Bisimplicial Complexes and Shortcut Graphs

Nima Hoda

Doctor of Philosopy

Department of Mathematics and Statistics

McGill University

Montreal, Quebec

April 14, 2019

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy.

Dedication

To Lena

Acknowledgements

I would like to thank my advisor, Daniel T. Wise, for his encouragement and for our many invaluable discussions over the course of this research. I thank Piotr Przytycki for carefully reading this thesis and providing detailed notes leading to many corrections. Thank you also to Brahim Abdenbi for correcting the french abstract. More generally, I would like to thank all the members of the McGill Geometric Group Theory Research Group for fostering an excellent environment for a young researcher to develop and thrive. I would like to thank the administrators of the McGill Department of Mathematics and Statistics for their kindness and efficacy. Finally, I would like to thank my dear family: mom, dad, Lena, Shiva, Reza, Arteen and Misha. Thank you for the many ways you have supported me and brought joy to my life throughout the years. This research was supported in part by an NSERC Postgraduate Scholarship, a Pelletier Fellowship, a Graduate Mobility Award, an ISM Scholarship and the NSERC grant of Daniel T. Wise.

Abstract

This thesis concerns groups acting on spaces of combinatorial nonpositive curvature. The first part of this thesis presents a discrete Morse-theoretic method for proving that a regular CW complex is homeomorphic to a sphere. We use this method to define *bisimplices*, the cells of a class of regular CW complexes we call *bisimplicial complexes*. The 1-skeleta of bisimplices are complete bipartite graphs making them suitable in constructing higher dimensional skeleta for bipartite graphs. We show that the flag bisimplicial completion of a finite bipartite bi-dismantlable graph is collapsible. We use this to show that the flag bisimplicial completion of a quadric complex is contractible and to construct a compact $K(G,1)$ for a torsion-free quadric group G.

The second part of this thesis introduces shortcut graphs and groups. *Shortcut graphs* are graphs in which long enough cycles cannot embed without metric distortion. *Shortcut groups* are groups which act properly and cocompactly on shortcut graphs. These notions unify a surprisingly broad family of graphs and groups of interest in geometric group theory and metric graph theory including: systolic and quadric groups (in particular finitely presented C(6) and C(4)-T(4) small cancellation groups), cocompactly cubulated groups, hyperbolic groups, Coxeter groups and the Baumslag-Solitar group $\mathrm{BS}(1,2)$. Most of these examples satisfy a strong form of the shortcut property. We show that shortcut groups are finitely presented, have exponential isoperimetric and isodiametric function and are closed under direct products and under HNN extensions and amalgamated products over finite subgroups. We show that groups satisfying the strong form of the shortcut property satisfy these properties and also have polynomial isoperimetric and isodiametric function.

Abrégé

Cette thèse concerne les groupes qui agissent sur des espaces combinatoires à courbure négative ou nulle. La première partie de cette thèse présente une méthode basée sur la théorie de Morse discrète pour démontrer qu'un CW-complexe régulier est homéomorphe à une sphère. Nous utilisons cette méthode pour définir les *bisimplexes*, les cellules d'une classe de CW-complexes réguliers que nous appelons les *complexes bisimpliciaux*. Les 1-squelettes des bisimplexes sont des graphes bipartis complets et donc utiles dans la construction des squelettes de haute dimension des graphes bipartis. Nous montrons que le complété de drapeau bisimplicial d'un graphe biparti fini bi-démontable est collapsible. Nous utilisons ce résultat pour montrer que le complété de drapeau bisimplicial d'un complexe quadrique est contractile et pour construire un $K(G, 1)$ pour un groupe quadrique sans torsion G.

La deuxième partie de cette thèse introduit les graphes et les groupes shortcuts. Les *graphes shortcuts* sont les graphes dans lesquelles les cycles assez long ne peuvent pas plonger sans distortion métrique. Les *groupes shortcuts* sont les groupes qui agissent proprement et cocompactement sur les graphes shortcuts. Ces notions unifient une famille assez large et intéressante de graphes et de groupes notamment dans la théorie des groupes géométriques et la théorie des graphes métriques comprenant: les groupes systoliques et quadriques (en particulier les groupes à petites simplifications C(6) et C(4)-T(4) finement présenté), les groupes cocompactement cubulés, les groupes hyperboliques, les groupes Coxeter et le groupe Baumslag-Solitar $\mathrm{BS}(1, 2)$. La plupart de ces exemples satisfont une forme forte de la propriété shortcut. Nous montrons également que les groupes shortcuts sont finement présentés, ont des fonctions isopérimétriques et isodiamétriques exponentielles et qu'ils sont fermés par rapport aux produits directes, aux extensions HNN et aux produits libres avec amalgame sur des sous-groupes finis. Nous montrons que les groupes satisfaisant la forme forte de la propriété shortcut satisfont ces propriétés et ont aussi des fonctions isopérimétriques et isodiamétriques polynomiales.

Table of Contents

List of Figures

Figure | page

Chapter 1

Introduction

A major current in group theory throughout the 20th century has been the study of groups acting on spaces of nonpositive curvature. This line of work traces its origins to the work of Max Dehn on the fundamental groups of surfaces of higher genus [11]. Dehn proved that these groups have decidable word problem by exploiting their actions on the hyperbolic plane, the prototypical space of negative curvature. Dehn's work was generalized to the study of small cancellation groups [18]. These are various classes of groups defined by presentations in which overlap between relators is restricted. This line of work was revolutionized by Gromov's study of hyperbolic groups and program of studying groups via their quasi-isometry types, leading to the field of geometric group theory [19].

In addition to the solvability of the word problem, many other properties of a group G may be determined given an approparite action of G on a space with controlled geometry or topology. For example, if G acts freely and cocompactly on a contractible CW complex then the quotient is a finite $K(G,1)$ from which homological invariants of G may be effectively computed [24]. By the Milnor-Schwarz Lemma, if G acts properly and cocompactly on a metric space X then G and X have the same quasi-isometry type [5] and so share geometric invariants such as isoperimetric and isodiametric functions [16]. These are examples that we touch upon directly in this thesis but there are many more such interactions which drive the search for spaces and actions with nice properties:

invariant subspace properties relate to subgroup structure, dynamics of automorphisms help us classify group elements and understand their interactions, boundary theories open avenues towards topological dynamics, etc.

Three major classes of spaces by which these interactions have been studied are CAT(0) cube complexes [19, 45, 52, 53, 37, 23, 36, 44, 38, 26, 21], systolic complexes [28, 20, 51, 9, 12, 22, 29, 40, 39] and quadric complexes [25]. One part of this thesis is focused on quadric complexes. *Quadric complexes* are a class of combinatorially nonpositively curved complexes which may in many ways be considered analogous to systolic complexes: whereas disk diagrams in systolic complexes are CAT(0) equilateral triangle complexes, disk diagrams in quadric complexes are CAT(0) square complexes; whereas systolic groups generalize finitely presented C(6) small cancellation groups, quadric groups generalize finitely presented C(4)-T(4) small cancellation groups; whereas the 1-skeleta of systolic complexes are precisely the graphs whose isometric cycles all have length 3, the 1-skeleta of quadric complexes are precisely the graphs all of whose isometric cycles have length 4; etc. However, as originally defined, quadric complexes have a major deficiency: they are not contractible. Hence it is not possible to directly compute homological invariants of a group using an action on a quadric complex.

In Part I of this thesis we remedy this deficiency of quadric complexes by adding higher dimensional cells to quadric complexes in a natural way. Specifically, we introduce a new class of CW complexes called *bisimplicial complexes*. The cells of these complexes, called *bisimplices*, are bipartite analogs of simplices since their 1-skeleta are complete bipartite graphs whereas the 1-skeleta of simplices are complete graphs. Their construction is challenging since they are not Euclidean polyhedra, as is the case for many other classes of cells frequently used in CW complexes: most notably simplices and cubes. In the course of the construction of bisimplices we develop discrete Morse-theoretic tools for proving that a given CW complex is homeomorphic to a sphere. We show that the flag bisimplicial completion of a finite bipartite bi-dismantlable graph is collapsible. We apply this to construct natural contractible supercomplexes for quadric complexes and thereby construct natural finite dimensional $K(G, 1)$ for torsion-free quadric groups.

Part II of this thesis draws inspiration from the remarkable fact that the 1-skeleta of all three of the major classes of spaces referred to above arose independently in the metric graph theory literature: CAT(0) complexes arose as median graphs [2, 34, 32, 44, 15, 8], systolic complexes arose as bridged graphs [49, 8] and quadric complexes arose as hereditary modular graphs [3, 25]. We aim to define broad classes of graphs that capture general metric properties of these and several other classes of spaces featured prominently in the geometric group theory literature. Specifically, we define *shortcut* graphs and groups and *strongly shortcut* graphs and groups. These are essentially classes of graphs and groups in which arbitrarily long cycles cannot embed without metric distortion. Many classes of graphs and groups satisfying various forms of nonpositive curvature conditions are shortcut, including the 1-skeleta of systolic and quadric complexes (in particular finitely presented C(6) and C(4)-T(4) small cancellation groups), 1-skeleta of finite dimensional CAT(0) cube complexes, hyperbolic graphs, standard Cayley graphs of finitely generated Coxeter groups and the standard Cayley graph of the Baumslag-Solitar group $\mathrm{BS}(1, 2)$. We show that most of these examples are strongly shortcut. We also derive consequences of the shortcut properties. We show that shortcut groups are finitely presented, have exponential isoperimetric and isodiametric function and are closed under direct products and under HNN extensions and amalgamated products over finite subgroups. We show that strongly shortcut groups satisfy these properties and also have polynomial isoperimetric and isodiametric function.

1.1 Contributions to original knowledge

The main novel results of Part I are summarized in Section 2.1 by Theorem A, Theorem B, Theorem C, Theorem D and Theorem E. The main results of Part II are summarized in Section 8.1 by Theorem F, Theorem G, Theorem H, Theorem I, Theorem J and Theorem K.

Part I

Bisimplicial Complexes and Asphericity

Chapter 2

Introduction

CW complexes are typically constructed by gluing together Euclidean polyhedra along faces. A *Euclidean polyhedron* is the convex hull of a finite point set in a Euclidean space, e.g., simplices and cubes. However, not all CW structures on cells of a CW complex arise as Euclidean polyhedra [31] and for some applications it is natural to use nonpolyhedral cells. In this part we construct an infinite family of nonpolyhedral CW balls called *bisimplices*. The 1-skeleta of bisimplices are connected complete bipartite graphs and so we consider them as bipartite analogs of simplices. Our motivation for this construction is to find a natural contractible higher dimensional skeleton for quadric complexes.

Quadric complexes are locally finite simply connected square complexes satisfying a certain combinatorial nonpositive curvature condition. A group is *quadric* if it acts properly and cocompactly on a quadric complex. Quadric complexes are examples of the generalized $(4,4)$-complexes of Wise [51] and were first studied in detail in the context of geometric group theory by the present author [25]. They generalize the folder complexes of Chepoi [8] and may be considered as square analogs of the 2-skeleta of systolic complexes [28]. They can be characterized by their 1-skeleta, which are precisely the hereditary modular graphs of metric graph theory [3].

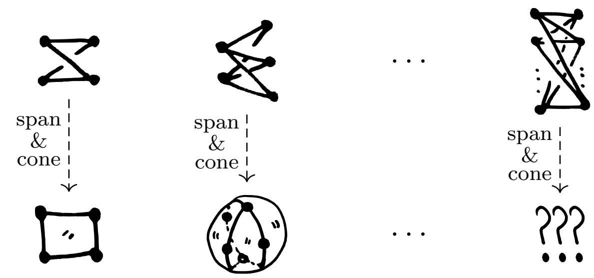

Figure 2–1: Bisimplices are essentially constructed by starting with a $K_{m,n}$, $m, n \geq 2$, inductively spanning a biclique on each proper $K_{m',n'}$ subgraph, $m', n' \geq 2$, and then taking the cone of the result. The difficulty lies in showing that the base of this cone is homeomorphic to $\mathbb{S}^{m+n-3}$ and hence that the cone has the structure of a regular CW complex with a single top dimensional cell of dimension $m + n - 2$. This is trivial for (m, n) equal to $(2, 2)$ or $(2, 3)$, as seen in the figure. To prove it for general (m, n) is not quite so easy.

2.1 Summary of results

In contrast to simplices which are indexed by dimension, bisimplices are indexed by two natural numbers $m, n \geq 1$. For each dimension $d \geq 2$ there are $\lceil \frac{d-1}{2} \rceil$ bisimplices of dimension d. Recall that a CW complex is *regular* if the characteristic maps of its cells are injective. See Figure 2–1.

Theorem A (Theorem 5.0.1)**.** *There exists a family* $\{⧖^{m,n}\}_{m,n \geq 1}$ *of regular CW complexes called bisimplices satisfying the following conditions.*

- $⧖^{m,n}$ *has a unique maximal cell and so* $⧖^{m,n}$ *is homeomorphic to a ball.*
- $⧖^{m,n}$ *has dimension* $m + n$*.*
- *The* 1*-skeleton of* $⧖^{m,n}$ *is the complete bipartite graph* $K_{m+1,n+1}$*.*

Moreover, the cells of a bisimplex $⧖$ are also bisimplices and these cells are precisely the full subcomplexes of $⧖$, aside from a few degenerate cases such as the $K_{0,\ell}$ and $K_{1,\ell}$ subgraphs of the 1-skeleton. We consider vertices and edges to be bisimplices also. These properties uniquely determine the cell posets of the bisimplices. However, proving that this family of posets is indeed a family of cell posets is not at all trivial and is an interesting application of the discrete Morse theory of Forman [13] and the Generalized Poincaré

Conjecture. Specifically, we prove Theorem A by inductively applying the following theorem, which we expect to have applications elsewhere.

