Algebraic & Geometric Topology 20 (2020) 49–134

Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups

ANTHONY GENEVOIS

The main technical result of this paper is to characterize the contracting isometries of a CAT(0) cube complex without any assumption on its local finiteness. Afterwards, we introduce the combinatorial boundary of a CAT(0) cube complex, and we show that contracting isometries are strongly related to isolated points at infinity, when the complex is locally finite. This boundary turns out to appear naturally in the context of Guba and Sapir's diagram groups, and we apply our main criterion to determine precisely when an element of a diagram group induces a contracting isometry on the associated Farley cube complex. As a consequence, in some specific cases, we are able to deduce a criterion to determine precisely when a diagram group is acylindrically hyperbolic.

20F65, 20F67

1 Introduction

A classical strategy in geometric group theory is to study groups by looking for actions on metric spaces with some "negatively curved behavior". In this article, we are interested in *contracting isometries*, which may be defined in arbitrary metric spaces and which have typically the same behavior as the isometries of a hyperbolic space.

Published: 23 February 2020 DOI: 10.2140/agt.2020.20.49

Formally, given a metric space X, an isometry $g \in X$ is *contracting* if

- g is *loxodromic*, i.e. there exists $x_0 \in X$ such that $n \mapsto g^n \cdot x_0$ defines a quasi-isometric embedding $\mathbb{Z} \to X$;
- if $C_g = \{g^n \cdot x_0 \mid n \in \mathbb{Z}\}$, then the diameter of the nearest-point projection of any ball disjoint from C_g onto C_g is uniformly bounded.

For instance, any loxodromic isometry of a Gromov-hyperbolic space is contracting. In order to justify that the existence of a contracting isometry confers to the group we are looking at a "hyperbolic behavior", let us mention the following statement, which is a consequence of Osin [26, Theorem 1.2] combined with Bestvina, Bromberg and Fujiwara [2, Theorem H] or Sisto [30, Theorem 1.5]:

Theorem 1.1 *If a group acts properly discontinuously on a geodesic space with at least one contracting isometry, then it is either virtually cyclic or acylindrically hyperbolic.*

A group G is *acylindrically hyperbolic* if it admits an action on a Gromov-hyperbolic space which is both *nonelementary* (i.e. with an infinite limit set) and *acylindrical*: for every $d \geq 0$, there exist some constants $R, N \geq 0$ such that, for every $x, y \in X$,

$$d(x,y) \geq R \implies \#\{g \in G \mid d(x,gx), d(y,gy) \leq d\} \leq N.$$

The class of acylindrically hyperbolic groups, introduced in [26], provides a powerful framework which allows us to study many groups of interest and which generalizes Gromov's hyperbolic groups and more generally relatively hyperbolic groups. We refer to [26, Appendix] for more information on acylindrically hyperbolic groups.

Here we focus on a specific class of geodesic spaces: the CAT(0) *cube complexes*, i.e. simply connected cellular complexes obtained by gluing cubes of different dimensions by isometries along their faces so that the link of every vertex is a simplicial flag complex. See Section 2 for more details. The motivation is twofold. Firstly, the combinatorics of hyperplanes available in CAT(0) cube complexes provides a powerful tool which helps us to answer questions about cube complexes (compared to arbitrary metric spaces). And, secondly, there exist many groups of interest which act on CAT(0) cube complexes, providing a large collection of potential applications.

So the main subject of this article is the study of contracting isometries in CAT(0) cube complexes. Previous works in this direction include Behrstock and Charney [1] and Charney and Sultan [6]. However, we emphasize that here we endow our cube complexes with their combinatorial metrics (i.e. the distances in their one-skeletons) and not with their CAT(0) metrics. This point of view usually allows us to work in a more general framework, typically without local finiteness conditions or without restrictions on the dimension.

The first question we are interested in is the following: when is an isometry of an arbitrary CAT(0) cube complex contracting? The answer we give to this question is the main technical result of our paper. Our criterion essentially deals with hyperplane configurations in the set $\mathcal{H}(\gamma)$ of the hyperplanes intersecting a *combinatorial axis* γ of a given loxodromic isometry g. We refer to Section 3 for precise definitions.

Theorem 1.2 *Let X be a* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry with a combinatorial axis γ. The following statements are equivalent*:

(i) *g is a contracting isometry*;

(ii) *joins of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{H}(\gamma)$ are uniformly thin*;

(iii) *there exists $C \geq 1$ such that*
- *$\mathcal{H}(\gamma)$ does not contain C pairwise transverse hyperplanes*,
- *any grid of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{H}(\gamma)$ is C–thin*;

(iv) *g skewers a pair of well-separated hyperplanes.*

It is worth noticing that the equivalence (i) $\Longleftrightarrow$ (iv) generalizes a similar criterion established in [6, Theorem 4.2], where the cube complex is supposed to be uniformly locally finite.

The second objective of this paper is to apply our characterization of contracting isometries to the study of Guba and Sapir's *diagram groups*. Loosely speaking, diagram groups are "two-dimensional free groups": in the same way that free groups are defined by concatenating and reducing words, diagram groups are defined by concatenating and reducing some two-dimensional objects, called *semigroup diagrams*. See Section 5.1 for precise definitions. Although these two classes of groups turn out to be quite different, the previous analogy can be pushed further. On the one hand, free groups act on their canonical Cayley graphs, which are simplicial trees; on the other hand, diagram groups act on the natural Cayley graphs of their associated groupoids, which are CAT(0) cube complexes, called *Farley complexes*. Moreover, in the same way that the boundary of a free group may be thought of as a set of infinite reduced words, it is natural to construct a boundary of a Farley cube complex as a set of infinite reduced diagrams.

This analogy suggests the introduction of a new boundary for general CAT(0) cube complexes, which we call the *combinatorial boundary*. In the definition below, if r is a combinatorial ray, $\mathcal{H}(r)$ denotes the set of the hyperplanes which intersect r.

Definition 1.3 Let X be a CAT(0) cube complex. Its *combinatorial boundary* is the poset $(\partial^c X, \prec)$, where $\partial^c X$ is the set of the combinatorial rays modulo the relation $r_1 \sim r_2$ if $\mathcal{H}(r_1) =_a \mathcal{H}(r_2)$ and where the partial order $\prec$ is defined by $r_1 \prec r_2$ whenever $\mathcal{H}(r_1) \subset_a \mathcal{H}(r_2)$.

Here, we used the following notation: given two sets X and Y, we say that X and Y are *almost equal*, denoted by $X =_a Y$, if the symmetric difference between X and Y is finite; we say that X is *almost included* into Y, denoted by $X \subset_a Y$, if X is almost equal to a subset of Y. In fact, this boundary is not completely new, since it admits strong relations with other boundaries; see the appendix for more details.

Now, it turns out that determining whether an isometry is contracting can be read naturally at infinity on the combinatorial boundary. Indeed, Theorem 1.2 implies the following statement:

Theorem 1.4 *Let X be a locally finite* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry with a combinatorial axis γ. Then g is a contracting isometry if and only if $\gamma(+\infty)$ is isolated in $\partial^c X$.*

A point in the combinatorial boundary is *isolated* if it is not $\prec$–comparable with any other point. Thus, the existence of a contracting isometry implies the existence of an isolated point in the combinatorial boundary. Of course, the converse cannot hold without an additional hypothesis on the action, but in some specific cases we are able to prove partial converses. For instance:

Theorem 1.5 *Let G be a group acting on a locally finite* CAT(0) *cube complex X with finitely many orbits of hyperplanes. Then $G \curvearrowright X$ contains a contracting isometry if and only if $\partial^c X$ has an isolated vertex.*

Theorem 1.6 *Let G be a countable group acting on a locally finite* CAT(0) *cube complex X. Suppose that the action $G \curvearrowright X$ is minimal* (i.e. X does not contain a proper G–invariant combinatorially convex subcomplex) *and G does not fix a point of $\partial^c X$. Then G contains a contracting isometry if and only if $\partial^c X$ contains an isolated point.*

These statements should be compared with Caprace and Sageev [5]. Notice that here we are not assuming that our cube complexes are finite-dimensional.

Now, let us focus on diagram groups, and more precisely on their elements inducing contracting isometries on the corresponding Farley complexes. If g is an element of a diagram group, which is *absolutely reduced* (i.e. the product g^n is reduced for every $n \geq 1$), let g^∞ denote the infinite diagram obtained by concatenating infinitely many copies of g. (See Section 5.2 for a precise description.) Then, thanks to Theorem 1.4 and a precise description of the combinatorial boundaries of Farley cube complexes, we are able to deduce the following criterion:

Theorem 1.7 *Let G be a diagram group and $g \in G \setminus \{1\}$ be absolutely reduced. Then g is a contracting isometry of the associated Farley complex if and only if the following two conditions are satisfied*:

- *g^∞ does not contain any infinite proper prefix*;
- *for any infinite reduced diagram Δ containing g^∞ as a prefix*, *all but finitely many cells of Δ belong to g^∞.*

Of course, although it is sufficient for elements with few cells, this criterion may be difficult to apply in practice, because we cannot draw the whole g^∞. This is why we give a more "algorithmic" criterion in Section 5.3.

Thus, we are able to determine precisely when a given element of a diagram group is a contracting isometry. In general, if there exist such isometries, it is not too difficult to find one of them. Otherwise, it is possible to apply Theorem 1.5 or Theorem 1.6; we emphasize that Theorem 1.5 may be particularly useful since diagram groups often act on their cube complexes with finitely many orbits of hyperplanes. Conversely, we are able to state that no contracting isometry exists if the combinatorial boundary does not contain isolated points; see for instance Example 5.26.

Therefore, we get powerful tools to pursue the cubical study of negatively curved properties of diagram groups we initialized in [10]. In view of Theorem 1.1, our study of contracting isometries leads to (at least partial) answers to the following question:

Question 1.8 When is a diagram group acylindrically hyperbolic?

Section 5.4 provides families of acylindrically hyperbolic diagram groups which we think to be of interest. On the other hand, a given group may be represented as a diagram group in different ways, and an acylindrically hyperbolic group (eg a nonabelian free group) may be represented as a diagram group so that there is no contracting isometry on the associated cube complex. We do not know if it is always possible to a find a "good" representation.

Nevertheless, a good class of diagram groups, where the problem mentioned above does not occur, corresponds to the case where the action on the associated complex is cocompact. Using the notation introduced below, they correspond exactly to the diagram groups $D(\mathcal{P}, w)$ where the class of w modulo the semigroup $\mathcal{P}$ is finite. We call them *cocompact diagram groups*. Focusing on this class of groups, we prove (see Theorem 5.48 for a precise statement):

Theorem 1.9 *A cocompact diagram group decomposes naturally as a direct product of a finitely generated free abelian group and finitely many acylindrically hyperbolic diagram groups.*

On the other hand, acylindrically hyperbolic groups cannot split as a direct product of two infinite groups (see [26, Corollary 7.3]), so we deduce a complete answer of Question 1.8 in the cocompact case:

Corollary 1.10 *A cocompact diagram group is acylindrically hyperbolic if and only if it is not cyclic and it does not split as a nontrivial direct product.*

Notice that this statement is no longer true without the cocompact assumption. Indeed, Thompson's group F and the lamplighter group $\mathbb{Z} \wr \mathbb{Z}$ are diagram groups which are not acylindrically hyperbolic (since they do not contain a nonabelian free group) and they do not split nontrivially as a direct product. Another consequence of Theorem 1.9 is that a cocompact diagram group either is free abelian or has a quotient which is acylindrically hyperbolic. Because acylindrically hyperbolic groups are *SQ–universal* (i.e. any countable group can be embedded into a quotient of the given group; see [26, Theorem 8.1]), we get the following dichotomy:

Corollary 1.11 *A cocompact diagram group is either free abelian or SQ–universal.*

In our opinion, cocompact diagram groups turn out to be quite similar to (finitely generated) right-angled Artin groups. In this context, Theorem 1.9 should be compared to the similar but more general [1, Theorem 5.2] (it must be noted that, although the statement is correct, the proof contains a mistake; see Minsyan and Osin [24, Remark 6.21]). Compare also [1, Lemma 5.1] with Proposition 5.53. Notice that there exist right-angled Artin groups which are not (cocompact) diagram groups (see Guba and Sapir [14, Theorem 30]), and conversely there exist cocompact diagram groups which are not right-angled Artin groups (see Example 5.46).

The paper is organized as follows. In Section 2, we give the prerequisites on CAT(0) cube complexes needed in the rest of the paper. Section 3 is essentially dedicated to the proof of Theorem 1.2, and in Section 4, we introduce combinatorial boundaries of CAT(0) cube complexes and we prove Theorem 1.4, as well as Theorems 1.5 and 1.6. Finally, in Section 5, we introduce diagram groups and we apply the results of the previous sections to deduce the various statements mentioned above. We added an appendix to compare the combinatorial boundary with other known boundaries of CAT(0) cube complexes.

Acknowledgement I am grateful to Jean Lécureux for having suggested to me a link between the combinatorial boundary and the Roller boundary; to Elia Fioravanti for having pointed out to me a few mistakes; to my advisor, Peter Haïssinsky, for all our discussions; and to the anonymous referee for comments which led to several improvements of the article.

2 Preliminaries

A *cube complex* is a CW complex constructed by gluing together cubes of arbitrary (finite) dimension by isometries along their faces. Furthermore, it is *nonpositively curved* if the link of any of its vertices is a simplicial *flag* complex (i.e. $n+1$ vertices span a n–simplex if and only if they are pairwise adjacent), and CAT(0) if it is nonpositively curved and simply connected. See [3, page 111] for more information.

Alternatively, CAT(0) cube complexes may be described by their 1–skeletons. Indeed, Chepoï notices in [7] that the class of graphs appearing as 1–skeletons of CAT(0) cube complexes coincides with the class of *median graphs*, which we now define.

Let Γ be a graph. If $x,y,z \in \Gamma$ are three vertices, a vertex m is called a *median point of* x*,* y *and* z whenever

$$\begin{aligned} d(x,y) &= d(x,m)+d(m,y), \\ d(x,z) &= d(x,m)+d(m,z), \\ d(y,z) &= d(y,m)+d(m,z). \end{aligned}$$

Note that, for any geodesics $[x,m]$, $[y,m]$ and $[z,m]$, the concatenations $[x,m]\cup[m,y]$, $[x,m]\cup[m,z]$ and $[y,m]\cup[m,z]$ are also geodesics; furthermore, if $[x,y]$, $[y,z]$ and $[x,z]$ are geodesics, then any vertex of $[x,y]\cap[y,z]\cap[x,z]$ is a median point of x, y and z.

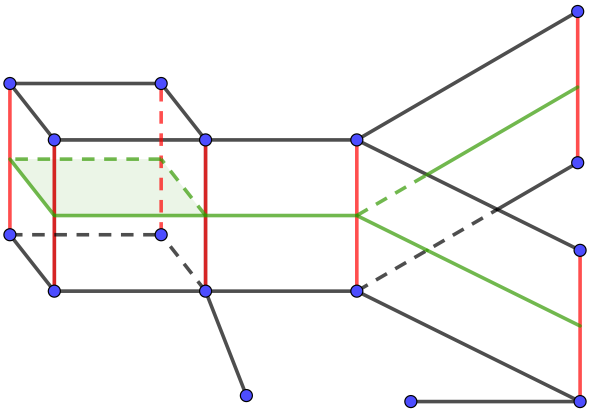

Figure 1: A hyperplane (in red) and its geometric realization (in green).

The graph Γ is *median* if every triple (x,y,z) of pairwise distinct vertices admits a unique median point, denoted by $m(x,y,z)$.

Theorem 2.1 [7, Theorem 6.1] *A graph is median if and only if it is the* 1*–skeleton of a* $\mathrm{CAT}(0)$ *cube complex.*

A fundamental feature of cube complexes is the notion of *hyperplane*. Let X be a nonpositively curved cube complex. Formally, a *hyperplane* J is an equivalence class of edges, where two edges e and f are equivalent whenever there exists a sequence of edges $e=e_0, e_1, \dots, e_{n-1}, e_n = f$, where e_i and e_{i+1} are parallel sides of some square in X. Notice that a hyperplane is uniquely determined by one of its edges, so, if $e \in J$, we say that J is the *hyperplane dual to* e. Geometrically, a hyperplane J is rather thought of as the union of the *midcubes* transverse to the edges belonging to J. See Figure 1. The *neighborhood* $N(J)$ of a hyperplane J is the smallest subcomplex of X containing J, i.e. the union of the cubes intersecting J. In the following, $\partial N(J)$ will denote the union of the cubes of X contained in $N(J)$ but not intersecting J, and $X \backslash\backslash J = (X \backslash N(J)) \cup \partial N(J)$. Notice that $N(J)$ and $X \backslash\backslash J$ are subcomplexes of X.

Theorem 2.2 [28, Theorem 4.10] *Let* X *be a* $\mathrm{CAT}(0)$ *cube complex and* J *a hyperplane. Then* $X \backslash\backslash J$ *has exactly two connected components.*

The two connected components of $X \backslash\backslash J$ will be referred to as the *halfspaces* associated to the hyperplane J.

Distances ℓ_p There exist several natural metrics on a CAT(0) cube complex. For example, for any $p \in (0, +\infty)$, the ℓ_p–norm defined on each cube can be extended to a distance defined on the whole complex, the *ℓ_p–metric*. Usually, the ℓ_1–metric is referred to as the *combinatorial distance* and the ℓ_2–metric as the CAT(0) *distance*. Indeed, a CAT(0) cube complex endowed with its CAT(0) distance turns out to be a CAT(0) space [22, Theorem C.9], and the combinatorial distance between two vertices corresponds to the graph metric associated to the 1–skeleton $X^{(1)}$. In particular, *combinatorial geodesics* are edge-paths of minimal length, and a subcomplex is *combinatorially convex* if it contains any combinatorial geodesic between two of its points.

Although Section 3 deals with arbitrary CAT(0) cube complexes, our cube complexes in Section 4 will be required to be *complete*. Roughly speaking, a complete CAT(0) cube complex may be infinite-dimensional but it has to remain finite-dimensional locally. More precisely:

Definition 2.3 A CAT(0) cube complex is *complete* if it does not contain an infinite increasing sequence of cubes.

The terminology is justified by [22, Theorem A.6], which shows that a CAT(0) cube complex is complete if and only if it is complete with respect to the CAT(0)–metric.

It turns out that the combinatorial metric and the hyperplanes are strongly linked together: the combinatorial distance between two vertices corresponds exactly to the number of hyperplanes separating them [18, Theorem 2.7], and:

Theorem 2.4 [18, Corollary 2.16] *Let X be a* CAT(0) *cube complex and J a hyperplane. The two components of $X \setminus\setminus J$ are combinatorially convex, as well as the components of $\partial N(J)$.*

This result is particularly useful when it is combined with the well-known Helly property, which follows from [12] (see also [27, Theorem 2.2; 18, Corollary 2.22]):

Theorem 2.5 *If $\mathcal{C}$ is a finite collection of pairwise intersecting combinatorially convex subcomplexes, then the intersection $\bigcap_{C \in \mathcal{C}} C$ is nonempty.*

Convention We emphasize that, in the whole article, our CAT(0) cube complexes are endowed with the combinatorial metric. So, when dealing with distances and geodesics, these always refer to combinatorial distances and combinatorial geodesics.

Combinatorial projection In CAT(0) spaces, and so in particular in CAT(0) cube complexes with respect to the CAT(0) distance, the existence of a well-defined projection onto a given convex subspace provides a useful tool. Similarly, with respect to the combinatorial distance, it is possible to introduce a *combinatorial projection* onto a combinatorially convex subcomplex. The notion of combinatorial projection has been introduced independently in several different places, including [10], but the first reference on the subject seems to be [19]. The combinatorial projection onto a combinatorially convex subcomplex is defined by the following proposition:

Proposition 2.6 [19, Lemma 13.8] *Let X be a* CAT(0) *cube complex, $C \subset X$ be a combinatorially convex subspace and $x \in X \setminus C$ be a vertex. Then there exists a unique vertex $y \in C$ minimizing the distance to x. Moreover, for any vertex of C, there exists a combinatorial geodesic from it to x passing through y.*

The following result makes precise how the combinatorial projection behaves with respect to the hyperplanes:

Proposition 2.7 *Let X be a* CAT(0) *cube complex, C a combinatorially convex subspace, $p\colon X \to C$ the combinatorial projection onto C and $x, y \in X$ two vertices. The hyperplanes separating $p(x)$ and $p(y)$ are precisely the hyperplanes separating x and y which intersect C. In particular, $d(p(x), p(y)) \leq d(x, y)$.*

Proof A hyperplane separating $p(x)$ and $p(y)$ separates x and y according to [9, Lemma 2.10]. Conversely, let J be a hyperplane separating x and y and intersecting C. Notice that, according to Lemma 2.8 below, if J separates x and $p(x)$, or y and $p(y)$, then necessarily J must be disjoint from C. Therefore, J has to separate $p(x)$ and $p(y)$. □

Lemma 2.8 [19, Lemma 13.8] *Let X be a* CAT(0) *cube complex and $N \subset X$ a combinatorially convex subspace. Let $p\colon X \to N$ denote the combinatorial projection onto N. Then every hyperplane separating x and $p(x)$ separates x and N.*

The following lemma will be particularly useful in this paper:

Lemma 2.9 *Let X be a* CAT(0) *cube complex and $C_1 \subset C_2$ two subcomplexes with C_2 combinatorially convex. Let $p_2\colon X \to C_2$ denote the combinatorial projection onto C_2, and p_1 the nearest-point projection onto C_1, which associates to any vertex the set of the vertices of C_1 which minimize the distance from it. Then $p_1 \circ p_2 = p_1$.*

Proof Let $x \in X$. If $x \in C_2$, clearly $p_1(p_2(x)) = p_1(x)$, so we suppose that $x \notin C_2$. If $z \in C_1$, then, according to Proposition 2.6, there exists a combinatorial geodesic between x and z passing through $p_2(x)$. Thus,

$$d(x,z) = d(x, p_2(x)) + d(p_2(x), z).$$

In particular, if $y \in p_1(x)$,

$$d(x, C_1) = d(x,y) = d(x, p_2(x)) + d(p_2(x), y).$$

The previous two equalities give

$$d(p_2(x), z) - d(p_2(x), y) = d(x,z) - d(x, C_1) \geq 0.$$

Therefore, $d(p_2(x), y) = d(p_2(x), C_1)$, i.e. $y \in p_1(p_2(x))$. We have proved $p_1(x) \subset p_1(p_2(x))$.

It is worth noting that we have also proved that $d(x, C_1) = d(x, p_2(x)) + d(p_2(x), C_1)$. Thus, for every $y \in p_1(p_2(x))$, once again because there exists a combinatorial geodesic between x and y passing through $p_2(x)$ according to Proposition 2.6,

$$d(x,y) = d(x, p_2(x)) + d(p_2(x), y) = d(x, p_2(x)) + d(p_2(x), C_1) = d(x, C_1),$$

i.e. $y \in p_1(x)$. We have proved that $p_1(p_2(x)) \subset p_1(x)$, concluding the proof. $\square$

We conclude with a purely technical lemma which will be used in Section 3.3.

Lemma 2.10 *Let X be a* CAT(0) *cube complex, $N \subset X$ a combinatorially convex subspace and $x, y \notin N$ two vertices. Fix a combinatorial geodesic $[x,y]$ and choose a vertex $z \in [x,y]$. Then $d(z,N) \leq d(x,N) + d(y,N)$.*

Proof For convenience, if $A, B \subset X$ are two sets of vertices, let $\Delta(A,B)$ denote the number of hyperplanes separating A and B. According to Lemma 2.8, $\Delta(\{a\}, N) = d(a,N)$ for every vertex $a \in X$. Notice that, because $z \in [x,y]$ implies that no hyperplane can separate z and $\{x,y\} \cup N$, we have

$$d(z,N) = \Delta(z,N) = \Delta(\{x,y,z\}, N) + \Delta(\{x,z\}, N \cup \{y\}) + \Delta(\{y,z\}, N \cup \{x\}),$$
$$d(x,N) = \Delta(x,N) = \Delta(x, N \cup \{y,z\}) + \Delta(\{x,z\}, N \cup \{y\}) + \Delta(\{x,y,z\}, N),$$
$$d(y,N) = \Delta(y,N) = \Delta(y, N \cup \{x,z\}) + \Delta(\{y,z\}, N \cup \{x\}) + \Delta(\{x,y,z\}, N).$$

Now, it is clear that $d(z,N) \leq d(x,N) + d(y,N)$. $\square$

Disc diagrams A fundamental tool to study CAT(0) cube complexes is the theory of *disc diagrams*. For example, they were extensively used by Sageev [28] and Wise [32]. The rest of this section is dedicated to basic definitions and properties of disc diagrams.

Definition 2.11 Let X be a nonpositively curved cube complex. A *disc diagram* is a continuous combinatorial map $D \to X$, where D is a finite contractible square complex with a fixed topological embedding into $\mathbb{S}^2$; notice that D may be *nondegenerate*, i.e. homeomorphic to a disc, or may be *degenerate*. In particular, the complement of D in $\mathbb{S}^2$ is a 2–cell, whose attaching map will be referred to as the *boundary path* $\partial D \to X$ of $D \to X$; it is a combinatorial path. The *area* of $D \to X$, denoted by $\mathrm{Area}(D)$, corresponds to the number of squares of D.

Given a combinatorial closed path $P \to X$, we say that a disc diagram $D \to X$ is *bounded* by $P \to X$ if there exists an isomorphism $P \to \partial D$ such that the following diagram is commutative:

$$\begin{tikzcd} \partial D \arrow[r] & X \\ P \arrow[u] \arrow[ur] & \end{tikzcd}$$

According to a classical argument due to van Kampen [20] (see also [23, Lemma 2.17]), there exists a disc diagram bounded by a given combinatorial closed path if and only if this path is null-homotopic. Thus, if X is a CAT(0) cube complex, then any combinatorial closed path bounds a disc diagram.

As a square complex, a disc diagram contains hyperplanes; they are called *dual curves*. Equivalently, they correspond to the connected components of the reciprocal images of the hyperplanes of X. Given a disc diagram $D \to X$, a *nogon* is a dual curve homeomorphic to a circle; a *monogon* is a subpath, of a self-intersecting dual curve, homeomorphic to a circle; an *oscugon* is a subpath of a dual curve whose endpoints are the midpoints of two adjacent edges; a *bigon* is a pair of dual curves intersecting into two different squares. See Figure 2.

Theorem 2.12 [32, Lemma 2.2] *Let X be a nonpositively curved cube complex and $D \to X$ a disc diagram. If D contains a nogon, a monogon, a bigon or an oscugon, then there exists a new disc diagram $D' \to X$ such that*

(i) *D' is bounded by ∂D,*

(ii) $\mathrm{Area}(D') \leq \mathrm{Area}(D) - 2$.

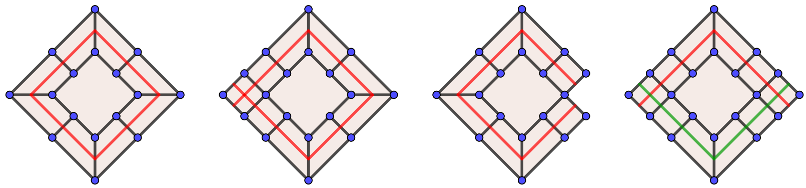

Figure 2: From left to right: a nogon, a monogon, an oscugon and a bigon.

Let X be a CAT(0) cube complex. A *cycle of subcomplexes* $\mathcal{C}$ is a sequence of subcomplexes $C_1, \dots, C_r$ such that $C_1 \cap C_r \neq \varnothing$ and $C_i \cap C_{i+1} \neq \varnothing$ for every $1 \leq i \leq r-1$. A disc diagram $D \to X$ is *bounded* by $\mathcal{C}$ if $\partial D \to X$ can be written as the concatenation of r combinatorial geodesics $P_1, \dots, P_r \to X$ such that $P_i \subset C_i$ for every $1 \leq i \leq r$. The *complexity* of such a disc diagram is defined by the pair $(\mathrm{Area}(D), \mathrm{length}(\partial D))$, and a disc diagram bounded by $\mathcal{C}$ will be of *minimal complexity* if its complexity is minimal with respect to the lexicographic order among all the possible disc diagrams bounded by $\mathcal{C}$ (allowing modifications of the paths P_i). It is worth noticing that such a disc diagram does not exist if our subcomplexes contain no combinatorial geodesics. On the other hand, if our subcomplexes are combinatorially geodesic, then a disc diagram always exists.

Our main technical result on disc diagrams is the following, which we proved in [9, Theorem 2.13]:

Theorem 2.13 *Let X be a* CAT(0) *cube complex, $\mathcal{C} = (C_1, \dots, C_r)$ a cycle of subcomplexes and $D \to X$ a disc diagram bounded by $\mathcal{C}$. For convenience, write ∂D as the concatenation of r combinatorial geodesics $P_1, \dots, P_r \to X$ with $P_i \subset C_i$ for every $1 \leq i \leq r$. If the complexity of $D \to X$ is minimal, then*

(i) *if C_i is combinatorially convex, two dual curves intersecting P_i are disjoint;*

(ii) *if C_i and C_{i+1} are combinatorially convex, no dual curve intersects both P_i and P_{i+1}.*

In general, a disc diagram $D \to X$ is not an embedding. However, the proposition below, which we proved in [9, Proposition 2.15], determines precisely when it is an isometric embedding.

Proposition 2.14 *Let X be a* CAT(0) *cube complex and $D \to X$ a disc diagram which does not contain any bigon. With respect to the combinatorial metrics, $\varphi\colon D \to X$ is an isometric embedding if and only if every hyperplane of X induces at most one dual curve of D.*

We mention two particular cases below.

Corollary 2.15 *Let $\mathcal{C} = (C_1, C_2, C_3, C_4)$ be a cycle of four subcomplexes. Suppose that C_2, C_3 and C_4 are combinatorially convex subcomplexes. If $D \to X$ is a disc diagram of minimal complexity bounded by $\mathcal{C}$, then $D \to X$ is an isometric embedding.*

Proof First, we write $\partial D \to X$ as the concatenation of four combinatorial geodesics $P_1, P_2, P_3, P_4 \to X$, where $P_i \subset C_i$ for every $1 \leq i \leq 4$. Suppose that there exist two dual curves c_1 and c_2 of D induced by the same hyperplane J. Because a combinatorial geodesic cannot intersect a hyperplane twice, the four endpoints of c_1 and c_2 belong to four different sides of ∂D. On the other hand, because C_2, C_3 and C_4 are combinatorially convex, it follows from Theorem 2.13 that any dual curve intersecting P_3 must intersect P_1; say that c_1 intersects P_1 and P_3. Therefore, c_2 must intersect P_2 and P_4, and we deduce that c_1 and c_2 are transverse. But this implies that J self-intersects in X, which is impossible. $\square$

Corollary 2.16 [9, Corollary 2.17] *Let X be a* CAT(0) *cube complex and $\mathcal{C}$ a cycle of four combinatorially convex subcomplexes. If $D \to X$ is a disc diagram of minimal complexity bounded by $\mathcal{C}$, then D is combinatorially isometric to a rectangle* $[0, a] \times [0, b] \subset \mathbb{R}^2$ *and $D \to X$ is an isometric embedding.*

Contracting isometries We gave in the introduction the definition of a contracting isometry. It turns out that an isometry is contracting if and only if its axis is contracting, in the following sense:

Definition 2.17 Let (S, d) be a metric space, $Y \subset C$ a subspace and $B \geq 0$ a constant. We say that Y is B*–contracting* if, for every ball disjoint from Y, the diameter of its nearest-point projection onto Y is at most B. A subspace will be *contracting* if it is B–contracting for some $B \geq 0$.

In fact, we may slightly modify this definition in the following way:

Lemma 2.18 *Let X be a geodesic metric space, $S \subset X$ a subspace and $L > 0$ a constant. Then S is contracting if and only if there exists $C \geq 0$ such that, for all vertices $x, y \in X$ satisfying $d(x, y) < d(x, S) - L$, the projection of $\{x, y\}$ onto S has diameter at most C.*

Proof The implication is clear: if S is B–contracting, $\{x, y\}$ is included into the ball $B(x, d(x, S))$ whose projection onto S has diameter at most B.

