(RELATIVE) HYPERBOLICITY IN NON-METRIC CUBICAL SMALL CANCELLATION

MACARENA ARENAS, KASIA JANKIEWICZ, AND DANIEL T. WISE

ABSTRACT. We prove that the group defined by a cubical presentation $\langle X \mid Y_1, \ldots, Y_s \rangle$ satisfying sufficient non-metric cubical small cancellation conditions is (relatively) hyperbolic if the complex X is (relatively) hyperbolic. This can be thought of as a generalization of the fact that classical $C(7)$-small cancellation groups are hyperbolic.

1. INTRODUCTION

A cubical presentation, introduced in [Wis21], is a higher dimensional generalization of a standard group presentation in terms of generators and relators. A nonpositively curved cube complex X plays the role of the generator, and the "relators" are local isometries of nonpositively curved cube complexes $Y_i \hookrightarrow X$. The associated group is defined as the quotient of the fundamental group $\pi_1 X$ by the images of the $\pi_1 Y_i$. Just 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 is a generalization of classical small cancellation theory (for classical small cancellation see e.g. [LS77]), also introduced in [Wis21]. In both the classical and cubical cases, the small cancellation conditions are expressed in terms of *pieces*, and we consider a *metric small-cancellation condition* $C'(\alpha)$ where $\alpha \in (0,1)$ and *non-metric small cancellation condition* $C(p)$ where $p > 1$. It is a fundamental result of classical small cancellation theory that a group presentation satisfying the classical $C'(\frac{1}{6})$ of $C(7)$ condition is Gromov-hyperbolic.

Cubical small cancellation has proven to be a fruitful tool in the study of groups acting on CAT(0) cube complexes. It was used by Wise as a step in the proof of the Malnormal Special Quotient Theorem [Wis21]. Cubical presentations and cubical small cancellation theory were also studied in [Jan17, AH22, JW22, Are22].

In this note we prove that under suitable assumptions a cubical presentation satisfying the cubical non-metric small cancellation condition $C(p)$ for sufficiently large p defines a hyperbolic or relative-hyperbolic group.

Theorem 4.3. *Let $X^* = \langle X \mid Y_1, \ldots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ cubical small cancellation condition for $p \geq 9$, where $X, Y_1, \ldots, Y_s$ are compact, and $\widetilde{X}$ is δ-hyperbolic. Additionally, assume that there is a uniform bound on the length of pieces. Then $\pi_1 X^*$ is hyperbolic.*

The relatively hyperbolic version is as follows.

2010 *Mathematics Subject Classification.* 20F06, 20F67, 20F65.

Key words and phrases. Small Cancellation, Cube Complexes, Hyperbolic Groups, Relatively Hyperbolic Groups.

Theorem 4.5. *Let X be a compact nonpositively curved cube complex and $\{X_j \looparrowright X\}_{i=1}^r$ a collection of of local isometries of compact nonpositively curved cube complexes such that $\pi_1 X$ is hyperbolic relative to the collection $\{\pi_1 X_j\}_{i=1}^n$. Let $X^* = \langle X \mid Y_1, \ldots, Y_s \rangle$ be a cubical presentation satisfying the cubical small cancellation condition $C(p)$ for $p \geq 9$, where Y_i's are compact. Moreover, assume pieces have uniformly bounded length. Let $p : \pi_1 X \to \pi_1 X^*$ be the natural quotient. Then $\pi_1 X^*$ is hyperbolic relative to $\{p(\pi_1 X_j)\}_{i=1}^r$.*

We remark that a bound on the size of pieces is a naturally occurring hypothesis. If $X^* = \langle X \mid Y_1, \ldots, Y_s \rangle$ is a cubical presentation where $\pi_1 X$ is (relatively) hyperbolic, X is compact (sparse) and each Y_i is compact, then a bound on the size of cone-pieces exists when the collection $\{\pi_1 Y_1, \ldots, \pi_1 Y_r\}$ is malnormal, and a bound on the size of wall-pieces exists when each $\pi_1 Y_i$ is a (relatively) quasiconvex (full) subgroup of $\pi_1 X$. See for instance [Wis21, 2.40, 3.52, and 5.49] or [Wis12, 3.28 and 3.29].