Theorem B (Theorem 4.0.6 and Remark 4.0.7)**.** *Let P be a poset such that the order complexes of the under sets of P are PL-triangulated spheres. If P and all of its over sets admit spherical matchings then the order complex of P is a PL-triangulated sphere.*

A *spherical matching* is a combinatorial structure on the Hasse diagram of the cell poset of a regular CW complex X. This combinatorial structure is essentially a discrete Morse function on X having exactly two critical cells and so, by the Sphere Theorem of Forman [13], implies that X is homotopy equivalent to a sphere. See Chapter 4 for an introduction to discrete Morse theory and the definition of spherical matching. The other terminology of Theorem B is defined in Chapter 3.

Having constructed the family of bisimplices, we may construct regular CW complexes having bisimplices as cells. We call these *bisimplicial complexes* when the intersection of any two bisimplices is a full subcomplex. Given a bipartite graph Γ there is a natural bisimplicial complex $\mathbb{X}(\Gamma)$ called the *flag bisimplicial completion* of Γ. The flag bisimplicial completion is defined analogously to the flag simplicial completion, also known as the clique complex, of a graph.

Our primary motivation for the definition of the flag bisimplicial completion is to apply it to the bipartite 1-skeleta of quadric complexes. *Quadric complexes* may be defined as simply connected 2-dimensional CW complexes whose minimal area disk diagrams are CAT(0) square complexes. We would like a natural way to glue higher dimensional cells to a quadric complex to obtain a contractible supercomplex. The 1-skeleton of a quadric complex X is bipartite and may contain $K_{2,3}$ so it is not possible to extend X simplicially or cubically. However, X equals the 2-skeleton of $\text{⧖}(X^1)$, so a natural candidate for a contractible supercomplex is $\text{⧖}(X^1)$.

Theorem C (Theorem 7.0.6)**.** *Let X be a nonempty quadric complex. Then the flag bisimplicial completion $\text{⧖}(X^1)$ is contractible.*

Metric balls in X^1 induce finite quadric subcomplexes of X and finite quadric complexes have bi-dismantlable 1-skeleta [3, 25]. A finite bipartite graph is *bi-dismantlable* if it can be reduced to a nonempty connected complete bipartite graph by successively deleting a vertex whose neigbourhood is contained in the neighborhood of another vertex. Theorem C then follows from Theorem D below whose proof is another application of the discrete Morse theory of Forman.

Theorem D (Theorem 6.0.2)**.** *Let X be a flag, nonempty finite bisimplicial complex. If X^1 is bipartite and bi-dismantlable then X is collapsible.*

This method of proving contractibility mirrors that of Chepoi and Osajda for weakly systolic complexes [9] via *LC-contractibility* [10, 33].

As pointed out to the present author by Damian Osajda, a quadric complex X may also naturally be made contractible by extending each connected complete bipartite subgraph of X^1 to a complete subgraph and then taking the flag simplicial completion of the resulting graph. However, this operation preseves neither the 1-skeleton nor the 2-skeleton of X. Moreover, the resulting complex has higher dimension than the flag bisimplicial completion $\text{⧖}(X^1)$.

If X is a compact locally quadric complex, the construction of the bisimplicial completion of the universal cover $\widetilde{X}$ has a corresponding construction in the base. We obtain from X a compact complex X^+ whose 2-skeleton is X and whose higher cells are obtained by successively gluing in higher dimensional bisimplices along immersions of their boundaries. Then applying Theorem C we obtain the following.

Theorem E (Theorem 7.1.4)**.** *Let X be a compact locally quadric complex. If $\pi_1(X)$ is torsion-free then X^+ is a compact $K(\pi_1(X), 1)$.*

Note that $\pi_1(X)$ in Theorem E is torsion-free if and only if the automorphism group of every immersion of the 2-skeleton of a bisimplex into X is trivial. This is a consequence of the invariant biclique theorem for quadric complexes [25]. Moreover, every torsion-free quadric group is the fundamental group of some locally quadric complex.

2.2 Structure of Part I

In Chapter 3 we give some background on posets, regular CW complexes and PL-triangulated spheres. In Chapter 4 we present basic theorems of the discrete Morse theory of Forman. We apply these theorems and the topological Generalized Poincaré Conjecture to prove a discrete Morse-theoretic sphere recognition theorem. We use our sphere recognition theorem in Chapter 5 to construct the infinite family of bisimplices. We then prove some basic facts about bisimplices. In Chapter 6 we introduce bisimplicial complexes and prove that flag finite bisimplicial complexes with dismantlable 1-skeleta are collapsible, again making use of discrete Morse theory. Finally, in Chapter 7 we prove that the flag bisimplicial completion of a quadric complex is contractible and describe how to construct a $K(G, 1)$ for a torsion-free quadric group G.

Chapter 3

Posets and regular CW complexes

Let P be a poset. The *covering relation* C_P on P is the following binary relation.

$$C_P(x, y) \iff x < y \text{ and there is no } z \text{ satisfying } x < z < y$$

A poset P is *graded* if every element $x \in P$ is assigned a grade $|x| \in \mathbb{N}$ such that the following conditions hold.

$$\begin{aligned} C_P(x, y) &\implies |x| + 1 = |y| \\ x < y &\implies |x| < |y| \end{aligned}$$

Let P be a poset. For $x \in P$, the *over set* O_x and *under set* U_x *of* P *at* x are the following subsets of P.

$$O_x = \{y \in P : y > x\} \qquad U_x = \{y \in P : y < x\}$$

We may write O_x^P and U_x^P if the poset is not clear from the context. Note that making the inequalities in the definitions of O_x and U_x nonstrict would give what are usually referred to as the *principal ideal* and *principal filter* having *principal element* x. For $x, y \in P$, the *strict interval* (x, y) of P *between* x and y is the subset of P defined as follows.

$$(x, y) = O_x \cap U_y$$

The over sets, under sets and strict intervals of P are themselves posets by restricting the order relation. If P is graded then the over sets, under sets and strict intervals of P are likewise themselves graded.

Let P be a poset. The set of nonempty chains of P form an abstract simplicial complex. Its associated simplicial complex is the *order complex* Δ_P of P.

A CW complex is *regular* if the characteristic maps of its cells are embeddings. Let X be a regular CW complex. The cells of X are regular CW subcomplexes. Viewing a cell x as a ball, we denote its boundary by ∂x and its interior by x°. The *k-skeleton* X^k of X is the subcomplex of X formed by the union of the cells of X of dimension at most k. The *cell poset* P_X of X is the set of cells of X ordered by inclusion. Cell posets are equipped with a natural grading, namely dimension: $|x| = \dim x$. A cell poset P uniquely determines its regular CW complex X_P. The order complex of the cell poset of a regular CW complex X is isomorphic to the barycentric subdivision of X. A subset Q of P is the cell poset of a subcomplex of X_P iff Q is *downward closed*, meaning the following.

$$x \in Q \text{ and } y < x \quad \implies \quad y \in Q$$

The following theorem of Björner characterizes the cell posets.

Theorem 3.0.1 (Björner [4, Proposition 3.1]). *Let P be a poset. Then P is a cell poset iff the order complexes of its under sets are homeomorphic to spheres.*

Proof. If P is the cell poset of a regular CW complex X then its under sets are the cell posets of the boundaries of its cells. The order complexes of these under sets are the barycentric subdivisions of these cells and so are homeomorphic to spheres.

To prove the converse, we construct X_P inductively on dimension. Define the height function $h\colon P \to \mathbb{N}$ as follows.

$$h(x) = \max\bigl\{|C| - 1 : C \text{ is a chain in } P \text{ with maximum } x\bigr\}$$

Note that $h(x)$ is finite because, otherwise, the order complex of the under set at x would be infinite dimensional. We have that $h(x) = 0$ for minimal elements of P and that $h(x) < h(y)$ for $x < y$. We will show that h is a grading on P. Let $x \in P$. Since the under set U_x has order complex homeomorphic to a sphere $\mathbb{S}^\ell$, the maximal chains of U_x must all have size $\ell + 1$. Hence $h(x) = \ell + 1$ and $h(y) = \ell$ for any element of P that is covered by x.

Let $P^k \subseteq P$ be defined by the following.

$$P^k = \{x \in P : h(x) \leq k\}$$

We will construct X_P such that its k-skeleton X_P^k has cell poset P^k. We begin the induction by letting $X_P^0 = P^0$. Suppose we have constructed X_P^k having cell poset P^k. Let $x \in P$ with $h(x) = k + 1$. Since h is a grading, we have $U_x \subseteq P^k$. So U_x is the cell poset of a subcomplex A of X_P^k. The barycentric subdivision of A_x is isomorphic to the order complex of U_x which, by hypothesis, is homeomorphic to a sphere. This sphere has dimension k since that is the height of a maximal chain in U_x. We construct X_P^{k+1} from X_P^k by attaching a $(k+1)$-ball along its boundary to each A_x with $h(x) = k+1$. Then P^{k+1} is the cell poset of X_P^{k+1}. Having inductively defined the skeleta

$$X_P^0 \subseteq X_P^1 \subseteq X_P^2 \subseteq \cdots$$

we obtain X_P as the colimit. $\square$

We consider the empty space to be the sphere of dimension -1.

Let P be a cell poset. The under sets of P are also cell posets. More precisely, the under set U_x at a cell x is the cell poset of the regular CW complex structure on the boundary of x.

The order complex of the over set O_x is isomorphic to the link of the barycenter of x in the barycentric subdivision of X_P. However, because of the existence of homology spheres, O_x need not be a cell poset. A homology sphere is a manifold with the homology of a sphere but which is not homeomorphic to a sphere. The double suspension of a homology sphere is homeomorphic to a sphere, as first proved in full generality by Cannon [6]. The Poincaré homology sphere X, also known as the spherical dodecahedron space, is a homology 3-sphere that has a simplicial triangulation [47, Section 62]. Let B be the regular CW complex with a single cell of dimension 6 and whose boundary ∂B has the structure of the simplicial double suspension of X. Let e be a 1-simplex of B joining two of the suspension points of ∂B. Then the link $\operatorname{lk} e$ is isomorphic to X and so the over set O_e of e in P_B is the cell poset of X augmented with a new maximum element corresponding to the top-dimensional cell of B. The under set of B in O_e then has order complex homeomorphic to X and not a sphere and hence O_e is not a cell poset.

If X_P is a simplicial complex, then the over set O_x at a cell x is also a cell poset. In fact, O_x is the cell poset of the link $\operatorname{lk} x$ of x. Theorem 3.0.5 characterizes the cell posets in which this holds.

Proposition 3.0.2. *Let P be a cell poset and suppose X_P is connected. If the over sets of P at its minimal elements have order complexes homeomorphic to spheres then X_P is a manifold.*

Proof. The order complexes of the over sets of P at minimal elements are isomorphic to the links of vertices of the order complex of P. Hence the links of vertices of the barycentric subdivision of X_P are homeomorphic to spheres.

That X_P is connected ensures that these spheres all have the same dimension, say $d-1$, and that X_P is second countable. Then X_P is a d-manifold. $\square$

A *PL-triangulated manifold* is a simplicial complex X that is homeomorphic to a manifold such that the link of every simplex of X is homeomorphic to a sphere [27]. PL-triangulated manifolds are referred to as *combinatorial manifolds* in the PL-topology literature.

Proposition 3.0.3. *Let P be a cell poset and suppose X_P is connected. Then the order complex of P is a PL-triangulated manifold iff the order complexes of the strict intervals and over sets of P are homeomorphic to spheres.*

Proposition 3.0.3 is an immediate consequence of the following lemma.

Lemma 3.0.4. *Let P be a cell poset and suppose X_P is connected. Let P' be the cell poset of the order complex Δ_P. The order complexes of the over sets of P' are homeomorphic to spheres iff the order complexes of the strict intervals and over sets of P are homeomorphic to spheres.*

Indeed, the over sets of P' are the cell posets of the links of the simplices of Δ_P. If Δ_P is a PL-triangulated manifold then these are all homeomorphic to spheres. Conversely, if the links are all homeomorphic to spheres then since X_P is connected, Proposition 3.0.2 ensures that Δ_P is homeomorphic to a manifold. Then, by definition, X_P is a PL-triangulated manifold.

