Combinatorics of CAT(0) cubical complexes, crossing complexes and co-skeletons

Rowan Rowlands

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2023

Reading committee:

Isabella Novik, Chair

Steven Klee

Rekha Thomas

Program authorized to offer degree:

Department of Mathematics

University of Washington

Abstract

Combinatorics of CAT(0) cubical complexes, crossing complexes and co-skeletons

Rowan Rowlands

Chair of the Supervisory Committee:

Isabella Novik

Department of Mathematics

This thesis consists of three papers about cubical complexes: Chapter 1 is [Row22a], Chapter 2 is [Row23], and Chapter 3 is [Row22b].

Chapter 1 extends a result by Dancis [Dan84] to cubical complexes: Dancis proved that any d-dimensional simplicial manifold can be reconstructed from its $(\lfloor d/2 \rfloor + 1)$-skeleton, and we prove an analogous result for d-dimensional cubical manifolds that can be embedded as a subcomplex into a cube I^N.

Chapter 2 studies CAT(0) cubical complexes, using the framework of a poset with inconsistent pairs developed by Ardila et al. [AOS12]. We introduce a simplicial complex called the "crossing complex" associated to each CAT(0) cubical complex, and study its properties. We deduce that this crossing complex holds much of the combinatorial information contained in the cubical complex: our main results relate their f-vectors, hyperplane/link structure, and balancedness.

Finally, Chapter 3 studies the topology of complements of skeletons in polytopal complexes: we derive a long exact sequence involving homology of skeleton complements and links, and we characterise various topological properties of spaces in terms of skeleton complements. Our main application of this machinery is to CAT(0) cubical complexes: we conclude that these complexes also share several topological properties with their crossing complexes.

Acknowledgments

Thanks to Marge Bayer, Martina Juhnke-Kubitzke, Steve Klee, Raman Sanyal and various anonymous referees for helpful comments on portions of this research. Thanks also to Jayadev Athreya, Sara Billey, Bennet Goeckner and other UW faculty members for insightful conversations and lessons on cubes, counting and other topics. And thanks to the members of my committee, Shayan Oveis Gharan (GSR), Steve Klee and Rekha Thomas, for dedicating their time to review my research.

Thanks to my parents for their constant support and encouragement, and to my grandparents for inspiring me to join the world of academia.

I am also grateful to my friends, Albert Artiles, Kristine Hampton and Adam Kapilow, for keeping me (relatively) sane while completing a PhD during a pandemic. Similar thanks are also owed to Cougar Mountain Premium Cookies, the creators of Elden Ring, and Weaving Works (RIP).

But my biggest thanks, of course, go to my advisor, Isabella Novik. Her encyclopaedic knowledge of combinatorial research, attention to detail and ability to ask just the right questions combine to make her an advisor without equal. I could not have written this thesis without her support and guidance.

This research was partially supported by a graduate fellowship from NSF grants DMS-1664865 and DMS-1953815.

Contents

List of Figures

Introduction

> *"I don't want to say too many words about the magic of the Cube, because it's basically a mystery. It's like the Mona Lisa smile. It's both complex and very simple at the same time. And, well, people like it. Even today."*
>
> — Erno Rubik

This thesis consists of three papers on cubical complexes written over the course of the author's degree: Chapter 1 is [Row22a], Chapter 2 is [Row23], and Chapter 3 is [Row22b]. These papers are reproduced as they appear in their published or preprint forms, with some minor changes.

Chapter 1, [Row22a], adapts a proof by Dancis to show that any cubical manifold that can be embedded as a subcomplex of a cube can be completely determined from its $(\lfloor d/2 \rfloor + 1)$-skeleton. Chapter 2, [Row23], studies CAT(0) cubical complexes using combinatorial tools developed by Ardila, Owen, and Sullivant, and defines a flag simplicial complex that captures much of the combinatorial information of a given CAT(0) cubical complex. Chapter 3, [Row22b], investigates the topological properties that can be deduced about a polytopal complex from the topological properties of complements of its skeletons, with applications to cubical complexes.

0.1 What is a cubical complex?

There are several common themes among these three papers — e.g. polytopal complexes, skeletons, homology, the nerve theorem — but the strongest common thread is cubical complexes. So, let us begin by asking: what is a cubical complex? Informally, cubical complexes are geometric structures built from cubes (i.e. hypercubes, $[0,1]^d$ for some d), just like how simplicial complexes are built from simplices. This seemingly straightforward idea turns out to be surprisingly fiddly to define: depending on how you adapt the definition of a simplicial complex, there are at least three non-equivalent definitions of a cubical complex you can reasonably write down, which are all relevant to different parts of this thesis.

First, let us start with simplicial complexes. A simplicial complex Δ is a finite, non-empty collection of sets that is closed under taking subsets: that is, if $\sigma \in \Delta$ and $\tau \subseteq \sigma$, then $\tau \in \Delta$.

If we view Δ as a poset, with its component sets σ ordered by inclusion, this definition says that the downset generated by an element σ (i.e. the subsets of σ that are contained in Δ) should be order-isomorphic to a boolean lattice (i.e. the collection of all subsets of σ). Boolean lattices are also isomorphic to the poset of faces of a simplex: this, then, is something we can modify to define a cubical complex.

Definition 0.1.1. A *cubical complex (viewpoint 1)* is a finite collection $\square$ of sets, such that for any $\sigma \in \square$, the poset $\{\tau \in \square : \tau \subseteq \sigma\}$ is order-isomorphic to the poset of faces of a cube, and if σ and τ are in $\square$ then so is $\sigma \cap \tau$.

This is the definition used in Chapters 1 and 2, and it is also the viewpoint taken in e.g. [BBC97; EH95; Het95; JR07].[1]

This definition is very combinatorial and abstract, and lends itself to poset-theoretic approaches, like the derivative complex defined in [BBC97] and the exploration of Möbius functions in [Ath12].

However, there is one significant drawback: every simplicial complex can be realised as a collection of geometric (convex) simplices in $\mathbb{R}^N$ for some N, but the analagous cubical statement is not true for cubical complexes defined this way. A simple example, observed

[1] It is also possible to omit the "family of sets" part of this definition, and define a cubical complex to be a poset where every principle downset is isomorphic to the face poset of a cube. This definition allows cubes to intersect in non-trivial ways, and it is perhaps more accurate to call this a variant of a simplicial *poset* rather than a simplicial complex.

by Hetyei [Het95], is the Möbius strip formed by glueing the ends of a 1×3 rectangle with a half-twist. However, while cubical complexes under this definition cannot always be realised with geometric, convex cubes, they can be realised with topological cubes as a CW complex. (In fact, many authors — e.g. [AOS12; HW08; BH99; Hag12; BB98] — also define cubical complexes to be CW complexes built from cubes in various ways, which is a viewpoint well suited to topological and metric investigations.)

This motivates the second of our promised three definitions of a cubical complex:

Definition 0.1.2. A *polytopal complex* X is a collection of polytopes in $\mathbb{R}^N$ for some N such that every face of a polytope $\sigma \in X$ is also in X, and the intersection of any polytopes in X is a face of each (perhaps empty).

A *cubical complex (viewpoint 2)* is a polytopal complex where every polytope is (combinatorially equivalent to) a cube.

We use this definition in Chapter 3, and this viewpoint also appears in [AB20; SZ04; ABY14]. It is well suited to geometric, polytopal approaches. As Hetyei's Möbius strip example demonstrates, this definition is strictly stronger than the first.

Being realisable with polytopes means that the combinatorics is a little nicer than in the first definition; however, things are even nicer under the third definition we will give.

One property in the simplicial world is that every simplicial complex on n vertices embeds into the simplex on n vertices. This is far from true in either definition of cubical complexes we have seen so far: for example, the graph (i.e. 1-dimensional complex) $K_{2,3}$ does not embed as a subcomplex into any cube. Our third definition, again strictly stronger than the previous one, addresses this:

Definition 0.1.3. A *cubical complex (viewpoint 3)* is a subcomplex of a cube — that is, a set of faces of an n-cube for some n that is closed under taking sub-faces.

We do not use this definition explicitly in this thesis, but it runs implicitly through most of it: most results in Chapter 1 only apply to cubical complexes of this type, and the bulk of the results on cubical complexes in Chapters 2 and 3 are about CAT(0) cubical complexes, which are also always of this type. This definition, or the equivalent definition that replaces the n-cube with the standard cubical grid in $\mathbb{R}^n$, is used in [BH72; KMM04; DKM11] (and indeed [Wik23]). See Section 1.1.2 for a more in-depth discussion on the differences between this and the other definitions.

Among the combinatorial benefits of this definition is that the (non-empty) faces of the n-cube come with a natural labelling with vectors in $\{0, 1, *\}^n$, so any subcomplex of an n-cube inherits this labelling, making combinatorial statements easier to make and prove — see e.g. [DKM11; EH95]. Also, the hyperplanes (see Section 2.5) of cubical complexes under this definition are much tamer, and do not display any of the complicated self-intersections exhibited in e.g. [SZ04]. And although definitions 1 and 3 are at opposite ends of the generality spectrum, they both share a distinct poset flavour — they are both characterised in terms of downsets and the face poset of a cube — so this third definition is also well suited to poset tools.

0.2 Comparison with simplicial complexes

The connections between simplicial and cubical complexes extend past their similar definitions: much research into cubical complexes has aimed to reproduce ideas about simplicial complexes. For example, mathematicians have defined and studied cubical analogues of simplicial homology [EH95; KMM04], the Stanley–Reisner ring [Het95], neighbourly polytopes and spheres [BBC97; JR07; JZ00], h-vectors [Adi96; Kle11; Het96] simplicial spanning trees [DKM11] and bistellar flips [Fun99; AL20].

While some of these investigations have had great success, others have found that cubical results can be harder to pin down than their simplicial counterparts. Many of their results, although valuable contributions to mathematics, have caveats or limitations that aren't needed in the simplicial world: the cubical generalised lower bound conjecture has only been verified up to dimension 4 [Kle11]; not all PL-homeomorphic cubical manifolds are related by bistellar flips [Fun99; AL20]; the combinatorial definitions of a cubical h-vector by [Adi96] don't match the algebraic constructions by [Het95; Het96].

These difficulties with cubical complexes have a couple of possible explanations. Cubes are in a sense more "rigid" than simplicial complexes: in a simplicial complex, you can often safely make small, local changes (e.g. geometrically you can "wiggle" a vertex about without affecting the overall structure), but cubical complexes often have structural obstacles to work within (e.g. there are parity restrictions on the face numbers in certain classes of cubical complex — see [BC00] — which cause obstructions for bistellar flips [Fun99]). Cubes are also "bigger" than simplices: for example, the number of vertices of a d-dimensional simplex (i.e. $d+1$) is linear in d, but for a d-cube, the number of vertices (i.e. 2^d) is exponential.

One other explanation for the difference in mathematics's understanding of simplicial and cubical complexes is that there has simply been much less cubical research done. At the time of writing, arXiv contains 2095 papers containing the phrase "simplicial complex", but only 183 containing "cubical complex". The Google Books Ngram viewer [Goo23] shows that the usage of "simplicial complex" has increased almost linearly since 1930, but "cubical complex" only peaked about 15% as high as "simplicial complex" in its heyday around 1960.[2] It is clear that there is plenty of knowledge left to be uncovered in the world of cubical complexes.

0.3 This thesis

The work in this thesis continues in this tradition of using simplicial ideas to study cubical complexes.

In Chapter 1, we take Dancis's proof that simplicial d-manifolds are determined by their $(\lfloor d/2 \rfloor + 1)$-skeleton [Dan84], and adapt the homological proof techniques to apply to cubical complexes. Our results do come with a caveat: they only work for subcomplexes of cubes — although the discussions above and in Chapter 1 indicate that this class of cubical complexes is still a reasonable one to consider.

In Chapter 2, we study CAT(0) cubical complexes, and introduce a flag simplicial

[2] These statistics should be taken with a grain of salt — there are several other names used for cubical complexes, including "cube complex", "cubical polyhedron", "cubication", "quadrangulation", etc.

complex called the "crossing complex" associated to each CAT(0) cubical complex. Using the bijection between (rooted) CAT(0) cubical complexes and posets with inconsistent pairs developed by [AOS12; ABY14], we discover that much combinatorial information in the cubical complex (e.g. face numbers, connectivity, balancedness) is echoed in the comparatively simpler simplicial complex.

In Chapter 3, we investigate complements of skeletons from a topological perspective. Much of this chapter builds tools that apply to any polytopal complex — in particular, we derive a long exact sequence relating skeleton complements and links — but the chapter concludes by considering cubical complexes, especially CAT(0) ones, and investigating topological properties coming from the simplicial world. The capstone result shows that several topological properties are true for a CAT(0) cubical complex if and only if they are true of its crossing complex.

Chapter 1

Reconstructing d-manifold subcomplexes of cubes from their $(\lfloor d/2 \rfloor + 1)$-skeletons

1.1 Introduction

1.1.1 Polytope reconstruction

The problem of reconstructing polytopes or related structures from partial information has seen much interest. To reconstruct an arbitrary d-dimensional polytope, for example, one needs to know its $(d-2)$-skeleton; however, some special classes of polytopes are determined by much less information: for example, zonotopes and simple d-dimensional polytopes are determined by their graphs, i.e. their 1-skeletons. For a survey of this field, see the excellent paper [Bay18].

Simplicial polytopes, and more generally simplicial spheres and manifolds, lie in between these two extremes. In [Dan84], Dancis presents a neat homological proof that d-dimensional simplicial manifolds are determined by their $\left(\lfloor \frac{d}{2} \rfloor + 1\right)$-skeleton, generalising an idea for simplicial spheres attributed to Perles. In this chapter, we modify Dancis's argument to prove a similar result for a class of cubical manifolds, namely those that can be embedded as cubical complexes into the standard cube I^n for some (potentially large) n.

To reconstruct the $(k+1)$-skeleton of a manifold (simplicial or cubical) from its k-skeleton, we need to determine what the $(k+1)$-dimensional faces are. If F is a $(k+1)$-dimensional face, its boundary is a subcomplex of the k-skeleton isomorphic to the boundary of a simplex or cube. Therefore, determining the $(k+1)$-dimensional faces of the complex amounts to answering the following question: if S^k is a subcomplex of the k-skeleton that is isomorphic to the boundary of a $(k+1)$-dimensional simplex or cube, is it actually the boundary of a $(k+1)$-dimensional face of the complex, or did it just show up by coincidence? Dancis's key insight was that this question can be answered with knowledge of just the k-skeleton, if k is large enough compared to d, thanks to Poincaré duality. In this chapter, we adapt Dancis's argument to cubical manifolds that are embeddable in cubes. Our main results are Theorem 1.3.8, Corollary 1.3.10, and Theorem 1.3.12.

1.1.2 On subcomplexes of the cube

At first glance, restricting our attention to complexes that can be embedded as a subcomplex of a cube may seem like a strict and arbitrary condition. However, we present three reasons why it is natural to consider these types of complexes.

First, there is the analogy to the simplicial world. Every simplicial complex is a subcomplex of a simplex: a simplicial complex is a collection of *some* subsets of $\{1, \ldots, n\}$, while the $(n-1)$-simplex is the collection of *all* subsets of $\{1, \ldots, n\}$.

Second, cubical complexes often arise in applied settings as subcomplexes of the standard grid of unit cubes in $\mathbb{R}^n$ (see [KMM04]). Blass and Holsztyński [BH72, Theorem 2.7] prove that any finite subcomplex of this grid can be embedded in some cube I^N (with N potentially much larger than n).

Third, many common families of cubical complexes can be embedded in the cube. This list includes boundaries of cubical zonotopes [Zie95, Corollary 7.17], stacked cubical polytopes (also called "capped"), the neighbourly cubical spheres constructed by [BBC97], CAT(0) complexes [AOS12, p. 149], and cubical barycentric subdivisions of simplicial complexes [BH72, Theorem 1.1]. Similarly, many operations on cubical complexes preserve the property of being a subcomplex of a cube, like product and disjoint union, cubical fissuring as defined by [BBC97], and certain cubical versions of Pachner's bistellar moves considered by [Fun99, Proposition 3.7.3].

On the other hand, we are obliged to point out that there are many cubical complexes that cannot be embedded into any cube. Fortunately, though, there is a fairly straightforward criterion for determining whether a given cubical complex can be embedded in a cube, for which we will need the following under-appreciated lemma.

Lemma 1.1.1 ([EH95, Lemma 12]). *If $f : V(I^m) \to V(I^n)$ is an injection from the vertices of the cube I^m to vertices of I^n which maps edges to edges, that is, an embedding of the underlying graph $\mathcal{G}(I^m)$ into $\mathcal{G}(I^n)$, then the image of f is a face of I^n.*

Corollary 1.1.2. *A cubical complex $\square$ embeds into I^n if and only if its underlying graph $\mathcal{G}(\square)$ embeds into $\mathcal{G}(I^n)$.*

So the question of whether a complex can embed into a cube is a graph-theoretic question — which is answered by the following theorem:

Theorem 1.1.3 ([HM72, Proposition 1]). *A connected graph $\mathcal{H}$ can be embedded in $\mathcal{G}(I^n)$ if and only if there exists a labelling of the edges of $\mathcal{H}$ with labels $\{1, \ldots, n\}$ such that:*

- *in every path in $\mathcal{H}$, some label appears an odd number of times, and*
- *in every cycle in $\mathcal{H}$, no label appears an odd number of times.*

Note that this criterion is similar to a condition given by [DSS94, Theorem 2] for a cubical manifold to be embeddable in the standard grid of unit cubes in $\mathbb{R}^n$.

1.2 Definitions

Cubical complexes are a variation on the well-studied notion of simplicial complexes, where instead of simplices, we use hypercubes (or "cubes" for short).

Definition 1.2.1. A *cubical complex* consists of a finite *vertex set* V and a collection $\square$ of subsets of V, called *faces*, satisfying the following conditions:

- $\varnothing$ is *not* in $\square$;
- for each $v \in V$, $\{v\}$ is in $\square$;
- for each $F \in \square$, the set $\widehat{F} \coloneqq \{G \in \square : G \subseteq F\}$ is isomorphic (as a poset ordered by inclusion) to the poset of non-empty faces of a cube; and
- if F and G are in $\square$, $F \cap G$ is either empty or in $\square$.

We usually abuse notation and just write $\square$ (pronounced "square") to denote the cubical complex, instead of $(V, \square)$. We will typically use capital Greek letters with horizontal and vertical lines to denote cubical complexes, e.g. $\square, \Gamma, \Pi, \Xi, \dots$, and other capital Greek letters for simplicial complexes, e.g. $\Delta, \Lambda, \dots$.

The choice of whether to include or exclude the empty set in the definition of $\square$ is arbitrary, and many authors use the opposite convention, for which the word "non-empty" should also be removed from the third condition. We make the choice to exclude the empty set because the set of non-empty faces of a cube is slightly simpler to describe than the set of all faces.

This definition should seem reminiscent of the definition of a simplicial complex. The only differences are the exclusion of the empty set, the fourth condition (which is redundant in the simplicial version), and "non-empty faces of a cube" should be replaced by "faces of a simplex" in the third condition: the face poset of a simplex is a Boolean lattice, isomorphic to the power set of $\{1, \dots, n\}$.

Much of the terminology for simplicial complexes carries over to cubical complexes as well. The *dimension* of a face F, denoted $\dim F$, is the dimension of the cube whose face poset $\widehat{F}$ is isomorphic to, and the dimension of a complex is the maximum of the dimensions of its faces. The *k-skeleton* of $\square$ is the complex $\operatorname{Skel}_k \square$ whose faces are $\{F \in \square : \dim F \leq k\}$. The 1-skeleton of $\square$ is also called its *graph*, $\mathcal{G}(\square)$, and the vertex set is sometimes denoted $V(\square)$.

If Π and Ξ are any cubical complexes, define $\Pi \times \Xi$ to be the complex with vertex set $V(\Pi) \times V(\Xi)$ and face set $\{f \times g : f \in \Pi, g \in \Xi\}$.

The prototypical example of a cubical complex is the *unit interval* $I := \big\{\{0\}, \{1\}, \{0,1\}\big\}$. This lets us write the poset of non-empty faces of the n-dimensional cube as $I^n = I \times \cdots \times I$. Define I^0 to be the complex with a single vertex. We will sometimes use the alternative notation $I := \{0, 1, *\}$, where $0 \preceq *$ and $1 \preceq *$ but 0 and 1 are incomparable. In this notation, I^n has vertex set $\{0,1\}^n$ and face set $\{0,1,*\}^n$, with the partial order where $(p_1, \ldots, p_n) \preceq^n (q_1, \ldots, q_n)$ if and only if $p_i \preceq q_i$ for all i. The *boundary* of a cube, denoted ∂I^n, is the complex with face set $I^n \setminus \{(*, \ldots, *)\}$.

A *map of cubical complexes* is a function $f : V(\Pi) \to V(\Xi)$ such that the image of any face of Π is a face of Ξ. If f is injective, the map is called an *embedding*; if f is bijective and f^{-1} is also a map of cubical complexes, we say f is an *isomorphism*, and we write $\Pi \cong \Xi$. A *subcomplex* of $\square$ consists of a subset $W \subseteq V$ and a subset $\Gamma \subseteq \square$ such that (W, Γ) is a cubical complex. In this situation, define $\square \setminus \Gamma$ to be the complex whose face set is $\{F \in \square : F \cap V(\Gamma) = \varnothing\}$; in other words, $\square \setminus \Gamma$ is the complex obtained by deleting all vertices of Γ from $\square$, along with all faces containing them. We say that Γ is a *full* subcomplex of $\square$ (sometimes called an "induced subcomplex") if it contains every face of $\square$ involving only vertices in W.

Any cubical complex can be given a geometric realisation as a CW complex. If $\square$ is a cubical complex, construct the geometric realisation $|\square|$ by taking the disjoint union of one copy of $[0,1]^k$ for each k-dimensional face of $\square$, identifying the subfaces $f \subseteq F$ with subfaces of the associated cube in agreement with the cubical structure, and for each inclusion $f \subseteq F$ of faces, glueing the cubes for f and F along the appropriate face. The geometric realisation of the complex I^n can thus be identified with the usual cube $[0,1]^n \subseteq \mathbb{R}^n$. Note that $|\Pi \times \Xi|$ is homeomorphic to $|\Pi| \times |\Xi|$.

Using the CW structure of a cubical complex, we can compute its homology and cohomology groups using cellular homology and cohomology, denoted $H_i(\square)$ and $H^i(\square)$ respectively. Crucially for this chapter, the cellular homology groups of a CW complex agree with the singular homology of the complex as a topological space — that is, $H_i(\square) \cong H_i(|\square|)$ — and the ith cellular homology can be computed from just the j-skeleton when $i < j$ — that is, $H_i(\square) \cong H_i(\operatorname{Skel}_j \square)$ — and similar statements for cohomology [Hat02, Theorem 2.35, Lemma 2.34(c) and Theorem 3.5]. In this chapter, if no coefficients are written, assume that all coefficients of homology are $\mathbb{Z}/2\mathbb{Z}$ (although many of the results can be easily modified for any sensible choices of coefficients), and we only consider non-reduced homology.

We say that a cubical complex M of dimension d is a *cubical homology manifold* if $|\mathrm{M}|$ is homeomorphic to a homology manifold (without boundary), that is, for every point $p \in |\mathrm{M}|$, the relative homology $H_i(|\mathrm{M}|, |\mathrm{M}| \setminus p; \mathbb{Z})$ is $\mathbb{Z}$ if $i = d$ and trivial otherwise.

1.3 The main result

The argument in Dancis's paper begins with some preliminary facts about simplicial complexes, then proceeds to a mostly topological proof. The topological reasoning applies to cubical complexes without much modification, but the beginning of the argument needs some changes before it will work in the cubical world.

1.3.1 Cubical modifications

First, Dancis uses the following lemma:

Lemma 1.3.1 (see [Mun84a, lemma 70.1])**.** *If Δ is a simplicial complex and Λ is a full (sometimes called "induced") subcomplex, then $|\Delta| \setminus |\Lambda|$ and $|\Delta \setminus \Lambda|$ are homotopy equivalent.*

To make this lemma work in the cubical world, we need a suitable analogue of fullness.

Definition 1.3.2. Let $\square$ be a cubical complex. A subcomplex $\Gamma \subseteq \square$ is *face-like* if for every face $F \in \square$, the intersection $V(\Gamma) \cap F$ is empty or a face of Γ.

For example, for any face G of a cubical complex, the subcomplex $\widehat{G}$ induced by G is a face-like subcomplex (hence the name).

Figure 1.1: An example of the deformation retraction in Proposition 1.3.3. The subcomplex Γ to be deleted is shown with dashed lines in Fig. 1.1a.

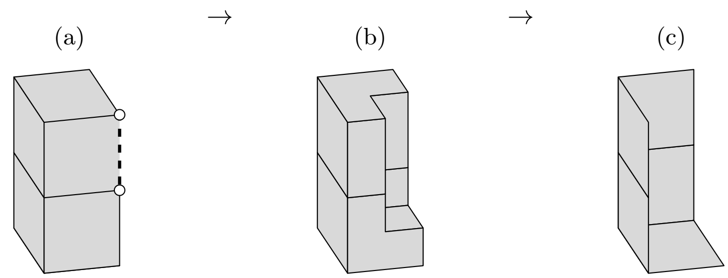

The analogous condition for simplicial complexes is equivalent to fullness. For cubical complexes, face-like-ness implies fullness, but the reverse implication is not true: for example, if $\square$ is a solid square and Γ is a pair of diagonally opposite vertices, then Γ is a full subcomplex but not face-like, since the intersection of Γ with the whole square is not a face.

Proposition 1.3.3. *Let $\square$ be a cubical complex and $\Gamma \subseteq \square$ a face-like subcomplex. Then $|\square| \setminus |\Gamma|$ is homotopy equivalent to $|\square \setminus \Gamma|$.*

Proof. We will prove the statement by first constructing a strong deformation retraction H_F from $|\widehat{F}| \setminus |\Gamma \cap \widehat{F}|$ to $|\widehat{F} \setminus (\Gamma \cap \widehat{F})|$ for each face F of $\square$, and then observing that wherever faces F and F' overlap, the deformation retractions H_F and $H_{F'}$ agree.

Let $G = V(\Gamma) \cap F$, so $\widehat{G} = \Gamma \cap \widehat{F}$. Since Γ is face-like, G is either empty or a face of Γ, and thus of F. If G is empty or $G = F$, then $|\widehat{F}| \setminus |\widehat{G}| = |\widehat{F} \setminus \widehat{G}|$, so we can take H_F to be the constant homotopy. Otherwise, by exploiting the bountiful symmetry of the cube, we may write $\widehat{F} = I^r$ and assume $\widehat{G}$ is the subcomplex $I^k \times 0^{r-k}$. Then

$$\begin{aligned} |\widehat{F}| \setminus |\widehat{G}| &= |I^r| \setminus |I^k \times 0^{r-k}| \\ &= |I^k| \times (|I^{r-k}| \setminus |0|), \end{aligned}$$

and

$$
\begin{aligned}
|\widehat{F} \setminus \widehat{G}| &= |I^r \setminus (I^k \times 0^{r-k})| \\
&= |I^k| \times |I^{r-k} \setminus 0|.
\end{aligned}
$$

The space $|I^{r-k}| \setminus |0|$ is a topological cube with a single corner point removed, while $|I^{r-k} \setminus 0|$ is a cube where every face containing that corner is removed, which is the subset of the cube consisting of points where some coordinate is 1.

Now, define a strong deformation retraction H_F from $|I^k| \times (|I^{r-k}| \setminus |0|)$ to $|I^k| \times |I^{r-k} \setminus 0|$ by

$$
H_F\big(t,(x,y)\big) \coloneqq \left(x,\, (1-t)y + \frac{ty}{\displaystyle\max_{i=1,\ldots,r-k} y_i}\right).
$$

Note that $H_F\big(0,(x,y)\big) = (x,y)$ and $H_F\big(1,(x,y)\big) = \big(x, \frac{y}{\max y_i}\big)$, which is in $|I^{r-k} \setminus 0|$ since some coordinate of $\frac{y}{\max y_i}$ is 1. We leave it to the reader to check the remaining details that this is a well defined deformation retraction.

Now that we have constructed homotopies on the faces F of $\square$, it remains to argue that they agree where faces overlap. By the definition of a cubical complex, the intersection of any two overlapping faces F and F' is a sub-face f of each. Think of $\widehat{F}$ as $I^r = \{0,1,*\}^r$, so the face F is $(*,\ldots,*)$, and G is $(*,\ldots,*,0,\ldots,0)$, with $*$ appearing k times. Thus the set of subfaces of F that meet Γ is $\{0,1,*\}^k \times \{0,*\}^{r-k}$.

Suppose f is one such face, so all of the last $r-k$ coordinates of f are 0 or $*$, and let ι be the inclusion $|\widehat{f}| \hookrightarrow |\widehat{F}|$. The homotopy H_F restricted to $|\widehat{f}|$ is constantly zero in the coordinates where f is 0, so $H_F|_{|\widehat{f}|}$ agrees with $\iota \circ H_f$. On the other hand, if f is a subface of F that does not meet Γ, at least one of the last $(r-k)$ coordinates of f must be 1, so

$$
\begin{aligned}
(1-t)y + \frac{ty}{\max_i y_i} &= (1-t)y + \frac{ty}{1} \\
&= y.
\end{aligned}
$$

Thus H_f and the restriction of H_F are both the constant homotopy on $|\widehat{f}|$, so they also agree.

Therefore, for any faces F and F', H_F and $H_{F'}$ agree on their intersection $F \cap F' = f$. Thus the homotopies glue together to form a deformation retraction from $|\square| \setminus |\Gamma|$ to $|\square \setminus \Gamma|$. ∎

The second modification we make to Dancis's argument is the following lemma, which explains why we only consider subcomplexes of the cube — this lemma is not true for arbitrary cubical complexes.

Lemma 1.3.4. *Suppose $\square$ is a subcomplex of I^n, and S^k is a subcomplex of $\square$ isomorphic to ∂I^{k+1}, with $k \geq 1$. Then S^k is a face-like subcomplex of $\square$ if and only if S^k is not the boundary of a $(k+1)$-dimensional face of $\square$, that is, S^k is not the complex $\partial \widehat{F}$ for any $F \in \square$.*

Proof. If S^k is the boundary of a face $F \in \square$, then $V(S^k) \cap F = V(S^k)$ is not a face of S^k, so S^k is not face-like.

On the other hand, even if S^k is not the boundary of any face of $\square$, Lemma 1.1.1 implies that it is still the boundary of a face F of I^n. Suppose $G \in I^n$ is any face. Since S^k and $\widehat{F}$ only differ in the face F, we have either $G \cap S^k = S^k$, if G contains F, or $G \cap S^k = G \cap F$, which is a face of S^k. Hence if $\square$ does not contain F, and thus does not contain any face containing F, then S^k is a face-like subcomplex of $\square$. ∎

1.3.2 Topological theorems

From this point on, Dancis's argument is mostly topological, so it carries over almost unchanged for cubical complexes. We restate Dancis's argument here, starting with the following version of Poincaré duality (see [Mun84a, Theorem 70.2]). Remember that we assume the coefficients of homology are $\mathbb{Z}/2\mathbb{Z}$ if none are written.