The note is organized as follows. In Section 2, we give background on cube complexes, cubical group presentations, and cubical small cancellation. In Section 3, we recall criteria for hyperbolicity and relative hyperbolicity for groups acting on graphs. In Section 4, we prove the main results.

Acknowledgement. The first author was supported by a Cambridge Trust & Newnham College Scholarship, and by the Denman Baynes Senior Research Studentship. The second author was supported by the NSF grant DMS-2203307.

2. Cubical background

2.1. **Nonpositively 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 [BH99, Sag95, Lea13, Wis21]. A *nonpositively curved cube complex* is a cell-complex X whose universal cover $\widetilde{X}$ is a CAT(0) cube complex. A *hyperplane* $\widetilde{U}$ in $\widetilde{X}$ 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 $\widetilde{U}$ of $\widetilde{X}$, we let $N(\widetilde{U})$ denote its *carrier*, which is the union of all closed cubes intersecting $\widetilde{U}$. The *combinatorial metric* d on the 0-skeleton of a nonpositively 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 $\phi : Y \to X$ between nonpositively 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 concatenatable 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 and disc diagrams in X^*.**

We recall the definition of cubical presentation, and cubical small cancellation conditions from [Wis21].

Definition 2.1 (Cubical Presentations)**.** A *cubical presentation* $\langle X \mid Y_1, \ldots, Y_m \rangle$ consists of a nonpositively curved cube complex X, and a set of local isometries $Y_i \looparrowright X$ of nonpositively 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.

For example, a standard presentation complex associated with a group presentation $G = \langle S \mid R \rangle$ can be viewed as a cubical presentation, where the nonpositively curved cube complex X is just a wedge of circles, one corresponding to each generators 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.

A *disc diagram* is a compact, contractible 2-complex D with a fixed planar embedding $D \subseteq \mathbb{R}^2$. A *disc diagram in* X^* is a combinatorial map $(D, \partial D) \to (X^*, X^{(1)})$ 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 edges contained in Y_i. The vertices in D which 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 *cone-cells*, and treat a whole such family as a single 2-cell.

A disc diagram D in X^* is *minimal* if the complexity of D, defined as $\mathrm{Comp}(D) = (\#\text{cone-cells}, \#\text{squares})$ 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 $\mathrm{Comp}(D) = (0, 0)$. 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 disc diagram* is a disc diagram with 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 which intersect each closed square either trivially, or along a mid-interval, i.e. 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 [Wis21, 2e].

2.3. Cubical small cancellation and cubical Greendlinger's Lemma.

Definition 2.2 (Pieces)**.** 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 an elevation of Y_i to the universal cover $\widetilde{X^*}$, excluding the case where $i = j$. An *abstract contiguous wall-piece* of X^* in Y_i is a component of $\widetilde{Y}_i \cap N(\widetilde{U})$, where $\widetilde{U}$ is a hyperplane that is disjoint from $\widetilde{Y}_i$. An *abstract contiguous piece* is either an abstract contiguous cone-piece or an abstract contiguous wall-piece. A *contiguous piece* is a path in an abstract contiguous piece. A path q in $\widetilde{X^*}$ is a *piece* if there exists a disc diagram $D \to \widetilde{X^*}$ and a contiguous piece p in D such that the edges of q are dual to the same dual curves as the edges of p. Concretely, if $p = e_1 \cdots e_k$ and $q = f_1 \cdots f_k$, and $d_i(e_i)$ and $d_i(f_i)$ are the dual curves corresponding to e_i and f_i respectively, then q is a piece if $d_i(e_i) = d_i(f_i)$ for each $i \in \{1, \ldots, k\}$.

The difference between contiguous pieces and pieces is illustrated in Figure 1.

Definition 2.3 (Cubical small cancellation conditions)**.** 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 $\alpha > 0$,

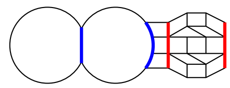

FIGURE 1. Blue paths are contiguous pieces, red paths are pieces but not contiguous pieces.

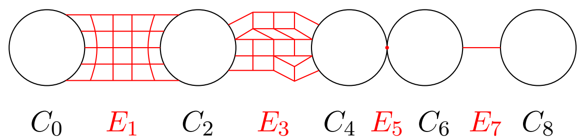

FIGURE 2. Example of a ladder.

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 [Wis21, Thm 4.1] for $p \geq 12$, and in [Jan17] for $p \geq 9$.

