

Hyperbolicity in non-metric cubical small-cancellation

Macarena Arenas^{1,2}  | Kasia Jankiewicz³ | Daniel T. Wise⁴

¹Clare College, University of Cambridge, Cambridge, UK

²Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK

³Department of Mathematics, University of California, Santa Cruz, California, USA

⁴Department of Mathematics & Statistics, McGill University, Montreal, Quebec, Canada

Correspondence

Macarena Arenas, Clare College, University of Cambridge, Cambridge CB2 1TL, UK.

Email: mcr59@cam.ac.uk

Funding information

Cambridge Trust & Newnham College Scholarship; Denman Baynes Junior Research Fellowship; NSF, Grant/Award Numbers: DMS-2203307, DMS-2238198; NSERC

Abstract

Given a non-positively curved cube complex X , we prove that the quotient of $\pi_1 X$ defined by a cubical presentation $\langle X \mid Y_1, \dots, Y_s \rangle$ satisfying sufficient non-metric cubical small-cancellation conditions is hyperbolic provided that $\pi_1 X$ is hyperbolic. This generalises the fact that finitely presented classical $C(7)$ small-cancellation groups are hyperbolic.

MSC 2020

20F06, 20F67, 20F65 (primary)

1 | INTRODUCTION

A cubical presentation is a higher dimensional generalisation of a classical group presentation in terms of generators and relators. A non-positively curved cube complex X plays the role of the ‘generators’, and the ‘relators’ are local isometries of non-positively curved cube complexes $Y_i \hookrightarrow X$. The associated group is the quotient of $\pi_1 X$ by the normal closure $\langle\langle \{ \pi_1 Y_i \} \rangle\rangle_{\pi_1 X}$ of $\pi_1 Y_i$. As in the classical setting, this group is the fundamental group of X with the Y_i ’s coned off. Likewise, cubical small-cancellation theory, introduced in [16], is a generalisation of classical small-cancellation theory (see, e.g., [12]). In both the classical and cubical cases, the small-cancellation conditions are

expressed in terms of pieces. A *piece* in a classical presentation is a word that appears in two different places among the relators. The *non-metric small-cancellation condition* $C(p)$ where $p > 1$ asserts that no relator is a concatenation of fewer than p pieces. The *metric small-cancellation condition* $C'(\alpha)$ $\alpha \in (0, 1)$ asserts that $|P| < \alpha|R|$ whenever P is a piece in a relator R . Note that $C'(\frac{1}{p}) \Rightarrow C(p+1)$. Pieces in cubical presentation are defined similarly, and the same implication holds in the cubical case.

Cubical small-cancellation has proven to be a fruitful tool in the study of groups acting on $\text{CAT}(0)$ cube complexes. It was used by Wise as a step in his proof of the Malnormal Special Quotient Theorem [16], and as such, played a crucial role in the proofs of the Virtual Haken and Virtual Fibering conjectures [1]. Cubical presentations and cubical small-cancellation theory were also studied and utilised in [2–4, 6–9, 14]. Although classical small cancellation groups have virtual cohomological dimension ≤ 2 [11], there exist cubical small cancellation groups with arbitrarily large virtual cohomological dimension, which is moreover controlled by $\text{cd } \pi_1 X$ and $\text{cd } \pi_1 Y_i$ (this is proven in a forthcoming paper by Arenas, which is the continuation of 2).

To illustrate the difference between metric and non-metric conditions, consider the following presentation: $\langle a, b \mid a^n w \rangle$. When w is a long messy word (read: small cancellation) starting and ending in b , then the $C(p)$ condition holds for all n . However a^{n-1} is a piece! So, $C'(\alpha)$ fails for sufficiently large n . Similar examples can be produced in the cubical setting. For instance, let $X = S \vee A$ where S is a cubulated surface and A is a circle. Let w be a small-cancellation path in X whose initial and terminal edges lie in S . Let \tilde{w} be the lift of w to \tilde{X} , and let \tilde{W} be the combinatorial convex hull of \tilde{w} . Let a^n be a length n arc that immerses onto A . Let Y be the quotient of $a^n \cup \tilde{W}$, identifying the endpoints of a^n and \tilde{w} . Then $\langle X \mid Y \rangle$ is a cubical presentation satisfying $C(p)$, but not $C'(\frac{1}{p})$, when $n \gg 0$. The pumping lemma shows that for any X with $\pi_1 X$ non-elementary hyperbolic, there are presentations X^* that are non-metric small cancellation, but not metric small cancellation.

It is a fundamental result of classical small-cancellation theory that a group admitting a finite presentation satisfying the classical $C'(\frac{1}{6})$ or $C(7)$ condition is hyperbolic. In analogy with the metric classical small-cancellation case, a cubical $C'(\frac{1}{14})$ presentation $\langle X \mid Y_1, \dots, Y_s \rangle$ yields a hyperbolic group if $\pi_1 X$ is hyperbolic and the Y_i are compact [16, Theorem 4.7]. However, the proof of that result does not extend to the non-metric case. The goal of this paper is to prove the following statement that recovers and generalises the result from the $C'(\frac{1}{14})$ setting.

Theorem 5.1. *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ cubical small-cancellation condition for $p \geq 14$, where X, Y_1, \dots, Y_s are compact, and $\pi_1 X$ is hyperbolic. Then $\pi_1 X^*$ is hyperbolic.*

1.1 | Proof strategy

The main ingredients of the proof are the notion of the piece metric (Definition 4.1) and Papasoglu's thin bigon criterion for hyperbolicity (Proposition 3.1).

The most immediate way of proving hyperbolicity for finitely presented $C'(\frac{1}{6})$ groups is to show that a linear isoperimetric inequality holds for their Cayley complexes. This follows from the fact that $C'(\frac{1}{6})$ presentations satisfy Dehn's algorithm by Greendlinger's Lemma (see for instance [12, V.4.5]). In the $C(7)$ setting, one is no longer guaranteed to have a Dehn presentation (consider the $\langle a, b \mid a^n w \rangle$ examples described above). Instead, the usual way of proving hyperbolicity in

this generality relies on the combinatorial Gauss–Bonnet Theorem. Another way is to realise that reduced disc diagrams satisfy $C'(\frac{1}{6})$ if we regard all pieces as having length 1. In fact, this viewpoint leads to the piece-metric.

To illustrate the basic idea behind our strategy, we sketch a proof of hyperbolicity in the $C(7)$ case using the piece-metric and the thin bigon criterion. The definition of the classical $C(n)$ condition and of all the diagrammatic notions introduced for cubical presentations in Subsection 2.2 can be particularised to this setting, and a version of Greendlinger’s Lemma also holds (see, for instance, [12, V.4.5]).

Illustrative theorem. *Let X be a 2-complex satisfying the $C(7)$ small-cancellation condition. Then X is hyperbolic with the piece metric.*

This implies hyperbolicity of finitely presented $C(7)$ groups, as in that case the piece-metric and the usual combinatorial metric on the Cayley graph are quasi-isometric (see Proposition 4.3).

Proof. We check that all bigons in $X^{(0)}$ are ϵ -thin in the piece metric for some $\epsilon > 0$, and apply Proposition 3.1 to conclude that X is hyperbolic.

Let γ_1, γ_2 be piece-geodesics forming a bigon in X , and let $D \rightarrow X$ be a reduced disc diagram with $\partial D = \gamma_1 \bar{\gamma}_2$. We claim that D is a (possibly degenerate) ladder, and hence that the bigon γ_1, γ_2 is 1-thin, as by definition any two cells in a ladder intersect along at most one piece.

Indeed, by Greendlinger’s Lemma, D is either a ladder, or contains at least three shells and/or spurs. First note that D cannot have spurs, as these can be removed to obtain paths γ'_1, γ'_2 with the same endpoints as γ_1, γ_2 , and which are shorter in the piece metric, thus contradicting that γ_1, γ_2 are piece-geodesics.

If D has at least three shells, then at least one shell S must have its outerpath Q contained in either γ_1 or γ_2 . As both cases are analogous, assume $Q \subset \gamma_1$, and let R be the innerpath of S . As S is a shell and X satisfies $C(7)$, then R is the concatenation of at most 3 pieces, so $|R|_p < |Q|_p$, and the path γ'_1 obtained from γ_1 by traversing R instead of Q is the concatenation of less pieces than γ_1 , contradicting that γ_1 is a piece-geodesic.

