Groups Geom. Dyn. (Online first)
DOI 10.4171/GGD/822

Published by EMS Press

Translation lengths in crossing and contact graphs of quasi-median graphs

Anthony Genevois

Abstract. Given a quasi-median graph X, the crossing graph ΔX and the contact graph ΓX are natural hyperbolic models of X. In this article, we show that the asymptotic translation length in ΔX or ΓX of an isometry of X is always rational. Moreover, if X is hyperbolic, these rational numbers can be written with denominators bounded above uniformly; this is not true in full generality. Finally, we show that, if the quasi-median graph X is constructible in some sense, then there exists an algorithm computing the translation length of every computable isometry. Our results encompass contact graphs in CAT(0) cube complexes and extension graphs of right-angled Artin groups.

Contents

1. Introduction

Given a group G and a finite generating set S, the *translation number* of an element $g \in G$ is the limit

$$\tau_S(g) := \lim_{n \to +\infty} \frac{\|g\|_S}{n},$$

where $\|\cdot\|_S$ denotes the word length associated with S. The limit exists by subadditivity, and it depends heavily on the chosen generating set. Introduced in [27] with the purpose of studying subgroups in biautomatic groups, the structure of the *translation spectrum*

Mathematics Subject Classification 2020: 20F65 (primary); 20F10, 20F67 (secondary).
Keywords: quasi-median graphs, crossing graph, contact graph, translation length, graph products of groups.

$\mathrm{TSpec}_S(G) := \{\tau_S(g) \mid g \in G \text{ infinite order}\}$ has been investigated for many families of groups, including hyperbolic and CAT(0) groups [12, 14, 29, 44, 45, 50], solvable groups [9–11], small cancellation groups [3, 34], Artin groups [6], mapping class groups [16], outer automorphism groups of free groups [2], Garside groups [38–40], Coxeter groups [43], and hierarchically hyperbolic groups [1].

Some of these results can be better understood in a more geometric framework. Namely, given a metric space (X, d) and an isometry g, the (*asymptotic* or *stable*) *translation length* of g is the limit:

$$\tau(g) := \lim_{n\to+\infty} \frac{d(o, g^n \cdot o)}{n},$$

where $o \in X$ is a basepoint. The limit exists (by subadditivity) and does not depend on the chosen basepoint (by triangular inequality). Given a finitely generated group acting on one of its Cayley graphs by left-multiplication, the translation length of an element coincides with the corresponding translation number. Now, the question becomes: Given our metric space X, what is the structure of the spectrum $\mathrm{TSpect}(X) := \{\tau(g) \mid g \in \mathrm{Isom}(X) \text{ unbounded orbits}\}$?

In several families of graphs, there exists a dichotomy in the possible behaviour of an isometry: Either it has bounded orbits or one of its powers acts as a translation on a bi-infinite geodesic. One consequence is that the translation spectrum is contained in the rational numbers $\mathbb{Q}$, or even in $\frac{1}{N}\mathbb{Z}$ for some integer $N \geq 1$ if one has a uniform control on the powers previously mentioned. Examples include locally finite hyperbolic graphs [14, 29, 45, 50]; median graphs, or equivalently one-skeleta of CAT(0) cube complexes [32]; quasi-median graphs with finite cliques (see Proposition 2.15); bridged graphs, or equivalently one-skeleta of systolic complexes [15]; and Helly graphs of finite combinatorial dimension [30].

Instead of considering all the isometries of a metric space, one can focus on a specific subgroup of isometries. One case of interest, which received a lot of attention, is given by mapping class groups acting on their curve graphs. Among the results available in the vast literature on the subject, let us mention that having a positive translation length characterises pseudo-Anosov elements [41]; the translation length is always a rational number, with a uniform control on the denominator [7,52]; and translation lengths can be computed algorithmically [37, 48, 52].

In order to emphasise the similarity between the *extension graph* introduced in [35] for right-angled Artin groups and the curve graph for mapping class groups, already motivated in [36], the recent work [4] initiates the study of translation spectra of right-angled Artin groups acting on their extension graphs. Their main result shows that, given a finite connected graph Γ, the translation length of an element of the right-angled Artin group $A(\Gamma)$ on the extension graph Γ^e is always rational. Moreover, if Γ has girth ≥ 6, then there is a uniform bound on the denominator.

In this article, our goal is to propose a more geometric and more general perspective on this result.

Our geometric framework is given by *quasi-median graphs*. Originated from the investigation in metric graph theory of graphs similar to median graphs (also known as one-skeleta of CAT(0) cube complexes), quasi-median graphs can be thought of as the "most median-like" family of graphs including median graphs and products of complete graphs. More formally, in the same way that median graphs can be defined as retracts of hypercubes, quasi-median graphs can be defined as retracts of Hamming graphs, (i.e., products of complete graphs). Quasi-median graphs have been introduced in geometric group theory by [17], and turn out to have applications in graph products of groups [17, 22, 23, 51] and their automorphism groups [21, 24], in Thompson's groups [17, 18], and in wreath products of groups [17, 25, 26].

As noticed in [17], the Cayley graph $\mathrm{QM}(\Gamma) := \mathrm{Cayl}(A(\Gamma), \bigcup_{v\in\Gamma}\langle v\rangle)$ of a right-angled Artin group $A(\Gamma)$ turns out to be a quasi-median graph. As such, its geometry is encoded in the combinatorics of its *hyperplanes*. Interestingly, the extension graph Γ^e, defined algebraically in [35], coincides with the *crossing graph* of $\mathrm{QM}(\Gamma)$, namely the graph whose vertices are the hyperplanes of $\mathrm{QM}(\Gamma)$ and whose edges connect two hyperplanes whenever they are transverse. Thus, right-angled Artin groups acting on their extension graphs can be thought of as a particular case of the more general study of crossing graphs of quasi-median graphs.

Crossing graphs have been initially introduced for median graphs (also known as one-skeleta of CAT(0) cube complexes) independently in [31, 46]. A related graph introduced in [31] is the *contact graph*. Given a (quasi-)median graph X, its contact graph is defined as the graph whose vertices are the hyperplanes of X and whose edges connect two hyperplanes whenever they are *in contact* (i.e., transverse or tangent). In the article, we denote by ΔX the crossing graph and by ΓX the contact graph.

Quasi-median graphs and their crossing and contact graphs provide our general geometric framework. Thus, the main question we are interested in is: Given an isometry of a quasi-median graph, what can be said about its translation length in the corresponding crossing and contact graphs? Our first main result is the following axis theorem.

Theorem 1.1. *Let X be a quasi-median graph, and let ΩX be the crossing or contact graph of X. In the former case, we assume that X has no cut-vertex, and, in any case, we assume that a vertex of X belongs to $\leq N$ cliques for some fixed $N \geq 1$. For every $g \in \mathrm{Isom}(X)$ with unbounded orbits in ΩX, there exists some $k \geq 1$ such that g^k admits an axis, that is it acts as a translation on some bi-infinite geodesic in ΩX.*

As already mentioned, such a dichotomy implies immediately that translation lengths must be rational numbers.

Corollary 1.2. *Let X be a quasi-median graph, and let ΩX be the crossing or contact graph of X. In the former case, we assume that X has no cut-vertex, and, in any case, we assume that a vertex of X belongs to $\leq N$ cliques for some fixed $N \geq 1$. For every $g \in \mathrm{Isom}(X)$, the translation length of g in ΩX is rational.*

Regarding Theorem 1.1, it is natural to ask whether there is a uniform upper bound on powers we can take, which would imply that the rational numbers from Corollary 1.2 can be written with uniformly bounded denominators. This turns out to be true when the underlying quasi-median graph is hyperbolic (see Corollary 4.3). However, this is not true in full generality. In Section 4.3, we show that [49] provides the example of a group acting geometrically on a median graph that contains elements with arbitrarily small translation lengths in the crossing graph.

As another natural question, are the translation lengths in crossing and contact graphs algorithmically computable? In order to answer this question, we introduce *constructible* quasi-median graphs and *computable* isometries (see Section 5), and we prove the second main result of this article.

Theorem 1.3. *Let X be a constructible quasi-median graph. There exists an algorithm that computes, given a computable isometry $g \in \mathrm{Isom}(X)$, the asymptotic translation length of g in ΩX.*

Theorems 1.1 and 1.3 are already new for median graphs (a.k.a. one-skeleta of CAT(0) cube complexes), providing a vast range of applications. As already mentioned, including quasi-median graphs also allows us to deduce information about right-angled Artin groups and their extension graphs. In fact, it is more natural to deal with *graph products of groups*. Given a graph Γ and a collection of groups $\mathcal{G} := \{G_u \mid u \in V(\Gamma)\}$ indexed by its vertex-set $V(\Gamma)$, the graph product $\Gamma\mathcal{G}$ is

$$\langle G_u\ (u \in V(\Gamma)) \mid [G_u, G_v] = 1\ (\{u, v\} \in E(\Gamma))\rangle,$$

where $E(\Gamma)$ denotes the edge-set of Γ and where $[G_u, G_v] = 1$ is a shorthand for $[g, h] = 1$ for all $g \in G_u$, $h \in G_v$. Usually, one says that graph products interpolate between free products (when Γ has no edge) and direct sums (when Γ is a complete graph). For instance, right-angled Artin groups coincide with graph products of infinite cyclic groups and right-angled Coxeter groups coincide with graph products of cyclic groups of order two. As shown in [17], the Cayley graph

$$\mathrm{QM}(\Gamma, \mathcal{G}) := \mathrm{Cayl}\Big(\Gamma\mathcal{G}, \bigcup_{u \in V(\Gamma)} G_u\Big)$$

is a median graph. We denote by $\Delta(\Gamma, \mathcal{G})$ its crossing graph. This is the natural generalisation of extension graphs for right-angled Artin groups. Algebraically, it can be defined as the graph whose vertices are the conjugates of vertex-groups and whose edges connect two subgroups whenever they commute (in the sense that every element of one subgroup commutes with every element of the other).

As an application of our general study of crossing graphs in quasi-median graphs, we get the following statement.

Theorem 1.4. *Let Γ be a finite connected graph and $\mathcal{G}$ a collection of groups indexed by $V(\Gamma)$. The following statements hold.*

- *For every $g \in \Gamma\mathcal{G}$, the translation length of g in $\Delta(\Gamma, \mathcal{G})$ is rational.*
- *Moreover, if Γ has no induced 4-cycle, then this rational number can be written with a denominator $\leq |V(\Gamma)|^{40 \cdot \mathrm{clique}(\Gamma)}$.*
- *If the groups in $\mathcal{G}$ have solvable word problems, then there exists an algorithm that computes, given a $g \in \Gamma\mathcal{G}$, the translation length of g in $\Delta(\Gamma, \mathcal{G})$.*

Theorem 1.4 improves the main result of [4] in several ways. First, it does not only apply to right-angled Artin groups but to arbitrary graph products, including for instance right-angled Coxeter groups. Next, even when restricted to right-angled Artin groups, Theorem 1.4 weakens the conditions required to control the denominators of the rational translation lengths: We replace the condition of having girth ≥ 6 with the condition of having no induced 4-cycle. Finally, the (explicit) algorithm we describe offers the possibility to investigate computably the structure of translation spectra of right-angled Artin groups acting on their extension graphs, which remains poorly understood. For instance, we do not if such a spectrum can contain a value less than 2, or when only integer values are taken.

However, it is worth mentioning that the upper bound given by the second item of Theorem 1.4 is far from being optimal in general (which is not surprising since it is obtained from an argument applying to arbitrary quasi-median graphs). For instance, it is proved in [4] that the upper bound for a right-angled Artin group defined by a graph Γ of girth ≥ 6 can be taken as the maximal degree of a vertex in Γ, which is optimal in some cases. This upper bound can be reproved geometrically in our framework, but the global picture remains unclear.

2. Quasi-median geometry

2.1. Generalities

There exist several equivalent definitions of quasi-median graphs; see, for instance, [5]. As mentioned in the introduction, quasi-median graphs can be defined as retracts of Hamming graphs (i.e., products of complete graphs). However, this definition, despite being conceptually simple, is difficult to use in practice. Below is the definition used in [17]. It sounds more technical, but it is easier to use in practice.

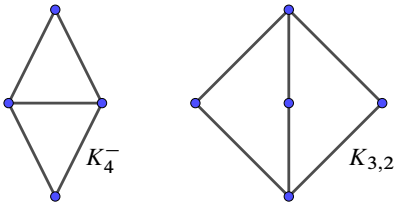

Recall that the graph $K_{3,2}$ is the bipartite complete graph, corresponding to two squares glued along two adjacent edges, and K_4^- is the complete graph on four vertices minus an edge, corresponding to two triangles glued along an edge.

Definition 2.1. A connected graph X is *quasi-median* if it does not contain K_4^- and $K_{3,2}$ as induced subgraphs, and if it satisfies the following two conditions:

(Triangle condition) For every vertex $a, x, y \in X$, if x and y are adjacent and if $d(a,x) = d(a,y)$, then there exists a vertex $z \in X$ which is adjacent to both x and y and which satisfies $d(a,z) = d(a,x) - 1$;

(Quadrangle condition) For every vertex $a, x, y, z \in X$, if z is adjacent to both x and y and if $d(a,x) = d(a,y) = d(a,z) - 1$, then there exists a vertex $w \in X$ which is adjacent to both x and y and which satisfies $d(a,w) = d(a,z) - 2$.

The triangle and quadrangle conditions are illustrated by Figure 1.

Recall that a *clique* is a maximal complete subgraph. According to [5], cliques in quasi-median graphs satisfy a strong convexity condition; namely, they are *gated* in the following sense.

Definition 2.2. Let X be a graph and $Y \subset X$ a subgraph. A vertex $y \in Y$ is a *gate* of an other vertex $x \in X$ if, for every $z \in Y$, there exists a geodesic between x and z passing through y. If every vertex of X admits a gate in Y, then Y is *gated*.

The fact quasi-median graphs do not contain induced copy of K_4^- implies that the intersection between two distinct cliques is always either empty or reduced to a single vertex; in particular, an edge belongs to a unique clique.

A *prism* is a subgraph which is a product of cliques. For instance, in the quasi-median graph given by Figure 2, there are nine maximal prisms: one edge, five "squares" (i.e., 4-cycles K_2^2), one "3-cube" (i.e., a product K_2^3), and two products $K_3 \times K_2$. We emphasise that, despite the fact that a square in our 3-cube also yields a prism (but no longer maximal), a 4-cycle $K_2 \times K_2$ in one of our $K_3 \times K_2$ cannot be a prism because it has a factor that is not a clique. In the same way that median graphs can be naturally thought of as made of cubes, quasi-median graphs can be thought of as made of prisms. In a quasi-median graph, the maximal number of factors of a prism is referred to as its *cubical dimension*.

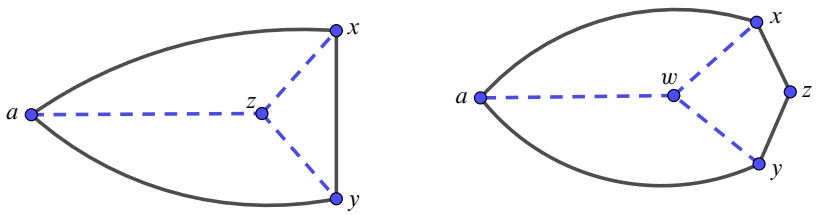

Figure 1. Triangle and quadrangle conditions.

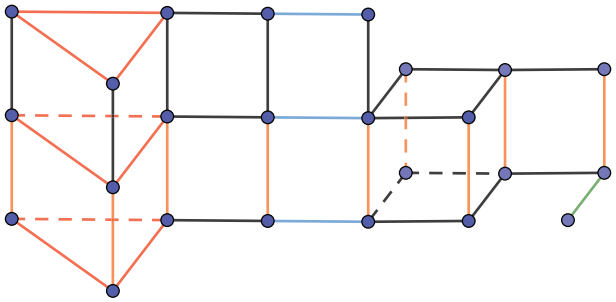

Figure 2. A quasi-median graph and some of its hyperplanes. The orange hyperplane is transverse to the red and blue hyperplanes. The green and orange hyperplanes are tangent. The red and blue hyperplanes are neither transverse nor tangent.

Hyperplanes. A fundamental tool in the study of quasi-median graphs is given by *hyperplanes*.

Definition 2.3. Let X be a graph. A *hyperplane* J is an equivalence class of edges with respect to the transitive closure of the relation saying that two edges are equivalent whenever they belong to a common triangle or are opposite sides of a square. We denote by $X\backslash\backslash J$ the graph obtained from X by removing the interiors of all the edges of J. A connected component of $X\backslash\backslash J$ is a *sector*. The *carrier* of J, denoted by $N(J)$, is the subgraph generated by all the edges of J. Two hyperplanes J_1 and J_2 are *transverse* if there exist two edges $e_1 \subset J_1$ and $e_2 \subset J_2$ spanning a 4-cycle in X, and they are *tangent* if they are not transverse but $N(J_1) \cap N(J_2) \neq \emptyset$.

See Figure 2 for examples of hyperplanes in a quasi-median graph.

The key point is that the geometry of a quasi-median graph reduces to the combinatorics of its hyperplanes. This idea is motivated by the following statement.

Theorem 2.4 ([17]). *Let X be a quasi-median graph.*

(i) *Every hyperplane J separates X, that is, $X\backslash\backslash J$ contains at least two connected components.*

(ii) *Carriers and sectors are gated subgraphs.*

(iii) *A path in X is geodesic if and only if it crosses every hyperplane at most once, that is, it contains at most once edge from each hyperplane.*

(iv) *The distance between two vertices coincides with the number of vertices separating them.*

We emphasise that two vertices, and more generally two sets of vertices, are separated by a hyperplane whenever they lie in two distinct sectors. In Section 2.2, we will introduce a weaker notion of separation.

Projections. As usual, a subgraph is *convex* if it contains every geodesic between any two of its vertices. A gated subgraph Y, as defined above, is for instance convex. It is worth noticing that the gate in Y of a vertex x, when it exists, is unique and minimises the distance to x in Y. As a consequence, it may be referred to as the *projection* of x onto Y. Because an intersection of gated subgraphs is gated, we can define the *gated hull* of a subset as the intersection of all the gated subgraphs containing in.

Our next result, which can be found in [17, Lemma 2.34], describes how projections and hyperplanes interact.

Proposition 2.5. *Let X be a quasi-median graph, $x \in X$ a vertex, and $Y \subset X$ a gated subgraph. Every hyperplane separating x from its projection on Y separates x from Y.*

Let us record the following consequence, which will be useful later.

Lemma 2.6. *Let X be a quasi-median graph and $Y, Z \subset X$ two gated subgraphs. For every $y \in Y$, $\mathrm{proj}_Z(y)$ and $\mathrm{proj}_Y(\mathrm{proj}_Z(y))$ minimise the distance between Y and Z. Moreover, a hyperplane separates these two vertices if and only if separates Y and Z.*

Proof. Let p denote the projection of y on Z, and q the projection of p on Z. According to Proposition 2.5, a hyperplane J separating p and q separates p from Y. A fortiori, it separates y and p. Applying Proposition 2.5 again, it follows that J also separates y from Z. Thus, every hyperplane separating p and q separates Y and Z. Conversely, every hyperplane separating Y and Z has to separate p and q. Clearly, the distance between Y and Z is bounded below by the number of hyperplanes separating them, so $d(p, q)$ must be equal to $d(Y, Z)$. ∎

Median triangles. Given a graph X and a triple of vertices (x_1, x_2, x_3), the triple (a_1, a_2, a_3) is a *median triangle* if it satisfies

$$d(x_i, x_j) = d(x_i, a_i) + d(a_i, a_j) + d(a_j, x_j) \quad \text{for all } i \neq j$$

and if it minimises the quantity $d(a_1, a_2) + d(a_2, a_3) + d(a_1, a_3)$ under this condition. In median graphs, median triangles coincide with median points.