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 k is the minimal number such that the inner path P of C can be expressed as a concatenation of k pieces. 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 *spur* is a vertex in partial D of valence 1 in D. If D contains a spur, then D is singular.

Definition 2.4 (Ladder)**.** A *pseudo-grid* between paths μ and ν is a square disc diagram E where the boundary path ∂E is a concatenation $\mu\rho\nu^{-1}\eta^{-1}$ 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.

A *grid* is a pseudo-grid isometric to a rectangle with the obvious square disc diagram structure. If a pseudo-grid E is degenerate then either $\mu = \nu$ or $\rho = \eta$.

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

(1) the boundary path ∂D is a concatenation $P_1P_2^{-1}$ where the initial points of P_1, P_2 lie in C_0, and the terminal points of P_1, P_2 lie in C_{2n},
(2) $P_1 = \alpha_0\rho_1\alpha_2 \cdots \alpha_{2n-2}\rho_{2n-1}\alpha_{2n}$ and $P_2 = \beta_0\eta_1\beta_2 \cdots \beta_{2n-2}\eta_{2n-1}\beta_{2n}$,
(3) the boundary path $\partial C_i = \nu_{i-1}\alpha_i\mu_{i+1}^{-1}\beta_i^{-1}$ (where ν_{-1} and μ_{2n+1} are trivial), and
(4) the boundary path $\partial E_i = \mu_i\rho_i\nu_i^{-1}\eta_i^{-1}$.

See Figure 2.

Lemma 2.5 (Cubical Greendlinger's Lemma [Wis21, Jan17])**.** *Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the $C(9)$ condition, and let $D \to X^*$ be a 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 [Wis21, Jan17], so that a single cone-cell and a single vertex are ladders.

3. (Relatively) hyperbolic background

We explain the convention we will follow. A pair (Y, d) is a *metric graph*, if there exists a graph Γ such 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, γ' in Y with the same endpoints, i.e. such that $\gamma(0) = \gamma'(0)$ and $\gamma(\ell) = \gamma'(\ell)$ where ℓ is the length of γ. A bigon γ, γ' is *ε-thin* if $d(\gamma(t), \gamma'(t)) < \epsilon$ 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 $\operatorname{im}\gamma \subseteq N_{\varepsilon'}(\operatorname{im}\gamma')$ and $\operatorname{im}\gamma' \subseteq N_{\varepsilon'}(\operatorname{im}\gamma)$ for some $\varepsilon' > 0$. Indeed, suppose that for every $t \in (0, \ell)$ there exists $t' \in (0, \ell)$ such that $d(\gamma(t), \gamma'(t')) < \varepsilon'$. Then $|t - t'| < \varepsilon$, as otherwise γ and γ' are not geodesic segments. That implies that $d(\gamma(t), \gamma'(t)) \leq d(\gamma(t), \gamma'(t')) + d(\gamma'(t'), \gamma'(t)) < 2\varepsilon'$.

The following is a hyperbolicity criterion for graphs, due to Papasoglu [Pap95, Thm 1.4] (see also [Wis21, Prop 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, there is also a converse.

Proposition 3.2. *If a graph Y is δ-hyperbolic, then every bigon in Y is δ-thin.*

3.2. Relative hyperbolicity.

We recall the definition of a relatively hyperbolic group from [Bow12]. A graph X is *fine* if for every integer ℓ, each edge of X is contained in finitely many simple cycles of length ℓ in X. A subset $A \subseteq X$ is *locally finite*, if every bounded subset of A is finite. By [Bow12, Prop 2.1], a graph X is fine if and only if for each vertex $x \in X$, the set of vertices adjacent to x is locally finite in $X \setminus \{x\}$.

Definition 3.3. A group G is *hyperbolic relative to a collection of subgroups $\{H_i\}_i$* if G admits an action on a connected, δ-hyperbolic graph X where

(1) X is fine,
(2) there are finitely many G-orbits of edges, and each edge stabilizer is finite,
(3) the subgroups H_i are precisely the infinite vertex stabilizers of X,
(4) every H_i is finitely generated.

