| \leq 4\delta + 3K$, $|\gamma| \leq 6\delta + 4K$, $\beta u_1 \gamma^{-1} =_G v_2$ and $\gamma u_2 \beta^{-1} =_G v_1$.*

Proposition 17 *Let u be a geodesic word in a δ-hyperbolic group G with $\delta \geq 1$. Then we have $u \equiv u_1u_2u_3$, where $u_3u_1 =_G \alpha$ for some word α with $|\alpha| \leq \delta$, and $u_2\alpha$ is $(1, 3\delta + 1)$-quasi fully reduced.*
In other words, the word $u' \equiv u_1u_2\alpha\alpha^{-1}u_3$ obtained by insertion of $\alpha\alpha^{-1}$ into u, can be split as $u'_1u'_2u'_3$ such that $u'_3u'_1 =_G 1$ and $u'_2 \equiv u_2\alpha$ is $(1, 3\delta+1)$-quasi fully reduced.

Proof: Let w be a geodesic word representing u^2, and consider the geodesic triangle with sides A, B and C, labelled by u, u, and w^{-1}, respectively. For simplicity, we shall assume that the meeting points of this triangle coincide with vertices in the Cayley graph. The argument in the other case, in which the meeting points lie midway between two vertices, is similar, and gives rise to the $3\delta + 1$ rather than 3δ in the proposition statement.

Let u_3 be the suffix of u that labels the part of the side A after the meeting point of the triangle on A, and let u_1 be the prefix of u that labels the part of the side B before the meeting point of the triangle on B. See diagram below.

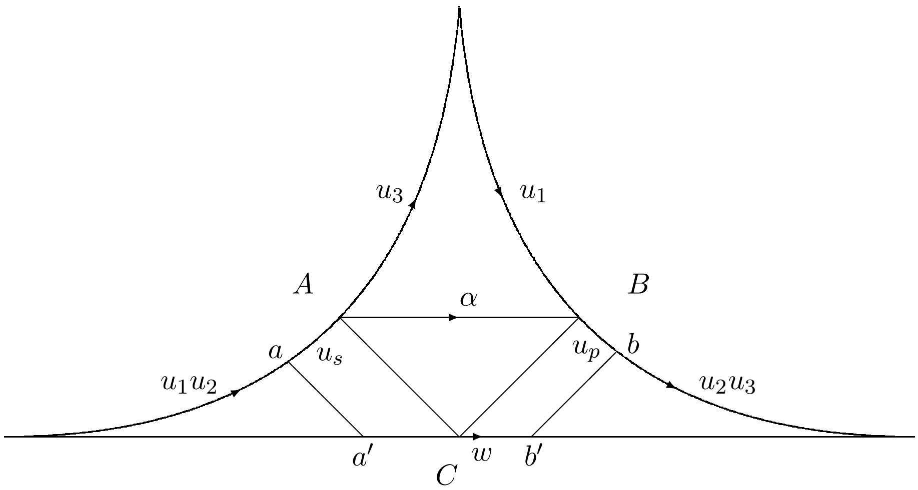

By definition of meeting points, we have $|u_1| = |u_3|$. By δ-hyperbolicity, we have $u_3u_1 =_G \alpha$ with $|\alpha| \leq \delta$. Suppose first that $|u_1| \geq |u|/2$. Then we replace u_1 and u_3 respectively by the prefix and suffix of u of length $|u|/2$ or $(|u| - 1)/2$, and we still have $u_3u_1 =_G \alpha$ with $|\alpha| \leq \delta$, and also $u \equiv u_1u_2u_3$ with $|u_2| \leq 1$. Since $\delta \geq 1$, we have $|u_2\alpha| \leq 3\delta$, so $u_2\alpha$ is $(1, 3\delta)$-quasi fully reduced.

Hence we assume that $|u_1| < |u|/2$, and we can write $u \equiv u_1u_2u_3$ with $|u_2| > 0$, and we need to show that $u_2\alpha$ is $(1, 3\delta)$-quasi fully reduced. To do this, we must show that the length of any subword σ of any cyclic conjugate of $u_2\alpha$ exceeds the geodesic length $|\sigma|_G$ of σ in G by at most 3δ. We shall denote prefixes and suffixes of α and of u_2 by α_p, α_s, u_p, u_s, respectively. If σ has one of the forms α_su_p, $u_s\alpha_p$, or $\alpha_su_2\alpha_p$ then, since u_2 (as a subword of u) is geodesic and $|\alpha_s| + |\alpha_p| \leq |\alpha| \leq \delta$, we see that $|\sigma|_G \geq |\sigma| - 2\delta$. Otherwise, we have $\sigma \equiv u_s\alpha u_p$ with $|u_s|, |u_p| > 0$. Referring again to the diagram above, by δ-hyperbolicity, the first point a on the path labelled u_s on A is at distance at most δ from the point a' on C at distance $|u_s|$ from the meeting point on C and coming before the meeting point on C. Similarly, the last point b on the path labelled u_p on B is at distance at most δ from the point b' on C at distance $|u_p|$ from the meeting point on C and coming after that meeting point. So the geodesic distance between a' and b' is $|u_s| + |u_p|$ and hence $|u_s| + |u_p| \leq 2\delta + |\sigma|_G$, and so