Before proving Lemma 3.0.4 we study the over sets of cell posets of order complexes. Let P be a poset and let P' be the cell poset of the order complex of P. Every $c \in P'$ is a nonempty chain

$$c = \{c_0, c_1, \dots, c_k\} \subset P$$

with

$$c_0 < c_1 < \cdots < c_k$$

in P and these chains are ordered by inclusion. Each element of the over set O_c is a chain in P containing c and so is determined by its intersections with U_{c_0}, with O_{c_k} and with the strict intervals (c_{i-1}, c_i). It follows that O_c embeds in the componentwise product order

$$(U_{c_0})'_{\perp} \times (c_0, c_1)'_{\perp} \times (c_1, c_2)'_{\perp} \times \cdots \times (c_{k-1}, c_k)'_{\perp} \times (O_{c_k})'_{\perp}$$

where $Q'_{\perp}$ denotes the poset of *all* chains (including the empty chain) in a poset Q. Aside from the presence of a minimum element corresponding to the empty simplex, this product is isomorphic to the cell poset of the simplicial join of the order complexes of U_{c_0}, O_{c_k} and the (c_{i-1}, c_i). The complement of this minimum element is the image of O_c under its embedding in the product. Hence, X_{O_c} is isomorphic to the simplicial join

$$X_{O_c} \cong O_{U_{c_0}} \bowtie O_{(c_0,c_1)} \bowtie O_{(c_1,c_2)} \cdots \bowtie O_{(c_{k-1},c_k)} \bowtie O_{O_{c_k}}$$

of the order complexes of U_{c_0}, O_{c_k} and the (c_{i-1}, c_i).

Proof of Lemma 3.0.4. The joins of spheres are spheres so, by the discussion above, the "if" part of Lemma 3.0.4 has been established. It remains to prove the "only if" part.

Assume that the order complexes of the over sets of P' are homeomorphic to spheres. Let $x, y, z \in P$ with $x < y$. Let $c^x \in P'$ and $c^z \in P'$ be maximal chains of P that have x and z as their maximums. Let $c_y \in P'$ be a maximal chain of P that has y as its minimum. Then we have

$$X_{O_{c_y \cup c^x}} \cong O_{(x,y)}$$

and

$$X_{O_{c^z}} \cong \Delta_{O_z}$$

and so $O_{(x,y)}$ and O_{O_z} are homeomorphic to spheres. □

A *PL-triangulated sphere* is a PL-triangulated manifold that is homeomorphic to a sphere.

Theorem 3.0.5. *Let P be a cell poset. The following conditions are equivalent.*

1. *The order complexes of the under sets of P are PL-triangulated spheres.*
2. *The order complexes of the strict intervals of P are PL-triangulated spheres.*
3. *The order complexes of the strict intervals of P are homeomorphic to spheres.*
4. *The over sets of P are cell posets.*

Theorem 3.0.5 characterizes the regular CW complexes X for which each d-cell x may be associated a link having the structure of a regular CW complex in which the $(k-d-1)$-cells naturally correspond to the k-cells of X that are incident to x. Theorem 3.0.5 says that this holds precisely when the boundaries of the cells of X are PL-triangulated spheres.

Proof of Theorem 3.0.5. (1) $\Longrightarrow$ (2) Let (x, y) be a strict interval of P. Then (x, y) is the over set of U_y at x. So, by the forward implication of Proposition 3.0.3, the order complex of (x, y) is homeomorphic to a sphere. Every over set and strict interval of (x, y) is a strict interval of P and so, by the same argument, must have order complex homeomorphic to a sphere. Then, by the reverse implication of Proposition 3.0.3, the order complex of (x, y) is a PL-triangulated sphere.

(2) $\Longrightarrow$ (3) This is clear.

(3) $\Longrightarrow$ (4) Let O_x be a over set of P. By Theorem 3.0.1 we need only show that the under set of O_x at any $y \in O_x$ has order complex homeomorphic to a sphere. But the under set of O_x at y is the strict interval (x, y) of P and so this holds.

(4) $\Longrightarrow$ (1) Let U_z be a under set of P. By Theorem 3.0.1, the order complex of U_z is homeomorphic to a sphere. To prove that it is a PL-triangulated sphere it suffices, by Proposition 3.0.3, to show that, for $x < y < z$, the strict interval (x, y) and the over set $O_x^{U_y}$ of U_y at x have order complexes homeomorphic to spheres. But (x, y) and $O_x^{U_y}$ are equal to the under sets $U_z^{O_x}$ and $U_y^{O_x}$ of O_x and so, by Theorem 3.0.1, they have order complexes homeomorphic to spheres. $\square$

Chapter 4

Forman Morse theory and the recognition of spheres

Let P be a finite graded poset. The *Hasse diagram* Γ_P of P is the covering relation C_P viewed as a directed graph. A *matching* M on P is a set of pairwise disjoint closed edges of Γ_P. An element $x \in P$ is *matched* by M if it is contained in an edge of M. A matching M on P is *acyclic* if the directed graph Γ_P^M obtained from Γ_P by reversing the direction on the edges of M has no directed cycles. An element $x \in P$ is a *critical element* of M if x is not matched by M. Acyclic matchings are also known as *Morse matchings*. If P is a cell poset then an acyclic matching M on P determines a Forman discrete Morse function [13] on X_P with the same set of critical cells [7]. The language of acyclic matchings for discrete Morse theory is due to Chari [7].

Proposition 4.0.1. *Let P be a finite graded poset, let M be a matching on P and let $(\gamma_0, \gamma_1, \gamma_2, \ldots, \gamma_k)$ be the sequence of edges of a directed cycle γ of Γ_P^M. Then no consecutive pair of edges (γ_i, γ_{i+1})—indices modulo k—has both edges in M or both edges in the complement of M.*

Proof. Following an edge of M causes a unit decrease in the grading. Following an edge not in M causes a unit increase in the grading. Hence γ must contain the same number of edges in M as it does edges not in M. Since M is a matching, there is no consecutive pair of edges of γ both contained in M. Suppose we have a consecutive pair of edges (γ_i, γ_{i+1}) neither of which are contained in M. Then there are two more M-edges in $(\gamma_{i+2}, \gamma_{i+3}, \ldots, \gamma_{i+k})$

then there are non-M-edges. Hence there is a consecutive pair of edges of γ both contained in M, contradicting the hypothesis that M is a matching. $\square$

We require the following basic theorems of Forman discrete Morse theory.

Theorem 4.0.2 (Forman [13]). *Let P be a cell poset. Let M be an acyclic matching on P and let Q be the set of critical cells of M. If Q is downward closed, then X_Q can be obtained from X_P by a sequence of elementary collapses. In particular, X_Q is homotopy equivalent to X_P.*

Theorem 4.0.3 (Forman [13, Corollary 3.5]). *Let P be a cell poset and let M be an acyclic matching on P. Then X_P is homotopy equivalent to a CW complex with as many cells of each dimension as M has critical cells of that dimension.*

Let P be a finite graded poset. A *spherical matching* on P is an acyclic matching M on P with two critical cells.

Theorem 4.0.4 (Sphere Theorem of Forman [13, Theorem 5.1(1)]). *Let P be a cell poset. If P has a spherical matching then X_P is homotopy equivalent to a sphere.*

We also require the Generalized Poincaré Conjecture for topological manifolds.

Theorem 4.0.5 (Topological Generalized Poincaré Conjecture). *A closed topological manifold X is homotopy equivalent to the d-sphere iff it is homeomorphic to the d-sphere.*

The first breakthrough in the proof of the Generalized Poincaré Conjecture was made by Smale, who proved that a PL-triangulated manifold X that is homotopy equivalent to the d-sphere is homeomorphic to the d-sphere, for $d \geq 5$ [48]. Stallings gave a different proof of this fact for $d \geq 7$ using an "engulfing" method [50]. This method was later extended by Zeeman to prove the cases $d = 5$ and $d = 6$ [54]. Newman generalized the engulfing method to topological manifolds and thus completed the proof of the Topological Generalized Poincaré Conjecture for $d \geq 5$ [35, Theorem 7]. In dimension 4 the conjecture was proved by Freedman [14]. In dimension 3 it was proved by Perelman [41, 43, 42]. In dimensions at most 2, it follows from the classification of manifolds.

Theorem 4.0.6. *Let P be a poset such that the order complexes of the under sets of P are PL-triangulated spheres. The order complex of P is a PL-triangulated sphere iff the order complexes of P and all of its over sets are homotopy equivalent to spheres.*

Remark 4.0.7. Let P be as in Theorem 4.0.6. By Theorem 3.0.1 and Theorem 3.0.5, P and its over sets are all cell posets. So if Q is equal to P or to one of its over sets then to show that the order complex of Q is homotopy equivalent to a sphere it suffices to show that Q has a spherical matching. This holds by Theorem 4.0.4 and the fact that the order complex of Q is the barycentric subdivision of X_Q.

Proof of Theorem 4.0.6. We prove the "if" part since the "only if" part is immediate. The proof is by induction on the maximum size k of a chain in P. If $k = 0$ then P is the empty poset and so has the PL-triangulated -1-sphere (i.e. the empty simplicial complex) as its order complex.

Suppose $k > 0$ and that the theorem holds for all lesser values of k. Take $x \in P$. We show that O_x satisfies the conditions of the theorem. The under sets of O_x are strict intervals of P and so, by Theorem 3.0.5, the order complexes of the under sets of O_x are PL-triangulated spheres. By assumption O_x is homotopy equivalent to a sphere as are its over sets since they are also over sets of P. Hence, by the inductive hypothesis, the order complex of O_x is a PL-triangulated sphere. By Theorem 3.0.5, the strict intervals of P are also PL-triangulated spheres and so, by Proposition 3.0.3, the order complex O_P of

P is a PL-triangulated manifold. By assumption, O_P is homotopy equivalent to a sphere so, by Theorem 4.0.5, O_P is homeomorphic to a sphere and so is a PL-triangulated sphere. $\square$

Chapter 5

Bisimplices

A *bipartitioned set* is a set S along with a bipartition $S = A \sqcup B$ into subsets that are possibly empty. A subset $T \subseteq S$ is considered to be a bipartitioned set with its induced bipartition $T = (T \cap A) \sqcup (T \cap B)$. Our goal is to span cells on certain subsets of a bipartitioned set, just as simplices are spanned on subsets of vertices of a simplicial complex. In our case, however, not all subsets are eligible to span a cell so we introduce the term spanworthy. A bipartitioned set $S = A \sqcup B$ is *spanworthy* if $S \neq \emptyset$ and the following holds.

$$|A| \leq 1 \Longleftrightarrow |B| \leq 1$$

Spanworthiness excludes precisely the following cases.

- $\emptyset \sqcup \emptyset$
- $A \sqcup \emptyset$ with $|A| \geq 2$
- $\emptyset \sqcup B$ with $|B| \geq 2$
- $A \sqcup \{b\}$ with $|A| \geq 2$
- $\{a\} \sqcup B$ with $|B| \geq 2$

In particular, it excludes any S of cardinality 3.

Theorem 5.0.1. *Let $S = A \sqcup B$ be a spanworthy bipartitioned set and let P be the collection of spanworthy proper subsets of S ordered by inclusion. Then P is a cell poset and O_P is a PL-triangulated sphere.*

Proof. Let $A = \{a_0, a_1, \ldots, a_m\}$ and $B = \{b_0, b_1, \ldots, b_n\}$.

If $m = -1$ or $n = -1$ then, by spanworthiness, S is a singleton and so P is empty. Then P is the cell poset of the empty simplicial complex, i.e., the PL-triangulated -1-sphere. The PL-triangulated -1-sphere is equal to its own barycentric subdivision O_P so the theorem holds in this case.

Assume that $m \geq 0$ and $n \geq 0$. If $m = 0$ or $n = 0$ then, by spanworthiness, $m = n = 0$ and P is the poset with two incomparable elements. This is the cell poset of the two point simplicial complex, i.e., the PL-triangulated 0-sphere. The PL-triangulated 0-sphere is equal to its own barycentric subdivision O_P so the theorem holds in this case.

Assume that $m \geq 1$ and $n \geq 1$. If $m = n = 1$ then the elements of P are the singletons and the $\{a_i\} \sqcup \{b_j\}$ for i and j ranging over 0 and 1. So P is isomorphic to the cell poset of the 4-cycle and so the theorem holds.

So, by symmetry, we may assume that $m \geq 1$ and $n > 1$. Assume that the theorem holds for all S of lesser cardinality. Then the order complexes of the under sets of P are PL-triangulated spheres. So, by Theorem 3.0.1, P is a cell poset and, by Theorem 4.0.6, it suffices to show that the order complexes of P and all of its over sets are homotopy equivalent to spheres.

Let $T \in P$. Spanworthiness implies that $|T| \neq 3$. If $|T| > 3$ then every proper subset of S containing T is spanworthy and so O_T is isomorphic to the cell poset of the boundary of a simplex and so has order complex homeomorphic to a sphere. So it remains only to show that the order complexes of P and O_T for $|T| \in \{1, 2\}$ are homotopy equivalent to spheres. By Remark 4.0.7 it suffices to show that P and such O_T have spherical matchings. By symmetry we need only consider the cases $T = \{a_0\} \sqcup \emptyset$, $T = \emptyset \sqcup \{b_0\}$ and $T = \{a_0\} \sqcup \{b_0\}$.

Consider the following families of edges of the Hasse diagram Γ_P of P.

$$M_1 = \left\{ \quad \{a_i\} \sqcup \emptyset \to \{a_i\} \sqcup \{b_n\} \quad \right\}$$

$$M_2 = \left\{ \quad \emptyset \sqcup \{b_j\} \to \{a_m\} \sqcup \{b_j\} \quad : \quad j \neq n \right\}$$

$$M_3 = \left\{ \{a_i\} \sqcup \{b_j\} \to \{a_i, a_m\} \sqcup \{b_j, b_n\} \quad : \quad i \neq m,\ j \neq n \right\}$$

$$M_4 = \left\{ A' \sqcup \{b_j, b_n\} \to \big(A' \cup \{a_m\}\big) \sqcup \{b_j, b_n\} : \quad \begin{array}{r} a_m \notin A',\ j \neq n \\ |A'| \geq 2 \end{array} \right\}$$

$$M_5 = \left\{ \quad A' \sqcup B' \to A' \sqcup \big(B' \cup \{b_n\}\big) \quad : \quad \begin{array}{r} b_n \notin B' \\ A' \sqcup B' \neq A \sqcup \big(B \setminus \{b_n\}\big) \\ |A'| \geq 2,\ |B'| \geq 2 \end{array} \right\}$$

Recall that $B = \{b_0, b_1, \ldots, b_n\}$ and we have assumed $n > 1$. Thus $|B| > 2$ and so the terminal endpoints of edges in M_3 and M_4 are proper subsets of $A \sqcup B$ and hence are elements of P.