Thus, D is a ladder, and the proof is complete. □

1.2 | Structure of the paper

The paper is organised as follows. In Section 2, we give background on cube complexes, cubical group presentations and cubical small-cancellation. In Section 3, we recall a criterion for hyperbolicity for groups acting on graphs. In Section 4, we define and analyse the piece metric. In Section 5, we prove Theorem 5.1.

2 | CUBICAL BACKGROUND

2.1 | Non-positively curved cube complexes

We assume that the reader is familiar with *CAT(0) cube complexes*, which are *CAT(0)* spaces having cell structures where each cell is isometric to a cube. We refer the reader to [5, 10, 15, 16]. A *non-positively curved cube complex* is a cell-complex X whose universal cover \tilde{X} is a *CAT(0)* cube

complex. A *hyperplane* \tilde{H} in \tilde{X} is a subspace whose intersection with each n -cube $[0, 1]^n$ is either empty or consists of the subspace where exactly one coordinate is restricted to $\frac{1}{2}$. For a hyperplane \tilde{H} of \tilde{X} , we let $N(\tilde{H})$ denote its *carrier*, which is the union of all closed cubes intersecting \tilde{H} . The *combinatorial metric* d on the 0-skeleton of a non-positively curved cube complex X is a length metric where the distance between two points is the length of the shortest combinatorial path connecting them. A map $\phi : Y \rightarrow X$ between non-positively curved cube complexes is a *local isometry* if ϕ is locally injective, ϕ maps open cubes homeomorphically to open cubes, and whenever a, b are concatenable edges of Y , if $\phi(a)\phi(b)$ is a subpath of the attaching map of a 2-cube of X , then ab is a subpath of a 2-cube in Y .

2.2 | Cubical presentations

We recall the notion of a cubical presentation, and the cubical small-cancellation conditions from [16].

A *cubical presentation* $\langle X \mid Y_1, \dots, Y_m \rangle$ consists of a non-positively curved cube complex X , and a set of local isometries $Y_i \looparrowright X$ of non-positively curved cube complexes. We use the notation X^* for the cubical presentation above. As a topological space, X^* consists of X with a cone on Y_i attached to X for each i . The vertices of the cones on Y_i 's will be referred to as *cone-vertices* of X^* . The cellular structure of X^* consists of all the original cubes of X , and the ‘pyramids’ over cubes in Y_i with a cone-vertex for the apex.

As mentioned in the introduction, cubical presentations generalise classical group presentations. Indeed, a classical presentation complex associated with a group presentation $G = \langle S \mid R \rangle$ can be viewed as a cubical presentation where the non-positively curved cube complex X is just a wedge of circles, one corresponding to each generator in S . The complexes Y_i correspond to relators r_i in R . Each cycle Y_i has length $|r_i|$, and the local isometry $Y_i \looparrowright X$ is defined by labelling the edges of Y_i with the letters of r_i .

The universal cover \widetilde{X}^* consists of a cube complex \widehat{X} with cones over copies of Y_i 's. The complex \widehat{X} is a covering space of X . A *combinatorial geodesic* in \widetilde{X}^* is a combinatorial geodesic in \widehat{X} , viewed as a path in \widetilde{X}^* .

2.3 | Disc diagrams in X^*

Throughout this paper, we will be analysing properties of *disc diagrams*, which we introduce below together with some associated terminology:

A map $f : X \longrightarrow Y$ between 2-complexes is *combinatorial* if it maps cells to cells of the same dimension. A complex is *combinatorial* if all attaching maps are combinatorial, possibly after subdividing the cells.

A *disc diagram* is a compact, contractible 2-complex D with a fixed planar embedding $D \subseteq \mathbb{S}^2$. The embedding $D \hookrightarrow \mathbb{S}^2$ induces a cell structure on \mathbb{S}^2 , consisting of the 2-cells of D together with an additional 2-cell, which is the 2-cell at infinity when viewing \mathbb{S}^2 as the one point compactification of \mathbb{R}^2 . The *boundary path* ∂D of D is the attaching map of the 2-cell at infinity. Similarly, an *annular diagram* is a compact 2-complex A with a fixed planar embedding $A \subseteq \mathbb{S}^2$ and the homotopy type of \mathbb{S}^1 . The annular diagram A has two boundary cycles $\partial_{\text{in}} A, \partial_{\text{out}} A$. A *disc diagram* in X^* is a combinatorial map $(D, \partial D) \rightarrow (X^*, X^{(1)})$ of a disc diagram. The 2-cells of a disc diagram D in X^* are of two kinds: squares mapping onto squares of X , and triangles mapping onto cones over

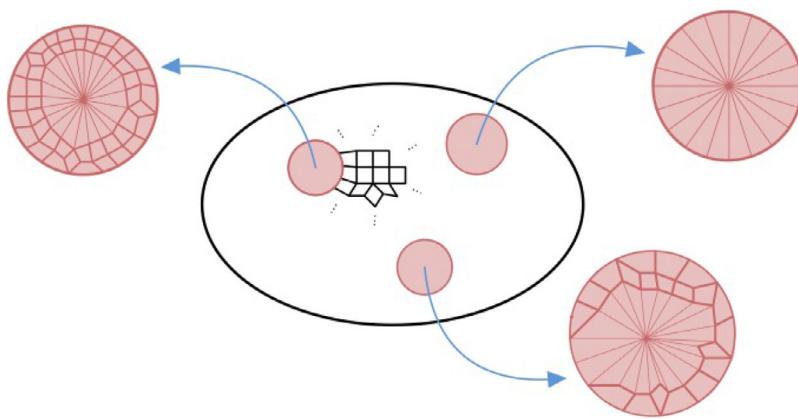


FIGURE 1 Cone-cells in a disc diagram. In figures, we will often omit the cell structure of cone-cells, unless needed.

edges contained in Y_i . The vertices in D that are mapped to the cone-vertices of X^* are also called the *cone-vertices*. Triangles in D are grouped into cyclic families meeting around a cone-vertex. We refer to such families as *cones*, and treat a whole such family as a single 2-cell. A *cone-cell* C is the union of an annular square diagram $A \rightarrow D$ whose interior embeds in D , together with a cone over $\partial_{in} A$. See Figure 1.

We emphasise that this definition differs slightly from the definition of a cone-cell in the literature, where A is simply a circle; allowing A to be an arbitrary annular diagram, D implicitly comes equipped with a *choice* of cone-cells.

The *square part* D_{\square} of D is a subdiagram which is the union of all the squares that are not contained in cone-cells.

A *square disc diagram* is a disc diagram whose square part is the whole diagram, that is, it contains no cone-cells. A *mid-interval* in a square, viewed as $[0, 1] \times [0, 1]$, is an interval $\{\frac{1}{2}\} \times [0, 1]$ or $[0, 1] \times \{\frac{1}{2}\}$. A *dual curve* in a square disc diagram D is a curve that intersect each closed square either trivially, or along a mid-interval, that is, a dual curve is a restriction of a hyperplane in X to D . We note that for each 1-cube of D , there exists a unique dual curve crossing it [16, 2e].

The *complexity* of a disc diagram D in X^* is defined as

$$\text{Comp}(D) = (\#\text{cone-cells}, \#\text{squares in } D_{\square}).$$

We say that D has *minimal complexity* if $\text{Comp}(D)$ is minimal in the lexicographical order among disc diagrams with the same boundary path as D . A disc diagram D in X^* is *degenerate* if $\text{Comp}(D) = (0, 0)$. A disc diagram D , in X^* is *singular* if D is not homeomorphic to a closed ball in \mathbb{R}^2 . This is equivalent to D either being a single vertex or an edge, or containing a cut vertex. In particular, every degenerate disc diagram is singular.

A square s is a *cornsquare* on a cone-cell C if a pair of dual curves emanating from consecutive edges a, b of c terminates on consecutive edges a', b' of ∂D .

Definition 2.1 (Reduction moves). We define six types of reduction moves. See Figure 2.