An equivalent definition of a group G hyperbolic relative to $\{H_i\}_i$ is in terms of a coned-off Cayley graph of G [Far98]. The *coned-off Cayley graph* $\hat{\Gamma}(G, \{H_i\}_i)$ of G with respect to $\{H_i\}_i$ is obtained from the Cayley graph $\Gamma(G)$ (with respect to some/any generating set of G) by adding extra vertex $v(gH_i)$ for each coset gH_i and for each element h of gH_i adding an edge of length $\frac{1}{2}$ between $v(gH_i)$ and h. The group G is hyperbolic relative to $\{H_i\}_i$ if and only if $\hat{\Gamma}(G, \{H_i\}_i)$ is δ-hyperbolic and fine.

4. Hyperbolicity of cubical small cancellation quotients

4.1. **Small cancellation quotients over a hyperbolic cube complex.** Let $X^* = \langle X \mid Y_1, \ldots, Y_s \rangle$ be a cubical presentation. Let us recall that we use the notation 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 $\widehat{X}$, and so it is denoted by $\widehat{X}^{(0)}$.

The *piece length* of a combinatorial path γ in $\widehat{X}^{(0)}$ is the smallest n such that $\gamma = \alpha_1 \cdots \alpha_n$ where each α_k is a 1-cube or a piece. The *piece metric* d_p on $\widehat{X}^{(0)}$ is defined as $\mathsf{d}_p(a,b) = n$ 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 a graph with all edges of length 1, obtained from the 1-skeleton $\widehat{X}^{(1)}$ of $\widehat{X}$ by adding extra edges between vertices at contained in a single piece. We will denote this graph by $(\widehat{X}^{(0)}, \mathsf{d}_p)$.

Remark 4.1. We point out the connection between the piece metric on $\widehat{X}^{(0)}$ introduced here and the notion of m-proximity studied in [Wis21, 5.l]: A hyperplane U is *m-proximate* to a 0-cube v if there is a path $P = P_1, \ldots, P_m$ such that each P_i is either a single edge or a piece, v is the initial vertex of P_1 and U is dual to an edge in P_m. So in the language of this note, a hyperplane U is m-proximate to a 0-cube v if there exists w in the carrier of U having $\mathsf{d}_p(v,w) = m-1$.

Proposition 4.2. *If there is a uniform bound on the length of pieces in X^*, then $(\widehat{X}^{(0)}, \mathsf{d}_p)$ is quasi-isometric to $(\widehat{X}^{(0)}, \mathsf{d})$ where d is the standard combinatorial metric.*

Proof. Indeed $\mathsf{d}_p(a,b) \le \mathsf{d}(a,b)$ for all $a, b \in \widehat{X}^{(0)}$, and since there is an upper bound M on the size of pieces, we also have that $\mathsf{d}(a,b) \le M\mathsf{d}_p(a,b)$. □

In the proof of the next theorem, we show that, under suitable assumptions, $(\widehat{X}^{(0)}, \mathsf{d}_p)$ is a δ-hyperbolic graph to deduce that $\pi_1 X^*$ is hyperbolic. The proof is inspired by the proof of [Wis21, Thm 4.7].

Theorem 4.3. *Let $X^* = \langle X \mid Y_1, \ldots, Y_s \rangle$ be a cubical presentation satisfying the $C(p)$ cubical small cancellation condition for $p \ge 9$, where $X, Y_1, \ldots, Y_s$ are compact, and $\widetilde{X}$ is δ-hyperbolic. Additionally, assume that there is a uniform bound on the length of pieces. Then $\pi_1 X^*$ is hyperbolic.*

Proof. We will verify that $(\widehat{X}^{(0)}, \mathsf{d}_p)$ satisfies the Thin Bigon Criterion (Proposition 3.1), and therefore is a δ-hyperbolic graph. Since $(\widehat{X}^{(0)}, \mathsf{d}_p)$ and $(\widehat{X}^{(0)}, \mathsf{d})$ are quasi-isometric

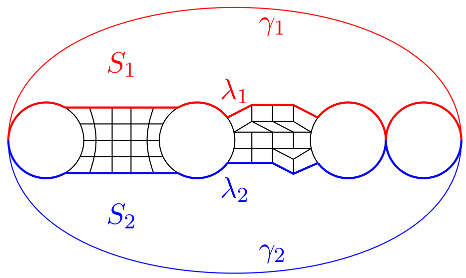

FIGURE 3. A decomposition of D into a union of S_1, S_2, D'.

by Proposition 4.2, and the action of $\pi_1 X^*$ on $(\widehat{X}^{(0)}, \mathsf{d})$ is proper and cocompact, then $\pi_1 X^*$ is hyperbolic.