The endpoints of these edges from different families or from different ends of edges in the same family can be distinguished by the cardinality of their parts and by the presence of a_m and b_n, as shown in Table 5–1. Moreover the initial endpoints of two edges from the same family are equal if and only if their terminal endpoints are equal. Hence we see that $M = \cup_i M_i$ forms a matching on P.

Let Γ_P^M be the directed graph obtained from Γ_P by reversing the direction of each edge in M. By Proposition 4.0.1, to show that M is an acyclic matching we need only show that Γ_P^M does not contain any directed cycles whose edges alternate between being contained and not contained in M. To do this it suffices to define a function $\alpha\colon P \to \mathbb{N}$ such that $\alpha(T_2) < \alpha(T_0)$ for any

Family	Initial Endpoint				Terminal Endpoint			
	Cardinalities		a_m	b_n	Cardinalities		a_m	b_n
M_1	1	0		$\perp$	1	1		$\top$
M_2	0	1	$\perp$	$\perp$	1	1	$\top$	$\perp$
M_3	1	1	$\perp$	$\perp$	2	2	$\top$	$\top$
M_4	≥ 2	2	$\perp$	$\top$	> 2	2	$\top$	$\top$
M_5	≥ 2	≥ 2		$\perp$	≥ 2	> 2		$\top$

Table 5–1: Distinguishing characteristics of the endpoints of edges in the families of edges described in the proof of Theorem 5.0.1. Under a_m or b_n, the symbol $\top$ indicates that this element is present in every member of the family and the symbol $\perp$ indicates that this element is not present in any member of the family.

directed path

$$T_0 \xrightarrow{e_0} T_1 \xrightarrow{e_1} T_2$$

of Γ_P^M with $e_0 \in M$ and $e_1 \notin M$. Note that in Γ_P, e_0 is directed from T_1 to T_0 so we have the following inclusions.

$$T_0 \supsetneq T_1 \subsetneq T_2$$

We define α as follows.

$$\alpha(T) = \begin{cases} 0, & a_m \notin T \text{ and } b_n \notin T \\ 1, & a_m \in T \text{ and } b_n \notin T \\ 2, & a_m \notin T \text{ and } b_n \in T \\ 3, & a_m \in T \text{ and } b_n \in T \end{cases}$$

We may think of α as a function summing the weights on the elements of T, where a_m is assigned a weight of 1 and b_n is assigned a weight of 2 and all remaining elements have zero weight. Since $T_1 = T_0 \setminus A$ for some nonempty $A \subseteq \{a_m, b_n\}$, we have $\alpha(T_1) < \alpha(T_0)$. Suppose $\alpha(T_2) = \alpha(T_0)$. Then $T_1 \cup A \subseteq T_2$ and, since $|T_0| - |T_1| = |T_2| - |T_1|$, we have $T_0 = T_2$. This is a contradiction since a pair of vertices of Γ_P^M may be joined by at most one edge and this edge is directed in a unique way. Suppose $\alpha(T_2) > \alpha(T_0)$. Then $T_1 = T_0 \setminus \{a_m\}$ and $T_2 = T_1 \cup \{b_n\}$. The equality $T_1 = T_0 \setminus \{a_m\}$ implies that e_0 is the reverse of an edge in M_2 or M_4. The equality $T_2 = T_1 \cup \{b_n\}$ and the fact that $T_2 \neq T_1$ implies that $b_n \notin T_1$. This rules out the possibility that e_0 is the reverse of an edge in M_4. Thus we have

$$\{a_m\} \sqcup \{b_j\} \xrightarrow{e_0} \emptyset \sqcup \{b_j\}$$

and so $T_1 = \emptyset \sqcup \{b_j\}$ and $T_2 = \emptyset \sqcup \{b_j, b_n\}$ which is not spanworthy, a contradiction. We have established that M is an acyclic matching.

Let $T \in P$. Since the Hasse diagram of O_T is an induced subgraph of Γ_P, the subset $M_T \subset M$ consisting of all edges both of whose endoints are contained in O_T is an acyclic matching on O_T. It remains only to show that M is spherical on P and that M_T is spherical on O_T for $T = \{a_0\} \sqcup \emptyset$, $T = \emptyset \sqcup \{b_0\}$ and $T = \{a_0\} \sqcup \{b_0\}$. In fact it will suffice to prove that M is spherical on P with critical elements $\emptyset \sqcup \{b_n\}$ and $A \sqcup \big(B \setminus \{b_n\}\big)$. Indeed, in this case the only critical element of M contained in O_T would be $A \sqcup \big(B \setminus \{b_n\}\big)$. The only other possible critical elements of O_T would arise from edges of M having one endpoint in O_T and the other endoint in $P \setminus O_T$. But there is a unique such edge of M, namely the edge with initial endpoint T. Hence M_T would have two critical elements.

We now prove that M is spherical with critical elements $\emptyset \sqcup \{b_n\}$ and $A \sqcup \big(B \setminus \{b_n\}\big)$. First we verify that these elements are indeed unmatched by M. Singletons in B appear as endpoints only in M_2 where $\emptyset \sqcup \{b_n\}$ is not present so $\emptyset \sqcup \{b_n\}$ is critical. The element $A \sqcup \big(B \setminus \{b_n\}\big)$ is maximal in P and so may only appear as a terminal endpoint of an edge of M. These all contain b_n except those in M_2 where they have the form $\{a_m\} \sqcup \{b_j\}$. Such an element cannot be equal to $A \sqcup \big(B \setminus \{b_n\}\big)$ since then $A \sqcup B = \{a_m\} \sqcup \{b_j, b_n\}$ which is not spanworthy.

Now, suppose $T = A' \sqcup B'$ is an element of P that is not equal to $\emptyset \sqcup \{b_n\}$ or $A \sqcup \big(B \setminus \{b_n\}\big)$. We will show that T is matched in M. We consider the following cases separately: (I) $|T| = 1$, (II) $|T| = 2$, (III) $|T| = 4$, (IV) $|T| > 4$ and $|B'| = 2$, (V) $|T| > 4$ and $|B'| > 2$.

Case I. $|T| = 1$. If $|A'| = 1$ then T is an initial endpoint in M_1. Otherwise $|B'| = 1$ and T is an initial endpoint in M_2.

Case II. $|T| = 2$. Then $|A'| = |B'| = 1$ by spanworthiness. If $b_n \in T$ then T is a terminal endpoint of M_1. Otherwise T is a terminal endpoint in M_2 if $a_m \in T$ and T is an initial endpoint in M_3 if $a_m \notin T$.

Case III. $|T| = 4$. Then $|A'| = |B'| = 2$ by spanworthiness. If $b_n \notin T$ then T is an initial endpoint of M_5. Otherwise T is a terminal endpoint in M_3 if $a_m \in T$ and T is an initial endpoint in M_4 if $a_m \notin T$.

Case IV. $|T| > 4$ and $|B'| = 2$. Then $|A'| > 2$. If $b_n \notin T$ then T is an initial endpoint of M_5. Otherwise T is a terminal endpoint in M_4 if $a_m \in T$ and T is an initial endpoint in M_4 if $a_m \notin T$.

Case V. $|T| > 4$ and $|B'| > 2$. Then $|A'| \geq 2$ by spanworthiness. If $b_n \in T$ then T is a terminal endpoint in M_5. Otherwise T is an initial endpoint in M_5. $\square$

Corollary 5.0.2. *Let $S = A \sqcup B$ be a spanworthy bipartitioned set and let P' be the collection of spanworthy subsets of S ordered by inclusion. Then P' is the cell poset of a regular CW complex homeomorphic to a ball.*

Note that the difference between P' in Corollary 5.0.2 and P in Theorem 5.0.1 is that P' contains S.

Proof of Corollary 5.0.2. By Theorem 3.0.1, it suffices to show that the order complexes of the under sets of P' are homeomorphic to spheres. Since S is the maximum in P', the under sets of P' are $P = P' \setminus \{S\}$ and the under sets of P. By Theorem 5.0.1, we know that P is the cell poset of a regular CW complex X_P that is homeomorphic to a sphere. Hence, by Theorem 3.0.1, the under sets of P have order complexes homeomorphic to spheres and the order complex of P is the barycentric subdivision of X_P and so is also homeomorphic to a sphere. □

Let $S = A \sqcup B$ be a spanworthy bipartitioned set and let P' be the collection of spanworthy subsets of S ordered by inclusion. By Corollary 5.0.2, P' is the cell poset of a regular CW complex $X_{P'}$ that is homeomorphic to a ball. A regular CW complex isomorphic to $X_{P'}$ is an (m,n)-*bisimplex* where $m = |A| - 1$ and $n = |B| - 1$. We let $⧖^{m,n}$ denote an (m,n)-bisimplex. See Figure 2–1.

Proposition 5.0.3. *There is an isomorphism* $⧖^{m,n} \cong ⧖^{n,m}$.

Proof. This is clear from the symmetry of the definition. □

Proposition 5.0.4. *The cells of a bisimplex are all bisimplices.*

Proof. The cell poset P of a bisimplex $⧖$ is isomorphic to the poset of spanworthy subsets of a spanworthy set $A \sqcup B$. The cell poset of a cell x of $⧖$ corresponds to $P' = U_T \cup \{T\}$ for some spanworthy T. Hence, the cell poset of x is isomorphic to the poset of spanworthy subsets of the spanworthy set T. □

Proposition 5.0.5. *The* 1*-skeleton of* $⧖^{m,n}$ *is isomorphic to the complete bipartite graph* $K_{m+1,n+1}$ *on* $m+1$ *and* $n+1$ *vertices.*

Proof. let P be the cell poset of $⧖^{m,n}$ viewed as the poset of spanworthy subsets of $A \sqcup B$ with $|A| = m+1$ and $|B| = n+1$. The 0-cells of $⧖^{m,n}$ correspond to the minimal elements of P. These are precisely the singletons in $A \sqcup B$ and so we may identify the 0-skeleton of $⧖^{m,n}$ with $A \sqcup B$. The 1-cells of $⧖^{m,n}$ correspond to those elements of P that cover singletons. These are precisely the sets $\{a\} \sqcup \{b\}$ with $a \in A$ and $b \in B$. $\square$

Proposition 5.0.6. *Let $m \geq 1$ and let $n \geq 1$. The dimension of $⧖^{m,n}$ is $m + n$.*

Proof. Let P be the poset of spanworthy subsets of $A \sqcup B$ with

$$A = \{a_0, a_1, \ldots, a_m\}$$

and

$$B = \{b_0, b_1, \ldots b_n\}$$

and identify P with the cell poset of $⧖^{m,n}$. Since $⧖^{m,n}$ is homeomorphic to a ball of some dimension k, the maximal chains of P all have cardinality $k + 1$. One such maximal chain is the following.

$$\begin{aligned}
&\{a_0\} \sqcup \emptyset \subsetneq \{a_0\} \sqcup \{b_0\} \subsetneq \{a_0, a_1\} \sqcup \{b_0, b_1\} \subsetneq \{a_0, a_1, a_2\} \sqcup \{b_0, b_1\} \\
&\subsetneq \{a_0, a_1, a_2, a_3\} \sqcup \{b_0, b_1\} \subsetneq \cdots \subsetneq \{a_0, \ldots, a_m\} \sqcup \{b_0, b_1\} \\
&\subsetneq \{a_0, \ldots, a_m\} \sqcup \{b_0, b_1, b_2\} \subsetneq \cdots \subsetneq \{a_0, \ldots, a_m\} \sqcup \{b_0, \ldots, b_n\}
\end{aligned}$$

This chain has cardinality $|A| + |B| - 1$ where the -1 is due to the jump $\{a_0\} \sqcup \{b_0\} \subsetneq \{a_0, a_1\} \sqcup \{b_0, b_1\}$. Hence $⧖^{m,n}$ has dimension $k + 1 - 1 = |A| + |B| - 1 - 1 = m + n$. $\square$

Chapter 6

Bisimplicial complexes

A full subcomplex Y of a regular CW complex X is *full* if $\partial x \subset Y$ implies $x \subset Y$ for any cell x of X. A *bisimplicial complex* is a regular CW complex X such that each cell x of X is isomorphic to a bisimplex and, for any two bisimplices x and y of X, the intersection $x \cap y$ is a full subcomplex of X. Note that this implies that the bisimplices themselves are full subcomplexes and, furthermore, that any finite intersection of bisimplices is full.