Conversely, suppose that there exists $C \geq 0$ such that, for all vertices $x, y \in X$ satisfying $d(x,y) < d(x,S) - L$, the projection of $\{x,y\}$ onto S has diameter at most C. Let $B = B(x,R)$ be a ball disjoint from S. If B is disjoint from the L–neighborhood of S, then our assumption implies that the projection of B onto S has diameter at most $2C$. From now on, we suppose that B intersects the L–neighborhood of S. As a consequence, $d(x,S) \leq R + L$. If $R \leq L$, then

$$\begin{aligned} d(x',y') &\leq d(x',x) + d(x,y) + d(y,y') = d(x,S) + d(x,y) + d(y,S) \\ &\leq 2(d(x,S) + d(x,y)) \leq 2(R + L + R) \leq 6L \end{aligned}$$

for every point $y \in B$ and all projections of x' and y' of x and y onto S, respectively. It follows that the diameter of the projection of B onto S is at most $12L$. Now, suppose that $R > L$ and fix a point $y \in B$ and two projections $x', y' \in S$ respectively of x and y onto S. We claim that $d(x',y') \leq 2(6L + C)$. If $y \in B(x, R - L)$, the inequality follows from our initial assertion, so suppose that $y \in B \setminus B(x, R - L)$. As a consequence, there exists a point $z \in B(y, 2L) \cap B(x, R - L)$. Fix one of its projections z' onto S. We distinguish two cases. If y is not in the $4L$–neighborhood of S, then

$$d(y',x') \leq d(y',z') + d(z',x') \leq C + C = 2C,$$

where the inequality $d(y',z') \leq C$ is justified by $d(y,z) \leq 2L < 3L \leq d(y,S) - L$. Next, if y belongs to the $4L$–neighborhood of S, then

$$\begin{aligned} d(y',x') &\leq d(y',y) + d(y,z) + d(z,z') + d(z',x') \leq d(y,S) + d(y,z) + d(z,S) + C \\ &\leq 2(d(y,S) + d(y,z)) + C \leq 2(4L + 2L) + C = 12L + C. \end{aligned}$$

Thus, we have proved that the diameter of B onto S has diameter at most $4(6L + C)$, concluding the proof of our lemma. □

When dealing with CAT(0) cube complexes, it is often more convenient to work with combinatorially convex subspaces. Of course, it is always possible to look at the combinatorial convex hull of the subspace which is considered, and the following lemma provides a sufficient condition which ensures that a subspace and its combinatorial convex hull are either both contracting or both noncontracting.

Lemma 2.19 *Let X be a* CAT(0) *cube complex, $S \subset X$ a subspace and $N \supset S$ a combinatorially convex subspace. Suppose that the Hausdorff distance $d_H(S,N)$ between S and N is finite. Then S is contracting if and only if N is contracting as well.*

Proof We know from Lemma 2.9 that the projection onto S is the composition of the projection of X onto N with the projection of N onto S. Consequently, the Hausdorff distance between the projections onto N and onto S is at most $d_H(N, S)$. Our lemma follows immediately. □

Typically, the previous lemma applies to quasiconvex combinatorial geodesics:

Lemma 2.20 *Let X be a* CAT(0) *cube complex and γ an infinite combinatorial geodesic. Let $N(\gamma)$ denote the combinatorial convex hull of γ. Then γ is quasiconvex if and only if the Hausdorff distance between γ and $N(\gamma)$ is finite.*

Proof We claim that (the 1–skeleton of) $N(\gamma)$ is the union $\mathcal{L}$ of the combinatorial geodesics whose endpoints are on γ. First, it is clear that any such geodesic must be included into $N(\gamma)$, whence $\mathcal{L} \subset N(\gamma)$. Conversely, given some $x \in N(\gamma)$, we want to prove that there exists a geodesic between two vertices of γ passing through x. Fix some $y \in \gamma$. Notice that, because x belongs to $N(\gamma)$, no hyperplane separates x from γ, so that any hyperplane separating x and y must intersect γ. As a consequence, for some $n \geq 1$ sufficiently large, every hyperplane separating x and y separates $\gamma(-n)$ and $\gamma(n)$; furthermore, we may suppose without loss of generality that y belongs to the subsegment of γ between $\gamma(-n)$ and $\gamma(n)$. We claim that the concatenation of a geodesic from $\gamma(-n)$ to x with a geodesic from x to $\gamma(n)$ defines a geodesic between $\gamma(-n)$ and $\gamma(n)$, which proves that $x \in \mathcal{L}$. Indeed, if this concatenation defines a path which is not a geodesic, then there must exist a hyperplane J intersecting it twice. Notice that J must separate x from $\{\gamma(n), \gamma(-n)\}$, so that we deduce from the convexity of halfspaces that J must separate x and y. By our choice of n, we deduce that J separates $\gamma(-n)$ and $\gamma(n)$. Finally, we obtain a contradiction since we already know that J separates x from $\gamma(-n)$ and $\gamma(n)$. Thus, we have proved that $N(\gamma) \subset \mathcal{L}$.

It follows that γ is quasiconvex if and only if the Hausdorff distance between γ and $N(\gamma)$ is finite. □

It is worth noticing that any contracting geodesic is quasiconvex. This is a particular case of [31, Lemma 3.3], which we record for future use:

Lemma 2.21 *Let X be a geodesic metric space. Every contracting geodesic of X is quasiconvex.*

So far, we have considered contracting subspaces. Now, let us focus on *contracting isometries*.

Definition 2.22 Let X be a metric space and $g \in \mathrm{Isom}(X)$ an isometry. Then g is *contracting* if there exists a point $x \in X$ such that the following two conditions hold:

- the map $n \mapsto g^n \cdot x$ induces a quasi-isometric embedding $\mathbb{Z} \to X$;
- the subspace $\{g^n \cdot x \mid n \in \mathbb{Z}\}$ is contracting.

It is worth noticing that this definition does not depend on the point we are looking, as justified by our next lemma.

Lemma 2.23 *Let G be a group acting on a geodesic space X. Fix two points $x_1, x_2 \in G$ and set $\gamma_1 = G \cdot x_1$ and $\gamma_2 = G \cdot x_2$. Suppose that the map $n \mapsto g^n \cdot x_1$ induces a quasi-isometric embedding $\mathbb{Z} \to X$ and that γ_1 is a B–contracting subspace. Then the map $n \mapsto g^n \cdot x_2$ induces a quasi-isometric embedding $\mathbb{Z} \to X$ and γ_2 is a contracting subspace.*

Proof By noticing that $d(g^n x_1, g^n x_2) = d(x_1, x_2)$ for every $n \in \mathbb{Z}$, we deduce that the Hausdorff distance between γ_1 and γ_2 is finite, say D, so that it follows easily that the map $n \mapsto g^n \cdot x_2$ induces a quasi-isometric embedding $\mathbb{Z} \to X$.

Now, fix a point $x \in X$. Let $y \in \gamma_1$ be a point of γ_1 minimizing the distance to x and $z \in \gamma_2$ a point of γ_2 minimizing the distance to x. Our goal is to show that $d(y,z) \leq B + 5D$. If $x = y$ then $d(y,z) = d(x,z) = d(x,\gamma_2) \leq D$ and we are done. From now on, we suppose that x is different from y. Since the Hausdorff distance between γ_1 and γ_2 is D, there exists a point $w \in \gamma_1$ such that $d(w,z) \leq D$. Fix a geodesic $[x,y]$ from x to y and a geodesic $[x,w]$ from x to w. Let $\epsilon > 0$ be a small real and a the unique point of $[x,y]$ satisfying $d(x,a) = d(x,y) - \epsilon$. Because $d(x,w) \geq d(x,\gamma_1) = d(x,y)$, the point b of $[x,w]$ satisfying $d(x,b) = d(x,a)$ is well defined. Finally, let c be a point of γ_1 minimizing the distance to b. Notice that

$$d(y,c) \leq B \tag{1}$$

because the nearest-point projection of the ball $B(x, d(x,a))$, which contains y and c, must have diameter at most B. Next,

$$d(x,w) \leq d(x,z) + d(z,w) \leq d(x,\gamma_2) + D \leq d(x,y) + d_H(\gamma_1,\gamma_2) + D,$$

hence

$$d(x,w) \leq d(x,y) + 2D. \tag{2}$$

We also have

$$\begin{aligned} d(c,w) &\leq d(b,c)+d(b,w) \leq 2d(b,w) \leq 2(d(x,w)-d(x,b)) \\ &\leq 2(d(x,w)-d(x,y)+\epsilon), \end{aligned}$$

hence

$$d(c,w) \leq 4D \tag{3}$$

by applying (2). By combining (1) and (3), we conclude that

$$d(y,z) \leq d(y,c)+d(c,w)+d(w,z) \leq B+5D.$$

As a consequence, the nearest-point projection of a ball onto γ_2 remains in the $(B+5D)$–neighborhood of its nearest-point projection onto γ_1. It follows that γ_2 must be contracting as well. □

Roughly speaking, to determine whether an isometry is contracting, we want to know if its axes are contracting, or equivalently if one of its axes is contracting. In CAT(0) cube complexes, an isometry does not always admit a "true" axis, i.e. a bi-infinite geodesic on which it acts by translations. However, Haglund proved the following result:

Proposition 2.24 [17] *Let X be a* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry. Suppose that g acts stably without inversion on X. Then either g fixes a vertex, or it acts on a bi-infinite geodesic as a translation of positive length.*

An isometry acts *stably without inversion* if no power of this isometry inverts a hyperplane, i.e. stabilizes it and switches the two halfspaces it delimits. The key observation is that every isometry acts stably without inversion in the cubical subdivision of the cube complex.

Corollary 2.25 *Let X be a* CAT(0) *cube complex and let X' denote its cubical subdivision. Then any isometry of X with an unbounded orbit admits an axis in X'.*

This statement justifies the following convention:

Convention In Sections 3 and 4, we will always suppose that an isometry of a CAT(0) cube complex admits an axis. According to Corollary 2.25, this assumption holds up to subdividing the cube complex.

In Section 5, where we study diagram groups, we will see that the CAT(0) cube complexes on which they act satisfy this property, so that subdividing will not be necessary.

3 Contracting isometries from hyperplane configurations

3.1 Quasiconvex geodesics, I

Given a subcomplex Y of some CAT(0) cube complex, we denote by $\mathcal{H}(Y)$ the set of hyperplanes intersecting Y. In this section, the question we are interested in is to determine when a combinatorial geodesic γ is quasiconvex just from the set $\mathcal{H}(\gamma)$. We begin by introducing some definitions.

Definition 3.1 A *facing triple* in a CAT(0) cube complex is the data of three hyperplanes such that no two of them are separated by the third hyperplane.

Definition 3.2 A *join of hyperplanes* is the data of two families of hyperplanes $(\mathcal{H} = (H_1, \dots, H_r), \mathcal{V} = (V_1, \dots, V_s))$ which do not contain facing triples and such that any hyperplane of $\mathcal{H}$ is transverse to any hyperplane of $\mathcal{V}$. If H_i (resp. V_j) separates H_{i-1} and H_{i+1} (resp. V_{j-1} and V_{j+1}) for every $2 \leq i \leq r-1$ (resp. $2 \leq j \leq s-1$), we say that $(\mathcal{H}, \mathcal{V})$ is a *grid of hyperplanes*.
If $(\mathcal{H}, \mathcal{V})$ is a join or a grid of hyperplanes satisfying $\min(\#\mathcal{H}, \#\mathcal{V}) \leq C$, we say that $(\mathcal{H}, \mathcal{V})$ is C*–thin*.

Our main criterion is:

Proposition 3.3 *Let X be a* CAT(0) *cube complex and γ an infinite combinatorial geodesic. The following statements are equivalent*:

(i) *γ is quasiconvex.*

(ii) *There exists $C \geq 1$ satisfying*
   - *$\mathcal{H}(\gamma)$ does not contain C pairwise transverse hyperplanes*,
   - *any grid of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H}, \mathcal{V} \subset \mathcal{H}(\gamma)$ is C–thin.*

(iii) *There exists some constant $C \geq 1$ such that any join of hyperplanes $(\mathcal{H}, \mathcal{V})$ satisfying $\mathcal{H}, \mathcal{V} \subset \mathcal{H}(\gamma)$ is C–thin.*

This will be a direct consequence of the following lemma and a criterion of hyperbolicity we proved in [9]:

Lemma 3.4 *An infinite combinatorial geodesic γ is quasiconvex if and only if its combinatorial convex hull is a hyperbolic subcomplex.*

Proof According to Lemma 2.20, γ is quasiconvex if and only if the Hausdorff distance between γ and $N(\gamma)$ is finite. In particular, if γ is quasiconvex, then $N(\gamma)$ is quasi-isometric to a line, and a fortiori is hyperbolic. Conversely, if $N(\gamma)$ is a hyperbolic subcomplex, its bigons are uniformly thin. This implies the quasiconvexity of γ. □

Theorem 3.5 [9, Theorem 3.3] *A* CAT(0) *cube complex is hyperbolic if and only if it is finite-dimensional and its grids of hyperplanes of X are uniformly thin.*

Proof of Proposition 3.3 From Lemma 3.4 and Theorem 3.5, we deduce the equivalence (i) $\Longleftrightarrow$ (ii). Then, because a grid of hyperplanes or a collection of pairwise transverse hyperplanes gives rise to a join of hyperplanes, (iii) clearly implies (ii).

To conclude, we want to prove (ii) $\Longrightarrow$ (iii). Let C denote the constant given by (ii). Let $(\mathcal{H},\mathcal{V})$ be a join of hyperplanes satisfying $\mathcal{H},\mathcal{V}\subset\mathcal{H}(\gamma)$. Suppose by contradiction that $\#\mathcal{H},\#\mathcal{V}\geq\mathrm{Ram}(C)$. Because $\mathcal{H}(\gamma)$ does not contain C pairwise transverse hyperplanes, we deduce that $\mathcal{H}$ and $\mathcal{V}$ each contain a subfamily of C pairwise disjoint hyperplanes, say $\mathcal{H}'$ and $\mathcal{V}'$, respectively. Notice that any hyperplane of $\mathcal{H}'$ is transverse to any hyperplane of $\mathcal{V}'$. Since the hyperplanes of $\mathcal{H}'$ and $\mathcal{V}'$ are all intersected by γ, we conclude that $(\mathcal{H}',\mathcal{V}')$ defines a (C,C)–grid of hyperplanes in X, contradicting the definition of C. Therefore, $\min(p,q)<\mathrm{Ram}(C)$. □

3.2 Contracting convex subcomplexes, I

In the previous section, we showed how to recognize the quasiconvexity of a combinatorial geodesic from the set of the hyperplanes it intersects. Therefore, according to Proposition 3.3, we reduced the problem of determining when a combinatorial geodesic is contracting to the problem of determining when a combinatorially convex subcomplex is contracting. This is the main criterion of this section:

Theorem 3.6 *Let X be a* CAT(0) *cube complex and $Y\subset X$ a combinatorially convex subcomplex. Then Y is contracting if and only if there exists $C\geq 0$ such that any join of hyperplanes $(\mathcal{H},\mathcal{V})$ with $\mathcal{H}\subset\mathcal{H}(Y)$ and $\mathcal{V}\cap\mathcal{H}(Y)=\varnothing$ satisfies $\min(\#\mathcal{H},\#\mathcal{V})\leq C$.*

Proof Suppose that Y is B–contracting for some $B\geq 0$, and let $(\mathcal{H},\mathcal{V})$ be a join of hyperplanes satisfying $\mathcal{H}\subset\mathcal{H}(Y)$ and $\mathcal{V}\cap\mathcal{H}(Y)=\varnothing$. Because any hyperplane of $\mathcal{H}$ intersects Y and $\mathcal{H}$ does not contain facing triples, there exist two vertices $x_-,x_+\in Y$ separated by the hyperplanes of $\mathcal{H}$. We can use the vertices x_- and x_+ to fix an orientation of the halfspaces delimited by the hyperplanes of $\mathcal{H}$. More precisely, let

$H^{\pm}$ denote the halfspace delimited by an $H \in \mathcal{H}$ which contains $x_{\pm}$. Also, for any $V \in \mathcal{V}$, let V^+ denote the halfspace delimited by V which does not contain Y. Now, we set

$$\mathfrak{H}^{\pm} = \bigcap_{H \in \mathcal{H}} H^{\pm} \quad \text{and} \quad \mathfrak{V} = \bigcap_{V \in \mathcal{V}} V^{+}.$$

If $\mathcal{C}$ is the cycle of subcomplexes $(\mathfrak{H}^-, \mathfrak{V}, \mathfrak{H}^+, Y)$, then Corollary 2.16 states that a disc diagram of minimal complexity bounded by $\mathcal{C}$ defines a flat rectangle $F \subset X$. For convenience, identify F with the square complex $[0,n]\times[0,m]$ such that $[0,n]\times\{0\}\subset Y$ and $[0,n] \times \{m\} \subset \mathfrak{V}$. Notice that $n \geq \#\mathcal{H}$ and $m \geq \#\mathcal{V}$. As a consequence, there exist two vertices $x, y \in [0,n]\times\{m\}$ satisfying $d(x,y) = \min(\#\mathcal{V}, \mathcal{H}) - 1$. Notice that $d(x,Y) \geq \#\mathcal{V} > d(x,y)$ because Y and $\{x,y\}$ are separated by all the hyperplanes of $\mathcal{V}$. Moreover, if $p\colon X \to Y$ denotes the combinatorial projection onto Y, because x and y are separated by $\min(\#\mathcal{H}, \#\mathcal{V}) - 1$ hyperplanes which intersect Y, it follows from Proposition 2.7 that $d(p(x), p(y)) \geq \min(\#\mathcal{H}, \#\mathcal{V}) - 1$. Finally, since Y is B–contracting, we conclude that

$$\min(\#\mathcal{H}, \#\mathcal{V}) \leq 1 + d(p(x), p(y)) \leq 1 + B.$$

Conversely, suppose that there exists $C \geq 0$ such that any join of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{H}(Y)$ and $\mathcal{V} \cap \mathcal{H}(Y) = \varnothing$ satisfies $\min(\#\mathcal{H}, \#\mathcal{V}) \leq C$. We want to prove that Y is contracting by applying Lemma 2.18. Let $x, y \in X$ be two distinct vertices satisfying $d(x,y) < d(x,Y) - C$ and let $p\colon X \to Y$ denote the combinatorial projection onto Y. If $\mathcal{H}$ denotes the set of the hyperplanes separating $\{x,y\}$ and Y, and $\mathcal{V}$ the set of the hyperplanes separating $p(x)$ and $p(y)$, then $(\mathcal{H}, \mathcal{V})$ defines a join of hyperplanes: because any hyperplane of $\mathcal{H}$ separates $\{x,y\}$ and $\{p(x), p(y)\}$, and any hyperplane of $\mathcal{V}$ separates $\{x, p(x)\}$ and $\{y, p(y)\}$ according to Proposition 2.7 and Lemma 2.8, we deduce that any hyperplane of $\mathcal{H}$ is transverse to any hyperplane of $\mathcal{V}$. Because the assumption $d(x,y) < d(x,Y) - C$ implies $\#\mathcal{H} > C$, we conclude that

$$d(p(x), p(y)) \leq \#\mathcal{V} \leq C.$$

Therefore, Y is contracting. □

Combining Proposition 3.3 and Theorem 3.6, we get:

Corollary 3.7 *Let X be a* CAT(0) *cube complex and γ an infinite combinatorial geodesic. Then γ is contracting if and only if there exists a constant $C \geq 1$ such that any join of hyperplanes $(\mathcal{H}, \mathcal{V})$ satisfying $\mathcal{H} \subset \mathcal{H}(\gamma)$ is necessarily C–thin.*

Proof Suppose that γ is contracting. In particular, γ is quasiconvex so, according to Proposition 3.3, there exists a constant C_1 such that any join of hyperplanes $(\mathcal{H},\mathcal{V})$ satisfying $\mathcal{H},\mathcal{V}\subset\mathcal{H}(\gamma)$ is C_1–thin. Then, because $N(\gamma)$ is also contracting according to Lemmas 2.20 and 2.19, we deduce from Theorem 3.6 that there exists a constant C_2 such that any join of hyperplanes $(\mathcal{H},\mathcal{V})$ satisfying $\mathcal{H}\subset\mathcal{H}(\gamma)$ and $\mathcal{V}\cap\mathcal{H}(\gamma)=\varnothing$ is C_2–thin. Now, let $(\mathcal{H},\mathcal{V})$ be any join of hyperplanes satisfying $\mathcal{H}\subset\mathcal{H}(\gamma)$. If we set $\mathcal{V}_1=\mathcal{V}\cap\mathcal{H}(\gamma)$ and $\mathcal{V}_2=\mathcal{V}\setminus\mathcal{V}_1$, then $(\mathcal{H},\mathcal{V}_1)$ is C_1–thin and $(\mathcal{H},\mathcal{V}_2)$ is C_2–thin. Thus,

$$\min(\#\mathcal{H},\#\mathcal{V})\leq\min(\#\mathcal{H},\#\mathcal{V}_1)+\min(\#\mathcal{H},\#\mathcal{V}_2)\leq C_1+C_2,$$

i.e. $(\mathcal{H},\mathcal{V})$ is $(C_1{+}C_2)$–thin.

Conversely, suppose that there exists a constant $C\geq 1$ such that any join of hyperplanes $(\mathcal{H},\mathcal{V})$ satisfying $\mathcal{H}\subset\mathcal{H}(\gamma)$ is necessarily C–thin. According to Proposition 3.3, γ is quasiconvex, so that it is sufficient to prove that $N(\gamma)$ is contracting to conclude that γ is contracting as well according to Lemma 2.19. Finally, it follows from Theorem 3.6 that $N(\gamma)$ is contracting. □

3.3 Well-separated hyperplanes

Separation properties of hyperplanes play a central role in the study of contracting isometries. *Strongly separated* hyperplanes were introduced in [1] in order to characterize the rank-one isometries of right-angled Artin groups; they were also crucial in the proof of the rank rigidity theorem for CAT(0) cube complexes in [5]. In [6], Charney and Sultan used contracting isometries to distinguish, quasi-isometrically, some cubulable groups, and they introduced *k–separated* hyperplanes in order to characterize contracting isometries in uniformly locally finite CAT(0) cube complexes [6, Theorem 4.2]. In this section, we introduce *well-separated* hyperplanes in order to generalize this characterization to CAT(0) cube complexes without any local finiteness condition.

Definition 3.8 Let J and H be two disjoint hyperplanes in some CAT(0) cube complex and $L\geq 0$. We say that J and H are *L–well-separated* if any family of hyperplanes transverse to both J and H, which does not contain any facing triple, has cardinality at most L. Two hyperplanes are *well-separated* if they are L–well-separated for some $L\geq 1$.

Theorem 3.9 *Let γ be an infinite combinatorial geodesic. Then γ is contracting if and only if there exist constants $r,L\geq 1$, hyperplanes $\{H_i, i\in\mathbb{Z}\}$ and vertices*

$\{x_i \in \gamma \cap N(H_i) \mid i \in \mathbb{Z}\}$ *such that, for every* $i \in \mathbb{Z}$,

- $d(x_i, x_{i+1}) \leq r$,
- *the hyperplanes* H_i *and* H_{i+1} *are* L*–well-separated.*

The following lemma is a combinatorial analogue to [6, Lemma 4.3], although our proof is essentially different. This is the key technical lemma needed to prove Proposition 3.11, which is in turn the first step toward the proof of Theorem 3.9.

Lemma 3.10 *Let* γ *be an infinite quasiconvex combinatorial geodesic. There exists a constant* C *depending only on* γ *such that, if two vertices* $x, y \in \gamma$ *satisfy* $d(x, y) > C$, *then there exists a hyperplane* J *separating* x *and* y *such that the nearest-point projection of* $N(J)$ *onto* γ *is included into* $[x, y] \subset \gamma$.

Proof Let J be a hyperplane separating x and y whose projection onto γ is not included into $[x, y] \subset \gamma$. Say that this projection contains a vertex at the right of y; the case where it contains a vertex at the left of x is completely symmetric. Thus, we have the configuration, illustrated by Figure 3, where $b \notin [x, y]$ is a projection onto γ of some vertex $a \in N(J)$, $z \in \gamma$ is a vertex adjacent to J, and $m = m(a, b, z)$ is the median point of $\{a, b, z\}$. Because m belongs to a combinatorial geodesic between z and b, and these two points belong to γ, necessarily $m \in N(\gamma)$. On the other hand, m belongs to a combinatorial geodesic between a and b, and by definition b is a vertex of γ minimizing the distance to a, so $d(m, b) = d(m, \gamma)$; because, by the quasiconvexity of γ, the Hausdorff distance between γ and $N(\gamma)$ is finite, say L,

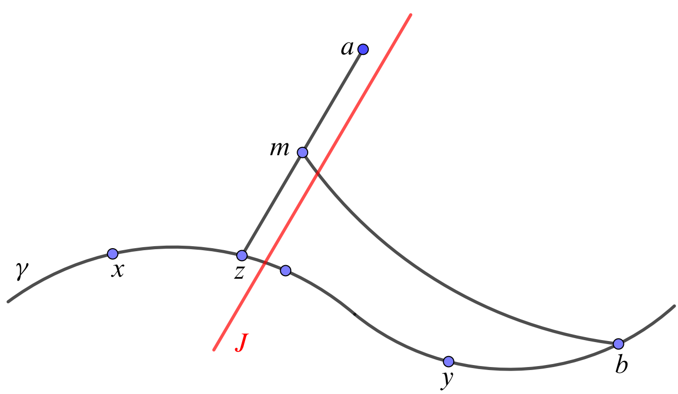

Figure 3: Configuration from Lemma 3.10.

we deduce that $d(m,b) = d(m,\gamma) \leq L$ since m belongs to a combinatorial geodesic between b and z (which implies that $m \in N(\gamma)$). Using Lemma 2.10, we get

$$d(y,N(J)) \leq d(z,N(J)) + d(b,N(J)) \leq d(b,m) \leq L.$$

Thus, J intersects the ball $B(y,L)$.

Let $\mathcal{H}$ be the set of hyperplanes separating x and y, and intersecting $B(y,L)$. Of course, because any hyperplane of $\mathcal{H}$ intersects γ, the collection $\mathcal{H}$ contains no facing triple. On the other hand, it follows from Proposition 3.3 that $\dim N(\gamma)$ is finite, so that if $\#\mathcal{H} \geq \mathrm{Ram}(s)$ for some $s \geq \dim N(\gamma)$ then the collection $\mathcal{H}$ must contain at least s pairwise disjoint hyperplanes. Since this collection of hyperplanes does not contain facing triples and intersects the ball $B(y,L)$, we deduce that $s \leq 2L$, hence $\#\mathcal{H} \leq \mathrm{Ram}\big(\max(\dim N(\gamma), 2L)\big)$.

Finally, we have proved that there are at most $2\,\mathrm{Ram}\big(\max(\dim N(\gamma), 2L)\big)$ hyperplanes separating x and y such that the projections of their neighborhoods onto γ are not included into $[x,y] \subset \gamma$. It clearly follows that $d(x,y) > 2\,\mathrm{Ram}\big(\max(\dim N(\gamma), 2L)\big)$ implies that there is a hyperplane separating x and y such that the projection of its neighborhood onto γ is included into $[x,y] \subset \gamma$. □

Proposition 3.11 *Let γ be an infinite combinatorial geodesic. If γ is quasiconvex then there exist constants $r, L \geq 1$, hyperplanes $\{H_i \mid i \in \mathbb{Z}\}$ and vertices $\{x_i \in \gamma \cap N(H_i) \mid i \in \mathbb{Z}\}$ such that, for every $i \in \mathbb{Z}$,*

- $d(x_i, x_{i+1}) \leq r$,
- *the hyperplanes H_i and H_{i+1} are disjoint.*

Proof Because γ is quasiconvex, we may apply the previous lemma; let C be the constant it gives. Let $\dots, y_{-1}, y_0, y_1, \dots \in \gamma$ be vertices along γ satisfying $d(y_i, y_{i+1}) = C+1$ for all $i \in \mathbb{Z}$. According to the previous lemma, for every $k \in \mathbb{Z}$, there exists a hyperplane J_k separating y_{2k} and y_{2k+1} whose projection onto γ is included into $[y_{2k}, y_{2k+1}] \subset \gamma$; let x_k be one of the two vertices in $\gamma \cap N(J_k)$. Notice that

$$\begin{aligned} d(x_i, x_{i+1}) &\leq d(y_{2i}, y_{2i+3}) \\ &\leq d(y_{2i}, y_{2i+1}) + d(y_{2i+1}, y_{2i+2}) + d(y_{2i+2}, y_{2i+3}) \\ &= 3(C+1). \end{aligned}$$

Finally, it is clear that, for every $k \in \mathbb{Z}$, the hyperplanes J_k and J_{k+1} are disjoint: if this were not the case, there would exist a vertex of $N(J_k) \cap N(J_{k+1})$ whose projection onto γ would belong to $[y_{2k}, y_{2k+1}] \cap [y_{2k+2}, y_{2k+3}] = \varnothing$. □

Proof of Theorem 3.9 Suppose γ contracting. In particular, γ is quasiconvex, so, by applying Proposition 3.11, we find a constant $L \geq 1$, a collection of pairwise disjoint hyperplanes $\{J_i, i \in \mathbb{Z}\}$ and a collection of vertices $\{y_i \in \gamma \cap N(J_i) \mid i \in \mathbb{Z}\}$ such that $d(y_i, y_{i+1}) \leq L$ for all $i \in \mathbb{Z}$. Let C be the constant given by Corollary 3.7 and set $x_i = y_{iC}$ for all $i \in \mathbb{Z}$. Notice that

$$\begin{aligned} d(x_i, x_{i+1}) = d(y_{iC}, y_{iC+C}) &\leq d(y_{iC}, y_{iC+1}) + \cdots + d(y_{iC+C-1}, y_{iC+C}) \\ &\leq (C+1)L. \end{aligned}$$

Now, we want to prove that the hyperplanes J_{nC} and $J_{(n+1)C}$ are C–well-separated for every $n \in \mathbb{Z}$. So let $\mathcal{H}$ be a collection of hyperplanes transverse to both J_{nC} and $J_{(n+1)C}$ which contains no facing triple. Because every hyperplane $H \in \mathcal{H}$ is transverse to any J_{nC+k} for $0 \leq k \leq C$, we obtain a join of hyperplanes $(V_{C+1}, V_{\#\mathcal{H}})$ satisfying $V_{C+1} \subset \mathcal{H}(\gamma)$. By definition of C, we deduce $\#\mathcal{H} \leq C$.

Conversely, suppose that there exist constants $l, L \geq 1$, hyperplanes $\{J_i \mid i \in \mathbb{Z}\}$ and vertices $\{x_i \in \gamma \cap N(J_i), i \in \mathbb{Z}\}$ such that, for every $i \in \mathbb{Z}$,

- $d(x_i, x_{i+1}) \leq l$,
- the hyperplanes J_i and J_{i+1} are L–well-separated.

Let $(\mathcal{V}_p, \mathcal{V}_q)$ be a join of hyperplanes with $\mathcal{V}_p = \{V_1, \dots, V_p\} \subset \mathcal{H}(\gamma)$. For convenience, we may suppose without loss of generality that each V_i intersects γ at y_i with the property that y_i separates y_{i-1} and y_{i+1} along γ for all $2 \leq i \leq p-1$. If $p > 3l + 2L$, there exist $L < r < s < p - L$ such that y_r and y_s are separated by J_k, J_{k+1}, J_{k+2} and J_{k+3} for some $k \in \mathbb{Z}$. Because J_k and J_{k+1} are L–well-separated, the hyperplanes $\{V_1, \dots, V_r\}$ cannot be transverse to both J_k and J_{k+1} since $r > L$, so there exists $1 \leq \alpha \leq r$ such that V_α and J_{k+1} are disjoint. Similarly, the hyperplanes $\{V_s, \dots, V_p\}$ cannot be transverse to both J_{k+2} and J_{k+3} since $p - s > L$, so there exists $1 \leq \omega \leq r$ such that V_ω and J_{k+2} are disjoint. Notice that the hyperplanes J_{k+1} and J_{k+2} separate V_α and V_ω, so that the hyperplanes of $\mathcal{V}_q$, which are all transverse to both V_α and V_ω, are transverse to both J_{k+1} and J_{k+2}. But J_{k+1} and J_{k+2} are L–well-separated, hence $q \leq L$.