We can also prove directly that $\pi_1 X^*$ acts properly and cocompactly on the graph $(\widehat{X}^{(0)}, \mathsf{d}_p)$. Indeed, the properness can be rephrased as finiteness of the vertex stabilizers, and follows directly from properness of the action of $\pi_1 X^*$ on $\widehat{X}^*$. To prove cocompactness of the action, it suffices to prove that the graph $(\widehat{X}^{(0)}, \mathsf{d}_p)$ has finitely many $\pi_1 X^*$-orbits of edges. The uniform bound on the size of pieces, and the compactness of X^* imply that the graph $(\widehat{X}^{(0)}, \mathsf{d}_p)$ is locally finite. That, together with the fact that there are only finitely many $\pi_1 X^*$-orbits of vertices in $\widehat{X}^{(0)}$, proves the claim.

Let γ_1, γ_2 form a bigon in $\widetilde{X}^{(0)}$, and let $D \to X^*$ be a minimal complexity disc diagram bounded by γ_1 and γ_2. Up to possibly replacing, γ_1 and γ_2 by their proper subpaths, we can assume that γ_1, γ_2 do not overlap on any edges. In particular, the diagram D is non-singular.

Amongst all possible choices of minimal complexity disc diagrams $D \to X^*$ bounded by γ_1 and γ_2, we moreover choose D so that the area of the square sub-diagrams bounded by $\gamma_i \lambda_i^{-1}$ is maximal amongst all possible choices of $S_1, S_2, \lambda_1, \lambda_2$, where the S_i and λ_i arise from a decomposition of D as a union of three subdiagrams $D' \cup S_1 \cup S_2$ satisfying that

- the interiors of D', S_1, S_2 are disjoint,
- for $i = 1, 2$, the diagrams S_i are the (possibly singular) square disk diagrams bounded by a bigon one of whose components is γ_i and whose other component we denote by λ_i, so the boundary path of S_i is $\gamma_i \lambda_i^{-1}$.

See Figure 3 for an example. Note that $\lambda_1 \lambda_2^{-1}$ is the boundary path of D'. By the maximality of S_1, S_2, and since λ_1 and λ_2 are geodesics in the piece metric, the diagram D' does not contain and corners and shells of degree ≤ 4, except possibly for the initial and terminal points of λ_1 and λ_2. Therefore by Lemma 2.5, D' is either a degenerate disc diagram, or a ladder.

If D' is degenerate, then D consists of squares only, and it can be viewed as a square disc diagram in $D \to \widetilde{X}$. The bigon (γ, γ') is δ-thin in the standard combinatorial metric on $\widetilde{X}^{(0)}$, where δ depends on the hyperbolicity constant of $\widetilde{X}$. Since $\mathsf{d}_p(a, b)$ is bounded above by the standard distance between a and b, we conclude that the bigon (γ, γ') is also δ-thin in the piece metric.

We note that, by compactness of Y_i's, there is a bound C on the diameter of each cone-cell in the piece metric. We show that if D' is a ladder, then for any $t \in (0, |\lambda_1|)$, we have $\mathsf{d}_p(\lambda_1(t), \operatorname{im} \lambda_2) \leq C$. Indeed, if $\lambda_1(t)$ belongs to some cone-cell R, then by the definition of the ladder $\operatorname{im} \lambda_2$ also intersects R, and the inequality holds. If $\lambda_1(t)$ lies in a pseudo-grid, then we actually have $\mathsf{d}_p(\lambda_1(t), \operatorname{im} \lambda_2) \leq 1$. Note that each pseudo-grid E_j in D' is actually a grid, since otherwise one could push a square in E_j to S_1 or S_2, hence increasing the area of S_1 or S_2 by 1, and contradicting the choices of D, S_1, S_2, λ_1 and λ_2. Thus $\lambda_1(t)$ lies in a single piece with some vertex of $\operatorname{im} \lambda_2$. Now note that (γ_i, λ_i) is a bigon in X, so we have already proven that it is δ-thin in the piece metric. We conclude that (γ_1, γ_2) is $(2\delta + C)$-thin.