$$|\sigma| = |u_s| + |u_p| + |\alpha| \leq |u_s| + |u_p| + \delta \leq |\sigma|_G + 3\delta,$$

as required. $\square$

Following these preparations, we are now ready to embark upon the proof that virtually free groups have asynchronously context-free inverse conjugacy problem. We start by describing some aspects of the proof. The push-down automaton will be non-deterministic, it will never return a negative answer, and will only return a positive answer when it has verified that the two input words are conjugate in G; so it cannot return an incorrect answer.

It will perform various transformations on the input words as it reads them. For example it may insert words of bounded length at various points, or it may replace an input symbol that it has read by a word of bounded length. It can do these things by using its states as memory. Similarly it can read ahead a bounded number of symbols and memorise the result. The purpose of doing this is to replace the input words by other words that satisfy various desirable properties. The automaton will not in general attempt to verify that these replacement words really do have these properties. But provided that we know theoretically that such replacements are possible, together with the fact that the automaton is non-deterministic and will never return an incorrect answer, we can effectively assume that these properties really are satisfied, and we shall do that during the proof.

Here is a summary of the proof. We suppose that words u, v are given as input.

1. Modify u, v to get (λ, ϵ)-quasigeodesics, for some fixed λ, ϵ.
2. Modify u, v to get geodesics.
3. Modify u, v as in Proposition 17, inserting into each a short word of the form $\alpha\alpha^{-1}$, so that $u \equiv u_1u_2u_3$, $v \equiv v_1v_2v_1$ where $u_1u_3 =_G v_1v_3 =_G 1$ and u_2, v_2 are $(1, 3\delta + 1)$-quasi fully reduced words for some fixed δ.
4. Using Proposition 16 verify that u_2 and v_2^{-1} are conjugate in G.
5. Verify that $u_1u_3 =_G v_1v_3 =_G 1$.

We emphasise that these are the steps in the proof of the theorem. In reality, the modifications of the input words described in the first three steps all take place together as the input words are read by the pushdown automaton. The reader will need to convince themselves that this is possible after studying the details of how these steps are carried out.

The constants λ, ϵ, δ depend only on G and its generating set X, so we may assume that they are known to the pushdown automaton. Only the first step in this proof makes use of the fact that the group is virtually free, rather than just being δ-hyperbolic for some δ, so we shall restate the remaining steps as a separate result that says essentially that the inverse conjugacy problem is solvable for hyperbolic groups on the assumption that the input words are (λ, ϵ)-quasigeodesics, for some fixed λ and ϵ.

Theorem 18 *The inverse conjugacy problem for a virtually free group is asynchronously context-free.*

PROOF: Let $G = \langle X \rangle$ be a virtually free group. We first demonstrate that Theorem 13 allows us to assume that, for some fixed $\lambda \geq 1$ and $\epsilon \geq 0$, our two input words are (λ, ϵ)-geodesics. Let w be one of the input words u, v^{-1}. By Theorem 13, we can replace certain subwords $w_1, \ldots w_k$ of w by G-equivalent words $\alpha_1, \alpha_2, \ldots, \alpha_k$ of length at most K, such that the resulting word is a (λ, ϵ)-quasigeodesic. The pushdown automaton $\mathcal{M}$ guesses non-deterministically which subwords should be replaced by which short words α_i. When it reaches the beginning of a subword w_i that is to be replaced by α_i, then it puts a marker on the stack followed by α_i^{-1}. It then reads w_i and behaves like a pushdown automaton that accepts the word problem for G by empty stack. If, after reading w_i, the marker is at the top of the stack, then it has verified that $w_i =_G \alpha_i$; otherwise it aborts. If the verification is successful, then $\mathcal{M}$ behaves as though it had read the input α_i which, since $|\alpha_i|$ is bounded, is possible by using states of $\mathcal{M}$ as memory.