A complete bipartite graph K is *spanworthy* if it is nonempty, connected and the bipartition on its vertex set is spanworthy. A spanworthy complete bipartite subgraph K of the 1-skeleton X^1 of a bisimplicial complex X *spans* a bisimplex $⧖$ of X if the 1-skeleton $(⧖)^1$ of $⧖$ is equal to K. Note that at most one bisimplex may span K since the intersection of two distinct bisimplices $⧖$ and $⧖'$ spanning K would be full in neither $⧖$ nor $⧖'$. A bisimplicial complex X is *flag* if every spanworthy complete bipartite subgraph K of X^1 spans a bisimplex $⧖$. We use the notation $⧖(A; B)$ to denote $⧖$, where $A \sqcup B$ is the bipartitioned vertex set of K.

Definition 6.0.1. Let Γ be a graph. The *flag bisimplicial completion* $⧖(\Gamma)$ of Γ is a flag bisimplicial complex defined inductively as follows. The 1-skeleton of $⧖(\Gamma)$ is Γ. Now, assume the $(k-1)$-skeleton of $⧖(\Gamma)$ has been defined. The k-skeleton is obtained by the following operation. To each subcomplex isomorphic to some $\partial ⧖^{m,n}$ with $\dim(⧖^{m,n}) = k$, glue in a copy of $⧖^{m,n}$ along the isomorphism.

Note that if X is a flag bisimplicial complex then $X = \boxtimes(X^1)$.

Let Γ be a finite bipartite graph. We view Γ^0 as a metric space with the shortest path metric. The *metric sphere* $S_r(u) \subseteq \Gamma^0$ of radius r about $u \in \Gamma^0$ is the set of vertices of Γ at distance r from u. If u and v are distinct vertices of Γ then u is *dominated* by v if there is an inclusion $S_1(u) \subset S_1(v)$ of neighbourhoods.

A finite bipartite graph Γ is *bi-dismantlable* if there exists a sequence $\Gamma = \Gamma_1, \Gamma_2, \ldots, \Gamma_n$ of graphs ending on a nonempty connected complete bipartite graph such that, for each $i < n$, Γ_{i+1} is a subgraph of Γ_i induced on the complement of $\{v_i\}$ for some v_i dominated in Γ_i.

Theorem 6.0.2. *Let X be a finite flag bisimplicial complex with X^1 bipartite. If X^1 is bi-dismantlable then X is collapsible.*

Proof. The proof is by induction on the length of the bi-dismantling sequence.

In the base case, X^1 is a nonempty connected complete bipartite graph on some bipartitioned vertex set $S = A \sqcup B$. Let $A = \{a_0, \ldots, a_m\}$ and $B = \{b_0, \ldots, b_n\}$. Without loss of generality $|A| \leq |B|$. If X^1 is not spanworthy then, as it is nonempty and connected, we have $|A| = 1$ and $|B| \geq 2$. Then the only spanworthy subgraphs of X^1 are its edges and vertices and so $X = X^1$. But X^1 is a tree and so X is collapsible.

Suppose now that X^1 is spanworthy. By flagness, X is a bisimplex $\boxtimes(A; B)$. If $|A| + |B| \leq 4$ then $\boxtimes(A; B)$ is isomorphic to a vertex, an edge or a square. These are collapsible. So we assume that $|A| \geq 2$ and $|B| > 2$. Identify the cell poset of $\boxtimes(A; B)$ with the poset P' of nonempty spanworthy subsets of $A \sqcup B$. Then the poset $P = P' \setminus \{A \sqcup B\}$ is the cell poset of $\partial \boxtimes(A; B)$. The proof of Theorem 5.0.1 gives a spherical matching M on P. Let M' be the matching obtained from M by adding the following edge.

$$A \sqcup \big(B \setminus \{b_n\}\big) \to A \sqcup B$$

Then M' is acyclic and leaves only $\emptyset \sqcup \{b_n\}$ unmatched. Hence, by Theorem 4.0.2, $X = \bowtie(A; B)$ is collapsible.

Now, suppose suppose the bi-dismantling sequence has nonzero length with v the first dominated vertex in the sequence. Let u be a dominator of v in X^1. Let P be the cell poset of X. Consider the downward closed subset Q of P defined as follows.

$$Q = \{x \in P : x \not\geq v\}$$

Then the subcomplex $X_Q = \bigcup Q$ is the full subcomplex of X induced on $X^0 \setminus \{v\}$ and so X_Q is flag. Moreover, X^1_Q is the induced subgraph of X^1 obtained from X^1 by deleting v and so X^1_Q is dismantlable. Hence X^1_Q is collapsible by induction. Therefore, by Theorem 4.0.2, it suffices to construct an acyclic matching M on P whose set of critical elements is Q.

Let w be any neighbour of v in X^1. Note that the vertices of any connected complete bipartite subgraph of X^1 containing v are at distance at most 2 from v. Consider the following families of edges in the Hasse diagram Γ_P of P.

$$M_1 = \Big\{ \bowtie\big(\{v\}; \emptyset\big) \to \bowtie\big(\{v\}; \{w\}\big) \Big\}$$

$$M_2 = \Big\{ \bowtie\big(\{v\}; \{x\}\big) \to \bowtie\big(\{u, v\}; \{w, x\}\big) \quad : \quad x \in S_1(v) \setminus \{w\} \Big\}$$

$$M_3 = \left\{ \bowtie\big(\{u, v\}; N\big) \to \bowtie\big(\{u, v\}; \{w\} \cup N\big) \quad : \quad \begin{array}{r} N \subseteq S_1(v) \setminus \{w\} \\ |N| \geq 2 \end{array} \right\}$$

$$M_4 = \left\{ \bowtie\big(T \cup \{v\}; N\big) \to \bowtie\big(T \cup \{u, v\}; N\big) \quad : \quad \begin{array}{r} T \subseteq S_2(v) \setminus \{u\} \\ N \subseteq \bigcap_{y \in T \cup \{v\}} S_1(y) \\ |T| \geq 1,\ |N| \geq 2 \end{array} \right\}$$

The union $M = \bigcup_i M_i$ is an acyclic matching on P whose set of critical elements is Q. The argument is very similar to that in the proof of Theorem 5.0.1 with w playing the role of a_m and u playing the role of b_n. $\square$

Chapter 7

Quadric complexes and asphericity

Definition 7.0.1. A *locally quadric complex* is a locally finite square complex X with immersed cells such that no reduced disk diagram in X has the form of Figure 7–1 and any immersed disk diagram of a form on the left-hand side of Figure 7–2 has a replacement on the right-hand side with the same boundary path. If, in addition, X is simply connected then X is *quadric*. A group G is *quadric* if it acts properly and cocompactly on a quadric complex.

For a full introduction to quadric complexes and groups see prior work of the present author [25].

A square complex X is *flag* if each square of X is bounded by an embedded 4-cycle and each embedded 4-cycle of X^1 bounds a unique square of X.

Proposition 7.0.2 ([25, Proposition 1.18])**.** *Let X be a connected square complex. Then X is quadric if and only if X is flag and every isometrically embedded cycle of X^1 has length* 4*.*

It follows from Proposition 7.0.2 that a quadric complex X is the 2-skeleton of the flag bisimplicial completion $\mathbb{X}(X^1)$. See Definition 6.0.1.

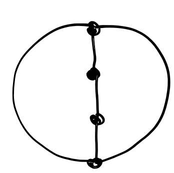

Figure 7–1: A disk diagram which can not be reduced in a locally quadric complex.

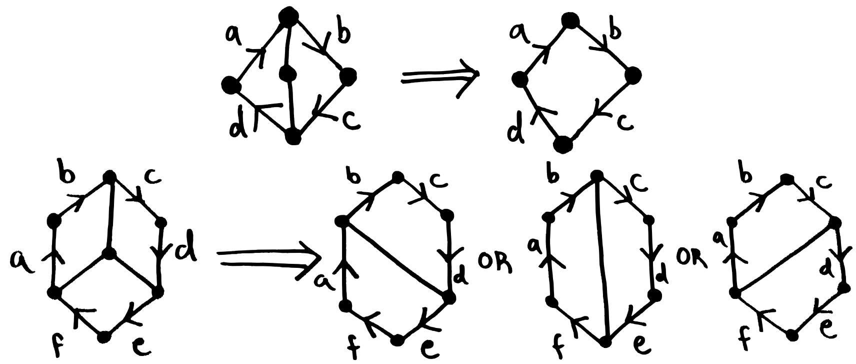

Figure 7–2: Replacement rules for disk diagrams in quadric complexes.

Theorem 7.0.3 (Bandelt [3, Theorem 1])**.** *A graph is hereditary modular if and only if it is connected and every isometrically embedded cycle has length* 4.

A graph is *modular* if for every triple of vertices u, v, w there exists a vertex x which lies on some geodesic between each pair of vertices in the triple. A graph is *hereditary modular* if each of its isometrically embedded subgraphs is modular.

The *metric ball* of radius $r \in \mathbb{N}$ centered at a vertex v of a graph (bisimplicial complex) is the induced (full) subgraph (subcomplex) on the set of vertices of distance at most r to v (in the 1-skeleton).

Remark 7.0.4. Let Γ be a modular graph. Then the metric balls of Γ are isometrically embedded. In particular, if Γ is hereditary modular then its metric balls are hereditary modular.

Theorem 7.0.5 (Bandelt [3, Theorem 2])**.** *Let Γ be a finite nonempty hereditary modular graph. Then Γ is bi-dismantlable.*

Theorem 7.0.6. *Let X be a nonempty quadric complex. Then the flag bisimplicial completion $\mathbb{X}(X^1)$ is contractible.*

Proof. The metric balls of X are quadric by Proposition 7.0.2, Theorem 7.0.3 and Remark 7.0.4. These balls are finite since quadric complexes are locally finite. Hence balls in X are collapsible by Proposition 7.0.2, Theorem 7.0.3, Theorem 7.0.5 and Theorem 6.0.2. The balls of X centered at a fixed vertex give an ascending exhaustion of X by contractible subcomplex and so X is contractible. $\square$

7.1 A $K(G,1)$ for torsion-free quadric groups

Let X be a locally quadric complex. Then the universal cover $\widetilde{X}$ is quadric and so $\pi_1 X$ is quadric. Let $\square_{m,n}$ denote the 2-skeleton of the bisimplex $\boxtimes^{m,n}$. Let $\square_{m,n} \to X$ be an immersion with $m, n \geq 2$. Since $\square_{m,n}$ is simply connected it lifts to $\widetilde{X}$. Since quadric complexes do not contain loops or bigons [25], every lift $\square_{m,n} \to \widetilde{X}$ is an embedding. Every torsion-free quadric group is the fundamental group of a compact locally quadric complex. However, a compact locally quadric complex may have a fundamental group with torsion. The following theorem allows to understand when this is the case.

Theorem 7.1.1 (Invariant Biclique Theorem [25])**.** *Let F be a finite group acting on a quadric complex $\widetilde{X}$. Then F stabilizes a nonempty connected complete bipartite subgraph of $\widetilde{X}$.*

To state the following corollary we need a definition. Let $\mathrm{Aut}(\square_{m,n})$ be the set of automorphisms of $\square_{m,n}$. Note that $\mathrm{Aut}(\square_{m,n})$ acts on the set of immersions $\{\square_{m,n} \to X\}$. For a given immersion $\square_{m,n} \to X$, we define $\mathrm{Aut}(\square_{m,n} \to X)$ as the stabilizer of $\square_{m,n} \to X$ in $\mathrm{Aut}(\square_{m,n})$.

Corollary 7.1.2. *Let X be a compact locally quadric complex. Then $\pi_1(X)$ has torsion if and only if $\mathrm{Aut}(\square_{m,n} \to X)$ is nontrivial for some immersion $\square_{m,n} \to X$ with $m, n \geq 1$.*

Proof. Let $g \in \pi_1(X) \setminus \{1\}$. Suppose g has finite order. Then $\langle g \rangle$ stabilizes a nonempty connected complete bipartite subgraph $K_{m+1,n+1}$ of $\widetilde{X}^1$. Since the action is free, we have $m, n \geq 1$. The full subcomplex induced by this $K_{m+1,n+1}$ is a $\square_{m,n}$. Restricting the covering map to $\square_{m,n}$ we have an immersion $\square_{m,n} \to$

X. Restricting the action of g to $\Box_{m,n}$ we obtain a nontrivial automorphism of $\Box_{m,n} \to X$.

Now, suppose there is a nontrivial automorphism $\varphi\colon \Box_{m,n} \to \Box_{m,n}$ of an immersion $\Box_{m,n} \to X$. Let $f\colon \Box_{m,n} \to \widetilde{X}$ be a lift of this immersion and identify $\Box_{m,n}$ with its image under f. Then φ extends to a nontrivial deck transformation which must have finite order. $\Box$

Let X be a compact locally quadric complex. We will construct a compact complex X^+ having X as its 2-skeleton such that X^+ is a $K\big(\pi_1(X), 1\big)$ if $\pi_1(X)$ is torsion-free. The construction is by induction on dimension. The 2-skeleton of X^+ is X. Now suppose we have already constructed the $(k-1)$-skeleton $(X^+)^{k-1}$ of X^+. To obtain the k-skeleton of X^+ perform the following operation. Along each immersion $\partial ⧖^{m,n} \to (X^+)^{k-1}$ with $\dim(⧖^{m,n}) = k$ glue in a copy of $⧖^{m,n}$. For the purposes of this operation, two immersions which are isomorphic over $(X^+)^{k-1}$ are considered identical and so result in only a single gluing. Since X is compact there is a bound on the size of a connected complete bipartite graph which can immerse in X. Hence X^+ is compact.

