Canad. J. Math. Vol. 71 (5), 2019 pp. 997–1018
http://dx.doi.org/10.4153/CJM-2018-036-7

Geometry of Infinitely Presented Small Cancellation Groups and Quasi-homomorphisms

Goulnara Arzhantseva and Cornelia Druţu

Abstract. We study the geometry of infinitely presented groups satisfying the small cancellation condition $C'(1/8)$, and introduce a standard decomposition (called the *criss-cross decomposition*) for the elements of such groups. Our method yields a direct construction of a linearly independent set of power continuum in the kernel of the comparison map between the bounded and the usual group cohomology in degree 2, without the use of free subgroups and extensions.

1 Introduction

An important direction of research in geometric group theory is the construction of infinite finitely generated groups with unusual properties, the so-called "infinite monsters". In this setting, one of the main techniques is small cancellation theory. It produces direct limits of Gromov hyperbolic groups that are therefore, in the class of infinitely presented groups, easier to deal with than other groups, and benefit in many ways from the techniques available in Gromov hyperbolic geometry. These, and other, infinite monsters are usually constructed with the goal of producing counter-examples to various conjectures in algebra, geometry, and analysis. Therefore, it is rather challenging to obtain positive results about them. Positive results have been proved for algebraic and geometric properties, and much less for analytic properties, until very recently. This paper and our subsequent work have been the first steps in this direction. The first arXiv version of this paper has already given an impetus to later works by various authors.

One of the first and easiest constructions of small cancellation groups is that of groups satisfying the classical small cancellation condition $C'(\lambda)$, where $\lambda \in (0, 1/6]$. It seems likely that most of the properties of Gromov hyperbolic groups are also satisfied by such groups, possibly for λ small enough, but this has not yet been proved for many of the analytic properties of hyperbolic groups.

Received by the editors October 18, 2017; revised July 3, 2018.
Published electronically December 26, 2018.
The first author was supported in part by the ERC grant ANALYTIC no. 259527, and by the Swiss NSF, under Sinergia grant CRSI22-130435.
The second author was supported in part by the EPSRC grant no. EP/K032208/1 entitled "Geometric and analytic aspects of infinite groups", by the project ANR Blanc ANR-10-BLAN 0116, acronym GGAA, and by the LABEX CEMPI.
AMS subject classification: **20F06**, 20F67.
Keywords: small cancellation theory, Greendlinger lemma, quasi-homomorphism, bounded cohomology.

In this paper, we provide a key tool for a systematic approach to the above-mentioned problem, the criss-cross decomposition. As an immediate application, we provide a new way of constructing an abundance of quasi-homomorphisms in these groups; see Section 1.2 and Theorem 4.6. Moreover, our quasi-homomorphisms are $\mathbb{Z}$-valued and have defect at most 2. This may be significant, because such a bound on the defect has consequences for the stable commutator length [Bav91].

Our methods very likely generalize to larger classes of groups.

1.1 Main Technical Tool: Criss-cross Decompositions

For several analytic and geometric properties of groups, including the ones discussed in this paper, it is crucial to understand if the group elements possess "standard decompositions" into products of certain "elementary" parts.

The main technical result of our paper is the construction of such a decomposition for elements of finitely generated groups defined by infinite presentations with the small cancellation condition $C'(\lambda)$, for $\lambda \leqslant \frac{1}{8}$ (the so-called $C'(1/8)$-groups). More precisely, given a pair of vertices in a Cayley graph of such a group, we obtain a detailed description of a set containing all the geodesics between the two vertices; see Theorem 3.15. The existence of such sets allows us to introduce a uniquely defined decomposition, called the *criss-cross decomposition*, of the elements of the given $C'(1/8)$-group; see Section 3.

It is worth noticing that our approach differs and cannot be deduced from the Rips–Sela canonical representatives [Sel92, RS95] in finitely presented small cancellation groups. Indeed, the Rips–Sela construction gives an equivariant choice of quasi-geodesic paths between pairs of vertices (with a view to reduce solving equations in a finitely presented small cancellation group, or more generally in a hyperbolic group, to solving equations in a free group). Their arguments are based on the existence of central points for geodesic triangles, granted by finite presentation only. The similarity in method between the Rips–Sela approach and ours does not go beyond the common use of the geometry of geodesic bigons [RS95, Theorem 5.1].

1.2 Application: the Bounded Versus the Usual Cohomology

An immediate application of the criss-cross decomposition is that infinitely presented $C'(1/12)$-groups are rich in quasi-homomorphisms (see Section 4 for definitions). This has a strong impact on the bounded cohomology of such groups [Gro82]. We thus deduce the following theorem.

Theorem 1.1 *Let G be a finitely generated group defined by an infinite presentation satisfying the small cancellation condition $C'(1/12)$. Then the kernel of the comparison map between the second bounded and the usual group cohomology*

$$H_b^2(G) \longrightarrow H^2(G),$$

is an infinite dimensional real vector space, with a basis of power continuum.

The above kernel can be identified with the real vector space $\widetilde{QH}(G)$ of quasi-homomorphisms modulo near-homomorphisms (where by a near-homomorphism we mean a function $\mathfrak{h}: G \to \mathbb{R}$ that differs from a homomorphism by a bounded function). The fact that the kernel is large has implications for the stable commutator length [Bav91], for the (lack of) bounded generation, for results of non-embeddability of higher rank lattices, etc.

The computation of $\widetilde{QH}(G)$ is therefore important, and it has been done for various classes of groups.

If G is amenable, then $\widetilde{QH}(G) = 0$ [Gro82]. Also, if G is an irreducible lattice in a semisimple Lie group of rank at least two and with finite center, then $\widetilde{QH}(G) = 0$ [BM99, BM02].

Groups that have a certain type of action on a hyperbolic space (in particular, subgroups of relatively hyperbolic groups, mapping class groups, etc.) have $\widetilde{QH}(G)$ infinite dimensional, with a basis of power continuum. This was proved by Brooks for non-abelian free groups [Bro81] and by Brooks and Series for non-amenable surface groups. In [Gro87] Gromov stated that all non-elementary hyperbolic groups have non-trivial second bounded cohomology. Epstein and Fujiwara proved that, in fact, for all non-elementary hyperbolic groups, $\widetilde{QH}(G)$ has a basis of power continuum [EF97]. Later, this result was extended to other types of groups acting on hyperbolic spaces and to their non-elementary subgroups [Fuj00, Fuj98]; in particular, to subgroups of mapping class groups of surfaces [BF02]. See also the survey of Fujiwara [Fuj09] and references therein. The same result was further extended to groups with free hyperbolically embedded subgroups by Hull and Osin [HO13].

On the whole, one can say that in all the cases where it was proved, up to now, that $\widetilde{QH}(G)$ has a basis of power continuum, the argument relied on the fact that the group considered contained a non-elementary hyperbolic subgroup, and an extension of the quasi-homomorphisms of that subgroup could be performed, if the subgroup was "well embedded" (*e.g.,* hyperbolically embedded, in the sense of [DGO17]). In particular, two years after the first arXiv version of this paper was posted, it was been proven by Gruber and Sisto in [GS14] that graphical small cancellation groups are acylindrically hyperbolic, and therefore, by work of Hull and Osin [HO13], contain hyperbolically embedded free non-abelian subgroups, and consequently the space of quasi-homomorphisms modulo near-homomorphisms has a basis of power continuum.

Our approach differs from all the previous ones in that we do not require the existence of "well embedded" non-elementary hyperbolic subgroups, and a potential extension of our methods may apply to groups satisfying other small cancellation conditions (*e.g.,* the Ol'shanskii graded small cancellation), in particular, to free Burnside groups of sufficiently large odd exponent or to various Tarski monsters.

Moreover, besides being by far the first proof of Theorem 1.1, our construction has the merit of by-passing the technicalities of [GS14] and [HO13] and of providing a direct explicit construction of a family of power continuum of linearly independent elements in $\widetilde{QH}(G)$, which is not an extension of a similar family for a non-elementary hyperbolic subgroup $H \leqslant G$, but is contained in the ℓ^1 infinite sum of all such extensions for all such subgroups H.

The following result is another immediate consequence of our theorem above.

Corollary 1.2 *Let G be a finitely generated group given by an infinite presentation satisfying the small cancellation condition $C'(1/12)$. Then G is not boundedly generated.*[1]

1.3 Plan of the Paper

The paper is organized as follows. Section 2 gives preliminary information on small cancellation groups. In Section 3, we describe the criss-cross decomposition of elements in infinitely presented small cancellation groups. We believe this description is of independent interest and can be applied to get further results on such groups. In Section 4, we focus on quasi-homomorphisms of $C'(1/12)$-small cancellation groups and prove Theorem 1.1.

2 Preliminaries on Infinite Small Cancellation Presentations

A set of words R in the alphabet A is said to be *symmetrized* if it contains r^{-1} and all the cyclic permutations of r and r^{-1}, whenever $r \in R$. Without loss of generality, we always assume that the set of group relators is symmetrized and that all relators $r \in R$ are reduced words in the alphabet A.

We focus on finitely generated groups with *infinite presentations*,

$$G = \langle A \mid r_1, \dots, r_k, \dots \rangle, \tag{2.1}$$

defined by a symmetrized family R of relators consisting of an infinite sequence of relators $r_1, \dots, r_k, \dots$.

We denote by R_k the set $\{r_1, \dots, r_k\}$ and by G_k the finitely presented group

$$G_k = \langle A \mid R_k \rangle = \langle A \mid r_1, \dots, r_k \rangle.$$

For two words u, v we write $u \sqsubset v$ when u is a subword of v. Let η be a constant in $(0, \frac{1}{2}]$. If in the preceding we have, moreover, that

$$\eta |v| \leqslant |u| \leqslant \frac{1}{2}|v|,$$

then we use the notation $u \sqsubset_\eta v$. We write $u \sqsubset R$ if there exists $v \in R$ such that $u \sqsubset v$, and similarly, with $\sqsubset$ replaced by $\sqsubset_\eta$.

