



High dimensional hyperbolic Coxeter groups that virtually fiber

Jean-François Lafont² · Barry Minemyer¹ · Gangotryi Sorcar³ ·
Matthew Stover⁴ · Joseph Wells⁵

Received: 26 February 2025 / Revised: 15 September 2025 / Accepted: 17 September 2025 /
Published online: 30 September 2025

© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2025

Abstract

This paper provides an iterative procedure for constructing hyperbolic Coxeter groups that virtually fiber over \mathbb{Z} that is flexible enough to yield infinitely many isomorphism classes in each virtual cohomological dimension (vcd) $n \geq 2$. Our procedure combines results of Jankiewicz, Norin, and Wise with a generalization of a construction due to Osajda involving a new simplicial thickening process. We also give a topological argument showing that the vcd of the right-angled Coxeter groups produced by our construction increases by exactly one with each iteration, guaranteeing that our process produces examples of every vcd.

1 Introduction

A group is said to *virtually algebraically fiber*, or just to *virtually fiber*, if it has a finite index subgroup that surjects onto \mathbb{Z} with finitely generated kernel. In a seminal paper,

✉ Barry Minemyer
bminemyer@commonwealthu.edu

Jean-François Lafont
jlafont@math.ohio-state.edu

Gangotryi Sorcar
gangotryi.sorcar@krea.edu.in

Matthew Stover
mstover@temple.edu

Joseph Wells
jwellsmath@protonmail.com

- 1 Commonwealth University of Pennsylvania, Bloomsburg, PA, USA
- 2 Ohio State University, Columbus, OH, USA
- 3 Krea University, Sricity, AP, India
- 4 Temple University, Philadelphia, PA, USA
- 5 Washington, DC, USA

Bestvina and Brady developed a combinatorial Morse Theory approach to determining whether certain groups virtually algebraically fiber [2]. More recently, Jankiewicz, Norin, and Wise gave a combinatorial criterion, based only on the existence of a certain state on the defining graph of a right-angled Coxeter group that forms a *legal orbit* under a certain set of moves, to certify that Bestvina–Brady Morse Theory applies to provide a virtual fibration [10]. This result motivated a large number of recent results in geometric group theory and hyperbolic geometry, including examples of higher-dimensional hyperbolic manifolds whose fundamental groups virtually fiber [1, 6, 8] and the first hyperbolic 5-manifold known to topologically fiber over the circle [9].

The purpose of this paper is to construct new examples of Gromov hyperbolic right-angled Coxeter groups that virtually fiber. Other than examples found in the above references, the only other infinite family of examples we know are due to Schesler and Zaremsky [14, Thm. 10.2]. As Schesler and Zaremsky remark after their proof, their examples all have cohomological dimension two. Our contribution is to construct infinitely many isomorphism classes in every cohomological dimension.

Main Theorem *For every $n \geq 2$, there exist infinitely many isomorphism classes of hyperbolic right-angled Coxeter groups that virtually algebraically fiber and have virtual cohomological dimension n .*

We now describe the construction. In [12], Osajda developed a recursive construction that builds hyperbolic right-angled Coxeter groups with arbitrarily large virtual cohomological dimension. In general, these Coxeter groups may not satisfy the hypotheses of [10, Thm. 6.14], and therefore the combinatorial systems of moves approach developed in that paper may not apply to provide a virtual algebraic fibration. We combine the recursive construction in [12] with a novel “ α -thickening” procedure generalizing the thickening procedure used by Osajda to construct hyperbolic right-angled Coxeter groups whose defining graphs admit a state with a legal orbit under an appropriate system of moves, so [10] then applies to provide a virtual fibration.

Given a finite cube complex X , our α -thickening procedure enlarges the vertex set by a surjection $\alpha: Y \rightarrow V(X)$ from a large finite set Y , inserts a high-dimensional simplex above each vertex of the original cube complex, and then applies the standard thickening procedure of [12]. This has the effect of inserting enough vertices and edges into the 1-skeleton of the complex so that one can develop a legal system of moves without changing the homotopy type of the resulting simplicial complex. This process may be of independent interest and is developed in Sect. 2.2 below.

We now describe the general strategy for producing our virtually fibered hyperbolic right-angled Coxeter groups in more detail. We work with cube complexes that are 5-large (see Sect. 2.2) and satisfy the conditions of having *no isolated corners* and *no disconnecting cubes*; see Sect. 3 for definitions.

1. Start with a finite 5-large cube complex X_n that has no isolated corners, no disconnecting cubes, and has (integral) cohomological dimension n .
2. Apply a carefully chosen α -thickening procedure to X_n to produce the simplicial complex $T_n := \text{Th}_\alpha(X_n)$. Since $\text{Th}_\alpha(X_n)$ is homotopy equivalent to X_n , the simplicial complex T_n also has cohomological dimension n .
3. Using properties of the α -thickening, we verify that the 1-skeleton of T_n is 5-large and admits a legal system of moves, and therefore the associated right-angled

- Coxeter group G_{n+1} is Gromov hyperbolic and virtually fibers. Generalizing and refining the Main Theorem in [12], we show that G_{n+1} has cohomological dimension $n + 1$. Let D_{n+1} denote the Davis complex for G_{n+1} .
4. Take a quotient of D_{n+1} by a torsion-free finite index subgroup of G_{n+1} of sufficiently high index, and call the resulting quotient cube complex X_{n+1} . We verify that X_{n+1} is 5-large, has no disconnecting cubes and no isolated corners, and has cohomological dimension $n + 1$.
 5. Return to Step 1, replacing X_n with X_{n+1} .

Remark 1.1 For each $k \geq 5$, one can start with a k -cycle for X_1 . Since α -thickening only increases the number of vertices, the resulting right-angled Coxeter group with vcd n has abelianization of order at least 2^k . Letting k go to infinity, one obtains examples in each vcd with abelianization of arbitrarily large order. In particular, one obtains infinitely many isomorphism classes with each fixed vcd.

Remark 1.2 The reader might naturally wonder whether probabilistic methods could be used to easily produce random graphs whose associated right-angled Coxeter groups satisfy the conclusions of our main theorem. The classical edge independent model $G(n, p)$ of random graphs considers graphs on n vertices, where each edge has probability p of being included. Davis and Kahle [5] showed that within the parameter range $n^{-1/k} \ll p \ll n^{-1/(k+1)}$ the right-angled Coxeter group associated to a randomly generated graph will almost surely have cohomological dimension k as n goes to infinity. However, within the parameter range $1/n \ll p \ll 1 - \frac{1}{n^2}$ the random graph will almost surely have associated right-angled Coxeter group that is not Gromov hyperbolic (see [3, Cor. 2.2]). Thus within the random model $G(n, p)$, the random graphs that will almost surely produce right-angled Coxeter groups of cohomological dimension at least 2 will almost surely correspond to right-angled Coxeter groups that are **not** Gromov hyperbolic.

2 Systems of moves and thickenings of cube complexes

In what follows we use some of the general theory of CAT(0) cube complexes. See [4, 13] for standard references.

2.1 Jankiewicz–Norin–Wise states and moves

The proof that our right-angled Coxeter groups virtually fiber relies on the *systems of moves* and *legal states* defined and developed by Jankiewicz et al. [10]. Here we quickly discuss their terminology and the main result that we will use in this paper. In what follows \mathbb{Z}_2 will denote the cyclic group of order two.

Let $\Gamma = \Gamma(V)$ be a finite simplicial graph with vertex set V . A *state* of Γ is a function $\varphi: V \rightarrow \mathbb{Z}_2$, which amounts to an assignment of 0 or 1 to every vertex in the graph. A state is said to be *legal* if the subgraphs of V induced by $\varphi^{-1}(0)$ and $\varphi^{-1}(1)$ are both connected and nonempty. A *move* at $v \in V$ is a state m_v with the following properties:

1. $m_v(v) = 1$
2. $m_v(u) = 0$ if u and v are adjacent in Γ .

A *system of moves* is a choice m_v for each $v \in V$. A move is typically not a legal state, except in the case where $m_v^{-1}(1) = \{v\}$ and v is not a cut point. Notice that it is allowable for two nonadjacent vertices to be assigned the same move, meaning that $m_v = m_w$ as functions from V to \mathbb{Z}_2 .

There is a natural bijection between the collection of states and the group \mathbb{Z}_2^V . This gives the collection of states the structure of an abelian group and, in particular, the subgroup $M \leq \mathbb{Z}_2^V$ generated by a system of moves is well-defined. A system of moves is called *legal* if there exists a legal state S such that every state in the set

$$M \cdot S = \{m + S : m \in M\}$$

is a legal state, and such an orbit $M \cdot S$ is called a *legal orbit*. Using this setup, Jankiewicz, Norin, and Wise applied Bestvina–Brady Morse Theory [2] to prove the following result.

Theorem 2.1 (Theorem 4.3 and Corollary 4.4 in [10]) *Suppose that Γ is a finite graph on which there exists a legal system of moves. Then the right-angled Coxeter group G with defining graph Γ virtually algebraically fibers.*

2.2 Simplicial thickenings of cube complexes

Let X denote a cube complex. Osajda [12] defines the *thickening* $\text{Th}(X)$ of X to be the simplicial complex whose vertices are the same as the vertices of X , and vertices v_1, \dots, v_k of $\text{Th}(X)$ span a simplex if and only if the vertices v_1, \dots, v_k are contained in a common cube of X . Osajda’s thickening is a special case of our more general α -thickening discussed below. To distinguish between the ideas, we refer to Osajda’s thickening using the notation $\text{Th}_1(X)$.

We recall some terminology related to simplicial complexes. A subcomplex $\Sigma' \subset \Sigma$ is *full* if every simplex of Σ whose vertices lie in Σ' is actually contained in Σ' . A simplicial complex Σ is *flag* if every $(k + 1)$ -tuple of pairwise adjacent vertices spans a k -simplex. Flag simplicial complexes are completely determined by their 1-skeleton. For $k \geq 5$, we say a simplicial complex Σ is *k-large* if it is flag, and any embedded cycle of length $< k$ fails to be full. Note that, for $k = 5$, this implies that for any embedded 4-cycle (v_1, v_2, v_3, v_4) , one of the diagonals must be present, i.e. either v_1 is joined to v_3 by an edge, or v_2 is joined to v_4 . The property of being 5-large is sometimes also called the “flag-no-square” condition.