Lemma 7.1.3. *Let X be a compact locally quadric complex. Suppose $\pi_1(X)$ torsion-free. Then the universal cover $\widetilde{X^+}$ is isomorphic to the bisimplicial completion $⧖\big(\widetilde{X}^1\big)$.*

Proof. The proof is by induction on skeleta. The 2-skeleta of X^+ and $⧖(\widetilde{X}^1)$ are X and $\widetilde{X}$ so the base case holds. Assume the statement holds for the $(k-1)$-skeleta: $\big(\widetilde{X^+}\big)^{k-1} \cong ⧖\big(\widetilde{X}^1\big)^{k-1}$. Each $\partial ⧖^{m,n}$ subcomplex, $\dim(⧖^{m,n}) = k$, of $⧖\big(\widetilde{X}^1\big)^{k-1}$ immerses into $(X^+)^{k-1}$ under the covering map and so spans a $⧖^{m,n}$ in $\big(\widetilde{X^+}\big)^k$. On the other hand, each immersion $\partial ⧖^{m,n} \to (X^+)^{k-1}$ with $\dim(⧖^{m,n}) = k$ lifts to an embedding in $⧖\big(\widetilde{X}^1\big)^{k-1}$ whose image thus spans a unique $⧖^{m,n}$ in $⧖\big(\widetilde{X}^1\big)^k$. So the set of boundaries of the k-dimensional bisimplices of $\big(\widetilde{X^+}\big)^k$ and $⧖\big(\widetilde{X}^1\big)^k$. No two k-dimensional bisimplices have the same boundary in $\text{⧖}\big(\widetilde{X}^1\big)^k$. The same holds for $\big(\widetilde{X^+}\big)^k$ by Corollary 7.1.2 and so $\big(\widetilde{X^+}\big)^k \cong \text{⧖}\big(\widetilde{X}^1\big)^k$. $\square$

The main theorem of this section follows immediately from Lemma 7.1.3 and Theorem 7.0.6.

Theorem 7.1.4. *Let X be a compact locally quadric complex. If $\pi_1(X)$ is torsion-free then X^+ is a compact $K(\pi_1(X), 1)$.*

Part II

Shortcut Graphs and Groups

Chapter 8

Introduction

One of the main currents in geometric group theory is the study of groups acting on spaces that satisfy various kinds of combinatorial nonpositive curvature properties. These spaces are typically associated with graphs having nice metric properties. For example, the 1-skeleta of CAT(0) cube complexes [44, 15, 8], systolic complexes [8] and quadric complexes [25], all of which arose independently in the geometric group theory and metric graph theory literature [19, 28, 20, 25, 2, 34, 32, 49, 3]. In this part we introduce a metric property which captures an aspect of nonpositive curvature which is shared by these graphs and a surprisingly large number of other graphs of importance in group theory and metric graph theory. We call the graphs satisfying this property shortcut graphs and those satisfying a stronger property strongly shortcut graphs. A *shortcut graph* is a graph for which there is a bound on the lengths of its isometrically embedded cycles. A *strongly shortcut graph* may be defined as a graph for which there is a bound on the lengths of its K-bilipschitz cycles, for some $K > 1$. A group is *(strongly) shortcut* if it acts properly and cocompactly on a (strongly) shortcut graph.

An initial motivation for the study of shortcut graphs and groups arose from systolic and quadric complexes. Chepoi characterized systolic complexes as those flag simplicial complexes whose 1-skeleta contain isometrically embedded cycles only of length three [8]. We similarly characterized quadric complexes as those square-flag square complexes whose 1-skeleta contain isometrically embedded cycles only of length four [25]. Hence the 1-skeleta of systolic and quadric complexes are shortcut. In particular, systolic and quadric groups, and thus finitely presented C(6) and C(4)-T(4) small cancellation groups, are shortcut [51, 25]. As we will see, many more prominent classes of graphs and groups satisfy the shortcut property.

8.1 Summary of results

The following theorems summarize our results.

Theorem F (Corollary 11.1.2)**.** *Shortcut groups are finitely presented.*

Theorem G (Theorem 11.2.1 and Theorem 11.3.1)**.** *Shortcut graphs and groups have exponential isoperimetric and isodiametric functions. Strongly shortcut graphs and groups have polynomial isoperimetric and isodiametric functions. In particular, shortcut groups have decidable word problem.*

Theorem H (Theorem 12.1.1)**.** *Products of (strongly) shortcut graphs and groups are (strongly) shortcut.*

Theorem I (Corollary 12.2.5)**.** *Graphs of (strongly) shortcut groups with finite edge groups are (strongly) shortcut. In particular, amalgamated free products and HNN extensions of (strongly) shortcut groups over finite subgroups are (strongly) shortcut.*

Theorem J (Theorems 13.1.2, 13.3.1, 13.4.2 and 13.4.4 and Corollary 13.2.3)**.** *The following classes of graphs are strongly shortcut.*

- *Hyperbolic graphs*
- 1*-skeleta of finite dimensional* CAT(0) *cube complexes*
- *Standard Cayley graphs of finitely generated Coxeter groups*
- *All Cayley graphs of* $\mathbb{Z}$ *and* $\mathbb{Z}^2$

In particular, hyperbolic groups, cocompactly cubulated groups and Coxeter groups are all strongly shortcut.

Theorem K (Theorem 14.2.2 and Theorem 14.3.4)**.** *The Baumslag-Solitar group* $\mathrm{BS}(1,2)$ *is shortcut but not strongly shortcut. Moreover,* $\mathrm{BS}(1,2)$ *has a Cayley graph which is shortcut and a Cayley graph which is not shortcut.*

8.2 Structure of Part II

In Chapter 9 we present the main definitions of Part II. In Chapter 10 we prove basic properties of shortcut graphs and groups. In Chapter 11 we construct disk diagrams for shortcut graphs and study their filling invariants. In Chapter 12 we describe ways of combining shortcut graphs and groups to obtain new shortcut graphs and groups. In Chapter 13 we prove that various classes of graphs and groups are strongly shortcut. In Chapter 14 we study the shortcut property in Cayley graphs of $\mathrm{BS}(1,2)$.

8.3 Conventions

A *graph* Γ is a 1-dimensional polyhedral complex whose edges are isometric to $[0,1] \subset \mathbb{R}$. In this way each graph is equipped with both the structure of a cellular complex with edges and vertices and the structure of a geodesic metric giving a distance between any pair of its points. A *combinatorial map* of graphs is a continuous map $\Gamma_1 \to \Gamma_2$ in which each vertex of Γ_1 maps onto a vertex of Γ_2 and each closed edge of Γ_1 maps onto a vertex or closed edge of Γ_2. A combinatorial map is *degenerate* if some closed edge maps onto a vertex. Otherwise it is *nondegenerate*.

Chapter 9

Definitions

9.1 Isometric and almost isometric cycles

Let Γ be a graph. A *combinatorial path in* Γ is a nondegenerate combinatorial map $P \to \Gamma$ from a graph P that is homeomorphic to a compact interval of $\mathbb{R}$. A *cycle* C is a graph homeomorphic to a circle. A *cycle in* Γ is a nondegenerate combinatorial map $C \to \Gamma$ from a cycle C. The *length* of a path or cycle, denoted $|P|$ or $|C|$, is the number of its edges.

A cycle $f\colon C \to \Gamma$ is *isometric* if it is an isometric embedding. Corollary 10.0.2 below shows that f is isometric if and only if

$$d_\Gamma\big(f(p), f(q)\big) \geq \frac{|C|}{2}$$

for every antipodal pair of points $p, q \in C$. With this in mind we give the following definition. A cycle $f\colon C \to \Gamma$ is ξ-*almost isometric*, for $\xi \in (0,1)$, if

$$d_\Gamma\big(f(p), f(q)\big) \geq \xi\frac{|C|}{2}$$

for every antipodal pair of points $p, q \in C$. One may imagine that if f is not isometric then there is a "shortcut" in Γ between a pair of its antipodes and if f is not ξ-almost isometric then there is such a "shortcut" which reduces the distance by a constant factor.

9.2 Shortcut graphs and groups

A connected simplicial graph Γ is *shortcut* if, for some $\theta \in \mathbb{N}$, every isometric cycle of Γ has length at most θ. A connected simplicial graph Γ

is *strongly shortcut* if, for some $\theta \in \mathbb{N}$ and some $\xi \in (0,1)$, every ξ-almost isometric cycle of Γ has length at most θ.

Proposition 10.0.5 below says that Γ is strongly shortcut if and only if there is a $K > 1$ and a bound on the lengths of the K-bilipschitz cycles of Γ. Since a 1-bilischitz map is the same thing as an isometric embedding, we can thus define both properties together: Γ is shortcut if there is a $K \geq 1$ and a bound on the lengths of the K-bilipschitz cycles of Γ; if this K can be chosen strictly greater than 1 then Γ is strongly shortcut.

A group G is *(strongly) shortcut* if it acts properly and cocompactly on a (strongly) shortcut graph.

Chapter 10

Basic properties

In this section we prove some basic properties of (strongly) shortcut graphs and groups. In particular, we prove that the strong shortcut property is equivalent to the existence of a $K > 1$ and a bound on the lengths of the K-bilipschitz cycles in a graph. We also show that a (strongly) shortcut group acts freely and cocompactly on a (strongly) shortcut graph.

Proposition 10.0.1. *Let Γ be a graph and let $\bar{\xi} \in (0, 1]$. A cycle $f\colon C \to \Gamma$ satisfies*

$$d_\Gamma\big(f(p), f(q)\big) \geq \bar{\xi}\frac{|C|}{2}$$

for every antipodal pair of points $p, q \in C$ if and only if

$$d_\Gamma\big(f(p), f(q)\big) \geq d_C(p, q) - (1 - \bar{\xi})\frac{|C|}{2}$$

for every pair of points $p, q \in C$.

Proof. The "if" part follows by applying the inequality to each antipodal pair of p and q. To prove the "only if" part, let $p, q \in C$. Let p' be the antipode of p. Then, since f is 1-Lipschitz, we have

$$\begin{aligned} d_\Gamma\big(f(p), f(p')\big) &\leq d_\Gamma\big(f(p), f(q)\big) + d_\Gamma\big(f(q), f(p')\big) \\ &\leq d_\Gamma\big(f(p), f(q)\big) + d_C(q, p') \\ &= d_\Gamma\big(f(p), f(q)\big) + d_C(p, p') - d_C(p, q) \\ &= d_\Gamma\big(f(p), f(q)\big) + \frac{|C|}{2} - d_C(p, q) \end{aligned}$$

but $\bar{\xi}\frac{|C|}{2} \leq d_\Gamma\big(f(p), f(p')\big)$ and so we have $d_\Gamma\big(f(p), f(q)\big) \geq d_C(p,q) - (1 - \bar{\xi})\frac{|C|}{2}$. $\square$

Corollary 10.0.2. *Let Γ be a graph and let $f\colon C \to \Gamma$ be a cycle. Then f is isometric if and only if*

$$d_\Gamma\big(f(p), f(q)\big) \geq \frac{|C|}{2}$$

for every antipodal pair of points $p, q \in C$.

Proof. This follows from the fact that f is 1-Lipschitz and by applying Proposition 10.0.1 with $\bar{\xi} = 1$. $\square$

Proposition 10.0.3. *Let Γ be a graph and let $f\colon C \to \Gamma$ be a cycle in Γ of length $|C| \geq 4$. If f is not isometric then*

$$d_\Gamma\big(f(u), f(v)\big) < d_C(u,v)$$

for some pair of vertices $u, v \in C^0$ with $d_C(u,v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor - 1$. If f is not ξ-almost isometric, for some $\xi \in (0,1)$, then

$$d_\Gamma\big(f(u), f(v)\big) < \xi d_C(u,v)$$

for some pair of vertices $u, v \in C^0$ with $d_C(u,v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor - 1$.

Proof. Let $\bar{\xi} \in (0,1]$. Suppose

$$d_\Gamma\big(f(p), f(q)\big) < \bar{\xi}\frac{|C|}{2}$$

for some pair of antipodal points $p, q \in C$. If $\bar{\xi} < 1$ then this is equivalent to f not being $\bar{\xi}$-almost isometric and, otherwise, this is equivalent to f not being isometric.

Let $d = d_\Gamma\big(f(p), f(q)\big)$. Let $\alpha\colon [0,d] \to \Gamma$ be a geodesic from $f(p)$ to $f(q)$. If $\alpha^{-1}(\Gamma^0) = \emptyset$ then $f(p)$ and $f(q)$ are contained in the interior of some common edge e of Γ. Then the edges e_1 and e_2 of C with $p \in e_1$ and

$q \in e_2$ map onto e. Then there are endpoints $u \in e_1$ and $v \in e_2$ such that $f(u) = f(v)$ and $d_C(p,u) + d_C(q,v) \leq 1$. So we have $d_\Gamma\big(f(u), f(v)\big) = 0$ and $d_C(u,v) \geq d_C(p,q) - d_C(p,u) - d_C(q,v) \geq \frac{|C|}{2} - 1$.