Notation 2.1 *We denote by $S(R)$ the set of words u such that $u \sqsubset R$ and by $S^\eta(R)$ the set of words u such that $u \sqsubset_\eta R$.*

Definition 2.2 ($C'(\lambda)$-condition) Let $\lambda \in (0,1)$. A symmetrized set R of words in the alphabet A is said to satisfy the $C'(\lambda)$-*condition* if the following hold:

(i) If u is a subword in a word $r \in R$ so that $|u| \geqslant \lambda |r|$, then u occurs only once in r;
(ii) If u is a subword in two distinct words $r_1, r_2 \in R$, then $|u| < \lambda \min\{|r_1|, |r_2|\}$.

We say that a group presentation $\langle A \mid R \rangle$ *satisfies the $C'(\lambda)$-condition* if R satisfies that condition.

[1] A group is *boundedly generated* if it can be expressed (as a set) as a finite product of cyclic subgroups.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

Our technical arguments use the language of van Kampen diagrams over a group presentation $\langle A \mid R\rangle$; for more details and terminology, see [LS77], and observe that the classical results below still hold for infinite group presentations.

The boundary of any van Kampen diagram (cell) Δ is denoted by $\partial\Delta$.

Lemma 2.3 (Greendlinger [LS77, Ch.V, Thm. 4.4]) *Every reduced van Kampen diagram Δ over the presentation (2.1) with small cancellation condition $C'(\lambda)$ for $\lambda \leqslant \frac{1}{6}$ contains a cell Π with $\partial\Pi$ labeled by a relator $r \in R$ such that $\partial\Delta \cap \partial\Pi$ has a connected component of length $> (1-3\lambda)|r|$.*

Definition 2.4 (n-gon) We call *n-gon* in a geodesic metric space a loop obtained by successive concatenation of n geodesics.

We say that the n-gon is *simple* if the loop thus obtained is simple, that is, if it does not have self-intersections.

Theorem 2.5 (*cf.* [GdlH91]) *Let Δ be a reduced van Kampen diagram over a group presentation $G = \langle A \mid R\rangle$ satisfying the $C'(\lambda)$-condition, with $\lambda \leqslant \frac{1}{8}$.*

(i) *Assume that $\partial\Delta$ is a simple bigon in the Cayley graph of G. Then it has the form of the bigon B in Figure 1.*

(ii) *Assume that $\partial\Delta$ is a simple triangle in the Cayley graph of G. Then it has one of the forms $T_1, \dots, T_4$ in Figure 1 and Figure 2.*

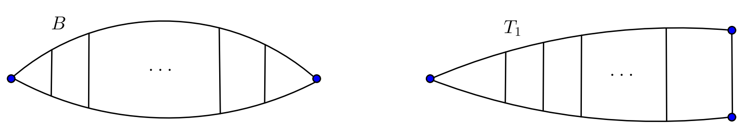

Figure 1: Simple bigon B and simple triangle T_1.

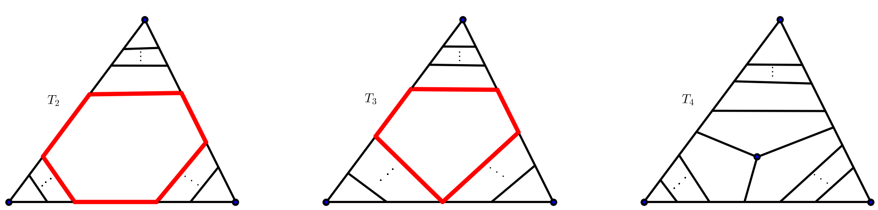

Figure 2: Simple triangles T_2, T_3, and T_4.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

3 Standard Decomposition of Elements in Small Cancellation Groups

This section is devoted to a thorough analysis of geodesics in Cayley graphs of infinitely presented small cancellation groups, and to the description of a set which, from many points of view, plays the part of the convex hull of a two points-set in irreducible buildings. We show here the main technical result of the paper, Theorem 3.15, and its algebraic counterpart Theorem 3.27.

Convention 3.1 Throughout this section, G denotes a finitely generated group with a (possibly infinite) presentation $\langle A \mid R \rangle$ satisfying the $C'(\lambda)$-condition with $\lambda \leqslant \frac{1}{8}$.

We only consider the Cayley graph of G with respect to the fixed (arbitrary) finite generating set A, and we omit mentioning A from now on. By "vertex", we shall always mean a vertex in that Cayley graph.

We call *contour* a loop in the Cayley graph of G labeled by a relator $r \in R$. By abuse of notation, given a contour t, we denote its length by $|t|$. Observe that a contour is always a simple loop (a non-trivial self-intersection leads to a contradiction with the small cancellation assumption by the Greendlinger lemma).

By an *arc* we mean a topological arc, that is the image of a topological embedding of an interval into a topological (in particular metric) space.

For every path $\mathfrak{p}$ in a metric space, we denote the initial point of $\mathfrak{p}$ by $\mathfrak{p}_-$ and the terminal point of $\mathfrak{p}$ by $\mathfrak{p}_+$. Given two points x, y on a geodesic $\mathfrak{g}$, we denote by $[x, y]$ the sub-geodesic of $\mathfrak{g}$ with endpoints x, y.

Lemma 3.2 *Let t be a contour labeled by a relator r and let a, b be two points on t.*

(i) *If one of the two arcs with endpoints a, b has length $< \frac{|t|}{2}$, then that arc is the unique geodesic with endpoints a, b in the Cayley graph.*
(ii) *If both arcs with endpoints a, b have length $\frac{|t|}{2}$, then these arcs are the only two geodesics with endpoints a, b in the Cayley graph.*
(iii) *The intersection of a geodesic with a contour is always composed of only one arc.*

Proof (i) Assume that there exists a geodesic joining a, b distinct from that arc. Then they compose at least one non-trivial simple bigon. Consider the minimal van Kampen diagram Δ with the same boundary label as this bigon. Let u be the label of the sub-arc of t and v the label of the sub-arc of the geodesic. According to Lemma 2.3, there exists a cell Π labeled by a relator intersecting the boundary $\partial\Delta$ in an arc of length $> 1 - 3\lambda$ of the length of $\partial\Pi$.

Assume first that $\partial\Pi$ does not coincide with t. By the small cancellation condition, the arc can have at most λ of the length of $\partial\Pi$ in common with the arc labeled by a subword of r, hence it has $> 1 - 4\lambda$ of the length of $\partial\Pi$ in common with the arc with the same label as the geodesic. As $\lambda \leqslant \frac{1}{8}$, this contradicts the fact that this is the label of a geodesic.

Now if $\partial\Pi$ coincides with t, then $\partial\Pi$ has at least $\frac{1}{2} - 3\lambda$ of its length in common with the arc labeled by v. In particular, it follows that $|u| \geqslant |v| > (\frac{1}{2} - 3\lambda)|r|$. Then $\partial\Delta \vartriangle \partial\Pi$ composes a new simple bigon with both sides of length at most $3\lambda|r|$. We

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

apply the argument above to this new bigon; the boundary of the cell provided by Lemma 2.3 cannot coincide with t this time, and we obtain a contradiction.

(ii) The argument to show that there exists no geodesic joining a, b and that is not entirely contained in t is as above.

(iii) It suffices to prove that this intersection is path connected. Indeed, let $\mathfrak{g}$ be a geodesic and let a, b be two points on $\mathfrak{g} \cap t$. The above arguments show that the part of $\mathfrak{g}$ between a and b must be contained in t. ∎

Definition 3.3 (Relator-tied geodesics and components) Let $\mathfrak{g}$ be a geodesic in the Cayley graph of G and let η be a number in $(0,1)$.

(i) $\mathfrak{g}$ is called *η-relator-tied* if it is covered by sub-geodesics labeled by words in $S^\eta(R)$.

(ii) an *η-relator-tied component* of $\mathfrak{g}$ is a maximal sub-geodesic of $\mathfrak{g}$ that is η-relator-tied.

Lemma 3.4 (i) *The η-relator-tied components of a geodesic $\mathfrak{g}$ are disjoint.*

(ii) *Assume that $\eta \leqslant \frac{1}{2} - 2\lambda$. If two points a, b are the endpoints of a geodesic $\mathfrak{g}$ with no η-relator-tied component, then $\mathfrak{g}$ is the unique geodesic with endpoints a, b.*

Proof Assertion (i) follows by definition, since two distinct η-relator-tied sub-geodesics that intersect compose a longer η-relator-tied sub-geodesic.

(ii) Any other geodesic $\mathfrak{g}'$ with endpoints a, b and distinct from $\mathfrak{g}$ would compose with $\mathfrak{g}$ simple geodesic bigons, therefore by Theorem 2.5(i), $\mathfrak{g}$ would contain a $(\frac{1}{2} - 2\lambda)$-relator-tied component. ∎

Definition 3.5 (η-compulsory geodesic) Given $0 < \eta \leqslant \frac{1}{2} - 2\lambda$, a geodesic as in Lemma 3.4(ii) is called an *η-compulsory geodesic*. A pair of endpoints a, b of an η-compulsory geodesic is called an *η-compulsory pair*.

We now proceed to analyze the η-relator-tied components of geodesics.

Lemma 3.6 *Let $\eta \geqslant 2\lambda$. Let $\mathfrak{g}$ be a η-relator-tied geodesic in the Cayley graph of G. Then there exists a unique sequence of successive vertices*

$$x_0 = a, x_1, y_0, x_2, y_1, \dots, x_{k+1}, y_k, y_{k+1} = b$$

such that the sub-geodesics with endpoints x_i, y_i with $i \in \{0, 1, \dots, k+1\}$ are labeled by words in $S^\eta(R)$, and are maximal with this property with respect to inclusion (see Figure 3).

Proof By hypothesis, $\mathfrak{g} \subseteq \bigcup_{i \in S_0} \mathfrak{g}_i$, where $\mathfrak{g}_i$ denotes a sub-geodesic of $\mathfrak{g}$ labeled by a word $u_i \in S^\eta(R)$ and the index set S_0 is finite (by compactness).

Without loss of generality, we assume that all the sub-geodesics $\mathfrak{g}_i$ in the covering above are maximal with respect to inclusion.