□

4.2. Small cancellation quotients over a relatively hyperbolic cube complex.

Let X be a nonpositively curved cube complex and let $\{X_j \looparrowright X\}_{i=1}^r$ be a collection of local isometries of nonpositively curved cube complexes. Suppose that $\pi_1 X$ is hyperbolic relative to $\{\pi_1 X_j\}_{i=1}^n$.

Let $\widetilde{X}_\wedge$ denote the 0-skeleton of the *coned-off* universal cover $\widetilde{X}$ of X, with a cone on each elevation $\widetilde{X}_j$. In particular, $\widetilde{X}_\wedge$ has two types of vertices: the vertices of the cubical part $\widetilde{X}$, and the cone-vertices in cones over elevations of $\widetilde{X}_j$. The *shortcut metric* d_s is a graph metric on $\widetilde{X}_\wedge$, where the edges of $\widetilde{X}_\wedge$ are of two kinds:

- edges of length 1, corresponding to the 1-cubes of $\widetilde{X}$, and
- edges of length $\frac{1}{2}$, joining a vertex of an elevation $\widetilde{X}_j$ with the corresponding cone-vertex.

In other words, the length of a combinatorial path γ with both endpoints in $\widetilde{X}$ is the smallest number n where $\gamma = \alpha_1 \cdots \alpha_n$ and each α_k is either a 1-cube in $\widetilde{X}$ or a path entirely contained in an elevation $\widetilde{X}_j$ of X_j for some $i \in \{1, \dots, r\}$.

Proposition 4.4. *The graph* $(\widetilde{X}_\wedge, \mathsf{d}_s)$ *is hyperbolic.*

Proof. Let x be a vertex in $\widetilde{X}_\wedge$. Then the union V of the orbit $(\pi_1 X^*)x$ and the set of all cone-vertices in $\widetilde{X}_\wedge$ is a $\pi_1 X^*$-set, which is cofinite. But V can also be viewed as the vertex set of a coned-off Cayley graph of Γ, so in particular it is a hyperbolic $\pi_1 X^*$-set (in the sense of [Bow12, p. 26]). By [Bow12, Lem 4.12] $\widetilde{X}_\wedge$ is also a hyperbolic $\pi_1 X^*$-set, and so the graph $(\widetilde{X}_\wedge, \mathsf{d}_s)$ is hyperbolic (in the "standard" sense of Gromov). □

Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation, where X is as above. Consider the cubical part $\widehat{X}$ of the universal cover $\widetilde{X}^*$, which is invariant under $\pi_1 X^*$. Note that $\widehat{X} \to X$ is a a covering spaces. For every elevation $\widehat{X}_j$ of X_j to $\widehat{X}$, add a cone over $\widehat{X}_j$. Denote the 0-skeleton of the resulting space by $\widehat{X}_\wedge$. It consists of two types of vertices: those that belong to the cubical part $\widehat{X}$ and the cone-vertices of cones over X_j's. We emphasize that the cone-vertices of cones over lifts of Y_i's are not included in $\widehat{X}$.

The *piece-shortcut metric* d_{ps} is a graph metric on $\widehat{X}_\wedge$ where $\widehat{X}_\wedge$ is viewed as a graph with

- edges of length 1 corresponding to the 1-cubes of $\widehat{X}$, and
- edges of length $\frac{1}{2}$ joining a vertex of an elevation $\widetilde{X}_j$ with the appropriate cone-vertex.
- edges of length 1 joining pairs of vertices belonging to a piece in $\widetilde{X}^*$.

Again, the length of a combinatorial path γ with both endpoints in $\widehat{X}$ is the smallest n such that $\gamma = \alpha_1 \cdots \alpha_n$ where each α_k is a 1-cube, a piece, or a path factoring through some $\widetilde{X}_j$ for some $i \in \{1, \dots, r\}$.

In fact, d_{ps} is a graph metric, and $(\widehat{X}_\wedge, \mathsf{d}_{ps})$ can be simultaneously thought of as:

- $(\widehat{X}_\wedge, \mathsf{d}_s)$ with added extra edges, where (abusing the notation) d_s denotes the natural metric on $\widehat{X}_\wedge$ induced by the shortcut metric d_s on $\widetilde{X}_\wedge$, and
- $(\widehat{X}^{(0)}, \mathsf{d}_p)$ with added cone-vertices and adjacent edges, where d_p is the piece metric as defined in Section 4.1.