Proposition 2.7. *In a quasi-median graph, every triple of vertices admits one and only one median triangle. Moreover, a hyperplane crossing this median triangle pairwise separates its vertices, which implies that the gated hull of the median triangle is a prism.*

See [17, Proposition 2.84 and Fact 2.90].

2.2. Convex subgraphs

Recall that a subgroup is *convex* if it contains all the geodesics between any two of its vertices. Because an intersection of convex subgraphs is again convex, we can define the

convex hull of a subset as the intersection of all the convex subgraphs containing it. The convex hull of two vertices a and b is referred to as the *interval between a and b*, denoted by $I(a,b)$. Alternatively, it is the union of all the geodesics connecting a and b. According to [17, Corollary 2.107], the convex hull of a finite subset is always finite, so as a particular case.

Lemma 2.8. *In a quasi-median graph, the interval between two vertices is always finite.*

In a quasi-median graph, a *multisector* is the subgraph induced by a union of sectors delimited by a given hyperplane. For instance, the complement of a single sector, referred to as a *cosector*, is a multisector. According to [17, Proposition 2.104], convex hulls can be characterised in terms of multisectors.

Proposition 2.9. *Let X be a quasi-median graph. The convex hull of a subset $S \subset X$ coincides with the intersection of all the multisectors containing S.*

Corollary 2.10. *Let X be a quasi-median graph and $S \subset X$ a set of vertices. Every sector intersecting the convex hull of S intersects S itself.*

Proof. If a sector does not intersect S, then its complement contains the convex hull of S according to Proposition 2.9, proving that our sector actually did not intersect the convex hull of S. ∎

It follows from Lemma 2.6 that, in a quasi-median graph, two disjoint gated subgraphs are separated by at least one hyperplane, in the sense that the two subgraphs lie into two distinct sectors. This is no longer true for convex subgraphs, but a weaker version of this property still holds. One says that two subsets in a quasi-median graph are *weakly separated* by a hyperplane if they lie in two disjoint multisectors. Then the following holds true.

Proposition 2.11. *Let X be a quasi-median graph and $A, B \subset X$ two convex subgraphs. If $A \cap B = \emptyset$ then there exists a hyperplane weakly separating A and B.*

As a first observation towards the proof of the proposition, notice that, even though gated subgraphs are known to satisfy the Helly property (see, for instance, [17, Propositon 2.8]), it clearly fails for convex subgraphs. (For instance, the three edges of a 3-cycle pairwise intersect but they do not globally intersect.) Nevertheless, the following holds true.

Lemma 2.12. *Let X be a quasi-median graph, let $Y \subset X$ be a convex subgraph, and let $Z_1, \dots, Z_n \subset X$ be gated subgraphs. If $Y, Z_1, \dots, Z_n$ pairwise intersect, then their total intersect is non-empty.*

Proof. It suffices to prove the lemma for $n = 2$, the general case following easily by induction or by the Helly property for gated subgraphs. Fix three vertices $a \in Y \cap Z_1$, $b \in Y \cap Z_2$, and $c \in Z_1 \cap Z_2$. Let (x, y, z) be the median triangle of (a, b, c). Because Z_1 is gated, we have $a, x, y, z, c \in Z_1$; because Z_2 is gated, we have $b, x, y, z, c \in Z_2$; and because Y is convex, we have $a, x, z, c \in Y$. Thus, $x, y \in Y \cap Z_1 \cap Z_2$ proving that $Y \cap Z_1 \cap Z_2$ is non-empty. ∎

Our next preliminary lemma is a weakened version of Proposition 2.5 satisfied by convex subgraphs.

Lemma 2.13. *Let X be a quasi-median graph, $Y \subset X$ a convex subgraph, $x \in X$ a vertex, and $M \subset Y$ the set of the vertices of Y minimising the distance to x. For every $z \in M$, a hyperplane separating x from z either separates x from Y or crosses M. Moreover, any two hyperplanes crossing M are transverse.*

Proof. Let J be a hyperplane separating x from some $z \in M$. If J does not cross Y, there is nothing to prove, so we assume that J does cross Y. Consequently, there exists some $y \in Y$ such that J separates z and y. Let (a, b, c) be the median triangle of (x, y, z). Because Y is convex, we have $y, b, c, z \in Y$. Necessarily $z = c$. Because J separates both z from y and z from x, necessarily J crosses the median triangle. A fortiori, it separates $c = z$ and b. But $d(x, b) = d(x, c) = d(x, Y)$, so b belongs to M. We conclude that J crosses M, as desired.

Let J and H be any two hyperplanes crossing M. Fix two vertices $y, z \in M$ separated by both J and H. Let (a, b, c) be the median triangle of (x, y, z). Because Y is convex, we have $y, b, c, z \in Y$. Necessarily, $y = b$ and $z = c$. Thus, J and H cross the median triangle, which implies that they must be transverse according to Proposition 2.7. ∎

Proof of Proposition 2.11. If there exists a hyperplane (strongly) separating A and B, then there is nothing to prove, so we assume that no hyperplane (strongly) separates A and B. Fix two vertices $a \in A$ and $b \in B$ minimising the distance between A and B, and let Z_0 denote the gated hull of $\{a, b\}$. The hyperplanes crossing Z_0 coincide with the hyperplanes separating a and b, say $J_1, \dots, J_n$. For every $1 \leq i \leq n$, assume that J_i delimits some sector Z_i that intersects both A and B.

Observe that $Z_1, \dots, Z_n$ pairwise intersect. Indeed, it follows from Lemma 2.13 that a hyperplane separating a and b belongs to one of the following families:

- the hyperplanes separating a from B and crossing the set M of the vertices of A minimising the distance to b;
- the hyperplanes separating b from A and crossing the set N of the vertices of B minimising the distance to a;
- the hyperplanes crossing both M and N.

Thus, each Z_i contains B and separates M, or contains A and separates N, or separates both M and N. Any two sectors of the first two types intersect every sector of any type.

And, because we also know from Lemma 2.13 that two hyperplanes crossing both M or both N are transverse, it follows that two sectors of the third type also intersect.

We are now in good position to apply Lemma 2.12 and conclude that the intersections $A \cap Z_0 \cap Z_1 \cap \cdots \cap Z_n$ and $B \cap Z_0 \cap Z_1 \cap \cdots \cap Z_n$ are both non-empty. But, because $Z_1, \dots, Z_n$ are sectors delimited by the hyperplanes $J_1, \dots, J_n$, which exhaust all the hyperplanes of Z_0, the intersection $Z_0 \cap Z_1 \cap \cdots \cap Z_n$ must be reduced to a single vertex. This vertex must then belong to both A and B, contradicting the assumption $A \cap B = \emptyset$.

Thus, we have proved that there exists a hyperplane separating a and b that weakly separates A and B. ∎

We conclude this subsection with a last easy observation, which we record for future use.

Lemma 2.14. *Let X be a quasi-median graph, $Y \subset X$ a convex subgraph, and $Z \subset X$ a gated subgraph. If $Y \cap Z \neq \emptyset$, then the projection on Z of a vertex in Y belongs to $Y \cap Z$.*

Proof. Fix two vertices $y \in Y$ and $p \in Y \cap Z$. Let z denote the projection of y on Z and let (ab, c) be the median triangle of (y, z, p). Because z has to belong to $I(y, p)$, necessarily $z = a \in I(y, p) \subset Y$. ∎

2.3. Loxodromic isometries

As a consequence of [32], an isometry of a median graph, as soon as it has unbounded orbits and has no power that is an inversion, admits an *axis*, that is, a bi-infinite geodesic on which it acts as a translation. This section is dedicated to an analogous statement for quasi-median graphs. In this broader context, an *inversion* is an isometry that stabilises a hyperplane and permutes non-trivially its sectors.

Proposition 2.15. *Let X be a quasi-median graph and $g \in \mathrm{Isom}(X)$ an isometry. If $\langle g \rangle$ acts on X without inversions and with unbounded orbits, then g admits an axis in X. More precisely, for every $x \in \mathrm{Min}(g)$ and for every geodesic $[x, gx]$ between x and gx, the concatenation of the $g^k[x, gx]$, $k \in \mathbb{Z}$, defines an axis of g.*

Here, $\mathrm{Min}(g) := \{x \in X \mid d(x, gx) = \min\{d(z, gz) \mid z \in X\}\}$ is the *minimising set* of g. An isometry admitting an axis is referred to as a *loxodromic isometry*. In fact, the second assertion of our proposition follows from the first one according to our next general observation.

Lemma 2.16. *Let X be a graph and $g \in \mathrm{Isom}(X)$ an isometry. Assume that there exists a bi-infinite geodesic γ on which g acts as a translation of length ℓ. Then $d(x, gx) \leq \ell$ for every $x \in X$. Moreover, if $x \in X$ is vertex satisfying $d(x, gx) = \ell$, then x belongs to an axis of g.*

Proof. Fix two vertices $x \in X$ and $y \in \gamma$. Then, for every $n \geq 1$, we have

$$n\ell - 2d(x,y) = d(y,g^n y) - 2d(x,y) \leq d(x,g^n x) \leq n d(x,gx),$$

hence $\ell - 2d(x,y)/n \leq d(x,gx)$. When $n \to +\infty$, we find $\ell \leq d(x,gx)$. This proves the first assertion of our lemma.

Now, fix a vertex $x \in X$ satisfying $d(x,gx) = \ell$. In order to show that x belongs to an axis of g, we need to prove that the $\langle g \rangle$-translates of x all lie on a common bi-infinite geodesic. It suffices to show that $d(x,g^n x) = \sum_{i=0}^{n-1} d(g^i x, g^{i+1} x)$ for every $n \geq 1$. Observe that

$$n\ell \leq d(x,g^n x) \leq n d(x,gx) = n\ell,$$

where the first inequality is obtained by applying our previous assertion to g^n. So

$$d(x,g^n x) = n\ell = \sum_{i=0}^{n-1} d(g^i x, g^{i+1} x),$$

as desired. ∎

Our strategy to prove Proposition 2.15 differs from [32] (thus providing an alternative proof of the main result of [32]). The idea is the following. Starting from an arbitrary vertex x in our quasi-median graph X, the orbit $\langle g \rangle \cdot x$ under our isometry g has to stay at bounded distance from the axis we are looking for. Loosely speaking, $\langle g \rangle \cdot x$ coincides with this axis up to some noise we have to remove. In order to smoothen the quasi-line $\langle g \rangle \cdot x$, we fix a large integer $N \geq 1$, define a reasonable centre c_k of $\{g^{k+i} x \mid -N \leq k \leq N\}$ for every $k \in \mathbb{Z}$, and show that the c_k all lie on a common bi-infinite geodesic.

Proof of Proposition 2.15. Given a vertex $x \in X$, the action of $\langle g \rangle$ on the gated hull of the orbit $\langle g \rangle x$ has only finitely many orbits of hyperplanes. Indeed, every hyperplane separating two vertices in the orbit $\langle g \rangle x$ admits a $\langle g \rangle$-translate separating x and $d(x,gx)$. Fix a vertex $x \in x$ such that the action of $\langle g \rangle$ on the gated hull of $\langle g \rangle x$ has the smallest possible number of orbits of hyperplanes.

Claim 1. For every hyperplane J, the set $\{k \in \mathbb{Z} \mid J \text{ separates } g^k x \text{ and } g^{k+1} x\}$ is finite.

Let $\mathcal{J}$ denote the set of the hyperplanes J separating $g^k x$ and $g^{k+1} x$ for infinitely many $k \in \mathbb{Z}$.

First, observe that $\mathcal{J}$ is finite. Indeed, it is clear that $\mathcal{J}$ is $\langle g \rangle$-invariant. And, because every hyperplane in $\mathcal{J}$ has a $\langle g \rangle$-translate separating x and gx while exactly $d(x,hx)$ hyperplanes separate x and gx, $\mathcal{J}$ contains only finitely many $\langle g \rangle$-orbits and they all have sizes $\leq d(x,gx)$.

Fix a hyperplane $J \in \mathcal{J}$ and a sector S delimited by J. If the $g^k S$, $k \in \mathbb{Z}$, pairwise intersect, then $\bigcap_{k \in \mathbb{Z}} g^k S$ yields a $\langle g \rangle$-invariant gated subgraph not crossed by any translate of J. This contradicts our choice of x. So there exists some $h_1 \in \langle g \rangle$ such that

$h_1 S \cap S = \emptyset$. Of course, $h_1 J$ and J cannot be transverse, but they also have to be distinct because h_1 is not an inversion. Let S' denote the sector delimited by J and containing $h_1 J$. Because $\mathcal{J}$ is $\langle g \rangle$-invariant and finite, we cannot have $h_1 S' \subsetneq S'$, so $h_1 S'$ must be the sector delimited by $h_1 J$ containing J. As before, if the $g^k S$, $k \in \mathbb{Z}$, pairwise intersect, we get a contradiction; so there must exist some $h_2 \in \langle g \rangle$ satisfying $h_2 S' \cap S' = \emptyset$. But then $h_2 S' \subset (S')^c \subsetneq h_1 S'$, hence $h_1^{-1} h_2 S' \subsetneq S'$. Again, this contradicts the fact that $\mathcal{J}$ is finite.

Thus, the only possibility is that $\mathcal{J}$ is empty, concluding the proof of Claim 1.

From now on, we fix a geodesic $[x, gx]$ between x and gx, and we denote by γ the concatenation of the $\langle g \rangle$-translates of $[x, gx]$. We enumerate the edges of γ as $\dots, e_{-1}, e_0, e_1, \dots$ such that e_i and e_{i+1} have a common endpoint for every $i \in \mathbb{Z}$. Claim 1 allows us to define, for every hyperplane J crossing $\langle g \rangle x$, the numbers

$$\ell(J) := \min\{k \in \mathbb{Z} \mid e_k \in J\} \quad \text{and} \quad r(J) := \max\{k \in \mathbb{Z} \mid e_k \in J\}.$$

In other words, $\ell(J)$ (resp. $r(J)$) indicates the leftmost (resp. rightmost) part of $\langle g \rangle x$ crossed by J. Notice that $\ell(gJ) = \ell(J) + 1$ and $r(gJ) = r(J) + 1$. As a consequence, the difference $r(J) - \ell(J)$ depends only on the $\langle g \rangle$-orbit of J. Because there are only finitely many orbits of hyperplanes crossing $\langle g \rangle x$, we can fix an integer $N \geq 1$ larger than any of these quantities.

We distinguish two types of hyperplanes crossing $\langle g \rangle x$. Such a hyperplane J is *one-sided* if all the e_k for $k \leq \ell(J)$ and $k \geq r(J)$ belong to the same sector delimited by J, which we denote by $S(J)$. Otherwise, J is *two-sided* and we denote by $L(J)$ (resp. $R(J)$) the sector delimited by J containing the e_k for $k \leq \ell(J)$ (resp. $k \geq r(J)$).

For every $k \in \mathbb{Z}$, set $B_k := \bigcup_{-N \leq i \leq N} e_{k+i}$. Notice that $gB_k = B_{k+1}$. Given a $k \in \mathbb{Z}$, we want to define a centre of B_k. For this purpose, given a hyperplane J, we define the sector $S_k(J)$ delimited by J as follows:

- if J is one-sided, $S_k(J) := S(J)$;
- if J is two-sided and $\ell(J) \geq k$, $S_k(J) := L(J)$;
- if J is two-sided and $\ell(J) < k$, $S_k(J) := R(J)$.

Roughly speaking, $S_k(J)$ represents the sector of J that contains the bigger part of B_k.

Claim 2. The intersection between the gated hull $\mathrm{GH}(B_k)$ of B_k and the $S_k(J)$ for J crossing B_k is reduced to a single vertex, which we denote by c_k.

It suffices to show that $\mathrm{GH}(B_k)$ and the $S_k(J)$ for J crossing B_k pairwise intersect. Then the Helly property for gated subgraphs assures that the total intersection is non-empty, and, because we are choosing one sector for each hyperplane crossing B_k, there cannot be two vertices in the intersection.

If J_1 is a one-sided hyperplane, then $S(J_1)$ clearly intersects $S(J_2)$ if J_2 is another one-sided hyperplane and both $L(J_2)$ and $R(J_2)$ if J_2 is a two-sided hyperplane. Moreover, if J_1 and J_2 are two two-sided hyperplanes, then $R(J_1) \cap R(J_2)$ and $L(J_1) \cap L(J_2)$ are both non-empty. So, given two hyperplanes J_1 and J_2 crossing B_k, the only

remaining possibility in order to have $S_k(J_1) \cap S_k(J_2) = \emptyset$ is that J_1, J_2 are two non-transverse two-sided hyperplanes with $S_k(J_1) = L(J_1)$ and $S_k(J_2) = R(J_2)$ (up to switching J_1 and J_2). Necessarily, from $S_k(J_1) \cap S_k(J_2)$, it follows that $\ell(J_2) < k \leq \ell(J_1)$, and because J_1 and J_2 are not transverse, we know that $r(J_2) < \ell(J_1)$. Therefore, $S_k(J_1) = L(J_1)$ and $S_k(J_2) = R(J_2)$ actually intersect.

Next, if J is a one-sided hyperplane crossing B_k, then B_k cannot lie in the complement of $S(J)$ by definition of N, so $S_k(J) \cap \mathrm{GH}(B_k) \neq \emptyset$. If J is two-sided with $L(J) \cap \mathrm{GH}(B_k) = \emptyset$, then $\{e_i \mid i \leq \ell(J)\} \cap B_k$ must be empty, which implies that $k - N > \ell(J)$. We deduce that $k > r(J)$, $g^k x \in R(J)$, and $S_k(J) = R(J)$. Thus, e_k belongs to $S_k(J) \cap B_k$, proving that $S_k(J)$ and $\mathrm{GH}(B_k)$ intersect. This concludes the proof of Claim 2.

Notice that $gc_k = c_{k+1}$ for every $k \in \mathbb{Z}$. We claim that the c_k all belong to a common bi-infinite geodesic. This proves that, fixing a geodesic $[c_0, c_1]$ between c_0 and $c_1 = gc_0$, the concatenation of the $\langle g \rangle$-translates of $[c_0, c_1]$ defines an axis for g. In order to prove our claim, it suffices to observe that c_k belongs to $S(J)$ for every $k \in \mathbb{Z}$ if J is a one-sided hyperplane, and, if J is two-sided, c_k belongs to $L(J)$ for $k \leq \ell(J)$ and to $R(J)$ for $k > \ell(J)$. Therefore, no hyperplane can separate a c_k from some c_i and c_j with $i < k < j$.

So far, we have proved the first assertion of Proposition 2.15. The second assertion follows from Lemma 2.16. ∎

2.4. Quasiconvex isometries

A loxodromic isometry is *quasiconvex* if it admits a quasiconvex axis. In order to quantify this property, we introduce the two following quantities.