A simplicial complex or cube complex is called *locally k-large* if the link of every vertex is a k -large simplicial complex. If X is a locally 5-large cubical complex, then $\text{Th}_1(X)$ is also a locally k -large simplicial complex [12, Lem. 3.2]. Since we will often have to keep track of whether or not vertices lie in a common cube, we will make use of the following:

Definition 2.2 Two distinct vertices v and w in a cube complex X have *cubical distance* 1 if and only if they are contained in a common cube. More generally, the *cubical*

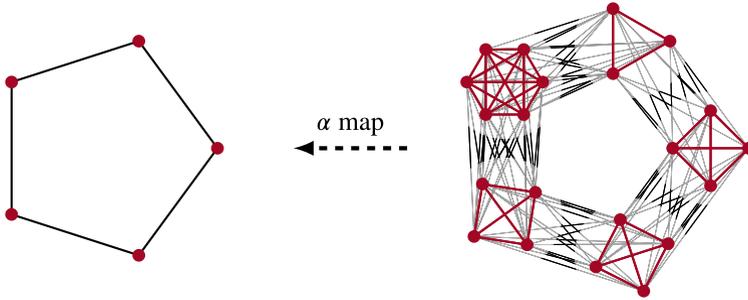


Fig. 1 A cube complex X on the left and a generic α -thickening $\text{Th}_\alpha(X)$ on the right. The α -thickenings can become extremely combinatorially complicated

distance $d_{\text{cube}}(v, w)$ between distinct vertices v and w is the minimum number n such that there is a sequence of vertices $v = v_0, \dots, v_n = w$ where $d_{\text{cube}}(v_i, v_{i+1}) = 1$ for each i .

Given a vertex v in a cube complex X , the cubical k -neighborhood of v is the subcomplex spanned by all vertices at cubical distance $\leq k$ from v . It will be convenient for our purposes to introduce the following terminology:

Definition 2.3 Let X be a cube complex. We say X is *5-large* if it is locally 5-large, and has the property that the cubical 2-neighborhood of every vertex is contractible.

Note that any locally 5-large cube complex is locally $\text{CAT}(0)$ by Gromov’s Lemma (see [4, App. I]). It follows that a simply connected locally 5-large cube complex is $\text{CAT}(0)$. In particular, if X is a 5-large cube complex, then the universal cover \tilde{X} is a $\text{CAT}(0)$ cube complex. The hypothesis on cubical 2-neighborhoods allows us to identify such neighborhoods in X with corresponding neighborhoods in the universal cover \tilde{X} .

In order to construct a legal system of moves as defined in Sect. 2.1, we will need a larger thickening than the one introduced by Osajda. The definition is as follows.

Definition 2.4 Let X be a finite cube complex with vertex set $V(X)$. Given a finite set Y equipped with a surjective map $\alpha : Y \rightarrow V(X)$, the α -thickening of X is the unique simplicial complex $\text{Th}_\alpha(X)$ with vertex set Y so that $y_0, \dots, y_k \in Y$ span a k -simplex of $\text{Th}_\alpha(X)$ if and only if the images $\alpha(y_0), \dots, \alpha(y_k)$ are contained in a common cube of X (Fig. 1).

Note that the vertex set Y is implicit in our notation $\text{Th}_\alpha(X)$ for the α -thickening. Two special cases of the thickening are noteworthy:

- If $X = \{x\}$ and $Y = \{y_0, \dots, y_n\}$, then $\alpha : Y \rightarrow X$ is the constant map and $\text{Th}_\alpha(X)$ is a single n -simplex.
- On any cube complex X , we can recover the thickening $\text{Th}_1(X)$ defined by Osajda [12], as it coincides with the α -thickening associated with the identity map $\alpha = \text{Id} : V(X) \rightarrow V(X)$.

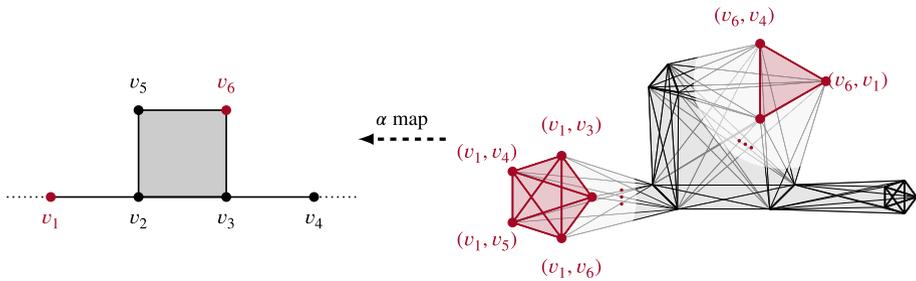


Fig. 2 A cube complex X on the left and the thickening $\mathbb{T}(X)$ on the right, highlighting the labeling of the preimages $\alpha^{-1}(v_1)$ and $\alpha^{-1}(v_6)$

We also note that there is an embedding of the Osajda thickening $\text{Th}_1(X) \hookrightarrow \text{Th}_\alpha(X)$ into the α -thickening associated with any section of the map $\alpha : Y \rightarrow V(X)$. Moreover, $\text{Th}_\alpha(X)$ deformation retracts to the image of $\text{Th}_1(X)$, and hence we have a homotopy equivalence $\text{Th}_\alpha(X) \simeq \text{Th}_1(X)$. Combining this with the homotopy equivalence $\text{Th}_1(X) \simeq X$ (see [12, Lem. 3.5]) we see that $\text{Th}_\alpha(X) \simeq X$. Next we show that the thickening inherits 5-largeness.

Lemma 2.5 *Suppose X is a 5-large cubical complex. Then for any surjective $\alpha : Y \rightarrow V(X)$, the simplicial complex $\text{Th}_\alpha(X)$ is 5-large.*

Proof Consider the universal cover \tilde{X} of X . Since \tilde{X} is a locally 5-large, simply connected cube complex, we have that $\text{Th}_1(\tilde{X})$ is 5-large; see [12, Prop. 3.4]. The covering map $\tilde{X} \rightarrow X$ naturally induces a covering map $\text{Th}_1(\tilde{X}) \rightarrow \text{Th}_1(X)$. Our hypothesis guarantees that cubical 2-neighborhoods in X lift to cubical 2-neighborhoods in \tilde{X} . It follows that the combinatorial 2-neighborhoods in $\text{Th}_1(X)$ lift isomorphically to combinatorial 2-neighborhoods in $\text{Th}_1(\tilde{X})$. However, the properties of being flag and 5-large are local: they only involve cycles of length at most 4. Any such cycle is contained in a combinatorial 2-neighborhood, so $\text{Th}_1(X)$ is also 5-large. For the general case where $\text{Th}_\alpha(X)$ is an arbitrary thickening, it is straightforward to use the embeddings $\text{Th}_1(X) \hookrightarrow \text{Th}_\alpha(X)$ to reduce to the $\text{Th}_1(X)$ case. We leave the details to the reader. □

3 Thickenings that produce legal systems of moves

This section contains the construction used to prove the Main Theorem. Throughout this section, X will denote a connected finite cube complex.

3.1 Thickening \mathbb{T} , states, and moves

We begin by defining a particular thickening $\mathbb{T}(X)$ of a cube complex X (shown in Fig. 2) that, when applied to a cube complex X with the appropriate properties, will admit a legal system of moves.

Definition 3.1 Let X be a cube complex for which every vertex is at cubical distance 2 from some other vertex. The thickening $\mathbb{T}(X)$ has vertices given by the following subset of $V(X) \times V(X)$:

$$V(\mathbb{T}(X)) = \{(v, w) : d_{\text{cube}}(v, w) \geq 2 \text{ in } X\}$$

The map $\alpha : V(\mathbb{T}(X)) \rightarrow V(X)$ is simply the projection onto the first factor.

In other words, $V(\mathbb{T}(X))$ consists of ordered pairs of vertices v, w of X that are not contained in a common cube. To simplify notation, we will denote the second coordinate as a subscript, so that the α -map just forgets the subscript, i.e., for the vertex $v_w \in V(\mathbb{T}(X))$ one has $\alpha(v_w) = v$. Also, if the vertex in the subscript is irrelevant, we will sometimes use the notation \bar{v} to denote a generic element of $\alpha^{-1}(v)$.

Next we define the initial state on $\mathbb{T}(X)$. We start by fixing once and for all a reference indexing $V(X) = \{x_1, \dots, x_k\}$ of the vertices of X . In terms of this reference indexing, we then assign

$$\mathcal{S}(v_w) = 1 \text{ if and only if } v = x_i, w = x_j \text{ and } i < j. \tag{1}$$

Suppose $v, w \in V(X)$ with $d_{\text{cube}}(v, w) \geq 2$, so there are associated vertices $v_w, w_v \in \mathbb{T}(X)$. The key property of our starting state \mathcal{S} is that it assigns opposite values to these two vertices. If $v = x_i, w = x_j$, and $i < j$, then $\mathcal{S}(v_w) = 1$ and $\mathcal{S}(w_v) = 0$. We now define the system of moves that, when paired with the state \mathcal{S} described in (1), produce a legal system on $\mathbb{T}(X)$.

Definition 3.2 Given any pair of vertices $v, w \in X$ at cubical distance ≥ 2 , we define the function $m_{vw} : V(\mathbb{T}(X)) \rightarrow \mathbb{Z}_2$ by setting $m_{vw}(v_w) = m_{vw}(w_v) = 1$ and m_{vw} to be identically zero on all other vertices.

Note that the vertices v_w and w_v are not adjacent in $\mathbb{T}(X)$, since they are at cubical distance at least 2 in X . This implies that each function m_{vw} is a move at the vertices v_w and w_v (and in fact, $m_{vw} = m_{wv}$). Informally, the move m_{vw} toggles the \mathbb{Z}_2 -assignment of both vertices v_w and w_v , and preserves the \mathbb{Z}_2 -assignment of all other vertices. Let M denote the subgroup generated by the system of moves.

Remark 3.3 Note that $m\mathcal{S}(v_w) \neq m\mathcal{S}(w_v)$ for any $m \in M$ and vertices $v, w \in V(X)$ with $d_{\text{cube}}(v, w) \geq 2$. Indeed, the initial state satisfies $\mathcal{S}(v_w) \neq \mathcal{S}(w_v)$, and every generator of M assigns the same value to v_w and w_v , and hence keeps the \mathbb{Z}_2 -assignment of v_w and w_v different.

3.2 Verifying the system is legal

The thickening $\mathbb{T}(X)$ can be applied to any finite cube complex X for which every vertex is at cubical distance 2 from some other vertex. We will now find suitable conditions on X to guarantee the initial state and moves on $\mathbb{T}(X)$ provide a legal system. The required assumptions on X are contained in the next two definitions.

Definition 3.4 A cube complex X has *no isolated corners* if, given any cube $\square \subset X$ and vertex $v \in \square$, there exists a vertex w adjacent to v and satisfying $w \notin \square$. X has *no disconnecting cubes* if, for every $\square \subset X$, the full subcomplex with vertex set $V(X) \setminus V(\square)$ is connected.

Our goal in this section will be to establish the following.

Proposition 3.5 *Let X be a 5-large cube complex. If X has no disconnecting cubes and no isolated corners, then the initial state S and system of moves $\{m_{vw}\}$ defined in this section form a legal system of moves on the 1-skeleton of $\mathbb{T}(X)$.*

The following definition and lemma will be used in the proof of this result.