This completes the first step of the summary of the proof presented above, and the result follows from the next theorem, which handles the remaining steps in the more general context of hyperbolic groups for which the input words are assumed to be quasigeodesics. It is a standard (and elementary) result that every virtually free group is δ-hyperbolic for some δ. $\square$

Theorem 19 *Let $G = \langle X \rangle$ be a δ-hyperbolic group and let $\lambda \geq 1$ and $\epsilon \geq 0$ be fixed. Then there is an asynchronous 2-variable pushdown automaton $\mathcal{M}$ over X that satisfies*

$$\begin{aligned} \{\,(u,v) \mid u, v\ (\lambda,\epsilon)\text{-quasigeodesics},\ u, v^{-1} \text{ conjugate in } G\,\} \quad &\subseteq \quad L(\mathcal{M}) \\ &\subseteq \quad \{\,(u,v) \mid u, v^{-1} \text{ conjugate in } G\,\}. \end{aligned}$$

PROOF: The pushdown automaton $\mathcal{M}$ to be constructed attempts to replace the input words u and v by G-equivalent geodesic words. Let K and M be the constants from Lemma 14, such that a (λ, ϵ)-geodesic u and a geodesic w with $u =_G w$ asynchronously K-fellow-travel with asynchronicity bound M. Now, $\mathcal{M}$ will maintain (by means of its states) a group element g of geodesic length at most K with $g =_G (u')^{-1}w'$, where u' is the prefix of u read so far, and w' is the prefix of the word w with which it is being replaced. On reading the next input letter x of u, $\mathcal{M}$ replaces x by a random word α of length at most M, and replaces g by $x^{-1}g\alpha$ if that group element has length at most K; otherwise the computation is aborted. If $g =_G 1$ after reading all of u, then we know that $u =_G w$; otherwise the computation is aborted. So, by means of this replacement, we can assume that the input word u is geodesic. The other input word v is similarly replaced by an equivalent word, which we shall assume to be geodesic.

From now on we shall use u and v to denote the equivalent words with which the actual input words have been replaced, and assume that they are geodesics. So this completes the second step of the proof summary above.

By Proposition 17, we have $u \equiv u_1 u_2 u_3$ where $u_3 u_1 =_G \alpha$ (and hence $u_1 \alpha^{-1} u_3 =_G 1$) with $|\alpha| \leq \delta$, and $u_2\alpha$ is $(1, 3\delta+1)$-quasi fully reduced. Our automaton $\mathcal{M}$ splits the input word non-deterministically into three parts u_1, u_2, u_3. It also replaces u_2 by $u_2\alpha$ and u_3 by $\alpha^{-1}u_3$ for some arbitrarily chosen word α with $|\alpha| \leq \delta$. So Proposition 17 says that it is possible to do this such that the replaced word u_2 is $(1, 3\delta + 1)$-quasi fully reduced, and $u_1 u_3 =_G 1$.

The input word v will be similarly split up and modified into $v_1 v_2 v_3$, where we know that it is possible to do this such that v_2 is $(1, 3\delta + 1)$-quasi fully reduced, and $v_1 v_3 =_G 1$. This completes the third step of the proof summary.

As in the proof of Theorem 11, $\mathcal{M}$ starts by putting u_1 on the stack followed by a marker μ_1, and then it puts v_1 on the stack followed by a marker μ_2. Now if u and v^{-1} are conjugate in G then so are u_2 and v_2^{-1}, and if our constructions of $u_1, u_2, u_3, v_1, v_2, v_3$ were successful, then u_2, v_2 and v_2^{-1} are $(1, 3\delta + 1)$-quasi fully reduced words. So in that case we can apply Proposition 16 to u_2 and v_2^{-1}. Note that the constant K in the statement of Proposition 16 depends only on δ.

Since $\lambda(8\delta+2K+\epsilon+1) = 11\delta+2K+2$ is a constant, $\mathcal{M}$ can read ahead and memorise up to $11\delta+2K+2$ symbols of u_2 and v_2 and thereby test whether $\max(|u_2|, |v_2|) \leq 11\delta + 2K + 2$. As there are only finitely many pairs of words satisfying this condition, we can assume that $\mathcal{M}$ knows already which of the pairs (u_2, v_2^{-1}) represent conjugate elements of G. If it finds that $\max(|u_2|, |v_2|) \leq 11\delta + 2K + 2$ with u_2 and v_2^{-1} conjugate, then it removes the marker μ_2 from the top of the stack and reads v_3 to verify (by empty stack) that $v_1 v_3 =_G 1$. If so, then the marker μ_1 is now at the top of the stack, and $\mathcal{M}$ removes it and reads u_3 to verify that $u_1 u_3 =_G 1$. If so, then it has proved that the two input words are conjugate and can return true. If any of these verifications fail, then the computation is aborted.