Theorem 1.3.5 (Poincaré duality)**.** *If M is a d-dimensional cubical homology manifold and Γ is a subcomplex, then $H_j(|\mathrm{M}|, |\mathrm{M}| \setminus |\Gamma|) \cong H^{d-j}(|\Gamma|)$. If M is orientable, the same* *holds with coefficients in $\mathbb{Z}$.*

Lemma 1.3.6 (cf. [Dan84, Lemma 6(b) and Lemma 8])**.** *Let $\mathrm{M} \subseteq I^n$ be a d-dimensional cubical homology manifold, and S^k a subcomplex of M isomorphic to ∂I^{k+1} with $k \geq 2$.*

- *If S^k is the boundary of a face in M, then $H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{M})$ for all $j \leq d-2$.*
- *If S^k is not the boundary of any face of M, then*
  - *$H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{M})$ when $j \leq d-2$ and $j \notin \{d-k, d-k-1\}$, and*
  - *$H_j(\mathrm{M} \setminus S^k) \ncong H_j(\mathrm{M})$ for some $j \in \{d-k, d-k-1\}$.*

Proof. First, suppose S^k is the boundary of a face F. Consider the long exact sequence of the pair $(\mathrm{M}, \mathrm{M} \setminus \widehat{F})$:

$$\cdots \longrightarrow H_{j+1}(\mathrm{M}, \mathrm{M} \setminus \widehat{F}) \longrightarrow H_j(\mathrm{M} \setminus \widehat{F}) \longrightarrow H_j(\mathrm{M}) \longrightarrow H_j(\mathrm{M}, \mathrm{M} \setminus \widehat{F}) \longrightarrow \cdots$$

Since $\widehat{F}$ is a face-like subcomplex, Proposition 1.3.3 and Theorem 1.3.5 imply that $H_j(\mathrm{M}, \mathrm{M} \setminus \widehat{F}) \cong H^{d-j}(\widehat{F})$, which is 0 when $j < d$. Therefore, away from $j = d$, the exact sequence breaks up into shorter exact sequences:

$$0 \longrightarrow H_j(\mathrm{M} \setminus \widehat{F}) \longrightarrow H_j(\mathrm{M}) \longrightarrow 0$$

But $\mathrm{M} \setminus \widehat{F}$ and $\mathrm{M} \setminus S^k$ are the same complex, so $H_j(\mathrm{M} \setminus S^k) = H_j(\mathrm{M} \setminus \widehat{F}) \cong H_j(\mathrm{M})$ when $j \leq d-2$.

Now, suppose S^k is not the boundary of a face. Consider the pair $(\mathrm{M}, \mathrm{M} \setminus S^k)$:

$$\cdots \longrightarrow H_{j+1}(\mathrm{M}, \mathrm{M} \setminus S^k) \longrightarrow H_j(\mathrm{M} \setminus S^k) \longrightarrow H_j(\mathrm{M}) \longrightarrow H_j(\mathrm{M}, \mathrm{M} \setminus S^k) \longrightarrow \cdots$$

In this case, Lemma 1.3.4 says that S^k is a face-like subcomplex of M, so Proposition 1.3.3 and Theorem 1.3.5 imply that $H_j(\mathrm{M}, \mathrm{M} \setminus S^k) \cong H^{d-j}(S^k)$, which is $\mathbb{Z}/2\mathbb{Z}$ when $j = d-k$ or $j = d$, and 0 otherwise. Therefore, away from $j = d-k$ and $j = d$, this exact sequence implies that $H_j(\mathrm{M}) \cong H_j(\mathrm{M} \setminus S^k)$. However, near $j = d-k$ we get the following exact sequence:

$$
\begin{aligned}
0 \longrightarrow H_{d-k}(\mathrm{M} \setminus S^k) \longrightarrow H_{d-k}(\mathrm{M}) \longrightarrow \mathbb{Z}/2\mathbb{Z} \\
\longrightarrow H_{d-k-1}(\mathrm{M} \setminus S^k) \longrightarrow H_{d-k-1}(\mathrm{M}) \longrightarrow 0 \qquad (1.1)
\end{aligned}
$$

So either $H_{d-k}(\mathrm{M} \setminus S^k) \ncong H_{d-k}(\mathrm{M})$ or $H_{d-k-1}(\mathrm{M} \setminus S^k) \ncong H_{d-k-1}(\mathrm{M})$. ■

Corollary 1.3.7 (cf. [Dan84, Main Lemma 10]). *If* $\mathrm{M} \subseteq I^n$ *is a* d*-dimensional cubical homology manifold and* $S^k \cong \partial I^{k+1}$ *a subcomplex with* $k \geq 2$*, then* S^k *is the boundary of a* $(k+1)$*-dimensional face of* M *if and only if* $H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{M})$ *for both* $j = d-k$ *and* $j = d-k-1$*.*

And since the homology groups $H_{d-k}(\mathrm{M})$, $H_{d-k-1}(\mathrm{M})$, $H_{d-k}(\mathrm{M} \setminus S^k)$ and $H_{d-k-1}(\mathrm{M} \setminus S^k)$ can be computed from the k-skeleton of M when $d-k < k$, we conclude:

Theorem 1.3.8 (cf. [Dan84, Theorem 11]). *Any* d*-dimensional cubical homology manifold* $\mathrm{M} \subseteq I^n$ *can be reconstructed from its* $\left(\lfloor \frac{d}{2} \rfloor + 1\right)$*-skeleton.*

Proof. Each $(k+1)$-dimensional face of M shows up in the k-skeleton as a subcomplex isomorphic to ∂I^{k+1}, so reconstructing the $(k+1)$-skeleton of M from its k-skeleton amounts to deciding which subcomplexes $S^k \cong \partial I^{k+1}$ are actually boundaries of faces. Corollary 1.3.7 answers this question when $k \geq \lfloor \frac{d}{2} \rfloor + 1$. ■

1.3.3 Some improvements on Theorem 1.3.8

Under some mild assumptions (for example, if M is a sphere), we can tighten this result to the $\lceil \frac{d}{2} \rceil$-skeleton (which differs from Theorem 1.3.8 when d is even).

Lemma 1.3.9 (cf. [Dan84, Lemma 12]). *Suppose* $\mathrm{M} \subseteq I^n$ *is a* d*-dimensional cubical homology manifold with* $d = 2r \geq 4$*, and either*

- $H_r(\mathrm{M}; \mathbb{Z}/2\mathbb{Z}) = 0$*, or*

- M *is orientable and* $H_r(\mathrm{M};\mathbb{Z})$ *is finite.*

Then a subcomplex $S^r \cong \partial I^{r+1}$ *bounds a face of* M *if and only if* $H_{r-1}(\mathrm{M} \setminus S^r) \cong H_{r-1}(\mathrm{M})$ *(with coefficients in* $\mathbb{Z}/2\mathbb{Z}$ *in the first case or* $\mathbb{Z}$ *in the second).*

Proof. In the first case, setting $k = r$ in Eq. (1.1) gives the following sequence:

$$0 \longrightarrow \mathbb{Z}/2\mathbb{Z} \longrightarrow H_{r-1}(\mathrm{M} \setminus S^r; \mathbb{Z}/2\mathbb{Z}) \longrightarrow H_{r-1}(\mathrm{M}; \mathbb{Z}/2\mathbb{Z}) \longrightarrow 0$$

In the second case, Poincaré duality holds with coefficients in $\mathbb{Z}$ for orientable manifolds, so we get this sequence:

$$H_r(\mathrm{M};\mathbb{Z}) \xrightarrow{\phi} \mathbb{Z} \longrightarrow H_{r-1}(\mathrm{M} \setminus S^r; \mathbb{Z}) \longrightarrow H_{r-1}(\mathrm{M}; \mathbb{Z}) \longrightarrow 0$$

in which the map ϕ must be zero, since there are no other homomorphisms from a finite group to $\mathbb{Z}$.

In both cases, we conclude that $H_{r-1}(\mathrm{M} \setminus S^r) \ncong H_{r-1}(\mathrm{M})$ (with appropriate coefficients) whenever S^r is not the boundary of a face in M. Conversely, if S^r is the boundary of a face, Lemma 1.3.6 still says that $H_{r-1}(\mathrm{M} \setminus S^r; \mathbb{Z}/2\mathbb{Z}) \cong H_{r-1}(\mathrm{M}; \mathbb{Z}/2\mathbb{Z})$, and a similar proof works with coefficients in $\mathbb{Z}$ when M is orientable. ∎

Corollary 1.3.10 (cf. [Dan84, Theorem 13])**.** *If* M *is a* $2r$*-dimensional cubical homology manifold satisfying the hypotheses of Lemma 1.3.9,* M *is determined by its* r*-skeleton.*

Theorem 1.3.8 and Corollary 1.3.10 let us reconstruct manifolds when the dimension is known. Theorem 1.3.12 will allow some ambiguity in the dimension. But first, a lemma:

Lemma 1.3.11. *If* M *is a connected,* d*-dimensional cubical homology manifold with* $d \geq 2$ *and* S^d *is a subcomplex homeomorphic to a* d*-dimensional sphere, then* $S^d = \mathrm{M}$.

Proof. The pair (M, S^d) gives the following long exact sequence:

$$\begin{array}{ccccccccc} H_{d+1}(\mathrm{M}, S^d) & \longrightarrow & H_d(S^d) & \longrightarrow & H_d(\mathrm{M}) & \longrightarrow & H_d(\mathrm{M}, S^d) & \longrightarrow & H_{d-1}(S^d) \\ \| & & \| & & \| & & \wr\| & & \| \\ 0 & & \mathbb{Z}/2\mathbb{Z} & & \mathbb{Z}/2\mathbb{Z} & & H^0(\mathrm{M} \setminus S^d) & & 0 \end{array}$$

Thus $H^0(\mathrm{M} \setminus S^d) = 0$, so $S^d = \mathrm{M}$. ■

Theorem 1.3.12 (cf. [Dan84, Theorem 14])**.** *Let* M *and* N *be cubical homology manifolds that are each subcomplexes of a cube. Let* $m = \dim \mathrm{M}$ *and* $n = \dim \mathrm{N}$*, and assume* $m \geq n \geq 3$*. Suppose* k *is an integer with*

- $k \geq \left\lfloor \frac{m}{2} \right\rfloor + 1$*; or*
- $k = \frac{m}{2}$ *and* $H_k(\mathrm{M}; \mathbb{Z}/2\mathbb{Z}) = H_k(\mathrm{N}; \mathbb{Z}/2\mathbb{Z}) = 0$*; or*
- $k = \frac{m}{2}$*,* M *and* N *are orientable, and* $H_k(\mathrm{M}; \mathbb{Z})$ *and* $H_k(\mathrm{N}; \mathbb{Z})$ *are finite.*

Then $\operatorname{Skel}_k \mathrm{M} \cong \operatorname{Skel}_k \mathrm{N}$ *implies* $\mathrm{M} \cong \mathrm{N}$.

(Our proof is a simplification of Dancis's simplicial version, but with slightly different hypotheses.)

Proof. The goal of this proof is to reduce to the case $k \geq n$. We will start by assuming $k < n$, and argue that the $(k+1)$-skeletons of M and N are isomorphic, so by induction, the n-skeletons must be isomorphic.

To begin with, let us assume $k < n$ and consider the case where $k \geq \left\lfloor \frac{m}{2} \right\rfloor + 1$. Under these assumptions,

$$n - k \leq m - k \leq m - \left\lfloor \frac{m}{2} \right\rfloor - 1 \leq k - 1 \leq n - 2 \leq m - 2.$$

The key facts are $n - k \leq k - 1$ and $m - k \leq k - 1$, which allow us to compute the $(n-k)$th, $(m-k)$th and lower homology groups from the k-skeleton, and $n - k \leq m - 2$ and $m - k \leq n - 2$, which will let us use Lemma 1.3.6.

Suppose $S^k \cong \partial I^{k+1}$ is a subcomplex of $\operatorname{Skel}_k \mathrm{M} \cong \operatorname{Skel}_k \mathrm{N}$. If S^k bounds a face in M, then Lemma 1.3.6 and the fact that $n - k \leq m - 2$ imply that $H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{M})$ when $j \in \{n-k, n-k-1\}$. Since $n - k \leq k - 1$, these homology groups depend only on the k-skeleton of M, thus $H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{N} \setminus S^k)$ and $H_j(\mathrm{M}) \cong H_j(\mathrm{N})$ when $j \in \{n-k, n-k-1\}$. Therefore, $H_j(\mathrm{N} \setminus S^k) \cong H_j(\mathrm{N})$ when $j \in \{n-k, n-k-1\}$, so Corollary 1.3.7 implies that S^k bounds a face in N.

Similarly, if S^k bounds a face in N, then $H_j(\mathrm{M} \setminus S^k) \cong H_j(\mathrm{N} \setminus S^k) \cong H_j(\mathrm{N}) \cong H_j(\mathrm{M})$ when $j \in \{m-k, m-k-1\}$; hence S^k bounds a face in M.

When $k = m/2$, the argument is similar. If $n = m$, the result follows from Corollary 1.3.10, so we may assume $n < m$. Then:

$$n - k \leq m - k - 1 = k - 1 \leq n - 2 \leq m - 2.$$

If S^k bounds a face in M, Lemma 1.3.6 and the inequality $n - k \leq m - 2$ imply that $H_j(\mathrm{N} \setminus S^k; \mathbb{Z}/2\mathbb{Z}) \cong H_j(\mathrm{M} \setminus S^k; \mathbb{Z}/2\mathbb{Z}) \cong H_j(\mathrm{M}; \mathbb{Z}/2\mathbb{Z}) \cong H_j(\mathrm{N}; \mathbb{Z}/2\mathbb{Z})$ when $j \in \{n - k, n - k - 1\}$, so S^k bounds a face in N by Corollary 1.3.7. On the other hand, if S^k bounds a face in N, Lemma 1.3.6 (or a modification with $\mathbb{Z}$ coefficients) implies $H_{k-1}(\mathrm{M} \setminus S^k) \cong H_{k-1}(\mathrm{N} \setminus S^k) \cong H_{k-1}(\mathrm{N}) \cong H_{k-1}(\mathrm{M})$ with coefficients in either $\mathbb{Z}/2\mathbb{Z}$ or $\mathbb{Z}$, and we conclude from Lemma 1.3.9 that S^k bounds a face in M.

Thus the $(k+1)$-dimensional faces in M and N are the same, in all cases when $k < n$; that is, the $(k+1)$-skeletons of M and N are isomorphic. Therefore, we can inductively increase k to reduce to the case where $k \geq n$. In this case, we have $\mathrm{Skel}_n \mathrm{M} \cong \mathrm{Skel}_n \mathrm{N} = \mathrm{N}$. However, we claim that the only skeletons of an m-dimensional manifold that are themselves manifolds are the 0-skeleton and the m-skeleton.

We may assume that M and $\mathrm{Skel}_n \mathrm{M}$ are connected, since the components of $\mathrm{Skel}_n \mathrm{M}$ are precisely the skeletons of the components of M when $n \geq 1$. If $\mathrm{Skel}_n \mathrm{M} \cong \partial I^{n+1}$, the only possible $(n+1)$-dimensional face of M is the cube whose boundary is $\mathrm{Skel}_n \mathrm{M}$. Then $\mathrm{Skel}_{n+1} \mathrm{M}$ is either ∂I^{n+1} or I^{n+1}, and there are no possible faces of higher dimension; but I^{n+1} is not a manifold (without boundary), so $\mathrm{M} \cong \partial I^{n+1} \cong \mathrm{Skel}_n \mathrm{M}$. Conversely, if $\mathrm{Skel}_n \mathrm{M} \ncong \partial I^{n+1}$, then Lemma 1.3.11 implies that there are no subcomplexes of $\mathrm{Skel}_n \mathrm{M}$ isomorphic to ∂I^{n+1}, and again $\mathrm{Skel}_n \mathrm{M} = \mathrm{M}$. Therefore, $\mathrm{M} \cong \mathrm{N}$. ■

Dancis gives an example of the cyclic $(d+1)$-polytopes, whose $\lfloor \frac{d-1}{2} \rfloor$-skeletons are the same as the skeleton of the boundary of a simplex, to show that the bounds on k in the simplicial version of Theorem 1.3.12 cannot be improved to $\lfloor \frac{d-1}{2} \rfloor$. The neighbourly cubical d-spheres and $(d+1)$-polytopes constructed by [BBC97; JZ00; JR07], whose $\lfloor \frac{d-1}{2} \rfloor$-skeletons match the boundary of a cube, serve the same purpose for cubical manifolds. (These complexes embed into a cube, since C_1 in [BBC97, Theorem 3.1] is itself a cube and fissuring preserves embeddability into a cube.)

This chapter relied on the cubical manifolds being embeddable in a cube. The question remains: what skeleton determines an arbitrary d-dimensional cubical manifold?

Chapter 2

Relating CAT(0) cubical complexes and flag simplicial complexes

2.1 Introduction

Simplicial complexes are well-known, well-studied objects used in many areas of math, from topology to optimisation, and their combinatorial aspects have been widely explored. Cubical complexes, while less well-known, are also important objects in several areas of math. In this chapter, we consider special cases of each, namely flag simplicial complexes and CAT(0) cubical complexes, and explore some of the many interesting connections between their combinatorial properties. We do this by defining the *crossing complex*, a flag simplicial complex associated to each CAT(0) cubical complex.

Flag simplicial complexes arise as the clique complex of a graph, i.e. the simplicial complex obtained by filling in every clique of a graph with a face. Much research has been done on their combinatorial structure, particularly their f-vectors and metric properties [AH16; Gal05; CD95; AB14; Fro08].

Similarly, CAT(0) cubical complexes first arose in metric space theory, where they were constructed to provide examples of non-positively curved metric spaces: for cubical complexes, the non-positive curvature condition reduces to a simple combinatorial criterion, according to a theorem by [Gro87] (see Theorem 2.2.3 below). More recently, CAT(0) cubical complexes are also commonly used in group theory and topology: researchers such as Davis, Niblo, Roller and Sageev have made discoveries involving various kinds of groups acting on CAT(0) and similar cubical complexes [DO01; NR98; Rol16; Sag95], and Agol, Haglund, Wise and others have used these ideas to make breakthroughs in low-dimensional topology [Ago13; HW08; Wis21].

In both metric geometry and group theory, the CAT(0) cubical complexes under consideration are usually infinite, and sometimes even infinite-dimensional. However, finite CAT(0) cubical complexes also have interesting structure and combinatorics, and often arise from applications, including state complexes of robotic systems [ABY14; AG04] and spaces of phylogenetic trees [BHV01].

Some previous research on the combinatorics of finite CAT(0) cubical complexes is worth noting here. Two particularly relevant examples are the work of Hagen [Hag14; Hag13; Hag12] and Ardila et al. [AOS12; ABY14; Ard19]. Hagen introduced the contact graph and crossing graph of a CAT(0) cubical complex (the latter of which will appear as the underlying graph of the crossing complex defined below), and also the simplicial boundary of a CAT(0) cubical complex; however, Hagen's results often only become interesting for infinite complexes. Ardila, Owen, and Sullivant give a useful bijection between (rooted) CAT(0) cubical complexes and a class of posets with extra structure, based on Roller and Sageev's notion of a halfspace system [Rol16; Sag95]: this bijection will form the basis for much of the work in this chapter. Other work that is less relevant to this chapter but still noteworthy includes [AB20; BC08; Bow14; CN05; Che00; Fio22], and other papers cited.

The structure of this chapter is as follows. In Section 2.2, we define the basic notions, and summarise Ardila, Owen, and Sullivant's theorem. In Section 2.3, we define the crossing complex of a finite CAT(0) cubical complex, and make some initial observations about the connections between the f-vectors of the cubical complex and its crossing complex: this section includes our main theorem, that the sets of f-vectors of flag simplicial complexes and CAT(0) cubical complexes are the same, up to an invertible linear

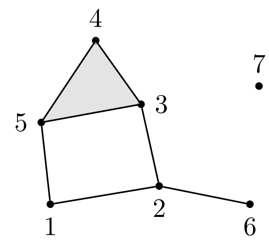

Figure 2.1: The geometric realisation of the flag simplicial complex $\Delta = \{\varnothing, \{1\}, \{2\}, \ldots, \{7\}, \{1,2\}, \{1,5\}, \{2,3\}, \{2,6\}, \{3,4\}, \{3,5\}, \{4,5\}, \{3,4,5\}\}$

transformation. Section 2.4 discusses ways of making new CAT(0) cubical complexes by combining old ones, and how this relates to the structure of the crossing complex. In Section 2.5, we take a closer look at the hyperplanes in CAT(0) cubical complexes: the major result of this section is that hyperplanes in the cubical complex correspond to vertex links in the crossing complex. Section 2.6 gives some results of a topological flavour, using the Nerve Theorem to find the crossing complex living as a subspace of the cubical complex, up to homotopy equivalence. Finally, in Section 2.7, we look at the special case of balanced complexes, and return to the f-vectors of our complexes.

While some results call upon lemmas in previous sections, most of the sections are fairly independent, and can be read in any order. The exceptions are Sections 2.2 and 2.3, which introduce important definitions and concepts, so these sections should be read first.

2.2 Definitions

We will begin by defining the two main objects of study in this chapter: flag simplicial complexes and CAT(0) cubical complexes. Readers who are familiar with simplicial complexes may wish to skip ahead to Section 2.2.2.

2.2.1 Flag simplicial complexes

Definition 2.2.1. A *simplicial complex* consists of a finite set V and a collection Δ of subsets of V, such that:

- $\varnothing \in \Delta$,

- for each $v \in V$, the set $\{v\}$ is in Δ, and

- if $\sigma \in \Delta$ and τ is any subset of σ, then $\tau \in \Delta$.

We usually abuse notation and just refer to Δ as the simplicial complex.

Elements of V are called *vertices*, and elements of Δ are called *faces*. The *dimension* of a face σ is $\dim \sigma := |\sigma| - 1$, and the dimension of Δ is the maximum dimension of its faces. Faces of dimension 0 are identified with the vertices, and faces of dimension 1 are called *edges*. Faces of Δ that are maximal under inclusion are called *facets*. If all facets of Δ have the same dimension, then Δ is *pure*. Take note that we only consider simplicial complexes with finitely many faces.

The collection of all vertices and edges of Δ forms a graph, called the *underlying graph of* Δ; conversely, any graph (which will always mean a finite, simple, undirected graph) can be thought of as a 1-dimensional (or smaller dimensional) simplicial complex.

The *link* of a face $\sigma \in \Delta$ is the following set:

$$\operatorname{link}_\Delta \sigma := \{\tau \in \Delta : \sigma \cup \tau \in \Delta \text{ and } \sigma \cap \tau = \varnothing\}.$$

We will sometimes just write $\operatorname{link} \sigma$ if the complex Δ is clear. The link is a simplicial complex. As a poset (ordered by inclusion), it is isomorphic to the *open star* of σ:

$$\operatorname{star}_\Delta \sigma := \{\rho \in \Delta : \rho \supseteq \sigma\},$$

which is not generally a simplicial complex. We also define the *closed star*:

$$\overline{\operatorname{star}}_\Delta \sigma := \{\rho \in \Delta : \rho \cup \sigma \in \Delta\},$$

which is the smallest simplicial complex containing all faces of the open star.

Any simplicial complex has a corresponding topological space, called its *geometric realisation* $\|\Delta\|$. If Δ has vertex set $\{1, \dots, n\}$, consider the standard basis $e_1, \dots, e_n$ in

$\mathbb{R}^n$. For each face $\sigma \in \Delta$, take the subset

$$\|\sigma\| \coloneqq \operatorname{conv}\{e_i : i \in \sigma\},$$

i.e. the convex hull of the basis vectors corresponding to elements of σ. The geometric realisation of Δ is then the union of these convex sets. We will often not make any distinction between the combinatorial object Δ and the associated topological space $\|\Delta\|$.

A *missing face* of Δ is a subset S of V such that S is not a face, but all proper subsets of S are faces of Δ. We say that Δ is *flag* if all missing faces of Δ have cardinality 2 (so they would be edges of Δ if they were indeed faces).

Given a graph G, there is an associated simplicial complex called the *clique complex* of G: the vertices of the clique complex are the vertices of G, and the faces are the cliques, i.e. the collections K of vertices such that every pair of vertices in K is connected by an edge in G. A clique complex is always flag; in fact, flag simplicial complexes are exactly those complexes that are clique complexes of some graph (namely the underlying graph of the complex). Similarly, the *anticlique complex* (or *independence complex*) of G is the simplicial complex whose vertices are the vertices of G, and whose faces are the independent sets, i.e. the collections of vertices with no pair connected by an edge. The anticlique complex of G is the clique complex of the complement graph, whose edges are precisely the non-edges of G.

2.2.2 CAT(0) cubical complexes

We next define a cubical complex. The definition is analogous to the definition of a simplicial complex — morally, a cubical complex is "like a simplicial complex but built from cubes instead of simplices".

Definition 2.2.2. A *cubical complex* consists of a finite set V and a collection $\square$ of subsets of V, such that:

- $\varnothing$ is *not* in $\square$,

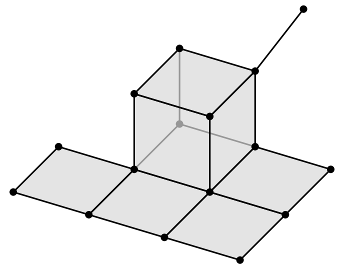

Figure 2.2: The geometric realisation of a 3-dimensional CAT(0) cubical complex

- for each $v \in V$, the set $\{v\}$ is in $\square$,
- if σ and τ are in $\square$, then $\sigma \cap \tau$ is either empty or in $\square$, and
- if $\sigma \in \square$, then the collection of elements of $\square$ that are contained in σ is isomorphic as a poset (ordered by inclusion) to the poset of non-empty faces of a cube.

This definition might appear different from the definition of a simplicial complex. However, note that the condition that $\sigma \cap \tau \in \Delta$ is redundant for simplicial complexes, and the condition that a simplicial complex is closed under taking subsets is equivalent to the requirement that for any $\sigma \in \Delta$, the collection of elements of Δ that are contained in σ is isomorphic to the face poset of a simplex. The requirement that $\varnothing \notin \square$ is a matter of preference — we could instead have required that $\varnothing \in \square$ and removed the word "non-empty" in the fourth condition. We chose the convention that $\varnothing \notin \square$ because the poset of non-empty faces of a cube is slightly simpler to describe than the poset of all faces, and because omitting the empty face makes some enumeration cleaner. Also, note that we only consider finite cubical complexes.

The dimension of $\sigma \in \square$ is $\log_2|\sigma|$ — the definition guarantees that $|\sigma|$ is a power of 2, so the dimension of σ will always be an integer. Otherwise, all the basic terminology for simplicial complexes (faces, vertices, edges, dimension of $\square$, facets, pureness, underlying graph) applies unchanged for cubical complexes.

We can also define links in cubical complexes, but we base the definition on the notion of a simplicial open star:

$$\operatorname{link}_{\square} \sigma := \{\rho \in \square : \rho \supseteq \sigma\}.$$

Unlike for simplicial complexes, a cubical link is not generally a subcomplex. As a poset (ordered by inclusion), the link of any face of a cubical complex is isomorphic to a *simplicial* complex. We also define closed stars, analogously to simplicial closed stars:

$$\overline{\operatorname{star}}_{\square}\,\sigma \coloneqq \{\rho \in \square : \rho \cup \sigma \text{ is contained in some face of } \square\}.$$

Like simplicial complexes, cubical complexes also have geometric realisations, although the construction is less clean — in general, the best we can get is a CW complex, rather than a union of convex sets in Euclidean space (although we will see that convex sets do work for CAT(0) complexes, defined below). For each i-dimensional face $\sigma \in \square$, take $\|\sigma\|$ to be an i-dimensional cube $[0,1]^i$, and identify the subfaces of σ with faces of $\|\sigma\|$. To construct $\|\square\|$, we topologically glue the cubes together along corresponding faces. However, even though this construction does not naturally live in $\mathbb{R}^n$ like for simplicial complexes, we can give a cubical complex the structure of a metric space, by endowing each cube $\|\sigma\|$ with the metric of the unit cube in $\mathbb{R}^i$, and choosing the gluings to be isometries.

An arbitrary metric space X is said to be CAT(0) if it has geodesics between any two points and it satisfies the "thin triangle" condition (see [BH99] for more details). However, we will not need the precise definition, thanks to the following theorem:

Theorem 2.2.3 ([Gro87, Section 4])**.** *A cubical complex $\square$ is CAT(0) if and only if it is simply connected and all vertex links are flag simplicial complexes. Moreover, any CAT(0) cubical complex is contractible.*[1]

This theorem gives the first inkling of the connections between flag simplicial complexes and CAT(0) cubical complexes.

One particularly useful feature of cubical complexes is their *hyperplanes*. Given any cubical complex $\square$, for each cube $[0,1]^i$ in $\square$, a subset $[0,1] \times \cdots \times \{\frac{1}{2}\} \times \cdots \times [0,1]$ is called a *midplane* of the cube. Two midplanes are *adjacent* if their intersection is a

[1] [AB20, Corollary II] go even further, and show that CAT(0) cubical complexes are *collapsible*.

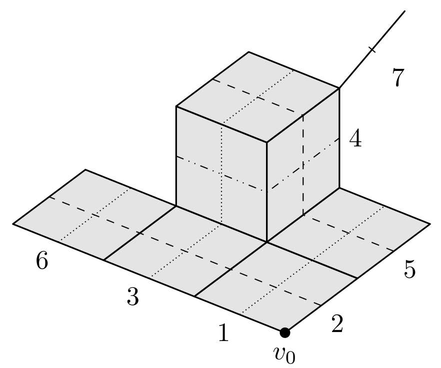

Figure 2.3: The CAT(0) cubical complex from Fig. 2.2 with its hyperplanes and a root vertex shown

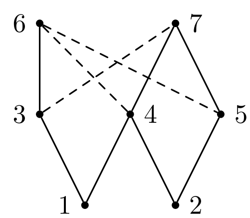

Figure 2.4: The Hasse diagram of a poset with inconsistent pairs

midplane of another cube; after taking the transitive closure of this relation, the union of all midplanes in any equivalence class is called a *hyperplane* of ❑. The midplanes give the hyperplane its own cubical complex structure. See Fig. 2.3 for an example.

One reason why hyperplanes are especially nice in CAT(0) cubical complexes is the following lemma:

Lemma 2.2.4 ([NR98, Lemma 2.7])**.** *In a CAT*(0) *complex* ❑*, the complement of any hyperplane has exactly two components.*

2.2.3 Posets with inconsistent pairs

Ardila, Owen, and Sullivant [AOS12] gave another, more combinatorial description of CAT(0) cubical complexes which will be much more useful to us. But first, we need some more definitions.

Definition 2.2.5. A *poset with inconsistent pairs* (or a PIP for short) is a triple $(P, \leq, \not\leftrightarrow)$ with the following properties:

- $(P, \leq)$ is a finite partially ordered set (poset), and
- $\not\leftrightarrow$ is a relation on P such that:
  - $a \not\leftrightarrow a$ is false for all a (that is, $\not\leftrightarrow$ is antireflexive),
  - $a \not\leftrightarrow b$ implies $b \not\leftrightarrow a$ (that is, $\not\leftrightarrow$ is symmetric), and
  - if $a \not\leftrightarrow b$, $a \leq a'$ and $b \leq b'$, then $a' \not\leftrightarrow b'$ (that is, $\not\leftrightarrow$ is inherited upwards).

If $a \not\leftrightarrow b$, then we say a and b form an *inconsistent pair*. A subset $S \subseteq P$ is *consistent* if it contains no inconsistent pairs.