Thus, we have proved that $\min(p, q) \leq 3l + 2L$. Finally, Corollary 3.7 implies that γ is contracting. □

3.4 Contracting isometries, I

Finally, we apply our criteria of contractivity to the axis of some isometry, providing a characterization of contracting isometries. We first need the following definition:

Definition 3.12 Let X be a CAT(0) cube complex and $g \in \mathrm{Isom}(X)$ an isometry. We say that g *skewers* a pair of hyperplanes (J_1, J_2) if $g^n D_1 \subsetneq D_2 \subsetneq D_1$ for some $n \in \mathbb{Z} \setminus \{0\}$ and for some halfspaces D_1 and D_2 delimited by J_1 and J_2, respectively.

Our main criterion is:

Theorem 3.13 *Let X be a* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry with a combinatorial axis γ. The following statements are equivalent*:

(i) *g is a contracting isometry.*

(ii) *There exists $C \geq 1$ such that any join of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{H}(\gamma)$ is C–thin.*

(iii) *g skewers a pair of well-separated hyperplanes.*

Proof The equivalence (i) $\iff$ (ii) is a direct consequence of Corollary 3.7.

Now, we want to prove (iii) $\Longrightarrow$ (i). So suppose that there exist two halfspaces D_1 and D_2 respectively delimited by two well-separated hyperplanes J_1 and J_2 such that $g^n D_1 \subsetneq D_2 \subsetneq D_1$ for some $n \in \mathbb{Z}$. Notice that

$$\cdots \subsetneq g^{2n} D_1 \subsetneq g^n D_2 \subsetneq g^n D_1 \subsetneq D_2 \subsetneq D_1 \subsetneq g^{-n} D_2 \subsetneq g^{-n} D_1 \subsetneq g^{-2n} D_2 \subsetneq \cdots .$$

We claim that J_1 intersects γ. Suppose by contradiction that this is not the case. Then, because γ is $\langle g \rangle$–invariant, $g^k J_1$ does not intersect γ for every $k \in \mathbb{Z}$. As a consequence, there exists some $m \in \mathbb{Z}$ such that $\gamma \subset g^{mn} D_1 \setminus g^{(m+1)n} D_1$. A fortiori, $g^m \gamma \subset g^{(m+1)n} D_1 \setminus g^{(m+2)n} D_1$, and because γ is $\langle g \rangle$–invariant, we deduce that

$$\gamma \subset (g^{mn} D_1 \setminus g^{(m+1)n} D_1) \cap (g^{(m+1)n} D_1 \setminus g^{(m+2)n} D_1) = \varnothing,$$

a contradiction. So J_1 intersects γ, and a fortiori, $g^{kn} J_1$ intersects γ for every $k \in \mathbb{Z}$. For every $k \in \mathbb{Z}$, the intersection between $g^{kn} N(J_1)$ and γ contains exactly two vertices; fix an orientation along γ and denote by y_k the first vertex along γ which belongs to $g^{kn} N(J_1)$. Clearly,

$$d(y_k, y_{k+1}) = d(g^{kn} y_0, g^{(k+1)n} y_0) = d(y_0, g^n y_0) = \|g\|^n$$

since y_0 belongs to the combinatorial axis of g, where $\|g\|$ denotes the combinatorial translation length of g (see [17]). Furthermore, J_1 and $g^n J_1$ are well-separated: indeed, because J_2 separates J_1 and $g^n J_1$, we deduce that any collection of hyperplanes transverse to both J_1 and $g^n J_1$ must also be transverse to both J_1 and J_2, so that we conclude that J_1 and $g^n J_1$ are well-separated since J_1 and J_2 are themselves well-separated. Therefore, $\{g^{kn} J_1 \mid k \in \mathbb{N}\}$ defines a family of pairwise well-separated hyperplanes intersecting γ at uniformly separated points. Theorem 3.9 implies that γ is contracting.

Conversely, suppose that g is contracting. According to Theorem 3.9, there exist three pairwise well-separated hyperplanes J_1, J_2 and J_3 intersecting γ. Say that they respectively delimit three halfspaces D_1, D_2 and D_3 satisfying $D_3 \subsetneq D_2 \subsetneq D_1$, where D_3 contains $\gamma(+\infty)$. We claim that g skewers the pair (J_1, J_2). Fix two vertices $x \in N(J_1) \cap \gamma$ and $y \in D_3 \cap \gamma$. Notice that, if I denotes the set of $n \in \mathbb{N}$ such that $g^n J_1$ and J_1 are transverse and $g^n J_1$ does not separate x and y, then $\{g^n J_1 \mid n \in I\}$ defines a set of hyperplanes (without facing triples, since they all intersect γ) transverse to both J_2 and J_3. Because J_2 and J_3 are well-separated, necessarily I has to be finite. A fortiori, there must exist at most $\#I + d(x,y)$ integers $n \in \mathbb{N}$ such that $g^n J_1$ and J_1 are transverse. As a consequence, there exists an integer $n \in \mathbb{N}$ such that $g^n J_1 \subset D_2$. Finally, we have proved (i) $\Longrightarrow$ (iii). □

4 Combinatorial boundary

4.1 Generalities

In this section, we will use the following notation: given two subsets of hyperplanes $\mathcal{H}_1$ and $\mathcal{H}_2$, we define the *almost-inclusion* $\mathcal{H}_1 \subset_a \mathcal{H}_2$ if all but finitely many hyperplanes of $\mathcal{H}_1$ belong to $\mathcal{H}_2$. In particular, $\mathcal{H}_1$ and $\mathcal{H}_2$ are *commensurable* provided that $\mathcal{H}_1 \subset_a \mathcal{H}_2$ and $\mathcal{H}_2 \subset_a \mathcal{H}_1$. Notice that commensurability defines an equivalence relation on the set of all the collections of hyperplanes, and the almost-inclusion induces a partial order on the associated quotient set. This allows the following definition.

Definition 4.1 Let X be a CAT(0) cube complex. Its *combinatorial boundary* is the poset $(\partial^c X, \prec)$, where $\partial^c X$ is the set of the combinatorial rays modulo the relation $r_1 \sim r_2$ if $\mathcal{H}(r_1)$ and $\mathcal{H}(r_2)$ are commensurable and where the partial order $\prec$ is defined by $r_1 \prec r_2$ whenever $\mathcal{H}(r_1) \subset_a \mathcal{H}(r_2)$.

Notice that this construction is equivariant, i.e. if a group G acts by isometries on X, then G acts on $\partial^c X$ by $\prec$–isomorphisms.

The following two lemmas essentially state that we will be able to choose a "nice" combinatorial ray representing a given point in the combinatorial boundary. They will be useful in the next sections.

Lemma 4.2 *Let X be a* CAT(0) *cube complex, $x_0 \in X$ a base vertex and $\xi \in \partial^c X$. There exists a combinatorial ray r with $r(0) = x_0$ and $r = \xi$ in $\partial^c X$.*

Proof Let r_n be a combinatorial path which is the concatenation of a combinatorial geodesic $[x_0, \xi(n)]$ between x_0 and $\xi(n)$ with the subray ξ_n of ξ starting at $\xi(n)$. If r_n is not geodesic, then there exists a hyperplane J intersecting both $[x_0, \xi(n)]$ and ξ_n. Thus, J separates x_0 and $\xi(n)$ but cannot separate $\xi(0)$ and $\xi(n)$ since otherwise J would intersect the ray ξ twice. We deduce that J necessarily separates x_0 and $\xi(0)$. Therefore, if we choose n large enough that the hyperplanes separating x_0 and $\xi(0)$ do not intersect the subray ξ_n, then r_n is a combinatorial ray. By construction, we have $r_n(0) = x_0$, and $\mathcal{H}(r_n)$ and $\mathcal{H}(\xi)$ are clearly commensurable, so $r_n = \xi$ in $\partial^c X$. $\square$

Lemma 4.3 *Let r and ρ be two combinatorial rays satisfying $r \prec \rho$. There exists a combinatorial ray p equivalent to r satisfying $p(0) = \rho(0)$ and $\mathcal{H}(p) \subset \mathcal{H}(\rho)$.*

Proof According to the previous lemma, we may suppose that $r(0) = \rho(0)$. By assumption, $\mathcal{H}(r) \subset_a \mathcal{H}(\rho)$. Let J be the last hyperplane of $\mathcal{H}(r) \setminus \mathcal{H}(\rho)$ intersected by r, i.e. there exists some $k \geq 1$ such that J is dual to the edge $[r(k), r(k+1)]$ and the hyperplanes intersecting the subray $r_k \subset r$ starting at $r(k+1)$ all belong to $\mathcal{H}(\rho)$. We claim that $r_k \subset \partial N(J)$. Indeed, if there existed some $j \geq k+1$ such that $d(r_k(j), N(J)) \geq 1$, then some hyperplane H would separate $r_k(j)$ and $N(J)$ according to Lemma 2.8; a fortiori, H would intersect r but not ρ, contradicting the choice of J. Let r' denote the combinatorial ray obtained from r by replacing the subray r_{k-1} by the symmetric of r_k with respect to J in $N(J) \simeq J \times [0,1]$. The set of hyperplanes intersecting r' is precisely $\mathcal{H}(r) \setminus \{J\}$. Because $|\mathcal{H}(r') \setminus \mathcal{H}(\rho)| < |\mathcal{H}(r) \setminus \mathcal{H}(\rho)|$, iterating the process will produce after finitely many steps a combinatorial ray p satisfying $p(0) = r(0) = \rho(0)$ and $\mathcal{H}(p) \subset \mathcal{H}(\rho)$. $\square$

Definition 4.4 Let X be a CAT(0) cube complex and $Y \subset X$ a subcomplex. We define the *relative combinatorial boundary* $\partial^c Y$ of Y as the subset of $\partial^c X$ corresponding the set of the combinatorial rays included into Y.

Lemma 4.5 *Let r be a combinatorial ray. Then $r \in \partial^c Y$ if and only if $\mathcal{H}(r) \subset_a \mathcal{H}(Y)$.*

Proof If $r \in \partial^c Y$, then by definition there exists a combinatorial ray $\rho \subset Y$ equivalent to r. Then

$$\mathcal{H}(r) \subset_a \mathcal{H}(\rho) \subset \mathcal{H}(Y).$$

Conversely, suppose that $\mathcal{H}(r) \subset_a \mathcal{H}(Y)$. According to Lemma 4.2, we may suppose that $r(0) \in Y$. Using the same argument as in the previous proof, we find an equivalent combinatorial ray ρ with $\mathcal{H}(\rho) \subset \mathcal{H}(Y)$ and $\rho(0) = r(0) \in Y$. Because Y is combinatorially convex, it follows that $\rho \subset Y$, hence $r \in \partial^c Y$. □

4.2 Quasiconvex geodesics, II

The goal of this section is to prove the following criterion, determining when a combinatorial axis is quasiconvex just from its endpoints in the combinatorial boundary:

Proposition 4.6 *Let X be a locally finite* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry with a combinatorial axis γ. A subray $r \subset \gamma$ is quasiconvex if and only if r is minimal in $\partial^c X$.*

We begin by proving the following preliminary lemma. (We recall that complete CAT(0) cube complexes are defined by Definition 2.3.)

Lemma 4.7 *Let X be a complete* CAT(0) *cube complex and r a combinatorial ray. Then $\mathcal{H}(r)$ contains an infinite sequence of pairwise disjoint hyperplanes.*

Proof We begin our proof with a completely general argument, without assuming that X is complete. First, we decompose $\mathcal{H}(r)$ as the disjoint union $\mathcal{H}_0 \sqcup \mathcal{H}_1 \sqcup \cdots$, where we define $\mathcal{H}_i = \{J \in \mathcal{H}(r) \mid d_\infty(r(0), J) = i\}$ for all $i \geq 0$. A priori, some $\mathcal{H}_i$ might be empty or infinite; moreover, $\mathcal{H}_i$ could be infinite and $\mathcal{H}_{i+1}$ nonempty. We want to understand the behavior of this sequence of collections of hyperplanes. Our first two claims state that each $\mathcal{H}_i$ is a collection of pairwise transverse hyperplanes and that $\mathcal{H}_i$ nonempty implies $\mathcal{H}_j$ nonempty for every $j < i$.

Claim 4.8 *For all $i \geq 0$ and $J_1, J_2 \in \mathcal{H}_i$, J_1 and J_2 are transverse.*

Suppose by contradiction that J_1 and J_2 are disjoint, and, for convenience, say that J_1 separates $r(0)$ and J_2. Because $d_\infty(r(0), J_1) = i$, there exist i pairwise disjoint hyperplanes $V_1, \dots, V_i$ separating $r(0)$ and J_1. Then $V_1, \dots, V_i, J_1$ are $i+1$ pairwise transverse hyperplanes separating $r(0)$ and J_2, hence $d_\infty(r(0), J_2) \geq i+1$, a contradiction.

Claim 4.9 *Let $J \in \mathcal{H}_i$, i.e. there exists a collection of i pairwise disjoint hyperplanes $V_0, \dots, V_{i-1}$ separating $r(0)$ and J. Then $V_j \in \mathcal{H}_j$ for every $0 \leq j \leq i-1$.*

First, since $V_0, \dots, V_{j-1}$ separate $r(0)$ and V_j, necessarily we have $d_\infty(r(0), V_j) \geq j$. Let $H_1, \dots, H_k$ be k pairwise disjoint hyperplanes separating $r(0)$ and V_j. Then $H_1, \dots, H_k, V_j, \dots, V_{i-1}$ define $k+i-j$ pairwise disjoint hyperplanes separating $r(0)$ and J, hence

$$k+i-j \leq d_\infty(r(0), J) = i,$$

i.e. $k \leq j$. We deduce that $d_\infty(r(0), V_j) \leq j$. Finally, we conclude that $d_\infty(r(0), V_j) = j$, that is, $V_j \in \mathcal{H}_j$.

Our third and last claim states that, if the first collections $\mathcal{H}_0, \dots, \mathcal{H}_p$ are finite, then there exists a sequence of cubes $Q_0, \dots, Q_p$ such that the intersection between two successive cubes is a single vertex and the hyperplanes dual to the cube Q_i are precisely the hyperplanes of $\mathcal{H}_i$. For this purpose, we define by induction the sequence of vertices $x_0, x_1, \dots$ by

- $x_0 = r(0)$;
- if x_j is defined and $\mathcal{H}_j$ is finite, x_{j+1} is the projection of x_j onto $C_j := \bigcap_{J \in \mathcal{H}_j} J^+$, where J^+ denotes the halfspace delimited by J which does not contain $r(0)$.

Notice that the intersection C_j is nonempty precisely because of Claim 4.8 and because we assumed that $\mathcal{H}_j$ is finite.

Claim 4.10 *For every $i \geq 0$, if $\mathcal{H}_{i-1}$ is finite then any hyperplane of $\mathcal{H}_i$ is adjacent to x_i, and, if $\mathcal{H}_i$ is finite, there exists a cube Q_i containing x_i and x_{i+1} as opposite vertices such that $\mathcal{H}_i$ is the set of hyperplanes intersecting Q_i.*

We argue by induction.

First, we notice that any hyperplane $J \in \mathcal{H}_0$ is adjacent to x_0. Suppose by contradiction that this is not the case, and let p denote the combinatorial projection of x_0 onto J^+. By assumption, $d(x_0, p) \geq 2$ so there exists a hyperplane H separating x_0 and p which is different from J. According to Lemma 2.8, H separates x_0 and J. This contradicts $J \in \mathcal{H}_0$. Therefore, any hyperplane $J \in \mathcal{H}_0$ is dual to an edge $e(J)$ with x_0 as an endpoint. If $\mathcal{H}_0$ is infinite, there is nothing to prove. Otherwise, $\{e(J) \mid J \in \mathcal{H}_0\}$ is a finite set of edges with a common endpoint and whose associated hyperplanes are pairwise transverse. Thus, because a CAT(0) cube complex does not contain interosculating hyperplanes, these edges span a cube Q_0. By construction, the set of hyperplanes intersecting Q_0 is $\mathcal{H}_0$. Furthermore, because the vertex of Q_0 opposite to x_0 belongs to C_0 and is at distance $\#\mathcal{H}_0$ from x_0, we deduce that it is precisely x_1.

Now suppose that $x_0, \dots, x_i$ are well defined and there exist some cubes $Q_0, \dots, Q_{i-1}$ satisfying our claim. We first want to prove that any hyperplane $J \in \mathcal{H}_i$ is adjacent to x_i. Suppose by contradiction that this is not the case, i.e. there exists some $J \in \mathcal{H}_i$ which is not adjacent to x_i. As a consequence, there must exist a hyperplane H separating x_i and J.

Case 1 (H does not separate x_0 and x_i) So H separates x_0 and J, and H does not belong to $\mathcal{H}_k$ for $k \leq i-1$. Noticing that $r(k) \in J^+$ for sufficiently large $k \geq 0$, we deduce that H necessarily intersects r, hence $H \in \mathcal{H}_j$ for some $j \geq i$. Therefore, there exist i pairwise disjoint hyperplanes $V_1, \dots, V_i$ separating $r(0)$ and H. A fortiori, $V_1, \dots, V_i, H$ define $i+1$ pairwise disjoint hyperplanes separating $r(0)$ and J, contradicting $J \in \mathcal{H}_i$.

Case 2 (H separates x_0 and x_i) In particular, H intersects a cube Q_j for some $0 \leq j \leq i-1$, i.e. $H \in \mathcal{H}_j$. Let p_0 (resp. p_i) denote the combinatorial projection of x_0 (resp. x_i) onto J^+. Notice that, because H is disjoint from J, it does not separate p_0 and p_i: these two vertices belong to the same halfspace delimited by H, say H^+. Because H separates x_i and J, it separates x_i and p_i, so x_i belongs to the second halfspace delimited by H, say H^-. Then, since H separates x_0 and x_i by hypothesis, we deduce that x_0 belongs to H^+; in particular, H does not separate x_0 and p_0. On the other hand, if $k \geq 0$ is sufficiently large that $r(k) \in J^+$, H must separate $r(0)$ and $r(k)$ since $H \in \mathcal{H}(r)$. According to Proposition 2.6, there exists a combinatorial geodesic between x_0 and $r(k)$ passing through p_0. Because H is disjoint from J^+, we conclude that H must separate x_0 and p_0, a contradiction.

We have proved that any hyperplane of $\mathcal{H}_i$ is adjacent to x_i. Thus, if $\mathcal{H}_i$ is finite, we can construct our cube Q_i as above. This concludes the proof of our claim.

From now on, we assume that X is a complete CAT(0) cube complex. First, as a consequence of the previous claim, we deduce that each $\mathcal{H}_i$ is finite. Indeed, suppose by contradiction that some $\mathcal{H}_i$ is infinite. Without loss of generality, we may suppose that $\mathcal{H}_j$ is finite for every $j < i$, so that the points $x_0, \dots, x_i$ and the cubes $Q_0, \dots, Q_{i-1}$ are well defined. According to our previous claim, any hyperplane of $\mathcal{H}_i$ is adjacent to x_i; thus, the infinite set of edges adjacent to x_i and dual to the hyperplanes of $\mathcal{H}_i$ define an infinite cube, contradicting the completeness of X.

Therefore, we have defined an infinite sequence of vertices $x_0, x_1, \dots$ and an infinite sequence of cubes $Q_0, Q_1, \dots$. Thanks to Claim 4.9, for every $i \geq 0$, there exists a sequence of pairwise disjoint hyperplanes $V_0^i, \dots, V_i^i$ with $V_j^i \in \mathcal{H}_j$. Because each $\mathcal{H}_k$ is finite, up to taking a subsequence, we may suppose that, for every $k \geq 0$, the sequence (V_k^i) is eventually constant to some hyperplane $V_k \in \mathcal{H}_k$. By construction, our sequence $V_0, V_1, \dots$ defines a collection of pairwise disjoint hyperplanes in $\mathcal{H}(r)$, concluding the proof of our lemma. □

Now we are ready to prove Proposition 4.6.

Proof of Proposition 4.6 Suppose that r is not quasiconvex. According to Proposition 3.3, for every $n \geq 1$ there exist a join of hyperplanes $(\mathcal{H}_n, \mathcal{V}_n)$ with $\mathcal{H}_n, \mathcal{V}_n \subset \mathcal{H}(r)$ and $\#\mathcal{H}_n, \#\mathcal{V}_n = n$. The hyperplanes of $\mathcal{H}_n \cup \mathcal{V}_n$ are naturally ordered depending on the order of their intersections along r. Without loss of generality, we may suppose that the hyperplane of $\mathcal{H}_n \cup \mathcal{V}_n$ which is closest to $r(0)$ belongs to $\mathcal{H}_n$; because $\langle g \rangle$ acts on $\mathcal{H}(\gamma)$ with finitely many orbits, up to translating by powers of g and taking a subsequence, we may suppose that this hyperplane $J \in \mathcal{H}_n$ does not depend on n. Finally, let V_n denote the hyperplane of $\mathcal{V}_n$ which is farthest from $r(0)$.

Thus, if J^- (resp. V_n^+) denotes the halfspace delimited by J (resp. V_n) which does not contain $r(+\infty)$ (resp. which contains $r(+\infty)$), then $\mathcal{C}_n = (J^-, V_n^+, r)$ defines a cycle of three subcomplexes. Let $D_n \to X$ be a disc diagram of minimal complexity bounded by $\mathcal{C}_n$. Because no hyperplane can intersect J, V_n or r twice, it follows from Proposition 2.14 that $D_n \to X$ is an isometric embedding, so we will identify D_n with its image in X. Let us write the boundary ∂D_n as a concatenation of three combinatorial geodesics

$$\partial D_n = \mu_n \cup \nu_n \cup \rho_n,$$

where $\mu_n \subset J^-$, $\nu_n \subset V_n^+$ and $\rho_n \subset r$. Because X is locally finite, up to taking a subsequence, we may suppose that the sequence of combinatorial geodesics (μ_n) converges to a combinatorial ray $\mu_\infty \subset J^-$. Notice that, if $H \in \mathcal{H}(\mu_\infty)$, then $H \in \mathcal{H}(\mu_k)$ for some $k \geq 1$; it follows from Theorem 2.13 that H does not intersect ν_k in D_k, hence $H \in \mathcal{H}(\rho_k)$. Therefore, $\mathcal{H}(\mu_\infty) \subset \mathcal{H}(r)$.

According to Lemma 4.7, there exists an infinite collection of pairwise disjoint hyperplanes $J_1, J_2, \dots \in \mathcal{H}(r)$. Because $\langle g \rangle$ acts on $\mathcal{H}(\gamma)$ with finitely many orbits, necessarily there exist some $r, s \geq 1$ and some $m \geq 1$ such that $g^m J_r = J_s$. Thus, $\{J_r, g^m J_r, g^{2m} J_r, \dots\}$ defines a collection of infinitely many pairwise disjoint hyperplanes in $\mathcal{H}(r)$ stable by the semigroup generated by g^m. For convenience, let $\{H_0, H_1, \dots\}$ denote this collection.

If J is disjoint from some H_k, then $H_k, H_{k+1}, \dots \notin \mathcal{H}(\mu_\infty)$ since $\mu_\infty \subset J^-$. On the other hand, $H_k, H_{k+1}, \dots \in \mathcal{H}(r)$, so we deduce that $\mu_\infty \prec r$ with $\mu_\infty \neq r$ in $\partial^c X$; this precisely means that r is not minimal in $\partial^c X$.

From now on, suppose that J is transverse to all the H_k. As a consequence, $g^{km} J$ is transverse to H_n for every $n \geq k$. For every $k \geq 1$, let H_k^+ denote the halfspace delimited by H_k which contains $r(+\infty)$ and let $p_k \colon X \to H_k^+$ be the combinatorial projection onto H_k^+. We define a sequence of vertices (x_n) by

$$x_0 = r(0) \qquad \text{and} \qquad x_{n+1} = p_{n+1}(x_n) \quad \text{for every } n \geq 0.$$

Because of Lemma 2.9, we have $x_n = p_n(x_0)$ for every $n \geq 1$. Therefore, it follows from Proposition 2.6 that, for every $n \geq 1$, there exists a combinatorial geodesic between x_0 and x_n passing through x_k for $k \leq n$. Thus, there exists a combinatorial ray ρ passing through all the x_k.

Let $H \in \mathcal{H}(\rho)$. Then H separates x_k and x_{k+1} for some $k \geq 0$, and it follows from Lemma 2.8 that H separates x_k and H_k. As a first consequence, we deduce that, for every $k \geq 1$, $g^{km} J$ cannot belong to $\mathcal{H}(\rho)$ since $g^{km} J$ is transverse to H_n for every $n \geq k$, hence $r \neq \rho$ in $\partial^c X$. On the other hand, H separates $x_0 = r(0)$ and H_k, and $r(n) \in H_k^+$ for n large enough, so $H \in \mathcal{H}(r)$. We have proved that $\mathcal{H}(\rho) \subset \mathcal{H}(r)$, i.e. $\rho \prec r$. Thus, r is not minimal in $\partial^c X$.

Conversely, suppose that r is not minimal in $\partial^c X$. Thus, there exists a combinatorial ray ρ with $\rho(0) = r(0)$ such that $\mathcal{H}(\rho) \subset \mathcal{H}(r)$ and $|\mathcal{H}(r) \setminus \mathcal{H}(\rho)| = +\infty$. Let $J_1, J_2, \dots \in \mathcal{H}(r) \setminus \mathcal{H}(\rho)$ be an infinite collection. Without loss of generality, suppose that, for every $i > j \geq 1$, the edge $J_j \cap r$ is closer to $r(0)$ than the edge $J_i \cap r$. Let $N \geq 1$, and let $d(N)$ denote the distance $d(r(0), \omega)$ where ω is the endpoint of the edge $J_N \cap r$ which is farthest from $r(0)$. Choose any collection of hyperplanes $\mathcal{V} \subset \mathcal{H}(\rho)$ with $\#\mathcal{V} \geq d(N) + N$. Then $\mathcal{V}$ contains a subcollection $\{V_1, \dots, V_N\}$ such that the edge $V_i \cap r$ is farther from $r(0)$ than the edge $J_N \cap r$. We know that each V_j separates $\{r(0), \omega\}$ and $\{r(k), \rho(k)\}$ for k large enough, and each J_i separates $\{r(k), \omega\}$ and $\{r(0), \rho(k)\}$ for k large enough (since J_i and ρ are disjoint); we deduce that J_i and V_j are transverse for any $1 \leq i, j \leq N$. Therefore, $(\{J_1, \dots, J_N\}, \{V_1, \dots, V_N\})$ defines a join of hyperplanes in $\mathcal{H}(r)$. In fact, we have proved:

Fact 4.11 *Let $r \prec \rho$ be two combinatorial rays with the same origin. If $\mathcal{J} \subset \mathcal{H}(r) \setminus \mathcal{H}(\rho)$ is an infinite collection, then for every $N \geq 1$ there exists a join of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{J}$, $\mathcal{V} \subset \mathcal{H}(\rho)$ and $\#\mathcal{H}, \#\mathcal{V} \geq n$.*

Since N can be chosen arbitrarily large, it follows from Proposition 3.3 that r is not quasiconvex. □

4.3 Contracting convex subcomplexes, II

Once again, according to Proposition 3.3, the previous section reduces the problem of determining when a combinatorial axis is contracting to the problem of determining when a (hyperbolic, cocompact) combinatorially convex subcomplex is contracting. To do so, we need the following definition:

Definition 4.12 A subset $S \subset \partial^c X$ is *full* if any element of $\partial^c X$ which is comparable with an element of S necessarily belongs to S.

Our main criterion is:

Theorem 4.13 *Let X be a locally finite* $\mathrm{CAT}(0)$ *cube complex and $Y \subset X$ a hyperbolic combinatorially convex* $\mathrm{Aut}(X)$*–cocompact subcomplex. Then Y is contracting if and only if $\partial^c Y$ is full in $\partial^c X$.*

Before turning to the proof of Theorem 4.13, recall that the n^{th} *Ramsey number* $\mathrm{Ram}(n)$ is defined as the least number R satisfying the following condition: if K is a complete graph with at least R vertices, then any labeling of the edges of K with two colors contains a unicolor subgraph containing at least n vertices. As a consequence, if $\mathcal{H}$ is a collection of $\mathrm{Ram}(n)$ hyperplanes of a given $\mathrm{CAT}(0)$ cube complex of dimension at most n, then $\mathcal{H}$ contains a subcollection of n pairwise disjoint hyperplanes. (See for instance [9, Lemma 3.7].)

Proof of Theorem 4.13 Suppose that Y is not contracting. Notice that, because Y is a cocompact subcomplex, its dimension, say d, must be finite. According to Theorem 3.6, for every $n \geq d$, there exists a join of hyperplanes $(\mathcal{H}_n, \mathcal{V}_n)$ with $\mathcal{H}_n \subset \mathcal{H}(Y)$, $\mathcal{V}_n \cap \mathcal{H}(Y) = \varnothing$ and $\#\mathcal{H}_n = \#\mathcal{V}_n \geq \mathrm{Ram}(n)$. Up to taking subcollections of $\mathcal{H}_n$, we may suppose thanks to Ramsey's theorem that $\mathcal{H}_n$ is a collection of pairwise disjoint hyperplanes of cardinality exactly n. For convenience, write $\mathcal{H}_n = (H_1^n, \dots, H_n^n)$, where H_i^n separates H_{i-1}^n and H_{i+1}^n for every $2 \leq i \leq n-1$. Finally, let D denote the intersection of all the halfspaces delimited by hyperplanes of $\mathcal{V}_n$ which are disjoint from Y.

Let $\mathcal{C}_n$ be the cycle of subcomplexes $(N(H_1^n), D, N(H_n^n), Y)$. By Corollary 2.16, a disc diagram of minimal complexity bounded by $\mathcal{C}_n$ defines a flat rectangle D_n. Say that a hyperplane intersecting the H_i^n in D_n is *vertical*, and a hyperplane intersecting the elements of $\mathcal{V}_n$ in D_n is *horizontal*.

Claim 4.14 *If the grids of hyperplanes of Y are all C–thin and $n > C$, then at most* $\mathrm{Ram}(C)$ *horizontal hyperplanes intersect Y.*

Let $\mathcal{V}$ be a collection of horizontal hyperplanes which intersect Y; notice that $\mathcal{V}$ does not contain any facing triple since there exists a combinatorial geodesic in D_n (which is a combinatorial geodesic in X) intersecting all the hyperplanes of $\mathcal{V}$. If $\#\mathcal{V} \geq \mathrm{Ram}(s)$ for some $s \geq \dim Y$, then $\mathcal{V}$ contains a subcollection $\mathcal{V}'$ with s pairwise disjoint hyperplanes, so that $(\mathcal{H}_n, \mathcal{V}')$ defines a (n,s)–grid of hyperplanes. If $n > C$, this implies $s \leq C$. Therefore, $\#\mathcal{V} \leq \mathrm{Ram}(C)$. This proves the claim.

Now, because Y is a cocompact subcomplex, up to translating the D_n, we may suppose that a corner of each D_n belongs to a fixed compact fundamental domain C; and because X is locally finite, up to taking a subsequence, we may suppose that, for every ball B centered in C, $(D_n \cap B)$ is eventually constant, so that (D_n) converges naturally to a subcomplex D_∞. Notice that D_∞ is isomorphic to the square complex $[0,+\infty) \times [0,+\infty)$ with $[0,+\infty) \times \{0\} \subset Y$; let ρ denote this combinatorial ray. Clearly, if $r \subset D_\infty$ is a diagonal combinatorial ray starting from $(0,0)$, then $\mathcal{H}(\rho) \subset \mathcal{H}(r)$, i.e. r and ρ are comparable in $\partial^c X$. Furthermore, the hyperplanes intersecting $[0,+\infty) \times \{0\} \subset D_\infty$ are horizontal hyperplanes in some D_n, and so infinitely many of them are disjoint from Y according to our previous claim; this implies that $r \notin \partial^c Y$. We conclude that $\partial^c Y$ is not full in $\partial^c X$.