- (0) Cancelling a pair of squares s, s' meeting at one edge e in the disc diagram, whose map to X^* factors through a reflection identifying them. That is, cutting out $e \cup \text{Int}(s) \cup \text{Int}(s')$ and then gluing together the paths $\partial s - e$ and $\partial s' - e$.

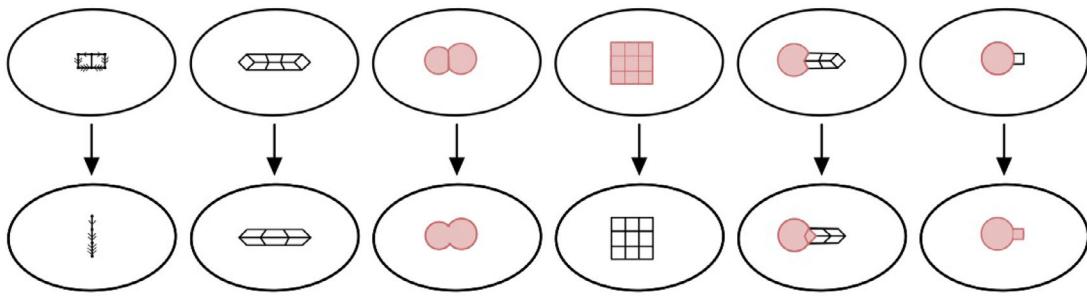


FIGURE 2 The six reduction moves from Definition 2.1.

- (1) Replacing a minimal bigon-diagram, that is, a disc subdiagram containing two dual curves intersecting each other twice, which is not contained in any other such subdiagram, with a lower complexity square disc diagram with the same boundary.
- (2) Replacing a pair of adjacent cone-cells with a single cone-cell.
- (3) Replacing a cone-cell with a square disc diagram with the same boundary.
- (4) Absorbing a cornsquare s to a cone-cell C , that is, replace a minimal subdiagram containing C and the two dual curves starting at C and ending in s with a lower complexity disc diagram with the same boundary and containing a cone-cell $C \cup s'$ for some square s' .
- (5) Absorbing a square with a single edge in a cone-cell into the cone-cell.

Definition 2.2 (Reduced and weakly reduced disc diagram). A disc diagram $D \rightarrow X^*$ in a cubical presentation is

- *reduced* if no moves (0)–(5) from Definition 2.1 can be performed in D ;
- *weakly reduced* if no moves (1)–(5) from Definition 2.1 can be performed in D .

Note that if D has minimal complexity then D is reduced, and that, in particular, each reduction move outputs a diagram D' with $\text{Area}(D') < \text{Area}(D)$ and $\partial D' = \partial D$. Consequently:

Lemma 2.3. *Let $D \rightarrow X^*$ be a disc diagram, then there exist disc diagrams $D' \rightarrow X^*$ and $D'' \rightarrow X^*$ satisfying:*

- (1) $\partial D = \partial D' = \partial D''$,
- (2) D' is weakly reduced and D'' is reduced,
- (3) D' is obtained from D after a finite number of moves of types (1)–(5), and D'' is obtained from D after a finite number of moves of type (0)–(5).

Remark 2.4. Many theorems about disc diagrams in the literature assume that the disc diagram is reduced or minimal complexity, but it is in fact sufficient to consider weakly reduced diagrams. For example, this is the case with Lemma 2.7 (the Cubical Greendlinger's Lemma).

2.4 | Cubical small-cancellation

We use the convention where $\bar{\rho}$ denotes the path ρ with the opposite orientation. A *grid* is a square disc diagram isometric to the product of two intervals. Let ρ and η be two combinatorial paths in

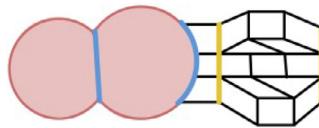


FIGURE 3 Blue paths are contiguous pieces, and yellow paths are pieces but not contiguous pieces.

\widetilde{X}^* . We say ρ and σ are *parallel* if there exists a grid $E \rightarrow \widetilde{X}^*$ with $\partial E = \mu\rho\overline{\nu}\overline{\eta}$, where the dual curves dual to edges of ρ , ordered with respect to its orientation, are also dual to edges of η , ordered with respect to its orientation. Concretely, if $\rho = e_1 \cdots e_k$ and $\sigma = f_1 \cdots f_k$, and $h(e_i)$ and $h(f_i)$ are the curves dual to e_i and f_i , respectively, then ρ is a piece if $h(e_i) = h(f_i)$ for each $i \in \{1, \dots, k\}$.

An *abstract contiguous cone-piece* of X^* in Y_i is a component of $\widetilde{Y}_i \cap \widetilde{Y}_j$, where \widetilde{Y}_i is a fixed elevation of Y_i to the universal cover \widetilde{X} , and either $i \neq j$ or where $i = j$ but $\widetilde{Y}_j \neq \widetilde{Y}_i$. Each abstract contiguous cone-piece P induces a map $P \rightarrow Y_i$ which is the composition $P \hookrightarrow \widetilde{Y}_i \rightarrow Y_i$, and a *contiguous cone-piece* of Y_j in Y_i is a combinatorial path $\rho \rightarrow P$ in an abstract contiguous cone-piece of Y_j in Y_i . An *abstract contiguous wall-piece* of X^* in Y_i is a component of $\widetilde{Y}_i \cap N(\widetilde{H})$, where \widetilde{H} is a hyperplane that is disjoint from \widetilde{Y}_i . Each abstract contiguous wall-piece P induces a map $P \rightarrow Y_i$, and a *contiguous wall-piece* of Y_i is a combinatorial path $\rho \rightarrow P$ in an abstract contiguous wall-piece of Y_i . A *piece* is a path parallel to a contiguous cone-piece or wall-piece.

The difference between contiguous pieces and pieces is illustrated in Figure 3.

For an integer $p > 0$, we say X^* satisfies the $C(p)$ *small-cancellation* condition if no essential combinatorial closed path in Y_i can be expressed as a concatenation of less than p pieces. For a constant $\alpha > 0$, we say X^* satisfies the $C'(\alpha)$ *small-cancellation* condition if $\text{diam}(P) < \alpha \|Y_i\|$ for every piece P involving Y_i .

Note that the $C'(\frac{1}{p})$ condition implies the $C(p+1)$ condition. When $p \geq 9$ and X^* is $C(p)$, then each immersion $Y_i \looparrowright X$ lifts to an embedding $Y_i \hookrightarrow \widetilde{X}^*$. This is proven in [16, Theorem 4.1] for $p \geq 12$, and in [8] for $p \geq 9$.

We record the following observation, a proof of which can be found in [2].

Lemma 2.5. *Let $X^* = \langle X \mid Y_1, \dots, Y_m \rangle$ be a cubical presentation where X and Y_1, \dots, Y_m are compact non-positively curved cube complexes. If X^* satisfies the cubical $C(p)$ condition for $p \geq 2$, then there is a bound on the combinatorial length of pieces of X^* .*

2.5 | Greendlinger's Lemma

A cone-cell C in a disc diagram D is a *boundary cone-cell* if C intersect the boundary ∂D along at least one edge. A non-disconnecting boundary cone-cell C is a *shell of degree k* if $\partial C = RQ$ where Q is the maximal subpath of ∂C contained in ∂D , and k is the minimal number such that R can be expressed as a concatenation of k pieces. We refer to R as the *innerpath* of C and Q as the *outerpath* of C .

A *corner* in a disc diagram D is a vertex v in ∂D of valence 2 in D that is contained in some square of D . A *cornsquare* is a square c and a pair of dual curves emanating from consecutive edges a, b of c that terminate on consecutive edges a', b' of ∂D . We abuse the notation and refer to the common vertex of a', b' as a cornsquare as well. A *spur* is a vertex in ∂D of valence 1 in D . If D contains a spur or a cut-vertex, then D is *singular*.

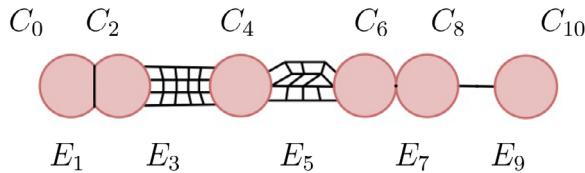


FIGURE 4 Example of a ladder.