Indeed, we begin with the sub-geodesics containing the vertex $\mathfrak{g}_-$. Consider two such sub-geodesics. If one is contained in the other, by the $C'(\lambda)$-condition and the fact that $\eta \geqslant 2\lambda$, it follows that both are subwords of the same relator $r_1 \in R$. Therefore, we take the longer of the two subwords, and we select it as the first term $\mathfrak{g}_1$ of the new

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

covering. The endpoint $(\mathfrak{g}_1)_+$ must be contained in another sub-geodesic $\mathfrak{g}_u$. The sub-geodesic $\mathfrak{g}_1 \sqcup \mathfrak{g}_u$ cannot be labeled by a word in $S^\eta(R)$, because this would contradict the maximality of $\mathfrak{g}_1$. We consider the maximal sub-geodesic $\mathfrak{g}_2$ labeled by a word in $S^\eta(R)$ and containing $\mathfrak{g}_u$. Continuing this argument, we obtain a cover $\bigcup_{i\in S_1} \mathfrak{g}_i$ of $\mathfrak{g}$ for some $S_1 \subseteq S_0$ such that $\mathfrak{g}_i$ are maximal sub-geodesics labeled by words in $S^\eta(R)$.

For an arbitrary small $\varepsilon > 0$ we have that $\mathfrak{g} \subseteq \bigcup_{i\in S_1} \mathfrak{g}_i^\varepsilon$, where $\mathfrak{g}_i^\varepsilon$ denotes the ε-neighborhood of $\mathfrak{g}_i$ in $\mathfrak{g}$. Since $\mathfrak{g}$ has topological dimension one, there exists $S_2 \subseteq S_1$ such that $\mathfrak{g} \subseteq \bigcup_{i\in S_2} \mathfrak{g}_i^\varepsilon$ and every point in $\mathfrak{g}$ is contained in at most two sets $\mathfrak{g}_i^\varepsilon$ with $i \in S_2$.

If an edge e in $\mathfrak{g}$ is not contained in $\bigcup_{i\in S_2} \mathfrak{g}_i$, then, for $\varepsilon < \frac{1}{2}$, this contradicts the fact that $\{\mathfrak{g}_i^\varepsilon \mid i \in S_2\}$ cover e. If a vertex in $\mathfrak{g}$ is not contained in $\bigcup_{i\in S_2} \mathfrak{g}_i$, then the edges adjacent to it are not contained in $\bigcup_{i\in S_2} \mathfrak{g}_i$, and we use the above.

We thus obtain that $\mathfrak{g} \subseteq \bigcup_{i\in S_2} \mathfrak{g}_i$, and every point in $\mathfrak{g}$ is contained in at most two sets $\mathfrak{g}_i$ with $i \in S_2$.

Assume that there exist two sequences

$$x_0 = a, x_1, y_0, x_2, y_1, \dots, x_{k+1}, y_k, y_{k+1} = b,$$
$$x_0' = a, x_1', y_0', x_2', y_1', \dots, x_{m+1}', y_m', y_{m+1}' = b,$$

and let $k \leqslant m$. We prove by induction on $0 \leqslant i \leqslant k+1$ that $[x_i, y_i] = [x_i', y_i']$.

First, consider the case $i = 0$. Then either $[x_0, y_0] \subseteq [x_0', y_0']$ or $[x_0', y_0'] \subseteq [x_0, y_0]$. The assumption $\eta \geqslant 2\lambda$ implies that both the label of $[x_0, y_0]$ and that of $[x_0', y_0']$ are subwords of the same relator r. The maximality condition implies that $[x_0, y_0] = [x_0', y_0']$.

We now assume that for some $j \geqslant 0$, we have $[x_i, y_i] = [x_i', y_i']$ for $0 \leqslant i \leqslant j$. We have that either $[y_j, y_{j+1}] \subseteq [y_j, y_{j+1}']$ or $[y_j, y_{j+1}'] \subseteq [y_j, y_{j+1}]$. By maximality and Lemma 3.2, the contour t_j containing the geodesic $[x_j, y_j]$ is distinct from the contour t_{j+1} containing the geodesic $[x_{j+1}, y_{j+1}]$, respectively the contour t_{j+1}' containing the geodesic $[x_{j+1}', y_{j+1}']$; see Figure 4. It follows that

$$\operatorname{dist}(x_{j+1}, y_j) < \lambda |r_{j+1}| \leqslant \frac{\lambda}{\eta} \operatorname{dist}(x_{j+1}, y_{j+1}),$$

whence

$$\operatorname{dist}(y_j, y_{j+1}) > \Big(1 - \frac{\lambda}{\eta}\Big) \operatorname{dist}(x_{j+1}, y_{j+1}) \geqslant \eta\Big(1 - \frac{\lambda}{\eta}\Big)|r_{j+1}|.$$

Similarly, we obtain that

$$\operatorname{dist}(y_j, y_{j+1}') > \eta\Big(1 - \frac{\lambda}{\eta}\Big)|r_{j+1}'|.$$

The hypothesis $\eta \geqslant 2\lambda$ implies that $\eta(1 - \frac{\lambda}{\eta}) \geqslant \lambda$; therefore, the inclusions

$$[y_j, y_{j+1}] \subseteq [y_j, y_{j+1}'] \text{ or } [y_j, y_{j+1}'] \subseteq [y_j, y_{j+1}]$$

imply that $t_{j+1} = t_{j+1}'$. The maximality of the sub-geodesics $[x_{j+1}, y_{j+1}]$ and $[x_{j+1}', y_{j+1}']$, and Lemma 3.2 allow to conclude that $[x_{j+1}, y_{j+1}] = [x_{j+1}', y_{j+1}']$. ■

Convention 3.7 For the rest of this section, let η be a fixed constant such that $\frac{1}{2} - 2\lambda \geqslant \eta \geqslant 2\lambda$ and $\eta' := \eta - \lambda$.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

Definition 3.8 (η-succession) We say that a sequence of contours $t_0, t_1, \dots, t_{k+1}$ is an *η-succession of contours* if, for every i, one of the endpoints of $t_{i-1} \cap t_i$ is at distance $\geqslant \eta |t_i|$ from at least one of the endpoints of $t_i \cap t_{i+1}$ (distance measured in t_i).

Corollary 3.9 *Let $\mathfrak{g}$ be an η-relator-tied geodesic. Then there exists a unique η-succession of contours $t_0, t_1, \dots, t_{k+1}$ such that for the decomposition described in Lemma 3.6 the sub-geodesic with endpoints x_i, y_i is contained in t_i.*

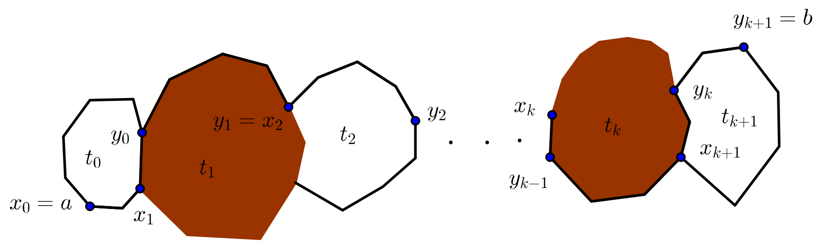

Figure 3: An η-relator-tied geodesic inside a succession of contours.

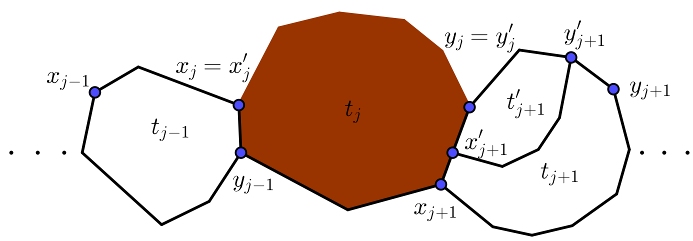

Figure 4: The uniqueness of the decomposition in Lemma 3.6.

Lemma 3.10 *Let a, b be two vertices joined by an η-relator-tied geodesic $\mathfrak{g}$. Then every geodesic $\mathfrak{g}'$ with endpoints a, b is η'-relator-tied, for $\eta' := \eta - \lambda$; moreover, $\mathfrak{g}'$ is contained in the η-succession of contours $t_0, t_1, \dots, t_{k+1}$ determined by $\mathfrak{g}$ according to Corollary 3.9.*

Proof There exist successive points in the intersection $\mathfrak{g} \cap \mathfrak{g}'$,

$$z_0 = a, z_1, \dots, z_{2m-1}, z_{2m}, z_{2m+1} = b$$

such that z_{2i}, z_{2i+1} are the endpoints of a connected component of $\mathfrak{g} \cap \mathfrak{g}'$, while z_{2i+1}, z_{2i+2} are the endpoints of two sub-geodesics of $\mathfrak{g}$ respectively $\mathfrak{g}'$, composing a simple bigon.

Let $x_0 = a, x_1, y_0, x_2, y_1, \dots, x_{k+1}, y_k, y_{k+1} = b$ be the unique sequence of points on $\mathfrak{g}$ provided by Lemma 3.6. For every $0 \leqslant i \leqslant m-1$, consider the endpoints z_{2i+1}, z_{2i+2} of a simple bigon. According to Theorem 2.5(i), the corresponding bigon is as in Figure 1. Consider τ, one of the contours appearing in this bigon; let α, β be the endpoints of the intersection of $\mathfrak{g}$ with τ. By the small cancellation condition and the fact that the two sides of the bigon are geodesics, it follows that the label of the sub-geodesic of $\mathfrak{g}$ limited by α, β is a sub-word of length $> (\frac{1}{2} - 2\lambda)|\tau|$.

Let m be the midpoint of the sub-geodesic of $\mathfrak{g}$ limited by α, β. Then there exist x_j, y_j separated by m. If the contour τ is distinct from the contour t_j, then

$$\frac{1}{2}\Big(\frac{1}{2} - 2\lambda\Big)|\tau| < \frac{1}{2}\operatorname{dist}(\alpha, \beta) < \lambda|\tau|,$$

whence $\lambda > \frac{1}{8}$, a contradiction. It follows that $\tau = t_j$, hence $\alpha = x_j$ and $\beta = y_j$. Thus, the endpoints of intersections of contours of the bigon with $\mathfrak{g}$ compose a subsequence of the sequence

$$x_0 = a, x_1, y_0, x_2, y_1, \dots, x_{k+1}, y_k, y_{k+1} = b,$$

with the property that $x_{i+1} = y_i$.