Assume now that p and q do not map to the same edge in Γ. Let $\alpha^{-1}(\Gamma^0) = \{x_1, x_2, \ldots, x_k\}$ with $0 \leq x_1 < x_2 < \cdots < x_k \leq d$. Then $\alpha|_{[0,x_1]}$ and $\alpha|_{[x_k,d]}$ factor through f so their images contain vertices $u, v \in C^0$ with $d_C(p,u) < 1$ and $d_C(q,v) < 1$ such that $f(u) = \alpha(x_1)$ and $f(v) = \alpha(x_k)$. Moreover

$$\begin{aligned}
d_\Gamma\big(f(u), f(v)\big) &= x_k - x_1 \\
&= d - x_1 - (d - x_k) \\
&= d - d_C(p,u) - d_C(q,v) \\
&< \bar{\xi} d_C(p,q) - d_C(p,u) - d_C(q,v) \\
&\leq \bar{\xi}\big(d_C(p,q) - d_C(p,u) - d_C(q,v)\big) \\
&\leq \bar{\xi} d_C(u,v)
\end{aligned}$$

and $d_C(u,v) \geq d_C(p,q) - d_C(p,u) - d_C(q,v) > \frac{|C|}{2} - 2$. Since $|C|$ and $d_C(u,v)$ are integers we obtain $d_C(u,v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor - 1$ $\square$

Remark 10.0.4. In fact, if $f \colon C \to \Gamma$ is not isometric then we can improve the u and v obtained from Proposition 10.0.3 so that $d_C(u,v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor$. This does not hold for f not ξ-almost isometric. (Consider the quotient map from C oriented and of even length that identifies two antipodal edges of C in an orientation reversing way.) In order to present a unified proof without additional case analysis, we give the weaker statement in Proposition 10.0.3.

Proposition 10.0.5. *Let Γ be a graph. Then Γ is strongly shortcut if and only if there exists $K > 1$ such that there is a bound on the length of K-bilipschitz cycles of Γ.*

Proof. If a cycle $f\colon C \to \Gamma$ is not ξ-almost isometric then, for some pair of antipodal points $p, q \in C$,

$$\xi d_C(p,q) = \xi \frac{|C|}{2} > d_\Gamma\big(f(p), f(q)\big)$$

and so f is not $\frac{1}{\xi}$-bilipschitz. This proves the "only if" part of the proposition.

To prove the "if" part of the proposition, suppose θ bounds the length of the K-bilipschitz cycles of Γ where $K > 1$. Let $1 - \frac{(K-1)^3}{13K^2(K+1)} < \xi < 1$. We will show that there is a bound on the lengths of the ξ-almost isometric cycles of Γ. Let $f\colon C \to \Gamma$ be a ξ-almost isometric cycle of Γ. We will define a sequence of open paths (P_i), a sequence of cycles $(C_i)_i$, a sequence of finite graphs $(\Gamma_i)_i$ and sequences of combinatorial maps as in the following commuting diagram.

$$\begin{tikzcd}
 & P_0 \arrow[d, hook] \arrow[dr, hook] & P_1 \arrow[d, hook] \arrow[dr, hook] & \cdots \arrow[dr, hook] & P_{n-1} \arrow[d, hook] \arrow[dr, hook] & \\
C \arrow[r, equal] \arrow[d, equal] & C_0 \arrow[d, hook] \arrow[dr] & C_1 \arrow[d, hook] \arrow[dr] & \cdots \arrow[dr] & C_{n-1} \arrow[d, hook] \arrow[dr] & C_n \arrow[d, hook] \\
C \arrow[r, equal] \arrow[drr, bend right, "f"'] & \Gamma_0 \arrow[r] \arrow[dr, "f_0"] & \Gamma_1 \arrow[r] \arrow[d, "f_1"] & \cdots \arrow[r] & \Gamma_{n-1} \arrow[r] \arrow[dll, "f_{n-1}"] & \Gamma_n \arrow[dlll, bend left, "f_n"] \\
 & & \Gamma & & &
\end{tikzcd}$$

Where it makes sense, we will use the same notation to refer to points and subspaces as we do to refer to their images under maps. We begin with $\Gamma_0 = C_0 = C$ and $f_0 = f$. Suppose we inductively have $C_i \hookrightarrow \Gamma_i \xrightarrow{f_i} \Gamma$. If the composition of these maps is K-bilipschitz then we terminate the sequence with $n = i$. Otherwise, let $u_i, v_i \in C_i^0$ be a furthest pair of vertices in C_i for which $d_\Gamma\big(f_i(u_i), f_i(v_i)\big) < \frac{1}{K} d_{C_i}(u_i, v_i)$. Let Q_i be a geodesic segment of C_i between u_i and v_i. Let P_i be the complement of Q_i. Let $R_i \to \Gamma$ be a geodesic from $f_i(u_i)$ to $f_i(v_i)$. We obtain $f_{i+1}\colon \Gamma_{i+1} \to \Gamma$ from f_i and $R_i \to \Gamma$

by identifying the endpoints of R_i with $\{u_i, v_i\}$. Let $C_{i+1} = P_i \cup R_i$ in Γ_{i+1}. The sequence always terminates since $|C_i|$ is strictly decreasing.

Our goal is to show that $\frac{|C_n|}{|C|}$ is uniformly bounded away from zero. Thus we will show that if we have arbitrarily long ξ-almost isometric cycles then we must also have arbitrarily long K-bilipschitz cycles.

Suppose $Q_i \subset \bigcap_{\iota=0}^{i-1} P_\iota$ for all $i < j$. Then the Q_i, with $i < j$, are pairwise disjoint in C and $C \setminus \left(\bigcap_{i=0}^{j-1} P_i\right) = \bigcup_{i=0}^{j-1} Q_i$. Since C_j is obtained from C by replacing the Q_i with R_i we see then that the R_i, with $i < j$, are pairwise disjoint in C_j and the complement in C_j of $\bigcap_{i=0}^{j-1} P_i$ is $\bigcup_{i=0}^{j-1} R_i$. Since f is ξ-almost isometric we have

$$|Q_i| - (1-\xi)\frac{|C|}{2} \leq |R_i| < \frac{1}{K}|Q_i|$$

for all $i < j$, by Proposition 10.0.1. Hence, if $i < j$ then we have the following inequality.

$$K|R_i| < |Q_i| < (1-\xi)\left(\frac{K}{(K-1)}\right)\frac{|C|}{2} \tag{$*$}$$

Moreover, we can find a pair of points p, q in the closure of $\bigcap_{i=0}^{j-1} P_i = C \setminus \left(\bigcup_{i=1}^{j-1} Q_i\right)$ at distance $d_C(p,q) \geq \frac{|C|}{2} - \frac{K(1-\xi)}{4(K-1)}|C|$. Let S_1 and S_2 be the two segments of C between p and q. Then

$$\begin{aligned}
d_\Gamma\big(f(p), f(q)\big) &\leq |S_i| - \Big|S_i \cap \bigcup_{i<j} Q_i\Big| + \Big|S_i \cap \bigcup_{i<j} R_i\Big| \\
&\leq |S_i| - \Big|S_i \cap \bigcup_{i<j} Q_i\Big| + \frac{1}{K}\Big|S_i \cap \bigcup_{i<j} Q_i\Big|
\end{aligned}$$

and so, by Proposition 10.0.1,

$$
\begin{aligned}
&|C| - \frac{K-1}{K}\sum_{i<j}|Q_i| \\
&\geq 2d_\Gamma\big(f(p), f(q)\big) \\
&\geq 2d_C(p,q) - (1-\xi)|C| \\
&\geq |C| - \frac{K(1-\xi)}{2(K-1)}|C| - (1-\xi)|C|
\end{aligned}
$$

which gives us the following inequality.

$$\sum_{i<j}|Q_i| \leq (1-\xi)\left(\frac{K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2} \tag{†}$$

We will now prove that $Q_i \subset \bigcap_{\iota=0}^{i-1} P_\iota$ for $1 < i < n$. For the sake of finding a contradiction, suppose $j \geq 1$ is the least integer with $Q_j \not\subset \bigcap_{i=0}^{j-1} P_i$. It is possible that, for some $i < j$, we have $Q_j \cap R_i \neq \emptyset$ but $R_i \not\subset Q_j$. This may happen for at most two R_i since such R_i must contain an endpoint of Q_j. Let Q_j^- be obtained from Q_j by subtracting the interiors of any such R_i and let $Q_j^+ \to C_j$ extend $Q_j \hookrightarrow C_j$ so as to include full copies of any such R_i. Let $Q^- \subset C$ be obtained from $Q_j^- \subset C_j$ by replacing any $R_i \subset Q_j^-$, where $i < j$, with $Q_i \subset C$. Let $Q^+ \to C$ be obtained from $Q_j^+ \to C_j$ by replacing any $R_i \hookrightarrow C_j$, where $i < j$, with $Q_i \hookrightarrow C$. Let $R^+ \to \Gamma$ be obtained from $Q_j^+ \to \Gamma$ by replacing $Q_j \to \Gamma$ with $R_j \to \Gamma$. Then $R^+ \to \Gamma$ and the composition

$Q^+ \to C \xrightarrow{f} \Gamma$ have the same endpoints in Γ and we have

$$\begin{aligned}
|R^+| &= |R_j| + |Q_j^+ \setminus Q_j| \\
&< \frac{1}{K}|Q_j| + |Q_j^+ \setminus Q_j| \\
&= \frac{1}{K}\big(|Q_j^-| + |Q_j \setminus Q_j^-|\big) + |Q_j^+ \setminus Q_j| \\
&\leq \frac{1}{K}|Q_j^-| + |Q_j \setminus Q_j^-| + |Q_j^+ \setminus Q_j| \\
&= \frac{1}{K}|Q_j^-| + |Q_j^+ \setminus Q_j^-| \\
&\leq \frac{1}{K}|Q_j^-| + \frac{1}{K}|Q^+ \setminus Q^-| \\
&= \frac{1}{K}|Q^+|
\end{aligned}$$

where the final inequality follows from the fact that $Q_j^+ \setminus Q_j^-$ consists of up to two copies of segments R_i which are replaced with corresponding Q_i in $Q^+ \setminus Q^-$. By assumption, Q_j nontrivially intersects at least one R_i, with $i < j$. Let i be minimal such that Q_j nontrivially intersects R_i. Since Q_j is not equal to this R_i we see that $|Q^+| > |Q_i|$. Hence, if $Q^+ \to C_i$ were an isometric embedding then this would contradict the choice of u_i and v_i. So, $Q^+ \to C_i$ is not an isometric embedding and so $|Q^+| > \frac{|C_i|}{2}$. But then

$$|Q^+| > \frac{|C_i|}{2} \geq \frac{|C|}{2} - \sum_{k<i} |Q_k| \geq \frac{|C|}{2} - (1-\xi)\left(\frac{K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2}$$

by (†) while

$$|Q_0| < (1-\xi)\left(\frac{K}{(K-1)}\right)\frac{|C|}{2}$$

by $(*)$. So $|Q_+| < |Q_0|$ would imply

$$1 - (1-\xi)\left(\frac{K(3K-2)}{(K-1)^2}\right) < (1-\xi)\left(\frac{K}{(K-1)}\right)$$

which after some manipulation gives $\xi < 1 - \frac{(K-1)^2}{K(4K-3)}$ which one can show contradicts our choice of $\xi > 1 - \frac{(K-1)^3}{13K^2(K+1)}$. Hence $|Q_+| \geq |Q_0|$ so if $Q^+ \to C$

were an isometric embedding then this would contradict the choice of u_0 and v_0. So, $Q^+ \to C$ is not an isometric embedding and so $|Q^+| > \frac{|C|}{2}$. On the other hand

$$\begin{aligned}
|Q^+| &\le |Q^-| + 2\max_{i<j}|Q_i| \\
&\le |Q_j^-| + \sum_{i<j}\big(|Q_i| - |R_i|\big) + 2\max_{i<j}|Q_i| \\
&\le \frac{|C_j|}{2} + \sum_{i<j}\big(|Q_i| - |R_i|\big) + 2\max_{i<j}|Q_i| \\
&\le \frac{|C|}{2} + 2\sum_{i<j}\big(|Q_i| - |R_i|\big) + 2\max_{i<j}|Q_i| \\
&\le \frac{|C|}{2} + 4\sum_{i<j}|Q_i| \\
&\le \frac{|C|}{2} + (1-\xi)\left(\frac{4K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2}
\end{aligned}$$

where the last inequality follows from (†). We have

$$1-\xi < \frac{(K-1)^3}{13K^2(K+1)} < \frac{(K-1)^2}{12K(K+1)} = \frac{(K-1)^2}{4K(3K+3)} < \frac{(K-1)^2}{4K(3K-2)} \qquad (‡)$$

and so $|Q^+| < |C|$ so Q^+ embeds in C and the endpoints u, v of Q^+ in C are at distance

$$d_C(u,v) \ge \frac{|C|}{2} - (1-\xi)\left(\frac{4K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2}.$$

But we also have

$$d_\Gamma\big(f(u), f(v)\big) \le |R^+| \le \frac{1}{K}|Q^+| \le \frac{1}{K}\left(\frac{|C|}{2} + (1-\xi)\left(\frac{4K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2}\right)$$

which, by Proposition 10.0.1, implies

$$\begin{aligned}
&\frac{1}{K}\left(\frac{|C|}{2} + (1-\xi)\left(\frac{4K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2}\right) \\
&\ge \frac{|C|}{2} - (1-\xi)\left(\frac{4K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2} - (1-\xi)\frac{|C|}{2}
\end{aligned}$$

which is equivalent to $1-\xi \geq \frac{(K-1)^3}{K(13K^2+2K-7)}$. But

$$1-\xi < \frac{(K-1)^3}{13K^2(K+1)} = \frac{(K-1)^3}{K(13K^2+13K)} < \frac{(K-1)^3}{K(13K^2+2K-7)}$$

so we have a contradiction. Therefore we have proved that $Q_i \subset \bigcap_{\iota=0}^{i-1} P_\iota$ for $1 < i < n$.