Definition 2.6 (Ladder). A *pseudo-grid* between paths μ and ν is a square disc diagram E where the boundary path ∂E is a concatenation $\mu\rho\bar{\nu}\eta$ such that

- (1) each dual curve starting on μ ends on ν , and vice versa,
- (2) no pair of dual curves starting on μ cross each other,
- (3) no pair of dual curves cross each other twice.

If a pseudo-grid E is degenerate then either $\mu = \nu$ or $\rho = \eta$.

A *ladder* is a disc diagram $(D, \partial D) \rightarrow (X^*, X^{(0)})$ which is an alternating union of cone-cells and/or vertices C_0, C_2, \dots, C_{2n} and (possibly degenerate) pseudo-grids $E_1, E_3, \dots, E_{2n-1}$, with $n \geq 0$, in the following sense:

- (1) the boundary path ∂D is a concatenation $\lambda_1\bar{\lambda}_2$ where the initial points of λ_1, λ_2 lie in C_0 , and the terminal points of λ_1, λ_2 lie in C_{2n} ,
- (2) $\lambda_1 = \alpha_0\rho_1\alpha_2 \dots \alpha_{2n-2}\rho_{2n-1}\alpha_{2n}$ and $\lambda_2 = \beta_0\eta_1\beta_2 \dots \beta_{2n-2}\eta_{2n-1}\beta_{2n}$,
- (3) the boundary path $\partial C_i = \nu_{i-1}\alpha_i\mu_{i+1}\beta_i$ for some ν_{i-1} and μ_{i+1} (where ν_{-1} and μ_{2n+1} are trivial), and
- (4) the boundary path $\partial E_i = \mu_i\rho_i\bar{\nu_i}\eta_i$.

See Figure 4.

Lemma 2.7 (Cubical Greendlinger's Lemma [8, 16]). *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the C(9) condition, and let $D \rightarrow X^*$ be a weakly reduced disc diagram. Then one of the following holds:*

- D is a ladder, or
- D has at least three shells of degree ≤ 4 and/or corners and/or spurs.

We note that our definition of ladder differs slightly from the definitions in [8, 16], so that a single cone-cell and a single vertex count as ladders here. Also, the statements in [8, 16] assume that the disc diagrams are reduced/minimal complexity, but the proofs work for weakly reduced disc diagrams.

3 | HYPERBOLIC BACKGROUND

We explain the convention we will follow. A pair (Y, d) is a *metric graph*, if there exists a graph Γ such that Y is the vertex set of Γ , and d is defined as follows. For each edge of Γ , we assign a positive number which is the *length* of that edge. The *length* of a simple path in Γ is the sum of the lengths of the edges in the path. A metric d on a set Y is a *graph metric*, if (Y, d) is a metric graph.

In this paper, all edges of metric graphs have one of two lengths: 1 or $\frac{1}{2}$.

3.1 | Thin bigon criterion for hyperbolicity

A *bigon* in a geodesic metric space Y is a pair of geodesic segments γ_1, γ_2 in Y with the same endpoints, that is, such that $\gamma_1(0) = \gamma_2(0)$ and $\gamma_1(\ell) = \gamma_2(\ell)$ where ℓ is the length of γ . A bigon γ_1, γ_2 is ε -thin if $d(\gamma_1(t), \gamma_2(t)) < \varepsilon$ for all $t \in (0, \ell)$. If we do not care about the specific value of ε , the above condition is equivalent to the condition that $\text{im } \gamma_1 \subseteq N_{\varepsilon'}(\text{im } \gamma_2)$ and $\text{im } \gamma_2 \subseteq N_{\varepsilon'}(\text{im } \gamma_1)$ for some $\varepsilon' > 0$. Indeed, suppose that for every $t \in (0, \ell)$ there exists $t' \in (0, \ell)$ such that $d(\gamma_1(t), \gamma_2(t')) < \varepsilon'$. Then $|t - t'| < \varepsilon$, as otherwise γ_1 and γ_2 are not geodesic segments. That implies that $d(\gamma_1(t), \gamma_2(t)) \leq d(\gamma_1(t), \gamma_2(t')) + d(\gamma_2(t'), \gamma_2(t)) < 2\varepsilon'$.

This generalises to paths γ_1, γ_2 whose endpoints are not necessarily the same. We say γ_1, γ_2 ε -fellow travel if $d(\gamma_1(t), \gamma_2(t)) < \varepsilon$ for all t .

The following is a hyperbolicity criterion for graphs, due to Papasoglu [13, Theorem 1.4] (see also [16, Proposition 4.6]).

Proposition 3.1 (Thin Bigon Criterion). *Let Y be a graph where all bigons are ε -thin for some $\varepsilon > 0$. Then there exists $\delta = \delta(\varepsilon)$ such that Y is δ -hyperbolic.*

Of course, the converse also holds.

4 | THE PIECE METRIC

Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation. As explained in Subsection 2.2, we write X^* to denote the complex X with cones over Y_i 's attached. In particular, X can be viewed as a subspace of X^* . The preimage of X in the universal cover \widetilde{X}^* of X^* is denoted by \widehat{X} . Note that \widehat{X} is a covering space of X . The preimage of the 0-skeleton of X in \widetilde{X}^* is also the 0-skeleton of \widehat{X} , so it is denoted by $\widehat{X}^{(0)}$.

Definition 4.1. The *piece length* of a combinatorial path γ in $\widehat{X}^{(0)}$ is the smallest n such that $\gamma = \nu_1 \dots \nu_n$ where each ν_k is a 1-cube or a piece. The *piece metric* d_p on $\widehat{X}^{(0)}$ is defined as $d_p(a, b) = n$ where n is the smallest piece length of a path from a to b .

We note that d_p is a graph metric when $\widehat{X}^{(0)}$ is viewed as the graph with all edges of length 1 obtained from the 1-skeleton $\widehat{X}^{(1)}$ of \widehat{X} by adding extra edges between vertices contained in a single piece. We will denote this graph by $(\widehat{X}^{(0)}, d_p)$.

A *piece decomposition* of a path γ is an expression $\gamma = \nu_1 \dots \nu_k$, where each ν_i is a piece or 1-cube. We make the following easy observation:

Lemma 4.2. *Let $\gamma, \gamma_1, \gamma_2$ be piece-metric geodesics in $\widehat{X}^{(0)}$ where $\gamma = \gamma_1 \gamma_2$. Then*

$$|\gamma_1|_p + |\gamma_2|_p - 1 \leq |\gamma|_p \leq |\gamma_1|_p + |\gamma_2|_p.$$

Proof. Any piece decomposition $\gamma = \nu_1 \dots \nu_k$ yields piece decompositions of both γ_1 and γ_2 , where at most one piece ν_i for $i \in \{1, \dots, k\}$ further decomposes into the concatenation of two pieces ν'_i, ν''_i , so $\gamma_1 = \nu_1 \dots \nu'_i$ and $\gamma_2 = \nu''_i \dots \nu_k$. Similarly, any two piece decompositions of γ_1 and γ_2 can be concatenated to obtain a piece decomposition of γ . \square

We now prove a few basic facts about the piece metric. First, it is quasi-isometric to the combinatorial metric under fairly weak hypotheses.

Proposition 4.3. *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ condition for $p \geq 2$, and where X, Y_1, \dots, Y_s are compact. Then $(\hat{X}^{(0)}, d_p)$ is quasi-isometric to $(\hat{X}^{(0)}, d)$ where d is the standard combinatorial metric. Moreover, there is a uniform bound on the d_p -diameters of cones.*

Proof. Indeed, $d_p(a, b) \leq d(a, b)$ for all $a, b \in \hat{X}^{(0)}$, and by Lemma 2.5 there is an upper bound M on the combinatorial length of pieces, so we also have that $d(a, b) \leq M d_p(a, b)$.