Let z_{2i+1} be an endpoint of a bigon. According to the above, z_{2i+1} equals some x_j such that t_j is the first contour in the bigon. Now consider x_{j-1}, y_{j-1} and the contour $t_{j-1} \neq t_j$. Then $\operatorname{dist}(x_j, y_{j-1}) < \lambda|t_{j-1}|$, whence $\operatorname{dist}(x_{j-1}, x_j) = \operatorname{dist}(x_{j-1}, y_{j-1}) - \operatorname{dist}(x_j, y_{j-1}) > \eta|t_{j-1}| - \lambda|t_{j-1}| = \eta'|t_{j-1}|$.

We have thus found that the sub-geodesic with endpoints z_{2i}, z_{2i+1} common to $\mathfrak{g}$ and $\mathfrak{g}'$ is η'-relator-tied. A sub-geodesic of $\mathfrak{g}'$ composing one of the simple bigons is easily seen to be η'-relator-tied as $\eta' \leqslant \frac{1}{2} - 2\lambda$, hence the entire of $\mathfrak{g}'$ is η'-relator-tied.

The fact that $\mathfrak{g}'$ is contained in the η-succession of contours $t_0, t_1, \dots, t_{k+1}$ is immediate from the argument above: the sub-arcs of $\mathfrak{g}'$ with endpoints z_{2i}, z_{2i+1} are contained in $\mathfrak{g}$, while the sub-arcs with endpoints z_{2i+1}, z_{2i+2} are covered by contours τ that are in the set $\{t_0, t_1, \dots, t_{k+1}\}$. ∎

The goal of the following two statements is to prepare the ground for the definition of the η-criss-cross decomposition for a pair of vertices a, b.

Lemma 3.11 *Let a and b be two arbitrary vertices. The endpoints of an η-relator-tied component in a geodesic joining a, b are contained in any other geodesic joining a, b.*

Proof Let $\mathfrak{g}, \mathfrak{g}'$ be two geodesics with endpoints a, b and let x, y be the endpoints of an η-relator-tied component on $\mathfrak{g}$. Assume that x is not on $\mathfrak{g}'$. Then x is in the interior of one of the sides of a bigon composed by $\mathfrak{g}$ and $\mathfrak{g}'$. On the other hand, this side is $(\frac{1}{2} - 2\lambda)$-relator-tied, hence the component of $\mathfrak{g}$ between x, y is not a maximal η-relator-tied sub-geodesic, a contradiction.

It follows that $x \in \mathfrak{g}'$ and a similar argument shows that $y \in \mathfrak{g}'$. ∎

Definition 3.12 (Geodesic sequences)

(i) We say that a vertex p is *between two vertices a and b* if $\operatorname{dist}(a, p) + \operatorname{dist}(p, b) = \operatorname{dist}(a, b)$. We do not exclude that $p = a$ or $p = b$.

(ii) We call *geodesic sequence* a finite sequence of vertices $p_1, \dots, p_m$ such that for every $1 \leqslant i \leqslant j \leqslant k \leqslant m$, p_j is between p_i and p_k.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

(iii) If a, b, c, d is a geodesic sequence then we write $(b, c) \Subset (a, d)$, and we say that the pairs (b, c) and (a, d) are *nested*.

Lemma 3.13 *Let p, a, q, b be a geodesic sequence such that p, q and respectively a, b are the endpoints of η-relator-tied geodesics. Then there exists an η-succession of contours that contains every geodesic joining p and b.*

Proof We denote by $[p, q]$ (resp. $[a, b]$) the η-relator-tied geodesics. Consider two arbitrary geodesics $[p, a]$ and $[q, b]$, not necessarily contained in $[p, q]$ (resp. $[a, b]$).

In the geodesic $[p, a] \cup [a, b]$, the sub-geodesic $[a, b]$ is contained in a maximal η-relator-tied component $[a', b]$. Lemma 3.11 applied to p, b and the geodesic joining them $[p, q] \cup [q, b]$ implies that $a' \in [p, q]$; moreover, a' is on every geodesic joining p, b. Thus, by possibly replacing a with a' we may assume that a is contained in every geodesic with endpoints p, b, in particular that $a \in [p, q]$. A similar argument allows to state that without loss of generality we can assume that q is contained in every geodesic joining p, b, in particular $q \in [a, b]$.

By Corollary 3.9, there exist two η-successions of contours,

$$t_0, t_1, \ldots, t_{k+1} \quad \text{and} \quad \tau_0, \tau_1, \ldots, \tau_{m+1}$$

such that every geodesic joining p, q is contained in $\bigcup_{i=0}^{k+1} t_i$, and every geodesic joining a, b is contained in $\bigcup_{j=0}^{m+1} \tau_j$.

Consider i maximal such that $a \in t_i$.

Assume that $i \neq k+1$. If $\tau_0 \neq t_i$, then $[a, b] \cap \tau_0$ intersects t_i in a sub-geodesic of length $< \lambda|\tau_0|$. Consequently it intersects t_{i+1} in a sub-geodesic of length either at least $\lambda|\tau_0|$ or at least $(\eta - \lambda)|t_{i+1}|$. In both cases it follows $\tau_0 = t_{i+1}$, whence $a \in t_{i+1}$, which contradicts the choice of i.

Thus, in this case, it follows that $\tau_0 = t_i$.

Let $\ell \geqslant 0$ be maximal such that $\tau_r = t_{i+r}$ for $0 \leqslant r \leqslant \ell$. It is immediate from the definition of an η-succession that the sequence

$$t_0, \ldots, t_i = \tau_0, \ldots, t_{i+\ell} = \tau_\ell, \tau_{\ell+1}, \ldots, \tau_{m+1}$$

is such a succession.

An arbitrary geodesic joining p and b must contain a and q; the sub-geodesic from p to a must be contained in $\bigcup_{j=0}^{i} t_j$, while the sub-geodesic from a to b must be contained in $\bigcup_{r=0}^{m+1} \tau_r$.

Assume now that $i = k+1$. Every geodesic joining a, q must be contained in t_{k+1}.

Suppose moreover that $\tau_0 \neq t_{k+1}$. Then $\operatorname{dist}(a, q) < \lambda \min\{|t_{k+1}|, |\tau_0|\}$. Since the distance from q to one of the endpoints of $t_k \cap t_{k+1}$ is at least $\eta|t_{k+1}|$, the same is true about one of the endpoints of $\tau_0 \cap \tau_1$, since it will be situated after q on a geodesic from a to b. Therefore, in this case

$$t_0, t_1, \ldots, t_{k+1}, \tau_0, \tau_1, \ldots, \tau_{m+1}$$

is an η-succession of contours.

Given an arbitrary geodesic joining p and b, the sub-geodesic from p to q is in $\bigcup_{j=0}^{k+1} t_j$, and the sub-geodesic from a to b is in $\bigcup_{r=0}^{m+1} \tau_r$.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

Suppose that $\tau_0 = t_{k+1}$. As before, the fact that q is at distance $\geqslant \eta|\tau_0|$ from one of the endpoints of $t_k \cap t_{k+1}$ implies that one of the endpoints of $\tau_0 \cap \tau_1$ satisfies the same. Therefore,

$$t_0, t_1, \dots, t_{k+1} = \tau_0, \tau_1, \dots, \tau_{m+1}$$

is an η-succession of contours, and an argument as above shows that it contains every geodesic joining p and b. ∎

Remark 3.14 The statement of Lemma 3.13 can be generalized as follows: if

$$p_0, p_1, q_0, p_2, q_1, \dots, p_{k+1}, q_k, q_{k+1}$$

is a geodesic sequence such that p_i, q_i are the endpoints of η-relator tied geodesics for $i \in \{0, 1, \dots, k+1\}$, then there exists an η-succession of contours containing every geodesic from p_0 to q_{k+1}.

The proof adapts the argument of Lemma 3.13, and we leave it as an exercise to the reader.

Theorem 3.15 *For every pair of vertices a, b in the Cayley graph of G, there exists a finite geodesic sequence*

$$z_0 = a, y_1, z_1, y_2, z_2, \dots, y_m, z_m, b = y_{m+1},$$

a sequence of η-compulsory geodesics $[z_0, y_1], [z_1, y_2], \dots, [z_i, y_{i+1}], \dots, [z_m, y_{m+1}]$, and a sequence of η-successions of contours

$$t_1^{(i)}, \dots, t_{k_i}^{(i)}, i \in \{1, 2, \dots, m\}$$

such that $y_i \in t_1^{(i)}$, $z_i \in t_{k_i}^{(i)}$ and every geodesic joining a, b is contained in

$$[a, y_1] \cup \bigcup_{j=1}^{k_1} t_j^{(1)} \cup [z_1, y_2] \cup \bigcup_{j=1}^{k_2} t_j^{(2)} \cup \cdots \cup [z_{i-1}, y_i] \cup \bigcup_{j=1}^{k_i} t_j^{(i)} \cup [z_i, y_{i+1}] \cup \cdots \cup [z_m, b]. \tag{3.1}$$

Definition 3.16 (η-criss-cross decomposition) We say that the sequence

$$(a, y_1), \backslash y_1, z_1/, (z_1, y_2), \backslash y_2, z_2/, \dots, \backslash y_m, z_m/, (z_m, b)$$

is *the η-criss-cross decomposition for the pair a, b.*

Notation 3.17 *For an arbitrary pair of vertices a, b, we denote by $\mathcal{G}^\eta(a, b)$ the set described in Theorem 3.15; see (3.1) and Figure 5.*

Proof of Theorem 3.15 If a, b is an η-compulsory pair, there is nothing to prove. Assume therefore that there exists a geodesic joining a, b with an η-relator-tied component. Let $p_1, q_1, \dots, p_h, q_h$ be all the pairs of points that appear as endpoints of η-relator-tied components in some geodesic joining a, b. Let $\mathfrak{g}$ be an arbitrary geodesic joining a, b. According to Lemma 3.11, $\mathfrak{g}$ contains all points $p_1, q_1, \dots, p_h, q_h$. The order in which these points appear is independent of the choice of $\mathfrak{g}$, since it is only determined by metric relations.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

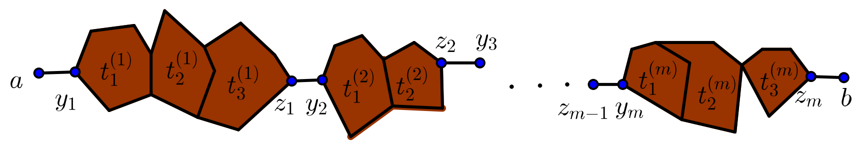

Figure 5: The η-criss-cross decomposition for the pair a, b.