Definition 3.6 Fix a cube complex X , a thickening $\text{Th}_\alpha(X)$, a state S on $V(\text{Th}_\alpha(X))$, and $\epsilon \in \mathbb{Z}_2$. We define the *detection function* corresponding to S by

$$\delta_\epsilon: V(X) \longrightarrow \{Y, N\}$$

as follows. If there exists $\bar{v} \in \alpha^{-1}(v)$ with $S(\bar{v}) = \epsilon$, then $\delta_\epsilon(v) = Y$. Otherwise $\delta_\epsilon(v) = N$.

The function δ_ϵ detects whether or not there is a vertex in the α -preimage of v that is assigned a value of ϵ by S . If such a vertex exists, then δ_ϵ assigns a value of Y for “yes”. Otherwise, it assigns N for “no”. Also note that, while we suppress this in the notation, the detection function δ_ϵ depends on the state S . In what follows the given state will always be clear from context.

Lemma 3.7 *Let X be a 5-large cube complex with no isolated corners and no disconnecting cubes, and let $\mathbb{T}(X)$ be the associated thickening. Fix $m \in M$ and $\epsilon \in \mathbb{Z}_2$. Let δ_ϵ denote the detection function corresponding to the state mS . Then there exists a cube \square_m in X such that $v \notin \square_m$ implies $\delta_\epsilon(v) = Y$.*

Note that the cube \square_m may not be unique. In fact, $\delta_\epsilon(v)$ could be Y for all $v \in V(X)$, in which case any cube in X vacuously satisfies the conclusion of the lemma. Further, it is possible that $\delta_\epsilon(v) = Y$ for some (or even all) of the vertices $v \in \square_m$.

Proof of Lemma 3.7 Suppose $v \in V(X)$ satisfies $\delta_\epsilon(v) = N$, and let $w \in V(X)$ be any vertex with $d_{\text{cube}}(v, w) \geq 2$. Then there are associated vertices v_w and w_v in $V(\mathbb{T}(X))$, and Remark 3.3 implies that $mS(v_w) \neq mS(w_v)$. Since $\delta_\epsilon(v) = N$, this means by definition that $mS(v_w) \neq \epsilon$, and therefore $mS(w_v) = \epsilon$ and so $\delta_\epsilon(w) = Y$. This shows that all vertices x with $\delta_\epsilon(x) = N$ must be at cubical distance 1 from each other.

It now remains to argue that all such vertices $\{v_1, \dots, v_k\}$ lie in a common cube in X . We know from Lemma 2.5 that $\mathbb{T}(X)$ is 5-large, and hence flag. By the above reasoning, the preimages of all the v_i form a collection of pairwise adjacent vertices in $\mathbb{T}(X)$. Since $\mathbb{T}(X)$ is flag, this collection of vertices span a simplex in $\mathbb{T}(X)$. By the definition of a thickening, this implies the vertices $\{v_1, \dots, v_k\}$ are all contained in a common cube of X . \square

With the Lemma in hand, we can now establish Proposition 3.5.

Proof Choose $m \in M$ and $\epsilon \in \mathbb{Z}_2$, and let \square_m be the cube provided by Lemma 3.7. Then fix $\bar{x}, \bar{y} \in V(\mathbb{T}(X))$ such that $m\mathcal{S}(\bar{x}) = m\mathcal{S}(\bar{y}) = \epsilon$. We must show that \bar{x} and \bar{y} are connected by a path within the subgraph induced by the set of vertices $(m\mathcal{S})^{-1}(\epsilon)$.

Set $x = \alpha(\bar{x})$ and $y = \alpha(\bar{y})$. We first assume that neither x nor y are contained in \square_m . Since X has no disconnecting cubes, there is a (cubical) path $x * v_1 * v_2 * \dots * v_k * y$ that does not meet \square_m . Then $\delta_\epsilon(v_i) = Y$ for each i , since all vertices with an output of N are contained in \square_m . Hence, for each i there exists $\bar{v}_i \in \alpha^{-1}(v_i)$ such that $m\mathcal{S}(\bar{v}_i) = \epsilon$. The path $\bar{x} * \bar{v}_1 * \dots * \bar{v}_k * \bar{y}$ thus connects \bar{x} to \bar{y} via vertices in $(m\mathcal{S})^{-1}(\epsilon)$.

Now, suppose that either x or y is in \square_m . If both x and y are in \square_m then the desired conclusion is clear, since all vertices in $\alpha^{-1}(x)$ and $\alpha^{-1}(y)$ are connected by an edge in $\mathbb{T}(X)$. Without loss of generality assume y is in \square_m and x is not. Since X has no isolated corners, there exists $z \in V(X)$ connected to y by an edge, with $z \notin \square_m$. The reasoning in the previous paragraph connects \bar{x} to some vertex $\bar{z} \in \alpha^{-1}(z)$ with a path contained in $(m\mathcal{S})^{-1}(\epsilon)$. We then append \bar{y} to the end of the path since it is connected by an edge to \bar{z} by assumption. □

Remark 3.8 As shown above, the assumptions that there are no disconnecting cubes and no isolated corners on X are sufficient for X to admit an α -thickening with a legal system of moves. For large cube complexes, one would expect that these conditions are also necessary by the following reasoning. Suppose X does not satisfy the no disconnecting cube condition, and let \square denote a maximal cube whose removal disconnects X . Then one can apply a sequence of moves making every vertex $v \in \square$ satisfy $\delta_\epsilon(v) = N$. In this way, \square would serve as an “ ϵ -blockade” between the two components of $X \setminus \square$, and thus the corresponding system of moves would not be legal. If X does not satisfy the no isolated corner condition, then any such corner vertex $v \in \square$ could be similarly blocked from the rest of the cube complex, by applying a sequence of moves that change the vertices above all the vertices in $\square \setminus v$. For small cube complexes (say, for example, three vertices connected as a line segment), these assumptions are not quite necessary, but one would expect these conditions to be necessary for larger complexes.

4 Proof of the main theorem

We now prove our main theorem following the inductive scheme outlined in the introduction. We start with a finite cube complex X_n having the following four properties:

- X_n is 5-large;
- X_n has no disconnecting cubes;
- X_n has no isolated corners;
- X_n has cohomological dimension n .

We start our process at the initial stage $n = 1$, by taking X_1 to be a cycle of length at least five. Note that all four of the above properties are then immediate.

Now consider the thickening $T_n := \mathbb{T}(X_n)$ of X_n defined in the preceding section. The 1-skeleton of the simplicial complex T_n defines a right-angled Coxeter group G_{n+1} . The proof now has the following three steps.

Step 1: Verify that the 1-skeleton of T_n has the sufficient graph theoretic properties for G_{n+1} to be Gromov hyperbolic and virtually fiber.

Step 2: Check that G_{n+1} has virtual cohomological dimension $n + 1$.

Step 3: If D_{n+1} is the Davis complex associated with the Coxeter group G_{n+1} , then verify that there is a suitable torsion-free finite index subgroup $\Lambda \leq G_{n+1}$ so that $X_{n+1} = D_{n+1}/\Lambda$ satisfies the four properties in our inductive scheme.

The first step follows immediately from the arguments in Sect. 3. Indeed, Lemma 2.5 shows that the simplicial complex T_n is 5-large, so the Coxeter group G_{n+1} will be Gromov hyperbolic. Moreover, Proposition 3.5 shows that the 1-skeleton of T_n has a legal system, so the Coxeter group G_{n+1} will virtually fiber by [10, Thm. 6.14]. This completes verification of the first step.

For Step 2, we want to compute the virtual cohomological dimension of the right-angled Coxeter group G_{n+1} associated with the complex T_n . A result of Davis (see [4, Cor. 8.5.5]) allows us to compute the virtual cohomological dimension from the topology of the simplicial complex T_n :

$$\text{vcd}(G_{n+1}) = \max \left\{ k : \overline{H}^{k-1}(T_n \setminus \sigma) \neq 0 \text{ for some } \sigma \subset T_n \right\}. \quad (2)$$

To calculate the cohomological dimension of spaces of the form $T_n \setminus \sigma$ where σ is a simplex of T_n , we will use a Mayer–Vietoris sequence. This will require certain auxiliary subsets of T_n that we now introduce.

Given a simplicial complex Σ , we can metrize Σ by making each simplex isometric to a spherical all right simplex of the appropriate dimension. Therefore each edge has length $\pi/2$, each triangle is isometric to the portion of the 2-sphere in the positive octant, and so forth. Given a simplex $\sigma \subset \Sigma$, consider the set $S(\sigma, \Sigma)$ of points at distance $\epsilon > 0$ from σ , thought of as a sphere centered on σ . Note that the topology of this space is independent of the choice of small positive $\epsilon > 0$. For example, if $v \in \Sigma$ is any vertex, then $S(v, \Sigma)$ is homeomorphic to the link $\text{lk}(v, \Sigma)$. However, for higher dimensional simplices, $S(\sigma, \Sigma)$ is a more complicated space built up from the links of σ as well as the links of the faces of σ . In the next section, we will establish the following:

Proposition 4.1 *For every nonempty simplex $\sigma \subset T_n$, we have that $\text{cd}(S(\sigma, T_n)) \leq n - 1$.*

Here the cohomological dimension $\text{cd}(Z)$ of a space Z is the largest integer k such that $\overline{H}^k(Z)$ is nontrivial, which is a slight abuse of terminology.¹ The proof of

¹ This is not the standard definition, as cohomological dimension is usually defined as the largest k so that $\overline{H}^k(Z, C) \neq 0$ for some closed subset $C \subset Z$. Throughout our paper we are working with the special case where $C = \emptyset$. Nevertheless, it will be convenient to use the notation $\text{cd}(Z)$, which could be smaller than the classically defined cohomological dimension of Z . We hope this will not cause confusion for the reader.

Proposition 4.1 will be postponed until the next section, as it is somewhat involved. Assuming this result, we can now continue our verification of the second step.

Proposition 4.2 *The group G_{n+1} has virtual cohomological dimension $= n + 1$.*

Proof Equation (2) implies the virtual cohomological dimension of G_{n+1} can be computed from the complex T_n . In particular, its vcd is one more than the maximal cohomological dimension $\text{cd}(T_n \setminus \sigma)$, where σ ranges over all simplices $\sigma \subset T_n$. In the special case where $\sigma = \emptyset$, we see that $\text{cd}(T_n) = \text{cd}(X_n) = n$, establishing the lower bound $\text{vcd}(G_{n+1}) \geq n + 1$.