As there are only finitely many Y_i 's and each Y_i is compact, there must be an upper bound on the diameter of a simple essential curve in Y_i with respect to d and thus with respect to d_p , which implies the second statement. \square

Corollary 4.4. *Suppose that $\pi_1 X$ is hyperbolic. Let $D \rightarrow X^*$ be a square diagram with boundary $\partial D = \gamma \bar{\lambda}$ where γ is a d_p -geodesic, and λ is a d -geodesic. Then the bigon D is M -thin for a uniform constant M .*

Proof. First note that $D \rightarrow X^*$ is a square diagram in \hat{X} , but it also lifts to \tilde{X} . The metric d_p also lifts to $\tilde{X}^{(0)}$, and by Proposition 4.3 d, d_p are quasi-isometric on $\hat{X}^{(0)}$, and therefore on $\tilde{X}^{(0)}$. The statement then follows from the uniform bound on the Hausdorff distance between geodesics and quasi-geodesics in hyperbolic spaces. \square

We note that ladders are thin with respect to the piece metric.

Proposition 4.5. *Suppose that $\pi_1 X$ is hyperbolic. Let $D \rightarrow X^*$ be a ladder with boundary $\partial D = \lambda_1 \bar{\lambda}_2$ as in Definition 2.6 where each subpath of λ_i contained in a single pseudo-grid is a geodesic. Then the bigon λ_1, λ_2 is ϵ -thin with respect to d_p for a uniform constant $\epsilon > 0$ dependent only on X^* .*

Proof. We only show that $\lambda_1 \subseteq N_\epsilon(\lambda_2)$, as the argument for $\lambda_2 \subseteq N_\epsilon(\lambda_1)$ is analogous. Let $x \in \lambda_1$. We want to show that $d_p(x, \lambda_2) \leq \epsilon$. If x belongs to a cone-cell C , then by the definition of the ladder, λ_2 also intersects C , so $d_p(x, \lambda_2)$ is bounded by the piece-metric diameter of C , which is uniformly bounded by some constant ϵ_1 by Proposition 4.3.

Otherwise x lies in a pseudo-grid. Let ρ, η be subpaths of λ_1, λ_2 , respectively, contained in the pseudo-grid that contains x . The paths ρ, η are both combinatorial geodesics by the assumption. By Proposition 4.3 ρ, η start and end at a uniform distance, as they lie in the same cone-cell. By hyperbolicity of \tilde{X} , there exists $\epsilon_2 > 0$ such that ρ, η ϵ_2 -fellow travel. The conclusion follows with $\epsilon = \max\{\epsilon_1, \epsilon_2\}$. \square

In the proof of Theorem 5.1, we will use the following technical lemma.

Lemma 4.6. *Let X^* be a cubical presentation, and let $E \rightarrow X^*$ be a square diagram, with the induced metric d_p . Suppose that $\partial E = \ell Q \bar{r}$ where γ is a piece geodesic, no dual curve in E crosses $\ell Q \bar{r}$ twice, and each of ℓ, Q, r contains no cornsquares of E in its interior. Moreover, assume that $|Q|_p \geq 3$. Then $|\gamma|_p \geq |\ell|_p + |Q|_p + |r|_p - 3$.*

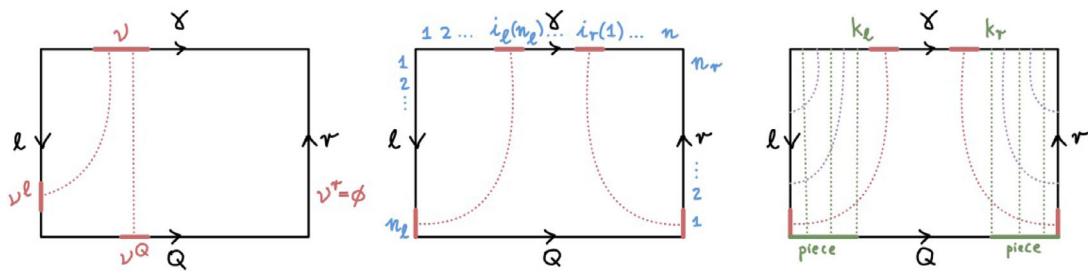


FIGURE 5 Steps of the proof of Lemma 4.6.

Proof. See Figure 5 for a diagram E with $\partial E = \ell Q r \bar{\gamma}$. By the assumptions, every dual curve of E starting at $\ell Q r$ must exit the diagram in γ . Thus, each edge of $\ell Q r$ is naturally paired with an edge of γ . For every piece ν in γ , we consider all the dual curves h_1, \dots, h_n starting at ν that exit E in $\ell Q r$.

These define a collection of edges in $\ell Q r$, and every subcollection of such consecutive edges forms a path that is a piece, as it is parallel to some path contained in one of Y_i . By grouping consecutive edges into maximal subpaths contained in one of ℓ , Q , or r , we get pieces ν^1, \dots, ν^k whose interiors are pairwise disjoint (ordered consistently with the orientation of $\ell Q r$), and say that ν projects to ν^1, \dots, ν^k .

First we claim that each of ℓ, Q, r contains at most one piece ν^i . Suppose to the contrary that ν^i, ν^{i+1} are both contained in ℓ (and the same argument applies to Q, r). Then each dual curve starting at an edge of ℓ lying between ν^i and ν^{i+1} must intersect at least one dual curve starting at edges of ν^i, ν^{i+1} , as otherwise it would also lie in a projection of ν , yielding a cornsquare in ℓ .

Thus, we can denote the projection of ν by ν^ℓ, ν^Q, ν^r where each piece is a possibly empty projection onto ℓ, Q, r , respectively. See left diagram in Figure 5. We will assume that they are oriented consistently with ℓ, Q, r , respectively, not necessarily consistently with ν .

Let $\gamma = \nu_1 \dots \nu_n$ be a minimal piece decomposition of γ , that is, $|\gamma|_p = n$. Let $\ell = \nu_{i_\ell(1)}^\ell \dots \nu_{i_\ell(n_\ell)}^\ell$ be the induced piece-decomposition where we only write non-trivial pieces. In particular, $i_\ell : \{1, \dots, n_\ell\} \rightarrow \{1, \dots, n\}$ is an injective function. We now claim that i_ℓ is monotone. Suppose to the contrary, that $1 \leq j < k \leq n_\ell$ but $i_\ell(j) > i_\ell(k)$. Then there must exist a cornsquare in the connected subpath of ℓ containing $\nu_{i_\ell(k)}^\ell$ and $\nu_{i_\ell(j)}^\ell$, which is a contradiction. Analogously, we get $Q = \nu_{i_Q(1)}^Q \dots \nu_{i_Q(n_Q)}^Q$ and $r = \nu_{i_r(1)}^r \dots \nu_{i_r(n_r)}^r$, and the functions i_Q, i_r are monotone. These are not necessarily the minimal piece decompositions, but certainly we have $|\ell|_p \leq n_\ell$, $|Q|_p \leq n_Q$ and $|r|_p \leq n_r$. To prove the lemma, we will show that $|Q|_p + n_\ell + n_r \leq n + 3$.

Note that $i_\ell(n_\ell)$ is the largest index in $\{1, \dots, n\}$ such that $\nu_{i_\ell(n_\ell)}$ has non-trivial projection onto ℓ , and similarly $i_r(1)$ is the lowest index in $\{1, \dots, n\}$ such that $\nu_{i_r(1)}$ has non-trivial projection onto r . See middle diagram in Figure 5. As i_ℓ, i_r are monotone, $n_\ell \leq i_\ell(n_\ell)$ and $n_r \leq n - i_r(1)$. Thus, it remains to prove that $|Q|_p \leq i_r(1) - i_\ell(n_\ell) + 3$.

Let k_ℓ be the largest number such that $i_Q(k_\ell) < i_\ell(n_\ell)$. We claim that $\nu_{i_Q(1)}^Q \dots \nu_{i_Q(k_\ell)}^Q$ is a single piece in Q . Indeed, the dual curves starting in $\nu_{i_Q(1)}^Q \dots \nu_{i_Q(k_\ell)}^Q$ must all intersect a dual curve starting in $\nu_{i_\ell(n_\ell)}$ and exiting the diagram in ℓ . See right diagram in Figure 5. Similarly, let k_r be the smallest number such that $i_Q(k_r) > i_r(1)$ and note that $\nu_{i_Q(k_r)}^Q \dots \nu_{i_Q(n_Q)}^Q$ is a single piece in Q . By assumption, $|Q|_p \geq 3$, so the subpath $\nu_{i_Q(k_\ell+1)}^Q \dots \nu_{i_Q(k_r-1)}^Q$ is non-empty. In particular, $k_r > k_\ell + 1$.