Definition 2.17. Let X be a quasi-median graph and γ a bi-infinite geodesic. Define

- $\mathrm{QC}(\gamma)$ as the largest Hausdorff distance between two bi-infinite geodesics in the convex hull of γ.
- $\mathrm{HQC}(\gamma)$ as the maximal $n \geq 0$ for which there exist two transverse collections of hyperplanes of size n crossing γ.

Thus, γ is quasiconvex if and only if the quantity $\mathrm{QC}(\gamma)$ is finite. The quantity $\mathrm{HQC}(\gamma)$ is a more median-friendly version of $\mathrm{QC}(\gamma)$ which coarsely coincides with $\mathrm{QC}(\gamma)$ if γ is an axis according to Lemma 2.20. Our first observation is that $\mathrm{HQC}(\cdot)$ does not depend on a particular choice of an axis. Indeed, the following lemma holds true.

Lemma 2.18. *Let X be a quasi-median graph and $g \in \mathrm{Isom}(X)$ a loxodromic isometry. Any two axes of g cross exactly the same hyperplanes.*

Proof. Let γ_1, γ_2 be two axes of g. First, we want to show that γ_1 and γ_2 fellow-travel in the following sense.

Claim 3. There exist $C, D \geq 0$ such that, for every $t_1 \in \mathbb{Z}$, there exists some $t_2 \in \mathbb{Z}$ such that $|t_1 - t_2| \leq C$ and $d(\gamma_1(t_1), \gamma_2(t_2)) \leq D$.

Fix two $s_1, s_2 \in \mathbb{Z}$. Set $C := \|g\|$ and $D := \|g\| + d(\gamma_1(s_1), \gamma_2(s_2))$. For every $t_1 \in \mathbb{Z}$, there exists some $r \in \mathbb{Z}$ such that $d(\gamma_1(t_1), g^r\gamma_1(s_1)) \leq \|g\|$. Set $t_2 := s_1 + r\|g\|$. The latter inequality implies that

$$|t_1 - t_2| = d(\gamma_1(t_1), \gamma_1(s_1 + r\|g\|)) \leq \|g\| = C.$$

Moreover, we have

$$\begin{aligned} d(\gamma_1(t_1), \gamma_2(t_2)) &\leq d(\gamma_1(t_1), \gamma_1(s_1 + r\|g\|)) + d(\gamma_1(s_1 + r\|g\|), \gamma_2(t_2)) \\ &\leq d(\gamma_1(t_1), g^r\gamma_1(s_1)) + d(g^r\gamma_1(s_1), g^r\gamma_2(s_2)) \\ &\leq \|g\| + d(\gamma_1(s_1), \gamma_2(s_2)) = D, \end{aligned}$$

which concludes the proof of Claim 3.

Now, let J be a hyperplane crossing γ_1. If J does not cross γ_2, then there exists some $k_0 \in \mathbb{Z}$ such that J separates $\gamma_1(k)$ from γ_2 for every $k \geq k_0$. Given an $s \geq 0$ larger than the Hausdorff distance between γ_1 and γ_2, which is finite as a consequence of Claim 3, we deduce that $J, gJ, \dots, g^sJ$ separate $\gamma_1(k + s\|g\|)$ from γ_2 for every $k \geq k_0$, which is impossible. Thus, every hyperplane crossing γ_1 has to cross γ_2 as well. Symmetrically, every hyperplane crossing γ_2 has to cross γ_1. ∎

Thus, we can safely define the following.

Definition 2.19. Let X be a quasi-median graph and $g \in \mathrm{Isom}(X)$ a loxodromic isometry. Define $\mathrm{HQC}(g)$ as $\mathrm{HQC}(\gamma)$ for an arbitrary axis γ of g.

We conclude this subsection by proving that $\mathrm{HQC}(g)$ coarsely coincides with $\mathcal{QC}(\gamma)$ for any choice of an axis γ of g.

Lemma 2.20. *Let X be a quasi-median graph of finite cubical dimension and γ an axis of some isometry $g \in \mathrm{Isom}(X)$. The inequality*

$$\mathrm{HQC}(g) \leq \mathrm{QC}(\gamma) \leq 2 \cdot \mathrm{HQC}(g)$$

holds.

Proof. For short, set $M := \mathrm{HQC}(g)$. We begin by proving the following observation.

Claim 4. Let $\alpha, \beta \subset \mathrm{QC}(\gamma)$ be two bi-infinite geodesics. Assume that every hyperplane crossing γ also crosses β. Then α lies in the $2M$-neighbourhood of β.

Given two vertices $p \in \alpha$ and $q \in \beta$, fix two vertices $q^-, q^+ \in \beta$ such that all the hyperplanes separating p and q (which must cross γ, and a fortiori β) cross β between q^- and q^+. Observe that p belongs to $I(q^-, q^+)$, because a hyperplane separating p from $\{q^-, q^+\}$ would have to cross β elsewhere than between q^- and q^+.

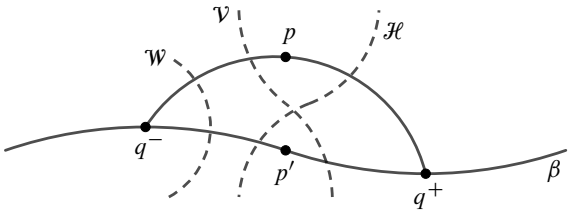

Let p' denote the vertex of β between q^- and q^+ satisfying $d(q^-, p) = d(q^+, p')$. Let $\mathcal{H}$ denote the set of the hyperplanes separating $\{p, q^-\}$ and $\{p', q^+\}$, and $\mathcal{V}$ the set of the hyperplanes separating $\{q^-, p'\}$ and $\{p, q^+\}$.

If $\mathcal{W}$ denotes the set of the hyperplanes separating q^- from $\{p, p', q^+\}$, then

$$|\mathcal{W}| + |\mathcal{H}| = d(q^-, p') = d(q^-, p) = |\mathcal{W}| + |\mathcal{V}|,$$

so $\mathcal{H}$ and $\mathcal{V}$ have the same size. Moreover, $\mathcal{H}$ and $\mathcal{V}$ are clearly transverse. Observe that $\mathcal{H} \cup \mathcal{V}$ coincides with the set of the hyperplanes separating p and p', hence

$$d(p, \beta) \leq d(p, p') = |\mathcal{H}| + |\mathcal{V}| \leq 2M.$$

Thus, we have proved that α lies in the $2M$-neighbourhood of β. Claim 4 is proved.

If $\mathrm{QC}(\gamma)$ is infinite, then applying Claim 4 to $\beta = \gamma$ and α arbitrary shows that M must be infinite as well. A fortiori, the inequality $\mathrm{QC}(\gamma) \leq 2M$ holds.

Now, assume that $\mathrm{QC}(\gamma)$ is finite, which amounts to saying that γ is quasi-dense in its convex hull $\mathrm{CH}(\gamma)$. Given a bi-infinite geodesic ζ in $\mathrm{CH}(\gamma)$, we know from Corollary 2.10 that every hyperplane crossing ζ also crosses γ. It turns out that, because γ is an axis, the converse also holds, that is, every hyperplane crossing γ also crosses ζ. Indeed, if there exists some hyperplane J crossing γ but not ζ, then J separates ζ from some infinite subray $\gamma_0 \subset \gamma$. Up to replacing g with its inverse, we assume that $g\gamma_0 \subset \gamma_0$. Because X has finite cubical dimension, there exists some $1 \leq d \leq \dim(X)$ such that J and $g^d J$ are not transverse. Then, for every $k \geq 1$, the hyperplanes $J, g^d J, \dots, g^{kd} J$ separate ζ from some infinite subray of γ_0. Taking k larger than the Hausdorff distance between γ and ζ yields a contradiction. Thus, we have proved that every bi-infinite geodesic in $\mathrm{CH}(\gamma)$ is crossed exactly by the hyperplanes crossing γ.

Consequently, given two bi-infinite geodesics α, β in $\mathrm{CH}(\gamma)$, we deduce by applying Claim 4 twice that the Hausdorff distance between α and β is at most $2M$. Therefore, $\mathrm{QC}(\gamma) \leq 2M$.

Finally, let $\mathcal{H}, \mathcal{V}$ be two finite transverse collections of hyperplanes of the same size crossing γ. Fix two vertices $a, b \in \gamma$ such that the hyperplanes in $\mathcal{H} \cup \mathcal{V}$ cross γ between a and b. For every $J \in \mathcal{H} \cup \mathcal{V}$, let J^- (resp. J^+) denote the sector delimited by J containing a (resp. b). Fix two arbitrary vertices

$$p \in \bigcap_{J \in \mathcal{H}} J^- \cap \bigcap_{J \in \mathcal{V}} J^+ =: A \quad \text{and} \quad q \in \bigcap_{J \in \mathcal{H}} J^+ \cap \bigcap_{J \in \mathcal{V}} J^- =: B.$$

Let (x, y, z) be the median triangle of (a, b, p). We deduce from the convexity of sectors that z belongs to A. For the same reason, we know that no hyperplane in $\mathcal{H}$ (resp. in $\mathcal{V}$)

can separate x and z (resp. y and z). But the hyperplanes separating x and z coincide with the hyperplanes separating y and z according to Proposition 2.7, so no hyperplane in $\mathcal{H} \cup \mathcal{V}$ can separate x and z, which implies that $x \in I(a,b) \cap D$. Thus, we have proved that there exists a geodesic ζ_0 between a and b passing through A. Let ζ denote the bi-infinite geodesic obtained from γ by replacing the subsegment between a and b with ζ_0. Similarly, there exists a geodesic ξ_0 between a and b passing through B. Let ξ denote the bi-infinite geodesic obtained from γ by replacing the subsegment between a and b with ξ_0. The key observation is that a vertex in $\zeta \cap A$ is separated from a given vertex in ζ either by all the hyperplanes in $\mathcal{H}$ or by all the hyperplanes in $\mathcal{V}$, so the Hausdorff distance between α and β must be at least $|\mathcal{H}| = |\mathcal{V}|$. Thus, we have proved that $\mathrm{QC}(\gamma) \geq M$. ∎

Remark 2.21. It can be proved that, given a loxodromic isometry g of a quasi-median graph X, the quantity $\mathrm{QC}(\gamma)$ does not depend on the axis γ we choose. More precisely, the convex hulls of any two axes turn out to be always isometric. However, this observation will be useful in the sequel as $\mathrm{HQC}(g)$ will be a quantity easier to compute.

3. Crossing and contact graphs

Recall that, given a quasi-median graph X, the *crossing graph* ΔX is the graph whose vertices are the hyperplanes of X and whose edges connect two hyperplanes whenever they are transverse, and the *contact graph* ΓX is the graph whose vertices are the hyperplanes of X and whose edges connect two hyperplanes whenever they are *in contact* (i.e., when their carriers intersect). In this section, we record basic properties of crossing and contact graphs. These properties are well known for median graphs, and can be found in [19, 31]; crossing graphs of some specific quasi-median graphs are also studied in [17, 21]. The arguments used in this section are straightforward adaptations.

3.1. Comparison of the two graphs

Even though crossing and contact graphs have similar definitions, there are a few fundamental differences. The first one is that the contact graph is always connected when the crossing graph may be disconnected. However, the connectedness of the crossing graph can be detected by the existence of cut-vertices. (For median graphs, this observation can be found in [42].)

Lemma 3.1. *The contact graph of a quasi-median graph is always connected.*

Proof. Let X be a quasi-median graph and A, B two hyperplanes. Fix an arbitrary path from $N(A)$ to $N(B)$ and let $J_1, \dots, J_n$ denote the hyperplanes successively crossed. Then the hyperplanes $A, J_1, \dots, J_n, B$ are successively in contact, producing a path from A to B in the contact graph. ∎

Given a quasi-median graph X, its *crossing graph* ΔX is the graph whose vertices are the hyperplanes of X and whose edges link two hyperplanes whenever they are transverse. The crossing graph is not always connected, but we understand the default of connectedness.

Lemma 3.2. *Let X be a quasi-median graph. Its crossing graph is disconnected if and only if there exists a cut-vertex in X.*

Proof. Given a vertex $x \in X$, define its clique-link as the graph whose vertices are the cliques containing x and whose edges connect two vertices if the corresponding cliques span a prism (or equivalently, if the hyperplanes containing the cliques are transverse). We can use clique-links in order to characterise cut-vertices.

Claim 5. A vertex $x \in X$ is a cut-vertex if and only if its clique-link is disconnected.

First, assume that the clique-link of x is connected. Given two vertices $a, b \in X$, fix a path between a and b, say a geodesic $[a, b]$. If x does not belong to $[a, b]$, then x does not separate a and b. Otherwise, let $e_1, \dots, e_n$ be a sequence of edges having x as an endpoint such that the cliques containing $e_1, \dots, e_n$ defines a path in the clique-link of x, such that e_1 is the first edge of $[a, b]$ that belongs to a clique containing x, and such that e_n is the last edge of $[a, b]$ that belongs to a clique containing x. For every $1 \leq i \leq n-1$, the edges e_i and e_{i+1} span a 4-cycle, say $e_i \cup e_{i+1} \cup \varepsilon_{i+1} \cup \varepsilon_i$. Then $\varepsilon_1, \dots, \varepsilon_n$ defines a path connecting the endpoints of e_1 and e_n distinct from x that does not pass through x. Thus, replacing the subsegment of $[a, b]$ between e_1 and e_n (including e_1 and e_n themselves) with $\varepsilon_1 \cup \cdots \cup \varepsilon_n$ produces a path between a and b that does not pass through x. This proves that x is not a cut-vertex.

Conversely, assume that the clique-link of x is disconnected. Let $\varepsilon_1, \varepsilon_2$ be two edges containing x that belong to two cliques lying in distinct connected components of the clique-link of x. Let a_1, a_2 denote the endpoints of $\varepsilon_1, \varepsilon_2$ distinct from x. We claim that a_1 and a_2 are separated by x. The key point is that any path between a_1 and a_2 can be obtained from $\varepsilon_1 \cup \varepsilon_2$ by some elementary operations. More precisely, given an oriented path γ in our graph X, which we decompose as a concatenation of oriented edges $e_1 \cdots e_n$, one says that γ' is obtained from γ by

- *flipping a square*, if there exists some $1 \leq i \leq n-1$ such that
$$\gamma' = e_1 \cdots e_{i-1} \cdot a \cdot b \cdot e_{i+2} \cdots e_n,$$
where e_i, e_{i+1}, b, a define a 4-cycle in X;
- *shortening a triangle*, if there exists some $1 \leq i \leq n-1$ such that
$$\gamma' = e_1 \cdots e_{i-1} \cdot a \cdot e_{i+2} \cdots e_n,$$
where e_i, e_{i+1}, a define a 3-cycle in X;
- *removing a backtrack*, if there exists some $1 \leq i \leq n-1$ such that
$$\gamma' = e_1 \cdots e_{i-1} \cdot e_{i+2} \cdots e_n,$$

where e_{i+1} is the inverse of e_i.

Then, given two paths with the same endpoints, one can be obtained from the other by flipping squares, shortening triangles, removing backtracks, and inverses of these operations. (See, for instance, [23, Lemma 2.2].)

Now, given a path from a_1 to a_2, we colour its edges successively such that the colour changes exactly when we pass through x from one clique to a second clique that lies in a distinct connected component in the clique-link of x. The number of colours turns out to be an invariant: Applying any of our elementary operations to a path does not modify the number of colours. Because $\varepsilon_1 \cup \varepsilon_2$ has two colours, it turns out that there cannot exist a path from a_1 to a_2 that does not pass through x because such a path would have only one colour. Thus, we have proved that x is a cut-vertex, concluding the proof of Claim 5.

We are finally ready to prove our lemma. If X has a cut-vertex x, then two hyperplanes lying in distinct components of $X\backslash\{x\}$ cannot be transverse, so the crossing graph is clearly not connected. Conversely, assume that X has no cut-vertex. Let A, B be two hyperplanes. We fix an arbitrary path from $N(A)$ to $N(B)$, and we denote by $J_1, \dots, J_n$ the hyperplanes successively crossed. Observe that $A, J_1, \dots, J_n, B$ are successively in contact. But the fact that the clique-link of every vertex is connected, which follows from Claim 5, implies that two hyperplanes in contact are connected by a path in the crossing graph. Therefore, there exists path in ΔX connecting A and B. This proves that ΔX is connected. ∎

Even when a quasi-median graph has no cut-vertex, the crossing and contact graphs may be quite different. In fact, according to [31, Proposition 2.19], given any graph, it is possible to construct a median graph whose crossing graph is this particular graph. On the other hand, the contact graph is always a quasi-tree [31]. Nevertheless, the following holds true.

Proposition 3.3. *Let X be a quasi-median graph. Assume that X has no cut-vertex and that every vertex in X belongs to $\leq N$ cliques for some fixed N. Then the canonical map $\Delta X \to \Gamma X$ is a quasi-isometry.*

The same statement for median graphs can be found in [19]. Our proposition follows from Lemmas 3.9 and 3.10, but it will not be used in this paper.

3.2. Hyperbolicity

As already mentioned, the contact graph of a median graph is always a quasi-tree [31]. The same statement, using the same arguments, can be proved for quasi-median graphs. However, this fact will not be necessary in the rest of the paper. Instead, we show that crossing and contact graphs are hyperbolic and compute explicit hyperbolicity constants. We start with the contact graph.

Lemma 3.4. *Let X be a quasi-median graph and A, B two hyperplanes. If $d_{\Gamma X}(A, B) \geq 2$, then every vertex of every geodesic in ΓX between A and B lies at distance ≤ 2 from a hyperplane separating A and B.*

Proof. Let $H_1, \dots, H_m$ be a maximal collection of pairwise non-transverse hyperplanes separating A and B; for convenience, we set $H_0 := A$ and $H_{m+1} := B$. Up to re-indexing our hyperplanes, we assume that H_i separates H_{i-1} and H_{i+1} for every $1 \leq i \leq m$. Notice that $m \geq 1$ because $d_{\Gamma X}(A, B) \geq 2$. Fix a geodesic $J_1, \dots, J_n$ from A to B in ΓX and an index $1 \leq i \leq n$. We claim that J_i lies at distance ≤ 2 from H_j for some $1 \leq j \leq m$.

If J_i is in contact with H_j for some $0 \leq j \leq m+1$, then there exists some $1 \leq k \leq m$ such that $d_{\Gamma X}(J_i, H_k) \leq 2$. From now on, we assume that J_i is not in contact with any H_j. As a consequence, J_i lies between H_j and H_{j+1} for some $0 \leq j \leq m$. Let r be the first index $\ell < i$ such that J_ℓ is in contact with H_j and let s be the last index $\ell > i$ such that J_ℓ is in contact with H_{j+1}. We have

$$d_{\Gamma X}(J_r, J_s) \leq 2 + d_{\Gamma X}(H_j, H_{j+1}) = 3.$$

Consequently, J_k lies at distance $\leq 3/2 + 1$ from H_j or H_{j+1}. ∎

Corollary 3.5. *Let X be a quasi-median graph. Its contact graph ΓX is 3-hyperbolic.*

Proof. Let A, B, C be three hyperplanes. Fix three geodesics $[A, B]$, $[B, C]$, $[A, C]$ in ΓX and a vertex $H \in [A, B]$. If $d_{\Gamma X}(A, B) \leq 1$, then H lies in $[B, C] \cup [A, C]$. Otherwise, we know from Lemma 3.4 that $d_{\Gamma X}(H, J) \leq 2$ for some hyperplane J separating A and B. If J is transverse to C, then A lies at distance ≤ 3 from $[B, C] \cup [A, C]$. Otherwise, J separates B and C or A and C (not exclusively). Necessarily, there exists some hyperplane in $[B, C]$ or $[A, C]$ equal or transverse to J, so H lies at distance ≤ 3 from $[B, C] \cup [A, C]$. Thus, we have proved that $[A, B]$ lies in the 3-neighbourhood of $[B, C] \cup [A, C]$. ∎

Next, we turn to the crossing graph. Recall that, in full generality, the crossing graph may be any graph, and in particular it may not be hyperbolic. Nevertheless, it becomes hyperbolic (and in fact a quasi-tree) under the good assumptions.