To show that $\text{vcd}(G_{n+1}) = n + 1$ we must verify that $\text{cd}(T_n \setminus \sigma) \leq n$ for each nonempty simplex $\sigma \subset T_n$. Given such a simplex σ , cover T_n by two open sets U, V , by setting U to be an ϵ -neighborhood of σ , and V to be the complement $T_n \setminus \sigma$. The intersection $U \cap V$ is then homotopy equivalent to $S(\sigma, T_n)$. Since U is contractible, the Mayer–Vietoris sequence in reduced cohomology reduces to the long exact sequence:

$$\dots \longrightarrow \bar{H}^i(T_n) \longrightarrow \bar{H}^i(T_n \setminus \sigma) \longrightarrow \bar{H}^i(S(\sigma, T_n)) \longrightarrow \bar{H}^{i+1}(T_n) \longrightarrow \dots$$

Since $\text{cd}(T_n) = n$, the sequence gives us isomorphisms $\bar{H}^i(T_n \setminus \sigma) \cong \bar{H}^i(S(\sigma, T_n))$ in all degrees $i \geq n + 1$. However we know from Proposition 4.1 that $\text{cd}(S(\sigma, T_n)) \leq n - 1$, so the groups $\bar{H}^i(T_n \setminus \sigma)$ then vanish when $i \geq n + 1$. We conclude that $\text{cd}(T_n \setminus \sigma) \leq n$, as desired. This completes the proof of the proposition. \square

We now proceed to the third step, namely checking that the inductive hypotheses are satisfied for a suitable quotient X_{n+1} of the Davis complex D_{n+1} . We first prove that the analogous properties for the associated Davis complex are satisfied, then verify that they descend to a suitably chosen quotient of the Davis complex.

Lemma 4.3 *The Davis complex D_{n+1} is 5-large, has no isolated corners, has no disconnecting cubes, and the group G_{n+1} is one-ended.*

Proof Links of the vertices of D_{n+1} are copies of the 5-large simplicial complex T_n , so local 5-largeness is immediate. Moreover, cubical neighborhoods of vertices in CAT(0) cube complexes (such as D_{n+1}) are known to be convex (see [7, Lem. 13.15]) hence contractible. This verifies that D_{n+1} is 5-large. Recall that the 1-skeleton of the Davis complex is canonically identified with the Cayley graph of the Coxeter group with respect to its standard generators. To see that D_{n+1} has no isolated corners, just note that at an isolated corner, the link would have to be a single simplex. This is impossible, since the link is T_n , which is not contractible since $\text{cd}(T_n) \neq 0$. It remains to check that D_{n+1} has no disconnecting cubes, and it is sufficient to argue that G_{n+1} is one-ended. This can be determined from T_n by verifying that the complement $T_n \setminus \sigma$ is connected for every simplex $\sigma \subset T_n$; see [4, Thm. 8.7.2]. However, this is clear since T_n is a thickening of X_n and, by induction, X_n has no disconnecting cubes and no isolated corners. Noting that the simplex σ lies above a cube, the argument in the proof of Proposition 3.5 shows the complement of σ is connected. \square

Corollary 4.4 *For a suitable finite index torsion-free subgroup $\Lambda \leq G_{n+1}$, the finite cube complex $X_{n+1} := D_{n+1}/\Lambda$ is 5-large, has no disconnecting cubes, has no isolated corners, and has cohomological dimension $= n + 1$.*

Proof Concerning the cohomological dimension, we have the series of equalities

$$\text{cd}(X_{n+1}) = \text{cd}(\Lambda) = \text{vcd}(G_{n+1}) = n + 1.$$

The first equality follows from the fact that the Davis complex D_{n+1} is CAT(0) and thus is contractible, hence X_{n+1} is an Eilenberg–MacLane space $K(\Lambda, 1)$. Since Λ is a finite index torsion-free subgroup of G_{n+1} , the second equality is a well-known result of Serre [4, Thm. F.3.4]. The third equality follows directly from Proposition 4.2. This establishes the statement on the cohomological dimension.

Next we observe that the properties of being locally 5-large and having no isolated corners are *purely local*, in that they can be determined by looking at the cubical 2-neighborhood of vertices. From Lemma 4.3 we see that the local requirements are verified for cubical 2-neighborhoods of vertices in the Davis complex D_{n+1} . Therefore, if the 2-cubical neighborhoods in X_{n+1} lift to cubical 2-neighborhoods in the universal cover, we would obtain the desired properties for X_{n+1} . It is now a very standard consequence of residual finiteness of the Coxeter group G_{n+1} that there is a finite index subgroup Λ of G_{n+1} with this lifting property. Lastly, we need to verify that X_{n+1} has no disconnecting cubes. This is again standard, following from Lemma 4.3 since D_{n+1} has no disconnecting cubes and G_{n+1} is one-ended. \square

Remark 4.5 Note that the cube complex X_{n+1} has the property that there exists another vertex at cubical distance ≥ 2 from v for every vertex v . This follows from the fact that X_{n+1} is a quotient of the unbounded cube complex D_{n+1} and that cubical 2-neighborhoods in X_{n+1} lift to cubical 2-neighborhoods in D_{n+1} . This confirms that we can apply our thickening process to X_{n+1} (see Definition 3.1).

5 Calculating the virtual cohomological dimension

5.1 Topology of links, spheres, and balls

In this subsection we establish some useful results on links, spheres, and balls in simplicial complexes. The reader who is primarily interested in the proof of Proposition 4.1 can skip over this subsection, referring back to results as needed. We merely observe that the simplicial complexes T_n are 5-large (see Lemma 2.5), as are all the links in T_n . Therefore all results in this section apply to them.

Recall that $\text{lk}(\sigma, \Sigma)$ is an abstract simplicial complex with vertices corresponding to simplices containing σ as a codimension one face, and where a collection of vertices spans a simplex in $\text{lk}(\sigma, \Sigma)$ if and only if the corresponding simplices in Σ are contained in a common larger simplex in Σ . This abstract definition holds for general simplicial complexes. In the special case where Σ is a flag simplicial complex, the link of a simplex $\sigma \subset \Sigma$ can be identified with the full subcomplex $\text{lk}(\sigma, \Sigma)$ of Σ spanned by all vertices at combinatorial distance one from all vertices of σ . Note that the property of being 5-large and flag is inherited by full subcomplexes, and thus it is inherited by links.

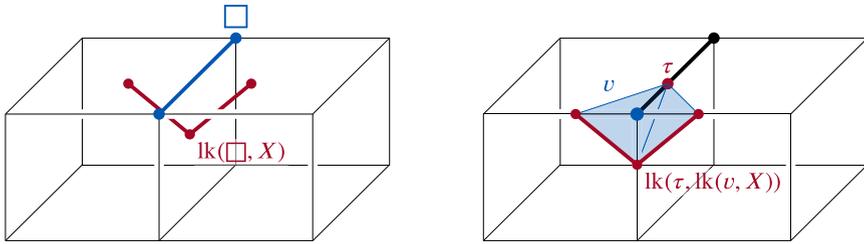


Fig. 3 The combinatorial isomorphism in Lemma 5.1

Similarly, if \square is a cube in the CAT(0) cube complex X , we can define the abstract *combinatorial link* of the cube. Vertices are the cubes containing \square as a codimension one face, and a collection of vertices spans a simplex if and only if the corresponding cubes lie in a common larger cube. There is also an alternative geometric interpretation of the link. We can define the *geometric link* of \square by choosing a point p in the interior of \square , and looking at the space of unit tangent vectors at p that are orthogonal to the cube \square . Thus each cube containing \square contributes a single simplex to the geometric link. The two notions of link are closely related, as they give homeomorphic spaces. The combinatorial link is purely combinatorial, whereas the geometric link can also be endowed with a metric structure under which each simplex is isometric to an all right spherical simplex. We will shift between the geometric link and the combinatorial link as convenient, and denote this common topological space by $\text{lk}(\square, X)$.

Lemma 5.1 *Let X be an arbitrary cube complex, $\square \subset X$ be a cube, and $v \in \square$ be a vertex of the cube. The cube \square defines a simplex $\tau \in \text{lk}(v, X)$, and there is a combinatorial isomorphism $\text{lk}(\square, X) \cong \text{lk}(\tau, \text{lk}(v, X))$.*

Proof Recall that vertices in $\text{lk}(\square, X)$ are given by cubes \square' satisfying $\square \subset \square'$ for which $\dim(\square') = \dim(\square) + 1$. Considering the simplex $\tau \in \text{lk}(v, X)$ corresponding to \square , we see that each such cube \square' gives rise to a simplex $\tau' \subset \text{lk}(v, X)$ containing the simplex τ as a codimension one face. In other words, this defines a vertex in $\text{lk}(\tau, \text{lk}(v, X))$ (see Fig. 3). Conversely, a simplex τ' in $\text{lk}(v, X)$ containing τ as a codimension one face arises from a cube \square' containing \square as a codimension one face. This establishes a bijection between the vertices of the two simplicial complexes $\text{lk}(\square, X)$ and $\text{lk}(\tau, \text{lk}(v, X))$.

Finally, we check that the actual simplicial complexes coincide. For the simplicial complex $\text{lk}(\square, X)$, the set of vertices $\{\square'_0, \dots, \square'_k\}$ span a k -simplex if and only if the cubes $\{\square'_0, \dots, \square'_k\}$ are all contained in common larger cube \square'' . Considering $\text{lk}(v, X)$, this collection of cubes correspond to a collection of simplices $\{\tau'_0, \dots, \tau'_k\}$ in $\text{lk}(v, X)$ all containing τ as a codimension one face. The larger cube \square'' leads to a larger simplex $\tau'' \subset \text{lk}(v, X)$ containing all of the τ'_i as faces. Viewing the τ'_i as a collection of vertices in $\text{lk}(\tau, \text{lk}(v, X))$, it follows that they span a simplex in $\text{lk}(\tau, \text{lk}(v, X))$. This shows that simplices in $\text{lk}(\square, X)$ give rise to simplices in $\text{lk}(\tau, \text{lk}(v, X))$. One can reverse this argument to show the converse. Thus the bijection on vertex sets described above also bijectively identifies simplices of the two complexes, establishing that the two simplicial complexes are isomorphic. \square

An identical argument can be used to show the following result.

Lemma 5.2 *Let Y be a simplicial complex, $\sigma \subset Y$ be a simplex, and $\tau \subset \sigma$ be a face of the simplex. Then σ defines a simplex $\hat{\sigma} \subset \text{lk}(\tau, Y)$, and there is a combinatorial isomorphism between $\text{lk}(\sigma, Y)$ and $\text{lk}(\hat{\sigma}, \text{lk}(\tau, Y))$. In particular, $\text{lk}(\sigma, Y) \cong \text{lk}(\hat{\sigma}, \text{lk}(v, Y))$ for any vertex v of σ .*

Next we turn our attention to the topology of the spherical sets $S(\sigma, \Sigma)$. Recall that this set is a small ϵ -sphere centered around the simplex σ , and it inherits a natural polyhedral structure from the ambient simplicial complex Σ . Indeed, each point $p \in S(\sigma, \Sigma)$ is contained in the interior of a unique simplex in Σ that splits as a join $\sigma_p * \tau_p$ where $\sigma_p \subset \sigma$ is a face of σ and τ_p is a simplex disjoint from σ . The intersection of $S(\sigma, \Sigma)$ with the simplex $\sigma_p * \tau_p$ is a copy of the polyhedron $\sigma_p \times \tau_p$, which we can think of as the “ ϵ -slice” $\sigma_p \times \tau_p \times \{\epsilon\}$ in the join:

$$\sigma_p * \tau_p = (\sigma_p \times \tau_p \times [0, 1]) / \sim$$

As we range over all simplices incident to σ , this yields the polyhedral decomposition of $S(\sigma, \Sigma)$.