As by definition of k_ℓ, k_r we have $i_Q(k_\ell + 1) \geq i_\ell(n_\ell)$ and $i_Q(k_r - 1) \leq i_r(1)$, we conclude that

$$\begin{aligned} & |v_{i_Q(k_\ell+1)}^Q v_{i_Q(k_\ell+2)}^Q \cdots v_{i_Q(k_r-1)}^Q|_p \\ & \leq |v_{i_Q(k_\ell+1)} v_{(i_Q(k_\ell+1)+1)} \cdots v_{i_Q(k_r-1)}|_p \\ & \leq |v_{i_\ell(n_\ell)} v_{(i_\ell(n_\ell)+1)} \cdots v_{i_r(1)}|_p \\ & \leq i_r(1) - i_\ell(n_\ell) + 1. \end{aligned}$$

This proves that $|Q|_p \leq i_r(1) - i_\ell(n_\ell) + 3$ and completes the proof. \square

5 | PROOF OF HYPERBOLICITY

In the proof of the next theorem, we show that, under suitable assumptions, $(\hat{X}^{(0)}, d_p)$ is a δ -hyperbolic graph to deduce that $\pi_1 X^*$ is hyperbolic. The basic strategy is similar to [16, Theorem 4.7], but the details in this case are significantly more involved.

Theorem 5.1. *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ cubical small-cancellation condition for $p \geq 14$, where X, Y_1, \dots, Y_s are compact, and $\pi_1 X$ is hyperbolic. Then $\pi_1 X^*$ is hyperbolic.*

Before we proceed with the proof of the above theorem we introduce a construction that is used in the proof. Let $Y \subset X$. The *cubical convex hull* of Y in X is the smallest cubically convex subcomplex of X contained in Y . That is, it is the smallest subcomplex $Hull(Y)$ satisfying that whenever a corner of an n -cube c with $n \geq 2$ lies in $Hull(Y)$, then $c \subset Hull(Y)$.

Construction 5.2 (Square pushes). *Let D be a minimal complexity disc diagram, and let $\gamma\rho = \partial D$. Let λ be a path with the same endpoints as γ and lying in the cubical convex hull of γ , such that $\gamma\lambda$ bounds a disc subdiagram D_0 of D of maximal area. In particular, D_0 is a square disc diagram, and $D' = D - D_0$ is a disc diagram with $\partial D' = \lambda\rho$, which has no corners contained in the interior of the path λ . The diagram D_0 can be obtained via a finite sequence of square pushes, that is, a sequence of subdiagrams*

$$\gamma = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_{n-1} \subseteq K_n = D_0,$$

where for each $i = 0, \dots, n-1$ the subdiagram K_{i+1} contains K_i and an additional square s such that at least two consecutive edges of s are contained in K_i . Choosing a square s and adding it to K_i to obtain K_{i+1} will be referred to as pushing a square.

Note that the sequence of diagrams K_0, \dots, K_n is indeed finite, as $\text{Area}(K_{i+1}) = \text{Area}(K_i) + 1$ for each i , and thus $\text{Area}(D - K_{i+1}) = \text{Area}(D - K_i) - 1$, so $n-1 \leq \text{Area}(D)$.

By construction, every dual curve h in D_0 starting in γ' must exit in γ . Indeed, every square S that is being pushed has at least two consecutive edges on γ (in the first step) or on some K_i (in general). Thus, the 2 dual curves emanating from S either directly terminate on γ or enter K_i , crossing some of the previously added squares. By induction on the area of K_i , we can thus conclude that these dual curves terminate on γ .

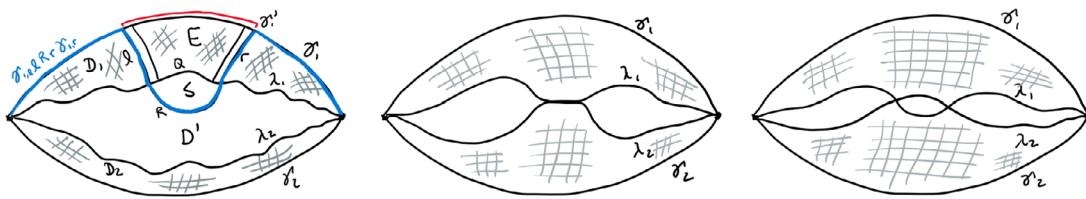


FIGURE 6 On the left, notation in the proof of Lemma 5.4; on the centre, possible overlapping paths in the proof of 5.1; on the right, the impossible transversal intersections described in the proof of 5.1.

Construction 5.3 (Sandwich decomposition of a bigon). *Let γ_1, γ_2 be paths forming a bigon. Let $D \rightarrow X^*$ be a reduced disc diagram with $\partial D = \gamma_1 \bar{\gamma}_2$. We define a decomposition of D into three (possibly singular) subdiagrams $D_1 \cup D' \cup D_2$ by applying Construction 5.2 twice as follows:*

- *we first apply it to the subpath $\gamma_1 \subseteq \partial D$ to obtain a decomposition $D = D_1 \cup D''$ where $\partial D_1 = \gamma_1 \bar{\lambda}_1$ and $\partial D'' = \lambda_1 \bar{\gamma}_2$;*
- *then we apply it to the subpath γ_2 of $\partial D''$ and we obtain a decomposition $D'' = D_2 \cup D'$ where $\partial D' = \lambda_1 \bar{\lambda}_2$ and $\partial D_2 = \lambda_2 \bar{\gamma}_2$.*

See Figure 6 for an example. We note that D_1, D_2 are square diagrams.

Lemma 5.4. *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ cubical small-cancellation condition for $p \geq 14$, where X, Y_1, \dots, Y_s are compact, and $\pi_1 X$ is hyperbolic.*

Then for any weakly reduced disc diagram $(D, \partial D) \rightarrow (X^, X)$ with $\partial D = \gamma_1 \bar{\gamma}_2$ where γ_1, γ_2 are d_p -geodesics and the subdiagram D' obtained from its sandwich decomposition $D = D_1 \cup D' \cup D_2$ is a ladder.*

The idea of the proof is as follows. Proceeding by contradiction, if D' is not a ladder, then D' contains a shell whose outerpath is disjoint from the endpoints q, q' of γ_1, γ_2 . Using this shell and Lemma 4.6, we construct a path γ with endpoints q and q' and with shorter piece-length than γ_1 , which contradicts the fact that γ_1 is a piece geodesic (see Figure 6).

Proof. Suppose to the contrary that D' is not a ladder. We will derive a contradiction with the fact that γ_1, γ_2 are d_p -geodesics. By Lemma 2.7, D' has at least three exposed cells, that is, shells of degree ≤ 4 , corners and/or spurs. Two of those exposed cells might contain q and q' , but there still must be at least one other exposed cell whose boundary path is disjoint from both q and q' . By construction of D' in Construction 5.3, there are no corners or spurs contained in the interior of the paths γ_1 and γ_2 , so we conclude that there must be a shell S of degree ≤ 4 in D' with the outerpath Q contained in γ_1 or γ_2 . Up to switching names of γ_1 and γ_2 , we can assume that Q is contained in γ_1 . Let R denote the innerpath of S in D' .

Let e_ℓ and e_r be the leftmost (first) and the rightmost (last) edge of R , and let h_ℓ, h_r be their dual curves in D_1 . By Construction 5.2 h_ℓ, h_r exit D_1 in γ_1 . Let γ'_1 be the minimal subpath of γ_1 that contains the edges dual to h_ℓ, h_r .

Let H_ℓ, H_r be the hyperplanes of \hat{X} extending h_ℓ, h_r , respectively. Let ℓ, r be combinatorial paths in D_1 parallel to h_ℓ, h_r and starting at the two endpoints of the path Q , respectively.