We consider the union $\bigcup_{i=1}^{h}[p_i, q_i]$, where $[p_i, q_i]$ denotes the sub-geodesic of $\mathfrak{g}$ with endpoints p_i, q_i. The connected components of this union are sub-geodesics $[y_1, z_1], \dots, [y_m, z_m]$ appearing on $\mathfrak{g}$ in this order. Note that $y_i \in \{p_1, \dots, p_h\}$ and that $z_i \in \{q_1, \dots, q_h\}$. In particular, both the points and the order are independent of the choice of the geodesic $\mathfrak{g}$.

It remains to apply Lemma 3.13 and Remark 3.14. ∎

Corollary 3.18 *For every pair of points a, b at distance $d > 0$ and every $x \leqslant d$ there exist at most two points p with the property that a, p, b is a geodesic sequence and* $\operatorname{dist}(a, p) = x$.

See Figure 6 for an example where there exist two points q_1, q_2 between a and b, at distance $x - 3$ from a, and two points p_1, p_2 between a and b, at distance x from a.

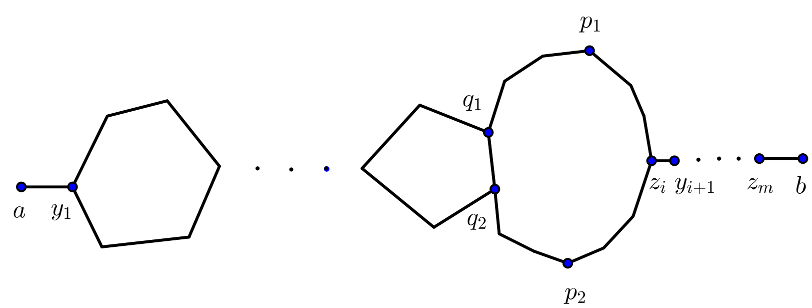

Figure 6: Example of pair a, b with two points between them at distance x from a.

Remark 3.19 (i) According to the above, every geodesic with endpoints y_m, z_m is η'-relator-tied, in particular, it is non-trivial.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

(ii) Due to the maximality condition defining the pairs y_i, z_i, we have that $z_i \neq y_{i+1}$ for every $1 \leqslant i \leqslant m-1$.

(iii) On the other hand, in the pairs $(a, y_1), (z_m, b)$ the endpoints may coincide.

Remark 3.20 If $(p, q) \Subset (x, y) \Subset (a, b)$ with p, q, x, y points in $\{p_1, q_1, \dots, p_k, q_k\}$, and if p, q are endpoints of an η-relator-tied component in a geodesic joining a, b, then p, q are endpoints of an η-relator-tied component in a geodesic joining x, y.

This simply follows from the fact that a geodesic $\mathfrak{g}$ joining a, b and on which p, q bound an η-relator-tied component must also contain x, y; see Lemma 3.11.

Definition 3.21 (η-relator covered pair) If the η-criss-cross decomposition of a pair a, b is $\backslash a, b/$, then we call such a pair an *η-relator covered pair*.

Definition 3.22 (Compulsory vertices) Given an η-criss-cross decomposition

$$(a, y_1), \backslash y_1, z_1/, (z_1, y_2), \backslash y_2, z_2/, \dots, \backslash y_m, z_m/, (z_m, b)$$

of a pair a, b, we call the vertices between z_i, y_{i+1} for some $i \in \{1, 2, \dots, m-1\}$ *η-compulsory vertices*.

Clearly, every geodesic with endpoints a and b must contain all the compulsory vertices.

Definition 3.23 (Prefixes and suffixes) Given an element $h \in G$, we denote by $P(h)$ (standing for *prefixes of* h) all the elements between $1, h$ and by $S(h)$ (standing for *suffixes of* h) all the elements of the form $x^{-1}h$ for $x \in P(h)$.

Note that the two sets $P(h)$ and $S(h)$ depend on the fixed generating set A.

Definition 3.24 (Compulsory and η-relator-covered elements) Let $h \in G$.

- If h is joined to 1 by at least one η-relator-tied geodesic , then we call h an *η-relator-tied element*.
- If the pair $1, h$ has the η-criss-cross decomposition $(1, h)$ (hence, there exists only one geodesic joining $1, h$, composed of compulsory vertices), then we call h an *η-compulsory element*.
- If the pair $1, h$ has the η-criss-cross decomposition $\backslash 1, h/$, then we call h an *η-relator-covered element*.

Notation 3.25 *We denote by RT^η the set of η-relator-tied elements. We denote by $\mathcal{C}^\eta$ the set of η-compulsory elements in G and by RC^η the set of η-relator-covered elements.*

Remark 3.26 The fact that h is η-relator-covered does not mean that there exists an η-relator-tied geodesic labeled by h; it only means that every geodesic $[a, b]$ labeled by h contains a family of successive vertices $y_0 = a, y_1, z_0, y_2, z_1, \dots, y_m, z_{m-1}, z_m = b$ such that for every i, there exists an η-relator-tied geodesic with endpoints y_i, z_i. In particular, by Lemma 3.10, every geodesic labeled by h is η'-relator-tied.

The following theorem is an algebraic version of Theorem 3.15.

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

Theorem 3.27 *Every element $g \in G$ can be written uniquely as a product*

$$g = \alpha_1 \beta_1 \alpha_2 \beta_2 \dots \alpha_m \beta_m \alpha_{m+1}$$

such that

- *β_i are non-trivial η-relator-covered elements;*
- *α_i are compulsory elements (non-trivial with the possible exception of α_1, α_{m+1});*
- *the vertices $y_i = \alpha_1 \beta_1 \dots \alpha_i$ and $z_i = \alpha_1 \beta_1 \dots \alpha_i \beta_i$ compose the geodesic sequence determining the η-criss-cross decomposition of the pair $1, g$.*

Notation 3.28 *Given an arbitrary element $h \in G$, we denote by $\mathcal{G}^\eta(h)$ the set $\mathcal{G}^\eta(1, h)$ as described in Notation 3.17.*

Notation 3.29 *Given $i \in \mathbb{N}$, $i \geqslant 2$, we denote by $\mathcal{D}_i$ the set of i-tuples*

$$(a_1, a_2, \dots, a_{i-1}, b)$$

such that for the element $g = a_1 a_2 \cdots a_{i-1} b$ the elements $a_1, a_2, \dots, a_{i-1}$ are the first $i-1$ elements in the criss-cross decomposition of g as described in Theorem 3.27.

The following lemma will be crucial for the results in Section 4 on quasi-homomorphisms.

Lemma 3.30 *Let $\lambda \leqslant \frac{1}{10}$ and let $\eta \geqslant 3\lambda$.*

Every η-succession of contours $t_0, t_1, \dots, t_{k+1}$ is totally geodesic: if a, b are two vertices in $\bigcup_{i=0}^{k+1} t_i$, then every geodesic joining a and b is contained in $\bigcup_{i=0}^{k+1} t_i$.

Proof Without loss of generality we assume that $a \in t_1 \setminus t_2$ and that $b \in t_{k+1} \setminus t_k$. Otherwise, assuming that a appears before b in the succession, we consider the largest i such that t_i contains a and the smallest j such that t_j contains b and take the succession $t_i, t_{i+1}, \dots, t_{j-1}, t_j$ instead of the initial one.

Let $\mathfrak{g}$ be a geodesic joining a and b. We argue for a contradiction and assume that $\mathfrak{g}$ is not contained in $\bigcup_{i=0}^{k+1} t_i$. Without loss of generality, we assume that $\mathfrak{g}$ intersects $\bigcup_{i=0}^{k+1} t_i$ only in its endpoints (otherwise, we replace $\mathfrak{g}$ by a sub-geodesic with this property).

Let $\mathfrak{p}$ be a topological arc joining a and b in $\bigcup_{i=0}^{k+1} t_i$ and of minimal length. By the Greendlinger Lemma, there exists a contour τ such that one of the connected components of its intersection with $\mathfrak{p} \cup \mathfrak{g}$ has length $> (1-3\lambda)|\tau|$. If $\tau = t_i$ for some i, then by the hypothesis on $\mathfrak{g}$, τ intersects $\mathfrak{p}$ in a connected component of length $> (1-3\lambda)|\tau|$. Then $\mathfrak{p}$ can be shortened by a length of $(1-6\lambda)|\tau|$, which contradicts the choice of $\mathfrak{p}$ as an arc of minimal length joining a and b in $\bigcup_{i=0}^{k+1} t_i$.

We therefore assume that $\tau \notin \{t_0, t_1, \dots, t_{k+1}\}$. Since $\mathfrak{g}$ is a geodesic, it follows that τ intersects $\mathfrak{p}$ in a subarc of length $> (\frac{1}{2} - 3\lambda)|\tau|$.

On the other hand, $\mathfrak{p}$ contains a succession of vertices

$$x_0 = a, x_1, y_0, x_2, y_1, \dots, x_{k+1}, y_k, y_{k+1} = b$$

such that the sub-arcs with endpoints x_i, y_i with $i \in \{0, 1, \dots, k+1\}$ are labeled by words in $S(R)$, which are moreover in $S^\lambda(R)$ if $i \neq 0, k+1$. Therefore, the connected

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

component of the intersection $\tau \cap \mathfrak{p}$ cannot contain a pair x_i, y_i with $i \in \{1, \ldots, k\}$. It follows that it can intersect at most two consecutive sub-arcs with endpoints x_i, y_i with $i \in \{0, 1, \ldots, k+1\}$, hence it is of length $< 2\lambda|\tau|$. We thus obtain that $\frac{1}{2} - 3\lambda < 2\lambda$, whence $\lambda > \frac{1}{10}$, a contradiction. ∎

The following results are not used in an essential manner in our arguments, but they complete nicely the description of geodesics in small cancellation groups.