It will be more convenient for our purposes to replace the small ϵ -sphere centered at σ by the combinatorial 1-sphere centered at σ . This is the full subcomplex $\hat{S}(\sigma, \Sigma)$ spanned by all vertices of Σ that are at combinatorial distance one from σ , i.e., that are not in σ but are connected by an edge to a vertex in σ .

Lemma 5.3 *Assume Σ is 5-large. Then in terms of the description of $S(\sigma, \Sigma)$ in terms of joins, $\hat{S}(\sigma, \Sigma)$ is the union of all the simplices τ_p .*

Proof Clearly the union of the simplices τ_p is contained in $\hat{S}(\sigma, \Sigma)$. To show the converse, we need to argue that for any simplex $\tau \subset \hat{S}(\sigma, \Sigma)$, we can find a vertex $w \in \sigma$ with the property that $\tau * w$ is a simplex of Σ . This is done via induction on $\dim(\tau)$, with the base case $\dim(\tau) = 0$ being immediate.

Now given τ , express it as a join $v * \tau'$ where $v \in \hat{S}(\sigma, \Sigma)$ is a vertex and $\tau' \subset \hat{S}(\sigma, \Sigma)$ is a codimension one face. By the inductive hypothesis, there exists a vertex $w' \subset \sigma$ with the property that $w' * \tau'$ is a simplex in Σ . If v is connected to w' , then $\tau * w'$ is a simplex since Σ is flag, and we are done. If v is not connected to w' , then v is connected to some vertex $v' \in \sigma$, with $v' \neq w'$. For each vertex $w \in \tau'$, consider the ordered 4-tuple $\{v, v', w', w\}$. This defines a square in Σ , and we know by construction that v is not connected to w' . Then 5-largeness forces w to be connected to v' . Since this holds for arbitrary $w \in \tau'$, we conclude that all the vertices of τ' are connected to v' . Again, the fact that Σ is flag implies that $\tau * v'$ is a simplex, thus completing the proof. □

We also introduce the notation $\hat{B}(\sigma, \Sigma)$, to denote the combinatorial 1-ball centered at σ . This is the union of all simplices that have nontrivial intersection with σ . Note that, if Σ is 5-large, Lemma 5.3 implies that $\hat{B}(\sigma, \Sigma)$ is a full subcomplex of Σ . The vertices of $\hat{B}(\sigma, \Sigma)$ naturally partition into vertices in σ and vertices in $\hat{S}(\sigma, \Sigma)$. Our next lemma describes the topology of the combinatorial 1-sphere $\hat{S}(\sigma, \Sigma)$ and the combinatorial 1-ball $\hat{B}(\sigma, \Sigma)$.

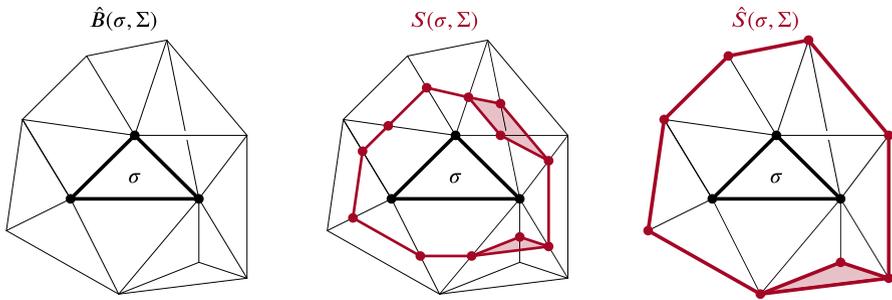


Fig. 4 The subsets described in Lemma 5.4

Lemma 5.4 *Assume Σ is 5-large, and let $\sigma \subset \Sigma$ be an arbitrary simplex. Then*

1. $\hat{S}(\sigma, \Sigma)$ and $S(\sigma, \Sigma)$ are homotopy equivalent, and
2. $\hat{B}(\sigma, \Sigma)$ is contractible.

Proof We first consider the topology of $\hat{S}(\sigma, \Sigma)$. There is a natural map $\rho : S(\sigma, \Sigma) \rightarrow \hat{S}(\sigma, \Sigma)$, given by projecting away from the simplex σ (see Fig. 4). More explicitly, on each simplex

$$\sigma_p * \tau_p = (\sigma_p \times \tau_p \times [0, 1]) / \sim$$

incident to σ , the intersection with $S(\sigma, \Sigma)$ is the polyhedral slice $\sigma_p \times \tau_p \times \{\epsilon\}$, while the intersection with $\hat{S}(\sigma, \Sigma)$ is the top slice $\tau_p \cong (\sigma_p \times \tau_p \times \{1\}) / \sim$. The map ρ is then just the projection to $\{1\}$ in the interval coordinate. In terms of the polyhedral structure on $\hat{S}(\sigma, \Sigma)$, ρ collapses the σ_p factor in each polyhedron $\sigma_p \times \tau_p$. Note that ρ is surjective by Lemma 5.3. Thus the intersection of the preimage under ρ of any point with any given polyhedron of $S(\sigma, \Sigma)$ is contractible or empty.

Now let $x \in \hat{S}(\sigma, \Sigma)$ be an arbitrary point, and let us consider $\rho^{-1}(x)$. There is a unique simplex $\tau_x \subset \hat{S}(\sigma, \Sigma)$ containing x in its interior. Lemma 5.3 implies that there is a maximal face $\sigma_x \subset \sigma$ associated with τ_x such $\sigma_x * \tau_x$ is a simplex of Σ , and σ_x is unique since Σ is flag. Then from the discussion in the previous paragraph, $\rho^{-1}(x)$ is a copy of $\sigma_x \times \{x\}$ inside the slice $\sigma_x \times \tau_x$, hence it is contractible. A continuous map with contractible point preimages is a cell-like map, and it is well known that a cell-like map between compact spaces homeomorphic to retracts of open subsets of \mathbb{R}^n is a homotopy equivalence (see [11, Cor. 1.3]). This proves that ρ is a homotopy equivalence, as desired.

To verify the corresponding statement for $\hat{B}(\sigma, \Sigma)$, we note that $\hat{B}(\sigma, \Sigma)$ is the union of simplices of the form $\sigma_p * \tau_p$, where τ_p is a face in $\hat{S}(\sigma, \Sigma)$ and σ_p is a face of σ by Lemma 5.3. We now form an auxiliary space $B(\sigma, \Sigma)$ and a continuous quotient map $\phi : B(\sigma, \Sigma) \rightarrow \hat{B}(\sigma, \Sigma)$. Recall that the space $S(\sigma, \Sigma)$ is the ϵ -sphere centered at σ , identified with the ϵ -slice $\sigma_p \times \tau_p \times \{\epsilon\}$ in the join description of each simplex incident to σ . We let $B(\sigma, \Sigma)$ denote the closed ϵ -ball centered at σ . Thus $B(\sigma, \Sigma)$ gets identified with $(\sigma_p \times \tau_p \times [0, \epsilon]) / \sim$ in the join description of each simplex incident to σ ; recall that at the 0-level we collapse the τ_p coordinates to points. Since we can contract the interval factors down, we clearly have that $B(\sigma, \Sigma)$ deformation retracts to σ , hence it is contractible.

The space $B(\sigma, \Sigma)$ contains a copy of $S(\sigma, \Sigma)$ at level ϵ , and we now define ϕ to be the quotient map ρ on the subset $S(\sigma, \Sigma)$. As we can stretch out the $[0, \epsilon]$ coordinate to $[0, 1]$, we see that the quotient can naturally be identified with $\hat{B}(\sigma, \Sigma)$. Moreover point preimages under ϕ are either points when $x \notin \hat{S}(\sigma, \Sigma)$ or coincide with the point preimage under ρ when $x \in \hat{S}(\sigma, \Sigma)$, and hence point preimages are always contractible. Thus the map ϕ is cell-like, and hence it is a homotopy equivalence. We conclude that the space $\hat{B}(\sigma, \Sigma)$ is contractible, as desired. \square

Remark 5.5 An alternative proof is possible for the second statement of Lemma 5.4 by noting that the maximal dimensional faces on $\hat{S}(\sigma, \Sigma)$ are free faces in the complex $\hat{B}(\sigma, \Sigma)$ by Lemma 5.3. Thus one can perform simplicial collapses, reducing the maximal dimension of faces in $\hat{S}(\sigma, \Sigma)$. After finitely many steps, this simplicially collapses $\hat{B}(\sigma, \Sigma)$ to the complex σ .

5.2 Analysis of combinatorial spheres in thickenings

We now return to analyzing the thickenings T_n from Sect. 4 in order to prove Proposition 4.1. Adopting all notation from that section, recall that the goal is to prove that $\text{cd}(S(\sigma, T_n)) \leq n - 1$ for all nonempty simplices σ of T_n . Equivalently, we want to show that $\bar{H}^i(S(\sigma, T_n)) = 0$ for all $i \geq n$ and $\emptyset \neq \sigma \subset T_n$.

As a first step, let us describe the topology of links in a general thickening. For any simplex $\sigma \subset \text{Th}_\alpha(X)$, by definition of the α -thickening we have that $\alpha(V(\sigma))$ is contained in some cube of X . Consider a minimal such cube under containment. In a general cube complex X this minimal cube might not be unique. However we have the following:

Lemma 5.6 *Let X be a 5-large cube complex, $\text{Th}_\alpha(X)$ an arbitrary thickening, and $\sigma \subset \text{Th}_\alpha(X)$ a nonempty simplex. Then there exists a unique minimal cube $\square_\sigma \subset X$ containing the vertices $\alpha(V(\sigma))$.*

Proof We are given a set $\alpha(V(\sigma))$ of vertices in the cube complex X , which we know is contained in a cube, and want to argue that there is a unique minimal cube containing $\alpha(V(\sigma))$. This is a purely local statement, as it only depends on the combinatorial 1-neighborhood of a vertex in $\alpha(V(\sigma))$. From the 5-large hypothesis, such a neighborhood lifts to \tilde{X} . Therefore it suffices to consider this question in \tilde{X} , which is a CAT(0) cube complex. However any two cubes in a CAT(0) cube complex can only intersect in a common subcube. Now assume \square is another minimal cube containing the set $\alpha(V(\sigma))$. Then the intersection $\square \cap \square_\sigma$ is a subcube of both \square_σ , and \square . Minimality then implies that $\square_\sigma = \square \cap \square_\sigma = \square$, which implies the minimal cube is unique. \square

Notice that in our setting, the cube complexes we are thickening are always 5-large, therefore Lemma 5.6 applies. We will denote by \square_σ the unique minimal cube containing the set $\alpha(V(\sigma))$. If we have a proper containment of simplices $\sigma \subsetneq \tau$, then clearly $\square_\sigma \subseteq \square_\tau$. Note that the containment of the corresponding cubes might not be proper, i.e., it could be that $\square_\sigma = \square_\tau$ despite the fact that $\sigma \neq \tau$. Our next lemma will exploit this feature and use it to relate the topology of the link of σ with the topology of the link of \square_σ .