Consider a minimal complexity square disc diagram E with boundary $\partial E = \ell Q r \bar{\gamma}'_1$ where ℓ and r are combinatorial paths contained in $N(H_\ell), N(H_r)$. In particular, ℓ and r do not intersect H_ℓ and H_r , respectively. Such a diagram E exists because we can choose a subdiagram of D_1 . Among

all possible choices of ℓ, r and E we pick a diagram with minimal area. A feature of the choice of E is that it has no cornsquares in the interiors of ℓ and r , as otherwise we could push that cornsquare out and reduce the area. Up to possibly replacing Q with another path with the same endpoints contained in the same cone, we can assume that Q has no cornsquares either. We will assume that this is the case for the remainder of the proof.

We will be applying Lemma 4.6 to E , so we first verify that the assumptions are satisfied. By Lemma 2.7, $|Q|_p \geq p - 4 > 3$. Next, we claim that every dual curve starting in $\ell Q r$ exits E in γ'_1 . The cases of dual curves starting in ℓ and r are analogous, so we only explain the argument for ℓ . Consider the subdiagram $E' = E \cup S$ of D' . Let e be an edge in ℓ . Note that the dual curve h to e in E' cannot terminate on Q , as this would imply that there is a cornsquare on Q . If h terminates on r , then h is parallel to Q , and therefore Q is a single wall-piece, contradicting the $C(p)$ condition. Thus, h must terminate on γ'_1 . Let now e be an edge of Q , and h its dual curve in E . We already know that h cannot exit E in ℓ or r . If h exited E in Q , it would either yield a cornsquare in the interior of Q , contradicting the choice of Q , or it would yield a bigon formed from 2 squares glued along a pair of adjacent edges, contradicting the minimal complexity of E .

As no dual curve in E crosses $\ell Q r$ twice, there are no cornsquares in none of ℓ, Q and r , and $|Q|_p \geq 3$, Lemma 4.6 implies that $|Q|_p + |\ell|_p + |r|_p - 3 \leq |\gamma'_1|_p$. Recall that R is the innerpath of the shell S of degree ≤ 4 in D' . By definition $|R|_p \leq 4$, and so the $C(p)$ condition with $p \geq 14$ for X^* implies that $|Q|_p \geq p - 4 \geq 10 > |R|_p + 5$. Combining the two inequalities we get

$$|\bar{\ell} R r|_p \leq |R|_p + |\ell|_p + |r|_p < (|Q|_p - 5) + |\ell|_p + |r|_p \leq |\gamma'_1|_p - 2.$$

In particular, if we write $\gamma_1 = \gamma_{1,\ell} \gamma'_1 \gamma_{1,r}$, then using Lemma 4.2 we get

$$|\gamma_1|_p \geq |\gamma_{1,\ell}|_p + |\gamma'_1|_p + |\gamma_{1,r}|_p - 2 > |\gamma_{1,\ell}|_p + |\bar{\ell} R r|_p + |\gamma_{1,r}|_p \geq |\gamma_{1,\ell} \bar{\ell} R r \gamma_{1,r}|_p$$

which contradicts the fact that γ_1 is a piece-geodesic, completing the proof. \square

We now combine the previous ingredients to finish the proof of Theorem 5.1. Here is an outline of the proof. We consider a piece-geodesic bigon and want to prove that it is thin in the piece-metric. We apply the reduction moves from Definition 2.1 to obtain a new bigon. We then apply Lemma 5.4 to show that the middle layer of its sandwich decomposition is a ladder, and thus it is thin by Proposition 4.5. Finally, we use Corollary 4.4 to deduce that the other layers are also thin, and consequently the original bigon is thin.

Proof of Theorem 5.1. We prove that the coned-off space $(\hat{X}^{(0)}, d_p)$ is δ' -hyperbolic for some δ' by showing that it satisfies the bigon criterion (Proposition 3.1).

Let γ_1, γ_2 be d_p -geodesic segments forming a bigon. Pick combinatorial geodesics λ_1, λ_2 such that $\gamma_1 \bar{\lambda}_1$ and $\gamma_2 \bar{\lambda}_2$ bound square disc diagrams D_1 and D_2 , as they are path-homotopic in \hat{X} . Let $D' \rightarrow X^*$ be a reduced disc diagram with boundary $\partial D' = \lambda_1 \bar{\lambda}_2$. By gluing $D_1 \cup D' \cup D_2$ along λ_1 and λ_2 , respectively, we obtain a (possibly non-weakly reduced) disc diagram $D \rightarrow X^*$ with $\partial D = \gamma_1 \bar{\gamma}_2$.

By Lemma 2.3, there exists a sequence of reduction moves (1) - (5) from Definition 2.1 that turns D into a weakly reduced disc diagram. We describe how each reduction move transforms a quintuple $(D, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2)$ into a quintuple $(\bar{D}, \bar{\lambda}_1, \bar{\lambda}'_1, \bar{\lambda}_2, \bar{\lambda}'_2)$, both satisfying conditions below.

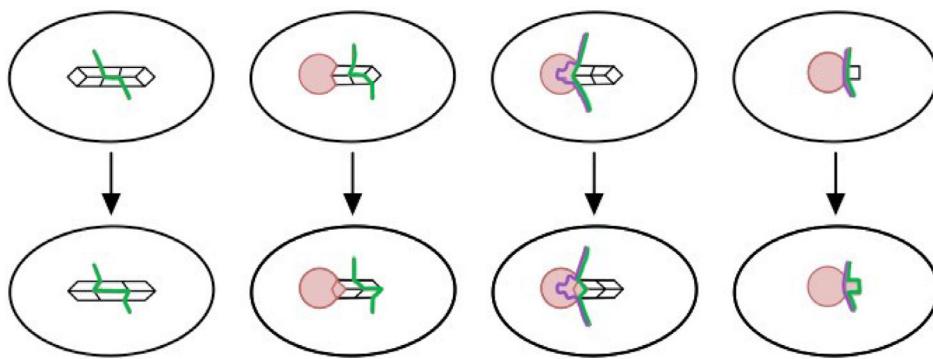


FIGURE 7 Reduction moves (1), (4), (4), (5). In the first two cases, the path $\lambda_i = \lambda'_i$ is reduced to the path $\ddot{\lambda}_i = \ddot{\lambda}'_i$. In the last two case, $\ddot{\lambda}'_i = \ddot{\lambda}'_i$ (purple paths), but $\ddot{\lambda}_i$ differs from λ_i .

Our sequence starts with $(D, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2) = (D, \lambda_1, \lambda_1, \lambda_2, \lambda_2)$. At each step $(D, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2)$ satisfies:

- (a) D is a disc diagram with $\partial D = \gamma_1 \bar{\gamma}_2$,
- (b) each λ_i is an embedded combinatorial path in D_{\square} , such that the three subdiagrams D_1, D_2 and D' of D with boundaries $\gamma_1 \bar{\lambda}_1, \lambda_2 \bar{\gamma}_2$ and $\lambda_1 \bar{\lambda}_2$, respectively, are reduced, and D_1, D_2 are square diagrams,
- (c) each λ'_i is an embedded combinatorial path in D (not necessarily in D_{\square}) such that λ'_i is a combinatorial geodesic in \tilde{X}^* , and $\lambda'_i \cap D_{\square} = \lambda_i \cap D_{\square}$,

and after applying a reduction move, we obtain a new quintuple $(\ddot{D}, \ddot{\lambda}_1, \ddot{\lambda}'_1, \ddot{\lambda}_2, \ddot{\lambda}'_2)$ satisfying the above conditions.

As the subdiagrams D_1, D_2 are square disc diagrams and D' is reduced, we never apply the reduction moves (3) and (2), as those would have to be performed within D' , contradicting that D' is reduced.

We now describe the transformation from λ_i to $\ddot{\lambda}_i$ and from λ'_i to $\ddot{\lambda}'_i$, for each reduction move. In each case the change will occur only within a subdiagram B that is transformed to \ddot{B} by the reduction. See Figure 7. We assume λ_i intersects the interior of B . We note that it might happen that both λ_1, λ_2 intersect B , in which case we apply the transformations to both λ_1, λ'_1 and λ_2, λ'_2 , according to the rules described below. It might happen that paths λ_1, λ_2 overlap, but at no step they intersect transversally, that is, at each step they yield a decomposition of the diagram D into $D_1 \cup D' \cup D_2$. See Figure 6.