We can draw the Hasse diagram of a PIP by drawing the usual Hasse diagram of the poset $(P, \leq)$ with solid lines, and adding dashed lines for each *minimal* inconsistent pair, that is, each pair $c \not\leftrightarrow d$ with no other pair $c' \not\leftrightarrow d'$ where $c \geq c'$ and $d \geq d'$. For example, the PIP shown in Fig. 2.4 has four inconsistent pairs, namely $3 \not\leftrightarrow 7$, $4 \not\leftrightarrow 6$, $5 \not\leftrightarrow 6$ and $6 \not\leftrightarrow 7$.

A definition is no good without examples, and fortunately there are two large sources of them. Firstly, if P is any poset, we can think of it as a PIP that has an order relation but happens to have no inconsistent pairs. Secondly, on the other extreme, we can construct PIPs with inconsistent pairs but no order relation: if G is any (finite, simple, undirected) graph, we can think of it as a PIP whose underlying set is the set of vertices of G, the inconsistent pairs are the edges of G, and there are no order relations (except the trivial relation $a \leq a$ for each a).

We will use a lot of standard poset terminology when talking about PIPs. Given a subset $S \subseteq P$, an element $x \in S$ is said to be *maximal* in S (respectively *minimal* in S) if there is no other element $y \in S$ with $x \leq y$ (resp. $x \geq y$). The set of maximal elements of S is denoted $\max S$, and the set of minimal elements is $\min S$. Every element of P is greater than or equal to some minimal element, and less than or equal to some maximal one.

A *downset* (or *order ideal*) is a subset $I \subseteq P$ with the property that if $a \in I$ and $a' \leq a$ then $a' \in I$. The *downset generated by* $S \subseteq P$ is the set

$$\downarrow S := \{x \in P : x \leq s \text{ for some } s \in S\}.$$

An *upset* (or *order filter*) is defined similarly with the inequalities reversed, and the *upset generated by* $S \subseteq P$ is denoted $\uparrow S$.

Two elements a and b of P are said to be *comparable* if either $a \leq b$ or $b \leq a$. A *chain* is a subset $C \subseteq P$ such that every two elements of C are comparable; an *antichain* is a subset $A \subseteq P$ where no two elements are comparable.

We will often discuss consistent downsets and consistent antichains. These two objects are related by the following lemma:

Lemma 2.2.6. *The map $I \mapsto \max I$ is a bijection from the set of consistent downsets of P to the set of consistent antichains, with inverse map given by $A \mapsto \downarrow A$.*

Proof. Apart from the word "consistent", this is a standard result about finite posets. Since I is consistent and $\max I$ is a subset of I, $\max I$ is consistent; the fact that $\downarrow A$ is consistent follows from the upward-inherited property of inconsistent pairs. ∎

At this point, the reader may be starting to wonder what PIPs have to do with cubical complexes. The answer lies in the following construction due to [AOS12], culminating in Theorem 2.2.9 (see also [ABY14; Ard19]).

Definition 2.2.7. Given a PIP P, we define an associated cubical complex $\square_P$ as follows. The vertex set of $\square_P$ is the set of consistent downsets of P. There is a face of $\square_P$ for each pair (I, M), where I is a consistent downset and M is some subset of $\max I$: the face is denoted $C(I, M)$, and it contains vertices

$$C(I, M) := \{I \setminus N : N \subseteq M\}.$$

The result is a CAT(0) cubical complex (though this is not obvious). The face $C(I, M)$

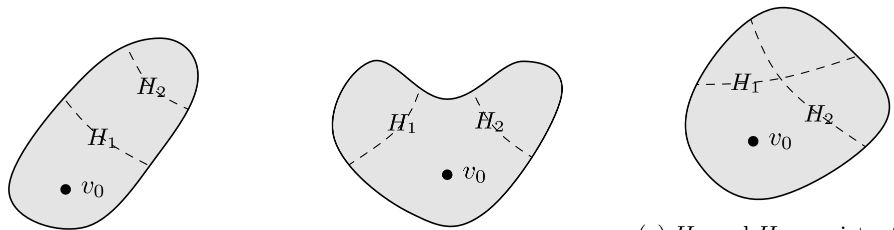

(a) $H_1 \leq H_2$ (b) $H_1 \nleftrightarrow H_2$ (c) H_1 and H_2 consistent and incomparable

Figure 2.5: How to turn hyperplanes into a PIP

contains $2^{|M|}$ vertices, one for each subset $N \subseteq M$, so it is an $|M|$-dimensional cube. Note that $\varnothing$ is always a consistent downset of any PIP, so a cubical complex constructed in this way has a natural choice of distinguished vertex, the one corresponding to $\varnothing$.

There is a natural embedding of $\square_P$ into the unit cube $[0,1]^{|P|}$: the vertices of $\square_P$ are associated with some subsets of P, whereas the vertices of $[0,1]^{|P|}$ correspond to the set of all subsets of P. The faces of $\square_P$ agree with the faces of $[0,1]^{|P|}$, so $\square_P$ is the induced subcomplex of $[0,1]^{|P|}$ generated by these vertices. Thus $\square_P$ can be realised as a union of convex sets in $\mathbb{R}^n$, faces of a cube.

Definition 2.2.8. A *rooted CAT*(0) *cubical complex* is a CAT(0) cubical complex $\square$ together with a choice of distinguished vertex v_0.

If we are given a rooted CAT(0) cubical complex, there is an associated PIP $P_\square$. The underlying set of $P_\square$ is the set of hyperplanes of $\square$. If H is a hyperplane, then the complement of H has two components, by Lemma 2.2.4: let H^- be the component that does not contain the root vertex v_0. We declare that $H_1 \leq H_2$ in $P_\square$ whenever $H_1^- \supseteq H_2^-$, and $H_1 \nleftrightarrow H_2$ when H_1^- and H_2^- are disjoint. Intuitively, "$H_1 \leq H_2$" means that hyperplane H_2 is entirely beyond H_1, to an observer standing at the root vertex, and "$H_1 \nleftrightarrow H_2$" means that this observer can see both hyperplanes without either obscuring any part of the other — see Fig. 2.5.

These constructions result in a theorem which forms the foundation for much of the rest of this chapter:

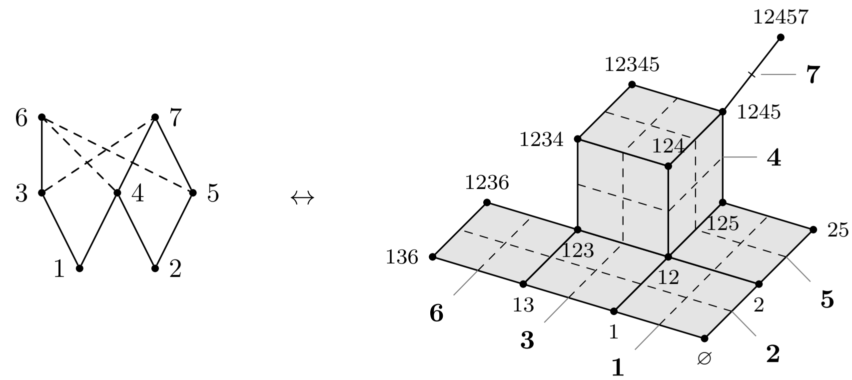

Figure 2.6: The bijection from Theorem 2.2.9, illustrated with the PIP from Fig. 2.4. For example: the set $I = \{1, 2, 3, 6\}$ is a consistent downset of P and thus a vertex of $\square_P$ (shown as "1236"), with maximal elements $\max I = \{2, 6\}$. The face $C(\{1, 2, 3, 6\}, \{2, 6\})$ is the left-most square in the figure, with vertex set $\big\{\{1, 2, 3, 6\}, \{1, 2, 3\}, \{1, 3, 6\}, \{1, 3\}\big\}$; the northwest edge of that square is the face $C(\{1, 2, 3, 6\}, \{2\}) = \big\{\{1, 2, 3, 6\}, \{1, 3, 6\}\big\}$.

Theorem 2.2.9 ([AOS12, Theorem 2.5])**.** *The map $P \mapsto \square_P$ is a bijection from the set of PIPs to the set of rooted CAT(0) cubical complexes, with inverse given by $\square \mapsto P_{\square}$.*

For example, if P is the PIP in Fig. 2.4, then $\square_P$ is the cubical complex in Figs. 2.2 and 2.3, as illustrated in Fig. 2.6. Several more examples are shown in the first two columns of Fig. 2.7.

Note that PIPs are sometimes also called "event structures". Also, a result similar to Theorem 2.2.9 was earlier proved (with very different terminology) in [BC93, Corollary 2.12].

2.3 The crossing complex

At last, we can define the main tool of this chapter: the crossing complex.

Definition 2.3.1. Given a PIP P, the *crossing complex* Δ_P is the simplicial complex whose vertex set is the underlying set of P, and whose faces are the consistent antichains of P.

For example, if P is the running example PIP shown in Figs. 2.4 and 2.6, the crossing complex Δ_P is the simplicial complex shown in Fig. 2.1. This example is reprinted in Fig. 2.7, along with several more examples.

Observe that Δ_P is a flag simplicial complex, since it is the anticlique complex of the graph whose edges are the comparable or inconsistent pairs in P. Our goal for the rest of this chapter is to demonstrate that the crossing complex Δ_P and the CAT(0) cubical complex $\square_P$ share many properties.

For a first connection, observe that hyperplanes h_1 and h_2 in $\square_P$ intersect if and only if h_1 and h_2 are consistent and incomparable elements in P. The underlying graph of Δ_P, which therefore has an edge between two hyperplanes precisely when they intersect in $\square_P$, was previously called the "crossing graph" of $\square_P$ by [Hag14] and the "transversality graph" by [Rol16].

The following fact about hyperplanes is well known:

Lemma 2.3.2 (see e.g. [Sag95, Theorem 4.14], [Hag14, Lemma 3.5])**.** *If $h_1, \ldots, h_k$ are hyperplanes of a CAT(0) cubical complex where $h_i \cap h_j \neq \varnothing$ for every i and j, then $\bigcap_{i=1}^k h_i \neq \varnothing$.*

In terms of the crossing complex, this means that the faces of Δ_P are exactly the sets of hyperplanes of $\square_P$ with a non-empty intersection. Consequently, we can compute the crossing complex of P directly from $\square_P$, and the result does not depend on the choice of root vertex of $\square_P$.

The following lemma was also essentially observed by [Hag14, Proposition 2.15] (among others), who showed that every graph is the crossing graph of some CAT(0) cubical complex. Using PIPs, we can simplify the proof.

Lemma 2.3.3. *Every flag simplicial complex is the crossing complex of some CAT(0) cubical complex.*

Proof. Recall that any graph G can be thought of as a PIP with no order relations and an inconsistent pair for each edge. In this interpretation, the consistent antichains are just

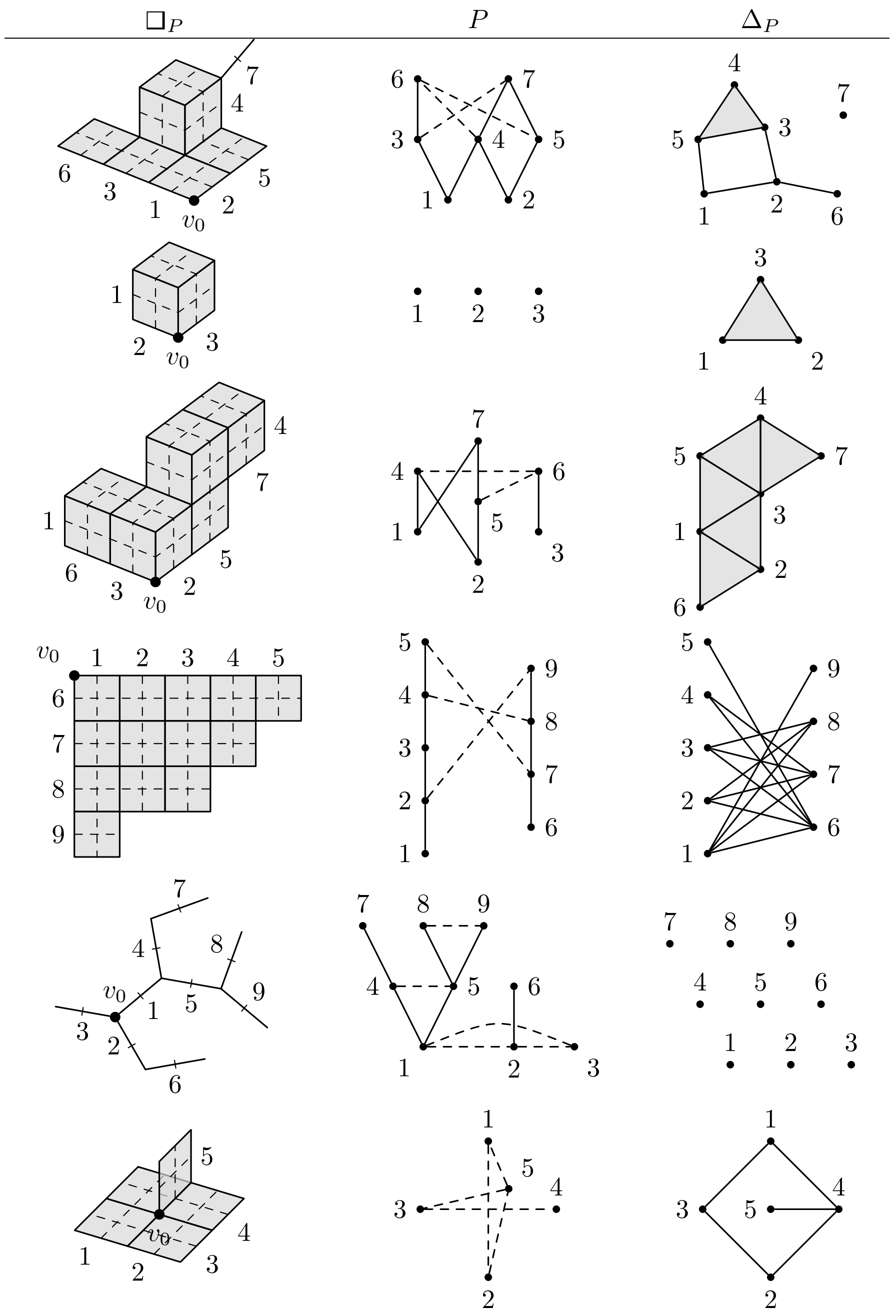

Figure 2.7: Some CAT(0) cubical complexes, their associated PIPs, and their crossing complexes

the consistent sets of vertices, i.e. the anticliques of G, and any flag simplicial complex is the anticlique complex of some graph. ■

However, note that different CAT(0) cubical complexes may have the same crossing complex — for example, any tree with n edges is a 1-dimensional CAT(0) cubical complex, and the crossing complex of such a tree always consists of n isolated points.

Next, we turn our attention to f-vectors.

Definition 2.3.4. If K is a complex (simplicial or cubical), define $f_i(K)$ to be the number of i-dimensional faces of K.

If Δ is a $(d-1)$-dimensional simplicial complex, the *f-vector* of Δ is the tuple $\big(f_{-1}(\Delta), \ldots, f_{d-1}(\Delta)\big)$, and the *$f$-polynomial* is:

$$f(\Delta, t) \coloneqq \sum_{i=0}^{d} f_{i-1}(\Delta)\, t^i.$$

If $\square$ is a d-dimensional cubical complex, its f-vector is $\big(f_0(\square), \ldots, f_d(\square)\big)$, and its f-polynomial is:

$$f(\square, t) \coloneqq \sum_{i=0}^{d} f_i(\square)\, t^i.$$

Note the difference in conventions between the simplicial and cubical cases! For instance, the constant term of $f(\Delta, t)$ is $f_{-1}(\Delta)$, which is always 1 to count the empty face; whereas the constant term of $f(\square, t)$ is $f_0(\square)$, the number of vertices of $\square$.

This brings us to the main result of this section.

Theorem 2.3.5. *Let P be a PIP. Then $f(\square_P, t) = f(\Delta_P, 1+t)$.*

Proof. Let "$A \trianglelefteq P$" mean "A is a consistent antichain of P".

By construction, the i-dimensional faces of $\square_P$ are in bijection with pairs (I, M) where $I \subseteq P$ is a consistent downset and $M \subseteq \max I$ with $|M| = i$. Lemma 2.2.6 says these pairs are in turn in bijection with pairs (A, M) where A is a consistent antichain of P and $M \subseteq A$ with $|M| = i$.

Therefore, using the convention that $f_i(K) = 0$ if $i > \dim K$,

$$\begin{aligned} f(\square_P, t) &= \sum_{i=0}^{\infty} f_i(\square_P) t^i \\ &= \sum_{i=0}^{\infty} \sum_{A \trianglelefteq P} \sum_{\substack{M \subseteq A \\ |M| = i}} t^i \\ &= \sum_{A \trianglelefteq P} \sum_{i=0}^{\infty} \binom{|A|}{i} t^i \\ &= \sum_{A \trianglelefteq P} (1+t)^{|A|} \\ &= \sum_{j=0}^{\infty} f_{j-1}(\Delta_P)(1+t)^j \\ &= f(\Delta_P, 1+t). \end{aligned}$$

■

This formula has some neat immediate consequences.

Corollary 2.3.6. $\dim \square_P = \dim \Delta_P + 1$.

Proof. The dimension of $\square_P$ is the degree of its f-polynomial, and the dimension of Δ_P is 1 less than the degree of its f-polynomial. (This fact also follows from Lemma 2.3.2.) ■

Corollary 2.3.7. *The Euler characteristic $\chi(\square_P)$ of $\square_P$ is* 1.

Proof. $\chi(\square_P) \coloneqq f(\square_P, -1) = f(\Delta_P, 0) = f_{-1}(\Delta_P) = 1$. (Of course, Theorem 2.2.3 says $\square_P$ is contractible, so we already knew this.) ■

Corollary 2.3.8. *The number of hyperplanes of $\square_P$ is $\sum_{i=0}^{d}(-1)^{i-1} i f_i(\square_P)$.*

Proof. The number of hyperplanes is $|P| = f_0(\Delta_P)$, which is the coefficient of the linear term in $f(\Delta_P, t) = f(\square_P, t-1)$. ■

But the most significant consequence of Theorem 2.3.5 comes when we combine it with Lemma 2.3.3:

Theorem 2.3.9. *The following sets are equal:*

$$\{p(t) : p \text{ is the } f\text{-polynomial of a } d\text{-dimensional } CAT(0) \text{ cubical complex}\},$$

and

$$\{q(t+1) : q \text{ is the } f\text{-polynomial of a } (d-1)\text{-dimensional flag simplicial complex}\}.$$

In other words, the sets of f-vectors of $CAT(0)$ cubical complexes and flag simplicial complexes are equal, up to the invertible linear transformation by the matrix T whose (i,j)th entry is $\binom{j-1}{i-1}$.

Proof. If $p(t)$ is the f-polynomial of some CAT(0) cubical complex $\square$, then Theorem 2.2.9 says $\square = \square_P$ for some P (after choosing a root arbitrarily), so $p(t+1)$ is the f-polynomial of Δ_P. Conversely, if $q(s)$ is the f-polynomial of some flag simplicial complex Δ, then Lemma 2.3.3 says $\Delta = \Delta_{P'}$ for some P', thus $q(s-1)$ is the f-polynomial of $\square_{P'}$.

The equivalence of the statement about f-vectors comes from this equality:

$$\begin{aligned}\sum_{i=0}^{d} f_i(\square_P)t^i &= \sum_{j=0}^{d} f_{j-1}(\Delta_P)(1+t)^j \\ &= \sum_{i=0}^{d}\sum_{j=0}^{d}\binom{j}{i} f_{j-1}(\Delta_P)t^i,\end{aligned}$$

noting that the indexing in T is shifted by 1. ■

We will revisit f-polynomials in Section 2.7.

2.4 Combining CAT(0) complexes

In this section, we aim to illustrate how the crossing complex and Ardila et al.'s bijection may be used to study $\square_P$, by examining some ways of building new complexes from old. We begin this section with three definitions.

First, a natural construction for combining cubical complexes is the product.

Definition 2.4.1. Given two cubical complexes $\square_1$ and $\square_2$, their *product* is the cubical complex $\square_1 \times \square_2$ whose vertices are pairs (v_1, v_2) with v_i a vertex of $\square_i$ for $i = 1, 2$, and whose faces are sets of the form $\sigma_1 \times \sigma_2$ for $\sigma_i \in \square_i$.

Second, there is a construction for simplicial complexes called the join.

Definition 2.4.2. If Δ_1 and Δ_2 are two simplicial complexes with vertex sets V_1 and V_2 respectively, their *join*[2] is the simplicial complex $\Delta_1 * \Delta_2$, whose vertex set is $V_1 \sqcup V_2$ and whose faces are sets of the form $\sigma_1 \sqcup \sigma_2$, where σ_i is a face of Δ_i for $i = 1, 2$.

Third, here is a way of combining PIPs.

Definition 2.4.3. If P and Q are two PIPs, define $P \underset{\leftrightarrow}{\sqcup} Q$ to be the PIP where:

- the underlying set is $P \sqcup Q$, the disjoint union of P and Q,
- if $p_1, p_2 \in P$, then the relations between p_1 and p_2 in $P \underset{\leftrightarrow}{\sqcup} Q$ are the same as in P, and similarly for $q_1, q_2 \in Q$, and
- if $p \in P$ and $q \in Q$, then p and q are incomparable and consistent in $P \underset{\leftrightarrow}{\sqcup} Q$.

In other words, $P \underset{\leftrightarrow}{\sqcup} Q$ is the PIP whose Hasse diagram is obtained by simply placing the Hasse diagrams for P and Q next to each other.

One may wonder what the connection between these constructions is — the answer is in the following lemma (also observed by [CS11, Lemma 2.5] and [Hag13, Proposition 1.3]):

Lemma 2.4.4. *$\square_P \times \square_Q \cong \square_{P \underset{\leftrightarrow}{\sqcup} Q}$, and $\Delta_{P \underset{\leftrightarrow}{\sqcup} Q} \cong \Delta_P * \Delta_Q$.*

Proof. The consistent downsets of $P \underset{\leftrightarrow}{\sqcup} Q$ are sets of the form $I \sqcup J$, where I and J are consistent downsets in P and Q respectively, and the set of maximal elements of $I \sqcup J$ is $\max I \sqcup \max J$. Therefore, vertices of $\square_{P \underset{\leftrightarrow}{\sqcup} Q}$ are in bijection with pairs (I, J) with I and J as above, which are in bijection with vertices of $\square_P \times \square_Q$; and faces of $\square_{P \underset{\leftrightarrow}{\sqcup} Q}$ are in bijection with tuples (I, M, J, N) with $M \subseteq \max I$ and $N \subseteq \max J$, which are in bijection with faces of $\square_P \times \square_Q$.

[2] Not to be confused with the join of two elements in a poset, $x \vee y$. We will not use this type of join in this chapter.

Note that the root vertices of $\square_P$ and $\square_Q$ correspond to the empty consistent downset, and $\varnothing \sqcup \varnothing = \varnothing$, so this isomorphism respects roots.

The consistent antichains of $P \underset{\leftrightarrow}{\sqcup} Q$ are sets of the form $A \sqcup B$ with A and B consistent antichains of P and Q respectively; the statement about $\Delta_{P \underset{\leftrightarrow}{\sqcup} Q}$ and $\Delta_P * \Delta_Q$ follows immediately. ∎

Here are some more constructions.

Definition 2.4.5. Suppose K_1 and K_2 are two complexes (both simplicial or both cubical, although the simplicial case will be more useful to us), with vertex sets V_1 and V_2 respectively. The disjoint union $K_1 \sqcup K_2$ of the complexes is defined to be the complex whose vertex set is $V_2 \sqcup V_2$ and whose face set is the set $K_1 \cup K_2$ (which is not quite a disjoint union if K_1 and K_2 are simplicial, since then they share the face $\varnothing$).

Definition 2.4.6. Given two rooted cubical complexes $\square_1$ and $\square_2$, the wedge sum $\square_1 \wedge \square_2$ is the cubical complex obtained by taking the disjoint union of $\square_1$ and $\square_2$ and identifying their root vertices together to make a single new root vertex.

Definition 2.4.7. Given two PIPs P and Q, the PIP $P \underset{\not\leftrightarrow}{\sqcup} Q$ is defined identically to $P \underset{\leftrightarrow}{\sqcup} Q$, except that each pair p, q with $p \in P$ and $q \in Q$ is incomparable and *inconsistent* in $P \underset{\not\leftrightarrow}{\sqcup} Q$. The Hasse diagram of $P \underset{\not\leftrightarrow}{\sqcup} Q$ is obtained by putting the Hasse diagrams of P and Q next to each other, then adding dotted lines from minimal element of P to each minimal element of Q.

Lemma 2.4.8. $\square_P \wedge \square_Q = \square_{P \underset{\not\leftrightarrow}{\sqcup} Q}$*, and* $\Delta_{P \underset{\not\leftrightarrow}{\sqcup} Q} = \Delta_P \sqcup \Delta_Q$.

Proof. There are three types of consistent downsets in $P \underset{\not\leftrightarrow}{\sqcup} Q$: they are either the empty set, a non-empty consistent downset in P, or a non-empty consistent downset in Q. The vertices of $\square_P \wedge \square_Q$ are either the root vertex (corresponding to the empty downset), or a non-root vertex in $\square_P$ or $\square_Q$ (corresponding to a non-empty downset in P or Q respectively). The faces of $\square_P \wedge \square_Q$ are in bijection with either $(\varnothing, \varnothing)$ (the root vertex) or (I, M) or (J, N) where I and J are non-empty consistent downsets of $\square_P$ and $\square_Q$ respectively, and $M \subseteq \max I$ and $N \subseteq \max J$.

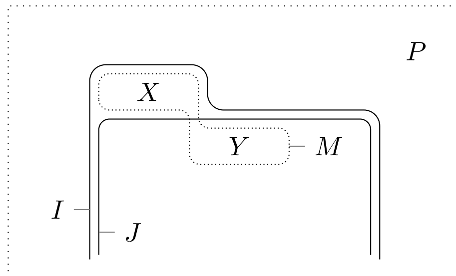

Figure 2.8: Illustration of the situation in Lemma 2.4.9

The consistent antichains of $P \underset{\leftrightsquigarrow}{\sqcup} Q$ are empty or a non-empty consistent antichain in P or Q; the faces of $\Delta_P \sqcup \Delta_Q$ are either the empty face or a non-empty face of Δ_P or Δ_Q. ■

The next lemma concerns vertex links in CAT(0) cubical complexes. Links are easier to describe and often more useful in the simplicial case — we will return to simplicial links in Section 2.5 — but for now, this lemma will be useful in proving Proposition 2.4.10.

Lemma 2.4.9. *Suppose a vertex v of $\square_P$ corresponds to the consistent downset $J \subseteq P$. Then the link of v in $\square_P$ is the crossing complex of the sub-PIP*

$$\max J \cup \min\{x \in P \setminus J : x \textit{ is consistent with all } j \in J\} \subseteq P$$

In particular, the link of the root vertex (where $J = \varnothing$) is the crossing complex of $\min P$.

Note that while no two elements of $\min P$ are ever comparable, they may be inconsistent.

Proof. Recall from Definition 2.2.7 that the faces of $\square_P$ are in bijection with pairs (I, M) where $I \subseteq P$ is a consistent downset and $M \subseteq \max I$. By definition, a face $C(I, M)$ in $\square_P$ contains v if and only if $I \setminus X = J$ for some $X \subseteq M$ — see Fig. 2.8.

If we define $Y = M \setminus X$, so $M = X \sqcup Y$ and $I = J \sqcup X$, then the pair (I, M) is determined by the choice of the pair (X, Y), and vice versa. There are some restrictions on what the sets X and Y may be — specifically, the conditions are as follows:

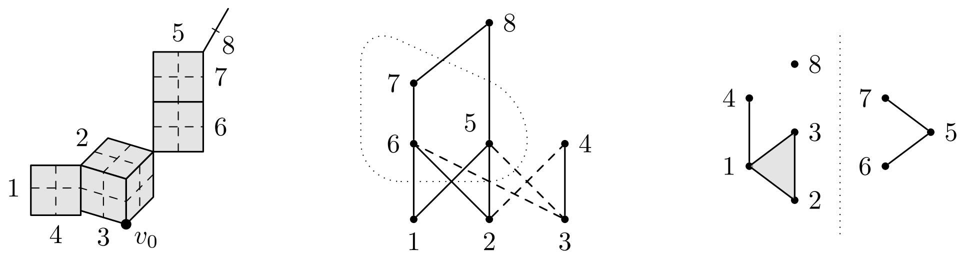

Figure 2.9: A CAT(0) cubical complex with a cut vertex, its corresponding PIP, and its crossing complex. Dotted lines show a disconnection of the crossing complex.

- $X \sqcup Y$ must be a consistent antichain (since $X \sqcup Y = M$);
- Y must be a subset of $\max J$ (since $Y \subseteq J$ and $Y \subseteq M \subseteq \max I$, and any element of J that is maximal in I must also be maximal in the subset J);
- every element of X must be consistent with every $j \in J$ (since $J \sqcup X = I$ is consistent); and
- X must be a subset of $\min(P \setminus J)$ (since $X \subseteq P \setminus J$, and the facts that X is an antichain and $J \sqcup X$ is a downset mean that any $y \in P$ with $y < x$ for some $x \in X$ must be in J), thus elements of X are in fact minimal in the subset $\{x \in P \setminus J : x \text{ is consistent with all } j \in J\}$.

Conversely, if X and Y satisfy these conditions, then $J \sqcup X$ is a consistent downset and $X \sqcup Y \subseteq \max(J \sqcup X)$. In other words, the faces of the link of v are in bijection with pairs (X, Y) such that $X \sqcup Y$ is a consistent antichain with $Y \subseteq \max J$ and $X \subseteq \min\{x \in P \setminus J : x \text{ is consistent with all } j \in J\}$. ■

With this lemma in hand, we can now give a combinatorial proof of the following result (which was also observed by [Hag13, Lemma 4.10]).

Proposition 2.4.10. *$\square_P$ has a cut vertex (i.e. a vertex v where $\|\square_P\| \setminus v$ is disconnected) if and only if Δ_P is disconnected.*

Proof. One direction follows immediately from the work above: if $\square_P$ has a cut vertex, then it is then a wedge sum of its two halves, and Lemma 2.4.8 says Δ_P is thus a disjoint union.

The other direction is less straightforward. The first part of this proof went smoothly because we could implicitly assume in Lemma 2.4.8 that the cut vertex is the root vertex of $\square_P$, since Δ_P does not depend on the choice of root; for the other direction of the proof, the cut vertex might be any vertex. The idea for the rest of the proof is to find the potential cut vertex in $\square_P$, set it to be the new root vertex, and then argue that it is indeed a cut vertex.

Suppose Δ_P is disconnected. This means we can partition the vertices of Δ_P into two non-empty sets, A and B, with no edges of Δ_P between the two sets; in terms of P, this means that every pair a, b with $a \in A$ and $b \in B$ is either inconsistent or comparable. If all such pairs are inconsistent, we can appeal to Lemma 2.4.8 again to conclude that $\square_P$ is a wedge sum; our next goal is thus to reduce to this case.

If not all pairs a, b are inconsistent, then some pair is comparable: without loss of generality, suppose $a_0 < b_0$ for some $a_0 \in A$ and $b_0 \in B$. By increasing $a_0 \in A$ and decreasing $b_0 \in B$ if necessary, we may even assume that $a_0 < b_0$ is a covering relation — that is, that there is no $x \in P$ with $a_0 < x < b_0$. Now, consider the downset $I \coloneqq (\downarrow b_0) \setminus b_0$, which is consistent since it is a subset of $\downarrow b_0$, and let v be the corresponding vertex in $\square_P$ — we will argue that v is a cut vertex.