Lemma 3.31 *Let $\mathfrak{g}$ be an η-relator-tied geodesic and let $\mathfrak{p}$ be a sub-geodesic in it. Then $\mathfrak{p}$ is either an η-compulsory geodesic, or it is the concatenation of three sub-geodesics $\mathfrak{p} = \mathfrak{p}_c \sqcup \mathfrak{p}_0 \sqcup \mathfrak{p}'_c$, where $\mathfrak{p}_c, \mathfrak{p}'_c$ are η-compulsory and contained in a contour (possibly either one of them or both trivial) and $\mathfrak{p}_0$ is an η-relator-tied component of $\mathfrak{p}$.*

Proof Let a, b be the endpoints of $\mathfrak{g}$. With the previous convention $\mathfrak{g} = [a, b]$.

Step 1. Let us first assume that $\mathfrak{p} = [a, \sigma]$, with a, σ, b a geodesic sequence.

Let $x_0 = a, x_1, y_0, x_2, y_1, \ldots, x_{k+1}, y_k, y_{k+1} = b$ be the unique sequence of points on $\mathfrak{g}$ defined by Lemma 3.6.

Assume that σ is in between a pair y_j, x_{j+2}. If the word labeling the geodesic $[x_{j+1}, \sigma]$ is contained in $S^\eta(R)$, then $\mathfrak{p}$ is η-relator-tied.

If the word labeling $[x_{j+1}, \sigma]$ is not in $S^\eta(R)$ (while it is still a sub-word of the relator labeling the contour t_{j+1}), then the pair x_{j+1}, σ is η-compulsory. This implies that the required decomposition is $\mathfrak{p} = [a, y_j] \sqcup [y_j, \sigma]$.

Assume now that σ is in between a pair x_{j+1}, y_j. If $[x_j, \sigma]$ is labeled by a word in $S^\eta(R)$, then $\mathfrak{p}$ is η-relator-tied; while in the opposite case the geodesic $[x_j, \sigma]$ is η-compulsory, and the conclusion holds with the decomposition $\mathfrak{p} = [a, y_{j-1}] \sqcup [y_{j-1}, \sigma]$.

Step 2. Assume now that $\mathfrak{p} = [\rho, \sigma]$, where a, ρ, σ, b is a geodesic sequence. According to Step 1, $[a, \sigma] = [a, \mu] \sqcup [\mu, \sigma]$, where $[a, \mu]$ is an η-relator-tied component and $[\mu, \sigma]$ is η-compulsory (possibly trivial) and contained in a contour. If $\rho \in [\mu, \sigma]$, then $\mathfrak{p}$ is η-compulsory. If $\rho \in [a, \mu]$, then by reversing the order on $[a, \mu]$ and applying Step 1, we obtain that $[\rho, \mu] = [\rho, \nu] \sqcup [\nu, \mu]$, where $[\rho, \nu]$ is η-compulsory (possibly trivial) and contained in a contour, and $[\nu, \mu]$ is an η-relator-tied component. It follows that

$$\mathfrak{p} = [\rho, \mu] \sqcup [\mu, \sigma] = [\rho, \nu] \sqcup [\nu, \mu] \sqcup [\mu, \sigma]$$

is the required decomposition. ∎

Lemma 3.32 *For each pair $\backslash y_j, z_j /$ in an η-criss-cross decomposition, there exists a geodesic sequence*

$$p'_1 = y_j, p'_2, q'_1, p'_3, q'_2, \ldots, p'_n, q'_{n-1}, q'_n = z_j, \text{ for some } n = n(j) \tag{3.2}$$

such that

- *(p'_s, q'_s) are maximal with respect to the partial order relation $\Subset$;*
- *$p'_{\ell+1}, q'_\ell$ bound η-relator-tied sub-geodesics both in the η-relator-tied geodesic joining p'_ℓ, q'_ℓ and in the η-relator-tied geodesic joining $p'_{\ell+1}, q'_{\ell+1}$.*

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

Proof By definition, $[y_j, z_j] = \bigcup_{i\in I_j}[p_i, q_i]$. Without loss of generality, we assume that there are no nested pairs among the (p_i, q_i) with $i \in I_j$; in other words, each pair (p_i, q_i) is maximal with respect to the partial order relation $\Subset$. Proceeding as in the proof of Lemma 3.6, we also assume that, after selecting a subset in I_j, every point on a (every) geodesic joining y_j, z_j is between at most two pairs (p_i, q_i). It then follows that the set of pairs indexed by I_j compose a geodesic sequence as in (3.2). We set $p'_i := p_i$ and $q'_i := q_i$.

Now consider two consequent pairs that overlap: two pairs (p'_i, q'_i) and (p'_{i+1}, q'_{i+1}) such that $p'_i, p'_{i+1}, q'_i, q'_{i+1}$ is a geodesic sequence.

By definition, there exists a geodesic $\mathfrak{g}$ joining a, b such that p'_i, q'_i are the endpoints of an η-relator-tied component of it. Given two points $x, y \in \mathfrak{g}$, we denote by $[x, y]$ the sub-geodesic of $\mathfrak{g}$ with endpoints x, y.

We likewise know that there exists a geodesic $\mathfrak{p}$ such that p'_{i+1}, q'_{i+1} bound an η-relator-tied component on $\mathfrak{p}$. According to the above, $\mathfrak{p}$ must contain q'_i. In what follows, for x, y in $\mathfrak{p}$, we denote by $\overline{x, y}$ the sub-arc of $\mathfrak{p}$ with endpoints x, y.

We have that $p'_{i+1} \in [p'_i, q'_i]$. Lemma 3.31 implies that either $[p'_{i+1}, q'_i]$ is an η-compulsory component contained in a contour, or $[p'_{i+1}, q'_i] = [p'_{i+1}, x] \sqcup [x, q'_i]$, where $[p'_{i+1}, x]$ is an η-compulsory component contained in a contour (possibly trivial) and $[x, q'_i]$ is an η-relator-tied component.

Assume that $[p'_{i+1}, q'_i]$ is an η-compulsory component contained in a contour. Then the geodesic $\mathfrak{p}$ must also contain $[p'_{i+1}, q'_i] \subset \mathfrak{g}$. By replacing on $\mathfrak{p}$ the sub-arc with endpoints a, p'_{i+1} by $[a, p'_{i+1}] \subset \mathfrak{g}$, we obtain a new geodesic $\mathfrak{r}$ joining a, b such that p'_i and q'_{i+1} are the endpoints of an η-relator-tied sub-geodesic. It follows that $(p'_i, q'_{i+1}) \Subset (\alpha, \beta)$, where α, β are the endpoints on $\mathfrak{r}$ of an η-relator-tied component. In particular $(\alpha, \beta) = (p_\ell, q_\ell)$ for some $\ell \in I_j$, and $(p'_i, q'_i) \Subset (p_\ell, q_\ell)$. This contradicts the fact that we have considered pairs maximal with respect to $\Subset$.

Assume that $[p'_{i+1}, q'_i] = [p'_{i+1}, x] \sqcup [x, q'_i]$, where $[p'_{i+1}, x]$ is an η-compulsory component contained in a contour (possibly trivial) and $[x, q'_i]$ is an η-relator-tied component. Since $q'_i \in \mathfrak{p}$ and $[x, q'_i]$ is an η-relator-tied component between p'_{i+1} and q'_i it follows that $x \in \mathfrak{p}$, hence $[p'_{i+1}, x] \subset \mathfrak{p}$. There exists $y \in \mathfrak{p}$ such that $\overline{p'_{i+1}, y}$ is labeled by a word in $S^\eta(R)$ and it is contained in a contour t. If $y \in \overline{p'_{i+1}, x} = [p'_{i+1}, x]$, then the contour t intersects a distinct contour in a sub-arc of length $\geqslant \eta|t|$, a contradiction. Hence, we must have that $x \in \overline{p'_{i+1}, y}$.

According to the small cancellation condition $\overline{x, y}$ has length $> (1-\frac{\lambda}{\eta})$ of the length of $\overline{p'_{i+1}, y}$, so at least $\eta(1-\frac{\lambda}{\eta})|t|$. This implies that if $\eta(1-\frac{\lambda}{\eta}) \geqslant \lambda$, equivalently $\eta \geqslant 2\lambda$, then by Lemma 3.10, t must be the first contour for the pair x, q'_i. But this implies that $p'_{i+1} = x$.

Similarly, we argue that $\overline{p'_{i+1}, q'_i}$ is an η-relator-tied geodesic. ∎

4 Quasi-homomorphisms on Small Cancellation Groups

Recall that a *quasi-homomorphism* (also called a *quasi-morphism* or a *pseudo-character*) on a group G is a function $\mathfrak{h}\colon G \to \mathbb{R}$ such that its defect

$$\mathfrak{d}(\mathfrak{h}) := \sup_{a,b\in G} \left|\mathfrak{h}(ab) - \mathfrak{h}(a) - \mathfrak{h}(b)\right|$$

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

is finite. The real vector space $\mathcal{Q}(G)$ of all quasi-homomorphisms of G has three important subspaces: the subspace $\ell^\infty(G)$ of bounded real functions on G, the subspace $\mathrm{Hom}(G,\mathbb{R}) = H^1(G,\mathbb{R})$ of homomorphisms on G, and the subspace $\ell^\infty(G) + \mathrm{Hom}(G,\mathbb{R})$ of functions that differ from a homomorphism by a bounded function. Consider the quotient spaces

$$QH(G) = \mathcal{Q}(G)/\mathcal{B}(G) \quad \text{and} \quad \widetilde{QH}(G) = \mathcal{Q}(G)/\big[\ell^\infty(G) + \mathrm{Hom}(G,\mathbb{R})\big].$$

The space $\widetilde{QH}(G)$ can be identified with the kernel of the comparison map

$$H^2_b(G) \longrightarrow H^2(G),$$

where $H^2_b(G)$ is the second bounded cohomology of G.

In this paper, as an application of our results on the geometry of small cancellation groups with the $C'(1/12)$-condition, we show that for such a group G, the space $\widetilde{QH}(G)$ is infinite dimensional with a basis of power continuum.

Following the work of Epstein and Fujiwara [EF97, Fuj00, Fuj98] as well as of Bestvina and Fujiwara [BF02], we prove the following proposition.