Fig. 5 Given a simplex σ (red) in $\text{Th}_\alpha(X)$, let \square_σ (black) be the minimal cube in X containing $\alpha(V(\sigma))$. The maximal simplex (green) in $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ is spanned by the vertices in $\alpha^{-1}(V(\square_\sigma)) \setminus V(\sigma)$ color figure online

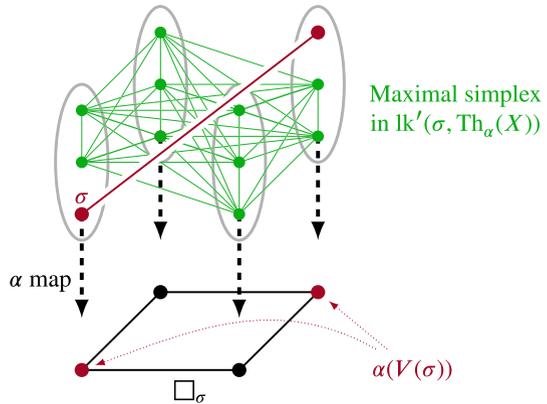
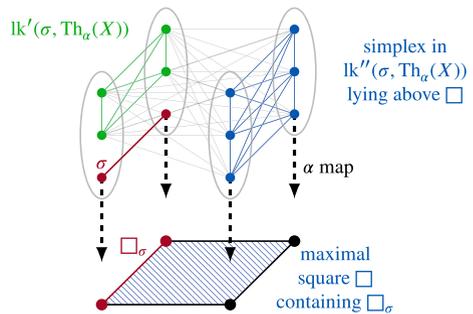


Fig. 6 A general simplex in $\text{lk}(\sigma, \text{Th}_\alpha(X))$ splits as a join of simplices in $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ (green) and $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ (blue) color figure online



Lemma 5.7 *Given any simplex $\sigma \subset \text{Th}_\alpha(X)$, the link $\text{lk}(\sigma, \text{Th}_\alpha(X))$ contains a pair of subcomplexes $\text{lk}'(\sigma, \text{Th}_\alpha(X))$, $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ with the following three properties:*

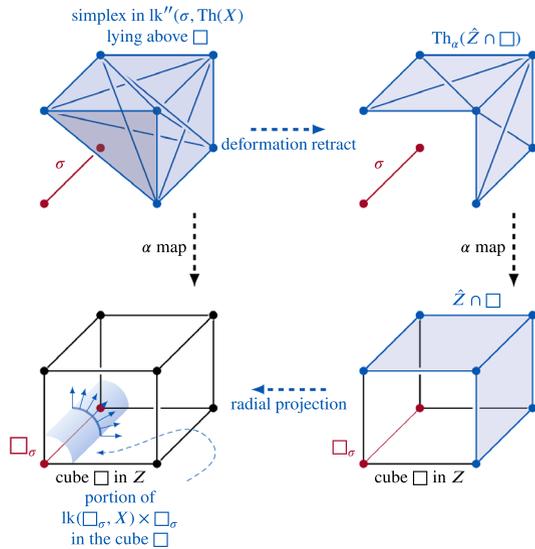
1. $\text{lk}(\sigma, \text{Th}_\alpha(X))$ decomposes as a join $\text{lk}'(\sigma, \text{Th}_\alpha(X)) * \text{lk}''(\sigma, \text{Th}_\alpha(X))$;
2. if $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ is nonempty, then it is contractible;
3. the complex $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ is homotopy equivalent to $\text{lk}(\square_\sigma, X)$.

Proof We start by defining the subcomplex $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ (see Fig. 5) and verifying property (2). Consider the vertices in the set $\alpha^{-1}(V(\square_\sigma))$. From the definition of the thickening $\text{Th}_\alpha(X)$, these span a simplex containing σ . Thus all vertices in this set are either at combinatorial distance zero or one from σ , according to whether they lie in σ or not. We let $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ be spanned by those vertices at distance one. This subcomplex is either empty or consists of a single simplex, and hence it is contractible. This establishes property (2).

Next let us define the complex $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ and establish property (1). This is the subcomplex of $\text{lk}(\sigma, \text{Th}_\alpha(X))$ spanned by vertices v with the property that $\alpha(v) \notin \square_\sigma$. In particular, the vertices of $\text{lk}(\sigma, \text{Th}_\alpha(X))$ are partitioned between $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ and $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ according to whether $\alpha(v) \in \square_\sigma$ or $\alpha(v) \notin \square_\sigma$ (see Fig. 6). This implies the subcomplexes $\text{lk}'(\sigma, \text{Th}_\alpha(X))$ and $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ are disjoint subcomplexes inside $\text{lk}(\sigma, \text{Th}_\alpha(X))$ and gives the containment

$$\text{lk}(\sigma, \text{Th}_\alpha(X)) \subseteq \text{lk}'(\sigma, \text{Th}_\alpha(X)) * \text{lk}''(\sigma, \text{Th}_\alpha(X)).$$

Fig. 7 Sequence of homotopy equivalences illustrating the argument for Property (3) in Lemma 5.7



For the reverse containment, given a pair of simplices $\tau' \subseteq \text{lk}'(\sigma, \text{Th}_\alpha(X))$ and $\tau'' \subseteq \text{lk}''(\sigma, \text{Th}_\alpha(X))$ we have that $\alpha(V(\tau')) \subseteq \square_\sigma$, while $\alpha(V(\tau'')) \subseteq \square$ for some cube $\square_\sigma \subsetneq \square$. It follows that $\alpha(V(\tau') \cup V(\tau'')) \subseteq \square$ and hence that $\tau' * \tau''$ defines a simplex in $\text{lk}(\sigma, \text{Th}_\alpha(X))$. This yields the reverse containment and establishes property (1).

To complete the proof of the lemma, it remains to verify property (3). Let Z denote the union of all the cubes containing \square_σ and \hat{Z} be the cube subcomplex of Z consisting of all cubes in Z that do not intersect \square_σ . There is a natural homeomorphism $\hat{Z} \cong \square_\sigma \times \text{lk}(\square_\sigma, X)$, so \hat{Z} is homotopy equivalent to $\text{lk}(\square_\sigma, X)$. From the definition of $\text{lk}''(\sigma, \text{Th}_\alpha(X))$, any simplex $\tau \subset \text{lk}''(\sigma, \text{Th}_\alpha(X))$ has the property that $\alpha(V(\tau))$ is disjoint from \square_σ , and hence it lies in $V(\hat{Z})$. In other words, the entire complex $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ lies “above” the cube complex \hat{Z} .

Now consider the thickening $\text{Th}_\alpha(\hat{Z})$ obtained by the restriction $\alpha : \alpha^{-1}(V(\hat{Z})) \rightarrow V(\hat{Z})$. This is a subcomplex of $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ that is homotopy equivalent to \hat{Z} , and hence to $\text{lk}(\square_\sigma, X)$. To complete the verification of property (3) it therefore suffices to show that $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ deformation retracts to $\text{Th}_\alpha(\hat{Z})$.

Recall that \hat{Z} is defined as a subcomplex of the cube complex Z . To construct the deformation retraction, we proceed by induction on the dimension of the cubes in Z . More precisely, we will consider a cube $\square \subset Z$ containing \square_σ and the restriction of the two complexes $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ and $\text{Th}_\alpha(\hat{Z})$ lying above $\hat{Z} \cap \square$, i.e., whose vertices α -project to vertices in $\hat{Z} \cap \square$. We induct on $\dim(\square)$ and show that in each dimension we can choose a deformation retraction compatible with the previously-constructed deformation retractions (see Fig. 7). This will allow the deformation retractions associated with each cube to glue compatibly and yield a global deformation retraction of $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ to $\text{Th}_\alpha(\hat{Z})$.

The bottom of the induction is the case of a cube $\square \subset Z$ with $\dim(\square) = \dim(\square_\sigma) + 1$. In this case, $\square = \square_\sigma \times I$, and $\hat{Z} \cap \square$ is the face of \square opposite \square_σ . This forces the two subcomplexes above $\hat{Z} \cap \square$ to coincide, and hence we can take the identity

map as the deformation retraction. Now for the inductive step, we assume that we are already provided the deformation retractions for the portions of the simplex that lie above the relevant lower dimensional faces of the cube $\hat{Z} \cap \square$. We would like to extend the deformation retraction to the rest of the simplicial complex lying above $\hat{Z} \cap \square$.

In general, if we have a simplicial complex A and a subcomplex $B \subseteq A$, the obstruction to deformation retracting A to B lies in the relative homotopy groups $\pi_*(A, B)$. In our setting, we already have a deformation retraction from some intermediate subspace $B \subseteq B' \subseteq A$, and would like to extend the deformation retraction to one from A to B . Again, the obstructions are the relative homotopy groups $\pi_*(A, B') \cong \pi_*(A, B)$. If these vanish, then one can construct the desired deformation retraction by composing a deformation retraction from A to B' with the given deformation retraction from B' to B . Thus the problem reduces to understanding the topology of the two sets we are studying.

To this end, let us identify the topology of the part of $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ lying above $\hat{Z} \cap \square$. From the definition, the vertices of this subcomplex are precisely those whose α -projection are contained in $V(\square) \setminus V(\square_\sigma)$. Since these vertices have image contained in a single cube, they span a simplex. Thus the restriction of the complex $\text{lk}''(\sigma, \text{Th}_\alpha(X))$ to the preimage under α of the cube \square is a single simplex Δ , and thus it is a contractible set.

On the other hand, the subcomplex $\text{Th}_\alpha(\hat{Z} \cap \square)$ is the α -thickening of the cube complex $\hat{Z} \cap \square$. However this cube complex is spanned by the vertices $V(\square) \setminus V(\square_\sigma)$, and thus is homotopic to the contractible set $\partial \square \setminus \square_\sigma$. We conclude that $\text{Th}_\alpha(\hat{Z} \cap \square)$ is also a contractible subset of Δ . Since both spaces are contractible, the relative homotopy groups all vanish and the desired deformation retraction exists. This completes the inductive step and verifies statement (3), completing the proof of the lemma. \square

Informally, Lemma 5.7 shows that, up to homotopy, the links of simplices in a thickening are either trivial or coincide with the links in the underlying cube complex. Note the following immediate consequence of the previous result.