Observe the following:

- a_0 is a maximal element in I, since b_0 covers a_0;
- b_0 is consistent with all elements of I; and
- b_0 is minimal in $P \setminus I$, and thus in $\{x \in P \setminus I : x \text{ is consistent with all } i \in I\}$.

Therefore, according to Lemma 2.4.9, both a_0 and b_0 are vertices of the link of v. Thus the vertices of $\operatorname{link} v$ can be partitioned into two *non-empty* sets, $A \cap \operatorname{link} v$ and $B \cap \operatorname{link} v$, with no edges between them, so $\operatorname{link} v$ is disconnected.

Now, choose v to be the new root vertex of $\square_P$. This new rooted CAT(0) complex will correspond to a different PIP, say P', but the geometry of $\square_{P'}$ is unchanged — in particular, the crossing complex $\Delta_{P'}$ is identical to Δ_P, and the link of v in $\square_{P'}$ is still disconnected by some partition A' and B'.

But now, since v is the root vertex of $\square_{P'}$, Lemma 2.4.9 says that $\operatorname{link}_{\square_{P'}} v = \Delta_{\min P'}$. No two elements of $\min P'$ can be comparable, so all elements of A' must instead be inconsistent with all elements of B'.

Every element $x' \in P'$ must be greater than or equal to some element of $\min P'$, but it cannot be the case that both $x' \geq a'$ and $x' \geq b'$ for some $a' \in A'$ and $b' \in B'$: if this were the case, the fact that $a' \nleftrightarrow b'$ together with the upward-inheriting property of inconsistent pairs would imply that $x' \nleftrightarrow x'$, which is forbidden. Therefore, the sets $\uparrow A'$ and $\uparrow B'$ form a setwise partition of P'. Moreover, if $x' \in \uparrow A'$ and $y' \in \uparrow B'$, then $x' \geq a'$ and $y' \geq b'$ for some $a' \in A'$ and $b' \in B'$, so since $a' \nleftrightarrow b'$, we must also have $x' \nleftrightarrow y'$; thus P' is the poset $A' \underset{\nleftrightarrow}{\sqcup} B'$. But Lemma 2.4.8 then implies that $\square_{P'}$ is a wedge sum, so $\square_P$ has a cut vertex. ∎

This proof used a lot of facts about posets — we will see an alternative, topological way to prove this result in Section 2.6.

2.4.1 Addendum

Note: This subsection does not appear in the published version of this chapter. However, there is one more fact that follows from the results above that is worth recording in this thesis.

Define a *free vertex* in $\square$ to be a vertex that is contained in exactly one facet. For example, if $\square$ is one-dimensional, it is a tree, so its free vertices are the leaf vertices.

The following result generalises the well-known fact that every non-trivial tree has at least two leaves.

Proposition 2.4.11. *If $\square_P$ is a d-dimensional CAT(0) cubical complex, then it has at least 2^d free vertices.*

Remember that we assume all complexes are finite — there are many infinite counterexamples to this Proposition.

Proof of Proposition 2.4.11. Let σ be a d-dimensional face of $\square_P$, and let $h_1, \ldots, h_d$ be the d hyperplanes passing through it. This face σ has 2^d vertices; let v be one of these vertices, and root $\square$ at v. Observe that since v is adjacent to $h_1, \ldots, h_d$, these hyperplanes are minimal elements of the corresponding PIP P, and since these hyperplanes all intersect in σ, they are consistent in P.

Let I be a consistent downset in P of maximal size such that I does not contain $h_1, \ldots, h_d$, and let w be the vertex of $\square$ corresponding to I: we will show that w is a free vertex. Since I does not contain $h_1, \ldots, h_d$, w is on the same side of the hyperplanes $h_1, \ldots, h_d$ as v, and is thus separated by at least one of these hyperplanes from each other vertex of σ; therefore, the 2^d different possible choices of v lead to 2^d distinct vertices w. (According to [ABY14, Proposition 7.4], the distance from the root vertex to the vertex corresponding to a consistent downset J is measured by $|J|$, so geometrically w is chosen to be as far away from v as you can get without crossing the hyperplanes $h_1, \ldots, h_d$.)

Now, consider the link of w in $\square$: by Lemma 2.4.9, this link is the crossing complex of the sub-PIP $\max I \cup S$, where

$$S \coloneqq \min\{x \in P \setminus I : x \text{ is consistent with every element of } I\}.$$

We claim that S is a subset of $\{h_1, \ldots, h_d\}$. Suppose x is an arbitrary element of S, so x is consistent with every element of I. If $x \geq h_i$ for some i, then by the inheritance property of consistency, h_i must also be consistent with every element of I; therefore, by minimality, x is h_i. On the other hand, if x is not greater than any h_i, then $I \cup {\downarrow} x$ is a consistent downset that does not contain $h_1, \ldots, h_d$ of strictly greater size than I, contradicting the choice of I. Thus S is a subset of $\{h_1, \ldots, h_d\}$.

Now, observe:

- Any two elements of $\max I$ are consistent, since I is consistent, and they are incomparable, since they are maximal elements of I.

- Any two elements of $S \subseteq \{h_1, \ldots, h_d\}$ are consistent and incomparable, because the hyperplanes $h_1, \ldots, h_d$ all intersect.

- Any element of $\max I$ and any element of S are consistent, by definition of S, and incomparable, since $h_1, \ldots, h_d$ are minimal elements of P that are not in I.

Therefore, the consistent antichain complex of $\max I \cup S$ — that is, the link of w — is a simplex, so w is a free vertex. There is one such vertex w for each vertex v of σ, so $\square$ contains 2^d free vertices. ■

2.5 Hyperplanes

In this section, we will take a closer look at hyperplanes in CAT(0) cubical complexes, through the lens of the *derivative complex* defined by [BC00]. This construction is heavily based on the poset structure of a cubical complex (where the faces are ordered by inclusion), so in order to use this construction, we must first say more about the poset structure of $\square_P$.

Lemma 2.5.1. *$C(I', M') \subseteq C(I, M)$ if and only if $M' \subseteq M$ and $I \setminus M \subseteq I' \subseteq I$.*

Proof. By definition, $C(I', M') \subseteq C(I, M)$ if and only if

$$\{I' \setminus N' : N' \subseteq M'\} \subseteq \{I \setminus N : N \subseteq M\}.$$

This is true if and only if it is true for the smallest and largest possible choices for N', namely $N' = \varnothing$ and $N' = M'$; therefore, $C(I', M') \subseteq C(I, M)$ if and only if

$$\{I', I' \setminus M'\} \subseteq \{I \setminus N : N \subseteq M\}.$$

Let S denote $\{I \setminus N : N \subseteq M\}$, for conciseness. The set I' is an element of S precisely when $I \setminus M \subseteq I' \subseteq I$. Similarly, $I' \setminus M'$ is an element of S precisely when

$I \setminus M \subseteq I' \setminus M' \subseteq I$; if we assume that $I' \in S$ already, then $I' \setminus M' \in S$ if and only if $M' \subseteq M$. ∎

Lemma 2.5.2. *The cubes $C(I_1, M_1)$ and $C(I_2, M_2)$ have a meet if and only if $(I_1 \setminus M_1) \cup (I_2 \setminus M_2) \subseteq I_1 \cap I_2$; if the meet exists, it is*

$$C(I_1, M_1) \cap C(I_2, M_2) = C(I_1 \cap I_2, M_1 \cap M_2).$$

Proof. The meet of $C(I_1, M_1)$ and $C(I_2, M_2)$, if it exists, is the maximal cube $C(J, N)$ such that $C(J, N) \subseteq C(I_i, M_i)$ for both $i = 1, 2$. According to Lemma 2.5.1, the set of faces satisfying this containment is the set of faces $C(J, N)$ satisfying

$$N \subseteq M_1 \cap M_2 \quad \text{and} \quad (I_1 \setminus M_1) \cup (I_2 \setminus M_2) \subseteq J \subseteq I_1 \cap I_2.$$

In order for this set to be non-empty, we need $(I_1 \setminus M_1) \cup (I_2 \setminus M_2) \subseteq I_1 \cap I_2$; if this is true, then the pair $(J, N) = (I_1 \cap I_2, M_1 \cap M_2)$ is in the set. This pair maximises (J, N) subject to the conditions that $J \subseteq I_i$ and $N \subseteq M_i$ for both $i = 1, 2$, so it must still be maximal given the extra condition $I_i \setminus M_i \subseteq J$. ∎

Now, let us state the definition of the derivative complex given by [BC00, Section 4].

Definition 2.5.3. Let $\square$ be a cubical complex. The *derivative complex* of $\square$, denoted $D\square$, is the poset where:

- the elements of $D\square$ are the sets $\{b, c\}$, where b and c are faces of $\square$ which have no meet but are both covered by the same face, and
- $\{b, c\} \preceq \{b', c'\}$ in $D\square$ if and only if $b \subseteq b'$ and $c \subseteq c'$ or $b \subseteq c'$ and $c \subseteq b'$.

See Fig. 2.10 for an example. Although the derivative complex is defined abstractly as a poset, in general it is isomorphic to the poset of faces of a cubical complex (as we define it, in terms of sets of vertices). This complex is not generally CAT(0) — it typically has many connected components. The components of $D\square$ are the hyperplanes of $\square$ (with

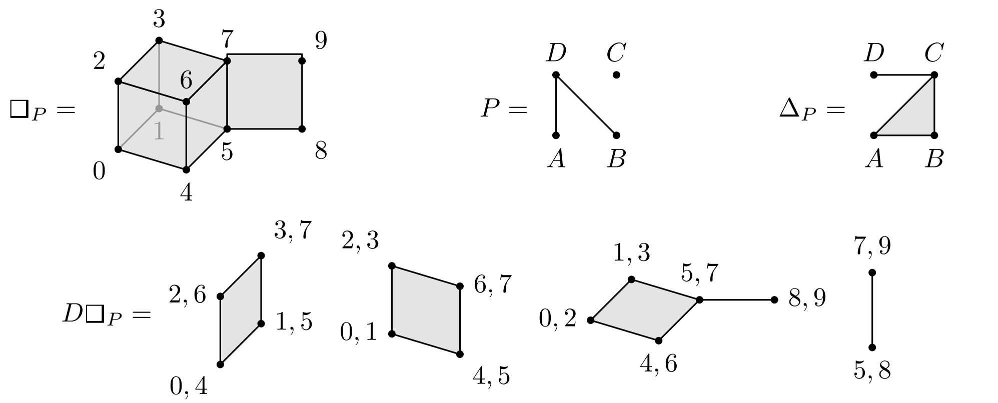

Figure 2.10: A CAT(0) cubical complex, its PIP and crossing complex, and its derivative complex

some caveats if the hyperplanes self-intersect — self-intersecting hyperplanes never occur in any subcomplex of a cube, though, so this issue does not arise for CAT(0) complexes). One reason for the name "derivative complex" is the following observation:

$$f(D\square, t) = \frac{d}{dt} f(\square, t).$$

And now we come to the main theorem for this section.

Theorem 2.5.4. *For each $x \in P$, let P_x be the sub-PIP*

$$P_x := \big\{ y \in P : y \text{ is consistent and incomparable with } x \big\} \subseteq P,$$

and let H_x be the corresponding CAT(0) cubical complex. Then $D(\square_P)$ is the disjoint union of the complexes H_x for $x \in P$.

Before we prove this theorem, let us note two of its consequences. First, since the components of $D(\square_P)$ are the hyperplanes of $\square_P$, this theorem gives a combinatorial proof of the following fact, which [NR98] observed from the metric space perspective (see also [Sag95, Theorem 4.11]):

Corollary 2.5.5. *The hyperplanes of a CAT(0) cubical complex are themselves CAT(0)* *cubical complexes.*

Second, in terms of the crossing complex, the underlying set of P_x is precisely the set of vertices of the link of x in Δ_P. Since Δ_P is flag, all links are induced subcomplexes, so we have the following corollary:

Corollary 2.5.6. *The crossing complex of the hyperplane H_x is $\Delta_{P_x} = \mathrm{link}_{\Delta_P} x$.*

We find it intriguing (and perhaps unsurprising) that links, one of the most important tools for studying simplicial complexes, correspond to hyperplanes, one of the most important tools for cubical complexes.

We conclude this section by proving the theorem.

Proof of Theorem 2.5.4. The first goal of this proof is to describe the elements of $D(\square_P)$, that is, the pairs $\{C(I_1, M_1), C(I_2, M_2)\}$ of faces that share a common cover but have no meet.

We will begin by describing the covering relations in $\square_P$. Suppose $C(I', M') \subset C(I, M)$ is a covering relation; then Lemma 2.5.1 says that $M' \subseteq M$ and $I \setminus M \subseteq I' \subseteq I$. The poset of faces of $\square_P$ is ranked by dimension, and the dimension of $C(I, M)$ is $|M|$; therefore, since ranks differ by 1 in a covering relation, the set M' must be $M \setminus x$ for some $x \in M$. By definition, I' must contain $M' = M \setminus x$; putting this together with Lemma 2.5.1, we must have

$$I \setminus x = (I \setminus M) \cup (M \setminus x) \subseteq I' \subseteq I.$$

Therefore, either $I' = I$ or $I' = I \setminus x$. Thus the covering relations in $\square_P$ take two forms: they are either

$$C(I, M \setminus x) \subset C(I, M)$$

or

$$C(I \setminus x, M \setminus x) \subset C(I, M)$$

for some $x \in M$.

Now, suppose two faces $C(J_1, N_1)$ and $C(J_2, N_2)$ of $\square_P$ share a common cover; the next question to ask is when these faces have a meet. If the common cover is $C(I, M)$, the possibilities for the two faces are:

$$\begin{aligned}
C(I, M \setminus x) \quad &\text{and} \quad C(I, M \setminus y) \quad \text{for some } x \neq y, \\
C(I \setminus x, M \setminus x) \quad &\text{and} \quad C(I \setminus y, M \setminus y) \quad \text{for some } x \neq y, \\
C(I, M \setminus x) \quad &\text{and} \quad C(I \setminus y, M \setminus y) \quad \text{for some } x \neq y, \text{ or} \\
C(I, M \setminus x) \quad &\text{and} \quad C(I \setminus x, M \setminus x) \quad \text{for some } x.
\end{aligned}$$

Lemma 2.5.2 says that $C(J_1, N_1)$ and $C(J_2, N_2)$ have a meet if and only if $(J_1 \setminus N_1) \cup (J_2 \setminus N_2) \subseteq J_1 \cap J_2$; therefore, in the four cases above, the containments we need to consider are the following:

$$\begin{aligned}
(I \setminus M) \cup \{x, y\} &\subseteq I, \\
I \setminus M &\subseteq I \setminus \{x, y\}, \\
(I \setminus M) \cup \{x\} &\subseteq I \setminus \{y\}, \quad \text{and} \\
(I \setminus M) \cup \{x\} &\not\subseteq I \setminus \{x\}
\end{aligned}$$

respectively. The containment holds in the first three cases, but fails in the fourth; therefore, the pairs of faces of $\square_P$ that have no meet but share a common cover — that is, the elements of $D(\square_P)$ — are the pairs of the form

$$C(I, M \setminus x) \quad \text{and} \quad C(I \setminus x, M \setminus x).$$

Now that we have described the underlying set of $D(\square_P)$, let us turn to its poset structure. Suppose that $F = \big\{C(I, M \setminus x), C(I \setminus x, M \setminus x)\big\}$ and $G = \big\{C(J, N \setminus y), C(J \setminus y, N \setminus y)\big\}$ are two elements of $D(\square_P)$. By the definition of $D(\square_P)$, we have $F \preceq G$ in

$D(\square_P)$ if and only if

$$C(I, M \setminus x) \subseteq C(J, N \setminus y) \quad \text{and} \quad C(I \setminus x, M \setminus x) \subseteq C(J \setminus y, N \setminus y),$$

or

$$C(I, M \setminus x) \subseteq C(J \setminus y, N \setminus y) \quad \text{and} \quad C(I \setminus x, M \setminus x) \subseteq C(J, N \setminus y)$$

in $\square_P$. According to Lemma 2.5.1, this happens if and only if

$$M \setminus x \subseteq N \setminus y, \quad J \setminus (N \setminus y) \subseteq I \subseteq J \quad \text{and} \quad J \setminus N \subseteq I \setminus x \subseteq J \setminus y,$$

or

$$M \setminus x \subseteq N \setminus y, \quad J \setminus N \subseteq I \subseteq J \setminus y \quad \text{and} \quad J \setminus (N \setminus y) \subseteq I \setminus x \subseteq J$$

in P. But the second of these two conditions is impossible: $I \subseteq J \setminus y$ implies that $y \notin I$, but $J \setminus (N \setminus y) \subseteq I \setminus x$ implies that $y \in I$. On the other hand, in the first of the two conditions, we have $J \setminus (N \setminus y) \subseteq I$, so y must still be an element of I, but $I \setminus x \subseteq J \setminus y$, so the only way this is possible is if $x = y$.

Therefore, putting this all together, we have $F \preceq G$ in $D(\square_P)$ if and only if $x = y$ and

$$M \setminus x \subseteq N \setminus x, \quad J \setminus (N \setminus x) \subseteq I \subseteq J \quad \text{and} \quad J \setminus N \subseteq I \setminus x \subseteq J \setminus x,$$

which happens if and only if $x = y$ and

$$M \subseteq N \quad \text{and} \quad J \setminus N \subseteq I \subseteq J,$$

which precisely means that $C(I, M) \subseteq C(J, N)$ in $\square_P$.

The fact that $F \preceq G$ only happens when $x = y$ means that we can partition the poset $D(\square_P)$ into disjoint, incomparable components, each determined by the choice of $x \in P$. By the preceding argument, the elements of the component corresponding to x are in order-preserving bijection with the faces $C(I, M)$ of $\square_P$ where $x \in M$.

For such a face, the requirement that $x \in M$ means the following:

- I must contain all elements of $\downarrow x$ since I is a downset;
- I must never contain any elements of $(\uparrow x) \setminus x$, since x is maximal in I; and
- I must never contain any element that is inconsistent with x.

The remaining elements of P are P_x. Therefore, I is determined by choosing a downward-closed subset J of P_x. Once J is chosen, I is determined as $I = J \cup \downarrow x$; conversely, $J = I \setminus \downarrow x$, so J is also determined by the choice of I. The maximal elements of J are $(\max I) \setminus x$, so M is determined by a choice of subset $N \subseteq \max J$, with $M = N \cup x$ and $N = M \setminus x$. Therefore, the faces $C(I, M)$ of $\square_P$ with $x \in M$ are in bijection with all faces of H_x. Moreover, this bijection is order-preserving, since

$$M' \subseteq M \quad \text{if and only if} \quad M' \setminus x \subseteq M \setminus x,$$

and

$$I \setminus M \subseteq I' \subseteq I \quad \text{if and only if} \quad (I \setminus \downarrow x) \setminus (M \setminus x) \subseteq I' \setminus \downarrow x \subseteq I \setminus \downarrow x.$$

Thus $D(\square_P)$ has one connected component for each element $x \in P$, and the component corresponding to x is isomorphic to the poset of faces of H_x. ■

Before we move on from this section, let us record a folklore result that follows from the work above. As we saw in this proof, the hyperplane of $\square_P$ corresponding to $x \in P$ is isomorphic (as a poset) to the set of faces of $\square_P$ with $x \in M$. Also, to switch tracks for a moment, recall from Definition 2.2.7 that $\square_P$ has a natural embedding into the cube $[0,1]^{|P|}$, where each face $C(I, M) \in \square_P$ is mapped to a face of $[0,1]^{|P|}$ parallel to the linear subspace spanned by the basis vectors corresponding to elements of M. Putting these two ideas together: the hyperplane of $\square_P$ corresponding to x is isomorphic as a poset to the faces of $\square_P$ that extend in the x direction inside $[0,1]^{|P|}$. In other words,

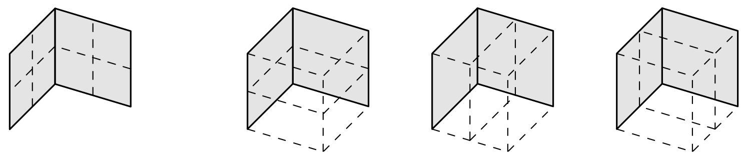

Figure 2.11: Hyperplanes in a CAT(0) complex as intersections with hyperplanes of the cube

Proposition 2.5.7. *In the standard embedding of $\square_P$ in $[0,1]^{|P|}$, the hyperplane corresponding to $x \in P$ is the intersection of $\square_P$ with the hyperplane of the cube $[0,1]^{|P|}$ perpendicular to the xth direction.*

See Fig. 2.11 for an example.

2.6 Topology

In this section, we will describe a way to find Δ_P as a subspace of $\square_P$, up to homotopy equivalence. First, we recall some general facts about topological spaces.

Definition 2.6.1. Suppose X is a topological space, and $\mathcal{Z} = \{Z_1, \ldots, Z_m\}$ is a collection of subspaces of X. The *nerve* of $\mathcal{Z}$ is the simplicial complex on vertex set $\{1, \ldots, m\}$, where a set $S \subseteq \{1, \ldots, m\}$ is a face if and only if $\bigcap_{i \in S} Z_i$ is non-empty.

For example, following the discussion on page 34, the crossing complex Δ_P is the nerve of the collection of hyperplanes of $\square_P$.

One of the most important results about nerves is the (helpfully named) Nerve Theorem:

Theorem 2.6.2 (see e.g. [Bjö95, Theorem 10.7])**.** *Suppose $\mathcal{Z} = \{Z_1, \ldots, Z_m\}$ is a collection of closed topological subspaces of a space X. If for all $I \subseteq \{1, \ldots, m\}$ the intersection $\bigcap_{i \in I} Z_i$ is contractible or empty, then the union $\bigcup_{i=1}^{m} Z_i$ is homotopy equivalent to the nerve of $\mathcal{Z}$.*

In the context of CAT(0) cubical complexes, we can apply the Nerve Theorem to conclude the following result.

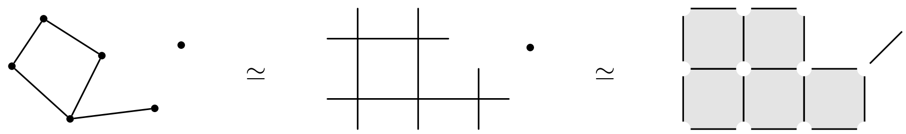

Figure 2.12: Three homotopy equivalent spaces

Theorem 2.6.3. *The following spaces are homotopy equivalent:*

- *The crossing complex* Δ_P*,*
- *The union of the hyperplanes of* $\square_P$ *(in the subspace topology inherited from* $\square_P$*), and*
- *The topological space* $\square_P$ *with all vertices removed — i.e.,* $\|\square_P\| \setminus \|V(\square_P)\|$ *(again with the subspace topology from* $\square_P$*).*

For example, see Fig. 2.12.

Proof. We want to apply the Nerve Theorem to the collection of hyperplanes of $\square_P$; however, before we can do this, we must argue that all non-empty intersections of hyperplanes are contractible.

Suppose H_1 and H_2 are two hyperplanes of $\square_P$. Proposition 2.5.7 says that H_1 and H_2 can be written as $\widetilde{H}_1 \cap \square_P$ and $\widetilde{H}_2 \cap \square_P$ respectively, where $\widetilde{H}_1$ and $\widetilde{H}_2$ are hyperplanes of the cube $[0,1]^{|P|}$. Note that $\widetilde{H}_1$ and $\widetilde{H}_2$ are themselves cubes, and the embeddings of H_1 and H_2 inside them agree with the standard CAT(0) complex embeddings of H_1 and H_2 into cubes. Therefore, $H_1 \cap H_2 = (\widetilde{H}_1 \cap \square_P) \cap (\widetilde{H}_2 \cap \square_P)$ is the intersection of H_1 with $\widetilde{H}_1 \cap \widetilde{H}_2$, which is a hyperplane of the cube $\widetilde{H}_1$ — but Proposition 2.5.7 says this is itself a hyperplane of H_1, if it is non-empty. That is, the non-empty intersection of two hyperplanes of $\square_P$ is a hyperplane of a hyperplane of $\square_P$. By induction, the non-empty intersection of any number of hyperplanes of $\square_P$ is an iterated hyperplane of $\square_P$.

Now, Corollary 2.5.5 says that hyperplanes are CAT(0) cubical complexes, so Theorem 2.2.3 implies that non-empty intersections of hyperplanes are contractible. Therefore,

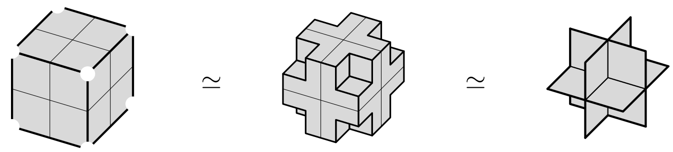

Figure 2.13: A deformation retraction from a cube with the vertices deleted to the union of its midcubes

we can apply the Nerve Theorem, and conclude that the union of the hyperplanes of $\square_P$ is homotopy equivalent to the nerve, Δ_P.

It only remains to show that the union of hyperplanes is homotopy equivalent to $\square_P$ with the vertices deleted. This idea is illustrated in Fig. 2.13.

Let $\square$ be an arbitrary cubical complex, and let $2\square$ denote the subdivision of $\square$ where each cube is replaced by a $2 \times \cdots \times 2$ stack of cubes. For each midcube M of a face F of $\square$, the subdivision $2M$ embeds as a subcomplex into $2F$. For each face f of $2F$ and for all i, either all points of f have the ith coordinate equal to $\frac{1}{2}$, or some points have ith coordinate equal to 0 or 1; therefore, each face of $2F$ either contains a vertex of F or is contained in a midcube of F. Thus the union of the midcubes of F is equal to the induced subcomplex of $2F$ obtained by deleting the vertices of F; for $\square$, this means that the union of the hyperplanes is the induced subcomplex of $2\square$ obtained by deleting all vertices of $\square$.

Now, let us use a result from Chapter 1. In the terminology of that chapter, the set V of vertices of $\square$ is a face-like subcomplex of $2\square$: no face of $2\square$ can contain more than one vertex in V. Therefore, by Proposition 1.3.3, $\|\square\| \setminus \|V\| \simeq \|\square \setminus V\|$, which is the union of the hyperplanes of $\square$. ■

(Note: See Proposition 3.6.7 for a generalisation of this result.)

This theorem gives an alternative way to prove Proposition 2.4.10 (which said that $\square_P$ has a cut vertex if and only if Δ_P is disconnected).

Alternative proof of Proposition 2.4.10. Observe that $\|\square_P\| \setminus \|V(\square_P)\|$ has more than one component if and only if some vertex of $\square_P$ is a cut vertex. Therefore, since connectedness is a homotopy invariant, $\square_P$ has a cut vertex if and only if Δ_P has more than one component. ∎

Remark 2.6.4 (observed by anonymous referee). Instead of applying the nerve theorem to the hyperplanes of a CAT(0) cubical complex $\square_P$, we could instead apply it to the hyperplane carriers (i.e. the sets of cubes containing a given hyperplane). By definition of Hagen's contact graph [Hag12], two hyperplane carriers intersect if and only if the corresponding vertices of the contact graph are adjacent; by the Helly property of convex subspaces [Hag12, Proposition 2.2.6], a set of hyperplane carriers thus has a common intersection if and only if the corresponding vertices form a clique in the contact graph. Thus the nerve of the hyperplane carriers is the clique complex of the crossing graph. The union of the hyperplane carriers is the entire cubical complex, so the nerve theorem implies that the clique complex of the contact graph is contractible.

In the remainder of this section, we present some more facts of a topological flavour, by considering the facets of $\square_P$ and Δ_P.

Proposition 2.6.5 (see [AOS12, Lemma 2.4])**.** *The maximal faces of $\square_P$ are exactly those of the form $C(\downarrow A, A)$ where A is a maximal consistent antichain of P (under inclusion). Thus $A \mapsto C(\downarrow A, A)$ is a bijection from facets of Δ_P to facets of $\square_P$.*

Proof. The facets of Δ_P are the maximal consistent antichains of P, by definition, so the second statement follows immediately from the first. To prove the first statement, there are two things to show: we need to check that $C(\downarrow A, A)$ is always a maximal face of $\square_P$, and that every maximal face has this form.

First, suppose A is a maximal consistent antichain, and $C(I, M)$ is a face of $\square_P$ with $C(\downarrow A, A) \subseteq C(I, M)$. In particular, Lemma 2.5.1 then says that $A \subseteq M \subseteq \max I$. Since A is a maximal consistent antichain, we must therefore have $A = M = \max I$, so $I = \downarrow \max I = \downarrow A$. Thus $C(\downarrow A, A) = C(I, M)$, so $C(\downarrow A, A)$ is indeed a maximal face of $\square_P$.

On the other hand, suppose $C(J, N)$ is an arbitrary maximal face of $\square_P$. Consider the face $C(J, \max J)$: since $N \subseteq \max J$ and $J \setminus \max J \subseteq J \subseteq J$, Lemma 2.5.1 says that

$C(J,N) \subseteq C(J, \max J)$, so the maximality of $C(J,N)$ means that $N = \max J$. Thus $J = {\downarrow} \max J = {\downarrow} N$, so $C(J,N) = C({\downarrow} N, N)$. Thus it only remains to prove that N is a maximal consistent antichain.

Suppose N is not maximal. The set

$$S := \{x \in P \setminus N : N \cup x \text{ is a consistent antichain}\}$$

is thus non-empty, so it has an element x_0 that is minimal with respect to the order in P. Since x_0 is minimal in S, any $y \in P$ with $y < x_0$ must be comparable or inconsistent with some $n \in N$. However, y cannot be inconsistent with n, as then x_0 and n would be inconsistent; and y cannot be greater than n, as then we would have $x_0 > n$. Thus y must be less than n, so $y \in {\downarrow} N$. Therefore, $({\downarrow} N) \cup x_0 = J \cup x_0$ is a consistent downset of P. We also have $N \subseteq N \cup x_0$ and $(J \cup x_0) \setminus (N \cup x_0) \subseteq J \subseteq J \cup x_0$, so Lemma 2.5.1 says that $C(J,N) \subsetneq C(J \cup x_0, N \cup x_0)$. But this contradicts the maximality of $C(J,N)$; therefore, N must be a maximal consistent antichain. ∎

Note that the dimension of A as a face of Δ_P is $|A| - 1$, whereas the dimension of $C({\downarrow} A, A)$ in $\square_P$ is $|A|$. This gives us an alternative proof of Corollary 2.3.6 (which said that $\dim \square_P = \dim \Delta_P + 1$), as well as the following corollary:

Corollary 2.6.6. *Δ_P is pure if and only if $\square_P$ is pure.*

We can also combine this proposition with Lemma 2.5.2, to obtain the following lemma:

Lemma 2.6.7. *If two facets $C({\downarrow} A, A)$ and $C({\downarrow} B, B)$ of $\square_P$ intersect in a face of dimension r, the corresponding facets A and B of Δ_P intersect in the face $A \cap B$, which has dimension $r - 1$.*

Proof. Lemma 2.5.2 says that the intersection of $C({\downarrow} A, A)$ and $C({\downarrow} B, B)$, if it exists, is

$C(\downarrow A \cap \downarrow B, A \cap B)$, which has dimension

$$\begin{aligned}\dim_{\square_P} C(\downarrow A \cap \downarrow B, A \cap B) &= |A \cap B| \\ &= \dim_{\Delta_P}(A \cap B) + 1. \end{aligned} \qquad \blacksquare$$

Recall that by convention, we require $\varnothing$ to be a face of every simplicial complex, but not a face of any cubical complex. Therefore, if $C(\downarrow A, A) \cap C(\downarrow B, B)$ is a face of dimension 0, then this lemma says that A and B intersect in the (-1)-dimensional empty face of Δ_P; however, if the intersection of $C(\downarrow A, A)$ and $C(\downarrow B, B)$ is empty, then we can say nothing about $A \cap B$.