Consider reduction move (1). Let B be a bigon-subdiagram of D , which is to be reduced. As this reduction move involves only squares, we have $\lambda_i \cap B = \lambda'_i \cap B$. Note that $\lambda_i \cap B$ cannot join the two corners of B because λ'_i is a combinatorial geodesic. Thus, $\lambda_i \cap B$ must be a combinatorial geodesic crossing both dual curves associated to B . For each $\lambda_i \cap B$ we set $\ddot{\lambda}_i \cap \ddot{B} = \ddot{\lambda}'_i \cap \ddot{B}$ to the combinatorial geodesic with the same endpoints in \ddot{B} and maximising the area of D_i . See the first diagram in Figure 7. Thus, the reduction move yields a new disc diagram \ddot{D} with new paths $\ddot{\lambda}_i, \ddot{\lambda}'_i$ for $i = 1, 2$ satisfying the required conditions. We note that in the case where both λ_1, λ_2 intersect B , the choices of $\ddot{\lambda}_1, \ddot{\lambda}_2$ ensure that $\ddot{\lambda}_1, \ddot{\lambda}_2$ do not intersect transversally.

Consider reduction move (4). Let B be a subdiagram associated to a cornsquare s and its dual curves ending on a cone-cell C . We set $\ddot{\lambda}_i \cap \ddot{B}$ to the combinatorial path with the same length and endpoints in \ddot{B} and maximising the area of D_i . If λ'_i coincides with λ_i in B we set $\ddot{\lambda}'_i \cap \ddot{B} = \ddot{\lambda}_i \cap \ddot{B}$.

See the second diagram in Figure 7. Otherwise, we set $\ddot{\lambda}'_i \cap \ddot{B} = \lambda'_i \cap B$. See the third diagram in Figure 7. Again, in the case where both λ_1, λ_2 intersect B , the choices of transformed paths ensure that no transversal intersection occurs.

Finally, consider reduction move (5). Let B be a subdiagram consisting of a square s overlapping with a cone-cell C along a single edge e . Thus, $\lambda_i \cap B = e$. We set $\ddot{\lambda}_i \cap \ddot{B}$ to the path $\partial s - e$. We also set $\ddot{\lambda}'_i \cap \ddot{B} = \lambda'_i \cap B$. See the last diagram in Figure 7.

Working under the assumption of weakly reduced: We now assume that $(D, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2)$ satisfies conditions (a)-(c) above and that D is weakly reduced. Following the notation in (b), we claim that either $D_1 \cup D' \cup D_2$ is the sandwich decomposition of D , or we can push squares into D_1 and D_2 modifying λ_i, λ'_i while preserving conditions (a)-(c). As $\lambda_i \cap D_{\square} = \lambda'_i \cap D_{\square}$ and λ'_i is a geodesic in \widehat{X}^* , no square s in D'_{\square} has three sides on λ_i . So, if a square s on λ_i can be pushed into D_i , then s must have two consecutive edges a, b on λ_i . Let $\lambda_i = \ell_i abr_i$ where ℓ_i, r_i are subpaths of $\lambda_i - ab$. Likewise, let $\lambda'_i = \ell'_i abr'_i$. Finally, define $\ddot{\lambda}_i = \ell_i cdr_i$ and $\ddot{\lambda}'_i = \ell'_i cdr'_i$ where c, d are the other two edges of s . The quintuple $(D, \ddot{\lambda}_1, \ddot{\lambda}'_1, \ddot{\lambda}_2, \ddot{\lambda}'_2)$ satisfies conditions (a)-(c). Indeed, as $|\lambda'_i| = |\ddot{\lambda}'_i|$, condition (c) is still satisfied. As this replacement does not affect D , nor the property of being (weakly) reduced, nor that D_1 and D_2 are square diagrams, conditions (a) and (b) are also preserved. We arrive at the sandwich decomposition after finitely many square pushes.

Working under the further assumption that $D_1 \cup D' \cup D_2$ is the sandwich decomposition of D : For the remainder of the proof, we assume that $(D, \lambda_1, \lambda'_1, \lambda_2, \lambda'_2)$ satisfies conditions (a)-(c) and that D is weakly reduced, and that the associated $D_1 \cup D' \cup D_2$ is the sandwich decomposition of D . By Lemma 5.4, the subdiagram D' with $\partial D' = \lambda_1 \bar{\lambda}_2$ is a ladder. Let D'' be the subdiagram of D' with $\partial D'' = \lambda'_1 \bar{\lambda}'_2$. Then D'' is also a ladder, as $\lambda_i \cap D' = \lambda'_i \cap D'$. By Proposition 4.5, the bigon λ'_1, λ'_2 is ϵ -thin for a uniform constant ϵ . By Corollary 4.4 there exists a uniform constant M such that γ_i, λ'_i is M -thin. Consequently, γ_1, γ_2 is $(\epsilon + 2M)$ -thin. \square

ACKNOWLEDGEMENTS

We thank the referee for helpful comments. The first author was supported by a Cambridge Trust & Newnham College Scholarship, and by the Denman Baynes Junior Research Fellowship. The second author was supported by the NSF Grants DMS-2203307 and DMS-2238198. The third author was supported by NSERC.

JOURNAL INFORMATION

The *Bulletin of the London Mathematical Society* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

ORCID

Macarena Arenas  <https://orcid.org/0000-0003-4965-8121>

REFERENCES

1. I. Agol, *The virtual Haken conjecture*, Doc. Math. **18** (2013), 1045–1087. (With an appendix by Agol, Daniel Groves, and Jason Manning.)
2. M. Arenas, *Asphericity of cubical presentations: The 2-dimensional case*, Int. Math. Res. Not. (2023), rna204. DOI [10.1093/imrn/rna204](https://doi.org/10.1093/imrn/rna204)

3. M. Arenas, *A cubical rips construction*, to appear in *Algebr. Geom. Topol.* (2023). <https://arxiv.org/pdf/2202.01048.pdf>
4. G. N. Arzhantseva and M. F. Hagen, *A cylindrically hyperbolicity of cubical small cancellation groups*, *Algebr. Geom. Topol.* **22** (2022), no. 5, 2007–2078.
5. M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*, Springer, Berlin, 1999.
6. D. Futer and D. T. Wise, *Cubulating random quotients of cubulated hyperbolic groups*, *Trans. Amer. Math. Soc. Ser. B* **11** (2024), 622–666.
7. J. Huang and D. T. Wise, *Virtual specialness of certain graphs of special cube complexes*, *Math. Ann.* **388** (2024), no. 1, 329–357.
8. K. Jankiewicz, *Lower bounds on cubical dimension of $C'(1/6)$ groups*, *Proc. Amer. Math. Soc.* **148** (2020), no. 8, 3293–3306.
9. K. Jankiewicz and D. T. Wise, *Cubulating small cancellation free products*, *Indiana Univ. Math. J.* **71** (2022), no. 4, 1397–1409.
10. I. J. Leary, *A metric Kan-Thurston theorem*, *J. Topol.* **6** (2013), no. 1, 251–284.
11. R. C. Lyndon, *On Dehn's algorithm*, *Math. Ann.* **166** (1966), 208–228.
12. R. C. Lyndon and P. E. Schupp, *Combinatorial group theory*, Springer, Berlin, 1977. *Ergebnisse der Mathematik und ihrer Grenzgebiete*, Band 89.
13. P. Papasoglu, *Strongly geodesically automatic groups are hyperbolic*, *Invent. Math.* **121** (1995), no. 2, 323–334.
14. P. Przytycki and D. T. Wise, *Mixed 3-manifolds are virtually special*, *J. Amer. Math. Soc.* **31** (2018), no. 2, 319–347.
15. M. Sageev, *Ends of group pairs and non-positively curved cube complexes*, *Proc. Lond. Math. Soc. (3)* **71** (1995), no. 3, 585–617.
16. D. T. Wise, *The structure of groups with a quasiconvex hierarchy*, *Annals of Mathematics Studies*, vol. 209, Princeton University Press, Princeton, NJ, 2021.