Lemma 3.6. *Let X be a quasi-median graph and A, B two hyperplanes. Assume that every X has no cut-vertex and that every vertex belongs to at most N cliques. If $d_{\Delta X}(A, B) \geq N$, then every vertex of every geodesic in ΔX between A and B lies at distance $\leq 2 + N/2$ from a hyperplane separating A and B.*

Proof. Let $H_1, \dots, H_m$ be a maximal collection of pairwise non-transverse hyperplanes separating A and B; for convenience, we set $H_0 := A$ and $H_{m+1} := B$. Up to re-indexing our hyperplanes, we assume that H_i separates H_{i-1} and H_{i+1} for every $1 \leq i \leq m$. Notice that $m \geq 1$ because $d_{\Gamma X}(A, B) \geq N$. (Indeed, two hyperplanes in contact lies at distance $\leq N$ in ΔX.) Fix a geodesic $J_1, \dots, J_n$ from A to B in ΔX and an index $1 \leq i \leq n$. We claim that J_i lies at distance ≤ 2 from H_j for some $1 \leq j \leq m$.

If J_i is equal or transverse to H_j for some $0 \leq j \leq m+1$, then there exists some $1 \leq k \leq m$ such that $d_{\Gamma X}(J_i, H_k) \leq 2$. From now on, we assume that J_i is distinct and not transverse to any H_j. As a consequence, J_i lies between H_j and H_{j+1} for some $0 \leq j \leq m$. Let r be the first index $\ell < i$ such that J_ℓ is transverse to H_j and let s be the last index $\ell > i$ such that J_ℓ is transverse to H_{j+1}. We have

$$d_{\Delta X}(J_r, J_s) \leq 2 + d_{\Delta X}(H_j, H_{j+1}) \leq 2 + N.$$

Consequently, J_k lies at distance $\leq (2+N)/2 + 1$ from H_j or H_{j+1}. ∎

Corollary 3.7. *Let X be a quasi-median graph. Assume that every X has no cut-vertex and that every vertex belongs to at most N cliques. The crossing graph ΔX is $(3 + N/2)$-hyperbolic.*

Proof. Let A, B, C be three hyperplanes. Fix three geodesics $[A, B], [B, C], [A, C]$ in ΔX and a vertex $H \in [A, B]$. If $d_{\Delta X}(A, B) < N$, then H lies in the $(N-1)/2$-neighbourhood of $[B, C] \cup [A, C]$. Otherwise, we know from Lemma 3.6 that $d_{\Gamma X}(H, J) \leq 2 + N/2$ for some hyperplane J separating A and B. If J is transverse to C, then A lies at distance $\leq 3 + N/2$ from $[B, C] \cup [A, C]$. Otherwise, J separates B and C or A and C (not exclusively). Necessarily, there exists some hyperplane in $[B, C]$ or $[A, C]$ equal or transverse to J, so H lies at distance $\leq 3 + N/2$ from $[B, C] \cup [A, C]$. Thus, we have proved that $[A, B]$ lies in the $(3 + N/2)$-neighbourhood of $[B, C] \cup [A, C]$. ∎

3.3. Strongly contracting isometries

Finally, we turn to the question of when a given isometry of a quasi-median graph induces an isometry with unbounded orbits in the corresponding crossing or contact graph. We refer to these isometries as *strongly contracting*.

Proposition 3.8. *Let X be a quasi-median graph and $g \in \mathrm{Isom}(X)$ an isometry admitting an axis γ. Let ΩX denote the crossing or contact graph of X. In the former case, assume that X has no cut-vertex and that every vertex belongs to $\leq N$ cliques. The following assertions are equivalent.*

- *g has unbounded orbits in ΩX.*
- *g is a loxodromic isometry of ΩX.*
- *There exist a power $r \geq 1$ and a hyperplane J crossing γ such that $\{g^{rk} J \mid k \in \mathbb{Z}\}$ is a collection of pairwise strongly separated hyperplanes.*

When these conditions are satisfied, we say that g is strongly separated.

Recall that two hyperplanes are *strongly separated* if no hyperplane can be transverse to both of them. Given two hyperplanes A and B, we denote by $\mathrm{ss}(A, B)$ the maximal number of pairwise strongly separated hyperplanes separating A and B.

Our proposition will be a straightforward consequence of the following estimates of distances in crossing and contact graphs.

Lemma 3.9. *Let X be a quasi-median graph. The inequalities*

$$\mathrm{ss}(A,B) \leq d_{\Gamma X}(A,B) \leq 3(1+\mathrm{ss}(A,B))$$

hold for all hyperplanes A, B.

Lemma 3.10. *Let X be a quasi-median graph with no cut-vertex and all of whose vertices belong to $\leq N$ cliques. The inequalities*

$$\mathrm{ss}(A,B) \leq d_{\Delta X}(A,B) \leq (2+N)(1+\mathrm{ss}(A,B))$$

hold for all hyperplanes A, B.

Proof of Lemmas 3.9 and 3.10. Let $J_1, \dots, J_n$ be a maximal collection of pairwise strongly separated hyperplanes separating A and B. For convenience, set $J_0 := A$ and $J_{n+1} := B$. Given an $0 \leq i \leq n$, observe that the last hyperplane in a maximal collection of pairwise non-transverse hyperplanes separating J_i from J_{i+1} cannot be strongly separated together with J_i (by maximality of our collection $J_1, \dots, J_n$), and it must be in contact with J_{i+1}. Hence, $d_{\Gamma X}(J_i, J_{i+1}) \leq 2+1 = 3$ and $d_{\Delta X}(J_i, J_{i+1}) \leq 2+N$. Therefore,

$$d_{\Gamma X}(A,B) \leq \sum_{i=0}^{n} d_{\Gamma X}(J_i, J_{i+1}) \leq 3(n+1) = 3(1+\mathrm{ss}(A,B))$$

and

$$d_{\Delta X}(A,B) \leq \sum_{i=0}^{n} d_{\Delta X}(J_i, J_{i+1}) \leq (2+N)(n+1) = (2+N)(1+\mathrm{ss}(A,B)).$$

Next, let $H_1, \dots, H_m$ be a geodesic in ΩX ($= \Delta X$ or ΓX) between A and B. For every $1 \leq i \leq n$, there must exist some $1 \leq j \leq m$ such that H_j is equal or transverse to J_i. Because the J_i are pairwise strongly separated, an H_j cannot appear twice, hence $m \geq n$. We conclude that $d_{\Omega X}(A,B) \geq \mathrm{ss}(A,B)$. ∎

Proof of Proposition 3.8. Let J be a hyperplane crossing γ. If g has unbounded orbits, then there exists some $r \geq 1$ such that $d_{\Omega X}(J, g^r J) \geq 3$. Necessarily, J and $g^r J$ are strongly separated, which implies that $\{g^{rk} J \mid k \in \mathbb{Z}\}$ is a collection of pairwise strongly separated hyperplanes. We deduce easily from Lemmas 3.9 and 3.10 that $k \mapsto g^{rk} J$ defines a quasi-isometric embedding $\mathbb{Z} \to \Omega X$, so g is loxodromic in ΩX. Of course, this implies that g has unbounded orbits in ΩX, concluding the proof of our proposition. ∎

4. Axes and rational translation lengths

4.1. General case

We are now ready to prove the first main theorem of this article, namely that isometries of quasi-median graphs induce isometries of the crossing / contact graph with rational

(asymptotic) translation lengths. Our next statement proves Theorem 1.1 and Corollary 1.2 from the introduction.

Theorem 4.1. *Let X be a quasi-median graph, and let ΩX be either the crossing graph or the contact graph of X. In the latter case, we assume that X has no cut-vertices, and, in any case, we assume that a vertex of X always belongs to $\leq N$ cliques. For every isometry $g \in \mathrm{Isom}(X)$ with unbounded orbits in ΩX, there exists a positive $k \leq N^{2\mathrm{QC}(g)+3}$ such that g^k admits an axis in ΩX. Moreover, there exists such an axis containing a hyperplane in contact with $B(o, 2\mathrm{QC}(\gamma)+1) \cap \mathrm{CH}(\gamma)$ where o is an arbitrary vertex of an arbitrary axis γ of g in X.*

Our proof is based on the next criterion, which essentially originates from [14] and which can be extracted from [4]. It is also an immediate consequence of the criterion given by Theorem A.2 which we record in the appendix.

Proposition 4.2. *Let Y be a graph and $g \in \mathrm{Isom}(Y)$ an isometry. Assume that there exists an isometrically embedded subgraph $A \subset Y$ on which $\langle g \rangle$ acts geometrically. Then g cyclically permutes $\leq \mathrm{width}(A)$ pairwise disjoint geodesics in A. Consequently, there exists some $k \leq \mathrm{width}(A)$ such that g^k has an axis in Y and the translation length of g is a rational number with k as a denominator.*

Here, the *width* of a two-ended graph refers to the minimal size of a set of vertices separating the two ends of the graph.

Proof of Theorem 4.1. According to Proposition 2.15, g admits an axis γ in X.

Claim 6. The convex hull $\mathrm{CH}(\gamma)$ of γ is locally finite and $\langle g \rangle$ acts cocompactly on it.

Assume for contradiction that $\mathrm{CH}(\gamma)$ is not locally finite, that is, there exists a vertex $x \in \mathrm{CH}(\gamma)$ with infinitely many neighbours in $\mathrm{CH}(\gamma)$. As a consequence, since every vertex in X belongs to only finitely many cliques, we can find two distinct neighbours $y, z \in \mathrm{CH}(\gamma)$ of x such that x, y, z all belong to a common clique C. Let J denote the hyperplane containing C. Because x, y, z belong to three pairwise distinct sectors delimited by J, it follows from Corollary 2.10 that γ intersects these sectors, which is impossible since we know from Theorem 2.4 that γ cannot intersect J twice. Thus, we have proved that $\mathrm{CH}(\gamma)$ is locally finite.

In order to prove that $\langle g \rangle$ acts cocompactly on $\mathrm{CH}(\gamma)$, it suffices to show that $\mathrm{CH}(\gamma)$ lies in a neighbourhood of γ. So we fix a vertex $x \in \mathrm{CH}(\gamma)$. According to Proposition 3.8, there exist a power $r \geq 1$ and a hyperplane J crossing γ such that $\{g^{rk}J \mid k \in \mathbb{Z}\}$ is a collection of pairwise strongly separated hyperplanes. Up to translating J with a power of g, we assume that x lies between J and $g^r J$. Let $y \in \gamma$ be an arbitrary vertex lying between J and $g^r J$. Observe that, as a consequence of Corollary 2.10, every hyperplane separating x and y must cross γ. On the other hand, a hyperplane crossing γ before $g^{-r}J$ (resp. after $g^{2r}J$) cannot separate x and y since otherwise it would be transverse to $g^{-1}J$

and J (resp. to $g^r J$ and $g^{2r} J$). Consequently, all the hyperplanes separating x and y must cross γ between $g^{-r} J$ and $g^{2r} J$. It follows from Theorem 2.4 that $d(x, y) \leq 3r\|g\|$ where $\|g\|$ denotes the translation length of g along γ. We conclude that $\mathrm{CH}(\gamma)$ lies in the $3r\|g\|$-neighbourhood of γ, as desired. The proof of Claim 6 is complete.

Let $A(g)$ denote the subgraph of ΩX induced by all the hyperplanes of X in contact with $\mathrm{CH}(\gamma)$ (i.e., whose carriers intersect $\mathrm{CH}(\gamma)$).

Claim 7. The action $\langle g \rangle \curvearrowright A(g)$ is cocompact.

It follows from the assumption that every vertex in X belongs to only finitely many cliques and from Claim 6 that $\langle g \rangle$ acts on $A(g)$ with only finitely many orbits of vertices. In order to show that there are only finitely many orbits of edges too, it is sufficient to observe that $A(g)$, as a subgraph of the crossing graph, is locally finite. So fix a hyperplane H in $A(g)$, that is, tangent to $\mathrm{CH}(\gamma)$. According to Proposition 3.8, there exist a power $r \geq 1$ and a hyperplane J crossing γ such that $\{g^{rk} J \mid k \in \mathbb{Z}\}$ is a collection of pairwise strongly separated hyperplanes. Up to translating J by a power of g, we assume that H lies between J and $g^{2r} J$. Clearly, a hyperplane in contact with H must lie between $g^{-r} J$ and $g^{3r} J$. Therefore, the neighbours of H in $A(g)$ are hyperplanes in contact with the (finite) piece of $\mathrm{CH}(\gamma)$ lying between $g^{-r} J$ and $g^{3r} J$, which yields the desired conclusion. The proof of Claim 7 is complete.

Claim 8. $A(g)$ is isometrically embedded in ΩX.

Let $M, N \in A(g)$ be two hyperplanes. Fix an arbitrary geodesic $J_1, \dots, J_k$ from M to N in ΩX and define its complexity as $d(N(J_1), \mathrm{CH}(\gamma)) + \cdots + d(N(J_k), \mathrm{CH}(\gamma))$. If the complexity is zero, then our geodesic lies in $A(g)$ and there is nothing to prove. From now on, we assume that the complexity is not zero, that is, there exists some $2 \leq i \leq k-1$ such that $N(J_i)$ and $\mathrm{CH}(\gamma)$ are disjoint. According to Proposition 2.11, there exists a hyperplane J weakly separating $N(J_i)$ and $\mathrm{CH}(\gamma)$. Necessarily, there exist $1 \leq a \leq i-1$ and $i+1 \leq b \leq k$ such that J is transverse to J_a and J_b. Observe that $b - a \leq 2$ since otherwise it would be possible to shorten our geodesic by replacing $J_a, J_{a+1}, \dots, J_{b-1}, J_b$ with J_a, J, J_b. In other words, J is transverse to J_{i-1} and J_{i+1}. By replacing J_i with J, we obtain a new path from M to N that is a geodesic, since the length remains the same, and whose complexity is smaller, since $d(\mathrm{CH}(\gamma), N(J)) < d(\mathrm{CH}(\gamma), N(J_i))$. By iterating, we eventually get a geodesic lying in $A(g)$, concluding the proof of Claim 8.

Claim 9. Fix a finite set of vertices $S \subset \mathrm{CH}(\gamma)$ separating the ends of $\mathrm{CH}(\gamma)$ and let CS denote the set of the hyperplanes in contact with S. Then CS separates the ends of $A(g)$.

Let A, B be two hyperplanes such that the intersections between their carriers and $\mathrm{CH}(\gamma)$ are non-empty and separated by S. Let $J_1, \dots, J_n$ be a path in ΩX from A to B all of whose hyperplanes are in contact with $\mathrm{CH}(\gamma)$. If one of the J_i is in contact with S, we are done. Otherwise, each J_i can be labelled by the component of $\mathrm{CH}(\gamma) \backslash S$ containing $N(J_i) \cap \mathrm{CH}(\gamma)$. Because $A = J_1$ and $B = J_n$ have different labels, there must exist some $1 \leq i \leq n-1$ such that J_i and J_{i-1} have different labels. Fix three vertices $a \in N(J_i) \cap \mathrm{CH}(\gamma)$, $b \in N(J_{i+1}) \cap \mathrm{CH}(\gamma)$, and $c \in N(J_i) \cap N(J_{i+1})$. Let (x, y, z) be

the median triangle of (a, b, c). Because $N(J_i)$ is gated, a, x, z, c, b all belong to $N(J_i)$. Similarly, b, y, z, c, a all belong to $N(J_{i+1})$. Thus, there exists a geodesic from a to b passing through x and y lying in $(N(J_i) \cup N(J_{i+1})) \cap \mathrm{CH}(\gamma)$. Because a and b are separated by S, it follows that either J_i or J_{i+1} is in contact with S. This concludes the proof of Claim 9.

So far, we know that $A(g)$ is isometrically embedded in ΩX and that $\langle g \rangle$ acts on $A(g)$ freely and cocompactly. It remains to find an upper bound on $\mathrm{width}(A(g))$. As a consequence of Claim 10, fixing an arbitrary vertex $o \in \gamma$, the set $S := B(o, 2\mathrm{QC}(\gamma) + 1) \cap \mathrm{CH}(\gamma)$ separates the ends of $\mathrm{CH}(\gamma)$. Then, Claim 9 implies that CS separates the ends of $A(g)$. We claim that $|CS| \leq N^{2\mathrm{QC}(\gamma)+3}$. This will imply the first assertion of our theorem according to Proposition 4.2. Also, an axis for a power of g in $A(g)$ will necessarily intersect CS, which proves the second assertion of our theorem.

In order to estimate the size of CS, we first estimate the size of S. Notice that every clique of X intersects S along at most one edge. Indeed, if S contains a 3-cycle, then we can find a hyperplane J delimiting three sectors intersecting S, and a fortiori $\mathrm{CH}(\gamma)$. According to Corollary 2.10, this implies that γ intersects three sectors delimited by J, which is impossible since γ is a geodesic (Theorem 2.4). Since every vertex of X belongs to at most N cliques, we deduce that every vertex of S has degree $\leq N$. Hence, $|S| \leq N^{2\mathrm{QC}(\gamma)+1}$. Since a hyperplane in contact with S contains an edge with an endpoint in S, and again because every vertex of X belongs to at most N cliques, we conclude that $|CS| \leq N^{2\mathrm{QC}(\gamma)+3}$, as desired.

Claim 10. Let Y be a graph, $Z \subset Y$ an isometrically embedded subgraph, and $\xi \subset Z$ a bi-infinite geodesic. Assume that the Hausdorff distance D between ξ and Z is finite. Given a vertex $o \in \xi$, the ball $B(o, 2D + 1) \cap Z$ separates the ends of Z.