2.7 Balancedness

In this final section, we will take a look at a special class of simplicial and cubical complexes, namely *balanced* complexes.

Definition 2.7.1. An *r-colouring* of a simplicial complex Δ is a map κ_s (with "s" for "simplicial") from the set $V(\Delta)$ of vertices of Δ to the set $\{1, \ldots, r\}$, with the property that for any two vertices connected by an edge (or equivalently, any two distinct vertices that lie in a common face), the images of the vertices under κ_s are different. If such an r-colouring exists, we say that Δ is *r-colourable*; if Δ is $(d-1)$-dimensional and d-colourable, we say it is *balanced*. Note that a $(d-1)$-dimensional simplicial complex must have a face with d vertices, by definition, so at least d colours are always necessary for colouring a $(d-1)$-dimensional complex: the balanced condition says that d colours are enough.

Here is another way of viewing colourings. The set $\{1, \ldots, r\}$ may be thought of as the set of vertices of the $(r-1)$-dimensional simplex Σ_{r-1}. From this viewpoint, an r-colouring of Δ is a map $\kappa_s : V(\Delta) \to V(\Sigma_{r-1})$, such that the restriction of κ_s to any face of Δ is a bijection to a face of Σ_{r-1}.

This idea motivates the following definition of colourings for cubical complexes: an r-

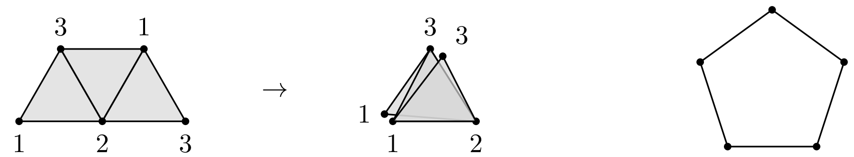

(a) Balanced (b) Not balanced

Figure 2.14: A balanced and a non-balanced simplicial complex

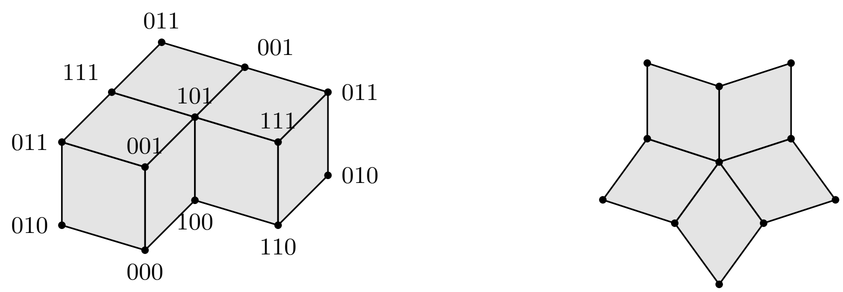

(a) Balanced (b) Not balanced

Figure 2.15: A balanced and a non-balanced cubical complex (whose crossing complexes are the ones in Fig. 2.14)

colouring of a cubical complex $\square$ is a map $\kappa_c : V(\square) \to V([0,1]^r)$ (with "c" for "cubical"), where $[0,1]^r$ is the r-dimensional cube, such that the restriction of κ_c to any face of $\square$ is a bijection to a face of $[0,1]^r$. If such a colouring exists, $\square$ is *r-colourable*, and if $\square$ is d-dimensional and d-colourable, it is called *balanced*. See Figs. 2.14 and 2.15 for some examples of balanced and non-balanced simplicial and cubical complexes.

Note that any simplicial complex with n vertices is n-colourable, by assigning a different colour to every vertex. Similarly, recall that any CAT(0) cubical complex with n hyperplanes can be embedded in the n-dimensional cube, so it is n-colourable. However, there are many cubical complexes (which cannot be CAT(0)) that are not r-colourable for any r — for example, any non-bipartite graph.

There are many connections between balanced and flag simplicial complexes. For example:

Theorem 2.7.2 ([Fro08])**.** *The f-vector of any flag simplicial complex is also the f-vector* *of some balanced simplicial complex.*

Now let us return to the world of CAT(0) cubical complexes, with the main result for this section:

Theorem 2.7.3. *Δ_P is r-colourable if and only if $\square_P$ is r-colourable. Hence Δ_P is balanced if and only if $\square_P$ is balanced.*

Proof. First, suppose Δ_P is r-colourable, so we have a colouring $\kappa_s : V(\Delta_P) \to \{1, \ldots, r\}$. Now, recall that the vertex set of $\square_P$ is the set of consistent downsets of P, and the vertex set of $[0,1]^r$ is $\{0,1\}^r$, so define a map $\kappa_c : V(\square_P) \to \{0,1\}^r$ by sending a downset I to the vector $w = (w_1, \ldots, w_r)$, where

$$w_j = \#\{i \in I : \kappa_s(i) = j\} \mod 2;$$

that is, the jth coordinate of w is 0 if there are an even number of elements of I with colour j, and it is 1 if this number is odd. We claim that this is a valid colouring.

Suppose $C(I, M) = \{I \setminus N : N \subseteq M\}$ is a face of $\square_P$. Since M is a subset of $\max I$, it is a consistent antichain of P, thus κ_s assigns different colours to all elements of M. Therefore, every vertex in $C(I, M)$ is assigned a different colour by κ_c, and these colours are precisely the set of vectors in $\{0,1\}^r$ where the jth coordinate may vary for all colours j appearing in M, and otherwise the jth coordinate matches the parity of colour j appearing in I. This set is a face of $[0,1]^r$, so κ_c is a valid colouring, and $\square_P$ is r-colourable.

Now, assume $\square_P$ is r-colourable, with colouring $\kappa'_c : V(\square_P) \to \{0,1\}^r$. Recall that the vertex set of Δ_P is in bijection with the set of hyperplanes of $\square_P$. Geometrically, since κ'_c is a bijection on each face, the image of a midcube under κ'_c is a midcube of a face of $[0,1]^r$, and if two midcubes meet at a common face in $\square_P$, their images also meet at a common face in $[0,1]^r$. Therefore, κ'_c takes each hyperplane of $\square_P$ to a subset of a hyperplane of $[0,1]^r$. There are r hyperplanes in $[0,1]^r$, so we can define a map $\kappa'_s : V(\Delta_P) \to \{1, \ldots, r\}$ by sending an element of $V(\Delta_P)$, which corresponds to a hyperplane of $\square_P$, to the hyperplane of $[0,1]^r$ containing its image under κ'_c.

Now, if two vertices of Δ_P are connected by an edge, then the corresponding hyperplanes of $\square_P$ must intersect. This intersection must meet some face of $\square_P$, so these two hyperplanes must involve distinct midcubes of this face. Since κ'_c is a bijection on this face, these two midcubes must be sent to different midcubes in a face of $[0,1]^r$; therefore, the images of the two hyperplanes in $\square_P$ must lie in different hyperplanes in $[0,1]^r$. Therefore, κ'_s is a valid colouring, so Δ_P is r-colourable.

Finally, Δ_P is balanced if and only if it is $(\dim \Delta_P + 1)$-colourable, which happens if and only if $\square_P$ is $(\dim \square_P)$-colourable by Corollary 2.3.6, which precisely means that $\square_P$ is balanced. ■

The crossing complex Δ_P is r-colourable if and only if the chromatic number of its underlying graph, the crossing graph, is at most r. The chromatic numbers of crossing graphs have been previously studied by Chepoi and Hagen [CH13], who found connections between the chromatic polynomial and the minimal number of trees whose product $\square_P$ embeds into.

The next corollary to Theorem 2.7.3 follows immediately by combining it with Theorem 2.3.5:

Corollary 2.7.4. *If $p(x)$ is the f-polynomial of a balanced CAT(0) cubical complex, then $p(x+1)$ is the f-polynomial of a balanced, flag simplicial complex.*

We will return to f-polynomials at the end of this section, but until then, let us take a detour.

One large class of balanced CAT(0) cubical complexes comes from taking P to be a poset, i.e. a PIP with no inconsistent pairs. Recall the following well-known fact about posets:

Theorem 2.7.5 (Dilworth's theorem [Dil50, Theorem 1.1])**.** *If P is a poset where the largest antichain has cardinality r, then P can be written as the union of r chains.*

Translating this into the language of balanced simplicial complexes gives us the following immediate consequence:

Corollary 2.7.6. *If P is a PIP with no inconsistent pairs, then Δ_P (and thus also $\square_P$) is balanced.*

Proof. If the size of the largest antichain of P is d, the antichain complex Δ_P has dimension $d-1$. Dilworth's theorem says that P is the union of d chains: we can therefore colour Δ_P by assigning colour i to the vertices in the ith chain. Two vertices of Δ_P are adjacent if and only if they are incomparable in P, which means they cannot be in the same chain: hence this colouring is valid. ■

These theorems are only useful if we can detect whether a given complex $\square_P$ comes from a PIP without inconsistent pairs: fortunately, there is the following result, which follows quickly from some observations by [AOS12; ABY14].

Lemma 2.7.7. *P has no inconsistent pairs if and only if there is some vertex v_∞ of $\square_P$ such that every vertex lies on some shortest edge path from v_0 to v_∞.*

If $\square_P$ is a complex with this property, we say that $\square_P = [v_0, v_\infty]$ is an *interval*.

Proof. First, [AOS12, Lemma 3.2] observed that if $\square_P$ is a complex of this form, then P is consistent.

Conversely, if P has no inconsistent pairs, then P itself is a consistent downset, so we may define v_∞ to be the vertex corresponding to P as a downset. Now, [ABY14, Proposition 7.4] noted that the shortest edge paths from the root vertex v_0 to the vertex v_∞ all have length $|P|$. Now, suppose w is an arbitrary vertex of $\square_P$ corresponding to the downset I. Construct a path from v_0 to v_∞ passing through w as follows: define $J_0 \coloneqq \varnothing$, and inductively take $J_i \coloneqq J_{i-1} \cup x_i$ where x_i is a minimal element of $I \setminus J_{i-1}$, until we reach the stage where $J_i = I$; from then on, do the same but take x_i to be a minimal element of $P \setminus J_{i-1}$ until $J_i = P$. Each element of P is taken as x_i once, so this path has length $|P|$, hence it is a shortest edge path. ■

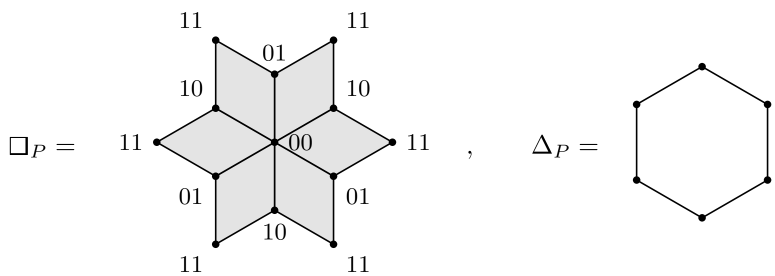

Figure 2.16: A balanced CAT(0) 2-dimensional cubical complex that does not embed into $\mathbb{R}^2$

Note that [AOS12, Theorem 3.5] observe a stronger property (using similar proof ideas): they show that if $\square_P = [v_0, v_\infty]$ is a d-dimensional interval, then it actually can be *embedded* into the unit grid structure in $\mathbb{R}^d$. This is not true in general for balanced complexes — for example, see Fig. 2.16.

While we are discussing PIPs with order relations but no inconsistent pairs, we may as well comment on the opposite end of the spectrum, namely PIPs with inconsistent pairs but no order relations. (Recall from the discussion after Definition 2.2.5 that any graph can be thought of as such a PIP.)

Lemma 2.7.8. *P has no order relations (except equality) if and only if every facet of $\square_P$ contains the root vertex; in other words, $\square_P = \overline{\operatorname{star}}\, v_0$.*

Proof. Recall from Proposition 2.6.5 that the facets of $\square_P$ are the faces of the form $C(\downarrow A, A)$, where A is a maximal consistent antichain of P. The vertex v_0 corresponds to the downset $\varnothing$, so a facet $C(\downarrow A, A)$ contains v_0 if and only if $(\downarrow A) \setminus A$ is empty, which happens if and only if all elements of A are minimal in P. Every element of P, minimal or not, is contained in some maximal consistent antichain, so all facets contain v_0 if and only if all elements of P are minimal. This precisely means that P has no order relations except equality. ■

Recall from Lemma 2.4.9 that the link of v_0 is the crossing complex of $\min P$. Therefore, if P has no order relations, so $\min P = P$, then $\operatorname{link} v_0$ is just the crossing complex of $\square_P = \overline{\operatorname{star}}\, v_0$. Thus $\square_P$ is essentially a kind of cubical cone over the crossing complex of P. The complexes shown in the last row of Fig. 2.7 as well as Figs. 2.15b and 2.16 are examples of this situation, if the vertex in the centre is chosen to be the root vertex. (The name "star" is particularly appropriate in these last two examples.)

Now, let us return to f-vectors and f-polynomials.

Definition 2.7.9. Suppose Δ is a simplicial complex with an r-colouring κ_s. For each subset $S \subseteq \{1, \ldots, r\}$, define $f_S(\Delta)$ to be the number of faces of Δ whose image under κ_s is exactly S. The tuple $\big(f_S(\Delta)\big)_{S\subseteq\{1,\ldots,r\}}$ is called the *coloured f-vector* or (confusingly) the *flag f-vector* of Δ (with no obvious connection to the notion of a flag simplicial complex).

The coloured f-vector is a refinement of the usual f-vector of Δ, since

$$f_i(\Delta) = \sum_{\substack{S\subseteq\{1,\ldots,r\} \\ |S|=i}} f_S(\Delta).$$

We can also define a refinement of the f-polynomial: define the *coloured f-polynomial* of Δ to be the following polynomial in r variables:

$$\begin{aligned} f(\Delta, x_1, \ldots, x_r) &\coloneqq \sum_{S\subseteq\{1,\ldots,r\}} f_S(\Delta) \prod_{j\in S} x_j \\ &= \sum_{\sigma\in\Delta} \prod_{i\in\sigma} x_{\kappa_s(i)}. \end{aligned}$$

Notice that the usual f-polynomial $f(\Delta, t)$ can be obtained from the coloured f-polynomial by setting all of the variables equal to t.

We can also define coloured f-vectors and f-polynomials for an r-coloured cubical complex $\square$, but now we need more information to specify the colour of a face. For each pair S, T of disjoint subsets of $\{1, \ldots, r\}$, let $f_{S,T}(\square)$ be the number of faces σ of $\square$ where:

- if the ith coordinate of $\kappa_c(v)$ is 1 for some vertex $v \in \sigma$ and 0 for some other vertex w, then $i \in T$, and

- if the ith coordinate of $\kappa_c(v)$ is 1 for all vertices v, then $i \in S$.

For example, if κ_c assigns the colours 1000, 1100, 1010 and 1110 to the vertices of σ, then this face contributes to the f-number $f_{\{1\},\{2,3\}}(\square)$, since all vertices have a 1 in the first position and they vary in the 2nd and 3rd positions.

Now, define the following polynomial in $2r$ variables:

$$f(\square, x_1, \ldots, x_r, y_1, \ldots, y_r) \coloneqq \sum_{S \subseteq \{1,\ldots,r\}} \sum_{T \subseteq \{1,\ldots,r\} \setminus S} f_{S,T}(\square) \prod_{k \in T} x_k \prod_{j \in S} y_j.$$

Note that the original f-polynomial $f(\square, t)$ can be recovered by setting $x_i = t$ and $y_i = 1$ for all i: in an s-dimensional face σ of $[0,1]^r$, exactly s coordinates vary, so setting $x_i = t$ and $y_i = 1$ means that a s-dimensional face of $\square$ contributes to the degree s term in the f-polynomial $f(\square, t)$.

Let us examine how this polynomial behaves for CAT(0) cubical complexes.

Lemma 2.7.10. *Suppose κ_s is a colouring of Δ_P, and let κ_c be the colouring of $\square_P$ constructed from κ_s in the proof of Theorem 2.7.3. Then*

$$f(\square_P, x_1, \ldots, x_r, y_1, \ldots, y_r) = \sum_{C(I,M) \in \square_P} \prod_{k \in \kappa_s(M)} x_k \prod_{j \in S_{I,M}} y_j$$

where $S_{I,M}$ is the set of colours appearing in I an odd number of times, but not in M.

Proof. The colours assigned to the face $C(I, M)$ are the vectors $w \in \{0,1\}^r$ where position i varies for each colour i appearing in M, and otherwise position i is always 1 if colour i appears an odd number of times in $I \setminus M$. Therefore, this face contributes to the f-number $f_{S_{I,M}, \kappa_s(M)}(\square_P)$. ∎

With this observation, we can refine the proof of Theorem 2.3.5 to get an r-colourable version.

Theorem 2.7.11. *If κ_s and κ_c are the related colourings from Lemma 2.7.10, then $f(\square_P, x_1, \ldots, x_r, 1, \ldots, 1) = f(\Delta_P, 1 + x_1, \ldots, 1 + x_r)$.*

Proof. As in the proof of Theorem 2.3.5, let "$A \trianglelefteq P$" mean "A is a consistent antichain of P".

Then,

$$\begin{aligned}
f(\square_P, x_1, \ldots, x_r, 1, \ldots, 1) &= \sum_{C(I,M) \in \square_P} \prod_{k \in \kappa_s(M)} x_k \\
&= \sum_{A \trianglelefteq P} \sum_{M \subseteq A} \prod_{m \in M} x_{\kappa_s(m)} \\
&= \sum_{A \trianglelefteq P} \prod_{m \in A} \left(1 + x_{\kappa_s(m)}\right) \\
&= f(\Delta_P, 1 + x_1, \ldots, 1 + x_r).
\end{aligned}$$

■

Chapter 3

Topology of complements of skeletons

3.1 Introduction

Polytopal complexes are important objects in topology and combinatorics, which include simplicial complexes, cubical complexes and polytopes. Their geometric realisations provide examples of a wide array of topological spaces, and much research has gone into studying the interplay between their combinatorial and topological aspects.

One important feature of a polytopal complex is its k-*skeleton*, the set of faces of the complex of dimension less than or equal to k. Skeletons act as a framework which the high-dimensional faces are attached to, so studying the structure of a skeleton can reveal a lot about a complex. For example, [Bay18] surveys many results about objects that can be reconstructed from their k-skeletons for certain values of k. Skeletons also play an important role in defining cellular homology, and in important topological results such as Poincaré duality.

In this chapter, however, we aim to approach this topic from the other direction, starting from the higher-dimensional faces instead of the low-dimensional ones. We define the *kth co-skeleton* of a complex to be the set of faces of dimension *higher* than k. We discover that there is a strong relationship between co-skeletons and *links*: a link of a face in a polytopal complex captures the local structure of the complex around that face, so in a sense, the co-skeletons give us a "global" summary of the "local" information of the complex.

Section 3.2 of this chapter sets out the basic definitions we will use. In Section 3.3, we examine some initial facts about co-skeletons of arbitrary complexes, and their connections with various forms of duality (see Proposition 3.3.3 and Corollary 3.3.5). In Section 3.4, we construct a long exact sequence relating the kth and $(k-1)$th co-skeletons and the links of k-dimensional faces (Theorem 3.4.1). In Section 3.5, we use this long exact sequence to study some families of complexes defined by properties of links — specifically, Cohen–Macaulay complexes, Leray complexes, and stacked balls — and give characterisations of these families in terms of the homology of their co-skeletons (Theorems 3.5.3, 3.5.10 and 3.5.11). Finally, Section 3.6 examines cubical complexes, particularly CAT(0) cubical complexes: we examine the "crossing complex" defined in [Row23] (Chapter 2), and show that a CAT(0) cubical complex has one of the properties discussed in Section 3.5 if and only if its crossing complex shares that property (Theorem 3.6.10).

3.2 Preliminaries

We begin with some definitions. Readers familiar with polytopal complexes and topology may skip most of this section, but beware that we give slightly non-standard definitions in a couple of places, specifically for geometric realisations and links.

3.2.1 Polytopal complexes

A *polytopal complex* X is a collection of polytopes in Euclidean space $\mathbb{R}^N$ with the following properties:

- If σ is in X and τ is a face of σ, then τ is in X, and
- If σ and σ' are polytopes in X, then $\sigma \cap \sigma'$ is a face of each (possibly the empty face).

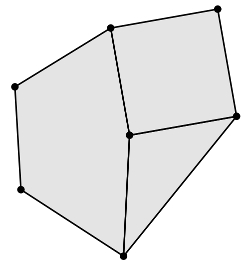

Figure 3.1: An example of a polytopal complex

See Fig. 3.1 for a small example. In this chapter, we will only consider finite polytopal complexes.

If every polytope in X is a simplex, then we say that $\Delta \coloneqq X$ is a *(geometric) simplicial complex*. If every polytope is a cube (that is, a polytope combinatorially equivalent to $[0,1]^i$ for some dimension $i \geq 0$), then $\square \coloneqq X$ is a *cubical complex*.

An *abstract simplicial complex* Δ is a poset isomorphic to the poset of faces of a geometric simplicial complex, ordered by inclusion. Equivalently, it is a collection of subsets of some finite set, with the property that $\sigma \in \Delta$ and $\tau \subseteq \sigma$ implies $\tau \in \Delta$. The dimension of a face σ is $\dim \sigma \coloneqq \#\sigma - 1$.

In any of these types of complex, faces of dimension 0 and 1 are called *vertices* and *edges* respectively, and a maximal face (by inclusion) is called a *facet*. If all facets of a complex have the same dimension, the complex is *pure*. The dimension of the complex is the maximum dimension of its faces. The number of faces of dimension k in a complex X is denoted $f_k(X)$.

If σ is a polytope, $|\sigma|$ will denote its relative interior. If S is a collection of faces in a polytopal complex, then their *geometric realisation* $|S|$ is the union of the relative interiors of the polytopes in S. This is slightly different from the usual definition of a geometric realisation — if σ is in S, we do not include the boundary of σ in the geometric realisation $|S|$ unless those boundary faces are also part of S, unlike other definitions. We will take care to distinguish between S as a set of polytopes and $|S|$ as a topological space.

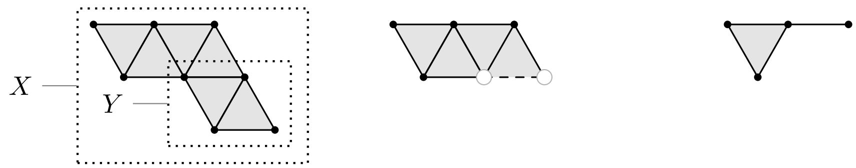

(a) X and Y (b) $|X \setminus Y|$ (c) $|X - Y|$

Figure 3.2: A comparison of the different types of deletion

Suppose Y is a subset of a polytopal complex X. We write $X \setminus Y$ to denote the difference of sets, so $X \setminus Y$ is the set of faces of X that are not in Y (which is not generally a polytopal complex). If Y is a polytopal subcomplex of X, then $X - Y$ will denote the polytopal complex consisting of all faces of X that do not intersect any faces of Y. We will reserve "$-$" to denote this "combinatorial" deletion, and use "$\setminus$" to denote deletion of sets or topological spaces. Note that $|X \setminus Y| = |X| \setminus |Y|$, but the space $|X - Y|$ is not in general the same; see Fig. 3.2, for an example. However, these spaces are sometimes related by the following lemma:

Lemma 3.2.1 ([Mun84a, Lemma 70.1])**.** *If Δ is a simplicial complex and Λ is an induced subcomplex (in other words, every face of Δ whose vertices are contained in the vertex set of Λ is a face of Λ), then $|\Delta - \Lambda|$ and $|\Delta| \setminus |\Lambda|$ are homotopy equivalent.*

Given a face σ of X, the *(open) star* of σ is the set of faces

$$\operatorname{star}_X \sigma \coloneqq \{\tau \in X : \sigma \subseteq \tau\}.$$

The geometric realisation of a star is always contractible, if $\sigma \neq \varnothing$. The *link* of σ is the set

$$\begin{aligned}\operatorname{link}_X \sigma &\coloneqq \{\tau \in X : \sigma \subset \tau,\ \tau \neq \sigma\} \\ &= \operatorname{star}_X \sigma \setminus \{\sigma\}.\end{aligned}$$

We will sometimes simply write "$\operatorname{star} \sigma$" and "$\operatorname{link} \sigma$" if the space X is clear from context.

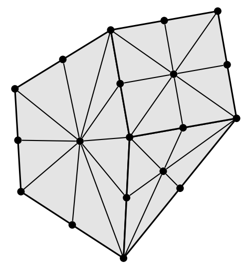

Figure 3.3: The barycentric subdivision of the polytopal complex in Fig. 3.1

Note that if Δ is an abstract simplicial complex, the usual definition of a link is slightly different: we will refer to the usual definition as the "simplicial link", defined by

$$\text{s-link}_\Delta \sigma \coloneqq \{\tau \in \Delta : \sigma \cup \tau \in \Delta,\ \sigma \cap \tau = \varnothing\}.$$

Although our link and the simplicial link are not the same, they are homotopy equivalent.

If X is a polytopal complex, its *barycentric subdivision* is the abstract simplicial complex $\operatorname{bary}(X)$ which has one vertex v_σ for each non-empty face σ of X, and a set $\{v_{\sigma_1}, \ldots, v_{\sigma_m}\}$ forms a face of $\operatorname{bary}(X)$ whenever $\{\sigma_1, \ldots, \sigma_m\}$ is a chain in the poset of faces of X ordered by inclusion (that is, $\sigma_1 \subset \cdots \subset \sigma_m$, up to reordering). If each vertex v_σ is positioned at the barycentre of the polytope σ, then the geometric realisation of $\operatorname{bary}(X)$ is exactly the geometric realisation of X, as a topological space. See Fig. 3.3 for an example.

Given a polytopal complex X, the set of k-dimensional faces of X will be denoted X_k. The *k-skeleton* of X is the subcomplex

$$\operatorname{Skel}_k X \coloneqq \{\sigma \in X : \dim \sigma \leq k\}.$$

The central definition in this chapter is the *kth co-skeleton* of X, which is the complement of the the k-skeleton:

$$\operatorname{Skel}_k^c X \coloneqq X \setminus \operatorname{Skel}_k X$$

$$= \{\sigma \in X : \dim \sigma > k\}.$$

We will mostly be interested in the topological properties of $|\mathrm{Skel}^c_k X| = |X| \setminus |\mathrm{Skel}_k X|$. Note that the co-skeleton is not in general a polytopal complex, except in two special cases: $\mathrm{Skel}^c_{-1} X$ is X itself (modulo the empty face, which makes no difference to the topology), and if X is d-dimensional, $\mathrm{Skel}^c_d X$ is the polytopal complex with no faces.

3.2.2 Topology

We will use the notation $A \simeq B$ to indicate that the spaces A and B are homotopy equivalent.

We assume that the reader is familiar with homology and cohomology; for background, refer to [Hat02; Mun84a]. While this chapter does focus on cell complexes, many spaces we consider are not themselves cell complexes, so $H_i(A)$ will denote singular homology, with coefficients in R where R is a field or $\mathbb{Z}$. Reduced homology is denoted by $\widetilde{H}_i(A)$. We use the convention that $\widetilde{H}_{-1}(\varnothing) = R$. Analogous statements apply to cohomology, denoted $H^i(A)$.

For reference, here are two important theorems from algebraic topology that we will use repeatedly.

Theorem 3.2.2 (Mayer–Vietoris, [Hat02, p. 149], [Mun84a, Theorem 33.1])**.** *If A and B are open subsets of a topological space, then there is a long exact sequence:*

$$\cdots \longrightarrow H_i(A \cap B) \longrightarrow H_i(A) \oplus H_i(B) \longrightarrow H_i(A \cup B) \longrightarrow H_{i-1}(A \cap B) \longrightarrow \cdots .$$

If $A \cap B \neq \varnothing$, then we may replace these unreduced homology groups with reduced homology throughout.

Theorem 3.2.3 (Nerve theorem, [Hat02, Corollay 4G.3], [Bjö95, Theorem 10.7])**.** *Suppose $\mathcal{U} = \{U_1, \ldots, U_n\}$ is a family of open sets whose union is a paracompact space (e.g. any subspace of $\mathbb{R}^n$), or a family of closed sets whose union is a triangulable space. Sup-*

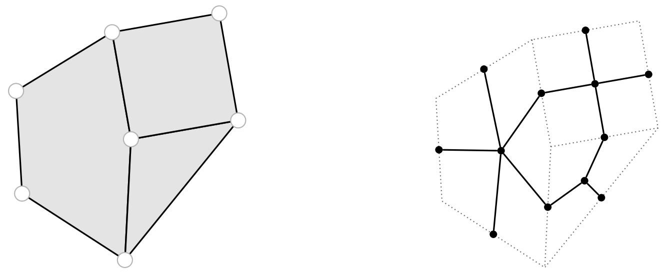

(a) $\mathrm{Skel}_0^c X$ (b) $\mathrm{bary}(X) - \mathrm{bary}(\mathrm{Skel}_0 X)$

Figure 3.4: Lemma 3.3.1 applied to Fig. 3.1, with $k = 0$

pose further that for every index set $I \subseteq \{1, \ldots, n\}$, the intersection $\bigcap_{i \in I} U_i$ is either contractible or empty.

Construct a simplicial complex $N(\mathcal{U})$ (called the "nerve" of $\mathcal{U}$) where the vertex set is $\{1, \ldots, n\}$ and the set I forms a face whenever $\bigcap_{i \in I} U_i$ is non-empty. Then

$$\bigcup_{i=1}^{n} U_i \simeq |N(\mathcal{U})|.$$

3.3 First results about co-skeletons

Let us begin to study the co-skeleton $\mathrm{Skel}_k^c X$.

This first lemma tells us that while $\mathrm{Skel}_k^c X$ is not itself a polytopal complex, it is homotopy equivalent to one. See Fig. 3.4.

Lemma 3.3.1. $|\mathrm{Skel}_k^c X| \simeq |\mathrm{bary}(X) - \mathrm{bary}(\mathrm{Skel}_k X)|$.

Proof. By definition,

$$\begin{aligned} |\mathrm{Skel}_k^c X| &= |X| \setminus |\mathrm{Skel}_k X| \\ &= |\mathrm{bary}(X)| \setminus |\mathrm{bary}(\mathrm{Skel}_k X)|. \end{aligned}$$

Now, $\mathrm{bary}(\mathrm{Skel}_k X)$ is an induced subcomplex of the simplicial complex $\mathrm{bary}(X)$: if $v_{\sigma_1}, \ldots, v_{\sigma_m}$ are vertices of $\mathrm{bary}(\mathrm{Skel}_k X)$ that form a face of $\mathrm{bary}(X)$, then the faces

$\sigma_1, \ldots, \sigma_m$ form a chain in the face poset of X, so they still form a chain in the face poset of $\operatorname{Skel}_k X$. Therefore, we may invoke Lemma 3.2.1:

$$|\operatorname{bary}(X)| \setminus |\operatorname{bary}(\operatorname{Skel}_k X)| \simeq |\operatorname{bary}(X) - \operatorname{bary}(\operatorname{Skel}_k X)|. \qquad \blacksquare$$

Corollary 3.3.2. *$H_i(\operatorname{Skel}_k^c \Delta) = 0$ when $i > d - k - 1$.*

Proof. The vertices of $\operatorname{bary}(X) - \operatorname{bary}(\operatorname{Skel}_k X)$ correspond to faces σ of X of dimensions between $k+1$ and d; therefore, the largest possible face of $\operatorname{bary}(X) - \operatorname{bary}(\operatorname{Skel}_k X)$ has $d-k$ vertices, so its dimension is $d-k-1$. ■

The name "co-skeleton" was chosen to suggest "complement", but also duality. Let us illustrate why.