Conversely, suppose that $\partial^c Y$ is not full in $\partial^c X$, i.e. there exist two $\prec$–comparable combinatorial rays r and ρ satisfying $\rho \subset Y$ and $r \notin \partial^c Y$.

Suppose that $r \prec \rho$. According to Lemma 4.3, we may suppose that $r(0) = \rho(0)$ and $\mathcal{H}(r) \subset \mathcal{H}(\rho)$. But $r(0) \in Y$ and $\mathcal{H}(r) \subset \mathcal{H}(Y)$ imply $r \subset Y$, contradicting the assumption $r \notin \partial^c Y$.

Suppose that $\rho \prec r$. According to Lemma 4.2, we may suppose that $r(0) = \rho(0)$. Because $r \notin \partial^c Y$, there exists an infinite collection $\mathcal{J} \subset \mathcal{H}(r) \setminus \mathcal{H}(Y)$; a fortiori,

$\mathcal{J} \subset \mathcal{H}(r) \setminus \mathcal{H}(\rho)$ since $\rho \subset Y$. It follows from Fact 4.11 that, for every $N \geq 1$, there exists a join of hyperplanes $(\mathcal{H}, \mathcal{V})$ with $\mathcal{H} \subset \mathcal{J}$, $\mathcal{V} \subset \mathcal{H}(\rho)$ and $\#\mathcal{H}, \#\mathcal{V} \geq N$. Therefore, Theorem 3.6 implies that Y is not contracting. □

Remark 4.15 Neither the hyperbolicity of the subcomplex Y nor its cocompactness was necessary to prove the converse of the previous theorem. Thus, the relative combinatorial boundary of a combinatorially convex subcomplex is always full.

4.4 Contracting isometries, II

Finally, we want to apply the results we have found in the previous sections to characterize contracting isometries from the combinatorial boundary. We begin by introducing the following definition:

Definition 4.16 A point $\xi \in \partial^c X$ is *isolated* if it is not comparable with any other element of $\partial^c X$.

Our main criterion is the following:

Theorem 4.17 *Let X be a locally finite* CAT(0) *cube complex and $g \in \mathrm{Isom}(X)$ an isometry with a combinatorial axis γ. Then g is a contracting isometry if and only if $\gamma(+\infty)$ is isolated in $\partial^c X$.*

Proof If g is a contracting isometry then a subray $r \subset \gamma$ containing $\gamma(+\infty)$ is contracting. Because $N(r)$ is a combinatorially convex subcomplex quasi-isometric to a line, it follows that $\partial^c N(r) = \{r\}$. On the other hand, since $N(r)$ is contracting, it follows from Theorem 4.13 that $\partial^c N(r)$ is full in $\partial^c X$. This precisely means that r is isolated.

Conversely, suppose that r is isolated in $\partial^c X$. It follows from Proposition 4.6 that r is quasiconvex. Thus, r is contracting if and only if $N(r)$ is contracting. The quasiconvexity of r also implies $\partial^c N(r) = \{r\}$. Because r is isolated in $\partial^c X$, $\partial^c N(r)$ is full in $\partial^c X$, and so $N(r)$ is contracting according to Theorem 4.13. □

Thus, containing an isolated point in the combinatorial boundary is a necessary condition in order to admit a contracting isometry. The converse is also true if we strengthen our assumptions on the action of our group. The first result in this direction is the following theorem, where we assume that there exist only finitely many orbits of hyperplanes:

Theorem 4.18 *Let G be a group acting on a locally finite* CAT(0) *cube complex X with finitely many orbits of hyperplanes. Then $G \curvearrowright X$ contains a contracting isometry if and only if $\partial^c X$ has an isolated vertex.*

This result will be essentially a direct consequence of our main technical lemma:

Lemma 4.19 *Let X be a locally finite* CAT(0) *cube complex. If $\partial^c X$ has an isolated point then, for every $N \geq 2$, there exists a collection of N pairwise well-separated hyperplanes of X which does not contain any facing triple.*

Proof Let $r \in \partial^c X$ be an isolated point. According to Lemma 4.7, there exists an infinite collection $V_0, V_1, \ldots \in \mathcal{H}(r)$ of pairwise disjoint hyperplanes.

Claim 4.20 *For every $k \geq 0$, $\mathcal{H}(V_k) \cap \mathcal{H}(r)$ is finite.*

Suppose by contradiction that $\mathcal{H}(V_k) \cap \mathcal{H}(r)$ is infinite for some $k \geq 0$. Fix a vertex $x_0 \in r$ adjacent to V_k and $H_1, H_2, \ldots \in \mathcal{H}(V_k) \cap \mathcal{H}(r)$. Now, let x_i be the combinatorial projection of x_0 onto $N(H_i)$ for every $i \geq 1$; notice that $d(x_0, x_i) \to_{i \to +\infty} +\infty$ since X is locally finite so that only finitely many hyperplanes intersect a given ball centered at x_0. Then, once again because X is locally finite, if we choose a combinatorial geodesic $[x_0, x_i]$ between x_0 and x_i for every $i \geq 1$, up to taking a subsequence, we may suppose that the sequence $([x_0, x_i])$ converges to some combinatorial ray ρ.

Let $J \in \mathcal{H}(\rho)$. Then J separates x_0 and x_i for some $i \geq 1$. From Lemma 2.8, we deduce that J separates x_0 and $N(H_i)$. On the other hand, we know that H_i intersects the combinatorial ray r which starts from x_0, so necessarily $J \in \mathcal{H}(r)$. We have proved that $\mathcal{H}(\rho) \subset \mathcal{H}(r)$ holds.

Next, notice that $x_i \in N(V_k)$ for every $i \geq 1$. Indeed, if p_i denotes the combinatorial projection of x_0 onto $N(V_k) \cap N(H_i)$ and m the median vertex of $\{x_0, x_i, p_i\}$, then $x_0, p_i \in N(V_k)$ implies $m \in N(V_k)$ and $x_i, p_i \in N(H_i)$ implies $m \in N(H_i)$, hence $m \in N(H_i) \cap N(V_k)$. Since x_i minimizes the distance to x_0 in $N(H_i)$ and that m belongs to a geodesic between x_0 and x_i, we conclude that $x_i = m \in N(V_k)$. A fortiori, $\rho \subset N(V_k)$. As a consequence, $V_{k+1}, V_{k+2}, \ldots \notin \mathcal{H}(\rho)$, so ρ defines a point of $\partial^c X$, different from r, which is $\prec$–comparable with r: this contradicts the fact that r is isolated in $\partial^c X$. Our claim is proved.

Claim 4.21 *For every $k \geq 0$, there exists some $i \geq 1$ such that V_k and V_{k+i} are $(2i)$–well-separated.*

Suppose by contradiction that V_k and V_{k+i} are not well-separated for every $i \geq 1$, so that there exists a collection of hyperplanes $\mathcal{H}^i \subset \mathcal{H}(V_k) \cap \mathcal{H}(V_{k+i})$ which does not contain any facing triple and satisfying $\#\mathcal{H}^i \geq 2i$. Let $x_0 \in r$ be a vertex adjacent to V_k. Because $\mathcal{H}^i$ does not contain any facing triple and $\mathcal{H}^i \subset \mathcal{H}(V_k)$, there exists a vertex $x_i \in N(V_k)$ such that x_0 and x_i are separated by at least i hyperplanes of $\mathcal{H}^i$; let $\mathcal{K}^i$ denote this set of hyperplanes. For every $i \geq 1$, let

$$C_i = \bigcap_{J \in \mathcal{K}^i} \{\text{halfspace delimited by } J \text{ containing } x_i\}.$$

Let $\mathcal{C}_i$ be the cycle of subcomplexes $(N(V_k), C_i, N(V_{k+i}), r)$, and let $D_i \to X$ be a disc diagram of minimal complexity bounded by $\mathcal{C}_i$. According to Corollary 2.15, $D_i \to X$ is an isometric embedding so that we will identify D_i with its image in X. Because $\mathcal{H}(V_k) \cap \mathcal{H}(r)$ is finite according to our previous claim, we know that, for sufficiently large $i \geq 1$, $\mathcal{K}^i$ will contain a hyperplane disjoint from r, so that $C_i \cap r = \varnothing$. Up to taking a subsequence, we will suppose that this is always the case. In particular, ∂D_i can be decomposed as the concatenation of four nontrivial combinatorial geodesics,

$$\partial D_i = \alpha_i \cup \beta_i \cup \gamma_i \cup r_i,$$

where $\alpha_i \subset N(V_k)$, $\beta_i \subset C_i$, $\gamma_i \subset N(V_{k+i})$ and $r_i \subset r$. Notice that the intersection between α_i and r_i is a vertex adjacent to x_0, so, if we use the local finiteness of X to define, up to taking a subsequence, a subcomplex D_∞ as the limit of (D_i), then D_∞ is naturally bounded by two combinatorial rays α_∞ and r_∞ which are the limits of (α_i) and (r_i), respectively; in particular, $\alpha_\infty \subset N(V_k)$ and r_∞ is a subray of r.

For each $i \geq 1$, we will say that a hyperplane of D_i intersecting β_i is *vertical*, and a hyperplane of D_i intersecting α_i is *horizontal*; a horizontal hyperplane $J \in \mathcal{H}(D_i)$ is *high* if $d(x_0, J \cap \alpha_i) > \#\mathcal{H}(V_k) \cap \mathcal{H}(r)$.

Fact 4.22 *Let $i \geq 1$. The vertical hyperplanes of D_i are pairwise disjoint and intersect r_i. The horizontal hyperplanes of D_i are pairwise disjoint. The high horizontal hyperplanes of D_i intersect γ_i.*

It follows directly from Theorem 2.13 that the vertical hyperplanes of D_i are pairwise disjoint and intersect r_i, and the horizontal hyperplanes of D_i are pairwise disjoint. Suppose by contradiction that a high horizontal hyperplane H does not intersect γ_i. According to Theorem 2.13, necessarily H intersects r. On the other hand, because $d(x_0, H \cap \alpha_i) > \#\mathcal{H}(V_k) \cap \mathcal{H}(r)$, there must exist a horizontal hyperplane J intersecting

α_i between H and x_0 which does not intersect r_i; once again, it follows from Theorem 2.13 that J has to intersect γ_i. A fortiori, J and H are necessarily transverse, contradicting the fact that two horizontal hyperplanes are disjoint. This proves the fact.

Now, we want to define a particular subdiagram $D_i' \subset D_i$. Let H_i be the nearest high horizontal hyperplane of D_i to $\alpha_i \cap r_i$. According to the previous fact, H_i intersects γ_i. Let D_i' be the subdiagram delimited by H_i which contains β_i. Notice that the hyperplanes intersecting D_i' are either vertical or high horizontal. Thus, from the decomposition of $\mathcal{H}(D_i')$ as horizontal and vertical hyperplanes given by Fact 4.22, we deduce that D_i' is isometric to some square complex $[0,a_i] \times [0,b_i]$. Notice that $b_i \to +\infty$ as $i \to +\infty$ since $V_{k+1}, \dots, V_{k+i-1}$ separate α_i and γ_i, and similarly $a_i \to +\infty$ as $i \to +\infty$ since $a_i = \mathrm{length}(\alpha_i) - |\mathcal{H}(V_k) \cap \mathcal{H}(r)|$ and $\#\mathcal{K}^i \geq i$. Therefore, if $D'_\infty \subset D_\infty$ denotes the limit of (D_i'), D'_∞ is naturally isomorphic to the square complex $[0,+\infty) \times [0,+\infty)$. Let $\rho_\infty \subset D'_\infty$ be the combinatorial ray associated to $[0,+\infty) \times \{0\}$, and $\mu \subset D'_\infty$ be a *diagonal* combinatorial ray, i.e. a ray passing through the vertices $\{(i,i) \mid i \geq 0\}$.

Notice that ρ_∞ is naturally the limit of $(\rho_i) := (\partial D_i' \cap N(H_i))$. Thus, any hyperplane intersecting ρ_∞ is a vertical hyperplane in some D_i. It follows from Fact 4.22 that $\mathcal{H}(\rho_\infty) \subset \mathcal{H}(r)$. Because r is isolated in $\partial^c X$, we deduce that $r = \rho_\infty$ in $\partial^c X$.

Now, we clearly have $\mathcal{H}(\rho_\infty) \subset \mathcal{H}(\mu)$, hence $\mathcal{H}(r) \subset_a \mathcal{H}(\mu)$. On the other hand, μ intersects infinitely many high horizontal hyperplanes, which are disjoint from r according to Fact 4.22. Therefore, μ is a point of $\partial^c X$ which is $\prec$–comparable with r, but different from it, contradicting the assumption that r is isolated in $\partial^c X$. This concludes the proof of our claim.

Now, we are able to conclude the proof of our lemma. Applying Claim 4.21 $N-1$ times, we find a sequence of hyperplanes $V_{i(1)}, \dots, V_{i(N)}$ such that $V_{i(k)}$ and $V_{i(k+1)}$ are $L(k)$–well-separated for some $L(k) \geq 1$. Thus, $\{V_{i(1)}, \dots, V_{i(N)}\}$ defines a collection of N pairwise L–well-separated hyperplanes, with $L = \max(L(1), \dots, L(N-1))$, which does not contain any facing triple. □

Proof of Theorem 4.18 If G contains a contracting isometry, then $\partial^c X$ contains an isolated point according to Theorem 4.17. Conversely, suppose that $\partial^c X$ contains an isolated point. Let N denote the number of orbits of hyperplanes for the action $G \curvearrowright X$. According to Lemma 4.19, X contains a family of $3N$ pairwise well-separated hyperplanes which does not contain any facing triple. By our choice of N, we deduce that there exist a hyperplane J and $g, h \in G$ such that J, gJ and hgJ are pairwise well-separated. Without loss of generality, suppose that hgJ separates J and gJ, and let J^+ be the halfspace delimited by J which contains gJ and hgJ. If $gJ^+ \subset J^+$ or $hgJ^+ \subset J^+$, we deduce that g or hg skewers a pair of well-separated hyperplanes, and we deduce from Theorem 3.13 that G contains a contracting isometry. Otherwise, $hgJ^+ \subset gJ^+$ since hgJ separates J and gJ. Therefore, h skewers a pair of well-separated hyperplanes, and we deduce from Theorem 3.13 that G contains a contracting isometry. □

Our second result, stating that the existence of an isolated point in the combinatorial boundary assures the existence of a contracting isometry, is the following:

Theorem 4.23 *Let G be a countable group acting on a locally finite* $\mathrm{CAT}(0)$ *cube complex X. Suppose that the action $G \curvearrowright X$ is minimal* (i.e. X does not contain a proper G–invariant combinatorially convex subcomplex) *and G does not fix a point of $\partial^c X$. Then G contains a contracting isometry if and only if $\partial^c X$ contains an isolated point.*

Given a hyperplane J, delimiting two halfspaces J^- and J^+, an *orientation* of J is the choice of an ordered pair (J^-, J^+). The following terminology was introduced in [5]:

Definition 4.24 Let G be a group acting on a $\mathrm{CAT}(0)$ cube complex. A hyperplane J is *G–flippable* if, for every orientation (J^-, J^+) of J, there exists some $g \in G$ such that $gJ^- \subsetneq J^+$.

Our main tool to prove Theorem 4.23 is:

Lemma 4.25 *Let G be a countable group acting minimally on a* $\mathrm{CAT}(0)$ *cube complex X without fixing any point of $\partial^c X$. Then any hyperplane of X is G–flippable.*

Proof Suppose that there exists a hyperplane J which is not G–flippable, i.e. J admits an orientation (J^-, J^+) such that $gJ^+ \cap J^+ \neq \varnothing$ for every $g \in G$. Thus, $\{gJ^+ \mid g \in G\}$ defines a family of pairwise intersecting halfspaces. If the intersection $\bigcap_{g \in G} gJ^+$ is nonempty, it defines a proper G–invariant combinatorially convex subcomplex, so that the action $G \curvearrowright X$ is not minimal.

From now on, suppose that $\bigcap_{g \in G} gJ^+ = \varnothing$. Fix an enumeration $G = \{g_1, g_2, \dots\}$ and let I_n denote the nonempty intersection $\bigcap_{i=1}^n g_i J^+$. In particular, we have $I_1 \supset I_2 \supset \cdots$ and $\bigcap_{n \geq 1} I_n = \varnothing$. Define a sequence of vertices (x_n) by

$$x_0 \notin I_1 \quad \text{and} \quad x_{n+1} = p_{n+1}(x_n) \quad \text{for every } n \geq 0,$$

where $p_k\colon X \to I_k$ denotes the combinatorial projection onto I_k. According to Proposition 2.6, for every $n \geq 1$, there exists a combinatorial geodesic between x_0 and x_n passing through x_k for any $1 \leq k \leq n$, so there exists a combinatorial ray r passing through all the x_n. Notice that, because we also have $x_n = p_n(x_0)$ for every $n \geq 1$ according to Lemma 2.9, $\mathcal{H}(r)$ is precisely the set of the hyperplanes separating x_0 from some I_n. In particular, if ρ is any combinatorial ray starting from x_0 and intersecting all the I_n, then $\mathcal{H}(r) \subset \mathcal{H}(\rho)$.

Therefore, $\bigcap_{k \geq 1} \partial^c I_k$ is a nonempty (since it contains r) subset of $\partial^c X$ with r as a minimal element, i.e. for any $\rho \in \bigcap_{k \geq 1} \partial^c I_k$, necessarily $r \prec \rho$.

Let us prove that $\bigcap_{k \geq 1} \partial^c I_k$ is G–invariant. Fix some $g \in G$, and, for every $k \geq 1$, let $L(k,g)$ be a sufficiently large integer so that $\{gg_1, \dots, gg_k\} \subset \{g_1, \dots, g_{L(k,g)}\}$. Of course, $L(k,g) \geq k$ and $gI_k = \bigcap_{i=1}^{k} gg_i J^+ \supset I_{L(k,g)}$. Consequently,

$$g \bigcap_{k \geq 1} \partial^c I_k = \bigcap_{k \geq 1} \partial^c (gI_k) \supset \bigcap_{k \geq 1} \partial^c I_{L(k,g)} = \bigcap_{k \geq 1} \partial^c I_k,$$

because $L(k,g) \to_{k \to \infty} +\infty$. Thus, we have proved that $g \bigcap_{k \geq 1} \partial^c I_k \supset \bigcap_{k \geq 1} \partial^c I_k$ for every $g \in G$. It follows that our intersection is G–invariant.

Finally, because G acts on $\partial^c X$ by order-automorphisms, we conclude that G fixes $r \in \partial^c X$. □

Corollary 4.26 *Let G be a countable group acting minimally on a* CAT(0) *cube complex X without fixing any point of $\partial^c X$. Then any pair of disjoint hyperplanes is skewered by an element of G.*

Proof Let (J,H) be a pair of disjoint hyperplanes. Fix two orientations $J = (J^-, J^+)$ and $H = (H^-, H^+)$ so that $J^+ \subsetneq H^+$. Now, by applying the previous lemma twice, we find two elements $g, h \in G$ such that $gJ^- \subsetneq J^+$ and $hgH^+ \subsetneq gH^-$. Thus,

$$hgH^+ \subsetneq gH^- = X \setminus gH^+ \subsetneq X \setminus gJ^+ = gJ^- \subsetneq J^+.$$

We deduce that hg skewers the pair (J,H). □

Remark 4.27 Lemma 4.25 may be thought of as a generalization of the flipping lemma proved in [5], where the Tits boundary is replaced with the combinatorial boundary in possibly infinite dimension. Corollary 4.26 corresponds to the double skewering lemma [5].

Proof of Theorem 4.23 If G contains a contracting isometry, then $\partial^c X$ contains an isolated point according to Theorem 4.17. Conversely, suppose that $\partial^c X$ contains an isolated point. It follows from Lemma 4.19 that X contains two well-separated hyperplanes, and it suffices to invoke Corollary 4.26 to find an element $g \in G$ which skewers this pair of hyperplanes: according to Theorem 3.13, this is a contracting isometry. □

5 Acylindrical hyperbolicity of diagram groups

5.1 Diagram groups and their cube complexes

We refer to [13, Sections 3 and 5] for a detailed introduction to *semigroup diagrams* and *diagram groups*.

For an alphabet Σ, let Σ^+ denote the free semigroup over Σ. If $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ is a semigroup presentation, where $\mathcal{R}$ is a set of pairs of words in Σ^+, the semigroup associated to $\mathcal{P}$ is the one given by the factor semigroup $\Sigma^+/\sim$, where $\sim$ is the smallest equivalent relation on Σ^+ containing $\mathcal{R}$. We will always assume that if $u = v \in \mathcal{R}$ then $v = u \notin \mathcal{R}$; in particular, $u = u \notin \mathcal{R}$.

A *semigroup diagram over* $\mathcal{P}$ is the analogue for semigroups to van Kampen diagrams for group presentations. Formally, it is a finite connected planar graph Δ whose edges are oriented and labeled by the alphabet Σ, satisfying the following properties:

- Δ has exactly one vertex-source ι (which has no incoming edges) and exactly one vertex-sink τ (which has no outgoing edges);
- the boundary of each cell has the form pq^{-1}, where $p = q$ or $q = p \in \mathcal{R}$;
- every vertex belongs to a positive path connecting ι and τ;
- every positive path in Δ is simple.

In particular, Δ is bounded by two positive paths: the *top path*, denoted by $\mathrm{top}(\Delta)$, and the *bottom path*, denoted by $\mathrm{bot}(\Delta)$. By extension, we also define $\mathrm{top}(\Gamma)$ and $\mathrm{bot}(\Gamma)$ for every *subdiagram* Γ. In the sequel, the notation $\mathrm{top}(\cdot)$ and $\mathrm{bot}(\cdot)$ will refer to the paths or to their labels depending on the context. Also, a (u, v)*–cell* (resp. a (u, v)*–diagram*) will refer to a cell (resp. a semigroup diagram) whose top path is labeled by u and whose bottom path is labeled by v.

Two words w_1 and w_2 in Σ^+ are *equal modulo* $\mathcal{P}$ if their images in the semigroup associated to $\mathcal{P}$ coincide. In particular, there exists a *derivation* from w_1 to w_2, i.e.

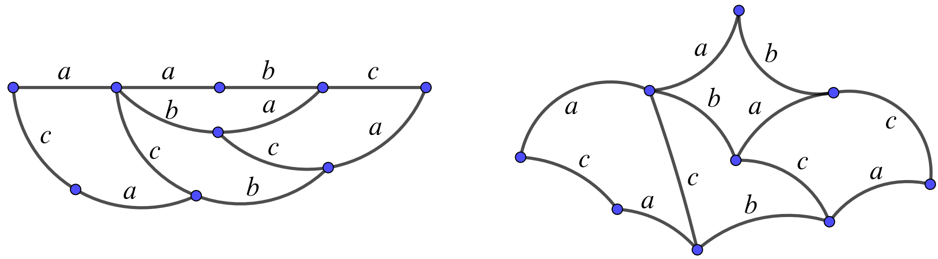

Figure 4: Two isotopic semigroup diagrams.

a sequence of relations of $\mathcal{R}$ allowing us to transform w_1 into w_2. To any such derivation is associated a semigroup diagram, or more precisely a (w_1, w_2)–diagram, whose construction is clear from the example below. As in the case of groups, the words w_1 and w_2 are equal modulo $\mathcal{P}$ if and only if there exists a (w_1, w_2)–diagram.

Example 5.1 Let $\mathcal{P} = \langle a, b, c \mid ab = ba,\ ac = ca,\ bc = cb \rangle$ be a presentation of the free abelian semigroup of rank three. In particular, the words a^2bc and $caba$ are equal modulo $\mathcal{P}$, with the possible derivation

$$aabc \xrightarrow{(a,ab\to ba,c)} abac \xrightarrow{(ab,ac\to ca,\varnothing)} abca \xrightarrow{(a,bc\to cb,a)} acba \xrightarrow{(\varnothing,ac\to ca,ba)} caba.$$

Then, the associated $(a^2bc, caba)$–diagram Δ is illustrated by Figure 4. On such a graph, the edges are supposed oriented from left to right. Here, the diagram Δ has nine vertices, twelve edges and four cells; notice that the number of cells of a diagram corresponds to the length of the associated derivation. The paths $\mathrm{top}(\Delta)$ and $\mathrm{bot}(\Delta)$ are labeled respectively by a^2bc and $caba$.

Since we are only interested in the combinatorics of semigroup diagrams, we will not distinguish isotopic diagrams. For example, the two diagrams of Figure 4 will be considered as equal. If $w \in \Sigma^+$, we define the *trivial diagram* $\epsilon(w)$ as the semigroup diagram without cells whose top and bottom paths, labeled by w, coincide. Any diagram without cells is trivial. A diagram with exactly one cell is *atomic*.

If Δ_1 is a (w_1, w_2)–diagram and Δ_2 a (w_2, w_3)–diagram, we define the *concatenation* $\Delta_1 \circ \Delta_2$ as the semigroup diagram obtained by identifying the bottom path of Δ_1 with the top path of Δ_2. In particular, $\Delta_1 \circ \Delta_2$ is a (w_1, w_3)–diagram. Thus, $\circ$ defines a partial operation on the set of semigroup diagrams over $\mathcal{P}$. However, restricted to the subset of (w, w)–diagrams for some $w \in \Sigma^+$, it defines a semigroup operation; such diagrams are called *spherical with base w*. We also define the *sum* $\Delta_1 + \Delta_2$ of

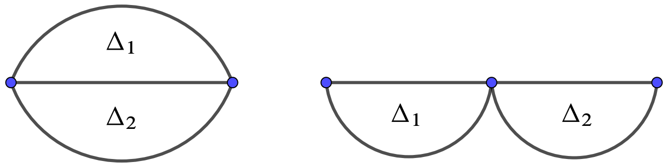

Figure 5: The concatenation and the sum of Δ_1 and Δ_2.

two diagrams Δ_1 and Δ_2 as the diagram obtained by identifying the rightmost vertex of Δ_1 with the leftmost vertex of Δ_2. See Figure 5.

Notice that any semigroup diagram can be viewed as a concatenation of atomic diagrams. In the following, if Δ_1 and Δ_2 are two diagrams, we will say that Δ_1 is a *prefix* (resp. a *suffix*) of Δ_2 if there exists a diagram Δ_3 satisfying $\Delta_2 = \Delta_1 \circ \Delta_3$ (resp. $\Delta_2 = \Delta_3 \circ \Delta_1$). Throughout this paper, the fact that Δ is a prefix of Γ will be denoted by $\Delta \leq \Gamma$.

Suppose that a diagram Δ contains a (u,v)–cell and a (v,u)–cell such that the top path of the first cell is the bottom path of the second cell. Then, we say that these two cells form a *dipole*. In this case, we can remove these two cells by first removing their common path, and then identifying the bottom path of the first cell with the top path of the second cell; thus, we *reduce the dipole*. A diagram is called *reduced* if it does not contain dipoles. By reducing dipoles, a diagram can be transformed into a reduced diagram, and a result of Kilibarda [21] proves that this reduced form is unique. If Δ_1 and Δ_2 are two diagrams for which $\Delta_1 \circ \Delta_2$ is well defined, let us denote by $\Delta_1 \cdot \Delta_2$ the reduced form of $\Delta_1 \circ \Delta_2$.

Thus, the set of reduced semigroup diagrams, endowed with the product $\cdot$ we defined above, naturally defines a groupoid $\mathcal{G}(\mathcal{P})$, i.e. , loosely speaking, a "group" where the product is only partially defined. The neutral elements correspond to the trivial diagrams, and the inverse of a diagram is constructed in the following way: if Δ is a (w_1, w_2)–diagram, its inverse Δ^{-1} is the (w_2, w_1)–diagram obtained from Δ by a mirror reflection with respect to $\mathrm{top}(\Delta)$. A natural way to deduce a group from this groupoid is to fix a baseword $w \in \Sigma^+$, and to define the *diagram group* $D(\mathcal{P}, w)$ as the set of reduced (w,w)–diagrams endowed with the product $\cdot$ we defined above.

Farley cube complexes The groupoid $\mathcal{G}(\mathcal{P})$ has a natural generating set, namely the set of atomic diagrams, and the group $D(\mathcal{P}, w)$ acts by left-multiplication on the connected component of the associated Cayley graph which contains $\epsilon(w)$. Surprisingly, this Cayley graph turns out to be naturally the 1–skeleton of a CAT(0) cube complex. Below, we detail the construction of this cube complex as explained in [8].

A semigroup diagram is *thin* whenever it can be written as a sum of atomic diagrams. We define the *Farley complex* $X(\mathcal{P},w)$ as the cube complex whose vertices are the reduced semigroup diagrams Δ over $\mathcal{P}$ satisfying $\mathrm{top}(\Delta)=w$, and whose n–cubes are spanned by the vertices $\{\Delta \cdot P \mid P \leq \Gamma\}$ for some vertex Δ and some thin diagram Γ with n cells. In particular, two diagrams Δ_1 and Δ_2 are linked by an edge if and only if there exists an atomic diagram A such that $\Delta_1 = \Delta_2 \cdot A$.

Theorem 5.2 [8, Theorem 3.13] *$X(\mathcal{P},w)$ is a* CAT(0) *cube complex. Moreover it is complete, i.e. there are no infinite increasing sequences of cubes in $X(\mathcal{P},w)$.*

There is a natural action of $D(\mathcal{P},w)$ on $X(\mathcal{P},w)$, namely $(g,\Delta)\mapsto g\cdot\Delta$. Then:

Proposition 5.3 [8, Theorem 3.13] *The action $D(\mathcal{P},w) \curvearrowright X(\mathcal{P},w)$ is free. Moreover, it is cocompact if and only if the class $[w]_{\mathcal{P}}$ of words equal to w modulo $\mathcal{P}$ is finite.*

In this paper, we will always suppose that the cube complex $X(\mathcal{P},w)$ is locally finite; in particular, this implies that the action $D(\mathcal{P},w) \curvearrowright X(\mathcal{P},w)$ is properly discontinuous. For instance, this is the case if $\mathcal{P}$ is a finite presentation, and because we may always suppose that $\mathcal{P}$ is a finite presentation if our diagram group is finitely generated, this assumption is not really restrictive. However, notice that $X(\mathcal{P},w)$ may be locally finite even if $\mathcal{P}$ is an infinite presentation. For example, the derived subgroup $[F,F]$ of Thompson's group F is isomorphic to the diagram group $D(\mathcal{P},a_0b_0)$, where

$$\mathcal{P}=\langle x,a_i,b_i,i\geq 0 \mid x=x^2,\ a_i=a_{i+1}x,\ b_i=xb_{i+1},\ i\geq 0\rangle.$$

See [14, Theorem 26] for more details. In this situation, $X(\mathcal{P},a_0b_0)$ is nevertheless locally finite.

Because $X(\mathcal{P},w)$ is essentially a Cayley graph, the combinatorial geodesics behave as expected:

Lemma 5.4 [10, Lemma 3] *Let Δ be a reduced diagram. If $\Delta = A_1\circ\cdots\circ A_n$, where each A_i is an atomic diagram, then the path*

$$\epsilon(w),\ A_1,\ A_1\circ A_2,\ldots,\ A_1\circ\cdots\circ A_n=\Delta$$

defines a combinatorial geodesic from $\epsilon(w)$ to Δ. Furthermore, every combinatorial geodesic from $\epsilon(w)$ to Δ has this form.

Corollary 5.5 [10, Corollary 3] *Let $A, B \in X(\mathcal{P}, w)$ be two reduced diagrams. Then we have*

$$d(A, B) = \#(A^{-1} \cdot B),$$

where $\#(\cdot)$ denotes the number of cells of a diagram. In particular, $d(\epsilon(w), A) = \#(A)$.