Let $p, q \in Z$ be two vertices separated by a large ball centred at o. Fix a geodesic ζ from p to q. Every vertex of ζ lies at distance $\leq D$ from a vertex of ξ. If such a vertex lies in the same component ξ_ℓ (resp. ξ_r) of $\xi \setminus B(o, D + 1)$ as p (resp. q), we say that it is a left (resp. right) vertex. Let $a \in \zeta$ denote the last left vertex along ζ and let b denote the first right vertex following a. If $x \in \xi_\ell$ (resp. $y \in \xi_r$) is vertex of ξ at distance $\leq D$ from a (resp. b), then

$$2(D + 1) \leq d(x, y) \leq d(x, a) + d(a, b) + d(b, y) \leq d(a, b) + 2D,$$

hence $d(a, b) \geq 2$. Consequently, there exists some vertex $c \in \zeta$ between a and b. By construction, c is neither left nor right, so there exists some $z \in \xi \cap B(o, D + 1)$ such that $d(c, z) \leq D$. Then

$$d(c, o) \leq d(c, z) + d(z, o) \leq D + D + 1 = 2D + 1,$$

proving that ζ intersects $B(o, 2D + 1)$, as desired. ∎

4.2. Hyperbolic case

Given an isometry g of some quasi-median graph X, the upper bound given by Theorem 4.1 depends only on the local structure of X and the quasiconvexity constant $\mathrm{QC}(g)$ of g. In full generality, the quantity $\mathrm{QC}(g)$ can take arbitrarily large finite values. But there is a worth mentioning situation where the quasiconvexity constants are uniformly bounded: in hyperbolic graphs. Indeed, if X is δ-hyperbolic, then the Hausdorff distance between any two bi-infinite geodesics having the same endpoints at infinity is at most 8δ (see, for instance, [13, Proposition 2.2.2]), hence $\mathrm{QC}(g) \leq 8\delta$ for every loxodromic isometry $g \in \mathrm{Isom}(X)$. Thus, Theorem 4.1 immediately implies the following.

Corollary 4.3. *Let X be a δ-hyperbolic quasi-median graph, and let ΩX be either the crossing graph or the contact graph of X. In the latter case, we assume that X has no cut-vertices, and, in any case, we assume that a vertex of X always belongs to $\leq N$ cliques. For every isometry $g \in \mathrm{Isom}(X)$ with unbounded orbits in ΩX, there exists a positive $k \leq N^{16\delta+3}$ such that g^k admits an axis in ΩX. As a consequence, the translation length of g in ΩX is a rational with a denominator $\leq N^{16\delta+3}$.*

4.3. Some examples

In this section, we describe examples where isometries of (quasi-)median graphs can have arbitrarily small translation lengths in the corresponding crossing graphs, thus contrasting with Corollary 4.3.

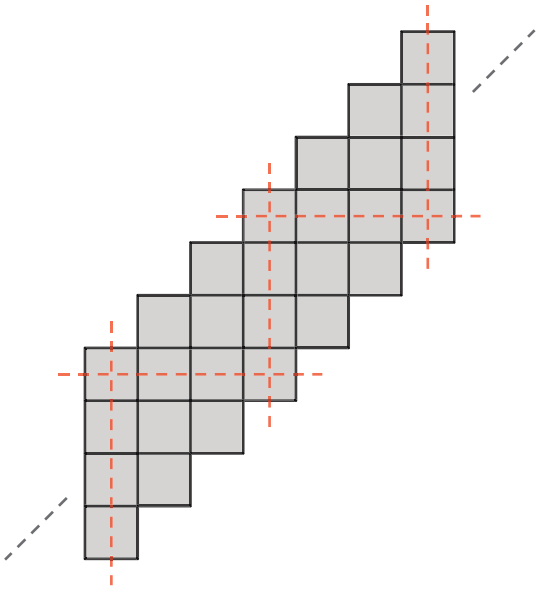

Fix an $n \geq 1$ and let X_n denote the median graph given by the subgraph in $\mathbb{Z}^2$ delimited by a regular staircase and its translate under $(0, n)$. There is an obvious isometry

$g \in \mathrm{Isom}(X_n)$ of translation length 2. The figure on the left represents X_3. If J denotes a vertical hyperplane in X_n, then J belongs to an axis of g^n in the crossing graph ΔX_n on which g^n acts as a translation of length 2. Consequently, g has translation length $2/n$ in ΔX_n.

From this example, we can artificially create a median graph admitting isometries with arbitrarily small but positive translation lengths in the crossing graph. For instance, for every $n \geq 1$ let Z_n be an infinite cyclic group. The free product $G := Z_1 * Z_2 * \cdots$ acts on its Bass–Serre tree T. We can blow up each vertex of T stabilised by a conjugate of Z_n and replace it with a copy of X_n on which Z_n acts. Then G acts on the median graph X thus obtained, and, for every $n \geq 1$, a generator of Z_n acts on the crossing graph of X with a translation length equal to $2/n$.

More interestingly, the example constructed in [49] produces a group acting properly and cocompactly on a median graph and containing elements with arbitrarily small translation length in the crossing graph. This observation is a straightforward consequence of the arguments in [49]. We describe the construction for the reader's convenience.

Let H be a group acting properly and cocompactly on a median graph Y. We assume that there exist elements $g \in H$ and $h_1, h_2, \ldots \in H$ such that:

- g admits a convex axis γ;
- for every $n \geq 1$, the projection of $h_n\gamma$ on γ has finite length $\ell_n \geq n$.

Such examples can be found, for instance, in uniform lattices of products of trees. Given an $n \geq 1$, up to conjugating h_n with a power of g, we can assume that $k_n \cdot \mathrm{proj}_\gamma(h_n\gamma)$ intersects $\mathrm{proj}_{h_n\gamma}(\gamma)$ along a proper subsegment of length $s_n > \ell_n - \|g\|$.

Consider the HNN extension $G := \langle H, t \mid tgt^{-1} = t\rangle$ acting on the median graph X obtained from the tree of spaces modelled on the Bass–Serre tree whose vertex-spaces are copies of Y and whose edge-spaces are copies of $\gamma \times [0,1]$. For every $n \geq 1$, set $k_n := th_n$. The configuration to keep in mind is illustrated by Figure 3.

Set $r := \left\lfloor \frac{\ell_n}{\ell_n - s_n} \right\rfloor$. One easily checks that the equality $d_{\Delta X}(J, k_n^r J) = 2$ holds and that the hyperplanes $k_n^{-r}J, J, k_n^r J, k_n^{2r} J, \ldots$ all lie on a bi-infinite geodesic in ΔX. We have

$$r = \left\lfloor \frac{\ell_n}{\ell_n - s_n} \right\rfloor > \left\lfloor \frac{n}{\|g\|} \right\rfloor, \quad \text{hence } \|k_n\|_{\Delta X} = \frac{2}{r} \leq \frac{2\|g\|}{n}.$$

Thus, the element k_n of G has smaller and smaller translation length in ΔX as n goes to infinity.

5. Constructible quasi-median graphs

In this section, our goal is to show that translation lengths in crossing and contact graphs are algorithmically computable. Of course, quasi-median graphs and isometries have to be given algorithmically in some sense. In order to clarify such an assumption, we introduce *constructible* quasi-median graphs and *computable* isometries.

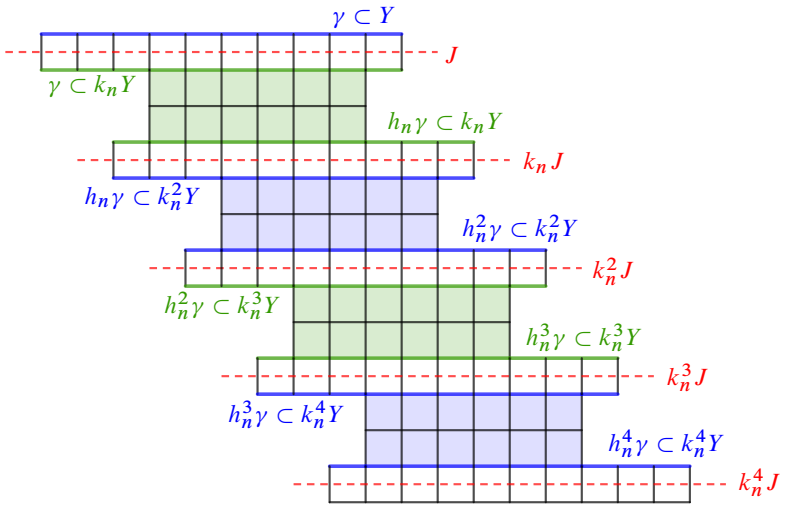

Figure 3. The isometry k_n has a small translation length in the crossing graph.

Definition 5.1. Let X be a quasi-median graph.

- X is *geodesically constructible* if there exists an algorithm that determines, given two vertices $x, y \in X$, the interval $I(x, y)$.
- X is *locally constructible* if there exists an algorithm that yields, given a vertex $x \in X$, a set of neighbours $\mathcal{N}(x)$ such that every clique containing x intersects $\mathcal{N}(x)$ along exactly one vertex.

A quasi-median graph is *constructible* if it is geodesically and locally constructible.

It could be tempted to define a graph X as being *metrically constructible* if there exists an algorithm that constructs, given a vertex x and an integer $R \geq 0$, the ball $B(x, R)$. However, despite the fact that metrically constructible are constructible, this notion only deals with locally finite graphs: Since an algorithm has to stop after a finite amount of time, its output must contain only a finite amount of data. Because we are also interested in quasi-median graphs that are not locally finite (e.g., quasi-median graphs associated with right-angled Artin groups), metrically constructible graphs are not relevant for us. Nevertheless, our quasi-median graphs do satisfy some weak property of local finiteness: Every vertex belongs to only finitely many cliques. This is the property exploited by our definition of locally constructible quasi-median graphs. It is worth noticing, however, that being constructible and being metrically constructible coincide for median graphs since their cliques are edges.

Definition 5.2. Let X be a graph. An isometry $g \in \mathrm{Isom}(X)$ is *computable* if there exists an algorithm that determines, given a vertex $x \in X$, the translate $g \cdot x$.

As a justification that our definition is well founded, let us mention that, given a compact *locally quasi-median prism complex* Q as defined in [17] (e.g., a compact non-positively curved cube complex), the one-skeleton X of the universal cover of Q is a constructible quasi-median graph and every $g \in \pi_1(Q)$ induces on X a computable isometry. This observation, which will not be used in our article, essentially follows from [17, Lemma 2.2] (already used during the proof of Lemma 3.2), which shows that any two homotopic paths in the one-skeleton of Q can be transformed into each other by a sequence of elementary transformations.

5.1. Basic algorithmic problems

In this section, we record a few elementary problems in constructible quasi-median graphs that can be solved algorithmically. These statements will be frequently used in the next sections.

Lemma 5.3. *Let X be a geodesically constructible quasi-median graph. There exists an algorithm that determines, given a vertex $x \in X$ and an edge $e \subset X$, the projection of x on the clique containing e.*

Proof. Let C denote the clique containing e and p, q the endpoints of e. The projection of x on C is the unique vertex in $I(x, p) \cap I(x, q) \cap C$. By assumption, $I(x, p)$ and $I(x, q)$ are algorithmically constructible. And a vertex belongs to C if and only if it lies at distance ≤ 1 from both p and q. ∎

Corollary 5.4. *Let X be a geodesically constructible quasi-median graph. There exists an algorithm that determines whether two given edges in X belong to the same hyperplane.*

Proof. Let $e, e' \subset X$ be two edges. As a consequence of Proposition 2.5, e belongs to the same hyperplane as e' if and only if the endpoints of e have distinct projections on the clique containing e'. Thus, our corollary follows from Lemma 5.3. ∎

Lemma 5.5. *Let X be a geodesically constructible quasi-median graph. There exists an algorithm that determines, given two edges $e, e' \subset X$, whether or not the hyperplanes containing e and e' are transverse.*

Proof. The key observation is that J and J' are transverse if and only if there exists a geodesic between two endpoints of e and e' with two consecutive edges spanning a 4-cyclic and belonging to J and J'.

Indeed, let p and p' be two endpoints of e and e' that are separated by both J and J'. Fix a geodesic γ_0 from p to p' and let $\varepsilon, \varepsilon' \subset \gamma_0$ denote the two edges in J, J'. If J and J' are transverse, then every hyperplane crossing γ_0 between ε and ε' are transverse to J or J'. If ε and ε' are adjacent, we are done. Otherwise, let $\eta \subset \gamma_0$ denote the closest edge from ε that lies between ε and ε' and whose hyperplane $J(\eta)$ is transverse to J. If ε and η

are not adjacent, then the hyperplane containing the edge η' between ε and η adjacent to η must be transverse to $J(\eta)$, so η and η' span a 4-cycle and we can modify γ_0 by flipping this 4-cycle in order to decrease the distance between η and ε. Thus, up to modifying the geodesic γ_0, we can assume that ε and η are adjacent. Necessarily, these two edges span a 4-cycle, because their hyperplanes are transverse, and we can again modify γ_0 by flipping this 4-cycle in order to decrease the distance between ε and ε'. After a few iterations, we get a geodesic with ε and ε' adjacent, providing the desired geodesic. The converse of our claim is clear.

Because geodesics are constructible, that spanning a 4-cycle is algorithmically decidable as a particular, and that belonging to a given hyperplane is algorithmically decidable according to Corollary 5.4, our lemma follows from the characterisation just proved. ∎

Lemma 5.6. *Let X be a constructible quasi-median graph. There exists an algorithm that decides, given a vertex $x \in X$ and an edge $e \subset X$, whether or not x belongs to the carrier of the hyperplane containing e.*

Proof. Because X is locally constructible, we get algorithmically a collection of edges $e_1, \dots, e_n$ containing x such that every hyperplane containing x in its carrier contains one of the e_i. So just need to check whether e belongs to the same hyperplane as one of the e_i, which is possible algorithmically according to Corollary 5.4. ∎

Lemma 5.7. *Let X be a constructible quasi-median graph. There exists an algorithm that gives, given two edges $e, e' \subset X$, two vertices $x, x' \in X$ minimising the distance between the carriers of the hyperplanes containing e, e'.*

Proof. Let N and N' denote the carriers of the hyperplanes containing e and e'. Fix two endpoints p and p' of e and e'. Let q denote the projection of p on N' and q' the projection of q on N. We know that q belongs to $I(p, p')$ and that q' belongs to $I(q, p)$, so there exists a geodesic between p and p' passing through q and q'. But it follows from Lemma 2.6 that p and p' are two vertices minimising the distance between N and N'.

Therefore, in order to get two vertices minimising the distance between N and N', it suffices, for every geodesic γ between p and p', to find the last vertex a_γ in γ belonging to N and the first vertex a'_γ in N' (algorithmically thanks to Lemma 5.6), and to choose a pair a_γ, a'_γ with $d(a_\gamma, a'_\gamma)$ minimal. Because geodesics are constructible, the whole process is algorithmic. ∎

Corollary 5.8. *Let X be a constructible quasi-median graph. There exists an algorithm that determines, given two edges $e, e' \subset X$, whether or not the hyperplanes containing e and e' are in contact.*

Proof. Two hyperplanes are in contact if and only if the distance between their carriers is zero, so the desired conclusion follows from Lemma 5.7. ∎

Lemma 5.9. *Let X be a constructible quasi-median graph. There exists an algorithm that determines, given two edges $e, e' \subset X$, whether or not the hyperplanes containing e and e' are strongly separated.*

Proof. Let J and J' denote the hyperplanes containing e and e', and N and N' their carriers. According to Lemma 5.7, we can find algorithmically two vertices $x \in N$ and $x' \in N'$ minimising the distance between N and N'. If J and J' are not strongly separated, then the projection of N' on N is not reduced to single vertex, so x must have a neighbour such that the hyperplane H separating these two vertices is transverse t both J and J'. But X is locally constructible, so we can get algorithmically edges $e_1, \dots, e_n$ containing x such that H contains one of the e_i and we can check thanks to Lemma 5.5 that H is transverse to both J and J'. Thus, the existence of a hyperplane transverse to both J and J' is decidable algorithmically. ∎

5.2. Strongly contracting isometries

This section is dedicated to the proof of the following theorem. As a consequence of Proposition 3.8, it can be thought of as a particular case of Theorem 1.3 since it recognises algorithmically the isometries with zero translation length in the crossing and contact graphs.

Theorem 5.10. *Let X be a constructible quasi-median graph. Let ΩX denote the crossing or contact graph of X. In the former case, we assume that X has no cut-vertex, and, in any case, we assume that every vertex of X belongs to $\leq N$ cliques. There exists an algorithm that determines, given a computable isometry $g \in \mathrm{Isom}(X)$, whether or not g is strongly contracting.*

Our argument has two steps. First, we show that crossing and contact graphs of constructible quasi-median graphs inherit weak algorithmic properties. More precisely, Proposition 5.12 shows geodesics in crossing and contact graphs can be algorithmically constructed in the following sense.

Definition 5.11. Let X be a graph. A *combing* is a map γ that associates with every pair of vertices $(x, y) \in X^2$ a path $\gamma(x, y)$ from x to y. It is *geodesic* if $\gamma(x, y)$ is a geodesic for all $x, y \in X$, and *constructible* if there exists an algorithm that, given two vertices $x, y \in X$, yields $\gamma(x, y)$.

Our second step is to show that, in a hyperbolic graph admitting a constructible geodesic combing, it can be decided whether a (non-parabolic) isometry is elliptic or loxodromic. This is Proposition 5.14.

Before turning to the proofs of our proposition, let us mention that, from the algorithmic point of view, we think about crossing and contact graphs slightly differently as before. Indeed, a vertex in such a graph is, by definition, a hyperplane. But a hyperplane, as an infinite collection of edges, is not algorithmic-friendly. Instead, we think of a

vertex in the crossing or contact graph as an edge of the quasi-median graph under consideration, and we define two edges as representing the same vertex if they belong to the same hyperplane (which can be decided algorithmically according to Lemma 5.5); two edges represent two adjacent vertices if the hyperplanes containing them are transverse or in contact (which can be decided algorithmically according to Corollary 5.8). Formally, this amounts to working with a pseudo-metric on the set of the edges of our quasi-median graph.

Proposition 5.12. *Let X be a constructible quasi-median graph. Let ΩX denote the crossing or contact graph of X; in the former case, we assume that X has no cut-vertex. Then ΩX admits a constructible geodesic combing.*

The proposition will essentially follow from the combination of our next lemma with the basic results of the previous section.

Lemma 5.13. *Let X be a quasi-median graph, A, B two hyperplanes, and $a \in N(A), b \in N(B)$ two vertices. There exist a geodesic $H_1, \dots, H_r$ between A and B in the crossing (resp. contact) graph and a geodesic γ between a and b in X such that $\gamma \cap N(H_i) \neq \emptyset$ for every $1 \leq i \leq r$.*

Proof. Let $H_0, \dots, H_r$ be a geodesic in the crossing (resp. contact) graph such that the quantity $d(a, N(H_0)) + \cdots + d(a, N(H_r))$ is as small as possible. Define the sequence $(x_i)_{0 \leq i \leq r+1}$ as follows:

- $x_0 := a$ and $x_{r+1} := b$;
- for every $1 \leq i \leq r$, x_i is the projection of x_{i-1} on $N(H_i)$.

For every $0 \leq i \leq r$, fix a geodesic $[x_i, x_{i+1}]$ between x_i and x_{i+1}, and let γ denote the concatenation $[x_0, x_1] \cup \cdots \cup [x_r, x_{r+1}]$. By construction, $x_i \in \gamma \cap N(H_i)$ for every $1 \leq i \leq r$. In order to conclude our lemma, it suffices to show that γ is a geodesic.