Lemma 5.8 *If σ is an arbitrary nonempty simplex in T_n , then $\text{cd}(\text{lk}(\sigma, T_n)) \leq n - 1$.*

Proof Since T_n is a thickening, we can apply Lemma 5.7 to see that $\text{lk}(\sigma, T_n)$ decomposes as a join $\text{lk}'(\sigma, T_n) * \text{lk}''(\sigma, T_n)$. If $\text{lk}'(\sigma, T_n) \neq \emptyset$, then $\text{lk}'(\sigma, T_n)$ is a nonempty simplex, hence it is contractible. This forces $\text{lk}(\sigma, T_n)$ to likewise be contractible, in which case we are done.

Alternatively, if $\text{lk}'(\sigma, T_n) = \emptyset$, then Lemma 5.7 tells us that

$$\text{lk}(\sigma, T_n) = \text{lk}''(\sigma, T_n) \simeq \text{lk}(\square_\sigma, X_n).$$

We let $v \in \square_\sigma$ be an arbitrary vertex of the cube \square_σ , and apply Lemma 5.1 to obtain a combinatorial isomorphism

$$\text{lk}(\square_\sigma, X_n) \cong \text{lk}(\tau, \text{lk}(v, X)) \cong \text{lk}(\tau, T_{n-1})$$

where τ is the (possibly empty) simplex in $\text{lk}(v, X)$ corresponding to \square_σ .

If the simplex τ is empty, this means \square_σ coincides with the vertex v and $\text{lk}(\tau, T_{n-1}) = T_{n-1}$. In that case we conclude by induction with respect to proving the Main Theorem that

$$\text{cd}(\text{lk}(\sigma, T_n)) = \text{cd}(T_{n-1}) = n - 1.$$

On the other hand, if τ is nonempty then we again obtain

$$\text{cd}(\text{lk}(\sigma, T_n)) = \text{cd}(\text{lk}(\tau, T_{n-1})) \leq n - 2$$

by induction with respect to both this lemma and the Main Theorem. Either way we obtain the desired bound, which concludes the proof of the lemma. \square

In order to compute the cohomological dimension of the polyhedral complex $S(\sigma, T_n)$, we will instead compute the cohomological dimension of the simplicial complex $\hat{S}(\sigma, T_n)$. Our next step is to give a partition of vertices in $\hat{S}(\sigma, T_n)$, according to their closest face of σ .

Note that every vertex v in $\hat{S}(\sigma, T_n)$ is at distance one from some vertex of σ by definition. For such a vertex v , we can look at all the vertices of σ that are adjacent to v . This defines a face of σ , and gives us a map $\pi : V(\hat{S}(\sigma, T_n)) \rightarrow \mathcal{F}(\sigma)$ from the set of vertices in $\hat{S}(\sigma, T_n)$ to the set $\mathcal{F}(\sigma)$ of faces of σ . We think of $\pi(v)$ as encoding the face of σ that is closest to the vertex v . We can now partition the vertices of $\hat{S}(\sigma, T_n)$ according to their image under π ; see Fig. 8. Next, use this partition to form a filtration of $\hat{S}(\sigma, T_n)$ by subcomplexes. Denote by S_i the subcomplex of $\hat{S}(\sigma, T_n)$ spanned by the preimage under π of all faces of codimension $\leq i$. We clearly have containments $S_0 \subset S_1 \subset \dots \subset S_k = \hat{S}(\sigma, T_n)$, where $k = \text{dim}(\sigma)$. This gives us a filtration of the complex $\hat{S}(\sigma, T_n)$ by the subcomplexes S_i ; again, see Fig. 8.

To compute the cohomological dimension of $\hat{S}(\sigma, T_n)$, we will analyze how the cohomological dimension changes as we work our way up this filtration. To do this, we now analyze how each simplicial complex S_i in the filtration is obtained from the previous simplicial complex S_{i-1} . The vertices of the subcomplex S_i are obtained from S_{i-1} by adding in all vertices $v \in \hat{S}(\sigma, T_n)$ with the property that $\pi(v)$ is a codimension i face of σ . Take a codimension i face $\tau \subset \sigma$, and consider the set of vertices $L_\tau := \pi^{-1}(\tau) \subset S_i$. As τ ranges over all the codimension i faces, the sets L_τ partition the vertices in $S_i \setminus S_{i-1}$. This gives us a description of the vertices that are added, as they are given by the disjoint union $\bigsqcup L_\tau$ as τ ranges over the codimension i faces. Next we need to describe the simplices that are added.

Lemma 5.9 *If τ, τ' are distinct codimension i faces of σ , then there are no edges joining vertices of L_τ to vertices of $L_{\tau'}$.*

Proof Consider a pair of vertices $v \in L_\tau$ and $v' \in L_{\tau'}$, where τ, τ' are distinct codimension i faces. Since $\tau \neq \tau'$ have the same codimension in σ , there is a vertex $w \in \tau \setminus \tau'$, and a vertex $w' \in \tau' \setminus \tau$. The vertices $v \in L_\tau$ and $w \in \tau$ are connected by an edge by the defining property of L_τ . Similarly v' and w' are connected by an edge. Then w, w' are both vertices in the simplex σ , so they are connected by an edge. For a contradiction, assume that v, v' are connected by an edge. Then the ordered

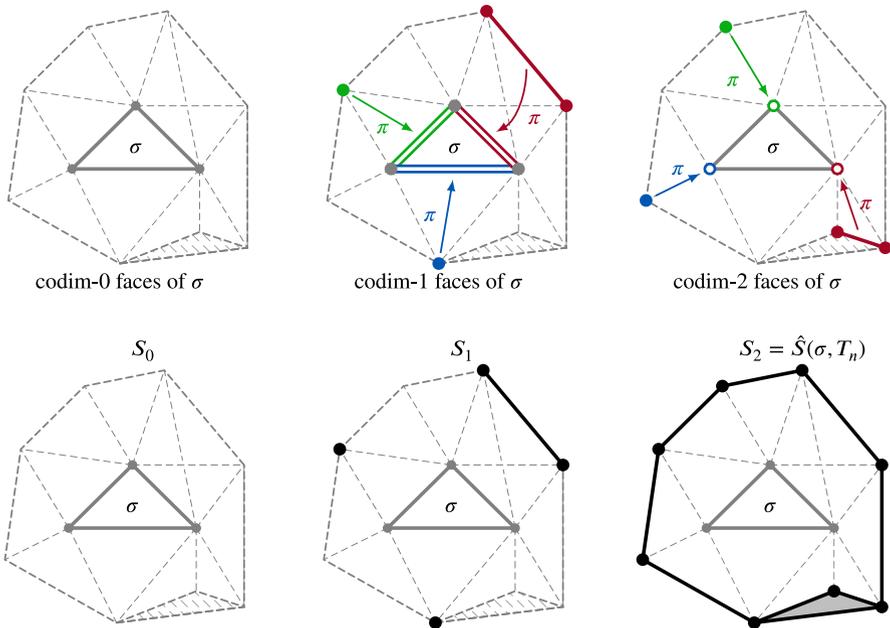


Fig. 8 The top row shows $\pi : V(\hat{S}(\sigma, T_n)) \rightarrow \mathcal{F}(\sigma)$ for faces of σ in each codimension. The bottom row shows the filtration $S_0 \subset S_1 \subset \dots \subset \hat{S}(\sigma, T_n)$

set of vertices v, w, w', v' defines a 4-cycle in T_n , which we know is 5-large, hence the 4-cycle must contain a diagonal. If v is connected by an edge to w' , we obtain a contradiction to the fact that $\pi(v) = \tau$ and $w' \notin \tau$. A similar contradiction occurs if v' is connected by an edge to w . We conclude that v, v' cannot be connected by an edge, establishing the lemma. \square

As a consequence of Lemma 5.9, any of the new simplices that are added to S_{i-1} either span vertices in a single L_τ or connect some vertices in L_τ to vertices in S_{i-1} . Denote the collection of vertices in $\text{lk}(\tau, T_n) \cap S_{i-1}$ by ∂L_τ .

Lemma 5.10 *The subcomplex of S_i spanned by the vertices $L_\tau \cup \partial L_\tau$ is naturally isomorphic to the simplicial complex $\text{lk}(\tau, T_n) \setminus \hat{\sigma}$, where $\hat{\sigma} \subset \text{lk}(\tau, T_n)$ is the simplex in $\text{lk}(\tau, T_n)$ corresponding to σ . Under this isomorphism, the set of vertices ∂L_τ corresponds to the vertices in $\text{lk}(\tau, T_n)$ at distance one from $\hat{\sigma}$. Moreover, this subcomplex contains all the simplices in S_i that contain a vertex of L_τ .*

Proof By construction, the vertices in L_τ can be identified with vertices in $\text{lk}(\tau, T_n)$ by viewing the link as a subcomplex in the flag complex T_n . We claim that the vertices of $\text{lk}(\tau, T_n) \setminus \{L_\tau \cup \hat{\sigma}\}$ are precisely the vertices of $\text{lk}(\tau, T_n)$ at combinatorial distance one from $\hat{\sigma}$. Indeed, if v is at combinatorial distance one from $\hat{\sigma}$, this means there is some vertex $w \in \hat{\sigma}$ connected by an edge to v . In terms of the embedding of $\text{lk}(\tau, T_n)$ into T_n , the simplex $\hat{\sigma} \subset \text{lk}(\tau, T_n)$ corresponds to the face of σ opposite τ , so $w \in \sigma \setminus \tau$. Since $v \in \text{lk}(\tau, T_n)$ is connected by an edge to w , we have that $\tau \cup \{w\} \subseteq \pi(v)$. This

forces the codimension of $\pi(v)$ to be at most $i - 1$, which means that $v \in S_{i-1}$. This argument can be reversed to show that the vertices in L_τ are precisely the vertices in $\text{lk}(\tau, T_n)$ whose combinatorial distance to $\hat{\sigma}$ is at least two. Of course, the vertices of $\hat{\sigma}$ in $\text{lk}(\tau, T_n)$ are precisely those in the face of σ opposite τ , so they do not lie in S_{i-1} . We conclude that the set ∂L_τ consists of precisely the vertices in $\text{lk}(\tau, T_n)$ that are exactly at distance 1 from $\hat{\sigma}$.

Thus all the vertices in $L_\tau \cup \partial L_\tau$ lie in $\text{lk}(\tau, T_n)$, and correspond to those vertices at combinatorial distance at least one from $\hat{\sigma} \subset \text{lk}(\tau, T_n)$. Since T_n is flag, the subcomplex they span is naturally isomorphic to the subcomplex of $\text{lk}(\tau, T_n)$ spanned by all its vertices that are not in $\hat{\sigma}$.