In [10], we describe the hyperplanes of $X(\mathcal{P}, w)$. We recall this description below. For any hyperplane J of $X(\mathcal{P}, w)$, we introduce the following notation:

- J^+ is the halfspace associated to J not containing $\epsilon(w)$,
- J^- is the halfspace associated to J containing $\epsilon(w)$,
- $\partial_\pm J$ is the intersection $\partial N(J) \cap J^\pm$.

Definition 5.6 A diagram Δ is *minimal* if its maximal thin suffix F has exactly one cell. (The existence and uniqueness of the maximal thin suffix is given by [8, Lemma 2.3].)

In the following, $\overline{\Delta}$ will denote the diagram $\Delta \cdot F^{-1}$, obtained from Δ by removing the suffix F. The following result uses minimal diagrams to describe hyperplanes in Farley complexes:

Proposition 5.7 [10, Proposition 2] *Let J be a hyperplane of $X(\mathcal{P}, w)$. Then there exists a unique minimal diagram Δ such that $J^+ = \{D \text{ diagram} \mid \Delta \leq D\}$. Conversely, if Δ is a minimal diagram and J the hyperplane dual to the edge $[\Delta, \overline{\Delta}]$, then $J^+ = \{D \text{ diagram} \mid \Delta \leq D\}$.*

Thanks to the geometry of CAT(0) cube complexes, we will use this description of the hyperplanes of $X(\mathcal{P}, w)$ in order to define the *supremum*, when it exists, of a collection of minimal diagrams. If Δ is a diagram, let $\mathcal{M}(\Delta)$ denote the set of its minimal prefixes. A set of minimal diagrams $\mathcal{D}$ is *consistent* if, for every $D_1, D_2 \in \mathcal{D}$, there exists some $D_3 \in \mathcal{D}$ satisfying $D_1, D_2 \leq D_3$.

Proposition 5.8 *Let $\mathcal{D}$ be a finite consistent collection of reduced minimal $(w, *)$–diagrams. Then there exists a unique $\leq$–minimal reduced $(w, *)$–diagram admitting any element of $\mathcal{D}$ as a prefix; let $\sup(\mathcal{D})$ denote this diagram. Furthermore, $\mathcal{M}(\sup(\mathcal{D})) = \bigcup_{D \in \mathcal{D}} \mathcal{M}(D)$.*

Proof Let $\mathcal{D} = \{D_1, \dots, D_r\}$. For every $1 \leq i \leq r$, let J_i denote the hyperplane associated to D_i. Because $\mathcal{D}$ is consistent, for every $1 \leq i, j \leq r$ the intersection $J_i^+ \cap J_j^+$ is nonempty, hence $C = \bigcap_{i=1}^r J_i^+ \neq \varnothing$. Let Δ denote the combinatorial projection of $\epsilon(w)$ onto C. For every $1 \leq i \leq r$, Δ belongs to J_i^+, hence $D_i \leq \Delta$. Then, if D is another diagram admitting any element of $\mathcal{D}$ as a prefix, necessarily $D \in C$, and it follows from Proposition 2.6 that there exists a combinatorial geodesic between $\epsilon(w)$ and D passing through Δ, hence $\Delta \leq D$. This proves the $\leq$–minimality and the uniqueness of Δ.

Now, let $D \in \mathcal{M}(\Delta)$. If J denotes the hyperplane associated to D, then J separates $\epsilon(w)$ and Δ, and according to Lemma 2.8 it must separate $\epsilon(w)$ and C. Therefore, either J belongs to $\{J_1, \dots, J_r\}$ or it separates $\epsilon(w)$ and some J_k; equivalently, either D belongs to $\{D_1, \dots, D_r\}$ or it is a prefix of some D_k. Hence $D \in \bigcup_{i=1}^r \mathcal{M}(D_i)$. Thus, we have proved that $\mathcal{M}(\Delta) \subset \bigcup_{D \in \mathcal{D}} \mathcal{M}(D)$; the inclusion $\bigcup_{D \in \mathcal{D}} \mathcal{M}(D) \subset \mathcal{M}(\Delta)$ is clear since any element of $\mathcal{D}$ is a prefix of Δ. □

As a consequence, we deduce that a diagram is determined by the set of its minimal prefixes.

Lemma 5.9 *Let Δ_1 and Δ_2 be two diagrams. If $\mathcal{M}(\Delta_1) = \mathcal{M}(\Delta_2)$ then $\Delta_1 = \Delta_2$.*

Proof We begin with some general results. Let Δ be a diagram.

Fact 5.10 $\#\mathcal{M}(\Delta) = \#\Delta$.

Indeed, there exists a bijection between $\mathcal{M}(\Delta)$ and the hyperplanes separating $\epsilon(w)$ and Δ, hence $\#\mathcal{M}(\Delta) = d(\epsilon(w), \Delta) = \#\Delta$. We deduce that:

Fact 5.11 $\sup(\mathcal{M}(\Delta)) = \Delta$.

Indeed, by definition any element of $\mathcal{M}(\Delta)$ is a prefix of Δ, so that $\sup(\mathcal{M}(\Delta))$ exists (because $\mathcal{M}(\Delta)$ is consistent) and $\sup(\mathcal{M}(\Delta)) \leq \Delta$ by the $\leq$–minimality of $\sup(\mathcal{M}(\Delta))$. On the other hand, we deduce from the previous fact that

$$\#\sup(\mathcal{M}(\Delta)) = \left|\mathcal{M}\big(\sup(\mathcal{M}(\Delta))\big)\right| = \left| \bigcup_{D \in \mathcal{M}(\Delta)} \mathcal{M}(D) \right| = \#\mathcal{M}(\Delta) = \#\Delta.$$

Therefore, we necessarily have $\sup(\mathcal{M}(\Delta)) = \Delta$.

Finally, our lemma follows easily, since $\Delta_1 = \sup(\mathcal{M}(\Delta_1)) = \sup(\mathcal{M}(\Delta_2)) = \Delta_2$. □

Axes of isometries In Sections 3 and 4, we always assumed that our CAT(0) cube complexes are such that any isometry admits an axis. According to Corollary 2.25, we may suppose that this is the case up to subdividing the cube complex. Here, we notice that such a subdivision is not necessary when looking at a diagram group acting on its Farley complex. More precisely:

Proposition 5.12 *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. Every nontrivial element of $D(\mathcal{P}, w)$ admits a combinatorial axis in $X(\mathcal{P}, w)$.*

Before turning to the proof of the proposition, we need the following definition:

Definition 5.13 Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. A (w, w)–diagram Δ is *cyclically reduced* if the concatenation $\Delta \circ \Delta$ is reduced.

The key lemma to prove our proposition is the following lemma, stating that any element of a diagram group decomposes as a conjugate (in the groupoid of semigroup diagrams) of a cyclically reduced diagram:

Lemma 5.14 [13, Lemma 15.4] *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. For every $g \in D(\mathcal{P}, w)$, there exist two semigroup diagrams A and Δ such that g is equal to the reduced concatenation $A \circ \Delta \circ A^{-1}$ and such that Δ is cyclically reduced.*

Proof of Proposition 5.12 Let $g \in D(\mathcal{P}, w)$ be a nontrivial element. According to Lemma 5.14, there exist two diagrams A and Δ such that g decomposes as the reduced concatenation $A \circ \Delta \circ A^{-1}$ and such that Δ is cyclically reduced. Notice that, as g is nontrivial, the diagram Δ contains at least one cell. Decompose Δ as a reduced concatenation of atomic diagrams $\Delta = \Delta_1 \circ \cdots \circ \Delta_n$, where $n \geq 1$. Let ℓ denote the oriented path

$$A,\ A \circ \Delta_1,\ A \circ \Delta_1 \circ \Delta_2,\ \ldots,\ A \circ \Delta_1 \circ \cdots \circ \Delta_n$$

in $X(\mathcal{P}, w)$, and L be the concatenation

$$\cdots \cup (\Delta^{-2} \cdot \ell) \cup (\Delta^{-1} \cdot \ell) \cup \ell \cup (\Delta \cdot \ell) \cup (\Delta^2 \cdot \ell) \cup \cdots .$$

The orientation of ℓ induces naturally an orientation on L. It is clear that g acts on the path L as a translation of length $\#\Delta$. We claim that L is a geodesic.

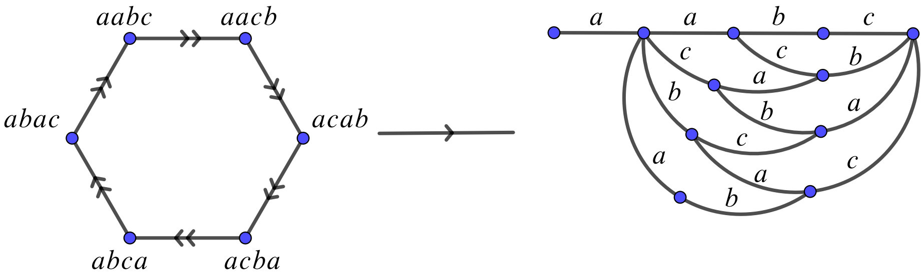

Figure 6: A loop in $S(\mathcal{P})$ and its associated semigroup diagram.

For every $k \geq 1$, let L_k denote the subpath

$$(\Delta^{-k} \cdot \ell) \cup \cdots \cup (\Delta^{-1} \cdot \ell) \cup \ell \cup (\Delta \cdot \ell) \cup \cdots \cup (\Delta^k \cdot \ell).$$

In order to show that L is a geodesic, it is sufficient to show that the L_k are geodesics. Notice that, for every $k \geq 1$, the path L_k is labeled by the concatenation $(\Delta_1 \circ \cdots \circ \Delta_n)^k$. Since Δ is cyclically reduced, we know that this concatenation is reduced, so that Lemma 5.4 implies that L_k is indeed a geodesic, which concludes the proof. □

Squier cube complexes Because the action $D(\mathcal{P}, w) \curvearrowright X(\mathcal{P}, w)$ is free and properly discontinuous, we know that $D(\mathcal{P}, w)$ is isomorphic to the fundamental group of the quotient space $X(\mathcal{P}, w)/D(\mathcal{P}, w)$. We give now a description of this quotient. We define the *Squier complex* $S(\mathcal{P})$ as the cube complex whose vertices are the words in Σ^+; whose (oriented) edges can be written as $(a, u \to v, b)$, where $u = v$ or $v = u$ belongs to $\mathcal{R}$, linking the vertices aub and avb; and whose n–cubes similarly can be written as

$$(a_1,\ u_1 \to v_1,\ \ldots,\ a_n,\ u_n \to v_n,\ a_{n+1}),$$

spanned by the set of vertices $\{a_1 w_1 \cdots a_n w_n a_{n+1} \mid w_i = v_i \text{ or } u_i\}$.

Then, there is a natural morphism from the fundamental group of $S(\mathcal{P})$ based at w to the diagram group $D(\mathcal{P}, w)$. Indeed, a loop in $S(\mathcal{P})$ based at w is just a series of relations of $\mathcal{R}$ applied to the word w so that the final word is again w, and such a derivation may be encoded into a semigroup diagram. Figure 6 shows an example, where the semigroup presentation is $\mathcal{P} = \langle a, b, c \mid ab = ba,\ bc = cb,\ ac = ca \rangle$.

Thus, this defines a map from the set of loops of $S(\mathcal{P})$ based at w to the set of spherical semigroup diagrams. In fact, the map extends to a morphism which turns out to be an isomorphism:

Theorem 5.15 [13, Theorem 6.1] $D(\mathcal{P}, w) \simeq \pi_1(S(\mathcal{P}), w)$.

For convenience, $S(\mathcal{P}, w)$ will denote the connected component of $S(\mathcal{P})$ containing w. Notice that two words $w_1, w_2 \in \Sigma^+$ are equal modulo $\mathcal{P}$ if and only if they belong to the same connected component of $S(\mathcal{P})$. Therefore, a consequence of Theorem 5.15 is:

Corollary 5.16 *If $w_1, w_2 \in \Sigma^+$ are equal modulo $\mathcal{P}$, then there exists a (w_2, w_1)–diagram Γ and the map*

$$\Delta \mapsto \Gamma \cdot \Delta \cdot \Gamma^{-1}$$

induces an isomorphism from $D(\mathcal{P}, w_1)$ to $D(\mathcal{P}, w_2)$.

As claimed, we note that the map $\Delta \mapsto \mathrm{bot}(\Delta)$ induces a universal covering $X(\mathcal{P}, w) \to S(\mathcal{P}, w)$, so that the action of $\pi_1(S(\mathcal{P}, w))$ on $X(\mathcal{P}, w)$ coincides with the natural action of $D(\mathcal{P}, w)$.

Lemma 5.17 [10, Lemma 2] *The map $\Delta \mapsto \mathrm{bot}(\Delta)$ induces a cellular isomorphism from the quotient $X(\mathcal{P}, w)/D(\mathcal{P}, w)$ to $S(\mathcal{P}, w)$.*

Changing the baseword When we work on the Cayley graph of a group, because the action of the group is vertex-transitive, we may always suppose that a fixed basepoint is the identity element up to a conjugation. In our situation, where $X(\mathcal{P}, w)$ is the Cayley graph of a groupoid, the situation is slightly different but it is nevertheless possible to make something similar. Indeed, if $\Delta \in X(\mathcal{P}, w)$ is a basepoint, then conjugating by Δ sends Δ to a trivial diagram, but which does not necessarily belong to $X(\mathcal{P}, w)$: in general, it will belong to $X(\mathcal{P}, \mathrm{bot}(\Delta))$. In the same time, $D(\mathcal{P}, w)$ becomes $D(\mathcal{P}, \mathrm{bot}(\Delta))$, which is naturally isomorphic to $D(\mathcal{P}, w)$ according to Corollary 5.16. Finally, we get a commutative diagram

$$\begin{tikzcd} D(\mathcal{P}, w) \arrow[r, "\text{isomorphism}"] \arrow[d] & D(\mathcal{P}, \mathrm{bot}(\Delta)) \arrow[d] \\ X(\mathcal{P}, w) \arrow[r, "\text{isometry}"] & X(\mathcal{P}, \mathrm{bot}(\Delta)) \end{tikzcd}$$

Thus, we may always suppose that a fixed basepoint is a trivial diagram up to changing the baseword, which does not disturb either the group or the cube complex. In particular, a contracting isometry stays a contracting isometry after the process.

Infinite diagrams Basically, an infinite diagram is just a diagram with infinitely many cells. To be more precise, if finite diagrams are thought of as specific planar graphs, then an infinite diagram is the union of an increasing sequence of finite diagrams with the same top path. Formally:

Definition 5.18 An *infinite diagram* is a formal concatenation $\Delta_1 \circ \Delta_2 \circ \cdots$ of infinitely many diagrams $\Delta_1, \Delta_2, \dots$ up to the following identification: two formal concatenations $A_1 \circ A_2 \circ \cdots$ and $B_1 \circ B_2 \circ \cdots$ define the same infinite diagram if, for every $i \geq 1$, there exists some $j \geq 1$ such that $A_1 \circ \cdots \circ A_i$ is a prefix of $B_1 \circ \cdots \circ B_j$, and conversely, for every $i \geq 1$, there exists some $j \geq 1$ such that $B_1 \circ \cdots \circ B_i$ is a prefix of $A_1 \circ \cdots \circ A_j$.

An infinite diagram $\Delta = \Delta_1 \circ \Delta_2 \circ \cdots$ is *reduced* if $\Delta_1 \circ \cdots \circ \Delta_n$ is reduced for every $n \geq 1$. If $\mathrm{top}(\Delta_1) = w$, we say that Δ is an *infinite w–diagram*.

Notice that, according to Lemma 5.4, a combinatorial ray is naturally labeled by a reduced infinite diagram. This explains why infinite diagrams will play a central role in the description of the combinatorial boundaries of Farley cube complexes.

Definition 5.19 Let $\Delta = \Delta_1 \circ \Delta_2 \circ \cdots$ and $\Xi = \Xi_1 \circ \Xi_2 \circ \cdots$ be two infinite diagrams. We say that Δ is a prefix of Ξ if, for every $i \geq 1$, there exists some $j \geq 1$ such that $\Delta_1 \circ \cdots \circ \Delta_i$ is a prefix of $\Xi_1 \circ \cdots \Xi_j$.

5.2 Combinatorial boundaries of Farley complexes

The goal of this section is to describe the combinatorial boundaries of Farley cube complexes as a set of infinite reduced diagrams. First, we need to introduce a partial order on the set of infinite diagrams. Basically, we will say that a diagram Δ_1 is *almost a prefix* of Δ_2 if there exists a third diagram Δ_0 which is a prefix of Δ_1 and Δ_2 such that all but finitely many cells of Δ_1 belong to Δ_0. Notice that, if we fix a decomposition of a diagram Δ as a concatenation of atomic diagrams, then to any cell π of Δ corresponds naturally a minimal prefix of Δ, namely the smallest prefix containing π (which is precisely the union of all the cells of Δ which are "above" π; more formally, using the relation introduced by Definition 5.35 below, this is the union of all the cells π' of Δ satisfying $\pi' \lll \pi$). Therefore, alternatively we can describe our order in terms of minimal prefixes.

Definition 5.20 Let Δ_1 and Δ_2 be two infinite reduced diagrams. Then

- Δ_1 is *almost a prefix* of Δ_2 if $\mathcal{M}(\Delta_1) \subset_a \mathcal{M}(\Delta_2)$;
- Δ_1 and Δ_2 are *commensurable* if $\mathcal{M}(\Delta_1) =_a \mathcal{M}(\Delta_2)$.

Notice that two infinite reduced diagrams are commensurable if and only if each one is almost a prefix of the other. Our model for the combinatorial boundaries of Farley cube complexes will be the following:

Definition 5.21 If $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ is semigroup presentation and $w \in \Sigma^+$ a baseword, let $(\partial(\mathcal{P}, w), <_a)$ denote the poset of the infinite reduced $(w,*)$–diagrams, up to commensurability, ordered by the relation $<_a$ of being almost a prefix.

Recall that to any combinatorial ray $r \subset X(\mathcal{P}, w)$ starting from $\epsilon(w)$ is naturally associated an infinite reduced $(w,*)$–diagram $\Psi(r)$.

Theorem 5.22 *The map Ψ induces a poset isomorphism $\partial^c X(\mathcal{P}, w) \to \partial(\mathcal{P}, w)$.*

The main tool to prove our theorem is provided by the following proposition, which is an infinite analogue of Proposition 5.8. Recall that a collection of (finite) diagrams $\mathcal{D}$ is consistent if, for every $\Delta_1, \Delta_2 \in \mathcal{D}$, there exists a third diagram Δ satisfying $\Delta_1, \Delta_2 \leq \Delta$.

Proposition 5.23 *Let $\mathcal{D}$ be a countable consistent collection of finite reduced minimal $(w,*)$–diagrams. Then there exists a unique $\leq$–minimal reduced $(w,*)$–diagram admitting any element of $\mathcal{D}$ as a prefix; let $\sup(\mathcal{D})$ denote this diagram. Furthermore, $\mathcal{M}(\sup(\mathcal{D})) = \bigcup_{D \in \mathcal{D}} \mathcal{M}(D)$.*

Proof Let $\mathcal{D} = \{D_1, D_2, \dots\}$, with $D_i \neq D_j$ if $i \neq j$. For each $i \geq 1$, let H_i denote the hyperplane of $X(\mathcal{P}, w)$ associated to D_i. Because $\mathcal{D}$ is consistent, the halfspaces H_i^+ and H_j^+ intersect for every $i \geq 1$. Therefore, for every $k \geq 1$, the intersection $I_k = \bigcap_{i=1}^k H_i^+$ is nonempty; let Δ_k denote the combinatorial projection of $\epsilon(w)$ onto I_k. As a consequence of Proposition 2.6, there exists a combinatorial ray starting from $\epsilon(w)$ and passing through each Δ_k. In particular, for every $k \geq 1$, Δ_k is a prefix of Δ_{k+1}, so that it makes sense to define the infinite diagram Δ as the limit of the sequence (Δ_k). Notice that, for every $k \geq 1$, $\Delta_k \in H_k^+$ so $D_k \leq \Delta_k$ and a fortiori $D_k \leq \Delta$.

Now, let Ξ be another reduced w–diagram admitting each D_k as a prefix. Fixing a decomposition of Ξ as an infinite concatenation of atomic diagrams, say $A_1 \circ A_2 \circ \cdots$, we associate to Ξ a combinatorial ray r starting from $\epsilon(w)$ defined as the path

$$\epsilon(w),\ A_1,\ A_1 \circ A_2,\ A_1 \circ A_2 \circ A_3,\ \dots.$$

Let $n \geq 1$. Because $D_1, \dots, D_n$ are prefixes of Ξ, there exists some $k \geq 1$ such that $r(k) \in I_n$. On the other hand, according to Proposition 2.6, there exists a combinatorial geodesic between $\epsilon(w)$ and $r(k)$ passing through Δ_n, hence $\Delta_n \leq r(k)$. Therefore, we have proved that, for every $n \geq 1$, Δ_n is a prefix of Ξ: this precisely means that Δ is a prefix of Ξ. The $\leq$–minimality and the uniqueness of Δ is thus proved.

Because the D_k are prefixes of Δ, the inclusion $\bigcup_{D\in\mathcal{D}}\mathcal{M}(D)\subset\mathcal{M}(\Delta)$ is clear. Conversely, let $D\in\mathcal{M}(\Delta)$; let J denote the associated hyperplane. Because D is a prefix of Δ, it is a prefix of Δ_k for some $k\geq 1$; and because Δ_k is the combinatorial projection of $\epsilon(w)$ onto I_k, we deduce that J separates $\epsilon(w)$ and I_k. On the other hand, $I_k=\bigcap_{i=1}^{k}H_i^+$ so there exists some $1\leq i\leq k$ such that either $J=H_i$ or J separates $\epsilon(w)$ and H_i; equivalently, D is a prefix of D_i. We have proved $\mathcal{M}(\Delta)\subset\bigcup_{D\in\mathcal{D}}\mathcal{M}(D)$. □

The following lemma states how the operation sup behaves with respect to the inclusion, the almost-inclusion and the almost-equality:

Lemma 5.24 *Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be two countable consistent collections of reduced $(w,*)$–diagrams.*

(i) *If $\mathcal{D}_1\subset\mathcal{D}_2$ then $\sup(\mathcal{D}_1)\leq\sup(\mathcal{D}_2)$.*

(ii) *If $\mathcal{D}_1\subset_a\mathcal{D}_2$ then $\sup(\mathcal{D}_1)$ is almost a prefix of $\sup(\mathcal{D}_2)$.*

(iii) *If $\mathcal{D}_1=_a\mathcal{D}_2$ then $\sup(\mathcal{D}_1)$ and $\sup(\mathcal{D}_2)$ are commensurable.*

Proof If $\mathcal{D}_1\subset\mathcal{D}_2$ then $\sup(\mathcal{D}_2)$ admits any element of $\mathcal{D}_1$ as a prefix, hence $\sup(\mathcal{D}_1)\leq\sup(\mathcal{D}_2)$ by the $\leq$–minimality of $\sup(\mathcal{D}_1)$.

If $\mathcal{D}_1\subset_a\mathcal{D}_2$ then we can write

$$\mathcal{M}(\sup(\mathcal{D}_1))=\bigcup_{D\in\mathcal{D}_1}\mathcal{M}(D)=\Big(\bigcup_{D\in\mathcal{D}_1\cap\mathcal{D}_2}\mathcal{M}(D)\Big)\sqcup\Big(\bigcup_{D\in\mathcal{D}_1\backslash\mathcal{D}_2}\mathcal{M}(D)\Big).$$

On the other hand, $\mathcal{D}_1\setminus\mathcal{D}_2$ is finite, because $\mathcal{D}_1\subset_a\mathcal{D}_2$, and $\mathcal{M}(D)$ is finite for any finite diagram D, so $\bigcup_{D\in\mathcal{D}_1\backslash\mathcal{D}_2}\mathcal{M}(D)$ is finite. This proves that $\mathcal{M}(\sup(\mathcal{D}_1))\subset_a\mathcal{M}(\sup(\mathcal{D}_2))$, i.e. $\sup(\mathcal{D}_1)$ is almost a prefix of $\sup(\mathcal{D}_2)$.

If $\mathcal{D}_1=_a\mathcal{D}_2$, we deduce by applying the previous point twice that $\sup(\mathcal{D}_1)$ is almost a prefix of $\sup(\mathcal{D}_2)$ and vice versa. Therefore, $\sup(\mathcal{D}_1)$ and $\sup(\mathcal{D}_2)$ are commensurable. □

Finally, Theorem 5.22 will be essentially a consequence of the previous lemma and the following observation. If r is a combinatorial ray in a Farley complex, let $\mathcal{D}(r)$ denote the set of the minimal diagrams which are associated to the hyperplanes of $\mathcal{H}(r)$. Then:

Lemma 5.25 $\Psi(r)=\sup(\mathcal{D}(r))$.

Proof Let $A_1,A_2,A_3,\dots$ be the sequence of atomic diagrams such that

$$r=(\epsilon(w),\ A_1,\ A_1\circ A_2,\ A_1\circ A_2\circ A_3,\ \dots).$$

In particular, $\Psi(r) = A_1 \circ A_2 \circ \cdots$ and $\mathcal{M}(\Psi(r)) = \bigcup_{n \geq 1} \mathcal{M}(A_1 \circ \cdots \circ A_n)$. On the other hand, any $D \in \mathcal{D}(r)$ must be a prefix of $A_1 \circ \cdots \circ A_n$ for some $n \geq 1$, and, conversely, any minimal prefix of some $A_1 \circ \cdots \circ A_n$ must belong to $\mathcal{D}(r)$. Therefore,

$$\mathcal{M}(\Psi(r)) = \bigcup_{n \geq 1} \mathcal{M}(A_1 \circ \cdots \circ A_n) = \mathcal{M}\big(\sup(\mathcal{D}(r))\big).$$

If Δ is a finite prefix of $\Psi(r)$, then $\mathcal{M}(\Delta) \subset \mathcal{M}(\Psi(r)) = \mathcal{M}\big(\sup(\mathcal{D}(r))\big)$, so $\Delta = \sup(\mathcal{M}(\Delta))$ must be a prefix of $\sup(\mathcal{D}(r))$ as well. Therefore, $\Psi(r) \leq \mathcal{M}\big(\sup(\mathcal{D}(r))\big)$. Similarly, we prove that $\mathcal{M}\big(\sup(\mathcal{D}(r))\big) \leq \Psi(r)$, hence $\Psi(r) = \mathcal{M}\big(\sup(\mathcal{D}(r))\big)$. □

Proof of Theorem 5.22 First, we have to verify that, if r_1 and r_2 are two equivalent combinatorial rays starting from $\epsilon(w)$, then $\Psi(r_1) =_a \Psi(r_2)$. Because r_1 and r_2 are equivalent, i.e. $\mathcal{H}(r_1) =_a \mathcal{H}(r_2)$, we know that $\mathcal{D}(r_1) =_a \mathcal{D}(r_2)$, which implies

$$\Psi(r_1) = \sup(\mathcal{D}(r_1)) =_a \sup(\mathcal{D}(r_2)) = \Psi(r_2)$$

according to Lemmas 5.24 and 5.25. Therefore, Ψ induces a map $\partial^c X(\mathcal{P}, w) \to \partial(\mathcal{P}, w)$. For convenience, this map will also be denoted by Ψ.

If $r_1 \prec r_2$, then $\mathcal{H}(r_1) \subset_a \mathcal{H}(r_2)$, hence

$$\Psi(r_1) = \sup(\mathcal{D}(r_1)) <_a \sup(\mathcal{D}(r_2)) = \Psi(r_2).$$

Therefore, Ψ is a poset morphism.

Let Δ be an infinite reduced $(w, *)$–diagram. Let $A_1, A_2, \dots$ be a sequence of atomic diagrams such that $\Delta = A_1 \circ A_2 \circ \cdots$. Let ρ denote the combinatorial ray $(\epsilon(w),\ A_1,\ A_1 \circ A_2,\ \dots)$. Now it is clear that $\Psi(\rho) = A_1 \circ A_2 \circ \cdots = \Delta$, so Ψ is surjective.

Let r_1 and r_2 be two combinatorial rays starting from $\epsilon(w)$ and satisfying $\Psi(r_1) =_a \Psi(r_2)$. It follows from Lemma 5.25 that $\sup(\mathcal{D}(r_1)) =_a \sup(\mathcal{D}(r_2))$. This means

$$\bigcup_{D \in \mathcal{D}(r_1)} \mathcal{M}(D) = \mathcal{M}\big(\sup(\mathcal{D}(r_1))\big) =_a \mathcal{M}\big(\sup(\mathcal{D}(r_2))\big) = \bigcup_{D \in \mathcal{D}(r_2)} \mathcal{M}(D).$$

Notice that $\mathcal{D}(r_2)$ is stable under taking a prefix. Indeed, if P is a prefix of some $\Delta \in \mathcal{D}(r_2)$ then the hyperplane J associated to P separates $\epsilon(w)$ and the hyperplane H associated to Δ; because $H \in \mathcal{H}(r_2)$ and $r_2(0) = \epsilon(w)$, necessarily $J \in \mathcal{H}(r_2)$, hence $P \in \mathcal{D}(r_2)$. As a consequence, we deduce that

$$\mathcal{D}(r_1) \setminus \mathcal{D}(r_2) \subset \Big(\bigcup_{D \in \mathcal{D}(r_1)} \mathcal{M}(D) \Big) \setminus \Big(\bigcup_{D \in \mathcal{D}(r_2)} \mathcal{M}(D) \Big).$$

We conclude that $\mathcal{D}(r_1) \setminus \mathcal{D}(r_2)$ is finite. We prove similarly that $\mathcal{D}(r_2) \setminus \mathcal{D}(r_1)$ is finite. Therefore, we have $\mathcal{D}(r_1) =_a \mathcal{D}(r_2)$, which implies $\mathcal{H}(r_1) =_a \mathcal{H}(r_2)$, so that r_1 and r_2 are equivalent. We have proved that Ψ is injective. □

Example 5.26 Let $\mathcal{P} = \langle x \mid x = x^2 \rangle$ be the semigroup presentation usually associated to Thompson's group F. An atomic diagram is called *positive* if the associated relation is $x \to x^2$, and *negative* if the associated relation is $x^2 \to x$; by extension, a diagram is *positive* (resp. *negative*) if it is a concatenation of positive (resp. negative) atomic diagrams. Notice that any reduced infinite $(x,*)$–diagram has an infinite positive prefix. On the other hand, there exists a reduced infinite positive $(x,*)$–diagram containing all the possible reduced infinite positive $(x,*)$–diagrams as prefixes: it corresponds to the derivation obtained by replacing at each step each x by x^2. Therefore, $\partial(\mathcal{P},x)$ does not contain any isolated vertex. We deduce from Theorem 4.17 that the automorphism group $\mathrm{Aut}(X(\mathcal{P},x))$ does not contain any contracting isometry.

Example 5.27 Let

$$\mathcal{P} = \langle a, b, p_1, p_2, p_3 \mid a = ap_1,\ b = p_1 b,\ p_1 = p_2,\ p_2 = p_3,\ p_3 = p_1 \rangle$$

be the semigroup presentation usually associated to the lamplighter group $\mathbb{Z} \wr \mathbb{Z}$. We leave as an exercise to notice that $\partial(\mathcal{P},ab)$ has no isolated point. Therefore, $\mathrm{Aut}(X(\mathcal{P},ab))$ does not contain any contracting isometry according to Theorem 4.17.

5.3 Contracting isometries

In this section, we fix a semigroup presentation $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ and baseword $w \in \Sigma^+$. Our goal is to determine precisely when a spherical diagram $g \in D(\mathcal{P},w)$ induces a contracting isometry on the Farley complex $X(\mathcal{P},w)$. For convenience, we will suppose that g is *absolutely reduced*, i.e. for every $n \geq 1$ the concatenation of n copies of g is reduced; this assumption is not really restrictive since, according to [13, Lemma 15.10], any spherical diagram is conjugate to an absolutely reduced spherical diagram (possibly for a different baseword). In particular, we may define the infinite reduced $(w,*)$–diagram g^∞ as $g \circ g \circ \cdots$. Our main criterion is:

Theorem 5.28 *Let $g \in D(\mathcal{P},w) - \{\epsilon(w)\}$ be absolutely reduced. Then g is a contracting isometry of $X(\mathcal{P},w)$ if and only if the following two conditions are satisfied:*

(i) *g^∞ does not contain any infinite proper prefix;*

(ii) *if Δ is a reduced diagram with g^∞ as a prefix, then Δ is commensurable to g^∞.*

This theorem will be essentially a consequence of the following two lemmas.