We now prove the relatively hyperbolic version of Theorem 4.3. The proof strategy is similar.

Theorem 4.5. *Let X be a compact nonpositively curved cube complex and $\{X_j \looparrowright X\}_{i=1}^r$ a collection of of local isometries of compact nonpositively curved cube complexes such that $\pi_1 X$ is hyperbolic relative to the collection $\{\pi_1 X_j\}_{i=1}^n$. Let $X^* = \langle X \mid Y_1, \dots, Y_s \rangle$ be a cubical presentation satisfying the cubical small cancellation condition $C(p)$ for $p \geq 9$, where Y_i's are compact. Moreover, assume pieces have uniformly bounded length. Let $p : \pi_1 X \to \pi_1 X^*$ be the natural quotient. Then $\pi_1 X^*$ is hyperbolic relative to $\{p(\pi_1 X_j)\}_{i=1}^r$.*

Proof. We will verify Definition 3.3 to prove that $\pi_1 X^*$ is hyperbolic relative to $\{p(\pi_1 X_j)\}$. The graph $(\widehat{X}_\wedge, \mathsf{d}_{ps})$ will play the role of the graph X in the definition. It is clearly a connected graph.

To verify that $\widehat{X}_\wedge$ is fine, we check the condition that for each vertex $x \in X$, the set of vertices adjacent to x is locally finite in $\widehat{X}_\wedge \setminus \{x\}$. If $x \in \widehat{X}$, then the set of vertices of adjacent to x is finite, since $\widehat{X}$ is locally finite, and x is adjacent to finitely many of cone-vertices. Now suppose that x is a cone-vertex. Then, by construction, the set of the vertices adjacent to x is contained in $\widehat{X}$, and so, is locally finite.

The group $\pi_1 X^*$ acts freely on the cubical part $\widehat{X}$, so the stabilizers of the vertices of $\widehat{X}$ are trivial. By construction the stabilizers of the cone-vertices of $\widehat{X}_\wedge$ are exactly $p(\pi_1 X_j)$. Since X_j's are compact, so $p(\pi_1 X_j)$ are finitely generated.

Next, we explain why $\widehat{X}_\wedge$ has finitely many $\pi_1 X^*$-orbits of edges. It has finitely many $\pi_1 X^*$-orbits of 1-cubes of $\widehat{X}$, since the action of $\pi_1 X^*$ on the cubical part $\widehat{X}$ is cocompact. All edges adjacent to a single cone-vertex are in the same $\pi_1 X^*$-orbit, and there are finitely many orbits of cone-vertices in $\widehat{X}_\wedge$, which proves that $\widehat{X}_\wedge$ has finitely many $\pi_1 X^*$-orbits of edges adjacent to some cone-vertex. By compactness of $X, Y_1, \dots, Y_s$ and the uniform bound on the size of pieces in X^*, there are finitely many $\pi_1 X^*$-orbits of edges of $\widehat{X}_\wedge$ corresponding to pieces.

It remains to prove that $(\widehat{X}_\wedge, \mathsf{d}_{ps})$ is a δ-hyperbolic graph. The inclusion $(\widehat{X}, \mathsf{d}_{ps|\widehat{X}}) \to (\widehat{X}_\wedge, \mathsf{d}_{ps})$ is a quasi-isometry, so it suffices to prove that $(\widehat{X}, \mathsf{d}_{ps|\widehat{X}})$ is δ-hyperbolic. We will abuse the notation and write d_{ps} for $\mathsf{d}_{ps|\widehat{X}}$.

Similarly as in the proof of Theorem 4.3 we will show that $(\widehat{X}, \mathsf{d}_{ps})$ satisfies the Thin Bigon Criterion (Propostion 3.1). Let γ_1, γ_2 form a bigon in $\widehat{X}$. We consider a minimal disc diagram $(D, \partial D) \to (X^*, X^{(0)})$ bounded by γ_1 and γ_2. As before, without loss of generality, by possibly passing to proper subpaths and a subdiagram, we assume that D is non-singular.