If γ is not a geodesic, then it crosses twice some hyperplane. Observe that, for every $0 \leq i \leq r-1$, no hyperplane crosses both $[x_i, x_{i+1}]$ and $[x_{i+1}, x_{i+2}]$ because every hyperplane separating x_i from its projection x_{i+1} has to separate x_i from $N(H_{i+1})$. Also, a hyperplane J cannot cross $[x_i, x_{i+1}]$ and $[x_j, x_{j+1}]$ for some $0 \leq i, j \leq r-1$ satisfying $|i - j| > 2$ since otherwise we could shorten $H_0, \dots, H_r$ by replacing $H_i, \dots, H_j$ with H_i, J, H_j. Therefore, if γ is not a geodesic, there must exist some $0 \leq i \leq r-2$ such that $[x_i, x_{i+1}]$ and $[x_{i+2}, x_{i+3}]$ are crossed by a common hyperplane J. As a consequence of what we already know, J cannot cross γ a third time, so it separates a from $N(H_{i+1})$. Because J must be transverse to both H_i and H_{i+2}, $H_0, \dots, H_i, J, H_{i+2}, \dots, H_r$ is also a geodesic in the crossing (resp. contact) graph of X with

$$\begin{aligned} &d(a, N(H_0)) + \cdots d(a, N(H_i)) + d(a, N(J)) + d(a, N(H_{i+2})) + \cdots d(a, N(H_r)) \\ &\qquad < d(a, N(H_0)) + \cdots + d(a, N(H_r)), \end{aligned}$$

contradicting the minimality satisfied by $H_0, \dots, H_r$. ∎

Proof of Proposition 5.12. Let $e, e' \subset X$ be two edges. Fix to endpoints p and p' of e and e'. For all geodesic γ from p to p' (which can be enumerated algorithmically because X is geodesically constructible), integer $1 \leq k \leq d(p, p')$, vertices $x_0 := p, x_1, \ldots, x_{k-1}, x_k := p'$ of γ, and hyperplanes $J_1, \ldots, J_{k-1}$ containing, respectively, $x_1, \ldots, x_{k-1}$ in their carriers (which can be enumerated algorithmically because X is locally constructible), check whether $A, J_1, \ldots, J_{k-1}, B$ defines a path in the crossing (resp. contact) graph (which can be done algorithmically according to Lemma 5.5 (resp. Corollary 5.8)). As a consequence of Lemma 5.13, the path in the crossing (resp. contact) graph with the smallest k yields a geodesic between A and B in the crossing (resp. contact) graph of X. We choose such a geodesic in order to define our constructible geodesic combing. ∎

Next, we turn to the second step of the proof of our theorem.

Proposition 5.14. *Let X be a δ-hyperbolic graph admitting a constructible geodesic combing. There exists an algorithm that determines, given a non-parabolic computable isometry $g \in \mathrm{Isom}(X)$, whether g is elliptic or loxodromic.*

Our proposition follows from our next two lemmas, which allow us to distinguish elliptic and loxodromic isometries in hyperbolic graphs.

Lemma 5.15. *Let X be a δ-hyperbolic graph and $g \in \mathrm{Isom}(X)$ a loxodromic isometry. For every $x \in X$ and every $k \in \mathbb{Z}$, $d(x, g^k x) \geq |k| + 16\delta$.*

Proof. According to [13, Proposition 10.6.4], we have

$$d(x, g^k x) \geq [g^k] \geq \|g^k\| + 16\delta = |k| \cdot \|g\| + 16\delta,$$

where, for every isometry h of X,

$$[h] := \min\{d(z, hz) \mid z \in X\} \quad \text{and} \quad \|h\| := \lim_{n \to +\infty} \frac{1}{n} d(o, h^n o)$$

for an arbitrary basepoint $o \in X$. But X is a graph, so we must have $\|g\| \geq 1$, concluding the proof of our lemma. ∎

Lemma 5.16. *Let X be a δ-hyperbolic graph and $g \in \mathrm{Isom}(X)$ an elliptic isometry. For every vertex $x \in X$, every integer $k \in \mathbb{Z}$, and every geodesic $[x, g^k x]$ between x and $g^k x$, there exists a vertex $y \in [x, g^k x]$ satisfying $d(y, g^k y) \leq 32\delta$.*

Proof. We begin by proving the following observation.

Fact 1. The subgraph $F := \{z \in X \mid \mathrm{diam}(\langle g \rangle \cdot z) \leq 2\delta\}$ is non-empty and is also 12δ-quasiconvex.

The fact that F is non-empty follows from [8]. To be precise, given a bounded subgraph $B \subset X$, define its radius $\mathrm{rad}(B)$ as $\min\{r \geq 0 \mid \exists c \in X, S \subset B(c, r)\}$. A vertex $c \in X$ is a centre of B if $B \subset B(c, \mathrm{rad}(B))$. According to [8, Lemma 2.1], the set of the centres of a bounded subgraph has diameter at most 2δ, so F contains the centre of every $\langle g \rangle$-orbit. (It is worth mentioning that [8, Lemma 2.1] assumes that the hyperbolic space is proper, but only in order to assure the existence of a centre, which is not needed here since our metric is discrete.) Now, let $a, b \in F$ be two vertices and $z \in I(a, b)$ a third vertex. Because the distance in X is 8δ-convex [13, Corollaire 10.5.3], we know that $\langle g \rangle \cdot z$ has diameter at most 12δ. Taking a centre c of this orbit, we have

$$d(z, F) \leq d(z, c) \leq \mathrm{rad}(\langle g \rangle \cdot z) \leq \mathrm{diam}(\langle g \rangle \cdot z) \leq 12\delta,$$

concluding the proof of our fact.

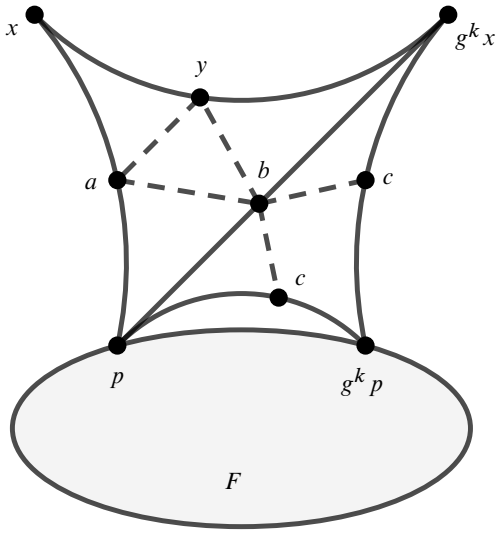

Now, fix a vertex $x \in X$, an integer $k \in \mathbb{Z}$, and a geodesic $[x, g^k x]$ between x and $g^k x$. Also, fix an arbitrary vertex $p \in F$ and some geodesics $[x, p]$, $[p, g^k x]$, and $[p, g^k p]$. Because X is δ-hyperbolic, there exist vertices $y \in [x, g^k x]$, $a \in [x, p]$, and $b \in [p, g^k x]$ pairwise at distance $\leq 2\delta$. Also, there exists $c \in [p, g^k p] \cup g^k[p, x]$ such that $d(b, c) \leq \delta$.

If $c \in [p, g^k p]$, then

$$d(y, F) \leq d(y, b) + d(b, c) + d(c, F) \leq 2\delta + \delta + 12\delta = 15\delta.$$

Therefore, $\mathrm{diam}(\langle g \rangle \cdot y) \leq 32\delta$.

Otherwise, if $c \in g^k[p, x]$, we have

$$|d(g^k p, c) - d(p, a)| \leq d(a, c) + d(p, g^k p) \leq 3\delta + 2\delta = 5\delta.$$

Because c and $g^k a$ lies on the same geodesic $g^k[p, x]$, this implies that $d(g^k a, c) \leq 5\delta$. Thus,

$$d(g^k y, y) \leq d(g^k y, g^k a) + d(g^k a, c) + d(c, y) \leq 2\delta + 5\delta + 3\delta = 10\delta.$$

Therefore, we have proved that $d(g^k y, y) \leq 32\delta$ in any case, concluding the proof of our lemma. ∎

Proof of Proposition 5.14. Fix an arbitrary vertex $o \in X$. Given a computable isometry $g \in \mathrm{Isom}(X)$, construct the geodesic $[o, g^{17\delta} o]$ given by our combing, and, for every vertex $x \in [o, g^{17\delta} o]$, compute $d(x, g^{17\delta} x)$. As a consequence of Lemmas 5.15 and 5.16, either all these values are greater than 33δ, and g must be loxodromic, or g is elliptic otherwise. ∎

Proof of Theorem 5.10. As a consequence of Propositions 5.12 and 5.14, there exists an algorithm that determines, given a computable isometry $g \in \mathrm{Isom}(X)$ that does not include a parabolic isometry on the crossing (resp. contact) graph, whether or not g is loxodromic in the crossing (resp. contact) graph, which amounts to saying that g is strongly contracting according to Proposition 3.8. Since we know from Proposition 3.8 that an isometry of X cannot induce a parabolic isometry of the crossing (resp. contact) graph, our theorem follows. ∎

5.3. Computing translation lengths

We are now finally ready to prove the second main result of this article, namely Theorem 1.3. In fact, we are going to prove the following stronger version of it.

Theorem 5.17. *Let X be a constructible quasi-median graph, and let ΩX be either the crossing graph or the contact graph of X. In the former case, we assume that X has not cut-vertex, and, in any case, we assume that a vertex of X belongs to at most N cliques for some fixed $N \geq 1$. There exists an algorithm that computes, given a computable isometry $g \in \mathrm{Isom}(X)$, the translation length of g in ΩX. Moreover, if the translation length is positive, then the algorithm provides an integer $k \geq 1$ and a hyperplane J such that J belongs to an axis of g^k in ΩX.*

In addition of the preliminary work done so far, we need to be able to algorithmically find axes of strongly contracting isometries in constructible quasi-median graphs.

Proposition 5.18. *Let X be a constructible quasi-median graph. Then there exists an algorithm that provides, given a computable strongly contracting isometry $g \in \mathrm{Isom}(X)$, a vertex $x \in X$ and an integer $N \geq 1$ such that x belongs to an axis of g^N.*

The following notion will be useful in order to find axes.

Definition 5.19. Let X be a quasi-median graph. An isometry $g \in \mathrm{Isom}(X)$ *skewers* an oriented edge $e := (x, y)$ if the sector containing y and delimited by the hyperplane containing e satisfies $gS \subsetneq S$.

As the typical example to keep in mind, let γ be an axis of our isometry g and let $e \subset \gamma$ be an edge. Orient γ such that g translates the vertices γ in the positive direction and endow e with the induced orientation. Then, if the hyperplanes containing e and ge are not transverse, then g skewers e.

First, we notice that skewering an edge can be detected algorithmically.

Lemma 5.20. *Let X be a quasi-median graph, $g \in \mathrm{Isom}(X)$ an isometry, and $e := (x, y)$ an oriented edge. Then g skewers e if and only if the following conditions hold.*

- *gy does not belong to $I(y, gx)$.*
- *The hyperplanes containing e and ge are distinct and non-transverse.*

Proof. First, assume that g skewers e. Let S denote the sector containing y delimited by the hyperplane J containing e. Clearly, $gS \subsetneq S$ implies that the hyperplanes J and gJ are distinct and non-transverse. Moreover, y and gx do not belong to gS, hence $I(y, gx) \subset gS^c$. This implies that gy, which belongs to gS, cannot belong to $I(y, gx)$. Conversely, assume that the two conditions given by our lemma hold. Because J and gJ are distinct and non-transverse, we have $gS \subsetneq S$ unless gS is the sector delimited by gJ containing J. If so, gy coincides with the projection of y on the clique containing gx and gy, hence $gy \in I(y, gx)$, contradicting our first condition. ∎

Corollary 5.21. *Let X be a geodesically constructible quasi-median graph. There exists an algorithm that determines, given a computable isometry $g \in \mathrm{Isom}(X)$ and an oriented edge e, whether or not g skewers e.*

Proof. The first condition of Lemma 5.20 is decidable algorithmically because X is geodesically constructible, as well as the second condition as a consequence of Corollary 5.4 and Lemma 5.5. ∎

Next, we show that powers of strongly contracting isometries skewer edges.

Lemma 5.22. *Let X be a quasi-median graph and $g \in \mathrm{Isom}(X)$ a strongly contracting isometry. There exists some $N \geq 1$ such that, for every vertex $x \in X$ and every geodesic $[x, g^N x]$ from x to $g^N x$, some oriented edge $e \subset [x, g^N x]$ is skewered by g^N and the hyperplanes containing e and $g^N e$ are strongly separated.*

Proof. According to Proposition 3.8, there exists some integer $N \geq 1$ and some hyperplane J crossing an axis of g such that $\{g^{kN} J \mid k \in \mathbb{Z}\}$ is a collection of pairwise strongly separated hyperplanes. Up to replacing J with a $\langle g^N \rangle$-translate, assume that our vertex x lies between $g^{-N} J$ and J. Of course, $g^N x$ lies between J and $g^N J$. Therefore, our

geodesic $[x, g^N x]$ from x to $g^N x$ must cross J; let e denote the corresponding oriented edge. The hyperplanes containing e and $g^N e$, namely J and $g^N J$, are strongly separated, and g^N clearly skewers e. ∎

Finally, we explain how skewered edges can be helpful in order to find axes of strongly contracting isometries.

Lemma 5.23. *Let X be a quasi-median graph, $g \in \mathrm{Isom}(X)$ a loxodromic isometry, and $e \subset X$ an oriented edge skewered by g. If the hyperplanes containing e and ge are strongly separated, then the two vertices minimising the distance between the carriers of these hyperplanes belong to an axis of g^2.*

Proof. Let $x, y \in X$ be two vertices such that $e = (x, y)$, let J denote the hyperplane containing e, and let S denote the sector delimited by J containing y. Given an axis γ of g, notice that J crosses γ. Indeed, otherwise either $\gamma \subset S$, and $J, gJ, g^2J, \dots$, yields infinitely many hyperplanes separating S^c and γ, which is impossible, or $S \cap \gamma = \emptyset$, and, given an arbitrary vertex $v \in S$, the hyperplanes $J, \dots, g^n J$ separate $g^n v$ from γ for every $n \geq 1$, which is also impossible since $d(g^n x, \gamma) = d(x, \gamma)$ for every $n \geq 1$. Thus, $\{g^k J \mid k \in \mathbb{Z}\}$ defines a collection of pairwise strongly separated hyperplanes crossing the axis γ of g. Fix a vertex $p \in \gamma$ lying between $g^{-1}J$ and J. Of course, $g^2 p$ lies between gJ and $g^2 J$. As a consequence of Proposition 2.15, g^2 admits an axis passing through any fixed geodesic between p and $g^2 p$. But, given two vertices $x \in N(J)$ and $y \in N(gJ)$ minimising the distance between $N(J)$ and $N(gJ)$, concatenating a geodesic from p to x with a geodesic from x to y and next a geodesic from y to $g^2 p$ yields a geodesic from p to $g^2 p$. (Indeed, no hyperplane can cross both $[p, x]$ and $[y, g^2 p]$ because J and gJ are strongly separated, and no hyperplane can cross both $[p, x]$ and $[x, y]$ or $[x, y]$ and $[y, g^2 p]$ since the hyperplanes separating x and y coincide with the hyperplanes separating $N(J)$ and $N(gJ)$. So no hyperplane can cross our path twice, proving that it is a geodesic.) ∎

Proof of Proposition 5.18. Fix an arbitrary vertex $o \in X$. For successive powers $k = 1, 2, \dots$, check whether some oriented edge e of some geodesic between o and $g^k o$ satisfies the following condition: g^k skewers e and the hyperplanes containing e and $g^k e$ are strongly separated. This can be done algorithmically according to Corollary 5.21 and Lemma 5.9. According to Lemma 5.22, we eventually find a power $N \geq 1$ and an oriented edge e for which it works. Applying Lemma 5.7 to the edges e and $g^N e$ yields two vertices that, according to Lemma 5.23, belong to an axis of g^{2N}. So our algorithm outputs one of these two vertices and the integer $2N$. ∎

Proof of Theorem 5.17. Let $g \in \mathrm{Isom}(X)$ be a computable isometry.

Step 1. Determine whether or not g is strongly contracting. If no, stop and output $\|g\|_{\Omega X} = 0$. If yes, pass to the next step.

Checking whether g is strongly contracting can be done algorithmically according to Theorem 5.10. If not, then it follows from Proposition 3.8 that g has bounded orbits in ΩX, hence $\|g\|_{\Omega X} = 0$. From now on, we assume that g is strongly contracting.

Step 2. Find an edge e and a power $L \geq 1$ such that e is contained in an axis of g^L and the hyperplanes containing e and $g^L e$ are strongly separated.

According to Proposition 5.18, we can find algorithmically a vertex $x \in X$ and a power $n \geq 1$ such that x belongs to an axis of g^n. Construct a geodesic $[x, g^n x]$. We know from Proposition 2.15 that the $\langle g^n \rangle$-translates of $[x, g^n x]$ yields an axis γ of g^n. Testing the successive values of $k \geq 0$, find an edge $e \subset [x, g^n x]$ such that the hyperplanes containing e and $g^{kn} e$ are strongly separated. This can be done algorithmically according to Lemma 5.9, and we know that we will eventually find such a k thanks to Proposition 3.8. Then e and $L := kn$ are the data we are looking for.

Step 3. Compute $\mathrm{HQC}(g)$.

Let $\mathcal{H}, \mathcal{V}$ be two finite transverse collection of hyperplanes crossing the axes of g. Up to translating $\mathcal{H}, \mathcal{V}$ by a power of g^L, we assume that at least one hyperplane H in $\mathcal{H} \cup \mathcal{V}$ lies between $g^{-L} J$ and $g^L J$, where J denotes the hyperplane containing our edge e. Without loss of generality, say that H belongs to $\mathcal{H}$. Because every hyperplane in $\mathcal{V}$ is transverse to H, they must lie between $g^{-2L} J$ and $g^{2L} J$. And because every hyperplane in $\mathcal{H}$ is transverse to every hyperplane in $\mathcal{V}$, they must lie between g^{-3L} and g^{3L}. Thus, every hyperplane in $\mathcal{H} \cup \mathcal{V}$ must separate $g^{-3L} e$ and $g^{3L} e$. We immediately deduce that $\mathrm{HQC}(g) \leq 6L\|g\|_X$, which is sufficient for the rest of the proof; but we can also compute precisely $\mathrm{HQC}(g)$ by studying the pattern of transverse hyperplanes separating $g^{-3L} e$ and $g^{3L} e$ (Lemma 5.5).

Step 4. Let x be an endpoint of e. Enumerate the hyperplanes in contact with

$$B\left(x, N^{2\mathrm{HQC}(g)}\right) \cap I\left(g^{-1-L(4\mathrm{HQC}(g)+1)} x, g^{1+L(4\mathrm{HQC}(g)+1)} x\right)$$

and, for each such hyperplane H, compute $d_{\Omega X}(H, g^{N^{2\mathrm{HQC}(g)}!} H)$. Pick a hyperplane K for which this quantity is minimal. Then K belongs to an axis of $g^{N^{2\mathrm{HQC}(g)}!}$ in ΩX and $\|g\|_{\Omega X} = d_{\Omega X}(K, g^{N^{2\mathrm{HQC}(g)}!} K)/N^{2\mathrm{HQC}(g)}!$.

We deduce from Theorem 4.1 and Lemma 2.20 that there exists some hyperplane H in contact with $F := B(x, 4\mathrm{HQC}(g) + 1) \cap \mathrm{CH}(\gamma)$ such that K belongs to an axis of g^k in ΩX where $k := N^{2\mathrm{HQC}(g)}!$.