Proposition 3.3.3. *Suppose P is a $(d+1)$-dimensional polytope. Then*

$$|\operatorname{Skel}_k^c \partial P| \simeq |\operatorname{Skel}_{d-k-1} \partial P^*|,$$

where P^* is the polar dual polytope to P.

Proof. By Lemma 3.3.1,

$$|\operatorname{Skel}_k^c \partial P| \simeq |\operatorname{bary}(\partial P) - \operatorname{bary}(\operatorname{Skel}_k \partial P)|.$$

The faces of $\operatorname{bary}(\partial P) - \operatorname{bary}(\operatorname{Skel}_k \partial P)$ are the chains in the poset of faces of ∂P of dimension greater than k. But the poset of non-empty faces of ∂P is isomorphic to the poset of non-empty faces of ∂P^*, flipped upside down. Under this flip, the faces of ∂P of dimension greater than k become faces of ∂P^* of dimension less than or equal to $d-k-1$. Therefore, $\operatorname{bary}(\partial P) - \operatorname{bary}(\operatorname{Skel}_k \partial P) = \operatorname{bary}(\operatorname{Skel}_{d-k-1} \partial P^*)$, so

$$\begin{aligned} |\operatorname{bary}(\partial P) - \operatorname{bary}(\operatorname{Skel}_k \partial P)| &= |\operatorname{bary}(\operatorname{Skel}_{d-k-1} \partial P^*)| \\ &= |\operatorname{Skel}_{d-k-1} \partial P^*|. \end{aligned} \qquad \blacksquare$$

Note that the fact that P is a polytope is not essential to this proposition, just that P has an associated "dual cell structure". In fact, if X is a d-dimensional homology manifold, the space $|\text{bary}(X) - \text{bary}(\text{Skel}_k X)|$ is exactly the $(d-k-1)$-skeleton of the dual cell structure used in some proofs of Poincaré duality — for example, see [Mun84a, §64], which calls this space the "dual $(d-k-1)$-skeleton" of the manifold.

A related duality result is the Alexander duality theorem, which has direct implications for co-skeletons.

Theorem 3.3.4 (Alexander duality, [Mun84a, §71, particularly Theorem 71.1 and Exercise 4])**.** *Suppose X is a polytopal complex where $\dim \Delta = d$ and $|X|$ is homeomorphic to a sphere, and suppose Y is a proper, non-empty subcomplex of X. Then*

$$\widetilde{H}^i(|Y|) \cong \widetilde{H}_{d-i-1}(|X| \setminus |Y|) \qquad \textit{and} \qquad \widetilde{H}_i(|Y|) \cong \widetilde{H}^{d-i-1}(|X| \setminus |Y|).$$

Taking $Y = \text{Skel}_k X$ gives us this corollary:

Corollary 3.3.5. *If X is a d-dimensional polytopal complex where $|X|$ is homeomorphic to a sphere, then for $k = 0, \dots, d-1$,*

$$\widetilde{H}^i(|\text{Skel}_k X|) \cong \widetilde{H}_{d-i-1}(|\text{Skel}^c_k X|) \qquad \textit{and} \qquad \widetilde{H}_i(|\text{Skel}_k X|) \cong \widetilde{H}^{d-i-1}(|\text{Skel}^c_k X|).$$

For a first example of an application for co-skeletons, let us consider neighbourly simplicial spheres. Suppose Δ is a simplicial sphere (that is, a simplicial complex where $|\Delta|$ is homeomorphic to a sphere) with n vertices. Then Δ is *t-neighbourly* if every set of t vertices is the vertex set of a face, or equivalently, if its $(t-1)$-skeleton is isomorphic to the $(t-1)$-skeleton of a simplex with n vertices. For example, the boundaries of d-dimensional cyclic polytopes are $\lfloor d/2 \rfloor$-neighbourly. A major reason why neighbourly spheres have been studied is their connection with the Upper Bound Conjecture (see e.g. [AK85]).

The following well-known fact lets us characterise skeletons of simplices by their homology.

Lemma 3.3.6. *Suppose Δ is a simplicial complex with n vertices, $n \geq 1$, and* $\dim \Delta = k < n$. *Then*

$$\operatorname{rank} \widetilde{H}_k(|\Delta|; \mathbb{Z}) \leq \binom{n-1}{k+1},$$

with equality if and only if Δ is isomorphic to the k-skeleton of a simplex with n vertices. If the coefficient ring is a field instead of $\mathbb{Z}$, the same holds with rank replaced by dimension.

This follows without much difficulty from [BK88, Theorem 3.1] and [Kal02, Equation (3.6)]. Alternatively, for a self-contained proof sketch, let v be a vertex of Δ and consider the Mayer–Vietoris long exact sequence with $A := |\operatorname{star} v|$, $B := |\Delta| \setminus \{v\} \simeq |\Delta - \{v\}|$. The sequence contains this segment:

$$\widetilde{H}_k(|\Delta - \{v\}|) \longrightarrow \widetilde{H}_k(|\Delta|) \longrightarrow \widetilde{H}_{k-1}(|\operatorname{link} v|).$$

The rank of the middle term is thus bounded by the sum of the ranks of the outer terms; now, use induction on n.

This lemma tells us that "being a skeleton of a simplex" is a homological property, if dimension and number of vertices are fixed. Combining this with Corollary 3.3.5, we obtain the following characterisations of neighbourly spheres in terms of their skeletons and their co-skeletons:

Corollary 3.3.7. *Suppose Δ is a d-dimensional simplicial sphere with n vertices. Then for $t = 1, \ldots, d$, the following are equivalent:*

- *Δ is t-neighbourly,*
- $\operatorname{rank} \widetilde{H}_{t-1}(|\operatorname{Skel}_{t-1} \Delta|; \mathbb{Z}) = \binom{n-1}{t}$,
- $\operatorname{rank} \widetilde{H}^{d-t}(|\operatorname{Skel}^c_{t-1} \Delta|; \mathbb{Z}) = \binom{n-1}{t}$.

Proof. By definition, Δ is t-neighbourly if and only if $\operatorname{Skel}_{t-1} \Delta$ is isomorphic to the $(t-1)$-skeleton of a simplex on n vertices. Lemma 3.3.6 says that this occurs if and only if $\operatorname{rank} \widetilde{H}_{t-1}(|\mathrm{Skel}_{t-1} \Delta|; \mathbb{Z}) = \binom{n-1}{t}$. And Corollary 3.3.5 says that

$$\widetilde{H}_{t-1}(|\mathrm{Skel}_{t-1} \Delta|; \mathbb{Z}) \cong \widetilde{H}^{d-t}(|\mathrm{Skel}^c_{t-1} \Delta|; \mathbb{Z}). \qquad \blacksquare$$

3.3.1 Addendum

This subsection was written for this thesis and does not appear in the preprint version of this chapter.

The barycentric subdivison of a polytopal complex X is an example of an *order complex*: given a poset P (e.g. the poset of non-empty faces of X), its order complex $\Delta(P)$ is the abstract simplicial complex whose vertices are the elements of the poset, and whose faces are chains $x_1 < \cdots < x_m$. Lemma 3.3.1 says that the geometric realisation of $\mathrm{Skel}^c_k X$ is homotopy equivalent to the order complex of $\mathrm{Skel}^c_k X$, as a sub-poset of the face poset of X. The following Proposition generalises this; it also justifies the claim from Section 3.2 that our definition of a link is homotopy equivalent to the usual definition in a simplicial complex.

Proposition 3.3.8. *Suppose S is any set of non-empty faces in a polytopal complex X. Then $|S|$ is homotopy equivalent to $\Delta(S)$.*

Proof. The goal of this proof is to show that $\Delta(S)$ is the nerve of a certain open cover of S.

First, let us consider barycentric subdivisions. The faces of $\mathrm{bary}(X)$ correspond to chains in the face poset of X, and the face corresponding to the chain $\sigma_1 < \cdots < \sigma_m$ lies inside the relative interior of the face σ_m, the maximal face of the chain.

Now, let us define our open cover. For each face $\sigma \in S$, define the set U_σ to be

$$U_\sigma \coloneqq |\mathrm{star}_{\mathrm{bary}(X)} v_\sigma| \cap |S|,$$

and let $\mathcal{U}$ denote the collection $\{U_\sigma : \sigma \in S\}$. Each U_σ is the intersection of an open set of $|X|$ with the subspace $|S|$, so U_σ is open in $|S|$. The set U_σ contains $|\sigma|$; therefore, $\mathcal{U}$

is an open cover of $|S|$.

For a set $\{\sigma_1, \dots, \sigma_m\} \subseteq S$, the intersection of the corresponding open sets is

$$\begin{aligned}\bigcap_{i=1}^{m} U_{\sigma_i} &= |S| \cap \bigcap_{i=1}^{m} |\mathrm{star}_{\mathrm{bary}(X)}\, v_{\sigma_i}| \\ &= |S| \cap \left| \bigcap_{i=1}^{m} \mathrm{star}_{\mathrm{bary}(X)}\, v_{\sigma_i} \right|.\end{aligned}$$

Define

$$I := \bigcap_{i=1}^{m} \mathrm{star}_{\mathrm{bary}(X)}\, v_{\sigma_i}.$$

Either this is empty or it is $\mathrm{star}_{\mathrm{bary}(X)}\{v_{\sigma_1}, \dots, v_{\sigma_m}\}$ if $\{v_{\sigma_1}, \dots, v_{\sigma_m}\}$ is a face of $\mathrm{bary}(X)$ — that is, if $\{\sigma_1, \dots, \sigma_m\}$ is a chain. Without loss of generality, assume that σ_m is the maximal element of this chain, so the face $\{v_{\sigma_1}, \dots, v_{\sigma_m}\}$ of $\mathrm{bary}(X)$ lies inside the relative interior of σ_m. By construction, σ_m is in S, so $|S| \cap |I|$ is non-empty. Every face of $\mathrm{star}_{\mathrm{bary}(X)}\{v_{\sigma_1}, \dots, v_{\sigma_m}\}$ that lies in $|S|$ is a simplex with v_{σ_m} as a vertex; therefore, for any point z in $|S| \cap |I|$, the straight line connecting z and v_{σ_m} is contained in $|S| \cap |I|$. Hence $|S| \cap |I|$ is star convex, and thus contractible.

So, $\mathcal{U}$ is an open cover of $|S|$ where every non-empty intersection is contractible. Therefore, the nerve theorem, Theorem 3.2.3, implies that $|S|$ is homotopy equivalent to the nerve $N(\mathcal{U})$. By the discussion above, an intersection $\bigcap_{i=1}^{m} U_{\sigma_i}$ is non-empty if and only if $\{\sigma_1, \dots, \sigma_m\}$ is a chain in S: therefore, $N(\mathcal{U})$ is the order complex of S. ■

3.4 The key long exact sequence

We now reach the main tool of this chapter, which we will make much use of in the upcoming sections.

Theorem 3.4.1. *If X is a d-dimensional polytopal complex and $0 \leq k \leq d$, then there is the following long exact sequence:*

$$\cdots \to \bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\mathrm{link}_X\, \sigma|) \to \widetilde{H}_i(|\mathrm{Skel}^c_k\, X|) \to \widetilde{H}_i(|\mathrm{Skel}^c_{k-1}\, X|) \to \bigoplus_{\sigma \in X_k} \widetilde{H}_{i-1}(|\mathrm{link}_X\, \sigma|) \to \cdots.$$

Proof. We will begin with the case $k = d$. The link of any d-dimensional face is empty, as is $\operatorname{Skel}^c_d X$. The $(d-1)$th co-skeleton $\operatorname{Skel}^c_{d-1} X$ consists only of the interiors of the d-dimensional faces, so it retracts onto $f_d(X)$ points. Therefore, the non-zero part of the claimed long exact sequence looks like this:

$$
\begin{array}{ccccccccc}
0 \longrightarrow & \widetilde{H}_0(\operatorname{Skel}^c_{d-1}) & \longrightarrow & \bigoplus_{\sigma \in X_d} \widetilde{H}_{-1}(\operatorname{link} \sigma) & \longrightarrow & \widetilde{H}_{-1}(\operatorname{Skel}^c_d X) & \longrightarrow 0. \\
& \wr\| & & \wr\| & & \wr\| & \\
& R^{f_d(X)-1} & & R^{f_d(X)} & & R &
\end{array}
$$

These homology groups do indeed form an exact sequence, with appropriate maps. (We won't need to know what the maps are in the rest of this chapter, just that they exist.)

For the remainder of this proof, we will assume that $k < d$. We will use the Mayer–Vietoris theorem. Define the sets A and B by:

$$
\begin{aligned}
A &:= \bigcup_{\sigma \in X_k} |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|, \\
B &:= |\operatorname{Skel}^c_k X| = |X| \setminus |\operatorname{Skel}_k X|.
\end{aligned}
$$

See Figs. 3.5 and 3.6 for an example. Both A and B are open in $|X|$: A is a union of open sets, and B is the complement of a closed set.

First, we claim that the union in A is in fact a *disjoint* union. If two stars, say $\operatorname{star}_{\operatorname{bary}(X)} v_\sigma$ and $\operatorname{star}_{\operatorname{bary}(X)} v_\tau$, were to intersect, they would have a face of $\operatorname{bary}(X)$ in common, and by the definition of an open star, this face would need to contain both v_σ and v_τ. But faces in $\operatorname{bary}(X)$ correspond to chains in the face poset of X; therefore, no face in $\operatorname{bary}(X)$ can contain both v_σ and v_τ, since σ and τ are both k-dimensional and are thus incomparable in the face poset. So A is a disjoint union of stars. Since these stars are open, each one thus forms a connected component of A.

Next, let us consider $A \cap B$. Since $\operatorname{star}_{\operatorname{bary}(X)} v_\sigma$ does not intersect any faces of X of

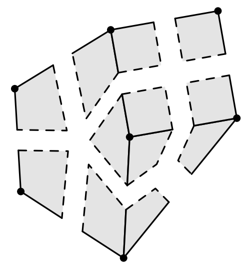

(a) $A = \bigsqcup_{\sigma \in X_0} |\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma|$

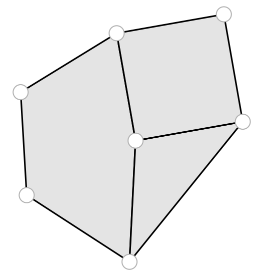

(b) $B = |\mathrm{Skel}_0^c\, X|$

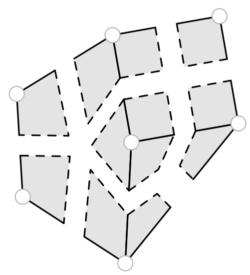

(c) $A \cap B = \bigsqcup_{\sigma \in X_0} |\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|$

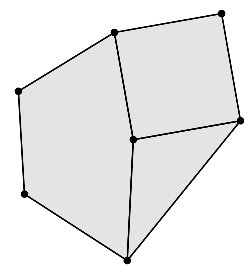

(d) $A \cup B = |\mathrm{Skel}_{-1}^c\, X| = |X|$

Figure 3.5: The sets A, B, $A \cap B$ and $A \cup B$ in the proof of Theorem 3.4.1, when $k = 0$ and X is the polytopal complex in Fig. 3.1

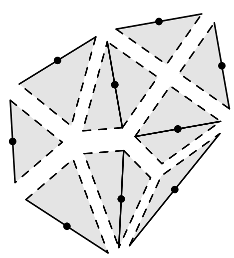

(a) $A = \bigsqcup_{\sigma \in X_1} |\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma|$

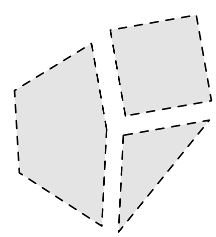

(b) $B = |\mathrm{Skel}_1^c\, X|$

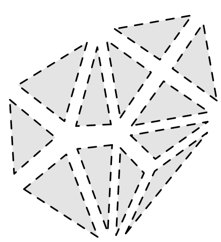

(c) $A \cap B = \bigsqcup_{\sigma \in X_1} |\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|$

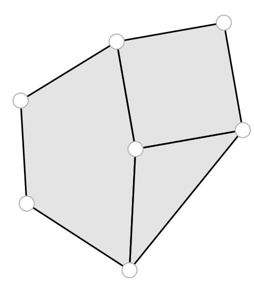

(d) $A \cup B = |\mathrm{Skel}_0^c\, X|$

Figure 3.6: The sets A, B, $A \cap B$ and $A \cup B$ in the proof of Theorem 3.4.1, when $k = 1$ and X is the polytopal complex in Fig. 3.1

dimension less than k, and the only k-dimensional face it intersects is σ, we conclude that

$$A \cap B = \bigsqcup_{\sigma \in X_k} |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|.$$

Now for $A \cup B$. Suppose τ is an arbitrary face of X. If $\dim \tau \leq k-1$, then $|\tau|$ is disjoint from both A and B; if $\dim \tau = k$, then $|\tau|$ is a subset of $|\operatorname{star}_{\operatorname{bary}(X)} v_\tau|$ hence $|\tau| \subseteq A$, and if $\dim \tau \geq k+1$ then $|\tau|$ is a subset of B. Therefore, $A \cup B$ consists of the relative interiors of all faces of dimension k and greater; that is,

$$A \cup B = |\operatorname{Skel}^c_{k-1} X|.$$

We can now apply the Mayer–Vietoris theorem, with unreduced homology. This gives us the following long exact sequence:

$$\begin{aligned} \cdots \longrightarrow H_i\Big(\bigsqcup_{\sigma \in X_k} |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|\Big) \longrightarrow \\ H_i\Big(\bigsqcup_{\sigma \in X_k} |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|\Big) \oplus H_i(|\operatorname{Skel}^c_k X|) \longrightarrow H_i(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \cdots . \end{aligned}$$

Since unreduced homology of a disconnected space can be decomposed as a direct sum over components, we can simplify this long exact sequence:

$$\begin{aligned} \cdots \longrightarrow \bigoplus_{\sigma \in X_k} H_i\big(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|\big) \longrightarrow \\ \Big(\bigoplus_{\sigma \in X_k} H_i(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|)\Big) \oplus H_i(|\operatorname{Skel}^c_k X|) \longrightarrow H_i(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \cdots . \end{aligned}$$

While this is a perfectly good long exact sequence, there are some steps we can do to "clean it up" and make it more convenient for later proofs. The first step is to convert from unreduced to reduced homology, which affects this sequence in degrees 0 and -1. For each face $\sigma \in X_k$, there are two cases to consider: either σ is not a facet, or it is.

If σ is not a facet, then $|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|$ is non-empty, so $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|) \neq 0$. The map from $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|)$ to $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|) = R$ is non-zero, so we can remove a summand of R from each of these two groups while preserving exactness; the result is that these unreduced homology groups become reduced.

On the other hand, if σ is a facet, then $|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|$ is empty, so the homology group $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|)$ is 0. Since $H_0(|\operatorname{Skel}^c_{k-1} X|)$ is non-zero, and the map from $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|)$ to $H_0(|\operatorname{Skel}^c_{k-1} X|)$ is non-zero, we can remove a summand of R from each of these two groups. If we also add a summand of R to each of $H_0(|\operatorname{Skel}^c_{k-1} X|)$ and $H_{-1}(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|) = 0$, then exactness is preserved, no net change has occurred to $H_0(|\operatorname{Skel}^c_{k-1} X|)$, and the two homology groups $H_{-1}(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|$ and $H_0(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|)$ become reduced.

There is one more change to make in converting from unreduced to reduced homology: the map from $H_0(|\operatorname{Skel}^c_k X|)$ to $H_0(|\operatorname{Skel}^c_{k-1} X|)$ is non-zero (both groups are non-zero because we assume $k < d$), so we may remove a summand of R from each to replace both of these unreduced homology groups with reduced homology.

With the observation that $\widetilde{H}_i(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|) = 0$ for all i, our long exact sequence now becomes:

$$\cdots \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_i\big(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|\big) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_k X|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \cdots .$$

There is one last thing to do to this long exact sequence: simplifying $\widetilde{H}_i\big(|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|\big)$. To do this, let us call on the Mayer–Vietoris theorem again. Assume that σ is not a facet, so $|\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \setminus |\sigma|$ is non-empty. Define two new spaces:

$$A' \coloneqq |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma|,$$
$$B' \coloneqq |\operatorname{star}_X \sigma| \setminus |\sigma| = |\operatorname{link}_X \sigma|$$

Since $|\sigma| \subseteq |\operatorname{star}_{\operatorname{bary}(X)} v_\sigma| \subseteq |\operatorname{star}_X \sigma|$ — for example, see Fig. 3.7 — the intersection

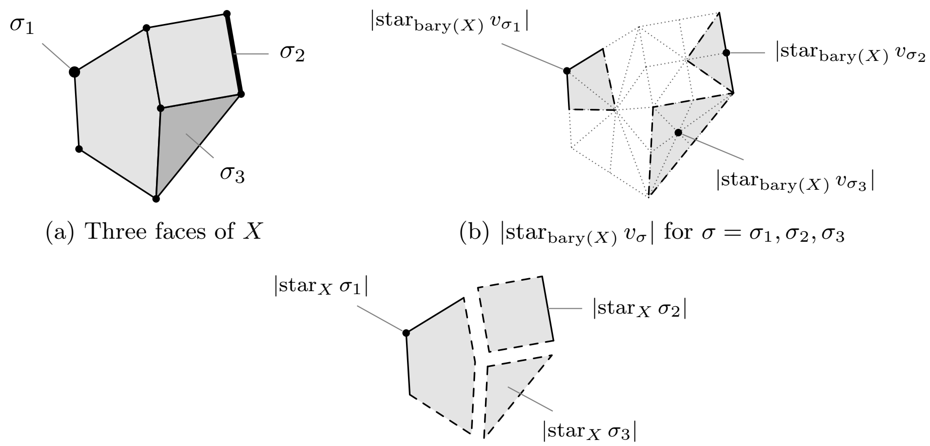

(a) Three faces of X

(b) $|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma|$ for $\sigma = \sigma_1, \sigma_2, \sigma_3$

(c) $|\mathrm{star}_X\, \sigma|$ for $\sigma = \sigma_1, \sigma_2, \sigma_3$

Figure 3.7: The spaces $|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma|$ and $|\mathrm{star}_X\, \sigma|$ in the proof of Theorem 3.4.1, for various choices of σ

$A' \cap B'$ is $|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|$, and the union $A' \cup B'$ is simply $|\mathrm{star}_X\, \sigma|$. Therefore, the Mayer–Vietoris long exact sequence with reduced homology for A' and B' is the following:

$$\begin{aligned} \cdots \longrightarrow \widetilde{H}_{i+1}(|\mathrm{star}_X\, \sigma|) &\longrightarrow \widetilde{H}_i(|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|) \\ &\longrightarrow \widetilde{H}_i(|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma|) \oplus \widetilde{H}_i(|\mathrm{star}_X\, \sigma| \setminus |\sigma|) \longrightarrow \widetilde{H}_i(|\mathrm{star}_X\, \sigma|) \longrightarrow \cdots . \end{aligned}$$

But stars of non-empty faces are contractible, so this becomes:

$$\cdots \longrightarrow 0 \longrightarrow \widetilde{H}_i(|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|) \longrightarrow \widetilde{H}_i(|\mathrm{star}_X\, \sigma| \setminus |\sigma|) \longrightarrow 0 \longrightarrow \cdots .$$

Therefore, $\widetilde{H}_i(|\mathrm{star}_{\mathrm{bary}(X)}\, v_\sigma| \setminus |\sigma|) \cong \widetilde{H}_i(|\mathrm{star}_X\, \sigma| \setminus |\sigma|) = \widetilde{H}_i(|\mathrm{link}_X\, \sigma|)$ for all non-facets σ. Finally, note that if σ is a facet, this isomorphism still holds, since all spaces involved are empty.

Hence we obtain the desired long exact sequence. ■

3.5 Complexes characterised by link conditions

Now that we have a long exact sequence relating co-skeletons and links, in this section we will consider some particular families of polytopal complexes that are characterised by links of faces having zero homology in some degrees.

3.5.1 Cohen–Macaulay complexes

A pure, d-dimensional polytopal complex X is *Cohen–Macaulay* if for every face σ, including $\sigma = \varnothing$, and every $i < d - \dim \sigma - 1$, the homology $\widetilde{H}_i(|\mathrm{link}_\Delta \sigma|)$ is 0. If $\Delta := X$ is a simplicial complex, this condition is equivalent to the Stanley–Reisner ring of Δ being a Cohen–Macaulay ring [Bjö95, p. 1855].

Remark 3.5.1. Cohen–Macaulayness is also a topological condition: X is Cohen–Macaulay if and only if $\widetilde{H}_i(|X|)$ and $\widetilde{H}_i(|X|, |X| \setminus \{p\})$ are both 0 for all $i < \dim X$ and all points $p \in |X|$. This was proved by Munkres [Mun84b, Corollary 3.4] for simplicial complexes, but a similar proof applies to polytopal complexes as well, with the following modified version of [Mun84b, Lemma 3.3]:

Lemma 3.5.2. *If X is a polytopal complex, σ a face of X and p a point in the relative interior of σ, then* $\widetilde{H}_i(|\mathrm{link}_X \sigma|) \cong \widetilde{H}_{i+\dim\sigma}(|\mathrm{link}_{\mathrm{bary}(X)} v_\sigma|) \cong \widetilde{H}_{i+\dim\sigma+1}(|X|, |X| \setminus \{p\})$.

Proof. We saw in the proof of Theorem 3.4.1 that

$$\begin{aligned}\widetilde{H}_i(|\mathrm{link}_X \sigma|) &\cong \widetilde{H}_i(|\mathrm{star}_{\mathrm{bary}(X)} v_\sigma| \setminus |\sigma|)\\ &= \widetilde{H}_i\big(|\mathrm{link}_{\mathrm{bary}(X)} v_\sigma| \setminus (|\sigma| \setminus \{v_\sigma\})\big)\\ &\cong \widetilde{H}_i(|\text{s-link}_{\mathrm{bary}(X)} v_\sigma| \setminus |\mathrm{bary}(\partial\sigma)|),\end{aligned}$$

where "s-link" denotes the simplicial link. Since $\mathrm{bary}(\partial\sigma)$ is an induced subcomplex of the simplicial complex $\text{s-link}_{\mathrm{bary}(X)} v_\sigma$, Lemma 3.2.1 says that

$$|\text{s-link}_{\mathrm{bary}(X)} v_\sigma| \setminus |\mathrm{bary}(\partial\sigma)| \simeq |\text{s-link}_{\mathrm{bary}(X)} v_\sigma - \mathrm{bary}(\partial\sigma)|.$$

It is not hard to check that

$$\text{s-link}_{\text{bary}(X)} v_\sigma = \big(\text{s-link}_{\text{bary}(X)} v_\sigma - \text{bary}(\partial\sigma)\big) * \text{bary}(\partial\sigma),$$

where $*$ denotes simplicial join [Mun84a, §62], so [Mun84a, Theorem 62.5] implies that

$$\widetilde{H}_i(|\text{s-link}_{\text{bary}(X)} v_\sigma - \text{bary}(\partial\sigma)|) \cong \widetilde{H}_{i+\dim\sigma}(|\text{s-link}_{\text{bary}(X)} v_\sigma|).$$

Finally, since $\text{bary}(X)$ is a simplicial complex, we can apply [Mun84b, Lemma 3.3] to conclude that

$$\widetilde{H}_{i+\dim\sigma}(|\text{s-link}_{\text{bary}(X)} v_\sigma|) \cong \widetilde{H}_{i+\dim\sigma+1}(|X|, |X| \setminus \{v_\sigma\}),$$

and note that although v_σ is typically defined to be the barycentre of σ, it can be chosen to be any point in the relative interior of σ. ■

The following theorem allows us to determine whether a complex is Cohen–Macaulay by considering its co-skeletons.

Theorem 3.5.3. *A d-dimensional polytopal complex X is Cohen–Macaulay if and only if it is pure and $\widetilde{H}_i(|\text{Skel}^c_k X|) = 0$ for all $i < d-k-1$ and all $k = -1, \ldots, d$.*

Note that this theorem combined with Corollary 3.3.2 means that the homology of $|\text{Skel}^c_k X|$ is entirely concentrated in degree $d-k-1$.

Proof of Theorem 3.5.3. To begin with, assume X is Cohen–Macaulay. We will prove that $\widetilde{H}_i(|\text{Skel}^c_k X|) = 0$ for $i < d-k-1$ by induction on k.

For the base case, when $k = -1$, we have

$$\widetilde{H}_i(|\text{Skel}^c_{-1} X|) = \widetilde{H}_i(|X|) = \widetilde{H}_i(|\text{link}_X \varnothing|),$$

which is zero for all $i < d$ since X is Cohen–Macaulay.

Now, suppose $k > -1$, and consider this part of the long exact sequence from Theorem 3.4.1:

$$\bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_k X|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_{k-1} X|).$$

Since X is Cohen–Macaulay and $i < d - k - 1$, we know that $\widetilde{H}_i(|\operatorname{link} \sigma|)$ is 0. Also, since i is less than $d - k - 1$, it is certainly less than $d - (k - 1) - 1$, so by induction, $\widetilde{H}_i(|\operatorname{Skel}^c_{k-1} X|)$ is 0. Therefore, $\widetilde{H}_i(|\operatorname{Skel}^c_k X|)$ must be 0 too, and we are done with this direction of the proof.

Now, for the reverse direction, assume that $\widetilde{H}_i(|\operatorname{Skel}^c_k X|) = 0$ for all $i < d-k-1$ and all $k = -1, \dots, d$. Consider this part of the long exact sequence for $k \geq 0$:

$$\widetilde{H}_{i+1}(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_k X|)$$

By our assumption, $\widetilde{H}_i(|\operatorname{Skel}^c_k X|) = 0$. Also, since $i < d - k - 1$, we have $i + 1 < d - (k - 1) - 1$, so our assumption tells us that $\widetilde{H}_{i+1}(|\operatorname{Skel}^c_{k-1} X|) = 0$ too. Therefore, $\bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|)$ is 0, so $\widetilde{H}_i(|\operatorname{link} \sigma|)$ must be 0 for all $i < d - k - 1 = d - \dim \sigma - 1$.

This is true for all faces σ of dimension $k \geq 0$, which leaves the case $\sigma = \varnothing$. In this case,

$$\widetilde{H}_i(|\operatorname{link} \sigma|) = \widetilde{H}_i(|X|) = \widetilde{H}_i(|\operatorname{Skel}^c_{-1} X|) = 0.$$

Therefore, X is Cohen–Macaulay. ■

Remark 3.5.4. This theorem has potential computational applications. Naïvely, to check whether some complex X is Cohen–Macaulay, one must compute the homology groups of the link of every face of X — if $\dim X = d$, there are at least 2^{d+1} faces. But with this result, one only needs to compute homologies of the kth co-skeletons of X for the $d + 2$ values of k between -1 and d, or equivalently (by Lemma 3.3.1) the simplicial complexes $\operatorname{bary}(X) - \operatorname{bary}(\operatorname{Skel}_k X)$. The tradeoff is that these simplicial complexes are much larger and more complicated than the links of faces: for instance, the number of vertices of $\operatorname{bary}(X) - \operatorname{bary}(\operatorname{Skel}_k X)$ is at least on the order of 2^d.