So assume that $\max(|u_2|, |v_2|) > 11\delta + 2K + 2$, and we are in the second case of Proposition 16. The automaton $\mathcal{M}$ now guesses words β, γ with $|\beta| \leq 4\delta + 3K$ and $|\gamma| \leq 6\delta + 4K$. As $\mathcal{M}$ reads the words u_2 and v_2 it splits them arbitrarily into two subwords $u_2 \equiv u_{21}u_{22}$, $v_2 \equiv v_{21}v_{22}$. It first reads u_{21}, but puts $\beta u_{21}\gamma^{-1}$ on the stack. It then reads v_{21} and tests (by empty stack) whether $\beta u_{21}\gamma^{-1}v_{21} =_G 1$. If so, then it reads u_{22}, but puts $\gamma u_{22}\beta^{-1}$ on the stack. Then it reads v_{22} and tests whether $\gamma u_{22}\beta^{-1}v_{22} =_G 1$. If so, then $\beta u_2 \beta^{-1} =_G v_{21}^{-1}v_{22}^{-1}$, which is conjugate in G to $v_2^{-1} \equiv v_{22}^{-1}v_{21}^{-1}$, so $\mathcal{M}$ has verified that u_2 and v_2^{-1} are conjugate in G. Furthermore, the marker μ_2 is now at the top of the stack, and $\mathcal{M}$ proceeds to verify that $v_1 v_3 =_G 1$ and that $u_1 u_3 =_G 1$, in the same way as above. $\square$

5 Indexed Languages are Richer

Here we shall only give the crucial idea for the proof of Theorem C, leaving the details to the reader.

Recall that an indexed language is accepted by a nested stack automaton which is a generalisation of a pushdown automaton, introduced by Aho in [1, 2]. More precisely, such a machine may move through the stack in read only mode and at any point it can open a nested stack. The top part of the previous (outer) stack is not accessible until this new stack is empty again. This enables checking whether two input words u and v define the same element of a free group as follows. The free reduction of u is read onto the stack and then the read-write head is moved down to the beginning of u (which is marked). The read-write head then moves back up through the stack letter by letter as v is read and freely reduced, comparing the free reduction of v with the freely reduced word that is on the stack. The reduction of v, as it is being read, is achieved on nested stacks that are inserted next to the symbol currently being read on the main stack. Since such a nested stack will necessarily be empty after verifying that a subword reduces freely to the empty word, the computation will continue as if this subword was missing. Using additional states this procedure also works in a virtually free group. Hence, for a virtually free group, one can adapt the strategy of the previous Section, after observing that all necessary replacements can be justified before they are carried out on a nested stack which will be empty after an affirmative check.

References

[1] A.V. Aho, ‘Indexed grammars-an extension of context-free grammars’, *J. Assoc. Comp. Mach.* 15 (1968) 647–671.

[2] A.V. Aho, ‘Nested stack automata’, *J. Assoc. Comp. Mach.* 16 (1969) 383–406.

[3] V.A. Anisimov, ‘The group languages’, *Kibernetika* 4 (1971) 18–24.

[4] M.R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Springer, 1999.

[5] J.E. Hopcroft and J.D. Ullman, *Introduction to automata theory, languages, and computation*, Addison-Wesley Publishing Co., Reading, Mass., 1979.

[6] D.F. Holt and C.E. Röver, ‘Groups with indexed co-word Problem’, *Internat. J. Algebra Comput.* 26 (2006), 985–1014.

[7] D.F. Holt and S. Rees, 'Regularity of quasigeodesics in a hyperbolic group', *Internat. J. Algebra Comput.* 12 (2002), 747–754.

[8] W. Magnus, A. Karrass and D. Solitar, *Combinatorial Group Theory*, Dover Publications Inc., New York, 1976.

[9] D.E. Muller and P.E. Schupp, 'Groups, the theory of ends, and context-free languages', *J. Comp. System Sci.* 26 (1983), 295–310.

Addresses

D. F. Holt
Mathematics Institute
University of Warwick
Coventry CV4 7AL
email: D.F.Holt@warwick.ac.uk

S. Rees
School of Mathematics and Statistics
University of Newcastle
Newcastle NE1 7RU
email: sarah.rees@newcastle.ac.uk

C. E. Röver
School of Mathematics, Statistics and Applied Mathematics
National University of Ireland, Galway
University Road
Galway
Ireland
email: class.roever@nuigalway.ie