A priori, the diagram D has two types of cone-cells: those mapping into cones over $\widetilde{X}_j$'s, which we will refer to as $\widetilde{X}_j$*-type*, and those mapping into cones over Y_i's, which we will refer to as Y_i*-type*. However, minimality of D implies that D has no $\widetilde{X}_j$-type cone-cells. Indeed, since $\widetilde{X}_j$'s are simply connected, any $\widetilde{X}_j$-type cone-cell in D could be replaced by squares, reducing the complexity of D.

We additionally make a similar assumption on D as in the proof of Theorem 4.3, and the remaining part of this proof is similar to the proof of Theorem 4.3. Amongst all possible choices of minimal complexity disc diagrams $D \to X^*$ bounded by γ_1 and γ_2, we choose D so that the area of the square subdiagrams bounded by $\gamma_i \lambda_i^{-1}$ is maximal amongst all possible choices of $S_1, S_2, \lambda_1, \lambda_2$, where the S_i and λ_i arise from a decomposition of D as a union of three subdiagrams $D' \cup S_1 \cup S_2$ satisfying:

- the interiors of D', S_1, S_2 are disjoint,
- for $i = 1, 2$, the diagrams S_i are the (possibly singular) square disk diagrams bounded by a bigon one of whose components is γ_i and whose other component we denote by λ_i, so the boundary path of S_i is $\gamma_i \lambda_i^{-1}$.

By the maximality of S_1, S_2, and since λ_1 and λ_2 are geodesics in the piece metric, the diagram D' does not contain any corners and shells of degree ≤ 4, except possibly for the initial and terminal points of λ_1 and λ_2. Therefore by Lemma 2.5, D' is a degenerate disc diagram, or a ladder.

If D' is degenerate, then D can be viewed as a disc diagram in $\widetilde{X}$, and the bigon (γ_1, γ_2) is δ-thin in the shortcut metric d_s. Since the piece-shortcut metric d_{ps} is bounded above by d_s, the bigon (γ_1, γ_2) is also δ-thin in d_{ps}.

Now suppose that D' is a ladder. As in the proof of Theorem 4.3, for any $t \in (0, |\lambda_1|)$, we have $\mathsf{d}_p(\lambda_1(t), \operatorname{im} \lambda_2) \leq C$, where C is the uniform bound on the diameter of Y_j-type cone-cells in the piece metric. Each bigon (γ_i, λ_i) is δ-thin in the shortcut metric, so we conclude that the bigon (γ_1, γ_2) is $(2\delta + C)$-thin.

□

References

[AH22] Goulnara N. Arzhantseva and Mark F. Hagen. Acylindrical hyperbolicity of cubical small-cancellation groups, 2022. Alg. Geom. Top., to appear.

[Are22] Macarena Arenas. A cubical Rips construction, 2022.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature.* Springer-Verlag, Berlin, 1999.

[Bow12] B. H. Bowditch. Relatively hyperbolic groups. *Internat. J. Algebra Comput.*, 22(3):1250016, 66, 2012.

[Far98] B. Farb. Relatively hyperbolic groups. *Geom. Funct. Anal.*, 8(5):810–840, 1998.

[Jan17] Kasia Jankiewicz. The fundamental theorem of cubical small cancellation theory. *Trans. Amer. Math. Soc.*, 369(6):4311–4346, 2017.

[JW22] Kasia Jankiewicz and Daniel T. Wise. Cubulating small cancellation free products. *Indiana Univ. Math. J.*, pages 1–11, 2022. To appear.

[Lea13] Ian J. Leary. A metric Kan-Thurston theorem. *J. Topol.*, 6(1):251–284, 2013.

[LS77] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Springer-Verlag, Berlin, 1977. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89.

[Pap95] P. Papasoglu. Strongly geodesically automatic groups are hyperbolic. *Invent. Math.*, 121(2):323–334, 1995.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[Wis12] Daniel T. Wise. *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, volume 117 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2012.

[Wis21] Daniel T. Wise. *The structure of groups with a quasiconvex hierarchy*, volume 209 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2021.

Centre for Mathematical Sciences, University of Cambridge, United Kingdom CB3 0WB
Email address: mcr59@cam.ac.uk

Department of Mathematics, University of California, Santa Cruz, USA 95064
Email address: kasia@ucsc.edu

Dept. of Math. & Stats., McGill University, Montreal, Quebec, Canada H3A 0B9
Email address: wise@math.mcgill.ca