(a) A Cohen–Macaulay simplicial complex

(b) A non-Cohen–Macaulay simplicial complex

Figure 3.8: Two simplicial complexes whose skeletons have isomorphic homology groups

Remark 3.5.5. We saw in Corollary 3.3.7 that neighbourly spheres are characterised by the homology groups of either their skeletons or their co-skeletons. However, for Cohen–Macaulay complexes, the homology groups of skeletons are not sufficient for a characterisation: for example, Fig. 3.8 shows two complexes, one Cohen–Macaulay and one not, whose k-skeletons have isomorphic ith homology groups for all k and i.

Theorem 3.5.3 says that if X is Cohen–Macauly, almost all of the long exact sequence from Theorem 3.4.1 is 0. The remaining non-zero terms are captured in the following corollary:

Corollary 3.5.6. *If X is a Cohen–Macaulay polytopal complex and $k = 0, \ldots, d$, we have the following short exact sequence:*

$$0 \longrightarrow \widetilde{H}_{d-k}(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_{d-k-1}(|\operatorname{link} \sigma|) \longrightarrow \widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_k X|) \longrightarrow 0.$$

Proof. Take the following segment of the long exact sequence from Theorem 3.4.1:

$$\widetilde{H}_{d-k}(|\operatorname{Skel}^c_k X|) \longrightarrow \widetilde{H}_{d-k}(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_{d-k-1}(|\operatorname{link} \sigma|) \longrightarrow$$
$$\widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_k X|) \longrightarrow \widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_{k-1} X|).$$

The first term in this segment, $\widetilde{H}_{d-k}(|\operatorname{Skel}^c_k X|)$, is 0 by Corollary 3.3.2. The last term, $\widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_{k-1} X|)$, is 0 by Theorem 3.5.3. ∎

We can combine these short exact sequences into a long exact sequence, using the following homological algebra fact:

Lemma 3.5.7. *Suppose we have two exact sequences:*

$$\cdots \longrightarrow A_2 \xrightarrow{a_2} A_1 \xrightarrow{a_1} C \longrightarrow 0 \qquad \textit{and} \qquad 0 \longrightarrow C \xrightarrow{b_0} B_{-1} \xrightarrow{b_{-1}} B_{-2} \longrightarrow \cdots .$$

Then the horizontal sequence in the following diagram is exact:

$$\begin{tikzcd}
 & & 0 \arrow[dr] & & 0 & & \\
 & & & C \arrow[ur] \arrow[dr, "b_0"] & & & \\
\cdots \arrow[r] & A_2 \arrow[r, "a_2"] & A_1 \arrow[ur, "a_1"] \arrow[rr, dashed, "b_0 \circ a_1"'] & & B_{-1} \arrow[r, "b_{-1}"] & B_{-2} \arrow[r] & \cdots .
\end{tikzcd}$$

Proof. Since a_1 is surjective, $\operatorname{im}(b_0 \circ a_1) = \operatorname{im}(b_0) = \ker(b_{-1})$. Since b_0 is injective, $\ker(b_0 \circ a_1) = \ker a_1 = \operatorname{im} a_2$. ∎

Corollary 3.5.8. *If X is Cohen–Macaulay, we have the following long exact sequence:*

$$0 \longrightarrow \widetilde{H}_d(|X|) \longrightarrow \bigoplus_{\sigma \in X_0} \widetilde{H}_{d-1}(|\operatorname{link}\sigma|) \longrightarrow \bigoplus_{\sigma \in X_1} \widetilde{H}_{d-2}(|\operatorname{link}\sigma|) \longrightarrow$$
$$\cdots \longrightarrow \bigoplus_{\sigma \in X_d} \widetilde{H}_{-1}(|\operatorname{link}\sigma|) \longrightarrow R \longrightarrow 0.$$

Proof. Use Lemma 3.5.7 to stitch together the short exact sequences from Corollary 3.5.6. The start of the sequence looks like this:

$$\begin{tikzcd}[column sep=small]
 & & 0 \arrow[dr] & & 0 & & & \\
 & & & \widetilde{H}_{d-1}(|\operatorname{Skel}^c_0 X|) \arrow[ur] \arrow[dr] & & & & \\
0 \arrow[r] & \widetilde{H}_d(|\operatorname{Skel}^c_{-1} X|) \arrow[r] & \bigoplus\limits_{\sigma \in X_0} \widetilde{H}_{d-1}(|\operatorname{link}\sigma|) \arrow[ur] \arrow[rr, dashed] & & \bigoplus\limits_{\sigma \in X_1} \widetilde{H}_{d-2}(|\operatorname{link}\sigma|) \arrow[dr] \arrow[rr, dashed] & & \bigoplus\limits_{\sigma \in X_2} \widetilde{H}_{d-3}(|\operatorname{link}\sigma|) \arrow[r, dashed] & \cdots \\
 & & & & & \widetilde{H}_{d-2}(|\operatorname{Skel}^c_1 X|) \arrow[ur] \arrow[dr] & & \\
 & & & & 0 \arrow[ur] & & 0 &
\end{tikzcd}$$

And the end like this:

$$\begin{tikzcd}
 & 0 \arrow[dr] & & 0 & & & & \\
 & & \widetilde{H}_1(|\operatorname{Skel}^c_{d-2} X|) \arrow[ur] \arrow[dr] & & & & & \\
\cdots \arrow[r, dashed] & \bigoplus\limits_{\sigma \in X_{d-2}} \widetilde{H}_1(|\operatorname{link} \sigma|) \arrow[ur] \arrow[rr, dashed] & & \bigoplus\limits_{\sigma \in X_{d-1}} \widetilde{H}_0(|\operatorname{link} \sigma|) \arrow[rr, dashed] \arrow[dr] & & \bigoplus\limits_{\sigma \in X_{d}} \widetilde{H}_{-1}(|\operatorname{link} \sigma|) \arrow[r] & \widetilde{H}_{-1}(|\operatorname{Skel}^c_{d} X|) \arrow[r] & 0 \\
 & & & & \widetilde{H}_0(|\operatorname{Skel}^c_{d-1} X|) \arrow[ur] \arrow[dr] & & & \\
 & & & 0 \arrow[ur] & & 0 & &
\end{tikzcd}$$

■

Note that this exact sequence is reminiscent of the "partition complex", a chain complex defined in [AY21, Definition 25].

Remark 3.5.9. If the definition of a Cohen–Macaulay polytopal complex is modified to allow $\widetilde{H}_i(|\operatorname{link}_X \sigma|)$ to be non-zero when $\sigma = \varnothing$, we get the definition of a *Buchsbaum* complex. Buchsbaum complexes can also be characterised by homological properties of their co-skeletons, although the statement is not as simple as for Cohen–Macaulay complexes in Theorem 3.5.3: a complex X is Buchsbaum if and only if the map $\widetilde{H}_i(|\operatorname{Skel}^c_k X|) \to \widetilde{H}_i(|X|)$ induced by the inclusion of topological spaces is an isomorphism for all $i < d-k-1$ and a surjection for $i = d-k-1$, for all $k = -1, \ldots, d$. This follows fairly easily from the long exact sequence in Theorem 3.4.1, although we leave the details to the reader.

3.5.2 Leray complexes

A simplicial complex Δ is *r-Leray* if every induced subcomplex Λ of Δ has $\widetilde{H}_i(|\Lambda|) = 0$ for $i \geq r$. For example, if Δ is the nerve of a family of convex open subsets of $\mathbb{R}^r$, then it follows from the nerve theorem (Theorem 3.2.3) that Δ is r-Leray [KM06].

Equivalently, Δ is r-Leray if the homology $\widetilde{H}_i(|\operatorname{link}_\Delta \sigma|)$ is 0 for every face $\sigma \in \Delta$ (including the empty face) and every $i \geq r$ [KM06, Proposition 3.1]. We will use this condition to generalise r-Leray-ness to polytopal complexes: we will say a complex X is

r-Leray if $\widetilde{H}_i(|\operatorname{link}_X \sigma|) = 0$ for all faces σ (including $\sigma = \varnothing$) and all $i \geq r$. Note that this condition is not equivalent to the condition on induced subcomplexes in the non-simplicial case: for example, if X is a square, every non-empty link is contractible so X is 0-Leray, but the subcomplex induced by a pair of diagonally opposite vertices is not connected.

The following theorem provides a characterisation of r-Leray complexes in terms of their co-skeletons.

Theorem 3.5.10. *A polytopal complex X is r-Leray if and only if $\widetilde{H}_i(|\operatorname{Skel}_k^c X|) = 0$ for all $i \geq r$ and all $k = -1, \ldots, d$.*

Proof. For the forward direction, assume X is r-Leray and let $i \geq r$. When $k = -1$,

$$\widetilde{H}_i(|\operatorname{Skel}_{-1}^c X|) = \widetilde{H}_i(|X|) = \widetilde{H}_i(|\operatorname{link}_X \varnothing|),$$

which is 0, since X is r-Leray.

For $k \geq 0$, consider the following subsequence of the long exact sequence from Theorem 3.4.1:

$$\bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}_k^c X|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}_{k-1}^c X|).$$

Since X is r-Leray, $\bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|)$ is 0. Therefore, the map from $\widetilde{H}_i(|\operatorname{Skel}_k^c X|)$ to $\widetilde{H}_i(|\operatorname{Skel}_{k-1}^c X|)$ is injective. Since this is true for all $k = 0, \ldots, d$, we get a series of injections:

$$\widetilde{H}_i(|\operatorname{Skel}_k^c X|) \hookrightarrow \widetilde{H}_i(|\operatorname{Skel}_{k-1}^c X|) \hookrightarrow \widetilde{H}_i(|\operatorname{Skel}_{k-2}^c X|) \hookrightarrow \cdots$$
$$\cdots \hookrightarrow \widetilde{H}_i(|\operatorname{Skel}_0^c X|) \hookrightarrow \widetilde{H}_i(|\operatorname{Skel}_{-1}^c X|).$$

But we saw above that $\widetilde{H}_i(|\operatorname{Skel}_{-1}^c X|) = 0$. Therefore, $\widetilde{H}_i(|\operatorname{Skel}_k^c X|)$ injects into 0, so it must itself be 0, for all $k = 0, \ldots, d$.

For the reverse direction, assume that $\widetilde{H}_i(|\operatorname{Skel}_k^c X|) = 0$ for all $k = -1, \ldots, d$ and all

$i \geq r$. When $k \geq 0$, consider this part of the long exact sequence:

$$\widetilde{H}_{i+1}(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_i(|\operatorname{link} \sigma|) \longrightarrow \widetilde{H}_i(|\operatorname{Skel}^c_k X|).$$

By our assumptions, the two outer terms here are both 0, hence the middle term is 0 as well, so $\widetilde{H}_i(|\operatorname{link} \sigma|) = 0$ for all k-faces σ. This works for all $k = 0, \ldots, d$, which only leaves $k = -1$: the only (-1)-dimensional face is $\varnothing$, and

$$\widetilde{H}_i(|\operatorname{link}_X \varnothing|) = \widetilde{H}_i(|X|) = \widetilde{H}_i(|\operatorname{Skel}^c_{-1} X|)$$

which we assumed to be 0.

Therefore, $\widetilde{H}_i(|\operatorname{link}_X \sigma|) = 0$ for all faces σ and all $i \geq r$, so X is r-Leray. ■

3.5.3 Stacked balls

A *homology sphere* is a polytopal complex X of dimension d where for every face σ including $\sigma = \varnothing$,

$$\widetilde{H}_i(|\operatorname{link} \sigma|) = \begin{cases} 0 & \text{if } i \neq d - \dim \sigma - 1 \\ R & \text{if } i = d - \dim \sigma - 1. \end{cases}$$

A *homology manifold* is defined similarly but only considering $\sigma \neq \varnothing$. A *homology ball* of dimension d is a complex X where:

- $\widetilde{H}_i(|X|) = 0$ for all i,
- for every non-empty face σ,

$$\widetilde{H}_i(|\operatorname{link} \sigma|) = \begin{cases} 0 & \text{if } i \neq d - \dim \sigma - 1 \\ R \text{ or } 0 & \text{if } i = d - \dim \sigma - 1, \end{cases}$$

  and
- the set of faces σ with $\widetilde{H}_{d-\dim \sigma - 1}(|\operatorname{link} \sigma|) = 0$ forms a subcomplex of X that is a

$(d-1)$-dimensional homology sphere.

If $\widetilde{H}_{d-\dim\sigma-1}(|\operatorname{link}\sigma|)$ is 0, σ is called a *boundary face*, and if this homology is R, σ is an *interior face*.

Suppose X is a homology ball of dimension d. If every face of X of dimension less than or equal to $d-s-1$ is a boundary face, then X is said to be *s-stacked*. For example, if X is 1-stacked, then the interior faces must all have dimension d or $d-1$; this condition is sometimes simply called "stacked" for simplicial complexes (e.g. [Kal87]), or "capped" for cubical complexes (e.g. [BB98]). Simplicial s-stacked balls are well studied due to their connection to the Lower Bound Conjecture — see e.g. [MN13].

Just like Cohen–Macaulay complexes, we can characterise s-stacked balls using their co-skeletons.

Theorem 3.5.11. *Suppose X is a homology ball with dimension d. Then X is s-stacked if and only if $\widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_k X|) = 0$ for all $k \leq d-s-1$ (or equivalently $\widetilde{H}_j(|\operatorname{Skel}^c_{d-j-1} X|) = 0$ for all $j \geq s$, where $j = d-k-1$).*

Proof. By definition, a k-face σ is a boundary face if and only if $\widetilde{H}_{d-k-1}(|\operatorname{link}\sigma|)$ is 0. Therefore, X is s-stacked if and only if

$$\bigoplus_{\sigma \in X_k} \widetilde{H}_{d-k-1}(|\operatorname{link}\sigma|) = 0$$

for all $k \leq d-s-1$.

We could apply this fact directly to the long exact sequence from Theorem 3.4.1; however, since X is a homology ball, it is Cohen–Macaulay, so we can take a shortcut by using the short exact sequence from Corollary 3.5.6:

$$0 \longrightarrow \widetilde{H}_{d-k}(|\operatorname{Skel}^c_{k-1} X|) \longrightarrow \bigoplus_{\sigma \in X_k} \widetilde{H}_{d-k-1}(|\operatorname{link}\sigma|) \longrightarrow \widetilde{H}_{d-k-1}(|\operatorname{Skel}^c_k X|) \longrightarrow 0.$$

The inner term of this short exact sequence is 0 for $k \leq d-s-1$ if and only if both outer terms are 0 for the same range of k, which is equivalent to the claimed condition. ∎

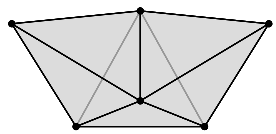

(a) A 1-stacked sphere, the boundary of the ball made from three stacked 3-simplices

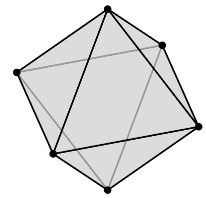

(b) A sphere that is not 1-stacked, the boundary of an octahedron

Figure 3.9: Two spheres with homeomorphic co-skeletons

Remark 3.5.12. A closely related notion is an *s-stacked sphere*, which is a complex that is the boundary of some s-stacked ball. Unfortunately, s-stacked spheres cannot be distinguished by topological features of their co-skeletons, at least for simplicial complexes. For example, the two spheres shown in Fig. 3.9 have homeomorphic kth co-skeletons for all k, but only one of the spheres is 1-stacked as a simplicial complex.

3.6 CAT(0) cubical complexes

For a further application of co-skeletons, let us consider some families of cubical complexes.

One of the key features of a cubical complex is its *hyperplanes*. If we associate an r-dimensional cube with the space $[0,1]^r$ in $\mathbb{R}^r$, then the *ith hyperplane* of that cube is the subspace where $x_i = \frac{1}{2}$ — see Fig. 3.10b. A hyperplane in a cubical complex $\square$ is a maximally connected cubical complex obtained by glueing together hyperplanes of its component cubes where they meet along faces — see Fig. 3.11. (We should note that in general, issues can arise when a hyperplane "intersects itself", in which case the hyperplane is not a true polytopal complex by our definitions. However, the focus of this section will be on CAT(0) cubical complexes, in which hyperplanes behave nicely — see e.g. Lemma 3.6.6 — so we will not dwell on this issue.)

Since each hyperplane of a cubical complex is itself a cubical complex, it also has its own hyperplanes. Therefore, we will define an *iterated hyperplane* of $\square$ to be a hyperplane of a hyperplane of ... of $\square$. We will say the original hyperplanes are the "1st iterated hyperplanes", the hyperplanes of the hyperplanes are the "2nd iterated hyperplanes",

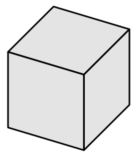

(a) A 3-cube

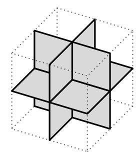

(b) The hyperplanes (aka 1st iterated hyperplanes) of a 3-cube

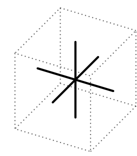

(c) The 2nd iterated hyperplanes of a 3-cube

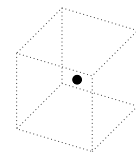

(d) The 3rd iterated hyperplane of a 3-cube

Figure 3.10: Hyperplanes and iterated hyperplanes of a cube

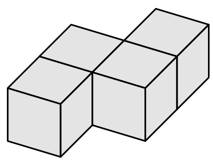

(a) $\square$

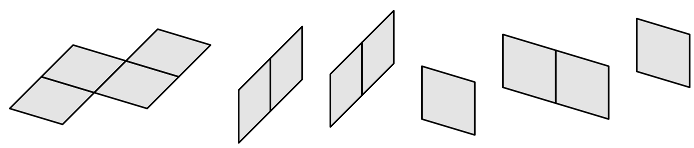

(b) The six hyperplanes of $\square$

Figure 3.11: A cubical complex and its hyperplanes

and so on. See Figs. 3.10a to 3.10d. As these figures show, an iterated hyperplane can sometimes be obtained in more than one way, when considered as a subspace of $\square$ — for example, each line segment in Fig. 3.10c is a hyperplane of two different squares in Fig. 3.10b.

The following proposition reveals the connection between co-skeletons and iterated hyperplanes.

Proposition 3.6.1. *Suppose* $\square$ *is a cubical complex. Then* $|\mathrm{bary}(\square) - \mathrm{bary}(\mathrm{Skel}_k \square)|$ *is equal to the union of the* $(k+1)$*th iterated hyperplanes of* $\square$*, considered as subspaces of* $\square$*.*

For example, compare Fig. 3.10b and Fig. 3.12. Note that in combination with Lemma 3.3.1, this proposition tells us that $|\mathrm{Skel}_k^c \square|$ is homotopy equivalent to the union of the $(k+1)$th iterated hyperplanes.

Proof of Proposition 3.6.1. We will prove this by showing that the intersections of $|\mathrm{bary}(\square) - \mathrm{bary}(\mathrm{Skel}_k \square)|$ and of the union of the $(k+1)$th iterated hyperplanes with each cube of

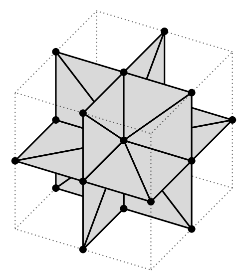

Figure 3.12: $\text{bary}([0,1]^3) - \text{bary}(\text{Skel}_0[0,1]^3)$

$\square$ are the same.

First, consider the union of the $(k+1)$th hyperplanes. In each cube, identified with $[0,1]^r$, a single $(k+1)$th iterated hyperplane is the result of setting $(k+1)$ of the coordinates to be $\frac{1}{2}$; therefore, the union of all $(k+1)$th iterated hyperplanes is the space

$$Z \coloneqq \left\{x \in |[0,1]^r| : \text{at least } k+1 \text{ coordinates of } x \text{ are } \tfrac{1}{2}\right\}.$$

Now, consider $|\text{bary}(\square) - \text{bary}(\text{Skel}_k \square)|$. The intersection of this space with some cube of $\square$, identified with $|[0,1]^r|$, is

$$B \coloneqq |\text{bary}([0,1]^r) - \text{bary}(\text{Skel}_k[0,1]^r)|.$$

For each face σ of $[0,1]^r$, we have a vertex v_σ of $\text{bary}([0,1]^r)$ at the barycenter of σ; the coordinates of v_σ will all be 0, 1 or $\frac{1}{2}$, and the number of coordinates equal to $\frac{1}{2}$ is $\dim \sigma$. The remaining vertices of $\text{bary}([0,1]^r)$ after deleting $\text{bary}(\text{Skel}_k[0,1]^r)$, therefore, are precisely those which have at least $k+1$ coordinates equal to $\frac{1}{2}$.

For the rest of this proof, we will show that $B = Z$. First, we claim that $B \subseteq Z$. For each face $\{v_{\sigma_1}, \ldots, v_{\sigma_m}\}$ of B, the set $\{\sigma_1, \ldots, \sigma_m\}$ is a chain, so it has a minimal element — without loss of generality, say this is σ_1. By construction, the dimension of σ_1 is at least $k+1$, so there is some index set I of at least $k+1$ elements so that $(v_{\sigma_1})_i = \frac{1}{2}$ for all $i \in I$. Since $\{\sigma_1, \ldots, \sigma_m\}$ is a chain, every v_{σ_j} for $j = 1, \ldots, m$ must also satisfy $(v_{\sigma_j})_i = \frac{1}{2}$ for all $i \in I$. Therefore, the entire set $\{v_{\sigma_1}, \ldots, v_{\sigma_m}\}$ lies in the affine subspace

$\{x \in \mathbb{R}^r : x_i = \frac{1}{2} \text{ for all } i \in I\}$, so the convex hull of this set is contained in Z. Thus $B \subseteq Z$.

Conversely, let z be a point in Z. By definition, there is an index set I of size $k+1$ so that $z_i = \frac{1}{2}$ for all $i \in I$. The set of all points of $[0,1]^r$ satisfying $x_i = \frac{1}{2}$ for $i \in I$ is an iterated hyperplane of $[0,1]^r$, and its barycentric subdivision is the subcomplex of $\mathrm{bary}([0,1]^r) - \mathrm{bary}(\mathrm{Skel}_k[0,1]^r)$ generated by vertices satisfying $(v_\sigma)_i = \frac{1}{2}$ for $i \in I$. Therefore, z is in B, so $Z \subseteq B$.

In summary, the space $|\mathrm{bary}(\square) - \mathrm{bary}(\mathrm{Skel}_k \square)|$ and the union of the $(k+1)$th iterated hyperplanes agree within each cube of $\square$, so they must be the same subspace. ■

There is a class of cubical complexes where hyperplanes and co-skeletons work particularly well. In general, the link of any vertex in a cubical complex is isomorphic as a poset to the non-empty faces of a simplicial complex. A cubical complex is *CAT(0)* if:

- it is simply connected, and
- the link of every vertex is isomorphic to a *flag* simplicial complex.

A simplicial complex is flag if the vertex set of every clique in its 1-skeleton is the vertex set of a face.

CAT(0) cubical complexes were first studied from the perspective of metric spaces: equivalently, a cubical complex is CAT(0) if and only if it is simply connected and has non-positive curvature at every point (see [BH99, Theorem 5.4 & Theorem 5.18]; see also [Gro87]). However, finite CAT(0) cubical complexes are also interesting from a purely combinatorial and topological perspective [AOS12; Hag14; Sag95; Rol16].

One particular construction associated to a CAT(0) cubical complex is its *crossing complex*, introduced in [Row23] (Chapter 2): given a CAT(0) cubical complex $\square$, its crossing complex, here denoted $\mathrm{Cross}(\square)$, is a simplicial complex whose vertices are in bijection with the hyperplanes of $\square$, and a set of vertices forms a face if the intersection of the corresponding hyperplanes is non-empty. (In other words, the crossing complex is the nerve of the hyperplanes, as defined in Theorem 3.2.3.)

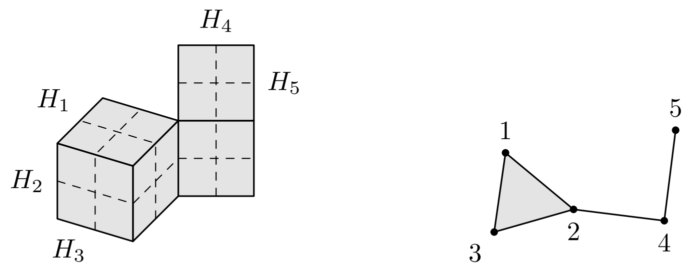

(a) $\square$, with hyperplanes labelled (b) $\mathrm{Cross}(\square)$

Figure 3.13: A CAT(0) cubical complex and its crossing complex

We record here some basic facts about CAT(0) cubical complexes and their crossing complexes.

Lemma 3.6.2 ([BH99, Corollary 1.5], [Gro87, Section 4])**.** *Any CAT(0) cubical complex is contractible (not just simply connected).*

Lemma 3.6.3 ([Hag14, Lemma 2.15])**.** $\mathrm{Cross}(\square)$ *is always a flag simplicial complex. In particular, it is the clique complex of the "crossing graph" defined by Hagen [Hag14, Definition 2.16].*

Lemma 3.6.4 (Proposition 2.6.5, Corollary 2.6.6)**.** *There is a bijection between the facets of $\square$ and the facets of* $\mathrm{Cross}(\square)$*. If a facet of $\square$ has dimension k, the corresponding facet of* $\mathrm{Cross}(\square)$ *has dimension $k-1$. As a consequence,* $\dim \mathrm{Cross}(\square) = \dim \square - 1$*, and $\square$ is pure if and only if* $\mathrm{Cross}\,\square$ *is pure.*

Lemma 3.6.5 ([Sag95, Theorem 4.11])**.** *Every hyperplane (and thus every iterated hyperplane) in a CAT(0) cubical complex is also CAT(0).*

Ardila et al. [AOS12, p. 9] described a way to embed a CAT(0) cubical complex with m hyperplanes into the cube $[0,1]^m$. This embedding is essentially canonical, up to the choice of a "root vertex" and a choice of labelling for the hyperplanes as $H_1, \ldots, H_m$.

Lemma 3.6.6 (Proposition 2.5.7)**.** *In this embedding of $\square$ into $[0,1]^m$, the hyperplane H_i of $\square$ is the intersection of $\square$ with the ith hyperplane of $[0,1]^m$.*

The resulting embedding of a hyperplane H_i into the ith hyperplane of $[0,1]^m$ agrees with Ardila et al.'s canonical embedding of H_i as a CAT(0) cubical complex into a cube; thus the iterated hyperplanes of $\square$ are themselves each an intersection of $\square$ with some set of hyperplanes of $[0,1]^m$. Therefore, if S is a subset of $\{1,\dots,m\}$ of size $(k+1)$, we define

$$\begin{aligned} H_S &\coloneqq \square \cap \bigcap_{i \in S} (i\text{th hyperplane of } [0,1]^m) \\ &= \bigcap_{i \in S} \big(\square \cap (i\text{th hyperplane of } [0,1]^m)\big) = \bigcap_{i \in S} H_i. \end{aligned}$$

Either this is a $(k+1)$th iterated hyperplane of $\square$, or it is empty, and all $(k+1)$th iterated hyperplanes of $\square$ are of this form. By definition of $\mathrm{Cross}(\square)$, H_S is an iterated hyperplane if and only if S is a face of $\mathrm{Cross}(\square)$; therefore, the $(k+1)$th iterated hyperplanes of $\square$ correspond to k-faces of $\mathrm{Cross}(\square)$, for $k = 0,\dots,d-1$.

We can now begin to prove some new results about CAT(0) cubical complexes. The following proposition generalises Theorem 2.6.3, which was the case of $k = 0$.

Proposition 3.6.7. *If $\square$ is a d-dimensional CAT(0) cubical complex and $\Delta \coloneqq \mathrm{Cross}(\square)$ is its crossing complex, then $|\mathrm{Skel}^c_k \square| \simeq |\mathrm{Skel}^c_{k-1} \Delta|$ for all $k = 0,\dots,d$.*

Proof. If $k = d$, then both $|\mathrm{Skel}^c_d \square|$ and $|\mathrm{Skel}^c_{d-1} \Delta|$ are empty, since $\dim \Delta = d-1$ by Lemma 3.6.4. Therefore, for the rest of this proof, assume that $k < d$.

We will apply the nerve theorem (Theorem 3.2.3) to each of $|\mathrm{Skel}^c_k \square|$ and $|\mathrm{Skel}^c_{k-1} \Delta|$, to show that they are both homotopy equivalent to the same nerve.

First, consider $|\mathrm{Skel}^c_k \square|$, and define the family $\mathcal{H} \coloneqq \{H_\sigma : \sigma \in \Delta_k\}$, where $H_\sigma \coloneqq \bigcap_{i \in \sigma} H_i$ as above. By Proposition 3.6.1 and the above discussion, $|\mathrm{Skel}^c_k \square|$ is homotopy equivalent to the union of $\mathcal{H}$. Each H_σ is a closed set in $\square$, and by Proposition 3.6.1 their union is a triangulable space. Any intersection of spaces in $\mathcal{H}$ is itself an iterated hyperplane if it is non-empty, which is CAT(0) by Lemma 3.6.5 and thus contractible by Lemma 3.6.2. Therefore, we may apply the nerve theorem to $\mathcal{H}$, and conclude that $|\mathrm{Skel}^c_k \square| \simeq |N(\mathcal{H})|$.

But what is $N(\mathcal{H})$? The vertices are in bijection with k-faces σ of Δ. A set of vertices, say $\{\sigma_1, \dots, \sigma_m\}$, forms a face of $N(\mathcal{H})$ whenever the intersection $H_{\sigma_1} \cap \cdots \cap H_{\sigma_m}$ is non-empty. This intersection is the set

$$\begin{aligned} H_{\sigma_1} \cap \cdots \cap H_{\sigma_m} &= \bigcap_{i \in \sigma_1} H_i \cap \cdots \cap \bigcap_{i \in \sigma_m} H_i \\ &= \bigcap_{i \in \sigma_1 \cup \cdots \cup \sigma_m} H_i \\ &= H_{\sigma_1 \cup \cdots \cup \sigma_m}. \end{aligned}$$

So $\{\sigma_1, \dots, \sigma_m\}$ is a face of $N(\mathcal{H})$ if and only if $\sigma_1 \cup \cdots \cup \sigma_m$ is a face of Δ.

Now, let us turn our attention to $\operatorname{Skel}^c_{k-1} \Delta$, the set of faces of Δ of dimension at least k. For each k-face σ of Δ, let S_σ denote $|\operatorname{star}_\Delta \sigma|$, and define $\mathcal{S} \coloneqq \{S_\sigma : \sigma \in \Delta_k\}$. Every face of Δ of dimension at least k is contained in $\operatorname{star}_\Delta \sigma$ for some k-face σ, and no face of dimension less than k is, so the union of $\mathcal{S}$ is $|\operatorname{Skel}^c_{k-1} \Delta|$. Each S_σ is open in Δ, and any non-empty intersection of stars is itself a star and thus contractible. Therefore, we may apply the nerve theorem to $\mathcal{S}$ as well.

Let us examine $N(\mathcal{S})$. The vertices are in bijection with k-faces σ of Δ, and a set $\{\sigma_1, \dots, \sigma_m\}$ forms a face of $N(\mathcal{S})$ whenever $\operatorname{star}_\Delta \sigma_1 \cap \cdots \cap \operatorname{star}_\Delta \sigma_m$ is non-empty. This happens if and only if there exists a face of Δ that contains all of $\sigma_1, \dots, \sigma_m$, which happens whenever $\sigma_1 \cup \cdots \cup \sigma_m$ is a face of Δ.