Lemma 5.29 *Let $g \in D(\mathcal{P}, w) - \{\epsilon(w)\}$ be absolutely reduced. Choose a combinatorial geodesic $[\epsilon(w), g]$ between $\epsilon(w)$ and g, and set $\gamma = \bigcup_{n \in \mathbb{Z}} g^n \cdot [\epsilon(w), g]$. Then γ is a bi-infinite combinatorial geodesic on which $\langle g \rangle$ acts by translation. Moreover, $\gamma(+\infty) = g^\infty$.*

Proof Notice that γ is a bi-infinite combinatorial path passing through g^n for every $n \in \mathbb{Z}$. Therefore, in order to deduce that γ is a geodesic, it is sufficient to show that, for every $n, m \in \mathbb{Z}$, the length of γ between g^n and g^{n+m} is equal to $d(g^n, g^{n+m})$; but this length is precisely

$$m \cdot \text{length}([\epsilon(w), g]) = m \cdot d(\epsilon(w), g) = m \cdot \#(g),$$

and, on the other hand, because g is absolutely reduced,

$$d(g^n, g^{n+m}) = \#(g^{-n} \cdot g^{n+m}) = \#(g^m) = m \cdot \#(g).$$

We conclude that γ is a combinatorial geodesic. The fact that $\langle g \rangle$ acts on γ by translation is clear. Finally, we deduce that $\gamma(+\infty) = g^\infty$ from the fact that, for any $k \geq 1$, we have $\gamma(k \cdot \#(g)) = g^k$. □

Definition 5.30 Let Δ be a possibly infinite $(w, *)$–diagram. The *support* of Δ is the set of maximal subpaths of w whose edges belong to the top path of some cell of Δ.

Lemma 5.31 *Let $\Delta_1, \dots, \Delta_n$ be copies of an absolutely reduced spherical (w, w)–diagram Δ, where $n > |w|$. If $e \subset \Delta_1 \circ \cdots \circ \Delta_n$ which belongs to the bottom path of some cell of Δ_1, then e belongs to the top path of some cell in $\Delta_1 \circ \cdots \circ \Delta_n$.*

Proof Say that an edge of Δ_i is a *cutting edge* if it belongs to the bottom path of some cell of Δ_i but does not belong to the top path of some cell in $\Delta_1 \circ \cdots \circ \Delta_n$. Suppose that there exists a cutting edge e_1 in Δ_1.

In particular, e_1 is an edge of the top path of Δ_2, so e_1 does not belong to the support of Δ_2. We deduce that Δ_2 decomposes as a sum $\Phi_2 + \epsilon + \Psi_2$, where ϵ is a trivial diagram with top and bottom paths equal to e_1. Let $\Phi_1 + \epsilon + \Psi_1$ denote the same decomposition of Δ_1. Because e_1 belongs to the bottom path of a cell of Δ_1, two cases may happen: either e_1 belongs to $\text{bot}(\Phi_1)$, so that $\text{top}(\Phi_2) \subsetneq \text{bot}(\Phi_1)$, or e_1 belongs to $\text{bot}(\Psi_1)$, so that $\text{top}(\Psi_2) \subsetneq \text{bot}(\Psi_1)$. Say we are in the former case, the latter case being similar. Now, because Δ_2 and Δ_1 are two copies of the same diagram, to the edge e_1 of Δ_1 corresponds an edge e_2 of Δ_2. Notice that, because e_1 belongs to $\mathrm{bot}(\Phi_1)$, necessarily $e_2 \neq e_1$. Moreover, the product $\Delta^{-1} \circ \Delta_1 \circ \cdots \circ \Delta_n$ naturally reduces to a copy of $\Delta_1 \circ \cdots \circ \Delta_{n-1}$, and this process sends the edge e_2 to the edge e_1. Therefore, because e_1 is a cutting edge in Δ_1, we deduce that e_2 is a cutting edge in Δ_2.

By iterating this construction, we find a cutting edge e_i in Δ_i for every $1 \leq i \leq n$, where $e_i \neq e_j$ provided that $i \neq j$. On the other hand, the path $\mathrm{bot}(\Delta_1 \circ \cdots \circ \Delta_n)$ necessarily contains all these edges, hence $n \leq |w|$. Consequently, Δ_1 cannot contain a cutting edge if $n > |w|$. □

Proof of Theorem 5.28 Let γ be the combinatorial axis of g given by Lemma 5.29. According to Theorem 4.17, g is a contracting isometry if and only if $\gamma(+\infty)$ is isolated in $\partial^c X(\mathcal{P}, w)$. Thus, using the isomorphism given by Theorem 5.22, g is a contracting isometry if and only if g^∞ is isolated in $\partial(\mathcal{P}, w)$, i.e.

- every infinite prefix of g^∞ is commensurable to g^∞;
- if Δ is a reduced diagram with g^∞ as a prefix, then Δ is commensurable to g^∞.

Therefore, to conclude, it is sufficient to prove that a proper infinite prefix Δ of g^∞ cannot be commensurable to g^∞. Because Δ is a proper prefix, there exists a cell π_1 of g^∞ which does not belong to Δ. Now, by applying Lemma 5.31 to a given edge of the bottom path of π_1, we find a cell π_2 whose top path intersects the bottom path of π_1 along at least one edge; by applying Lemma 5.31 to a given edge of the bottom path of π_2, we find a cell π_3 whose top path intersects the bottom path of π_2 along at least one edge; and so on. Finally, we find an infinite sequence of cells $\pi_2, \pi_3, \ldots$ such that, for every $i \geq 2$, the top path of π_i has a common edge with the bottom path of π_{i-1}. For every $i \geq 1$, let Ξ_i denote the smallest (finite) prefix of g^∞ which contains the cell π_i. This is a minimal diagram, and by construction $\Xi_1, \Xi_2, \ldots \notin \mathcal{M}(\Delta)$, which proves that Δ and g^∞ are not commensurable. □

Example 5.32 Let $\mathcal{P} = \langle a, b, p \mid a = ap,\ b = pb \rangle$ and let $g \in D(\mathcal{P}, ab)$ be the spherical diagram illustrated by Figure 7. Then (see Figure 8) g^∞ clearly contains a proper infinite prefix. Therefore, g is not a contracting isometry of $X(\mathcal{P}, ab)$.

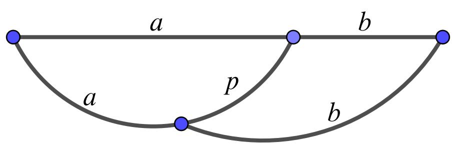

Figure 7: Spherical diagram from Example 5.32.

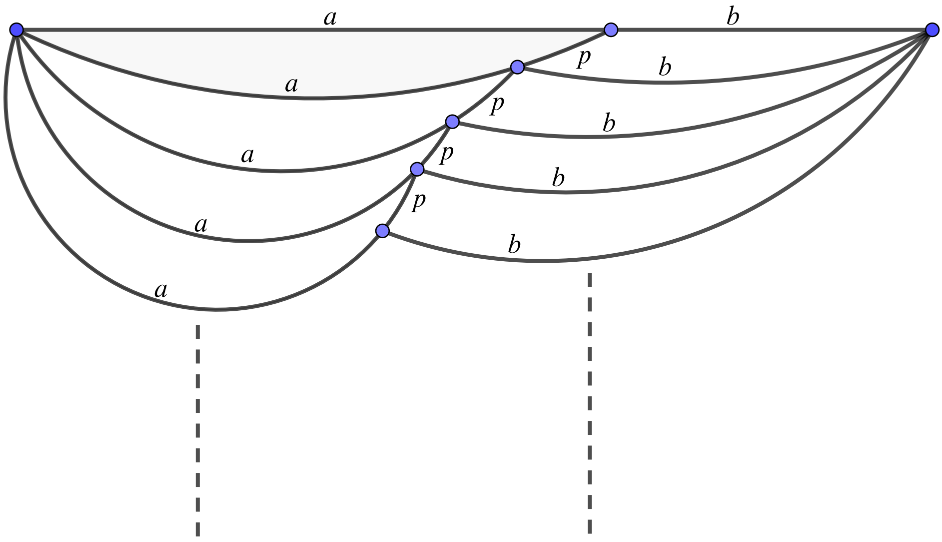

Figure 8: The infinite diagram g^∞ associated to the diagram g given by Figure 7.

Example 5.33 Let $\mathcal{P} = \langle a, b, c \mid a = b,\ b = c,\ c = a \rangle$ and let $g \in D(\mathcal{P}, a^2)$ be the spherical diagram illustrated by Figure 9. Then g^∞ is a prefix of the diagram Δ given by Figure 10, but g^∞ and Δ are not commensurable. Therefore, g is not a contracting isometry of $X(\mathcal{P}, a^2)$.

Example 5.34 Let $\mathcal{P} = \langle a, b, c, d \mid ab = ac,\ cd = bd \rangle$ and let $g \in D(\mathcal{P}, ab)$ be the spherical diagram illustrated by Figure 11. Then the infinite diagram g^∞ is given by Figure 11. We can notice that g^∞ does not contain a proper infinite prefix and that any diagram containing g^∞ as a prefix is necessarily equal to g^∞.

The criterion provided by Theorem 5.28 may be considered as unsatisfactory, because we cannot draw an infinite diagram in practice (although it seems to be sufficient for spherical diagrams with only few cells, as suggested by the above examples). We conclude this section by stating and proving conditions equivalent to the assertions (i)

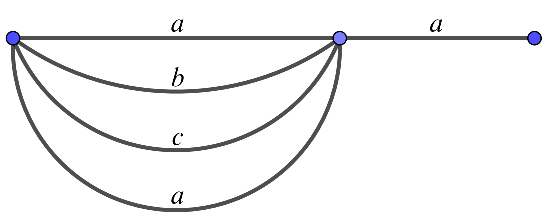

Figure 9: Spherical diagram from Example 5.33.

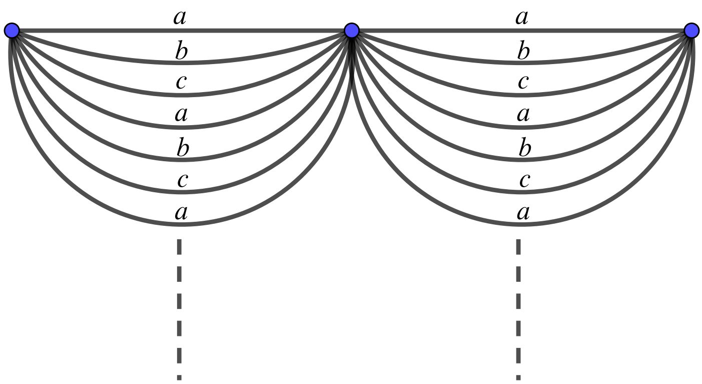

Figure 10: The infinite diagram g^∞ associated to the diagram g given by Example 5.33.

and (ii) of Theorem 5.28. For convenience, we will suppose our spherical diagram is absolutely reduced and *normal*, i.e. the factors of every decomposition as a sum are spherical; this assumption is not really restrictive since, according to [13, Lemma 15.13], every spherical diagram is conjugate to a normal absolutely reduced spherical diagram (possibly for a different baseword), and this new diagram can be found "effectively" [13, Lemma 15.14]. We begin with the condition (i). The following definition will be needed:

Definition 5.35 Given a diagram Δ, we introduce a partial relation $\prec$ on the set of its cells, defined by if π and π' are two cells of Δ, then $\pi \prec \pi'$ holds if the intersection between the bottom path of π and the top path of π' contains at least one edge. Let $\prec\!\prec$ denote the transitive closure of $\prec$.

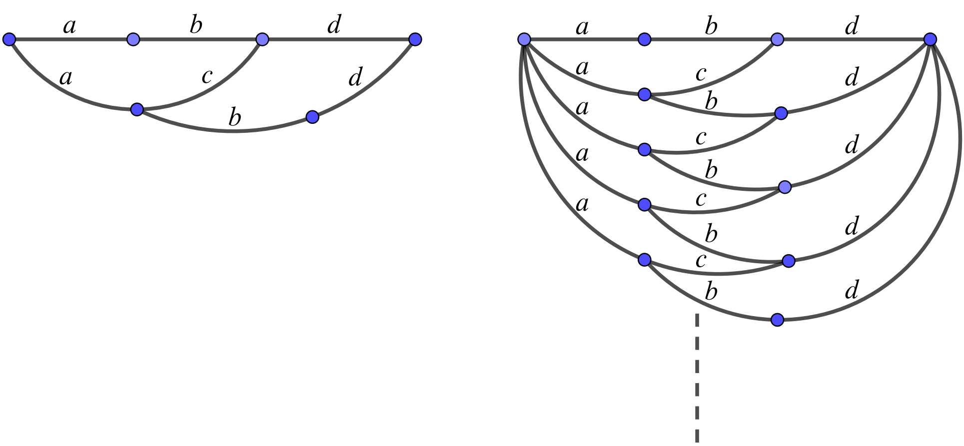

Figure 11: The diagram g and its associated infinite diagram g^∞ from Example 5.34.

Lemma 5.36 *Let $g \in D(\mathcal{P}, w) - \{\epsilon(w)\}$ be absolutely reduced. Suppose that we can write $w = xyz$, where $\operatorname{supp}(g) = \{y\}$. The following two conditions are equivalent:*

(i) *$\partial(\mathcal{P}, x)$ and $\partial(\mathcal{P}, z)$ are empty.*

(ii) *if Δ is a reduced diagram with g^∞ as a prefix, then Δ is commensurable to g^∞.*

Proof First, notice that $g = \epsilon(x) + g_0 + \epsilon(z)$ for some (y, y)–diagram g_0. This assertion clearly holds for some $(y, *)$–diagram g_0, but, because g is a (w, w)–diagram, the equality

$$xyz = w = \operatorname{top}(g) = \operatorname{bot}(g) = x \cdot \operatorname{bot}(g_0) \cdot z$$

holds in Σ^+. Therefore, $\operatorname{bot}(g_0) = y$, and we conclude that g_0 is indeed a (y, y)–diagram.

Now, we prove (i) $\Longrightarrow$ (ii). Let Δ be a reduced diagram containing g^∞ as a prefix. Suppose that π is a cell of Δ such that there exists a cell π_0 of g^∞ satisfying $\pi_0 \prec\!\!\prec \pi$. So there exists a chain of cells $\pi_0 \prec \pi_1 \prec \cdots \prec \pi_n = \pi$. By definition, π_1 share an edge with π_0. On the other hand, it follows from Lemma 5.31 that this edge belongs to the top path of a cell of g^∞, hence $\pi_1 \subset g^\infty$. By iterating this argument, we conclude that π belongs to g^∞. Therefore, if Ξ denotes the smallest prefix of Δ containing a given cell which does not belong to g^∞, then $\operatorname{supp}(\Xi)$ is included into x or z.

As a consequence, there exist a $(x, *)$–diagram Φ and $(z, *)$–diagram Ψ such that $\Delta = \Phi + g_0^\infty + \Psi$. On the other hand, because $\partial(\mathcal{P}, x)$ and $\partial(\mathcal{P}, y)$ are empty, necessarily $\#\Phi, \#\Psi < +\infty$. Therefore,

$$|\mathcal{M}(\Delta) \setminus \mathcal{M}(g^\infty)| = |\mathcal{M}(\Phi + \epsilon(yz)) \sqcup \mathcal{M}(\epsilon(xy) + \Psi)| = \#\Phi + \#\Psi < +\infty,$$

where we used Fact 5.10. Therefore, Δ is commensurable to g^∞.

Conversely, we prove (ii) $\Longrightarrow$ (i). Let Φ be a reduced $(x, *)$–diagram and Ψ a reduced $(z, *)$–diagram. Set $\Delta = \Phi + g_0^\infty + \Psi$. This is a reduced diagram containing g^∞ as a prefix. As above, we have

$$|\mathcal{M}(\Delta) \setminus \mathcal{M}(g^\infty)| = \#\Phi + \#\Psi.$$

Because Δ must be commensurable to g^∞, we deduce that Φ and Ψ are necessarily finite. We conclude that $\partial(\mathcal{P}, x)$ and $\partial(\mathcal{P}, z)$ are empty. □

Now, we focus on the condition (ii) of Theorem 5.28. We first introduce the following definition; explicit examples are given at the end of this section.

Definition 5.37 Let $g \in D(\mathcal{P}, w) - \{\epsilon(w)\}$ be absolutely reduced. A subpath u of w is *admissible* if there exists a prefix $\Delta \leq g$ such that $\mathrm{supp}(\Delta) = \{u\}$ and $\mathrm{bot}(\Delta) \cap \mathrm{bot}(g)$ is a nonempty connected subpath of $\mathrm{bot}(g)$; the subpath of w corresponding to $\mathrm{bot}(\Delta) \cap \mathrm{bot}(g)$ is a *final subpath* associated to u; it is *proper* if this is not the whole $\mathrm{bot}(g)$. Given a subpath u, we define its *admissible tree* $T_g(u)$ as the rooted tree constructed inductively as follows:

- the root of $T_g(u)$ is labeled by u;
- if v is a nonadmissible subpath of w labeling a vertex of $T_g(u)$, then v has only one descendant, labeled by the symbol $\varnothing$;
- if v is an admissible subpath of w labeling a vertex of $T_g(u)$, then the descendants of v correspond to the proper final subpaths associated to v.

If w has length n, we say that $T_g(u)$ is *deep* if it contains at least $\alpha(w)+1$ generations, where $\alpha(w) = \frac{1}{2}n(n+1)$ is chosen so that w contains at most $\alpha(w)$ subpaths.

Lemma 5.38 *Let $g \in D(\mathcal{P}, w) - \{\epsilon(w)\}$ be absolutely reduced. The following two conditions are equivalent*:

(i) *g^∞ contain an infinite proper prefix.*

(ii) *There exists a subpath u of w such that $T_g(u)$ is deep.*

Proof For convenience, let $g^\infty = g_1 \circ g_2 \circ \cdots$, where each g_i is a copy of g.

We first prove (i) $\Longrightarrow$ (ii). Suppose that g^∞ contains an infinite proper prefix Ξ. For every $i \geq 1$, Ξ induces a prefix Ξ_i of g_i. Up to taking an infinite prefix of Ξ, we can suppose that $\mathrm{supp}(\Xi_i)$ and $\mathrm{bot}(\Xi_i) \cap \mathrm{bot}(g_i)$ are connected subpaths of $\mathrm{top}(g_i)$ and $\mathrm{bot}(g_i)$, respectively. Thus, if ξ_i denotes the word labeling $\mathrm{supp}(\Xi_i)$ for every $i \geq 1$, then the sequence (ξ_i) defines an infinite descending ray in the rooted tree $T_g(\xi_1)$. A fortiori, $T_g(\xi_1)$ is deep.

Conversely, we prove (ii) $\Longrightarrow$ (i). Suppose that there exists a subword u of w such that $T_g(u)$ is deep. Because w contains at most $\alpha(w)$ subwords, necessarily $T_g(u)$ contains twice a proper subword ξ of $\mathrm{supp}(g)$ along a line from the root to a leaf. Let $\xi = \xi_0, \ldots, \xi_n = \xi$ be the geodesic between these two vertices in $T_g(u)$. For every $0 \leq i \leq n-1$, there exists a prefix Ξ_i of g such that $\mathrm{supp}(\Xi_i) = \{\xi_i\}$ and $\mathrm{bot}(X_i) \cap \mathrm{bot}(g)$ is a connected path labeled by ξ_{i+1}. In particular, $\Xi_i = \epsilon(x_i) + \widetilde{\Xi}_i + \epsilon(y_i)$ for some words $x_i, y_i \in \Sigma^+$ and some $(\xi_i, p_i\xi_{i+1}q_i)$–diagram $\widetilde{\Xi}_i$. Now, define

$$\Xi'_{i+1} = \epsilon(x_0 p_0 \cdots p_i) + \widetilde{\Xi}_{i+1} + \epsilon(q_i \cdots q_0 y_0) \quad \text{for } i \geq 0.$$

Note Ξ'_{i+1} is a $(x_0p_0\cdots p_i\xi_iq_i\cdots q_0y_0, x_0p_0\cdots p_{i+1}\xi_{i+1}q_{i+1}\cdots q_0y_0)$–diagram, so that the concatenation $\Xi = \Xi_0 \circ \Xi'_1 \circ \cdots \circ \Xi'_{n-1}$ is well defined.

We have proved g^n contains a prefix Ξ such that $\mathrm{supp}(\Xi)=\{\xi\}$ and $\mathrm{bot}(\Xi)\cap \mathrm{bot}(g^n)$ is the same subpath ξ.

Say that Ξ is a $(x\xi y, xp\xi qy)$–diagram for some words $x,y,p,q\in\Sigma^+$. In particular, Ξ decomposes as a sum $\epsilon(x)+\widetilde{\Xi}+\epsilon(y)$, where $\widetilde{\Xi}$ is a $(\xi, p\xi q)$–diagram. For every $i\geq 0$, set

$$\Delta_i = \epsilon(xp^i)+\widetilde{\Xi}+\epsilon(q^iy).$$

Finally, the concatenation $\Delta=\Delta_0\circ\Delta_1\circ\cdots$ defines an infinite prefix of g^∞. Moreover,

$$\mathrm{supp}(\Delta)=\mathrm{supp}(\Xi)=\{\xi\}\subsetneq \mathrm{supp}(g)\subset \mathrm{supp}(g^\infty),$$

so Δ is a proper prefix of g^∞. □

By combining Lemmas 5.36 and 5.38 with Theorem 5.28, we obtain the following new criterion:

Proposition 5.39 *Let $g\in D(\mathcal{P},w)-\{\epsilon(w)\}$ be normal and absolutely reduced. Then g is a contracting isometry of $X(\mathcal{P},w)$ if and only if the following two conditions are satisfied*:

- *The support of g is reduced to $\{w_1\}$ for some subword w_1 of w, and, if we write $w=xw_1y$, then $\partial(\mathcal{P},x)$ and $\partial(\mathcal{P},y)$ are empty.*
- *For every subpath u of w, $T_g(u)$ is not deep.*

Proof According to Theorem 5.28, g is contracting if and only if

(i) g^∞ does not contain any infinite proper prefix;

(ii) if Δ is a reduced diagram with g^∞ as a prefix, then Δ is commensurable to g^∞.

Now, Lemma 5.38 states that the condition (i) is equivalent to

(iii) for every subpath u of w, $T_g(u)$ is not deep;

and Lemma 5.36 implies that the condition (ii) is a consequence of

(iv) the support of g is reduced to $\{w_1\}$ for some subword w_1 of w, and, if we write $w=xw_1y$, then $\partial(\mathcal{P},x)$ and $\partial(\mathcal{P},y)$ are empty.

To conclude the proof, it is sufficient to show that (iv) is a consequence of (i) and (ii). So suppose that (i) and (ii) hold. If $\mathrm{supp}(g)$ is not reduced to a single subpath, then g

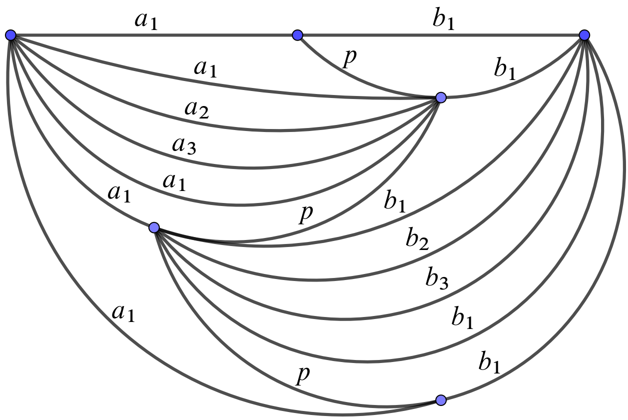

Figure 12: Spherical diagram from Example 5.41.

decomposes as a sum of nontrivial diagrams, and because g is normal, it decomposes as a sum of nontrivial spherical diagrams, say $g = g_1 + g_2$. Now, since g_1 and g_2 are spherical, $g^\infty = g_1^\infty + g_2^\infty$, so g^∞ clearly contains an infinite proper prefix, contradicting (i). Therefore, $\mathrm{supp}(g)$ reduces to a single subpath, and Lemma 5.36 implies that (iv) holds. $\square$

Remark 5.40 The first bullet of Proposition 5.39 is clearly satisfied if $\mathrm{supp}(g) = \{w\}$. Because the case often happens, we will say that spherical diagrams satisfying this property are *full*.

Example 5.41 The group $\mathbb{Z} \bullet \mathbb{Z} = \langle a, b, t \mid [a, b^{t^n}] = 1,\ n \geq 0 \rangle$, introduced in [13, Section 8], is isomorphic to the diagram group $D(\mathcal{P}, ab)$, where

$$\mathcal{P} = \big\langle a_1, a_2, a_3, p, b_1, b_2, b_3 \mid a_1 = a_1 p,\ a_1 = a_2,\ a_2 = a_3,\ a_3 = a_1, \\ b_1 = pb_1,\ b_1 = b_2,\ b_2 = b_3,\ b_3 = b_1 \big\rangle.$$

Let $g \in D(\mathcal{P}, ab)$ be the spherical diagram illustrated by Figure 12. This is a full, normal and absolutely reduced diagram. Moreover, $a_1 b_1$ is the only admissible subpath, and $T_g(a_1 b_1)$ is

$$\begin{array}{c} a_1 b_1 \\ | \\ b_1 \end{array}$$

Therefore, g is a contracting isometry of $X(\mathcal{P}, ab)$. In particular, $\mathbb{Z} \bullet \mathbb{Z}$ is acylindrically hyperbolic.

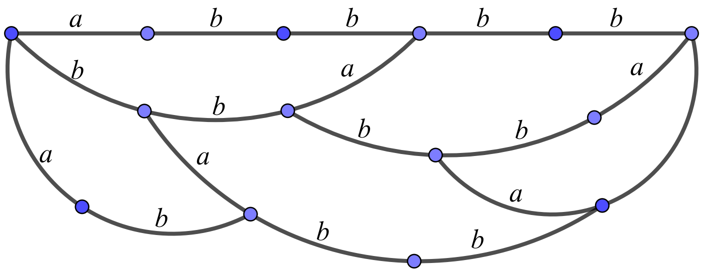

Figure 13: Spherical diagram from Example 5.42.

Example 5.42 Let $\mathcal{P} = \langle a, b \mid ab = ba,\ ab^2 = b^2a \rangle$ and let $g \in D(\mathcal{P}, ab^4)$ be the spherical diagram illustrated by Figure 13. This is a full, normal and absolutely reduced diagram. Moreover, ab^4 is the only admissible subpath, and $T_g(ab^4)$ is

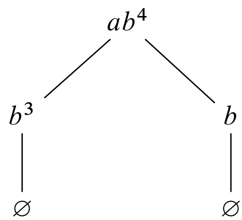

We conclude that g is a contracting isometry of $X(\mathcal{P}, ab^4)$.

5.4 Some examples

In this section, we exhibit some interesting classes of acylindrically hyperbolic diagram groups by applying the results proved in the previous section. In Example 5.43, we consider a family of groups $U(m_1, \dots, m_n)$ already studied in [13] and prove that they are acylindrically hyperbolic. Moreover, we notice that each $L_n = U(1, \dots, 1)$ turns out to be naturally a subgroup of a right-angled Coxeter group. In Example 5.44, we show that the $\bullet$–product, as defined in [13], of two nontrivial diagram groups is acylindrically hyperbolic but not relatively hyperbolic. Finally, in Example 5.46, we prove that some diagram products, as defined in [14], of diagram groups turn out to be acylindrically hyperbolic. As a byproduct, we also exhibit a cocompact diagram group which is not isomorphic to a right-angled Artin group.

Example 5.43 Let $\mathcal{P}_n = \langle x_1, \dots, x_n \mid x_i x_j = x_j x_i \text{ for } 1 \leq i < j \leq n \rangle$ and let $U(m_1, \dots, m_n)$ denote the diagram group $D(\mathcal{P}, x_1^{m_1} \cdots x_n^{m_n})$, where $m_i \geq 1$ for every $1 \leq i \leq n$. Notice that, for every permutation $\sigma \in S_n$, $U(m_{\sigma(1)}, \dots, m_{\sigma(n)})$ is isomorphic to $U(m_1,\dots,m_n)$; therefore, we will always suppose that $m_1 \leq \cdots \leq m_n$. For example, the groups $U(p,q,r)$ were considered in [13, Example 10.2] (see also [10, Example 2; 11, Example 5.11]). In particular, we know that

- $U(m_1,\dots,m_n)$ is trivial if and only if $n \leq 2$;
- $U(m_1,\dots,m_n)$ is infinite cyclic if and only if $n=3$ and $m_1=m_2=m_3=1$;
- $U(m_1,\dots,m_n)$ is a nonabelian free group if and only if $n=3$ and $m_1=1$, or $n=4$ and $m_1=m_2=m_3=1$, or $n=5$ and $m_1=m_2=m_3=m_4=m_5=1$.

We claim that, whenever it is not cyclic, the group $U(m_1,\dots,m_n)$ is acylindrically hyperbolic. For instance, if $g \in U(1,1,2)$ denotes the spherical diagram associated to the derivation

$$x_1x_2x_3x_3 \to x_2x_1x_3x_3 \to x_2x_3x_1x_3 \to x_3x_2x_1x_3 \to x_3x_1x_2x_3 \to x_1x_3x_2x_3 \to x_1x_2x_3x_3,$$

then g is a contracting isometry of the associated Farley cube complex. (The example generalizes easily to other cases although writing an argument in full generality is laborious.) In particular, we deduce that $U(m_1,\dots,m_n)$ does not split as a nontrivial direct product.

An interesting particular case is the group L_n, corresponding to $U(1,\dots,1)$ with n ones. A first observation is that L_n is naturally a subgroup of $U(m_1,\dots,m_n)$, and, conversely, $U(m_1,\dots,m_n)$ is naturally a subgroup of $L_{m_1+\cdots+m_n}$. Secondly, L_n can be described as a group of *pseudobraids*; this description is made explicit by the example of a diagram with its associated pseudobraid illustrated by Figure 14. Of course, we look at our pseudobraids up to the move illustrated by Figure 15, corresponding to the reduction of a dipole in the associated diagram. Finally, if C_n denotes the group of all pictures of pseudobraids, endowed with the obvious concatenation, then L_n corresponds to a "pure subgroup" of C_n. It is clear that, if σ_i corresponds to the element of C_n switching the i^{th} and $(i+1)^{\text{st}}$ braids, then a presentation of C_n is

$$\langle \sigma_1,\dots,\sigma_{n-1} \mid \sigma_i^2=1,\ [\sigma_i,\sigma_j]=1 \text{ if } |i-j|\geq 2\rangle.$$

Alternatively, C_n is isomorphic to the right-angled Coxeter group $C(\overline{P}_{n-1})$, where $\overline{P}_{n-1}$ is the complement of the graph P_{n-1} which is a segment with $n-1$ vertices, i.e. $\overline{P}_{n-1}$ is the graph with the same vertices as P_{n-1} but whose edges link two vertices precisely when they are not adjacent in P_{n-1}. In particular, L_n is naturally a finite-index subgroup of $C(\overline{P}_{n-1})$.

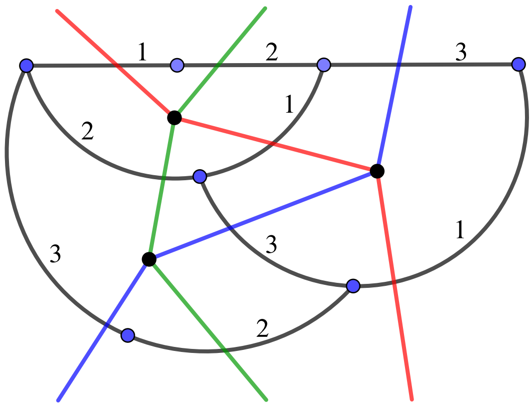

Figure 14: A semigroup diagram and its associated pseudobraid.