Observe that $F \subset I(g^{-Q} x, g^Q x)$ where $Q := 1 + L(4\mathrm{HQC}(g) + 1)$. Indeed, x lies between $g^{-L} J$ and $g^L J$, so $B(x, 4\mathrm{HQC}(g) + 1)$ must lie between $g^{-Q+1} J$ and $g^{Q-1} J$. Now, if p belongs to $\mathrm{CH}(\gamma)$ and lies between $g^{-Q+1} J$ and $g^{Q-1} J$, then p must lie in $I(g^{-Q} x, g^Q x)$ because otherwise there would exist some hyperplane separating p from $\{g^{-Q} x, g^Q x\}$ but also crossing γ, necessarily not between $g^{-Q} x$ and $g^Q x$, so such a hyperplane would be transverse to both $g^{-Q+1} J$ and $g^{-Q} J$ or to both $g^{Q-1} J$ and $g^Q J$, which is impossible since these two pairs are strongly separated.

Thus, there exists some hyperplane H in contact with $F' := B(x, 4\mathrm{HQC}(g) + 1) \cap I(g^{-Q}x, g^{Q}x)$ such that H belongs to an axis of g^k in ΩX. We can enumerate the hyperplanes in contact with F' because an interval is always finite (Lemma 2.8) and because X is locally constructible. Because distances in ΩX are computable according to Proposition 5.12, we can pick a hyperplane K in our collection for which $d_{\Omega X}(K, g^k K)$ is minimal. It follows from Lemma 2.16 that K belongs to an axis of g^k. We have

$$\|g\|_{\Omega X} = \frac{1}{k}\|g^k\| = \frac{1}{k}\lim_{n\to+\infty}\frac{d_{\Omega X}(K, g^{kn}K)}{n} = d_{\Omega X}(K, g^k K)/k,$$

concluding the proof of our theorem. ∎

6. Graph products of groups

Let Γ be a graph and $\mathcal{G} = \{G_u \mid u \in V(\Gamma)\}$ a collection of groups indexed by the vertex-set $V(\Gamma)$ of Γ. The *graph product* $\Gamma\mathcal{G}$ is

$$\langle G_u\ (u \in V(\Gamma)) \mid [G_u, G_v] = 1\ (\{u, v\} \in E(\Gamma))\rangle,$$

where $E(\Gamma)$ denotes the edge-set of Γ and where $[G_u, G_v] = 1$ is a shorthand for $[g, h] = 1$ for all $g \in G_u$, $h \in G_v$. The groups of $\mathcal{G}$ are referred to as *vertex-groups*. We emphasise that, unless explicitly stated, vertex-groups are not assumed to be finite.

Convention. We always assume that the groups in $\mathcal{G}$ are non-trivial. Notice that it is not a restrictive assumption, since a graph product with some trivial factors can be described as a graph product over a smaller graph all of whose factors are non-trivial.

A *word* in $\Gamma\mathcal{G}$ is a product $g_1 \cdots g_n$ where $n \geq 0$ and where, for every $1 \leq i \leq n$, $g_i \in G$ for some $G \in \mathcal{G}$; the g_is are the *syllables* of the word, and n is the *length* of the word. Clearly, the following operations on a word does not modify the element of $\Gamma\mathcal{G}$ it represents:

Cancellation: Delete the syllable g_i if $g_i = 1$.

Amalgamation: If $g_i, g_{i+1} \in G$ for some $G \in \mathcal{G}$, replace the two syllables g_i and g_{i+1} by the single syllable $g_i g_{i+1} \in G$.

Shuffling: If g_i and g_{i+1} belong to two adjacent vertex-groups, switch them.

A word is *graphically reduced* if its length cannot be shortened by applying these elementary moves. Every element of $\Gamma\mathcal{G}$ can be represented by a graphically reduced word, and this word is unique up to the shuffling operation. For more information on graphically reduced words, we refer to [28] (see also [20, 33]).

The connection between graph products and quasi-median graphs is made explicit by the following statement [17, Proposition 8.2, Corollary 8.7].

Theorem 6.1. *Let Γ be a graph and $\mathcal{G}$ a collection of groups indexed by $V(\Gamma)$. The Cayley graph*

$$\mathrm{QM}(\Gamma,\mathcal{G}) := \mathrm{Cayl}\Big(\Gamma\mathcal{G}, \bigcup_{G\in\mathcal{G}} G\backslash\{1\}\Big)$$

is a quasi-median graph of cubical dimension $\mathrm{clique}(\Gamma) = \max\{\#V(\Lambda) \mid \Lambda \subset \Gamma \text{ clique}\}$.

Notice that $\Gamma\mathcal{G}$ naturally acts by isometries on $\mathrm{QM}(\Gamma,\mathcal{G})$ by left-multiplication and that, as a Cayley graph, the edges of $X(\Gamma,\mathcal{G})$ are naturally labelled by generators, but also by vertices of Γ (corresponding to the vertex-group which contains the generator). It turns out that any two edges of $\mathrm{QM}(\Gamma,\mathcal{G})$ must be labelled by the same vertex of Γ if they belong to the same hyperplane (see [17, Lemma 8.9]), which implies that the hyperplanes of $\mathrm{QM}(\Gamma,\mathcal{G})$ are also naturally labelled by vertices of Γ. An easy observation that will be needed later is [17, Lemma 8.12], namely the following lemma.

Lemma 6.2. *Let Γ be a simplicial graph and $\mathcal{G}$ a collection of groups indexed by $V(\Gamma)$. Two transverse hyperplanes are labelled by adjacent vertices of Γ, and two tangent hyperplanes are labelled by distinct vertices of Γ.*

Essentially by construction of the quasi-median graph, we have the following description of its geodesics [17, Lemma 8.3].

Lemma 6.3. *Let Γ be a graph and $\mathcal{G}$ be a collection of groups indexed by $V(\Gamma)$. Fix two elements $g, h \in \Gamma\mathcal{G}$ and write $g^{-1}h$ as a graphically reduced word $u_1\cdots u_n$. Then the sequence of vertices*

$$g, gu_1, gu_1u_2, \dots, gu_1\cdots u_n = h$$

defines a geodesic between g and h in $\mathrm{QM}(\Gamma,\mathcal{G})$. Conversely, any geodesic between g and h is labelled by a graphically reduced word representing $g^{-1}h$.

Before turning to the proof of Theorem 1.4, we need a last preliminary lemma. Namely, we need to understand when the quasi-median graph of a graph product is hyperbolic. A characterisation is already given by [17, Fact 8.33], but we include a proof here with an estimate on the hyperbolicity constant.

Lemma 6.4. *Let Γ be a simplicial graph and $\mathcal{G}$ a collection of groups indexed by $V(\Gamma)$. The quasi-median graph $\mathrm{QM}(\Gamma,\mathcal{G})$ is hyperbolic if and only if $\mathrm{clique}(\Gamma)$ is finite and Γ has no induced 4-cycle. If so, the graph is $5\mathrm{clique}(\Gamma)$-hyperbolic.*

Proof. If Γ contains a clique with n vertices, then $\mathrm{QM}(\Gamma,\mathcal{G})$ contains an isometrically embedded product of n complete graphs, and a fortiori of $[0,1]^n$. Therefore, if $\mathrm{clique}(\Gamma)$ is infinite, that is, if n can be chosen arbitrarily large, then $\mathrm{QM}(\Gamma,\mathcal{G})$ cannot be hyperbolic. If Γ contains an induced 4-cycle (a,b,c,d), then, fixing non-trivial elements $p \in G_a$, $q \in G_b$, $r \in G_c$, and $s \in G_d$, the map $(i,j) \mapsto (pr)^i(qs)^j$ defines an isometric embedding $\mathbb{Z}^2 \to \mathrm{QM}(\Gamma,\mathcal{G})$. Therefore, $\mathrm{QM}(\Gamma,\mathcal{G})$ cannot be hyperbolic.

Conversely, assume that Γ has a finite clique number and has no induced 4-cycle. Under these assumptions, observe the following.

Fact 2. The Hausdorff distance between two geodesics with the same endpoints is at most $2\mathrm{clique}(\Gamma)$.

Let α, β be two geodesics with the same endpoints and $x \in \alpha$ a vertex. It follows from [17, Lemma 2.114] that there exists an isometrically embedded copy of $[0,a] \times [0,a]$ in $\mathrm{QM}(\Gamma, \mathcal{G})$ with $(0,0) = x$ and $(a,a) \in \beta$. If $a > \mathrm{clique}(\Gamma)$, then there are at least two non-transverse hyperplanes separating $(0,0)$ and $(0,a)$. Because the hyperplanes in a maximal collection of pairwise non-transverse hyperplanes separating two vertices are successively tangent, it follows that there exist two tangent hyperplanes separating $(0,0)$ and $(0,a)$. Similarly, there must exist two tangent hyperplanes separating $(0,0)$ and $(a,0)$. Because these two pairs of hyperplanes are transverse, we deduce from Lemma 6.2 that Γ contains an induced 4-cycle, a contradiction. Thus, we must have $a \leq \mathrm{clique}(\Gamma)$, which implies that $d(p,\beta) \leq 2\mathrm{clique}(\Gamma)$, and which finally prove our fact.

Now, fix three vertices a, b, c and three geodesics $[a,b]$, $[b,c]$, $[a,c]$. Let (x,y,z) be the median triangle of (a,b,c) and let $[a,x]$, $[b,y]$, $[c,z]$, $[x,y]$, $[y,z]$, $[x,z]$ be geodesics. Given a vertex $p \in [a,b]$, it follows from Fact 2 that there exists some $q \in [a,x] \cup [x,y] \cup [y,b]$ such that $d(p,q) \leq 2\mathrm{clique}(\Gamma)$. If $q \in [a,x] \cup [y,b]$, then we can apply Fact 2 again and deduce that p lies in the $4\mathrm{clique}(\Gamma)$-neighbourhood of $[b,c] \cup [c,a]$. Otherwise, if $q \in [x,y]$, then it follows from Proposition 2.7 and Lemma 6.2 that $d(q,y) \leq \mathrm{clique}(\Gamma)$. Applying Fact 2 again, we conclude that p lies in the $5\mathrm{clique}(\Gamma)$-neighbourhood of $[b,c] \cup [c,a]$. ∎

Proof of Theorem 1.4. The first assertion follows from Corollary 1.2, from the fact that every vertex in $\mathrm{QM}(\Gamma, \mathcal{G})$ belongs to exactly $|V(\Gamma)|$ cliques, and from the observation that $\mathrm{QM}(\Gamma, \mathcal{G})$ has no cut-vertex as justified by Claim 5 and by the fact that the clique-link of a vertex in $\mathrm{QM}(\Gamma, \mathcal{G})$ is isomorphic to Γ.

The second assertion follows from Corollary 4.3 and Lemma 6.4. The third assertion follows from Theorem 1.3 and our next observation.

Claim 11. If the groups in $\mathcal{G}$ have solvable word problems, then $\mathrm{QM}(\Gamma, \mathcal{G})$ is constructible.

Algorithmically, we think of $\mathrm{QM}(\Gamma, \mathcal{G})$ whose vertices are words in $\Gamma\mathcal{G}$, two words representing the same vertex if they are equal in $\Gamma\mathcal{G}$ (or equivalently, if they admit identical graphically reduced representatives, which can be checked algorithmically if the groups in $\mathcal{G}$ have solvable word problems), and whose edges connect two words if one can be obtained from the other by right-multiplying by a non-trivial element in some vertex-group. The fact that $\mathrm{QM}(\Gamma, \mathcal{G})$ is geodesically constructible follows from Lemma 6.3 (and from the fact that graphically reduced representatives of an element of $\Gamma\mathcal{G}$ can be easily enumerated because the groups in $\mathcal{G}$ have solvable word problems). In order to show that $\mathrm{QM}(\Gamma, \mathcal{G})$ is also locally constructible, fix a non-trivial element $s_u \in G_u$ for every $u \in V(\Gamma)$. Then, for every word g, $\{gs_u \mid u \in V(\Gamma)\}$ is a set of neighbours of g such

that every clique containing g contains exactly one of these neighbours. We conclude that $\mathrm{QM}(\Gamma, \mathcal{G})$ is constructible, as desired. ∎

A. Translation spectrum of Morse elements

In this appendix, our goal is to clarify the construction of axes for isometries in graphs and to mention an application to Morse elements. Before stating our main result, let us emphasise that, despite the fact that distances in graphs are integers, stable translation lengths may be irrational. Even finitely generated groups may contain elements with irrational translation lengths [9]. As an elementary illustration of this phenomenon, let us mention the following example.

Example A.1. Fix an irrational number $r \in (0, 1)$, which we describe as the limit of rational numbers p_n/q_n. Let X denote the graph obtained from a set of isolated vertices indexed by $\mathbb{Z}$ by gluing a path of length p_n (which we will refer to as an arc) between the vertices i and $i + q_n$ for all $i \in \mathbb{Z}$ and $n \geq 1$. Up to adding $(p_0, q_0) := (1, 1)$ to our sequence if necessary, X is connected. The graph X has an obvious isometry g that shifts the vertices indexed by $\mathbb{Z}$. We claim that, if we choose our diophantine approximation of r carefully, then the stable translation length of g is r.

As a consequence of Dirichlet's theorem, we can assume that

$$\left| r - \frac{p_n}{q_n} \right| < \frac{1}{q_n^2} \quad \text{for every } n \geq 1.$$

Up to extracting a subsequence and applying a reflection centred at r, we can assume that $r < p_n/q_n$ for every $n \geq 0$. For convenience, set for every $n \geq 0$ the error $e_n := p_n/q_n - r$ of our approximation. It follows from the centred inequality above that the sequence $(e_n q_n)_{n \geq 0}$ tends to zero as $n \to +\infty$. Up to extracting a subsequence, we can assume that $(e_n q_n)_{n \geq 0}$ is decreasing. Also, since q_n has to tend to infinity as $n \to +\infty$, we can assume, up to extracting a subsequence, that $q_n \geq e_m q_m / 2r$ for all $n, m \geq 1$.

We claim that $d(0, q_n) = p_n$ in X for every $n \geq 1$. This will imply that the translation length of g is

$$\lim_{n \to +\infty} \frac{1}{q_n} d(0, g^{q_n} \cdot 0) = \lim_{n \to +\infty} \frac{1}{q_n} d(0, q_n) = \lim_{n \to +\infty} \frac{p_n}{q_n} = r,$$

as desired.

Let γ be a geodesic in X connecting 0 to q_n. We can decompose γ as a concatenation of arcs. Let $I \subset \mathbb{N}$ be a multi-subset such that the multi-set of the lengths of these arcs is $\{p_i \mid i \in \mathbb{N}\}$. (Here, we are using multi-sets because lengths are counted with multiplicity.) We have

$$d(0, q_n) = \mathrm{length}(\gamma) = \sum_{i \in I} p_i = \sum_{i \in I} (r q_i + e_i q_i).$$

Notice that, because the endpoints of our arcs are vertices indexed by integers that differ successively by $\pm q_i$ for some $i \in I$, we can write

$$q_n = \sum_{i \in I} \eta_i q_i, \quad \text{where } \varepsilon_i \in \{\pm 1\} \text{ for every } i \in I.$$

Set $N := \{i \in I \mid \eta_i = -1\}$. Then

$$d(0, q_n) = rq_n + 2r \sum_{i \in N} q_i + \sum_{i \in I} e_i q_i = p_n - e_n q_n + 2r \sum_{i \in N} q_i + \sum_{i \in I} e_i q_i.$$

If there is some $i \in I$ satisfying $i \leq n$, then we deduce from $e_i q_i \geq e_n q_n$ that $d(0, q_n) \geq p_n$. Now, assume that $i > n$ for every $i \in I$. From our decomposition of q_n above, it follows that N is non-empty. Then, fixing an arbitrary $i \in N$, we deduce from $2rq_i \geq e_n q_n$ that $d(0, q_n)$. Thus, we have proved that $d(0, q_n) \geq p_n$. But, by construction of X, there is an arc of length p_n connecting 0 to q_n, hence the desired equality $d(0, q_n) = p_n$.

Now, our goal is to prove the following criterion. Like most of the similar statement available in the literature, our argument is an elaboration of Delzant's argument [14]. It is also inspired by Bowditch's application to mapping class groups [7] and by its quantitative improvement from [4].

Theorem A.2. *Let X be a graph and $g \in \mathrm{Isom}(X)$ an isometry. Assume that all the orbits of g are infinite and that there exists a $\langle g \rangle$-invariant collection of bi-infinite geodesics $\mathcal{L}$ and a finite set of vertices $S \subset X$ that intersects every bi-infinite geodesic contained in $\bigcup \mathcal{L}$. Then there exists $n \leq |S|$ such that g cyclically permutes n pairwise disjoint bi-infinite geodesics. In particular, g^n admits an axis (i.e., there is a bi-infinite geodesic on which g acts non-trivially as a translation).*

Proof. Our first goal is to show that g permutes a finite collection of bi-infinite geodesics, regardless of its cardinality. We start by noticing the following.

Claim 12. The orbits of g are coarsely embedded.

Assume for contradiction that this is not the case. Fixing some $s \in S$, we find an $R \geq 0$ such that $I := \{i \in \mathbb{Z} \mid d(g^i s, s) \leq R\}$ is infinite. Fix a line $L \in \mathcal{L}$, it follows from the $\langle g \rangle$-invariance of $\mathcal{L}$ that $g^i S \cap L \neq \emptyset$ for every $i \in \mathbb{Z}$. Therefore, for every $i \in I$, there exists $s_i \in S$ such that $g^i s_i \in L$. Notice that

$$d(g^i s_i, s) \leq d(g^i s_i, g^i s) + d(g^i s, s) \leq \mathrm{diam}(S) + R =: R^+,$$

so $g^i s_i \in L \cap B(s, R^+)$. Since S and $L \cap B(s, R^+)$ are finite but I is infinite, we can find an $s', s'' \in S$ such that $g^i s' = s''$ for infinitely many i. If i and j are two distinct such indices, then it follows that g^{i-j} fixes s', contradicting the fact that the orbits of g are infinite. This concludes the proof of Claim 12.

Claim 12 implies that, up to replacing g with a large power, we can assume that $d(S, gS) > \mathrm{diam}(S)$. As a consequence, given a line $L \in \mathcal{L}$, the intersections $L \cap g^i S$

lie in pairwise disjoint subsegments. The picture to keep in mind is that the $g^i S$ are checkpoints that are successively crossed by the lines from $\mathcal{L}$.

Order the $\langle g \rangle$-orbits of edges of $\bigcup \mathcal{L}$ by an ordinal (i.e., a totally ordered set all of whose subsets have a minimal element, for instance $\mathbb{N}$ if $\mathcal{L}$ is countable). An edge is naturally labelled by the element of the ordinal given by the orbit containing our edge. This allows us to order the oriented paths in $\bigcup \mathcal{L}$ of a given length with the lexicographic order. Say that an oriented bi-infinite geodesic L contained in $\mathcal{L}$ is *special* if every subsegment of L is minimal among the oriented paths connecting its initial point to its terminal point. Notice that, for all $x, y \in \bigcup \mathcal{L}$ contained in a common connected component, there exists exactly one minimal oriented geodesic from x to y (here, minimal among all the oriented geodesics from x to y). However, it is not clear that special lines even exist.