But this means that $N(\mathcal{S}) = N(\mathcal{H})$! Therefore,

$$|\operatorname{Skel}^c_k \square| \simeq |N(\mathcal{H})| = |N(\mathcal{S})| \simeq |\operatorname{Skel}^c_{k-1} \Delta|. \qquad \blacksquare$$

Next, we will examine some applications of this result, but first we need one more construction.

There is a standard bijection from the non-empty faces σ of the cube $[0,1]^r$ to vectors χ^σ in $\{0, 1, *\}^r$, where the ith coordinate of χ^σ is 0 or 1 if all points x in σ have $x_i = 0$ or $x_i = 1$ respectively, and χ^σ_i is $*$ otherwise. If σ and τ are two faces, then $\sigma \subseteq \tau$ if and only if $\chi_i^\sigma \preceq \chi_i^\tau$ for all i, where $0 \prec *$ and $1 \prec *$ are the only non-equality relations on 0, 1 and $*$.

Suppose Δ is an abstract simplicial complex with vertex set $\{1, \dots, n\}$. We define the *cubical cone* over Δ, denoted $\operatorname{cone}(\Delta)$, to be the subcomplex of $[0,1]^n$ consisting of all faces σ such that the set $\left\{i : \chi_i^\sigma \in \{1, *\}\right\}$ is a face of Δ (including the empty face). For example, see Fig. 3.14. The reason for the name "cubical cone" is the following lemma, analogous to properties of a simplicial cone.

Lemma 3.6.8. *The cubical cone over Δ is the unique subcomplex of a cube (up to isomorphism) with the property that there is a vertex v_0 (the "cone point", which is the point $(0, \dots, 0)$ in the standard labelling) such that:*

- *every facet contains v_0, and*
- *the link of v_0 is (isomorphic as a poset to the non-empty faces of) Δ.*

Proof. First, let us check that the cubical cone does satisfy these properties. If σ is a face of $\operatorname{cone}(\Delta)$, we can replace every 1 in χ^σ with $*$ to obtain the vector of a face of $\operatorname{cone}(\Delta)$ that contains σ; therefore, the vectors of facets of $\operatorname{cone}(\Delta)$ consist only of 0 and $*$, so every facet contains $v_0 = (0, \dots, 0)$. Conversely, the faces that properly contain v_0 are precisely the faces $\sigma \neq v_0$ such that χ^σ consists of 0 and $*$, and these faces are in bijection with non-empty faces of Δ, so the link of v_0 is indeed Δ.

Now, suppose $\square$ is an arbitrary subcomplex of a cube satisfying these properties. By the symmetry of the cube, we may assume that v_0 is the vertex $(0, \dots, 0)$. As before, the link of v_0 is then the set of faces of $\square$ whose vectors consist only of 0 and $*$; since we assume this set is isomorphic to Δ, we may again use the symmetry of the cube to permute the coordinates and assume that this set is equal to $\operatorname{link}_{\operatorname{cone}(\Delta)} v_0$. Since $\square$ contains these faces of the cube, it must also contain every sub-face of these faces, so $\square$ contains every face of $\operatorname{cone}(\Delta)$. Moreover, it cannot contain any other faces: any such face must be contained in some facet of $\square$, but that facet cannot be one that we have already accounted for in $\operatorname{link}_\square v_0$, so the facet would not contain v_0. Therefore, $\square \cong \operatorname{cone}(\Delta)$. ∎

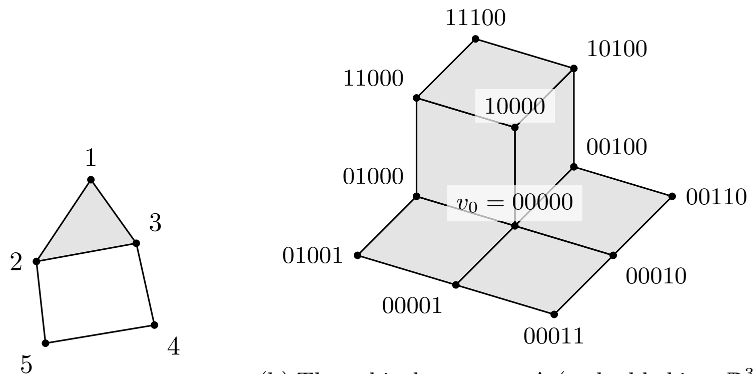

(a) Δ

(b) The cubical cone over Δ (embedded into $\mathbb{R}^3$), with vertices labelled

Figure 3.14: A cubical cone

While every flag simplicial complex is the crossing complex of some CAT(0) cubical complex (see [Hag14, Proposition 2.19] and Lemma 2.3.3, where the construction is the cubical cone), in general, there can be many CAT(0) cubical complexes with the same crossing complex. However, we can now prove that if Cross($\square$) is a connected manifold, the only possibility is $\square = \text{cone}(\Delta)$.

Corollary 3.6.9. *Let $\square$ be a d-dimensional CAT(0) cubical complex and Δ its crossing complex. If Δ is a connected homology manifold, then $\square$ is the cubical cone over Δ.*

Proof. Examine this section of the long exact sequence from Theorem 3.4.1, taking $k = 0$ and using coefficients in $\mathbb{Z}/2$:

$$\widetilde{H}_d(|\text{Skel}^c_{-1}\square|) \longrightarrow \bigoplus_{v\in\square_0} \widetilde{H}_{d-1}(|\text{link}\, v|) \longrightarrow \widetilde{H}_{d-1}(|\text{Skel}^c_0\square|) \longrightarrow \widetilde{H}_{d-1}(|\text{Skel}^c_{-1}\square|).$$

Note that $|\text{Skel}^c_{-1}\square|$ is simply $|\square|$, which is contractible since $\square$ is CAT(0), so the first and last terms of this sequence are 0. By Proposition 3.6.7, $\widetilde{H}_{d-1}(|\text{Skel}^c_0\square|)$ is isomorphic to $\widetilde{H}_{d-1}(|\text{Skel}^c_{-1}\Delta|) = \widetilde{H}_{d-1}(|\Delta|)$, and since Δ is a connected homology manifold of dimension $d-1$, this homology group is $\mathbb{Z}/2$. Therefore,

$$\bigoplus_{v\in\square_0} \widetilde{H}_{d-1}(|\operatorname{link} v|) \cong \mathbb{Z}/2.$$

Thus there must be one vertex v_0 for which $\widetilde{H}_{d-1}(|\operatorname{link} v_0|) = \mathbb{Z}/2$, and all other vertices have $\widetilde{H}_{d-1}(|\operatorname{link} v|) = 0$.

Every vertex link in a CAT(0) cubical complex is an induced subcomplex of the crossing complex (Lemma 2.4.9). However, the only induced subcomplex of a connected $(d-1)$-manifold with non-zero homology in degree $d-1$ is the entire manifold. Therefore, $\operatorname{link} v_0 = \Delta$.

Lemma 3.6.4 says that there is a bijection between the facets of the cubical complex $\square$ and the facets of its crossing complex Δ. Since $\operatorname{link} v_0 = \Delta$, we also have a bijection between the facets of Δ and the facets of $\square$ that contain v_0. Therefore, every facet of $\square$ must contain v_0. Hence $\square$ is a cubical cone. ■

We conclude this chapter by combining Proposition 3.6.7 with Theorems 3.5.3, 3.5.11 and 3.5.10, to compare the combinatorial properties discussed in Section 3.5 for $\square$ and $\mathrm{Cross}(\square)$.

Theorem 3.6.10. *Suppose $\square$ is a CAT(0) cubical complex and Δ is its crossing complex. Then:*

(a) $\square$ is Cohen–Macaulay if and only if Δ is Cohen–Macaulay.

(b) $\square$ is r-Leray if and only if Δ is r-Leray.

(c) If $\square$ and Δ are both homology balls, then $\square$ is s-stacked if and only if Δ is s-stacked.

Note that the assumption that both $\square$ and Δ are balls in (c) is essential: it is possible for one to be a ball and the other not. For example, if $\square$ consists of m squares glued in a $1 \times m$ rectangle, it is a 1-stacked ball, but its crossing complex is a star graph with m pendants, which is not a ball for $m > 2$.

Proof of Theorem 3.6.10. (a) Let $d = \dim \square$. By Theorem 3.5.3, $\square$ is Cohen–Macaulay if and only if it is pure and $\widetilde{H}_i(|\mathrm{Skel}^c_k \square|) = 0$ for all $k = -1, \ldots, d$ and all $i < d-k-1$. Any CAT(0) cubical complex is always contractible, so $\widetilde{H}_i(|\mathrm{Skel}^c_{-1} \square|)$ is guaranteed to be 0, hence we only need to consider $k = 0, \ldots, d$. Now, Proposition 3.6.7 says that $\widetilde{H}_i(|\mathrm{Skel}^c_k \square|) = \widetilde{H}_i(|\mathrm{Skel}^c_{k-1} \Delta|)$ for k in this range, and Lemma 3.6.4 says that $\square$ is pure if and only if Δ is pure; therefore, $\square$ is Cohen–Macaulay if and only if Δ is pure and $\widetilde{H}_i(|\mathrm{Skel}^c_{k-1} \Delta|) = 0$ for all $k = 0, \ldots, d$ and all $i < d - k - 1 = (d-1) - (k-1) - 1$. But since $\dim \Delta = d - 1$, this is exactly the condition given by Theorem 3.5.3 for Δ to be Cohen–Macaulay.

(b) By Theorem 3.5.10, $\square$ is r-Leray if and only if $\widetilde{H}_i(|\mathrm{Skel}^c_k \square|) = 0$ for all $i \geq r$ and all $k = -1, \ldots, d$. Because $\square$ is CAT(0), $|\mathrm{Skel}^c_{-1} \square|$ is guaranteed to be contractible, so we only need to check $k = 0, \ldots, d$. By Proposition 3.6.7, this condition is equivalent to $\widetilde{H}_i(|\mathrm{Skel}^c_j \Delta|)$ being 0 for all $i \geq r$ and all $j = -1, \ldots, d-1$, where $j = k - 1$. And this precisely means that Δ is r-Leray.

(c) By Theorem 3.5.11, $\square$ is s-stacked if and only if $\widetilde{H}_{d-k-1}(|\mathrm{Skel}^c_k \square|) = 0$ for all $k \leq d - s - 1$, where $d = \dim \square$. Proposition 3.6.7 tells us this is equivalent to $\widetilde{H}_{d-k-1}(|\mathrm{Skel}^c_{k-1} \Delta|)$ being 0 for $k \leq d - s - 1$. Equivalently, letting $j = k - 1$, we have $\widetilde{H}_{(d-1)-j-1}(|\mathrm{Skel}^c_j \Delta|) = 0$ for $j \leq (d-1) - s - 1$. Since $\dim \Delta = d - 1$, this is equivalent to Δ being s-stacked. ∎

Bibliography

[AG04] A. Abrams and R. Ghrist. “State Complexes for Metamorphic Robots”. In: *The International Journal of Robotics Research* 23.7–8 (2004), pp. 811–826. DOI: 10.1177/0278364904045468. eprint: https://doi.org/10.1177/0278364904045468. URL: https://doi.org/10.1177/0278364904045468.

[AH16] Michał Adamaszek and Jan Hladký. “Upper bound theorem for odd-dimensional flag triangulations of manifolds”. In: *Mathematika* 62.3 (Jan. 2016), pp. 909–928. ISSN: 2041-7942. DOI: 10.1112/s0025579316000115. URL: http://dx.doi.org/10.1112/S0025579316000115.

[Adi96] Ron M. Adin. “A new cubical h-vector”. In: *Proceedings of the 6th Conference on Formal Power Series and Algebraic Combinatorics (New Brunswick, NJ, 1994)*. Vol. 157. 1-3. 1996, pp. 3–14. DOI: 10.1016/S0012-365X(96)83003-2. URL: https://doi.org/10.1016/S0012-365X(96)83003-2.

[AB14] Karim Adiprasito and Bruno Benedetti. “The Hirsch conjecture holds for normal flag complexes”. In: *Math. Oper. Res.* 39.4 (2014), pp. 1340–1348. ISSN: 0364-765X. DOI: 10.1287/moor.2014.0661. URL: https://doi.org/10.1287/moor.2014.0661.

[AB20] Karim Adiprasito and Bruno Benedetti. “Collapsibility of CAT(0) spaces”. In: *Geom. Dedicata* 206 (2020), pp. 181–199. ISSN: 0046-5755. DOI: 10.1007/s10711-019-00481-x. URL: https://doi.org/10.1007/s10711-019-00481-x.

[AL20] Karim Adiprasito and Gaku Liu. *Normal crossing immersions, cobordisms and flips*. 2020. arXiv: 2001.01108 [math.GT].

[AY21] Karim Adiprasito and Geva Yashfe. "The partition complex: an invitation to combinatorial commutative algebra". In: *Surveys in combinatorics 2021*. Vol. 470. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge, 2021, pp. 1–41.

[Ago13] Ian Agol. "The virtual Haken conjecture". In: *Doc. Math.* 18 (2013). With an appendix by Agol, Daniel Groves, and Jason Manning, pp. 1045–1087. ISSN: 1431-0635.

[AK85] N. Alon and G. Kalai. "A simple proof of the upper bound theorem". In: *European J. Combin.* 6.3 (1985), pp. 211–214. ISSN: 0195-6698. DOI: 10.1016/S0195-6698(85)80029-9. URL: https://doi.org/10.1016/S0195-6698(85)80029-9.

[Ard19] Federico Ardila. "CAT(0) geometry, robots, and society". In: *arXiv e-prints*, arXiv:1912.10007 (Dec. 2019), arXiv:1912.10007. arXiv: 1912.10007 [math.HO].

[ABY14] Federico Ardila, Tia Baker, and Rika Yatchak. "Moving robots efficiently using the combinatorics of CAT(0) cubical complexes". In: *SIAM J. Discrete Math.* 28.2 (2014), pp. 986–1007. ISSN: 0895-4801. DOI: 10.1137/120898115. URL: https://doi.org/10.1137/120898115.

[AOS12] Federico Ardila, Megan Owen, and Seth Sullivant. "Geodesics in CAT(0) cubical complexes". In: *Adv. in Appl. Math.* 48.1 (2012), pp. 142–163. ISSN: 0196-8858. DOI: 10.1016/j.aam.2011.06.004. URL: https://doi.org/10.1016/j.aam.2011.06.004.

[Ath12] Christos A. Athanasiadis. "On the Möbius function of a lower Eulerian Cohen-Macaulay poset". In: *J. Algebraic Combin.* 35.3 (2012), pp. 373–388. ISSN: 0925-9899. DOI: 10.1007/s10801-011-0306-4. URL: https://doi.org/10.1007/s10801-011-0306-4.

[BBC97] Eric K. Babson, Louis J. Billera, and Clara S. Chan. "Neighborly cubical spheres and a cubical lower bound conjecture". In: *Israel J. Math.* 102 (1997), pp. 297–315. ISSN: 0021-2172. DOI: 10.1007/BF02773804. URL: https://doi.org/10.1007/BF02773804.

[BC00] Eric K. Babson and Clara S. Chan. "Counting faces of cubical spheres modulo two". In: vol. 212. 3. Combinatorics and applications (Tianjin, 1996). 2000, pp. 169–183. DOI: 10.1016/S0012-365X(99)00285-X. URL: https://doi.org/10.1016/S0012-365X(99)00285-X.

[BC08] Hans-Jürgen Bandelt and Victor Chepoi. "Metric graph theory and geometry: a survey". In: *Surveys on discrete and computational geometry.* Vol. 453. Contemp. Math. Amer. Math. Soc., Providence, RI, 2008, pp. 49–86. DOI: 10.1090/conm/453/08795. URL: https://doi.org/10.1090/conm/453/08795.

[BC93] Jean-Pierre Barthélémy and Julien Constantin. "Median graphs, parallelism and posets". In: vol. 111. 1-3. Graph theory and combinatorics (Marseille-Luminy, 1990). 1993, pp. 49–63. DOI: 10.1016/0012-365X(93)90140-O. URL: https://doi.org/10.1016/0012-365X(93)90140-O.

[Bay18] M. M. Bayer. "Graphs, skeleta and reconstruction of polytopes". In: *Acta Math. Hungar.* 155.1 (2018), pp. 61–73. ISSN: 0236-5294. DOI: 10.1007/s10474-018-0804-0. URL: https://doi.org/10.1007/s10474-018-0804-0.

[BHV01] Louis J. Billera, Susan P. Holmes, and Karen Vogtmann. "Geometry of the space of phylogenetic trees". In: *Adv. in Appl. Math.* 27.4 (2001), pp. 733–767. ISSN: 0196-8858. DOI: 10.1006/aama.2001.0759. URL: https://doi.org/10.1006/aama.2001.0759.

[Bjö95] A. Björner. "Topological methods". In: *Handbook of combinatorics, Vol. 1, 2.* Elsevier Sci. B. V., Amsterdam, 1995, pp. 1819–1872.

[BK88] Anders Björner and Gil Kalai. “An extended Euler-Poincaré theorem”. In: *Acta Math.* 161.3-4 (1988), pp. 279–303. ISSN: 0001-5962. DOI: 10.1007/BF02392300. URL: https://doi.org/10.1007/BF02392300.

[BH72] Jósef Blass and Włodzimierz Holsztyński. “Cubical polyhedra and homotopy. III”. In: *Atti Accad. Naz. Lincei Rend. Cl. Sci. Fis. Mat. Natur. (8)* 53 (1972), pp. 275–279.

[BB98] G. Blind and R. Blind. “The almost simple cubical polytopes”. In: *Discrete Math.* 184.1-3 (1998), pp. 25–48. ISSN: 0012-365X. DOI: 10.1016/S0012-365X(97)00159-3. URL: https://doi.org/10.1016/S0012-365X(97)00159-3.

[Bow14] Brian H. Bowditch. “Embedding median algebras in products of trees”. In: *Geom. Dedicata* 170 (2014), pp. 157–176. ISSN: 0046-5755. DOI: 10.1007/s10711-013-9874-x. URL: https://doi.org/10.1007/s10711-013-9874-x.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Vol. 319. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xxii+643. ISBN: 3-540-64324-9. DOI: 10.1007/978-3-662-12494-9. URL: https://doi.org/10.1007/978-3-662-12494-9.

[CS11] Pierre-Emmanuel Caprace and Michah Sageev. “Rank rigidity for CAT(0) cube complexes”. In: *Geom. Funct. Anal.* 21.4 (2011), pp. 851–891. ISSN: 1016-443X. DOI: 10.1007/s00039-011-0126-7. URL: https://doi.org/10.1007/s00039-011-0126-7.

[CD95] Ruth Charney and Michael Davis. “The Euler characteristic of a nonpositively curved, piecewise Euclidean manifold.” In: *Pacific J. Math.* 171.1 (1995), pp. 117–137. URL: https://projecteuclid.org:443/euclid.pjm/1102370321.

[CN05] Indira Chatterji and Graham Niblo. "From wall spaces to CAT(0) cube complexes". In: *Internat. J. Algebra Comput.* 15.5-6 (2005), pp. 875–885. ISSN: 0218-1967. DOI: `10.1142/S0218196705002669`. URL: `https://doi.org/10.1142/S0218196705002669`.

[Che00] Victor Chepoi. "Graphs of some CAT(0) complexes". In: *Adv. in Appl. Math.* 24.2 (2000), pp. 125–179. ISSN: 0196-8858. DOI: `10.1006/aama.1999.0677`. URL: `https://doi.org/10.1006/aama.1999.0677`.

[CH13] Victor Chepoi and Mark F. Hagen. "On embeddings of CAT(0) cube complexes into products of trees via colouring their hyperplanes". In: *J. Combin. Theory Ser. B* 103.4 (2013), pp. 428–467. ISSN: 0095-8956. DOI: `10.1016/j.jctb.2013.04.003`. URL: `https://doi.org/10.1016/j.jctb.2013.04.003`.

[Dan84] Jerome Dancis. "Triangulated n-manifolds are determined by their $[n/2]+1$-skeletons". In: *Topology Appl.* 18.1 (1984), pp. 17–26. ISSN: 0166-8641. DOI: `10.1016/0166-8641(84)90028-2`. URL: `https://doi.org/10.1016/0166-8641(84)90028-2`.

[DO01] Michael W Davis and Boris Okun. "Vanishing theorems and conjectures for the ℓ^2-homology of right-angled Coxeter groups". In: *Geom. Topol.* 5.1 (2001), pp. 7–74. DOI: `10.2140/gt.2001.5.7`. URL: `https://doi.org/10.2140/gt.2001.5.7`.

[Dil50] R. P. Dilworth. "A decomposition theorem for partially ordered sets". In: *Ann. of Math. (2)* 51 (1950), pp. 161–166. ISSN: 0003-486X. DOI: `10.2307/1969503`. URL: `https://doi.org/10.2307/1969503`.

[DSS94] N. P. Dolbilin, M. A. Shtan'ko, and M. I. Shtogrin. "Cubic manifolds in lattices". In: *Izv. Ross. Akad. Nauk Ser. Mat.* 58.2 (1994), pp. 93–107. ISSN: 1607-0046. DOI: `10.1070/IM1995v044n02ABEH001598`. URL: `https://doi.org/10.1070/IM1995v044n02ABEH001598`.

[DKM11] Art M. Duval, Caroline J. Klivans, and Jeremy L. Martin. "Cellular spanning trees and Laplacians of cubical complexes". In: *Adv. in Appl. Math.* 46.1-4 (2011), pp. 247–274. ISSN: 0196-8858. DOI: 10.1016/j.aam.2010.05.005. URL: https://doi.org/10.1016/j.aam.2010.05.005.

[EH95] Richard Ehrenborg and Gábor Hetyei. "Generalizations of Baxter's theorem and cubical homology". In: *J. Combin. Theory Ser. A* 69.2 (1995), pp. 233–287. ISSN: 0097-3165. DOI: 10.1016/0097-3165(95)90053-5. URL: https://doi.org/10.1016/0097-3165(95)90053-5.

[Fio22] Elia Fioravanti. "Automorphisms of contact graphs of CAT(0) cube complexes". In: *Int. Math. Res. Not. IMRN* 5 (2022), pp. 3278–3296. ISSN: 1073-7928. DOI: 10.1093/imrn/rnaa280. URL: https://doi.org/10.1093/imrn/rnaa280.

[Fro08] Andrew Frohmader. "Face vectors of flag complexes". In: *Israel J. Math.* 164 (2008), pp. 153–164. ISSN: 0021-2172. DOI: 10.1007/s11856-008-0024-3. URL: https://doi.org/10.1007/s11856-008-0024-3.

[Fun99] Louis Funar. "Cubulations, immersions, mappability and a problem of Habegger". In: *Ann. Sci. École Norm. Sup. (4)* 32.5 (1999), pp. 681–700. ISSN: 0012-9593. DOI: 10.1016/S0012-9593(01)80003-3. URL: https://doi.org/10.1016/S0012-9593(01)80003-3.

[Gal05] Swiatoslaw R. Gal. "Real root conjecture fails for five- and higher-dimensional spheres". In: *Discrete & Computational Geometry* 34.2 (May 2005), pp. 269–284. ISSN: 1432-0444. DOI: 10.1007/s00454-005-1171-5. URL: http://dx.doi.org/10.1007/s00454-005-1171-5.

[Goo23] Google. *Google Books Ngram Viewer*. Last accessed 10 April 2023. 2023. URL: https://books.google.com/ngrams/graph?content=simplicial+complex%2Ccubical+complex&year_start=1900.

[Gro87] M. Gromov. “Hyperbolic groups”. In: *Essays in group theory.* Vol. 8. Math. Sci. Res. Inst. Publ. Springer, New York, 1987, pp. 75–263. DOI: `10.1007/978-1-4613-9586-7_3`. URL: `https://doi.org/10.1007/978-1-4613-9586-7_3`.

[Hag12] Mark F. Hagen. *Geometry and combinatorics of cube complexes.* Thesis (Ph.D.)—McGill University (Canada). ProQuest LLC, Ann Arbor, MI, 2012, p. 192. ISBN: 978-0494-79046-5. URL: `https://www.proquest.com/dissertations-theses/geometry-combinatorics-cube-complexes/docview/1267897075/se-2`.

[Hag13] Mark F. Hagen. “The simplicial boundary of a CAT(0) cube complex”. In: *Algebr. Geom. Topol.* 13.3 (2013), pp. 1299–1367. ISSN: 1472-2747. DOI: `10.2140/agt.2013.13.1299`. URL: `https://doi.org/10.2140/agt.2013.13.1299`.

[Hag14] Mark F. Hagen. “Weak hyperbolicity of cube complexes and quasi-arboreal groups”. In: *J. Topol.* 7.2 (2014), pp. 385–418. ISSN: 1753-8416. DOI: `10.1112/jtopol/jtt027`. URL: `https://doi.org/10.1112/jtopol/jtt027`.

[HW08] Frédéric Haglund and Daniel T. Wise. “Special cube complexes”. In: *Geom. Funct. Anal.* 17.5 (2008), pp. 1551–1620. ISSN: 1016-443X. DOI: `10.1007/s00039-007-0629-4`. URL: `https://doi.org/10.1007/s00039-007-0629-4`.

[Hat02] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002, pp. xii+544. ISBN: 0-521-79160-X; 0-521-79540-0.

[HM72] Ivan Havel and Jaroslav Morávek. “B-valuations of graphs”. In: *Czechoslovak Math. J.* 22(97) (1972), pp. 338–351. ISSN: 0011-4642.

[Het95] G. Hetyei. “On the Stanley ring of a cubical complex”. In: *Discrete Comput. Geom.* 14.3 (1995), pp. 305–330. ISSN: 0179-5376. DOI: `10.1007/BF02570709`. URL: `https://doi.org/10.1007/BF02570709`.

[Het96] Gábor Hetyei. "Invariants des complexes cubiques". In: *Ann. Sci. Math. Québec* 20.1 (1996), pp. 35–52. ISSN: 0707-9109.

[JZ00] M. Joswig and G. M. Ziegler. "Neighborly cubical polytopes". In: *Discrete Comput. Geom.* 24.2-3 (2000). The Branko Grünbaum birthday issue, pp. 325–344. ISSN: 0179-5376. DOI: 10.1007/s004540010039. URL: https://doi.org/10.1007/s004540010039.

[JR07] Michael Joswig and Thilo Rörig. "Neighborly cubical polytopes and spheres". In: *Israel J. Math.* 159 (2007), pp. 221–242. ISSN: 0021-2172. DOI: 10.1007/s11856-007-0044-4. URL: https://doi.org/10.1007/s11856-007-0044-4.

[KMM04] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational homology*. Vol. 157. Applied Mathematical Sciences. Springer-Verlag, New York, 2004, pp. xviii+480. ISBN: 0-387-40853-3. DOI: 10.1007/b97315. URL: https://doi.org/10.1007/b97315.

[Kal87] Gil Kalai. "Rigidity and the lower bound theorem. I". In: *Invent. Math.* 88.1 (1987), pp. 125–151. ISSN: 0020-9910. DOI: 10.1007/BF01405094. URL: https://doi.org/10.1007/BF01405094.

[Kal02] Gil Kalai. "Algebraic shifting". In: *Computational commutative algebra and combinatorics (Osaka, 1999)*. Vol. 33. Adv. Stud. Pure Math. Math. Soc. Japan, Tokyo, 2002, pp. 121–163. DOI: 10.2969/aspm/03310121. URL: https://doi.org/10.2969/aspm/03310121.

[KM06] Gil Kalai and Roy Meshulam. "Intersections of Leray complexes and regularity of monomial ideals". In: *J. Combin. Theory Ser. A* 113.7 (2006), pp. 1586–1592. ISSN: 0097-3165. DOI: 10.1016/j.jcta.2006.01.005. URL: https://doi.org/10.1016/j.jcta.2006.01.005.

[Kle11] Steven Klee. “Lower bounds for cubical pseudomanifolds”. In: *Discrete Comput. Geom.* 46.2 (2011), pp. 212–222. ISSN: 0179-5376. DOI: 10.1007/s00454-011-9329-9. URL: https://doi.org/10.1007/s00454-011-9329-9.

[Mun84a] James R. Munkres. *Elements of algebraic topology.* Addison-Wesley Publishing Company, Menlo Park, CA, 1984, pp. ix+454. ISBN: 0-201-04586-9.

[Mun84b] James R. Munkres. “Topological results in combinatorics”. In: *Michigan Math. J.* 31.1 (1984), pp. 113–128. ISSN: 0026-2285. DOI: 10.1307/mmj/1029002969. URL: https://doi.org/10.1307/mmj/1029002969.

[MN13] Satoshi Murai and Eran Nevo. “On the generalized lower bound conjecture for polytopes and spheres”. In: *Acta Math.* 210.1 (2013), pp. 185–202. ISSN: 0001-5962. DOI: 10.1007/s11511-013-0093-y. URL: https://doi.org/10.1007/s11511-013-0093-y.

[NR98] G. A. Niblo and L. D. Reeves. “The geometry of cube complexes and the complexity of their fundamental groups”. In: *Topology* 37.3 (1998), pp. 621–633. ISSN: 0040-9383. DOI: 10.1016/S0040-9383(97)00018-9. URL: https://doi.org/10.1016/S0040-9383(97)00018-9.

[Rol16] Martin Roller. *Poc Sets, Median Algebras and Group Actions.* https://arxiv.org/abs/1607.07747. 2016. arXiv: 1607.07747 [math.GN].

[Row22a] Rowan Rowlands. “Reconstructing d-manifold subcomplexes of cubes from their $(\lfloor d/2 \rfloor + 1)$-skeletons”. In: *Discrete Comput. Geom.* 67.2 (2022), pp. 492–502. ISSN: 0179-5376. DOI: 10.1007/s00454-021-00321-4. URL: https://doi.org/10.1007/s00454-021-00321-4.

[Row22b] Rowan Rowlands. *Topology of complements of skeletons.* https://arxiv.org/abs/2209.08434. 2022. DOI: 10.48550/ARXIV.2209.08434. URL: https://arxiv.org/abs/2209.08434.

[Row23] Rowan Rowlands. *Relating CAT(0) cubical complexes and flag simplicial complexes.* 2023. DOI: 10.1016/j.aam.2023.102502. URL: https://doi.org/10.1016/j.aam.2023.102502.

[Sag95] Michah Sageev. "Ends of group pairs and non-positively curved cube complexes". In: *Proc. London Math. Soc. (3)* 71.3 (1995), pp. 585–617. ISSN: 0024-6115. DOI: 10.1112/plms/s3-71.3.585. URL: https://doi.org/10.1112/plms/s3-71.3.585.

[SZ04] Alexander Schwartz and Günter M. Ziegler. "Construction techniques for cubical complexes, odd cubical 4-polytopes, and prescribed dual manifolds". In: *Experiment. Math.* 13.4 (2004), pp. 385–413. ISSN: 1058-6458. URL: http://projecteuclid.org/euclid.em/1109106431.

[Wik23] Wikipedia contributors. *Cubical complex — Wikipedia, The Free Encyclopedia.* [Online; accessed 9-April-2023]. 2023. URL: https://en.wikipedia.org/w/index.php?title=Cubical_complex&oldid=1059213772.

[Wis21] Daniel T. Wise. *The structure of groups with a quasiconvex hierarchy.* Vol. 209. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 2021, pp. x+357. ISBN: [9780691170442]; [9780691170459]; [9780691213507].

[Zie95] Günter M. Ziegler. *Lectures on polytopes.* Vol. 152. Graduate Texts in Mathematics. Springer-Verlag, New York, 1995, pp. x+370. ISBN: 0-387-94365-X. DOI: 10.1007/978-1-4613-8431-1. URL: https://doi.org/10.1007/978-1-4613-8431-1.