With this interpretation, the group $U(m_1, \dots, m_n)$ corresponds to the subgroup of L_m, where $m = m_1 + \dots + m_n$, defined as follows: Color the m braids with n colors such that, for every $1 \leq i \leq n$, there are m_i braids with the color i. Now, $U(m_1, \dots, m_n)$ is the subgroup of the pure pseudobraids where two braids of the same color are never switched.

Therefore, the groups $U(m_1, \dots, m_n)$ turn out to be natural subgroups of right-angled Coxeter groups.

Example 5.44 Let G and H be two groups. We define the product $G \bullet H$ by the relative presentation

$$\langle G, H, t \mid [g, t^n h t^{-n}] = 1 \text{ for } g \in G,\ h \in H,\ n \geq 0 \rangle.$$

As proved in [13, Theorem 8.6], the class of diagram groups is stable under the $\bullet$–product. More precisely, let $\mathcal{P}_1 = \langle \Sigma_1 \mid \mathcal{R}_1 \rangle$ and $\mathcal{P}_2 = \langle \Sigma_2 \mid \mathcal{R}_2 \rangle$ be two semigroup presentations and w_1 and w_2 two basewords. For $i = 1, 2$, suppose that there do not

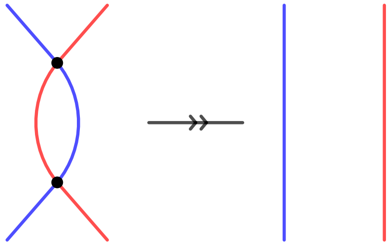

Figure 15: The reduction for pseudobraids associated to the reduction for diagrams.

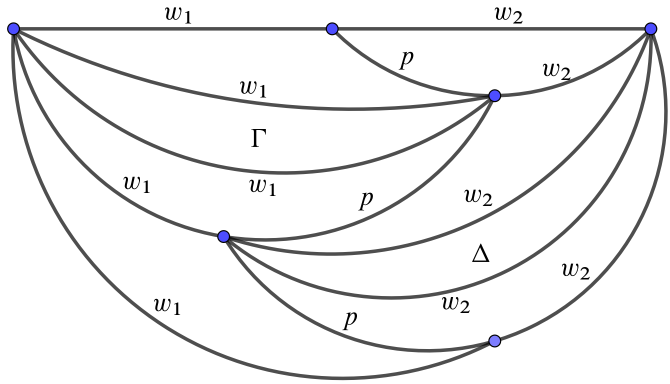

Figure 16: A contracting isometry.

exist two words x and y such that xy is nonempty and $w_i = xw_iy$; as explained in the proof of [13, Theorem 8.6], we may suppose that this condition holds up to a small modification of $\mathcal{P}_i$ which does not modify the diagram group. Then, if

$$\mathcal{P} = \langle \Sigma_1 \sqcup \Sigma_2 \sqcup \{p\} \mid \mathcal{R}_1 \sqcup \mathcal{R}_2 \sqcup \{w_1 = w_1 p,\ w_2 = pw_2\}\rangle,$$

then $D(\mathcal{P}, w_1w_2)$ is isomorphic to $D(\mathcal{P}_1, w_1) \bullet D(\mathcal{P}_2, w_2)$ [13, Lemma 8.5]. Now, we claim that, whenever $D(\mathcal{P}_1, w_1)$ and $D(\mathcal{P}_2, w_2)$ are nontrivial, $D(\mathcal{P}_1, w_1) \bullet D(\mathcal{P}_2, w_2)$ is acylindrically hyperbolic. Indeed, if $\Gamma \in D(\mathcal{P}_1, w_1)$ and $\Delta \in D(\mathcal{P}_2, w_2)$ are nontrivial, then the spherical diagram of $D(\mathcal{P}, w_1w_2)$ illustrated by Figure 16 is a contracting isometry. On the other hand, $D(\mathcal{P}_1, w_1) \bullet D(\mathcal{P}_2, w_2)$ contains a copy of $\mathbb{Z} \bullet \mathbb{Z}$, so it cannot be cyclic. We conclude that $D(\mathcal{P}_1, w_1) \bullet D(\mathcal{P}_2, w_2)$ is indeed acylindrically hyperbolic.

In fact, the product $G \bullet H$ can be described as an HNN extension: if H^∞ denotes the free product of infinitely many copies of H, then $G \bullet H$ is isomorphic to the HNN extension of $G \times H^\infty$ over H^∞ with respect to two monomorphisms $H^\infty \hookrightarrow G \times H^\infty$ associated to $H_i \mapsto H_i$ and $H_i \mapsto H_{i+1}$. Now, if t denotes the stable letter of the HNN extension and if $h \in H^\infty$ is a nontrivial element of the first copy of H in H^∞, then it follows from Britton's lemma that $H^\infty \cap (H^\infty)^{t^{-1}ht} = \{1\}$. Therefore, H^∞ is a weakly malnormal subgroup, and it follows from [24, Corollary 2.3] that, if G and H are both nontrivial, the product $G \bullet H$ is acylindrically hyperbolic. It is interesting that:

Fact 5.45 *If G and H are torsion-free, then any proper malnormal subgroup of $G \bullet H$ is a free group.*

Indeed, let K be a malnormal subgroup of $G \bullet H$. Suppose that K contains a nontrivial element of $G \times H^\infty$, say gh. Then $h \in K$ and $g \in K$ since

$$\langle gh \rangle = \langle gh \rangle \cap \langle gh \rangle^h \subset K \cap K^h \quad \text{and} \quad \langle gh \rangle = \langle gh \rangle \cap \langle gh \rangle^g \subset K \cap K^g.$$

If h is not trivial, then $G \subset K$ since

$$\langle h \rangle = \langle h \rangle \cap \langle h \rangle^G \subset K \cap K^G;$$

then we deduce that $H^\infty \subset K$ since

$$G = G \cap G^{H^\infty} \subset K \cap K^{H^\infty}.$$

Therefore, $G \times H^\infty \subset K$. Similarly, if h is trivial then g is nontrivial, and we deduce that $H^\infty \subset K$ from

$$\langle g \rangle = \langle g \rangle \cap \langle g \rangle^{H^\infty} \subset K \cap K^{H^\infty}.$$

We also notice that $G \subset K$ since

$$H^\infty = H^\infty \cap (H^\infty)^G \subset K \cap K^G.$$

Consequently, $G \times H^\infty \subset K$. Thus, we have proved that a malnormal subgroup of $G \bullet H$ either intersects trivially $G \times H^\infty$ or contains $G \times H^\infty$. If K intersects trivially $G \times H^\infty$, the action of K on the Bass–Serre tree of $G \bullet H$, associated to its decomposition as an HNN extension we mentioned above, is free, so K is necessarily free. Otherwise, K contains $G \times H^\infty$. If t denotes the stable letter of our HNN extension, then

$$H_2 * H_3 * \cdots = H^\infty \cap (H^\infty)^t \subset K \cap K^t,$$

where we used the notation $H^\infty = H_1 * H_2 * H_3 * \cdots$ where each H_i is a copy of H. Therefore, $t \in K$, and we conclude that K is not a proper subgroup. The fact is proved.

As a consequence, $\bullet$–products of two nontrivial diagram groups are good examples of acylindrically hyperbolic groups in the sense that they are not relatively hyperbolic. Indeed, if a group is hyperbolic relative to a collection of subgroups, then this collection must be malnormal (see for instance [25, Theorem 1.4]), so that a relatively hyperbolic $\bullet$–product of two torsion-free groups must be hyperbolic relative to a collection of free groups according to the previous fact. On the other hand, such a group must be hyperbolic (see for instance [25, Corollary 2.41]), which is impossible since we know that a $\bullet$–product of two nontrivial torsion-free groups contains a subgroup isomorphic to $\mathbb{Z}^2$.

Example 5.46 Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. Fixing a family of groups $\mathcal{G}_\Sigma = \{G_s \mid s \in \Sigma\}$, we define the *diagram product* $\mathcal{D}(\mathcal{G}_\Sigma, \mathcal{P}, w)$ as the fundamental group of the following 2–complex of groups:

- the underlying 2–complex is the 2–skeleton of the Squier complex $S(\mathcal{P}, w)$;
- to any vertex $u = s_1 \cdots s_r \in \Sigma^+$ is associated the group $G_u = G_{s_1} \times \cdots \times G_{s_r}$;
- to any edge $e = (a, u \to v, b)$ is associated the group $G_e = G_a \times G_b$;
- to any square is associated the trivial group;
- for every edge $e = (a, u \to v, b)$, the monomorphisms $G_e \to G_{aub}$ and $G_e \to G_{avb}$ are the canonical maps $G_a \times G_b \to G_a \times G_u \times G_b$ and $G_a \times G_b \to G_a \times G_v \times G_b$.

Guba and Sapir proved that the class of diagram groups is stable under diagram product [14, Theorem 4]. Explicitly:

Theorem 5.47 *If $G_s = D(\mathcal{P}_s, w_s)$ where $\mathcal{P}_s = \langle \Sigma_s \mid \mathcal{R}_s \rangle$ is a semigroup presentation and $w_s \in \Sigma_s^+$ a baseword, then the diagram product $\mathcal{D}(\mathcal{G}_\Sigma, \mathcal{P}, w)$ is the diagram group $D(\overline{\mathcal{P}}, w)$, where $\overline{\mathcal{P}}$ is the semigroup presentation*

$$\langle \Sigma \sqcup \overline{\Sigma} \sqcup \Sigma_0 \mid \mathcal{R} \sqcup \overline{\mathcal{R}} \sqcup \mathcal{S} \rangle,$$

where $\Sigma_0 = \{a_s \mid s \in \Sigma\}$ is a copy of Σ, $\overline{\Sigma} = \bigsqcup_{s \in \Sigma} \Sigma_s$, $\overline{\mathcal{R}} = \bigsqcup_{s \in \Sigma} \mathcal{R}_s$ and finally $\mathcal{S} = \bigsqcup_{s \in \Sigma} \{s = a_s w_s a s\}$.

Of course, any semigroup diagram with respect to $\mathcal{P}$ is a semigroup diagram with respect to $\overline{\mathcal{P}}$. Thus, if $D(\mathcal{P}, w)$ contains a normal, absolutely reduced, spherical diagram which does not contains a proper infinite prefix and whose support is $\{w\}$, then any diagram product over $(\mathcal{P}, w)$ will be acylindrically hyperbolic or cyclic. (Notice that such a diagram product contains $D(\mathcal{P}, w)$ as a subgroup, so it cannot be cyclic if $D(\mathcal{P}, w)$ is not cyclic itself.)

For instance, let $\mathcal{P} = \langle x, y, z, a, b \mid yz = az,\ xa = xb,\ b = y \rangle$. Then $D(\mathcal{P}, xyz)$ contains the spherical diagram illustrated by Figure 17, which is a contracting isometry of $X(\mathcal{P}, xyz)$. By the above observation, if G denotes the diagram product over $(\mathcal{P}, apb)$ with $G_x = G_z = \mathbb{Z}$ and $G_y = G_a = G_b = \{1\}$, then G is acylindrically hyperbolic. In this case, the Squier complex $S(\mathcal{P}, xyz)$ is just a cycle of length three. By computing a presentation of the associated graph of groups, we find

$$G = \langle x, y, t \mid [x, y] = [x, y^t] = 1 \rangle.$$

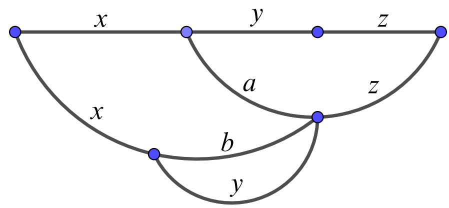

Figure 17: A diagram which implies the acylindrical hyperbolicity of a diagram product.

This example is interesting because this is a cocompact diagram group which is not a right-angled Artin group. Let us prove this assertion. From the presentation of G, it is clear that its abelianization is $\mathbb{Z}^3$ and we deduce from [4, Exercice II.5.5] that the rank of $H_2(G)$ is at most two (the number of relations of our presentation). On the other hand, if $A(\Gamma)$ denotes the right-angled Artin group associated to a given graph Γ, then $H_1(A(\Gamma)) = \mathbb{Z}^{\# V(\Gamma)}$ and $H_2(A(\Gamma)) = \mathbb{Z}^{\# E(\Gamma)}$, where $V(\Gamma)$ and $E(\Gamma)$ denote the set of vertices and edges of Γ, respectively. Therefore, the only right-angled Artin groups which might be isomorphic to G must correspond to a graph with three vertices and at most two edges. We distinguish two cases: either this graph is not connected or it is a segment of length two. In the former case, the right-angled Artin group decomposes as a free product, and in the latter, it is isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$. Thus, it is sufficient to show that G is freely irreducible and is not isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$ in order to deduce that G cannot be isomorphic to a right-angled Artin group. First, notice that, because G is acylindrically hyperbolic, its center must be finite (see [26, Corollary 7.3]), so G cannot be isomorphic to $\mathbb{F}_2 \times \mathbb{Z}$. Next, suppose that G decomposes as a free product. Because the centralizer of x is not virtually cyclic, we deduce that x belongs to a conjugate A of some free factor; in fact, A must contain the whole centralizer of x, hence $y, y^t \in A$. As a consequence, $y^t \in A \cap A^t$, so that $A \cap A^t$ must be infinite. Since a free factor is a malnormal subgroup, we deduce that $t \in A$. Finally, $x, y, t \in A$ hence $A = G$. Therefore, G does not contain any proper free factor, i.e. G is freely irreducible. This concludes the proof of our assertion.

To conclude, it is worth noticing that the whole diagram product may contain contracting isometries even if the underlying diagram group (i.e. the diagram group obtained from the semigroup presentation and the baseword, forgetting the factor groups) does not contain such isometries on its own Farley cube complex. For instance, if $\mathcal{P} = \langle a, b, p \mid a = ap,\ b = pb \rangle$, $w = ab$, $G_p = \{1\}$, and $G_a = D(\mathcal{P}_1, w_1)$ and $G_b = D(\mathcal{P}_2, w_2)$ are nontrivial, then our diagram product is isomorphic to $G_a \bullet G_b$ (see [14, Example 7]), and the description of $G_a \bullet G_b$ as a diagram group is precisely the one we mentioned in the previous example, where we saw that the action on the associated Farley cube complex admits contracting isometries. On the other hand, $X(\mathcal{P},ab)$ is a "diagonal halfplane", so that its combinatorial boundary does not contain any isolated point: a fortiori, $X(\mathcal{P},ab)$ cannot admit a contracting isometry.

5.5 Cocompact case

In this section, we focus on cocompact diagram groups, i.e. we will consider a semigroup presentation $\mathcal{P}=\langle \Sigma \mid \mathcal{R} \rangle$ and a baseword $w \in \Sigma^+$ whose class modulo $\mathcal{P}$ is finite. Roughly speaking, we want to prove that $D(\mathcal{P},w)$ contains a contracting isometry if and only if it does not split as a direct product in a natural way, which we describe now.

If there exists a $(w, u_1v_1\cdots u_nv_nu_{n+1})$–diagram Γ for some words $u_i, v_j, \in \Sigma^+$, then we have the natural map

$$\begin{aligned} &D(\mathcal{P},v_1)\times\cdots\times D(\mathcal{P},v_n)\to D(\mathcal{P},w),\\ &\qquad (V_1,\dots,V_n)\mapsto \Gamma\cdot(\epsilon(u_1)+V_1+\cdots+\epsilon(u_n)+V_n+\epsilon(u_{n+1}))\cdot\Gamma^{-1}. \end{aligned}$$

This is a monomorphism, and we denote by $\Gamma\cdot D(\mathcal{P},v_1)\times\cdots\times D(\mathcal{P},v_n)\cdot\Gamma^{-1}$ its image. A subgroup of this form will be referred to as a *canonical subgroup*; if furthermore at least two factors are nontrivial, then it will be a *large canonical subgroup*.

The main result of this section is the following decomposition theorem (see Theorem 1.9 in the introduction):

Theorem 5.48 *Let $\mathcal{P}=\langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword whose class modulo $\mathcal{P}$ is finite. Then there exist some words $u_0,u_1,\dots,u_m \in \Sigma^+$ and a $(w,u_0u_1\cdots u_m)$–diagram Γ such that*

$$D(\mathcal{P},w)=\Gamma\cdot(D(\mathcal{P},u_0)\times D(\mathcal{P},u_1)\times\cdots\times D(\mathcal{P},u_m))\cdot\Gamma^{-1},$$

where $D(\mathcal{P},u_i)$ is either infinite cyclic or acylindrically hyperbolic for every $1\le i\le m$, and $m\le \dim_{\mathrm{alg}} D(\mathcal{P},w)\le \dim X(\mathcal{P},w)$.

The acylindrical hyperbolicity will follow from the existence of a contracting isometry in the corresponding Farley cube complex. In order to understand what happens when such an isometry does not exist, we will use the following dichotomy, which is a slight modification, in the cocompact case, of the rank rigidity theorem proved in [5]:

Theorem 5.49 *Let G be a group acting geometrically on an unbounded* CAT(0) *cube complex X. Either G contains a contracting isometry of X or there exists a G–invariant convex subcomplex $Y \subset X$ which splits as the Cartesian product of two unbounded subcomplexes.*

Proof Recall that a hyperplane of X is *essential* if no G–orbit lies in a neighborhood of some halfspace delimited by this hyperplane; we will denote by $\mathcal{E}(X)$ the set of the essential hyperplanes of X. Let $Y \subset X$ denote the *essential core* associated to the action $G \curvearrowright X$, which is defined as the *restriction quotient* of X with respect to $\mathcal{E}(X)$. See [5, Section 3.3] for more details. According to [5, Proposition 3.5], Y is a G–invariant convex subcomplex of X.

It is worth noticing that, because the action $G \curvearrowright X$ is cocompact, a hyperplane J of X does not belong to $\mathcal{E}(X)$ if and only if one of the halfspaces it delimits lies in the $R(J)$–neighborhood of $N(J)$ for some $R(J) \geq 1$. Set $R = \max_{J \in \mathcal{H}(X)} R(J)$. Because there exist only finitely many orbits of hyperplanes, R is finite.

We claim that two hyperplanes of Y are well-separated in Y if and only if they are well-separated in X.

Of course, two hyperplanes of Y which are well-separated in X are well-separated in Y. Conversely, let $J_1, J_2 \in \mathcal{H}(Y)$ be two hyperplanes and fix some finite collection $\mathcal{H} \subset \mathcal{H}(X)$ of hyperplanes intersecting both J_1 and J_2 which does not contain any facing triple. We write $\mathcal{H} = \{H_0, \dots, H_n\}$ and we fix a halfspace H_i^+ delimited by H_i for every $0 \leq i \leq n$, so that $H_0^+ \supset H_1^+ \supset \cdots \supset H_n^+$. If $n \leq 2R$, there is nothing to prove. Otherwise, by our choice of R, it is clear that the hyperplanes $H_R, \dots, H_{n-R}$ belong to $\mathcal{E}(X)$, and a fortiori to $\mathcal{H}(Y)$. Consequently, if J_1 and J_2 are L–well-separated in Y, then they are $(L+2R)$–well-separated in X.

Now we are ready to prove our theorem. Because the action $G \curvearrowright Y$ is geometric and *essential* (i.e. every hyperplane of Y is essential), it follows from [5, Theorem 6.3] that either Y splits nontrivially as a Cartesian product of two subcomplexes or G contains a contracting isometry of Y. In the former case, notice that the factors cannot be bounded because the action $G \curvearrowright Y$ is essential; in the latter case, we deduce that the contracting isometry of Y induces a contracting isometry of X by combining our previous claim with Theorem 3.13. This concludes the proof. □

Therefore, the problem reduces to understand the subproducts of $X(\mathcal{P}, w)$. This purpose is achieved by the next proposition.

Proposition 5.50 *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. Let $Y \subset X(\mathcal{P}, w)$ be a convex subcomplex which splits as the Cartesian product of two subcomplexes $A, B \subset Y$. Then there exist a word $w' \in \Sigma^+$, a (w, w')–diagram Γ and words $u_1, v_1, \dots, u_k, v_k \in \Sigma^+$ for some $k \geq 1$ such that $w' = u_1 v_1 \cdots u_k v_k$ in Σ^+ and*

$$A \subset \Gamma \cdot X(\mathcal{P}, u_1) \times \cdots \times X(\mathcal{P}, u_k) \cdot \Gamma^{-1} \quad \text{and} \quad B \subset \Gamma \cdot X(\mathcal{P}, v_1) \times \cdots \times X(\mathcal{P}, v_k) \cdot \Gamma^{-1}.$$

Furthermore,

$$Y \subset \Gamma \cdot X(\mathcal{P}, u_1) \times X(\mathcal{P}, v_1) \times \cdots \times X(\mathcal{P}, u_k) \times X(\mathcal{P}, v_k) \cdot \Gamma^{-1}.$$

In the previous statement, Y is a Cartesian product of the subcomplexes A and B in the sense that there exists an isomorphism $\varphi\colon A \times B \to Y$ which commutes with the inclusions $A, B, Y \hookrightarrow X$. More explicitly, we have the commutative diagram

$$\begin{tikzcd}
 & A \arrow[dl, hook'] \arrow[dr, hook] & \\
A \times B \arrow[r, "\varphi"] & Y \arrow[r, hook] & X \\
 & B \arrow[ul, hook] \arrow[ur, hook'] &
\end{tikzcd}$$

Proof of Proposition 5.50 Up to conjugating by a diagram Γ of $A \cap B$, we may assume without loss of generality that $\epsilon(w) \in A \cap B$. First, we claim that, for every $\Delta_1 \in A$ and $\Delta_2 \in B$, $\operatorname{supp}(\Delta_1) \cap \operatorname{supp}(\Delta_2) = \varnothing$.

Suppose by contradiction that there exist some $\Delta_1 \in A$ and $\Delta_2 \in B$ satisfying $\operatorname{supp}(\Delta_1) \cap \operatorname{supp}(\Delta_2) \neq \varnothing$. Without loss of generality, we may suppose that we chose a counterexample minimizing $\#\Delta_1 + \#\Delta_2$. Because $\operatorname{supp}(\Delta_1) \cap \operatorname{supp}(\Delta_2) \neq \varnothing$, there exist two cells $\pi_1 \subset \Delta_1$ and $\pi_2 \subset \Delta_2$ and an edge $e \subset w$, where w is thought of as a segment representing both $\operatorname{top}(\Delta_1)$ and $\operatorname{top}(\Delta_2)$, such that $e \subset \operatorname{top}(\pi_1) \cap \operatorname{top}(\pi_2)$. Notice that π_i is the only cell of the maximal thin suffix of Δ_i for $i = 1, 2$. Indeed, if this was not the case, by taking the minimal prefixes P_1 and P_2 of Δ_1 and Δ_2 containing π_1 and π_2, respectively, we would produce two diagrams satisfying $\operatorname{supp}(P_1) \cap \operatorname{supp}(P_2) \neq \varnothing$ and $\#P_1 + \#P_2 < \#\Delta_1 + \#\Delta_2$. Moreover, because Y is convex, necessarily the factors A and B have to be convex as well, so that A and B, as sets of diagrams, have to be closed by taking prefixes since $\epsilon(w) \in A \cap B$ and according to the description of the combinatorial geodesics given by Lemma 5.4; this implies that $P_1 \in A$ and $P_2 \in B$, contradicting the choice of our counterexample Δ_1 and Δ_2. For $i=1,2$, let $\overline{\Delta}_i$ denote the diagram obtained from Δ_i by removing the cell π_i. We claim that

$$\operatorname{supp}(\overline{\Delta}_1) \cap \operatorname{suppp}(\Delta_2) = \operatorname{supp}(\Delta_1) \cap \operatorname{suppp}(\overline{\Delta}_2) = \varnothing.$$

Indeed, since we know that A and B are stable by taking prefixes, $\overline{\Delta}_1$ belongs to A and $\overline{\Delta}_2$ belongs to B. Therefore, because

$$\#\overline{\Delta}_1 + \#\Delta_2 = \#\Delta_1 + \#\overline{\Delta}_2 < \#\Delta_1 + \#\Delta_2,$$

having $\operatorname{supp}(\overline{\Delta}_1) \cap \operatorname{suppp}(\Delta_2) \neq \varnothing$ or $\operatorname{supp}(\Delta_1) \cap \operatorname{suppp}(\overline{\Delta}_2) \neq \varnothing$ would contradict our choice of Δ_1 and Δ_2.

Because $\operatorname{supp}(\overline{\Delta}_1) \cap \operatorname{suppp}(\overline{\Delta}_2) = \varnothing$, it is possible to write $w = a_1 b_1 \cdots a_p b_p$ in Σ^+ so that

$$\overline{\Delta}_1 = A_1 + \epsilon(b_1) + \cdots + A_p + \epsilon(b_p), \quad \overline{\Delta}_2 = \epsilon(a_1) + B_1 + \cdots + \epsilon(a_p) + B_p$$

for some $(a_i,*)$–diagrams A_i and $(b_i,*)$–diagrams B_i. Set

$$\Delta_0 = A_1 + B_1 + \cdots + A_p + B_p.$$

Notice that $\overline{\Delta}_1$ and $\overline{\Delta}_2$ are prefixes of Δ_0.

For $i=1,2$, let E_i denote the atomic diagram such that the concatenation $\Delta_0 \circ E_i$ corresponds to gluing π_i on $\overline{\Delta}_i$ as a prefix of Δ_0. Notice that the diagrams E_1 and E_2 exist precisely because

$$\operatorname{supp}(\overline{\Delta}_1) \cap \operatorname{suppp}(\Delta_2) = \operatorname{supp}(\Delta_1) \cap \operatorname{suppp}(\overline{\Delta}_2) = \varnothing.$$

As $\operatorname{top}(\pi_1) \cap \operatorname{top}(\pi_2) \neq \varnothing$, since the intersection contains e, the two edges $(\Delta_0, \Delta_0 \circ E_1)$ and $(\Delta_0, \Delta_0 \circ E_2)$ of $X(\mathcal{P}, w)$ do not span a square. Therefore, the hyperplanes J_1 and J_2 dual to these two edges, respectively, are tangent. For $i=1,2$, noticing that Δ_i is a prefix of $\Delta_0 \circ E_i$ whereas it is not a prefix of Δ_0, we deduce from Proposition 5.7 that the minimal diagram associated to J_i is precisely Δ_i. On the other hand, the hyperplanes corresponding to Δ_1 and Δ_2 are clearly dual edges of A and B, respectively. A fortiori, J_1 and J_2 must be transverse (in $Y = A \times B$). Therefore, we have found an interosculation in the CAT(0) cube complex $X(\mathcal{P}, w)$, which is impossible.

Thus, we have proved that, for every $\Delta_1 \in A$ and $\Delta_2 \in B$, $\operatorname{supp}(\Delta_1) \cap \operatorname{supp}(\Delta_2) = \varnothing$. In particular, we can write $w = u_1 v_1 \cdots u_k v_k$ in Σ^+ for some $u_1, v_1, \ldots, u_k, v_k \in \Sigma^+$ such that $\operatorname{supp}(\Delta_1) \subset u_1 \cup \cdots \cup u_k$ and $\operatorname{supp}(\Delta_2) \subset v_1 \cup \cdots \cup v_k$ for every $\Delta_1 \in A$

and $\Delta_2 \in B$. By construction, it is clear that

$$A \subset X(\mathcal{P}, u_1) \times \cdots \times X(\mathcal{P}, u_k) \quad \text{and} \quad B \subset X(\mathcal{P}, v_1) \times \cdots \times X(\mathcal{P}, v_k).$$

Now, let $\Delta \in Y = A \times B$ be vertex and let $\Delta_1 \in A$, $\Delta_2 \in B$ denote its coordinates; they are also the projections of Δ onto A and B, respectively. By the previous remark, we can write

$$\Delta_1 = U_1 + \epsilon(v_1) + \cdots + U_k + \epsilon(v_k), \quad \Delta_2 = \epsilon(u_1) + V_1 + \cdots + \epsilon(u_k) + V_k$$

for some $(u_i, *)$–diagrams U_i and $(v_i, *)$–diagrams V_i. Set $\Xi = U_1 + V_1 + \cdots U_k + V_k$. Noticing that $d(\Xi, A) \geq \#V_1 + \cdots + \#V_k$ and $d(\Xi, \Delta_1) = \#V_1 + \cdots + \#V_k$, we deduce that Δ_1 is the projection of Ξ onto A. Similarly, we show that Δ_2 is the projection of Ξ onto B. Consequently,

$$\Delta = \Xi \in X(\mathcal{P}, u_1) \times X(\mathcal{P}, v_1) \times \cdots \times X(\mathcal{P}, u_k) \times X(\mathcal{P}, v_k),$$

which concludes the proof. □

By combining Theorem 5.49 and Proposition 5.50, we are finally able to prove:

Corollary 5.51 *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword whose class modulo $\mathcal{P}$ is finite. Then either $D(\mathcal{P}, w)$ contains a contracting isometry of $X(\mathcal{P}, w)$ or there exists a (w, uv)–diagram Γ such that*

$$D(\mathcal{P}, w) = \Gamma \cdot D(\mathcal{P}, u) \times D(\mathcal{P}, v) \cdot \Gamma^{-1},$$

where $D(\mathcal{P}, u)$ and $D(\mathcal{P}, v)$ are nontrivial.

Proof According to Theorem 5.49, if $D(\mathcal{P}, w)$ does not contain a contracting isometry of $X(\mathcal{P}, w)$, then there exists a convex $D(\mathcal{P}, w)$–invariant subcomplex $Y \subset X(\mathcal{P}, w)$ which splits as the Cartesian product of two unbounded subcomplexes A and B. Up to conjugating by the smallest diagram of Y, we may assume without loss of generality that Y contains $\epsilon(w)$. Let Γ be the $(w, u_1 v_1 \cdots u_k v_k)$–diagram given by Proposition 5.50. Notice that, because $\epsilon(w) \in Y$ and that Y is $D(\mathcal{P}, w)$–invariant, necessarily any spherical diagram belongs to Y. Thus, we deduce from Lemma 5.52 below that

$$D(\mathcal{P}, w) = \Gamma \cdot D(\mathcal{P}, u_1) \times D(\mathcal{P}, v_1) \times \cdots \times D(\mathcal{P}, u_k) \times D(\mathcal{P}, v_k) \cdot \Gamma^{-1}.$$

Because $A \subset \Gamma \cdot X(\mathcal{P}, u_1) \times \cdots \times X(\mathcal{P}, u_k) \cdot \Gamma^{-1}$, we deduce from the unboundedness of A that $X(\mathcal{P}, u_1) \times \cdots \times X(\mathcal{P}, u_k)$ is infinite; since the class of w modulo $\mathcal{P}$ is finite, this implies that there exists some $1 \leq i \leq k$ such that $D(\mathcal{P}, u_i)$ is nontrivial. Similarly, we show that there exists some $1 \leq j \leq k$ such that $D(\mathcal{P}, v_j)$ is nontrivial.