Claim 13. There exists at least one special bi-infinite geodesic in $\bigcup \mathcal{L}$.

Fix a line $L \in \mathcal{L}$. We know that, for every $i \in \mathbb{Z}$, $g^i S$ intersects L. Fix a point $\ell_i \in g^i S \cap L$. We orient L so that, along L, ℓ_j comes after ℓ_i whenever $i < j$. For every $i \geq 1$, let $[\ell_{-i}, \ell_i]$ denote the unique minimal geodesic connecting ℓ_{-i} to ℓ_i and let L_i denote the line obtained from L by replacing the subsegment connecting ℓ_{-i} to ℓ_i with $[\ell_{-i}, \ell_i]$. Notice that $L_i \cap g^j S \neq \emptyset$ for all $i \geq 1$ and $j \in \mathbb{Z}$. Because S is finite, by a diagonal argument we can assume, up to extracting a subsequence from $(L_i)_{i \geq 1}$, that $L_i \cap g^j S$ is eventually constant for each j fixed, say to $s_j \in S$. By construction, for i sufficiently large compared to j, the subsegment of L_i connecting s_{-j} to s_j is the unique minimal geodesic connecting s_{-j} to s_j. Therefore, L_i converges to some bi-infinite geodesic L_∞, which is contained in $\bigcup \mathcal{L}$ and which is special. This concludes the proof of Claim 13.

Claim 14. There exist at most $|S|^2$ special bi-infinite geodesics in $\bigcup \mathcal{L}$.

Assume for contradiction that there exist more than $|S|^2$ special bi-infinite geodesics, say $L_1, \dots, L_r$. Because they are distinct, we can find a large ball B such that the intersections $L_i \cap B$ are pairwise distinct. Fix an $k \geq 1$ sufficiently large so that $g^{-k} S$ and $g^k S$ are disjoint from B. Because $r > |S|^2$, we can find two distinct indices $1 \leq i, j \leq r$ such that $L_i \cap L_j \cap g^{-k} S \neq \emptyset$ and $L_i \cap L_j \cap g^k S \neq \emptyset$. Let s_{-k} and s_k be two points in these intersections. Then, because L_i and L_j are special, they must agree between s_{-k} and s_k with the unique minimal geodesic connecting s_{-k} to s_k, which implies that L_i and L_j must agree in B, a contradiction. This proves Claim 14.

Because the action of g on $\bigcup \mathcal{L}$ preserves the set of special bi-infinite geodesics, we conclude from Claim 14 that there exists a collection of $\leq |S|^2$ bi-infinite geodesics that is preserved by $\langle g \rangle$. This already implies that some power of g admits an axis, but we do not have yet the optimal control on the size of this power.

Claim 15. If $\mathcal{M}$ is a $\langle g \rangle$-invariant collection of bi-infinite geodesics in $\bigcup \mathcal{L}$ of minimal (finite) size, then $\mathcal{M}$ is a collection of pairwise disjoint geodesics.

First of all, notice that such a collection $\mathcal{M}$ must have size $\leq |S|$, since all the geodesics of $\mathcal{M}$ must intersect S but are pairwise disjoint. Therefore, proving Claim 15 will conclude

the proof of our theorem. During its proof, we will use repeatedly the following elementary observation.

Fact 3. Let W be a graph and α a bi-infinite geodesic. Decompose α as a union of two rays α_- and α_+ having the same origin, say $o \in \alpha$. If ρ is a geodesic ray starting from o and intersecting α_+ infinitely many times, then $\alpha_- \cup \rho$ is a bi-infinite geodesic.

Let $r_0 = o, r_1, r_2, \dots$ denote the successive vertices of ρ that belong to α_+. Then $\alpha_- \cup \rho$ can be described as the limit of the lines obtained from α by replacing successively the subsegments connecting r_i to r_{i+1} with the subsegment of ρ having the same endpoints. It is clear that these operations yield a geodesic at each step. This implies that $\alpha_- \cup \rho$ is a geodesic, concluding the proof of Fact 3.

Now, assume for contradiction that our collection $\mathcal{M}$ contains two intersecting geodesics. By minimality, $\langle g \rangle$ acts transitively on $\mathcal{M}$, so there must exist $L \in \mathcal{M}$ and $1 \leq m < |\mathcal{M}|$ such that $L \cap g^m L \neq \emptyset$. Fix a vertex $x \in L \cap g^m L$. Notice that $\langle g^{|\mathcal{M}|} \rangle x \subset L \cap g^m L$ since $g^{|\mathcal{M}|}$ stabilises each line of $\mathcal{M}$. We have the following configuration:

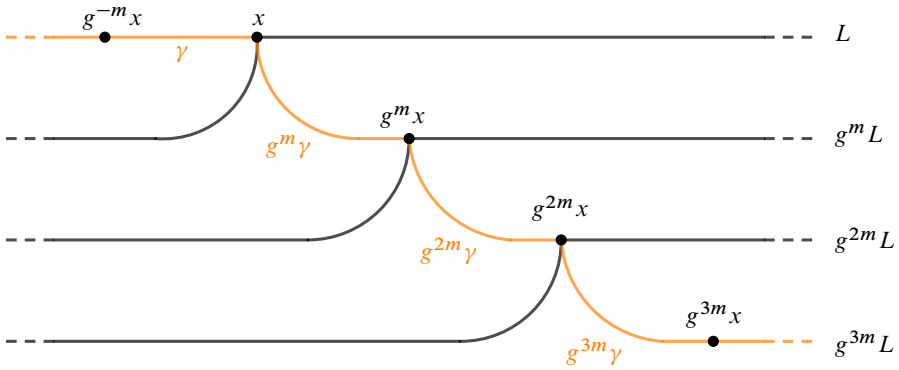

where γ denotes the subsegment of L connecting $g^{-m}x$ to x. By applying Fact 3 iteratively, we deduce that, for every $k \geq 1$, there exists a bi-infinite geodesic passing through $\gamma, g^m\gamma, \dots, g^{km}\gamma$. This implies that the concatenation $\bigcup_{i \in \mathbb{Z}} g^{im}\gamma$ yields a bi-infinite geodesic L'. By construction, g^m stabilises L', so $\mathcal{M}' := \{L', gL', \dots, g^{m-1}L'\}$ is a new $\langle g \rangle$-invariant set of lines of size $m < |\mathcal{M}|$, contradicting our choice of $\mathcal{M}$. ∎

Generalising the fact that elements in hyperbolic groups have rational translation lengths, we can deduce from Theorem A.2 that *Morse elements* in arbitrary finitely generated groups have rational translation lengths. Recall that, given a metric space X and a map $M : \mathbb{R}_+^2 \to \mathbb{R}_+$ (referred to as the *Morse gauge*), an isometry $g \in \mathrm{Isom}(X)$ is *M-Morse* if there exists a point $o \in X$ such that the orbit $\langle g \rangle \cdot o$ is quasi-isometrically embedded and such that, for all constants $A > 0$, $B \geq 0$ and for every (A, B)-quasigeodesic γ connecting two points in $\langle g \rangle \cdot o$, γ stays in the $M(A, B)$-neighbourhoods of $\langle g \rangle \cdot o$. An isometry is *Morse* if it is M-Morse for some Morse gauge M.

Corollary A.3. *For every integer $N \geq 1$ and every Morse gauge M, there exists a constant $D \geq 0$ such that the following holds. Let X be a graph all of whose vertices have degree $\leq N$. For every M-Morse isometry g, there exists $n \leq D$ such that g^n has an axis in X.*

Proof. Let $g \in \mathrm{Isom}(X)$ be an M-Morse isometry. Let $o \in X$ denote a vertex as given by the definition above. For every $k \geq 1$, fix a geodesic $[g^{-k}o, g^k o]$ connecting $g^{-k}o$ and $g^k o$. By local finiteness, the $[g^{-k}o, g^k o]$ subconverge to a bi-infinite geodesic L. Notice that L has to stay in the $M(1,0)$-neighbourhood of $\langle g \rangle \cdot o$. Let $\mathcal{L}$ denote the set of all the bi-infinite geodesics contained in the $M(1,0)$-neighbourhood of $\langle g \rangle \cdot o$. Clearly, $\mathcal{L}$ is $\langle g \rangle$-invariant. It is also clear that every bi-infinite geodesic contained in $\bigcup \mathcal{L}$ has to intersect the ball $B(o, M(1,0))$, which is finite with size bounded above in function of N and M. The desired conclusion follows from Theorem A.2. ∎

As an immediate consequence, we get the following statement (which improves [47, Theorem K]).

Corollary A.4. *Let G be a group and $S \subset G$ a finite generating set. For every Morse gauge M,*

$$\{\tau_S(g) \text{ where } g \text{ is } M\text{-Morse}\}$$

is a discrete set of rational numbers.

Acknowledgements. I am grateful to Hyungryul Baik and Donggyun Seo for interesting discussions regarding their work [4], and to Sam Shepherd for interesting discussions about his recent preprint [49]. I also thank the anonymous referee for their comments.

References

[1] C. Abbott, M. Hagen, H. Petyt, and A. Zalloum, Uniform undistortion from barycentres, and applications to hierarchically hyperbolic groups. 2023, arXiv:2305.09742v1

[2] E. Alibegović, Translation lengths in Out(F_n). *Geom. Dedicata* **92** (2002), 87–93 (Dedicated to J. Stallings on the occasion of his 65th birthday) Zbl 1041.20024 MR 1934012

[3] G. N. Arzhantseva, C. H. Cashen, D. Gruber, and D. Hume, Negative curvature in graphical small cancellation groups. *Groups Geom. Dyn.* **13** (2019), no. 2, 579–632 Zbl 1512.20128 MR 3950644

[4] H. Baik, D. Seo, and H. Shin, On the finiteness property of hyperbolic simplicial actions: the right-angled Artin groups and their extension graphs. *Geom. Dedicata* **217** (2023), no. 1, article no. 7, 29 pp. Zbl 07621815 MR 4505009

[5] H.-J. Bandelt, H. M. Mulder, and E. Wilkeit, Quasi-median graphs and algebras. *J. Graph Theory* **18** (1994), no. 7, 681–703 Zbl 0810.05057 MR 1297190

[6] M. Bestvina, Non-positively curved aspects of Artin groups of finite type. *Geom. Topol.* **3** (1999), 269–302 Zbl 0998.20034 MR 1714913

[7] B. H. Bowditch, Tight geodesics in the curve complex. *Invent. Math.* **171** (2008), no. 2, 281–300 Zbl 1185.57011 MR 2367021

[8] N. Brady, Finite subgroups of hyperbolic groups. *Internat. J. Algebra Comput.* **10** (2000), no. 4, 399–405 Zbl 1010.20030 MR 1776048

[9] G. R. Conner, A class of finitely generated groups with irrational translation numbers. *Arch. Math. (Basel)* **69** (1997), no. 4, 265–274 Zbl 0901.20023 MR 1466819

[10] G. R. Conner, Properties of translation numbers in nilpotent groups. *Comm. Algebra* **26** (1998), no. 4, 1069–1080 Zbl 0902.20015 MR 1612184

[11] G. R. Conner, Discreteness properties of translation numbers in solvable groups. *J. Group Theory* **3** (2000), no. 1, 77–94 Zbl 0956.20039 MR 1736519

[12] G. R. Conner, Translation numbers of groups acting on quasiconvex spaces. In *Computational and geometric aspects of modern algebra (Edinburgh, 1998)*, pp. 28–38, London Math. Soc. Lecture Note Ser. 275, Cambridge University Press, Cambridge, 2000 Zbl 0986.20040 MR 1776764

[13] M. Coornaert, T. Delzant, and A. Papadopoulos, *Géométrie et théorie des groupes. Les groupes hyperboliques de Gromov*. Lecture Notes in Math. 1441, Springer, Berlin, 1990 Zbl 0727.20018 MR 1075994

[14] T. Delzant, Sous-groupes distingués et quotients des groupes hyperboliques. *Duke Math. J.* **83** (1996), no. 3, 661–682 Zbl 0852.20032 MR 1390660

[15] T. Elsner, Isometries of systolic spaces. *Fund. Math.* **204** (2009), no. 1, 39–55 Zbl 1231.20037 MR 2507689

[16] B. Farb, A. Lubotzky, and Y. Minsky, Rank-1 phenomena for mapping class groups. *Duke Math. J.* **106** (2001), no. 3, 581–597 Zbl 1025.20023 MR 1813237

[17] A. Genevois, Cubical-like geometry of quasi-median graphs and applications to geometric group theory. 2017, arXiv:1712.01618v1

[18] A. Genevois, Embeddings into Thompson's groups from quasi-median geometry. *Groups Geom. Dyn.* **13** (2019), no. 4, 1457–1510 Zbl 1511.20159 MR 4033512

[19] A. Genevois, Hyperbolicities in CAT(0) cube complexes. *Enseign. Math.* **65** (2019), no. 1–2, 33–100 Zbl 1472.20090 MR 4057355

[20] A. Genevois, On the geometry of van Kampen diagrams of graph products of groups. 2019, arXiv:1901.04538v2

[21] A. Genevois, Automorphisms of graph products of groups and acylindrical hyperbolicity. [v1] 2018, [v4] 2022, arXiv:1807.00622v4, to appear in *Mem. Amer. Math. Soc.*

[22] A. Genevois, Quasi-isometrically rigid subgroups in right-angled Coxeter groups. *Algebr. Geom. Topol.* **22** (2022), no. 2, 657–708 Zbl 07570603 MR 4464462

[23] A. Genevois, Special cube complexes revisited: a quasi-median generalization. *Canad. J. Math.* **75** (2023), no. 3, 743–777 Zbl 1514.20151 MR 4586831

[24] A. Genevois and A. Martin, Automorphisms of graph products of groups from a geometric perspective. *Proc. Lond. Math. Soc. (3)* **119** (2019), no. 6, 1745–1779 Zbl 1461.20008 MR 4295519

[25] A. Genevois and R. Tessera, A note on morphisms to wreath products. [v1] 2021, [v2] 2022, arXiv:2110.09822v2

[26] A. Genevois and R. Tessera, Asymptotic geometry of lamplighters over one-ended groups. *Invent. Math.* **238** (2024) 1–67 Zbl 7921928 MR 4794592

[27] S. M. Gersten and H. B. Short, Rational subgroups of biautomatic groups. *Ann. of Math. (2)* **134** (1991), no. 1, 125–158 Zbl 0744.20035 MR 1114609

[28] E. Green, *Graph products of groups*. PhD thesis, University of Leeds, 1990

[29] M. Gromov, Hyperbolic groups. In *Essays in group theory*, pp. 75–263, Math. Sci. Res. Inst. Publ. 8, Springer, New York, 1987 Zbl 0634.20015 MR 0919829

[30] T. Haettel and D. Osajda, Locally elliptic actions, torsion groups, and nonpositively curved spaces. [v1] 2021, [v2] 2022, arXiv:2110.12431v2

[31] M. F. Hagen, Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topol.* **7** (2014), no. 2, 385–418 Zbl 1348.20044 MR 3217625

[32] F. Haglund, Isometries of $CAT(0)$ cube complexes are semi-simple. *Ann. Math. Qué.* **47** (2023), no. 2, 249–261 Zbl 1528.20069 MR 4645691

[33] T. Hsu and D. T. Wise, On linear and residual properties of graph products. *Michigan Math. J.* **46** (1999), no. 2, 251–259 Zbl 0962.20016 MR 1704150

[34] I. Kapovich, Small cancellation groups and translation numbers. *Trans. Amer. Math. Soc.* **349** (1997), no. 5, 1851–1875 Zbl 0885.20023 MR 1357396

[35] S.-h. Kim and T. Koberda, Embedability between right-angled Artin groups. *Geom. Topol.* **17** (2013), no. 1, 493–530 Zbl 1278.20049 MR 3039768

[36] S.-H. Kim and T. Koberda, The geometry of the curve graph of a right-angled Artin group. *Internat. J. Algebra Comput.* **24** (2014), no. 2, 121–169 Zbl 1342.20042 MR 3192368

[37] J. P. Leasure, *Geodesics in the complex of curves of a surface*. Ph.D. thesis, The University of Texas at Austin, 2002 MR 2705485

[38] S. J. Lee, Garside groups are strongly translation discrete. *J. Algebra* **309** (2007), no. 2, 594–609 Zbl 1155.20038 MR 2303195

[39] E.-K. Lee and S. J. Lee, Translation numbers in a Garside group are rational with uniformly bounded denominators. *J. Pure Appl. Algebra* **211** (2007), no. 3, 732–743 Zbl 1150.20022 MR 2344226

[40] E.-K. Lee and S.-J. Lee, Some power of an element in a Garside group is conjugate to a periodically geodesic element. *Bull. Lond. Math. Soc.* **40** (2008), no. 4, 593–603 Zbl 1171.20026 MR 2438075

[41] H. A. Masur and Y. N. Minsky, Geometry of the complex of curves I: Hyperbolicity. *Invent. Math.* **138** (1999), no. 1, 103–149 Zbl 0941.32012 MR 1714338

[42] G. A. Niblo, The singularity obstruction for group splittings. *Topology Appl.* **119** (2002), no. 1, 17–31 Zbl 0991.20024 MR 1881707

[43] G. A. Noskov, Asymptotic behavior of word metrics on Coxeter groups. *Doc. Math.* **16** (2011), 373–398 Zbl 1241.20046 MR 2804516

[44] A. Poleksić, Groups acting on quasiconvex spaces and translation numbers. *Proc. Amer. Math. Soc.* **128** (2000), no. 11, 3177–3183 Zbl 0986.20041 MR 1694875

[45] A. Poleksić, Translation numbers in negatively curved groups. *Topology Appl.* **102** (2000), no. 2, 203–205 Zbl 0961.20035 MR 1741488

[46] M. Roller, Poc sets, median algebras and group actions: an extended study of Dunwoody's construction and Sageev's theorem. 2016, arXiv:1607.07747v1

[47] J. Russell, D. Spriano, and H. C. Tran, The local-to-global property for Morse quasi-geodesics. *Math. Z.* **300** (2022), no. 2, 1557–1602 Zbl 1515.20247 MR 4363788

[48] K. J. Shackleton, Tightness and computing distances in the curve complex. *Geom. Dedicata* **160** (2012), 243–259 Zbl 1262.57020 MR 2970053

[49] S. Shepherd, A cubulation with no factor system. [v1] 2022, [v2] 2023, arXiv:2208.10421v2, to appear in *Algebr. Geom. Topol.*

[50] E. L. Swenson, Hyperbolic elements in negatively curved groups. *Geom. Dedicata* **55** (1995), no. 2, 199–210 Zbl 0834.20041 MR 1334214

[51] M. Valiunas, Acylindrical hyperbolicity of groups acting on quasi-median graphs and equations in graph products. *Groups Geom. Dyn.* **15** (2021), no. 1, 143–195 Zbl 07362826 MR 4235751

[52] R. C. H. Webb, Combinatorics of tight geodesics and stable lengths. *Trans. Amer. Math. Soc.* **367** (2015), no. 10, 7323–7342 Zbl 1339.57029 MR 3378831

Received 29 September 2022.

Anthony Genevois
Institut Mathématiques Alexander Grothendieck, University of Montpellier, Place Eugène Bataillon, 34090 Montpellier, France; anthony.genevois@umontpellier.fr