Proposition 4.1 *Let G be a finitely generated infinitely presented group and let $\langle S \mid R\rangle$ be a presentation such that R satisfies the $C'(1/12)$-condition. For a given $\eta \in [3\lambda, \frac{1}{2} - 2\lambda]$ appropriately chosen, there exists a sequence $\mathfrak{u}_n$ of elements in G and a sequence $\mathfrak{h}_{\mathfrak{u}_n} \colon G \to \mathbb{R}$ of quasi-homomorphisms, with $n \in \mathbb{N}, n \geqslant 1$, such that*

(i) *the set of word lengths $|\mathfrak{u}_n|$ diverges to ∞;*
(ii) *every group homomorphism $\phi\colon G \to \mathbb{R}$ has the property that $\phi(\mathfrak{u}_n) = 0$ for every $n \in \mathbb{N}, n \geqslant 1$;*
(iii) *the sequence of defects $\mathfrak{d}(\mathfrak{h}_{\mathfrak{u}_n})$ is bounded;*
(iv) *for every n and every $k \in \mathbb{N}, k \geqslant 1$, $\mathfrak{h}_{\mathfrak{u}_n}(\mathfrak{u}_n^k) = k$;*
(v) *for every $n \neq m$, and every $k \in \mathbb{N}, k \geqslant 1$, $\mathfrak{h}_{\mathfrak{u}_n}(\mathfrak{u}_m^k) = 0$.*

Proof We enumerate the relators $\{r_1, r_2, \dots\}$ in R so that their lengths compose a non-decreasing sequence. Consider the sequence of finite subsets of $\mathbb{N}$ defined by

$$I_n = [1 + 2 + \cdots + n, 1 + 2 + \cdots + n + 1) \cap \mathbb{N}.$$

Define two sequences of finite subsets A_n and B_n of R, described by $A_n = \{r_{2i-1} \mid i \in I_n\}$ and $B_n = \{r_{2i} \mid i \in I_n\}$.

To simplify the notation, in what follows, we denote the relator r_{2i-1} by α_i and r_{2i} by β_i, respectively. Thus, $A_n = \{\alpha_i \mid i \in I_n\}$ and $B_n = \{\beta_i \mid i \in I_n\}$.

Let X be a finite set of relators equal either to a set A_n or to a set B_n. We construct an element $x \in G$ corresponding to X, as follows. Assume X is composed of the relators $\rho_1, \dots, \rho_k$ enumerated in increasing order. For every $i \in \{1, 2, \dots, k\}$ let y_i be the prefix of ρ_i of length $\lfloor |\rho_i|/2 \rfloor$. Define the element $x = y_1 y_2 \cdots y_k$. An argument very similar to the one in Lemma 3.30 implies that x is an η-relator-tied element and that every geodesic joining 1 and x is contained in the η-succession of contours $t_1, y_1 t_2, y_1 y_2 t_3, \dots, [y_1 \cdots y_{k-1}] t_k$, where t_i is the loop through 1 in the Cayley graph, labeled by ρ_i.

When $X = A_n$ (resp. $X = B_n$) the corresponding element x is denoted by a_n (resp. by b_n).

We define $\mathfrak{u}_n = [a_n, b_n]$. This implies Proposition 4.1(ii).

Lemma 3.30 applied to geodesics joining 1 to $\mathfrak{u}_n$ implies that the length $|\mathfrak{u}_n|$ is at least the double of $(\frac{1}{2} - \lambda) \sum_{i \in I_n} [|\alpha_i| + |\beta_i|]$. It follows that Proposition 4.1(i) is also satisfied.

We now define the sequence of quasi-homomorphisms. We start with a general construction. Let $\mathfrak{v}$ be an η-relator-tied element in G.

Definition 4.2 (i) Let $(a, b) \in G \times G$. A *quasi-copy of* $\mathfrak{v}$ *nested inside* (a, b) is a pair of points $x, y \in \mathcal{G}^\eta(a, b)$ such that $y = x\mathfrak{v}$ and such that there exists an η-succession of contours $t_1, \dots, t_k$ contained in $\mathcal{G}^\eta(a, b)$ such that:

- x is either one of the endpoints of the intersection of t_1 with a contour t_0 such that $t_0, t_1, \dots, t_k$ is an η-succession contained in $\mathcal{G}^\eta(a, b)$, or the intersection of t_1 with a compulsory geodesic preceding $t_1, \dots, t_k$ in $\mathcal{G}^\eta(a, b)$;
- y is either one of the endpoints of the intersection of t_k with a contour t_{k+1} such that $t_1, \dots, t_k, t_{k+1}$ is an η-succession contained in $\mathcal{G}^\eta(a, b)$, or the intersection of t_k with a compulsory geodesic succeeding $t_1, \dots, t_k$ in $\mathcal{G}^\eta(a, b)$.

(ii) We say that two quasi-copies of $\mathfrak{v}$ nested inside a, b are *non-overlapping* if the corresponding η-successions of contours $t_1, \dots, t_k$, respectively $\tau_1, \dots, \tau_k$, are disjoint, as finite sets of contours.

(iii) When $(a, b) = (1, g)$ for some element $g \in G$ we speak about *quasi-copies of* $\mathfrak{v}$ *nested inside* g.

Note that according to the definition of $\mathcal{G}^\eta(a, b)$ and to Lemma 3.10, the pair of points x, y uniquely determines the η-succession $t_1, \dots, t_k$.

Lemma 4.3 *Let x, y be a pair of points in $\mathcal{G}^\eta(g)$ (with the Notation 3.28) composing a nested quasi-copy of $\mathfrak{v}$ in g, and let $t_1, \dots, t_k$ be the corresponding η-succession of contours. There exists no other pair of points p, q in $\bigcup_{i=1}^k t_i$ such that $q = p\mathfrak{v}$.*

Proof Lemma 3.30 can be easily generalized to pairs of points a, b contained in an η-succession of contours. Applied to the pair x, y, it implies that every geodesic joining x, y is η-relator-tied. This implies that $\mathfrak{v}$ is an η-relator-tied element. Let $\mathfrak{g}$ be an η-relator-tied geodesic joining 1 and $\mathfrak{v}$. It follows that $x\mathfrak{g}$ is contained in $\bigcup_{i=1}^k t_i$, whence the unique sequence of vertices on $\mathfrak{g}$ described in Lemma 3.6 contains k pairs x_i, y_i.

Assume that there exists another pair of points $p \in t_r$ and $q \in t_s$ with $1 \leqslant r \leqslant s \leqslant k$ such that p, q compose a nested quasi-copy of $\mathfrak{v}$ in g. The $p\mathfrak{g}$ is a geodesic joining p and q, which, according to Lemma 3.30, is contained in $\bigcup_{i=r}^s t_i$. The uniqueness of the sequence in Lemma 3.6 implies that $s - r + 1 = k$, whence $r = 1$ and $s = k$. The same uniqueness implies that each pair px_i, py_i, translate of the corresponding pair on $\mathfrak{g}$, is the pair of endpoints of the intersection $p\mathfrak{g} \cap t_i$.

The first pair in the unique sequence of vertices on $\mathfrak{g}$ as in Lemma 3.6 is of the form $1, h$, where h is represented by a word w_1 in $S^{\frac{1}{2}-2\lambda}(R)$, prefix of a relator ρ labeling a unique loop τ through 1 in the Cayley graph. By the above, $x\tau = p\tau = t_1$; therefore, $p^{-1}x\tau = \tau$. This and the small cancellation condition $C'(1/12)$ imply that the element $p^{-1}x$ is trivial in G. Indeed, the condition $C'(1/12)$ implies that the stabilizer in G of

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

any contour is trivial; otherwise, one could find two distinct copies of the same long sub-word in the label of that contour.

We conclude that $p = x$, and $q = p\mathfrak{v} = x\mathfrak{v} = y$. ∎

Definition 4.4 The point x is called the *initial point* of the nested quasi-copy, while y is called the *terminal point* of the nested quasi-copy.

We define $c_{\mathfrak{v}}: G \times G \to \mathbb{R}$ such that $c_{\mathfrak{v}}(a,b)$ is the maximal number of pairwise non-overlapping quasi-copies of $\mathfrak{v}$ nested inside (a,b).

By abuse of notation, we define $c_{\mathfrak{v}}: G \to \mathbb{R}$ such that $c_{\mathfrak{v}}(g)$ is the maximal number of pairwise non-overlapping quasi-copies of $\mathfrak{v}$ nested inside g.

Clearly, $c_{\mathfrak{v}}(a,b) = c_{\mathfrak{v}}(ha,hb)$ and $c_{\mathfrak{v}}(g) = c_{\mathfrak{v}}(h,hg)$, for every $h \in G$.

Proposition 4.5 *Let $\mathfrak{v}$ be one of the elements $\mathfrak{u}_n$ for $n \in \mathbb{N}$. The map $\mathfrak{h}_{\mathfrak{v}}: G \to \mathbb{R}$, $\mathfrak{h}_{\mathfrak{v}} = c_{\mathfrak{v}} - c_{\mathfrak{v}^{-1}}$ is a quasi-morphism with defect at most 2.*

Proof Let g and h be two arbitrary elements in G. Our goal is to show that

$$\left|\mathfrak{h}_{\mathfrak{v}}(gh) - \mathfrak{h}_{\mathfrak{v}}(g) - \mathfrak{h}_{\mathfrak{v}}(h)\right| \leqslant 2.$$

The study of geodesic triangles that was done in the preceding section implies that the intersection $\mathcal{G}^{\eta}(g) \cap \mathcal{G}^{\eta}(h) \cap \mathcal{G}^{\eta}(g,gh)$ is either a contour or a tripod (with some branches possibly reduced to a point) appearing as intersection of three contours, or a sub-path in a contour ω composed of three consecutive sub-paths (possibly reduced to a point) of lengths $< \lambda|\omega|$, for the first and third, and $< \eta|\omega|$ for the second. Note that whatever the geometric nature of the intersection, it splits each of the three sets $\mathcal{G}^{\eta}(g)$, $\mathcal{G}^{\eta}(h)$, $\mathcal{G}^{\eta}(g,gh)$, into two connected components.

We call the intersection $\mathcal{G}^{\eta}(g) \cap \mathcal{G}^{\eta}(h) \cap \mathcal{G}^{\eta}(g,gh)$ the *median object* for the triple g, h, gh, and we denote it $\mathfrak{m}(g,h)$.