Now, write $u_1v_1\cdots u_kv_k=uv$ in Σ^+, where $u,v\in\Sigma^+$ are two subwords such that one contains u_i and the other v_j. Then

$$D(\mathcal{P},w)=\Gamma\cdot D(\mathcal{P},u)\times D(\mathcal{P},v)\cdot\Gamma^{-1},$$

where $D(\mathcal{P},u)$ and $D(\mathcal{P},v)$ are nontrivial. □

Lemma 5.52 *Let $\mathcal{P}=\langle\Sigma\mid\mathcal{R}\rangle$ be a semigroup presentation and $w\in\Sigma^+$ a baseword whose class modulo $\mathcal{P}$ is finite. If a spherical diagram Δ decomposes as a sum $\Delta_1+\cdots+\Delta_n$, then each Δ_i is a spherical diagram.*

Proof Suppose by contradiction that one factor is not spherical. Let Δ_k be the leftmost factor which is not spherical, i.e. Δ_l is spherical, say a (x_l,x_l)–diagram, for every $l<k$. For $l\geq k$, say that Δ_l is a (x_l,y_l)–diagram. In Σ^+, we have the equality

$$x_1\cdots x_{k-1}x_kx_{k+1}\cdots x_n=\mathrm{top}(\Delta)=\mathrm{bot}(\Delta)=x_1\cdots x_{k-1}y_ky_{k+1}\cdots y_n,$$

hence $x_k\cdots x_n=y_k\cdots y_n$. Because the equality holds in Σ^+, notice that $\lg(x_k)=\lg(y_k)$ implies $x_k=y_k$, which is impossible since we supposed that Δ_k is not spherical. Therefore, two cases may happen: either $\lg(x_k)>\lg(y_k)$ or $\lg(x_k)<\lg(y_k)$. Replacing Δ with Δ^{-1} if necessary, we may suppose without loss of generality that we are in the latter case. Thus, x_k is a prefix of y_k: there exists some $y\in\Sigma^+$ such that $y_k=x_ky$. On the other hand, because Δ_k is a (x_k,y_k)–diagram, the equality $x_k=y_k$ holds modulo $\mathcal{P}$. A fortiori, $x_k=x_ky$ modulo $\mathcal{P}$. We deduce that

$$x_1\cdots x_{k-1}x_ky^nx_{k+1}\cdots x_n\in[w]_{\mathcal{P}}$$

for every $n\geq 1$. This implies that $[w]_{\mathcal{P}}$ is infinite, contradicting our hypothesis. □

Proof of Theorem 5.48 We argue by induction on the algebraic dimension of the considered diagram group, i.e. the maximal rank of a free abelian subgroup. If the algebraic dimension is zero, there is nothing to prove since the diagram group is trivial in this case. From now on, suppose that our diagram group $D(\mathcal{P},w)$ has algebraic dimension at least one. By applying Corollary 5.51, we deduce that either $D(\mathcal{P},w)$ contains a contracting isometry of $X(\mathcal{P},w)$, and we conclude that $D(\mathcal{P},w)$ is either cyclic or acylindrically hyperbolic, or there exist two words $u,v\in\Sigma^+$ and a (w,uv)–diagram Γ such that $D(\mathcal{P},w)=\Gamma\cdot(D(\mathcal{P},u)\times D(\mathcal{P},v))\cdot\Gamma^{-1}$, where $D(\mathcal{P},u)$ and $D(\mathcal{P},v)$ are nontrivial. In the first case, we are done; in the latter case, because $D(\mathcal{P},u)$ and $D(\mathcal{P},v)$ are nontrivial (and torsion-free, as any diagram group), their algebraic dimensions are strictly less than the one of $D(\mathcal{P},w)$, so that we can apply our induction hypothesis to find a $(u, u_1 \cdots u_r)$–diagram Φ and a $(v, v_1 \cdots v_s)$–diagram Ψ such that

$$D(\mathcal{P}, u) = \Phi \cdot D(\mathcal{P}, u_1) \times \cdots \times D(\mathcal{P}, u_r) \cdot \Phi^{-1}$$

and, similarly,

$$D(\mathcal{P}, v) = \Psi \cdot D(\mathcal{P}, v_1) \times \cdots \times D(\mathcal{P}, v_s) \cdot \Psi^{-1},$$

where, for every $1 \leq i \leq r$ and $1 \leq j \leq s$, $D(\mathcal{P}, u_i)$ and $D(\mathcal{P}, v_j)$ are either infinite cyclic or acylindrically hyperbolic. Now, if we set $\Xi = \Gamma \cdot (\Phi + \Psi)$, then

$$D(\mathcal{P}, w) = \Xi \cdot D(\mathcal{P}, u_1) \times \cdots \times D(\mathcal{P}, u_r) \times D(\mathcal{P}, v_1) \times \cdots \times D(\mathcal{P}, v_s) \cdot \Xi^{-1}$$

is the decomposition we are looking for.

Finally, the inequality $n \leq \dim_{\mathrm{alg}} D(\mathcal{P}, w)$ is clear since diagram groups are torsion-free, and the inequality $\dim_{\mathrm{alg}} D(\mathcal{P}, w) \leq \dim X(\mathcal{P}, w)$ is a direct consequence of [10, Corollary 5]. □

We conclude this section by showing that, in the cocompact case, being a contracting isometry can be characterized algebraically. For convenience, we introduce the following definition:

Proposition 5.53 *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword whose class modulo $\mathcal{P}$ is finite. If $g \in D(\mathcal{P}, w) - \{\epsilon(w)\}$, the following statements are equivalent*:

(i) *g is a contracting isometry of $X(\mathcal{P}, w)$*;

(ii) *the centralizer of g is infinite cyclic*;

(iii) *g does not belong to a large canonical subgroup.*

Proof First of all, notice that the centralizer of a contracting isometry in a group acting geometrically on a CAT(0) cube complex is automatically virtually cyclic. This observation can be proved directly by showing that our centralizer acts geometrically on the union of all the axes of the contracting isometry we are looking at and by noticing that this union is a quasiline. Instead, for short, we give an indirect argument. By comparing Theorem 3.13 and [6, Theorem 4.2], we know that a contracting isometry of our CAT(0) cube complex with respect to the combinatorial metric is also contracting with respect to the CAT(0) metric. The desired conclusion follows by combining [30] and [26, Theorem 1.4 and Corollary 6.9]. Because diagram groups are torsion-free, we deduce the implication (i) $\Longrightarrow$ (ii).

Suppose that there exists a (w, uv)–diagram Γ such that $g \in \Gamma \cdot D(\mathcal{P}, u) \times D(\mathcal{P}, v) \cdot \Gamma^{-1}$, where $D(\mathcal{P}, u)$ and $D(\mathcal{P}, v)$ are nontrivial. Let $g_1 \in D(\mathcal{P}, u)$ and $g_2 \in D(\mathcal{P}, v)$ be two spherical diagrams such that $g = \Gamma \cdot (g_1 + g_2) \cdot \Gamma^{-1}$. In particular, the centralizer of g contains: $\Gamma \cdot (\langle g_1 \rangle \times \langle g_2 \rangle) \cdot \Gamma^{-1}$ if g_1 and g_2 are nontrivial; $\Gamma \cdot (D(\mathcal{P}, u) \times \langle g_2 \rangle) \cdot \Gamma^{-1}$ if g_1 trivial and g_2 is not; $\Gamma \cdot (\langle g_1 \rangle \times D(\mathcal{P}, v)) \cdot \Gamma^{-1}$ if g_2 is trivial and g_1 is not. We deduce that (ii) $\Longrightarrow$ (iii).

Finally, suppose that g is not a contracting isometry. Up to taking a conjugate of g, we may suppose without loss of generality that g is absolutely reduced. According to Theorem 5.28 and Lemma 5.36, three cases may happen:

(a) $\operatorname{supp}(g)$ has cardinality at least two;

(b) $w = xyz$ in Σ^+ with $\operatorname{supp}(g) = \{y\}$ and $X(\mathcal{P}, x)$ or $X(\mathcal{P}, z)$ infinite;

(c) g^∞ contains a proper infinite prefix.

In case (a), g decomposes as a sum with at least two nontrivial factors, and we deduce from Lemma 5.52 that g belongs to a large canonical subgroup. In case (b), if $X(\mathcal{P}, x)$ is infinite then $D(\mathcal{P}, x)$ is nontrivial since $[x]_{\mathcal{P}}$ is finite, and we deduce that g belongs to the large canonical subgroup $D(\mathcal{P}, x) \times D(\mathcal{P}, yz)$; we argue similarly if $X(\mathcal{P}, z)$ is infinite. In case (c), suppose that g^∞ contains a proper infinite prefix Ξ. For convenience, write $g^\infty = g_1 \circ g_2 \circ \cdots$, where each g_i is a copy of g. Suppose by contradiction that g cannot be decomposed as a sum with at least two nontrivial factors. As a consequence, for each $i \geq 1$, the subdiagram g_i either is contained into Ξ or it contains a cell which does not belong to Ξ but whose top path intersects Ξ. Because Ξ is a proper prefix of g^∞, there exists some index $j \geq 1$ such that g_j is not included into Ξ, and it follows from Lemma 5.31 that, for every $i \geq j$, g_i satisfies the same property. Let Ξ_{j+r} denote the prefix of $g_1 \circ \cdots \circ g_{j+r}$ induced by Ξ. We know that, for every $0 \leq s \leq r$, the subdiagram g_{j+s} contains a cell which does not belong to Ξ but whose top path intersects Ξ. This implies that $\operatorname{bot}(\Xi_{j+r})$ has length at least r. On the other hand, the finiteness of $[w]_{\mathcal{P}}$ implies that the cardinality of $[\operatorname{bot}(\Xi_{j+r})]_{\mathcal{P}}$ is bounded by a constant which does not depend on r. Thus, we get a contradiction if r is chosen sufficiently large. We have proved that g decomposes as a sum with at least two nontrivial factors. As in case (a), we deduce from Lemma 5.52 that g belongs to a large canonical subgroup. □

Remark 5.54 The implication (iii) $\Longrightarrow$ (ii) also follows from the description of centralizers in diagram groups [13, Theorem 15.35], since they are canonical subgroups.

Appendix Combinatorial boundary vs other boundaries

In this appendix, we compare the combinatorial boundary of a CAT(0) cube complex with three other boundaries, namely the *simplicial boundary*, introduced in [16]; the *Roller boundary*, introduced in [27]; and its variation studied in [15]. In what follows, we will always assume that our CAT(0) cube complexes have countably many hyperplanes. For instance, this happens when they are locally finite.

A.1 Simplicial boundary

For convenience, we will say that a CAT(0) cube complex is *ω–dimensional* if it does not contain an infinite collection of pairwise transverse hyperplanes.

Let X be an ω–dimensional CAT(0) cube complex. A *unidirection boundary set* (or *UBS* for short) is an infinite collection of hyperplanes $\mathcal{U}$ not containing any facing triple which is

- *inseparable*, i.e. each hyperplane separating two elements of $\mathcal{U}$ belongs to $\mathcal{U}$;
- *unidirectional*, i.e. any hyperplane of $\mathcal{U}$ bounds a halfspace containing only finitely many elements of $\mathcal{U}$.

According to [16, Theorem 3.10], any UBS is commensurable to a disjoint union of *minimal* UBSs, where a UBS $\mathcal{U}$ is *minimal* if any UBS $\mathcal{V}$ satisfying $\mathcal{V} \subset \mathcal{U}$ must be commensurable to $\mathcal{U}$; the number of these minimal UBSs is the dimension of the UBS we are considering. Then, to any commensurability class of a k–dimensional UBS is associated a *k–simplex at infinity* whose faces correspond to its sub-UBSs. All together, these simplices at infinity define a simplicial complex, called the *simplicial boundary* $\partial_\triangle X$ of X. See [16] for more details.

The typical example of a UBS is the set of the hyperplanes intersecting a given combinatorial ray. A simplex arising in this way is called *visible*. Therefore, to any point in the combinatorial boundary naturally corresponds a simplex of the simplicial boundary. The following proposition is clear from the definitions:

Proposition A.1 *Let X be an ω–dimensional* CAT(0) *cube complex. Then $(\partial^c X, \leq)$ is isomorphic to the face poset of the visible part of $\partial_\triangle X$.*

In particular, when X is *fully visible*, i.e. when every simplex of the simplicial boundary is visible, the two boundaries essentially define the same object. Notice however that Farley cube complexes are not necessarily fully visible. For example, if

$\mathcal{P} = \langle a, b, p \mid a = ap,\ b = pb \rangle$, then $X(\mathcal{P}, ab)$ is a "diagonal halfplane" and is not fully visible. Moreover, although Farley cube complexes are complete, they are not necessarily ω–dimensional, so that the simplicial boundary may not be well defined. This is the case for the cube complex associated to Thompson's group F, namely $X(\mathcal{P}, x)$, where $\mathcal{P} = \langle x \mid x = x^2 \rangle$; or the cube complex associated to the lamplighter group $\mathbb{Z} \wr \mathbb{Z}$, namely $X(\mathcal{P}, ab)$, where $\mathcal{P} = \langle a, b, p, q, r \mid a = ap,\ b = pb,\ p = q,\ q = r,\ r = p \rangle$.

A.2 Roller boundary

A *pocset* $(\Sigma, \leq, ^*)$ is a partially ordered set $(\Sigma, \leq)$ with an order-reversing involution * such that a and a^* are not $\leq$–comparable for every $a \in \Sigma$. A subset $\alpha \subset \mathfrak{P}(\Sigma)$ is an *ultrafilter* if

- for every $a \in \Sigma$, exactly one element of $\{a, a^*\}$ belongs to α;
- if $a, b \in \Sigma$ satisfy $a \leq b$ and $a \in \alpha$, then $b \in \alpha$.

Naturally, to every CAT(0) cube complex X is associated a pocset: the set of halfspaces $\mathcal{U}(X)$ with the inclusion and the complementary operation. In this way, every vertex can be thought of as an ultrafilter, called a *principal ultrafilter*; see for example [29]. For convenience, let X° denote the set of ultrafilters of the pocset associated to our cube complex. Notice that X° is naturally a subset of $2^{\mathcal{U}(X)}$, and can be endowed with the Tykhonov topology.

Now, the *Roller boundary* $\mathfrak{R}X$ of X is defined as the space of nonprincipal ultrafilters of X° endowed with the Tykhonov topology. Below, we give an alternative description of $\mathfrak{R}X$.

Given a CAT(0) cube complex X, fix a base vertex $x \in X$, and let $\mathfrak{S}_x X$ denote the set of combinatorial rays starting from x up to the following equivalence relation: two rays r_1 and r_2 are equivalent if $\mathcal{H}(r_1) = \mathcal{H}(r_2)$. Now, if $\xi_1, \xi_2 \in \mathfrak{S}_x X$, we define the distance between ξ_1 and ξ_2 in $\mathfrak{S}_x X$ by

$$d(\xi_1, \xi_2) = 2^{-\min\{d(x,J) \mid J \in \mathcal{H}(r_1) \oplus \mathcal{H}(r_2)\}},$$

where $\oplus$ denotes the symmetric difference. In fact, as a consequence of the inclusion $A \oplus B \subset (A \oplus B) \cup (B \oplus C)$ for all sets A, B and C, the distance d turns out to be *ultrametric*, i.e.

$$d(\xi_1, \xi_3) \leq \max(d(\xi_1, \xi_2), d(\xi_2, \xi_3))$$

for every $\xi_1, \xi_2, \xi_3 \in \mathfrak{S}_x X$. Also, notice that this distance is well defined since its expression does not depend on the representatives we chose.

Given a combinatorial ray r, let $\alpha(r)$ denote the set of halfspaces U such that r lies eventually in U.

Proposition A.2 *The map $r \mapsto \alpha(r)$ induces a homeomorphism $\mathfrak{S}_x X \to \mathfrak{R}X$.*

Proof If r is a combinatorial ray starting from x, it is clear that $\alpha(r)$ is a nonprincipal ultrafilter, i.e. $\alpha(r) \in \mathfrak{R}X$. Now, if for every hyperplane J we denote by J^+ (resp. J^-) the halfspace delimited by J which does not contain x (resp. which contains x), then we have

$$\alpha(r) = \{J^+ \mid J \in \mathcal{H}(r)\} \sqcup \{J^- \mid J \notin \mathcal{H}(r)\}.$$

We first deduce that $\alpha(r_1) = \alpha(r_2)$ for any pair of equivalent combinatorial rays r_1 and r_2 starting from x, so that we get a well-defined induced map $\alpha\colon \mathfrak{S}_x X \to \mathfrak{R}X$; secondly, the injectivity of α follows, since $\alpha(r_1) = \alpha(r_2)$ clearly implies $\mathcal{H}(r_1) = \mathcal{H}(r_2)$.

Now, let $\eta \in \mathfrak{R}X$ be a nonprincipal ultrafilter, and let $\mathcal{H}^+(\eta)$ denote the set of the hyperplanes J satisfying $J^+ \in \eta$. Notice that, if $\mathcal{H}^+(\eta)$ is finite, then all but finitely many halfspaces of η contain x, which implies that η is principal; therefore, $\mathcal{H}^+(\eta)$ is infinite, say $\mathcal{H}^+(\eta) = \{J_1, J_2, \dots\}$. On the other hand, $J_i^+ \cap J_j^+$ is nonempty for every $i, j \geq 1$, so, for every $k \geq 1$, the intersection $C_k = \bigcap_{i=1}^k J_i^+$ is nonempty; let x_k denote the combinatorial projection of x onto C_k. According to Lemma 2.9, x_{k+1} is the combinatorial projection of x_k onto C_{k+1} for every $k \geq 1$. Therefore, by applying Proposition 2.6, we know that there exists a combinatorial ray ρ starting from x and passing through all the x_k. Now, if $J \in \mathcal{H}(\rho)$, there exists some $k \geq 1$ such that J separates x and x_k, and a fortiori x and C_k according to Lemma 2.8. We deduce that either J belongs to $\{J_1, \dots, J_k\}$ or J separates x and some hyperplane of $\{J_1, \dots, J_k\}$; equivalently, there exists some $1 \leq l \leq k$ such that $J_l^+ \subset J^+$. Because $J_l^+ \in \eta$, we deduce that $J^+ \in \eta$. Thus, we have proved that $\mathcal{H}(\rho) = \mathcal{H}^+(\eta)$. It follows that $\alpha(\rho) = \eta$. The surjectivity of α is proved.

Finally, it is an easy exercise of general topology to verify that α is indeed a homeomorphism. The details are left to the reader. □

Therefore, the combinatorial boundary $\partial^c X$ may be thought of as a quotient of the Roller boundary. This is the point of view explained in the next section.

It is worth noticing that this description of the Roller boundary allows us to give a precise description of the Roller boundary of Farley cube complexes:

Proposition A.3 *Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and $w \in \Sigma^+$ a baseword. The Roller boundary of $X(\mathcal{P}, w)$ is homeomorphic to the space of the infinite reduced w–diagrams endowed with the metric*

$$(\Delta_1, \Delta_2) \mapsto 2^{-\min\{\#\Delta \mid \Delta \text{ is not a prefix of both } \Delta_1 \text{ and } \Delta_2\}},$$

where $\#(\cdot)$ denotes the number of cells of a diagram.

A.3 Guralnik boundary

Given a pocset $(\Sigma, \leq, ^*)$, Guralnik defines the boundary $\mathfrak{R}^* \Sigma$ as the set of almost-equality classes of ultrafilters, partially ordered by the relation: $\Sigma_1 \leq \Sigma_2$ if $\Sigma_2 \cap \overline{\Sigma}_1 \neq \varnothing$, where $\overline{(\cdot)}$ denotes the closure in $\mathfrak{R}X$ with respect to the Tykhonov topology. See [15] for more details.

In particular, if X is a CAT(0) cube complex, it makes sense to define the boundary $\mathfrak{R}^* X$ as the previous boundary of the associated pocset. Notice that $\mathfrak{R}^* X$ contains a particular point, corresponding to the almost-equality class of the principal ultrafilters; let $\mathfrak{p}$ denote this point. Then $\mathfrak{R}^* X \setminus \{\mathfrak{p}\}$ is naturally the quotient of the Roller boundary $\mathfrak{R}X$ by the almost-equality relation.

Remark A.4 Although Guralnik considers only ω–dimensional pocsets in [15], the definition makes sense for every pocset, and in particular is well defined for every CAT(0) cube complex. Notice also that the boundary $\mathfrak{R}^*$ is called the *Roller boundary* there, whereas the terminology now usually refers to the boundary defined in the previous section.

Proposition A.5 *The posets $(\mathfrak{R}^* X \setminus \{\mathfrak{p}\}, \leq)$ and $(\partial^c X, \leq)$ are isomorphic.*

Proof If r is an arbitrary combinatorial ray, choose a second ray $\rho \in \mathfrak{S}_x X$ equivalent to it, and set $\varphi(r) \in \mathfrak{R}^* X$ as the almost-equality class of $\alpha(\rho)$, where $\alpha\colon \mathfrak{S}_x X \to \mathfrak{R}(X)$ is the isomorphism given by Proposition A.2. Notice that, if $\rho' \in \mathfrak{S}_x X$ is equivalent to r, then $\alpha(\rho)$ and $\alpha(\rho')$ are almost equal, so we deduce that $\varphi(r)$ does not depend on our choice of ρ. As another consequence, φ must be constant on the equivalence classes of rays, so that we get a well-defined map

$$\varphi\colon \partial^c X \to \mathfrak{R}^* X \setminus \{\mathfrak{p}\}.$$

Clearly, φ is surjective. Now, if $r_1, r_2 \in \partial^c X$ satisfy $\varphi(r_1) = \varphi(r_2)$, then $\alpha(r_1) =_a \alpha(r_2)$; using the notation of the proof of Proposition A.2, this implies that

$$\mathcal{H}(r_1) = \mathcal{H}^+(\alpha(r_1)) =_a \mathcal{H}^+(\alpha(r_2)) = \mathcal{H}(r_2),$$

hence $r_1 = r_2$ in $\partial^c X$. Thus, φ is injective.

Let $r_1, r_2 \in \partial^c X$ be two combinatorial rays satisfying $r_1 \prec r_2$, i.e. $\mathcal{H}(r_1) \subset_a \mathcal{H}(r_2)$. According to Lemmas 4.2 and 4.3, we may suppose without loss of generality that $r_1(0) = r_2(0) = x$ and $\mathcal{H}(r_1) \subset \mathcal{H}(r_2)$. Let $\mathcal{H}(r_2) = \{J_1, J_2, \dots\}$. For every $k \geq 1$, let $C_k = \bigcap_{i=1}^k J_i^+$ and let x_k denote the combinatorial projection of x onto C_k. Let $\mathcal{H}(r_1) = \{H_1, H_2, \dots\}$, and, for every $i \geq 1$, let $K_i = \bigcap_{j=1}^i H_j \cap C_k$; notice that, because $\mathcal{H}(r_1) \subset \mathcal{H}(r_2)$, K_i is nonempty. Fix some $k \geq 1$. Finally, for every $i \geq 1$, let y_i denote the combinatorial projection of x_k onto K_i. By combining Lemma 2.9 and Proposition 2.6, we know that there exists a combinatorial ray starting from x_k and passing through all the y_i; let ρ_k denote the concatenation of a combinatorial geodesic between x and x_k with the previous ray.

If ρ_k is not a combinatorial ray, then there exists a hyperplane J separating x and x_k, and x_k and some y_j. On the other hand, it follows from Lemma 2.8 that J must separate x and C_k, so that J cannot separate x_k and y_j since $x_k, y_j \in C_k$. Therefore, ρ_k is a combinatorial ray.

If $\mathcal{H}_k$ denotes the set of the hyperplanes separating x and x_k, then

$$\mathcal{H}(\rho_k) = \mathcal{H}_k \cup \mathcal{H}(r_1).$$

Indeed, if J is a hyperplane intersecting ρ_k which does not separate x and x_k, then J must separate x_k and some y_j; a fortiori, J must be disjoint from K_j according to Lemma 2.8. Thus, if z is any vertex of $r_1 \cap K_j$, we know that J separates x_k and y_j, but cannot separate x and x_k, or y_j and z; therefore, necessarily J separates x and z, hence $J \in \mathcal{H}(r_1)$. Conversely, let J be a hyperplane of $\mathcal{H}(r_1)$. In particular, $J \in \mathcal{H}(r_2)$, so ρ_k eventually lies in J^+. Because $\rho_k(0) = x \notin J^+$, we conclude that $J \in \mathcal{H}(\rho_k)$.

In particular, because $\mathcal{H}_k$ is finite, we deduce that $\mathcal{H}(\rho_k) =_a \mathcal{H}(r_1)$, i.e. $r_1 = \rho_k$ in $\partial^c X$. On the other hand, it is clear that, for any finite subset $\mathcal{H} \subset \mathcal{H}(r_2)$, there exists some $k \geq 1$ sufficiently large that $\mathcal{H} \subset \mathcal{H}_k$. Therefore, the sequence of ultrafilters $(\alpha(\rho_k))$ converges to $\alpha(r_2)$ in the Roller boundary.

Consequently, $\alpha(r_2)$ belongs to the closure of the almost-equality class of $\alpha(r_1)$. This precisely means that $\varphi(r_1) \leq \varphi(r_2)$ in $\mathfrak{R}^* X$. We have proved that φ is a poset morphism.

Now, let $r_1, r_2 \in \partial^c X$ be such that $\varphi(r_1) \leq \varphi(r_2)$ in $\mathfrak{R}^* X$. This means that there exists a combinatorial ray ρ such that $\alpha(\rho)$ belongs to the intersection between the almost-equality class of $\alpha(r_2)$ and the closure of the almost-equality class of $\alpha(r_1)$ in the Roller boundary. So there exists a sequence of combinatorial rays (ρ_k) such that $\rho_k = r_1$ in $\partial^c X$ for every $k \geq 1$ and $\alpha(\rho_k) \to \alpha(\rho)$. Let r_0 be the minimal element of the class of r_1 in $\mathfrak{S}_x X$ which is given by Lemma A.6 below. Then $\mathcal{H}(r_0) \subset \mathcal{H}(\rho_k)$ for every $k \geq 1$, hence $\mathcal{H}(r_0) \subset \mathcal{H}(\rho)$. Therefore,

$$\mathcal{H}(r_1) \subset_a \mathcal{H}(r_0) \subset \mathcal{H}(\rho) =_a \mathcal{H}(r_2),$$

hence $r_1 \prec r_2$ in $\partial^c X$. We have proved that φ^{-1} is a poset morphism as well. □

Lemma A.6 *Let $r \in \mathfrak{S}_x X$ be a combinatorial ray. There exists $r_0 \in \mathfrak{S}_x X$ equivalent to r such that, for any combinatorial ray $\rho \in \mathfrak{S}_x X$ equivalent to r, we have $\mathcal{H}(r_0) \subset \mathcal{H}(\rho)$.*

Proof Let $\mathcal{H} \subset \mathcal{H}(r)$ be the set of the hyperplanes $J \in \mathcal{H}(r)$ such that r does not stay in a neighborhood of J; say $\mathcal{H} = \{J_1, J_2, \dots\}$. For every $k \geq 1$, let $C_k = \bigcap_{i=1}^k J_i^+$ and let x_k denote the combinatorial projection of x onto C_k. By combining Lemma 2.9 and Proposition 2.6, we know that there exists a combinatorial ray r_0 starting from x and passing through all the x_k. We claim that r_0 is the ray we are looking for.

Let $\rho \in \mathfrak{S}_x X$ be a combinatorial ray equivalent to r, and let $J \in \mathcal{H}(r_0)$ be a hyperplane. By construction, J separates x and some x_k; a fortiori, it follows from Lemma 2.8 that J separates x and C_k. On the other hand, given some $1 \leq l \leq k$, because r does not stay in neighborhood of J_l, there exist infinitely many hyperplanes of $\mathcal{H}(r)$ which are included into J_l^+; therefore, since ρ is equivalent to r, we deduce that ρ eventually lies in J_l^+. A fortiori, ρ eventually lies in C_k. Consequently, because J separates $x = \rho(0)$ and C_k, we conclude that $J \in \mathcal{H}(\rho)$. □

References

[1] **J Behrstock**, **R Charney**, *Divergence and quasimorphisms of right-angled Artin groups*, Math. Ann. 352 (2012) 339–356 MR

[2] **M Bestvina**, **K Bromberg**, **K Fujiwara**, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. 122 (2015) 1–64 MR

[3] **M R Bridson**, **A Haefliger**, *Metric spaces of non-positive curvature*, Grundl. Math. Wissen. 319, Springer (1999) MR

[4] **K S Brown**, *Cohomology of groups*, Graduate Texts in Mathematics 87, Springer (1982) MR

[5] **P-E Caprace**, **M Sageev**, *Rank rigidity for CAT(0) cube complexes*, Geom. Funct. Anal. 21 (2011) 851–891 MR

[6] **R Charney**, **H Sultan**, *Contracting boundaries of* CAT(0) *spaces*, J. Topol. 8 (2015) 93–117 MR

[7] **V Chepoi**, *Graphs of some* CAT(0) *complexes*, Adv. in Appl. Math. 24 (2000) 125–179 MR

[8] **D S Farley**, *Finiteness and* CAT(0) *properties of diagram groups*, Topology 42 (2003) 1065–1082 MR

[9] **A Genevois**, *Coning-off* CAT(0) *cube complexes*, preprint (2016) arXiv

[10] **A Genevois**, *Hyperbolic diagram groups are free*, Geom. Dedicata 188 (2017) 33–50 MR

[11] **A Genevois**, *Hyperplanes of Squier's cube complexes*, Algebr. Geom. Topol. 18 (2018) 3205–3256 MR

[12] **V N Gerasimov**, *Semi-splittings of groups and actions on cubings*, from "Алгебра, геометрия, анализ и математическая физика" (Y G Reshetnyak, L A Bokut, S K Vodopyanov, I A Taĭmanov, editors), Izdat. Ross. Akad. Nauk Sib. Otd. Inst. Mat., Novosibirsk (1997) 91–109 MR In Russian; translated as *Fixed-point-free actions on cubings*, Siberian Adv. Math. 8 (1998) 36–58

[13] **V Guba**, **M Sapir**, *Diagram groups*, Mem. Amer. Math. Soc. 620, Amer. Math. Soc., Providence, RI (1997) MR

[14] **V S Guba**, **M V Sapir**, *On subgroups of the R Thompson group F and other diagram groups*, Mat. Sb. 190 (1999) 3–60 MR In Russian; translated in Sb. Math. 190 (1999) 1077–1130

[15] **D Guralnik**, *Coarse decompositions of boundaries for* CAT(0) *groups*, preprint (2006) arXiv

[16] **M F Hagen**, *The simplicial boundary of a* CAT(0) *cube complex*, Algebr. Geom. Topol. 13 (2013) 1299–1367 MR

[17] **F Haglund**, *Isometries of* CAT(0) *cube complexes are semi-simple*, preprint (2007) arXiv

[18] **F Haglund**, *Finite index subgroups of graph products*, Geom. Dedicata 135 (2008) 167–209 MR

[19] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR

[20] **E R V Kampen**, *On some lemmas in the theory of groups*, Amer. J. Math. 55 (1933) 268–273 MR

[21] **V Kilibarda**, *On the algebra of semigroup diagrams*, Internat. J. Algebra Comput. 7 (1997) 313–338 MR

[22] **I J Leary**, *A metric Kan–Thurston theorem*, J. Topol. 6 (2013) 251–284 MR

[23] **J P McCammond**, **D T Wise**, *Fans and ladders in small cancellation theory*, Proc. London Math. Soc. 84 (2002) 599–644 MR

[24] **A Minasyan**, **D Osin**, *Acylindrical hyperbolicity of groups acting on trees*, Math. Ann. 362 (2015) 1055–1105 MR

[25] **D V Osin**, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc. 843, Amer. Math. Soc., Providence, RI (2006) MR

[26] **D Osin**, *Acylindrically hyperbolic groups*, Trans. Amer. Math. Soc. 368 (2016) 851–888 MR

[27] **M Roller**, *Pocsets, median algebras and group actions: an extended study of Dunwoody's construction and Sageev's theorem*, PhD thesis, Universität Regensburg (1998) arXiv

[28] **M Sageev**, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995) 585–617 MR

[29] **M Sageev**, CAT(0) *cube complexes and groups*, from "Geometric group theory" (M Bestvina, M Sageev, K Vogtmann, editors), IAS/Park City Math. Ser. 21, Amer. Math. Soc., Providence, RI (2014) 7–54 MR

[30] **A Sisto**, *Contracting elements and random walks*, J. Reine Angew. Math. 742 (2018) 79–114 MR

[31] **H Sultan**, *Hyperbolic quasi-geodesics in* CAT(0) *spaces*, Geom. Dedicata 169 (2014) 209–224 MR

[32] **D Wise**, *The structure of groups with a quasiconvex hierarchy*, preprint (2012) Available at https://www.math.u-psud.fr/~haglund/Hierarchy29Feb2012.pdf

Université Paris-Sud
Orsay, France

anthony.genevois@math.u-psud.fr

Received: 28 July 2017 Revised: 16 September 2018