To check the last statement, we note that Lemma 5.9 implies that if $w \in S_i$ is a vertex connected to some vertex $v \in L_\tau$, then $w \in L_\tau \cup S_{i-1}$. It then suffices to check that any such $w \in S_{i-1}$ in fact lies in $\partial L_\tau = S_{i-1} \cap \text{lk}(\tau, T_n)$. In other words, it is sufficient to check that such a w lies in $\text{lk}(\tau, T_n)$. Since $w \in S_{i-1}$, we know that $\pi(w)$ is a face of codimension $\leq i - 1$. Since τ has codimension i , there exists a vertex $w' \in \sigma \setminus \tau$ connected to w . We now claim that $\tau \subset \pi(w)$, and hence that $w \in \text{lk}(\tau, T_n)$. Indeed, let $v' \in \tau$ be an arbitrary vertex, and observe that the ordered 4-tuple of vertices v, w, w', v' forms a 4-cycle. Since v is not connected to w' , since $w' \notin \tau = \pi(v)$, 5-largeness of T_n forces the w to be connected to v' . We conclude that any vertex $w \in S_{i-1}$ connected to a vertex of L_τ is also in $\text{lk}(\tau, T_n)$, and hence lies in the subset $\partial L_\tau = S_{i-1} \cap \text{lk}(\tau, T_n)$. This checks the last statement, and completes the proof of the lemma. □

Finally, we will require the following technical lemma.

Lemma 5.11 *Consider a pair of simplicial complexes Z, Z' that can be expressed as gluings $Z = A \cup_C B$ and $Z' = A \cup_C B'$, where each inclusion of C into A induces the same maps on reduced cohomology. Assume B is contractible, $\text{cd}(Z) \leq n - 1$, and $\text{cd}(B') \leq n - 1$. Then $\text{cd}(Z') \leq n - 1$.*

Proof We first consider the Mayer–Vietoris sequence associated with the gluing $Z = A \cup_C B$. Since B is contractible by hypothesis, the sequence reduces to

$$\dots \longrightarrow \bar{H}^i(Z) \longrightarrow \bar{H}^i(A) \longrightarrow \bar{H}^i(C) \longrightarrow \bar{H}^{i+1}(Z) \longrightarrow \dots$$

We also know by hypothesis that $\text{cd}(Z) \leq n - 1$, so $\bar{H}^i(Z) = 0$ for $i \geq n$. The long exact sequence now tells us that the map $\bar{H}^i(A) \rightarrow \bar{H}^i(C)$ induced by the inclusion $C \hookrightarrow A$ is an isomorphism when $i \geq n$ and a surjection when $i = n - 1$.

Now we turn our attention to $Z' = A \cup_C B'$, obtaining the Mayer–Vietoris sequence:

$$\dots \longrightarrow \bar{H}^{i-1}(C) \longrightarrow \bar{H}^i(Z') \longrightarrow \bar{H}^i(A) \oplus \bar{H}^i(B') \longrightarrow \bar{H}^i(C) \longrightarrow \dots$$

By hypothesis, we have that $\text{cd}(B') \leq n - 1$, so $\bar{H}^i(B') = 0$ for $i \geq n$. In particular, when $i \geq n$ the sequence above reduces to:

$$\dots \longrightarrow \bar{H}^{i-1}(C) \longrightarrow \bar{H}^i(Z') \longrightarrow \bar{H}^i(A) \longrightarrow \bar{H}^i(C) \longrightarrow \dots$$

From our earlier analysis, we also know that the map $\overline{H}^i(A) \rightarrow \overline{H}^i(C)$ is an isomorphism when $i \geq n$, and a surjection when $i = n - 1$. This forces $\overline{H}^i(Z') = 0$ in the range $i \geq n$, which proves the lemma. \square

With our preliminaries in place, we can now establish Proposition 4.1, and hence complete the proof of the Main Theorem.

Proof We want to show that $cd(S(\sigma, T_n)) \leq n - 1$ for every nonempty simplex $\sigma \subset T_n$. From Lemma 5.4, we have a homotopy equivalence $S(\sigma, T_n) \simeq \hat{S}(\sigma, T_n)$, so we instead show that $cd(\hat{S}(\sigma, T_n)) \leq n - 1$. As discussed above, the set $\hat{S}(\sigma, T_n)$ has a filtration

$$S_0 \subset S_1 \subset \dots \subset S_k = \hat{S}(\sigma, T_n).$$

We will consider the cohomological dimension of each stage in this filtration, and check that it always remains less than or equal to $n - 1$.

At the bottom of the filtration we have the space S_0 . This is the full subcomplex of $\hat{S}(\sigma, T_n)$ containing all vertices v that satisfy $\pi(v) = \sigma$. This means that the vertex v is adjacent to every vertex of σ . Since T_n is a flag complex, we see that the join $v * \sigma$ is a $(k + 1)$ -simplex containing σ , and hence v can be identified with a vertex in $lk(\sigma, T_n)$. It follows that the subcomplex S_0 is precisely the copy of $lk(\sigma, T_n)$ inside the complex $\hat{S}(\sigma, T_n)$. Applying Lemma 5.8 we see that $cd(S_0) = cd(lk(\sigma, T_n)) \leq n - 1$. This establishes the base case of the induction. Next we consider how the cohomological dimension changes when we go from some S_{i-1} to S_i .

By induction, we may now assume that $cd(S_{i-1}) \leq n - 1$. From Lemma 5.9 and Lemma 5.10, we see that S_i is obtained by gluing on a finite collection of subcomplexes to S_{i-1} . This collection is parametrized by the $\binom{k+1}{i}$ codimension i faces of σ . Let us now analyze how the topology changes each time we add one of these subcomplexes. We let A_1, A_2, \dots denote the collection of spaces we are adding on to S_{i-1} and B'_j denote the result of adding A_1, \dots, A_j to S_{i-1} . Note that once we have reached $j = \binom{k+1}{i}$, we have added on the subsets corresponding to each codimension i face, hence $B'_j = S_{i+1}$. Thus it will be sufficient to argue that each $cd(B'_j) \leq n - 1$, which we now do by a second inductive argument.

We will use Lemma 5.11 to see that the cohomological dimension cannot increase as we go from B'_j to B'_{j+1} . In the notation of Lemma 5.11, we are setting:

- $Z := lk(\tau, T_n)$ where τ is one of the codimension i faces of σ , so $cd(Z) \leq n - 1$ by Lemma 5.8;
- $A := A_{j+1}$ is the subcomplex of Z spanned by all vertices that are not in the subcomplex $\hat{\sigma}$ from Lemma 5.10 applied to τ ;
- B is the subcomplex of Z consisting of all simplices incident to $\hat{\sigma}$ (from Lemma 5.10), which coincides with $\hat{B}(\hat{\sigma}, Z)$, and is thus contractible by Lemma 5.4, since T_n , hence Z , is flag and 5-large;
- $B' := B'_j$ is the union of S_{i-1} with finitely many of the subcomplexes A_1, \dots, A_j already added on, so by induction we can assume that $cd(B') \leq n - 1$;
- $C := \partial L_\tau \subseteq S_{i-1}$ from the conclusions about ∂L_τ in Lemma 5.10, hence C is contained in B and B' .

Thus all the hypotheses of Lemma 5.11 are satisfied, which allows us to conclude that the space $Z' = A \cup_C B'$ satisfies $cd(Z') \leq n - 1$. However, observe that Z' is

exactly the result of adding on $A = A_{j+1}$ to $B' = B'_j$, hence it is the space B'_{j+1} , thus completing the inductive step. After performing this step for the finitely many codimension i faces, we see that $\text{cd}(S_i) \leq n - 1$. This completes the inductive step, and continuing this procedure up the filtration completes the proof that $\text{cd}(S(\sigma, T_n)) = \text{cd}(\hat{S}(\sigma, T_n)) \leq n - 1$. \square

Acknowledgements This work was completed as part of the ‘Collaborative Projects in Geometric Topology’ program organized by the Geometry and Topology group at the University of Virginia and supported by an RTG grant from the National Science Foundation (DMS-1839968). We thank the main organizers Sara Maloni and Thomas Mark for providing us the opportunity to work together on this project. We also thank Mike Davis, Jingyin Huang, and Boris Okun for some helpful comments and references. Lafont was partially supported by the NSF under grant DMS-2407438. Minemyer was partially supported by an AMS-Simons Research Enhancement Grant for Primarily Undergraduate Institution Faculty. Stover was partially supported by NSF grants DMS-2203555 and DMS-2506896, along with award SFI-MPS-TSM-00014184 from the Simons Foundation. The authors thank the referee for a very careful reading of the paper.

Data availability Data sharing is not applicable to this article as no datasets were generated or analyzed during the current study.

Declarations

Conflict of interest The authors have no conflict of interest to declare that are relevant to the content of this article.

References

- Battista, L., Martelli, B.: Hyperbolic 4-manifolds with perfect circle-valued morse functions. *Trans. Am. Math. Soc.* **375**, 2597–2625 (2022)
- Bestvina, M., Brady, N.: Morse theory and finiteness properties of groups. *Invent. Math.* **129**(3), 445–470 (1997)
- Charney, R., Farber, M.: Random groups arising as graph products. *Algebr. Geom. Topol.* **12**(2), 979–995 (2012)
- Davis, M.W.: *The Geometry and Topology of Coxeter Groups*. London Mathematical Society Monographs Series, vol. 32. Princeton University Press, Princeton, NJ (2008)
- Davis, M.W., Kahle, M.: Random graph products of finite groups are rational duality groups. *J. Topol.* **7**(2), 589–606 (2014)
- Fisher, S.P.: Algebraic fibering of a hyperbolic 7-manifold. *Bull. Lond. Math. Soc.* **55**(3), 1347–1357 (2023)
- Haglund, F., Wise, D.T.: Special cube complexes. *Geom. Funct. Anal.* **17**(5), 1551–1620 (2008)
- Italiano, G., Martelli, B., Migliorini, M.: Hyperbolic manifolds that fiber algebraically up to dimension 8. *J. Inst. Math. Jussieu*, pp. 1–38 (2022)
- Italiano, G., Martelli, B., Migliorini, M.: Hyperbolic 5-manifolds that fiber over \mathbb{S}^1 . *Invent. Math.* **231**, 1–38 (2023)
- Jankiewicz, K., Norin, S., Wise, D.T.: Virtually fibering right-angled coxeter groups. *J. Inst. Math. Jussieu* **20**(3), 957–987 (2021)
- Lacher, R.C.: Cell-like mappings. I. *Pac. J. Math.* **30**, 717–731 (1969)
- Osajda, D.: A construction of hyperbolic coxeter groups. *Comment. Math. Helv.* **88**(2), 353–367 (2013)
- Sageev, M.: CAT(0) cube complexes and groups. In: *Geometric Group Theory*, vol. 21 of IAS/Park City Math. Ser., pp. 7–54. American Mathematical Society, Providence, RI (2014)
- Schesler, E., Zaremsky, M.C.B.: Random subcomplexes of finite buildings, and fibering of commutator subgroups of right-angled Coxeter groups. *J. Topol.* **16**(1), 20–56 (2023)

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.