We say that $\mathfrak{m}(g,h)$ *separates a quasi-copy of* $\mathfrak{v}$ *nested inside* (a,b), where $(a,b) \in \{(1,g),(1,gh),(g,gh)\}$ if the two points x, y determining that quasi-copy are in two different connected components of $\mathcal{G}^{\eta}(a,b) \smallsetminus \mathfrak{m}(g,h)$.

Assume that the maxima $c_{\mathfrak{v}^{\pm 1}}(g)$, $c_{\mathfrak{v}^{\pm 1}}(gh)$ and $c_{\mathfrak{v}^{\pm 1}}(g,gh)$ are all attained only by considering nested quasi-copies that are not separated by $\mathfrak{m}(g,h)$. In that case one can easily see that $\mathfrak{h}_{\mathfrak{v}}(gh) - \mathfrak{h}_{\mathfrak{v}}(g) - \mathfrak{h}_{\mathfrak{v}}(h) = 0$.

Assume now that every counting that realizes the maximum $c_{\mathfrak{v}}(gh)$ must take into account a pair x, y separated by $\mathfrak{m}(g,h)$. Inside $\mathcal{G}^{\eta}(gh)$, one has then an η-succession of contours $t_1, \dots, t_k$ with $x \in t_1$ and $y \in t_k$. The choice of the labels of contours in $\mathcal{G}^{\eta}(\mathfrak{u}_n)$ implies that:

- no quasi-copy of $\mathfrak{v}^{-1}$ nested inside gh can contain a sub-sequence in the sequence of contours $t_1, \dots, t_k$;
- no initial point of a quasi-copy of $\mathfrak{v}$ nested inside g can be contained in $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^{\eta}(g)$;
- no terminal point of a quasi-copy of $\mathfrak{v}$ nested inside (g,gh) can be contained in $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^{\eta}(g,gh)$.

It is nevertheless possible that $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^{\eta}(g)$ contains an initial point of a quasi-copy of $\mathfrak{v}^{-1}$ nested inside g. But in that case, no terminal point of a quasi-copy of $\mathfrak{v}^{-1}$

https://doi.org/10.4153/CJM-2018-036-7 Published online by Cambridge University Press

nested inside (g, gh) can be contained in $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^\eta(g, gh)$. We thus obtain that

$$\mathfrak{h}_{\mathfrak{v}}(gh) - \mathfrak{h}_{\mathfrak{v}}(g) - \mathfrak{h}_{\mathfrak{v}}(h) = 2. \tag{4.1}$$

Similarly, $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^\eta(g, gh)$ may contain a terminal point of a quasi-copy of $\mathfrak{v}^{-1}$ nested inside (g, gh); in which case $\bigcup_{i=1}^{k} t_i \cap \mathcal{G}^\eta(g)$ cannot contain an initial point of a quasi-copy of $\mathfrak{v}^{-1}$ nested inside g, and (4.1) is still verified.

If none of the above two cases occurs, then the right-hand side in (4.1) is 1.

In the case when every counting that realizes the maximum $c_{\mathfrak{v}^{-1}}(gh)$ must take into account a pair x, y separated by $\mathfrak{m}(g, h)$ similar arguments work and give equalities as in (4.1), with the right hand side either -2 or -1.

The cases when $c_{\mathfrak{v}^{\pm 1}}(gh)$ is replaced by either $c_{\mathfrak{v}^{\pm 1}}(g)$ or $c_{\mathfrak{v}^{\pm 1}}(g, gh)$ are treated similarly and give equalities like in (4.1), with the right hand side ± 2 or ± 1. ∎

We now finish the proof of Proposition 4.1. Proposition 4.5 implies that all the quasi-homomorphisms $\mathfrak{h}_{\mathfrak{u}_n}$ have a defect bounded by 2. Properties (iv) and (v) follow from Corollary 3.9 and from the construction of the η-relator-tied elements $\mathfrak{u}_n$. ∎

The end of the proof now follows the standard argument in the work of Epstein–Fujiwara [EF97, Fuj00, Fuj98] and Bestvina–Fujiwara [BF02]. We repeat it here for the sake of completeness.

Theorem 4.6 *Let G be an infinitely presented finitely generated group given by a presentation satisfying the small cancellation condition $C'(1/12)$. Then there exists an injective linear map $\ell^1 \to \widetilde{QH}(G)$. In particular, the dimension of $\widetilde{QH}(G)$ is power continuum.*

Proof We consider the map $\ell^1 \to \mathcal{Q}(G)$ defined by $(a_n) \mapsto \sum_n a_n \mathfrak{h}_{\mathfrak{u}_n}$. Proposition 4.1(iii) implies that each image is indeed a quasi-morphism. Proposition 4.1(i) implies that when $a_n \mathfrak{h}_{\mathfrak{u}_n}$ is evaluated in some element $g \in G$, only finitely many terms take non-zero value, thus the sum is always finite.

The above map defines a linear map $\ell^1 \to \widetilde{QH}(G)$. We now prove that it is injective. Let $(a_n) \in \ell^1$ be such that $\mathfrak{h} = \sum_n a_n \mathfrak{h}_{\mathfrak{u}_n}$ is at bounded distance from a homomorphism. In particular, it follows by Proposition 4.1(ii) that for every n and k, $\mathfrak{h}(\mathfrak{u}_n^k)$ is uniformly bounded.

On the other hand, given $n \in \mathbb{N}$ such that $a_n \neq 0$, Proposition 4.1(iv) and (v) imply that $\mathfrak{h}(\mathfrak{u}_n^k) = a_n k$. This contradicts the fact that $\mathfrak{h}(\mathfrak{u}_n^k)$ is bounded uniformly in k. ∎

Acknowledgments The work on this paper was carried out during visits of the first author to the University of Oxford and of the second author to the University of Vienna. We thank these institutions for their support and hospitality. The first author is also grateful to the ETH Zürich and to the CRM in Barcelona, and both authors wish to thank the Isaac Newton Institute in Cambridge for its hospitality during the completion of the paper.

We thank Thomas Delzant, Denis Osin, and the anonymous referee for useful comments and corrections.

References

[Bav91] C. Bavard, *Longueur stable des commutateurs*. Enseign. Math. (2) **37**(1991), no. 1-2, 109–150.

[BF02] M. Bestvina and K. Fujiwara, *Bounded cohomology of subgroups of mapping class groups*. Geom. Topol. **6**(2002), 69–89 (electronic). http://dx.doi.org/10.2140/gt.2002.6.69

[BM99] M. Burger and N. Monod, *Bounded cohomology of lattices in higher rank Lie groups*. J. Eur. Math. Soc. (JEMS) **1**(1999), no. 2, 199–235. http://dx.doi.org/10.1007/s100970050007

[BM02] ______, *Continuous bounded cohomology and applications to rigidity theory*. Geom. Funct. Anal. **12**(2002), no. 2, 219–280. http://dx.doi.org/10.1007/s00039-002-8245-9

[Bro81] R. Brooks, *Some remarks on bounded cohomology*, Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978), Ann. of Math. Stud., vol. 97, Princeton Univ. Press, Princeton, N.J., 1981, pp. 53–63.

[DGO17] F. Dahmani, V. Guirardel, and D. Osin, *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces*. Mem. Amer. Math. Soc. **245**(2017), no. 1156. http://dx.doi.org/10.1090/memo/1156

[EF97] D. B. A. Epstein and K. Fujiwara, *The second bounded cohomology of word-hyperbolic groups*. Topology **36**(1997), no. 6, 1275–1289. http://dx.doi.org/10.1016/S0040-9383(96)00046-8

[Fuj98] K. Fujiwara, *The second bounded cohomology of a group acting on a Gromov-hyperbolic space*. Proc. London Math. Soc. (3) **76**(1998), no. 1, 70–94. http://dx.doi.org/10.1112/S0024611598000033

[Fuj00] ______, *The second bounded cohomology of an amalgamated free product of groups*. Trans. Amer. Math. Soc. **352**(2000), no. 3, 1113–1129. http://dx.doi.org/10.1090/S0002-9947-99-02282-5

[Fuj09] ______, *Quasi-homomorphisms on mapping class groups*. In: Handbook of Teichmüller theory. Vol. II, IRMA Lect. Math. Theor. Phys., 13, Eur. Math. Soc., Zürich, 2009, pp. 241–269.

[GdlH91] E. Ghys and P. de la Harpe, *Infinite groups as geometric objects (after Gromov)*. In: Ergodic theory, symbolic dynamics and hyperbolic spaces (Trieste, 1989) Oxford University Press, 1991, pp. 299–314.

[Gro82] M. Gromov, *Volume and bounded cohomology*. Inst. Hautes Études Sci. Publ. Math. (1982), no. 56, 5–99.

[Gro87] ______, *Hyperbolic groups*. In: Essays in group theory, Math. Sci. Res. Inst. Publ., 8, Springer, New York, 1987. http://dx.doi.org/10.1007/978-1-4613-9586-7_3

[GS14] D. Gruber and A. Sisto, *Infinitely presented graphical small cancellation groups are acylindrically hyperbolic*. Ann. Inst. Fourier., to appear. arxiv:1408.4488

[HO13] M. Hull and D. Osin, *Induced quasicocycles on groups with hyperbolically embedded subgroups*. Algebr. Geom. Topol. **13**(2013), no. 5, 2635–2665. http://dx.doi.org/10.2140/agt.2013.13.2635

[LS77] R. Lyndon and P. Schupp, *Combinatorial group theory*. Springer-Verlag, Berlin-New York, 1977.

[RS95] E. Rips and Z. Sela, *Canonical representatives and equations in hyperbolic groups*. Invent. Math. **120**(1995), no. 3, 489–512. http://dx.doi.org/10.1007/BF01241140

[Sel92] Z. Sela, *Uniform embeddings of hyperbolic groups in Hilbert spaces*. Israel J. Math. **80**(1992), no. 1–2, 171–181. http://dx.doi.org/10.1007/BF02808160

University of Vienna, Faculty of Mathematics, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria
Email: goulnara.arzhantseva@univie.ac.at

Mathematical Institute, AWB, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom
Email: drutu@maths.ox.ac.uk