Then the Q_i are all pairwise disjoint in C and $C \setminus \left(\bigcap_i P_i\right) = \bigcup_i Q_i$ so C_n is obtained from C by replacing $Q_i \subset C$ with R_i, for each i. Then since f_n is K-bilipschitz and by (†) and (‡), we have

$$\begin{aligned}\theta \geq |C_n| \geq |C| - \sum_i |Q_i| &\geq |C| - (1-\xi)\left(\frac{K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2} \\ &> |C| - \left(\frac{(K-1)^2}{4K(3K-2)}\right)\left(\frac{K(3K-2)}{(K-1)^2}\right)\frac{|C|}{2} \\ &= \frac{7}{8}|C|\end{aligned}$$

So $|C| < \frac{8}{7}\theta$ and we see that $\frac{8}{7}\theta$ bounds the lengths of ξ-almost isometric cycles of Γ. $\square$

Proposition 10.0.6. *Let Γ be a (strongly) shortcut graph. Then the graph obtained from Γ by subdividing each edge is (strongly) shortcut.*

Proof. Let Γ' be the barycentric subdivision of Γ. Then Γ' is isometric to Γ after scaling the metric by a factor of 2. Since isometric cycles are embedded, they have no backtracks and so every isometric cycle of Γ' is the subdivision of an isometric cycle of Γ. Hence, if θ bounds the lengths of the isometric cycles of Γ then 2θ bounds the lengths of the isometric cycles of Γ'. So if Γ is shortcut then Γ' is shortcut. Similary, if there is a bound on the K-bilipschitz cycles of Γ then there is a bound on the K-bilipschitz cycles of Γ'. So, by Proposition 10.0.5, if Γ is strongly shortcut then Γ' is strongly shortcut. $\square$

Proposition 10.0.7. *Let G be a (strongly) shortcut group. Then G acts freely and cocompactly on a (strongly) shortcut graph.*

Proof. Let G act properly and cocompactly on a (strongly) shortcut graph Γ. If Γ has a single vertex then G is finite and so acts freely on any Cayley graph of G which is strongly shortcut because it is connected and finite. So we may assume that Γ has more than one vertex.

By Proposition 10.0.6, we may assume that G acts on Γ without edge inversions. Let $\pi\colon G \times \Gamma^0 \to \Gamma^0$ be the projection onto the second factor. Define a graph $\tilde{\Gamma}$ on the vertex set $G \times \Gamma^0$ where $\tilde{\Gamma}$ has an edge joining (g, v) and (g', v') for each edge joining v and v'. Then the diagonal action $G \curvearrowright G \times \Gamma^0$ extends to $\tilde{\Gamma}$ and the projection π extends to a G-equivariant nondegenerate combinatorial map $\pi\colon \tilde{\Gamma} \to \Gamma$. That G acts on Γ without edge inversions rules out nontrivial fixed points of midpoints of edges of $\tilde{\Gamma}$ and so the action of G on $\tilde{\Gamma}$ is free. Let $\{v_1, v_2, \ldots, v_k\}$ be a set of orbit representatives of $G \curvearrowright \Gamma^0$. Let $\hat{\Gamma}$ be the induced subgraph of $\tilde{\Gamma}$ on $\bigcup_{g \in G} \big\{(g, gv_i)\big\}_i$. We will prove that $\hat{\Gamma}$ is (strongly) shortcut and that the action of G on $\hat{\Gamma}$ is cocompact.

Let $\hat{\pi}\colon \hat{\Gamma} \to \Gamma$ be the restriction of π to $\hat{\Gamma}$. For a vertex $v \in \Gamma^0$, we have $v = gv_i$ for some i and so $\hat{\pi}(g, gv_i) = v$. Moreover, since $\hat{\Gamma}$ is an induced subgraph of $\tilde{\Gamma}$, for each pair of vertices $(g, u), (h, v) \in \hat{\Gamma}^0$, we see that $\hat{\pi}$ induces a bijection between the set of edges between (g, u) and (h, v) and the set of edges between u and v. This implies that for any $(g, u), (h, v) \in \hat{\Gamma}^0$, we can lift any path of nonzero length $\alpha\colon P \to \Gamma$ between u and v to a path $\hat{\alpha}\colon P \to \hat{\Gamma}$ from (g, u) to (h, v). The lift is not unique since, if the sequence of vertices visited by α is $(u = u_0, u_1, u_2, \ldots, u_k = v)$ then, for $0 < i < k$, the lift of u_i in $\hat{\alpha}$ may be any $(g, u_i) \in \hat{\pi}^{-1}(u_i)$.

If Γ is strongly shortcut then let $\theta \geq 3$ bound the lengths of the $\bar{\xi}$-almost isometric cycles of Γ. Otherwise, let $\theta \geq 3$ bound the lengths of the isometric cycles of Γ and set $\bar{\xi} = 1$. Let $\hat{\xi} = \frac{1+\bar{\xi}}{2}$ and let $f\colon C \to \hat{\Gamma}$ be a ($\hat{\xi}$-almost) isometric a cycle of length $|C| > \theta$. By Proposition 10.0.3,

$$d_\Gamma\big(\hat{\pi} \circ f(u), \hat{\pi} \circ f(v)\big) < \bar{\xi} d_C(u,v)$$

for some $u, v \in C^0$ with $d_C(u,v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor - 1$. If $d_\Gamma\big(\hat{\pi} \circ f(u), \hat{\pi} \circ f(v)\big) > 0$ then let $\alpha\colon P \to \Gamma$ be a geodesic from $\hat{\pi} \circ f(u)$ to $\hat{\pi} \circ f(v)$. Otherwise, let $\alpha\colon P \to \Gamma$ be a path of length 2 from $\hat{\pi} \circ f(u)$ to $\hat{\pi} \circ f(v) = \hat{\pi} \circ f(u)$. This is always possble since Γ is a connected graph on more than one vertex. By the previous paragraph, we may lift α to a path $\hat{\alpha}\colon P \to \hat{\Gamma}$ from $f(u)$ to $f(v)$. So we see that either

$$d_{\hat{\Gamma}}\big(f(u), f(v)\big) < \bar{\xi} d_C(u,v)$$

or

$$d_{\hat{\Gamma}}\big(f(u), f(v)\big) \leq 2$$

and so, by Proposition 10.0.1, one of

$$d_C(u,v) - (1 - \hat{\xi})\frac{|C|}{2} < \bar{\xi} d_C(u,v) \tag{§}$$

or

$$d_C(u,v) - (1 - \hat{\xi})\frac{|C|}{2} \leq 2 \tag{¶}$$

must hold. Since $d_C(u,v) \geq \frac{|C|}{2} - \frac{3}{2}$ we see that (¶) gives the bound $|C| \leq \frac{7}{\hat{\xi}}$. On the other hand, (§) is equivalent to

$$(1 - \bar{\xi}) d_C(u,v) < (1 - \hat{\xi})\frac{|C|}{2}$$

and so gives

$$(\hat{\xi} - \bar{\xi})\frac{|C|}{2} < (1 - \bar{\xi})\frac{3}{2}$$

which is impossible if $\bar{\xi} = 1$ and otherwise gives the bound $|C| \leq \frac{3(1-\bar{\xi})}{\hat{\xi}-\bar{\xi}}$. $\square$

Chapter 11

Filling properties and disk diagrams

In this section we study disk diagrams and the isoperimetric and isodiametric functions of (strongly) shortcut graphs and groups. Let Γ be a graph and let $\theta \in \mathbb{N}$. For the purposes of the current discussion, a cycle C is always based and oriented. Hence two cycles $f_1, f_2 \colon C \to \Gamma$ may be distinct even if $f_1 = f_2 \circ \psi$ for some $\psi \in \mathrm{Aut}(C)$. Let

$$S_\theta = \big\{ f \colon C \to \Gamma : |C| \leq \theta \big\}$$

be the set of cycles in Γ. The *θ-filling* $F_\theta(\Gamma)$ is the 2-complex whose 1-skeleton is Γ and whose 2-skeleton has a unique 2-cell with attaching map $f \colon C \to \Gamma$ for each $f \in S_\theta$. If a group G acts on Γ then G acts on S_θ by $g \cdot f = \varphi_g \circ f$ where $\varphi_g \in \mathrm{Aut}(\Gamma)$ is the automorphism by which g acts on Γ. Thus the action of G on Γ extends to an action on $F_\theta(\Gamma)$ such that an element $g \in G$ stabilizes a 2-cell F if and only if g stabilizes F pointwise.

For $\theta, N \in \mathbb{N}$ with $3 \leq \theta \leq N$ and $\xi \in (0,1)$ consider the following property.

$$\text{Every cycle } C \to \Gamma \text{ with } \theta < |C| \leq N \text{ is not } \xi\text{-almost isometric.} \qquad (*)$$

Remark 11.0.1. If Γ is shortcut then Γ satisfies $(*)$ for θ bounding the lengths of isometric cycles of Γ, for any $N \geq \theta$ and for $\xi \in \big(\frac{N-2}{N}, 1\big)$. Of course, if Γ is strongly shortcut, then it satisfies $(*)$ for a fixed ξ not depending on N.

Construction 11.0.2. Let Γ be a graph satisfying $(*)$. Given a cycle $f\colon C \to \Gamma$ of length $|C| \leq N$ we will inductively construct a disk diagram $D_f \to F_\theta(\Gamma)$ for f. If $|C| \leq \theta$ then $D_f \to F_\theta(\Gamma)$ is just a single 2-cell mapping to the 2-cell of $F_\theta(\Gamma)$ whose attaching map is isomorphic to f. Otherwise f is not ξ-almost isometric and so, by Proposition 10.0.3,

$$d_\Gamma\big(f(u), f(v)\big) < \xi d_C(u, v)$$

for some pair of vertices $u, v \in C^0$ with $d_C(u, v) \geq \left\lfloor \frac{|C|}{2} \right\rfloor - 1$. Let P and Q be the two segments of C joining u and v. Let $g\colon R \to \Gamma$ be a geodesic path from $f(u)$ to $f(v)$ in Γ and note that $|R| < \xi \min\big\{|P|, |Q|\big\}$. Glue f and g together along $u \sim g^{-1}\big(f(u)\big)$ and $v \sim g^{-1}\big(f(v)\big)$ to obtain a combinatorial map $h\colon (C \sqcup R)/{\sim} \to \Gamma$ with $h|_{P\cup R}$ and $h|_{Q\cup R}$ cycles of length

$$|P| + |R| < |P| + \xi|Q| < |C|$$

and

$$|Q| + |R| < |Q| + \xi|P| < |C|$$

and so, by induction we have disc diagrams $D_{h|_{P\cup R}} \to F_\theta(\Gamma)$ and $D_{h|_{Q\cup R}} \to F_\theta(\Gamma)$ for $h|_{P\cup R}$ and $h|_{Q\cup R}$. Gluing $D_{h|_{P\cup R}} \to F_\theta(\Gamma)$ and $D_{h|_{Q\cup R}} \to F_\theta(\Gamma)$ together along R we obtain a disk diagram $D_f \to F_\theta(\Gamma)$ for f.

11.1 Simple connectedness

Theorem 11.1.1. *Let Γ be a shortcut graph and let $\theta \geq 3$ bound the lengths of the isometric cycles of Γ. Then $F_\theta(\Gamma)$ is simply connected.*

Proof. Let $f\colon C \to \Gamma$ be a cycle in Γ. Then, by Remark 11.0.1 Γ satisfies $(*)$ for θ as given, for $N = |C|$ and for $\xi = \frac{N-1}{N}$. Hence, we may apply Construction 11.0.2 to obtain a disk diagram for f. ☐

Corollary 11.1.2. *Let G be a (strongly) shortcut group. Then there is a compact 2-complex X with $G = \pi_1(X)$ such that the universal cover $\widetilde{X}$ of X has (strongly) shortcut 1-skeleton. In particular, we have that G is finitely presented.*

Proof. It suffices to consider the case where G is shortcut since every strongly shortcut group is shortcut. By Proposition 10.0.7, there is a free and cocompact action of G on a shortcut graph Γ. Let $\theta \geq 3$ bound the lengths of the isometric cycles of Γ. Then G acts freely and cocompactly on the filling $F_\theta(\Gamma)$. $\square$

11.2 Isoperimetric function

Theorem 11.2.1. *Let Γ be a graph. If Γ is shortcut then, for θ large enough, the filling $F_\theta(\Gamma)$ has an exponential isoperimetric function. If Γ is strongly shortcut then, for θ large enough, the filling $F_\theta(\Gamma)$ has a polynomial isoperimetric function.*

Proof. Suppose Γ is shortcut and let $\theta \geq 3$ bound the lengths of the isometric cycles of Γ. Let $\Delta\colon \mathbb{N} \to \mathbb{N}$ be the Dehn function of $F_\theta(\Gamma)$. We will prove, by induction on n that $\Delta(n) \leq 2^n$. If $n \leq \theta$ then this clearly holds since any cycle of length at most θ bounds a 2-cell in $F_\theta(\Gamma)$. Let $f\colon C \to \Gamma$ be a cycle of length $n > \theta$. Applying Construction 11.0.2 to f with $N = n$ and $\xi = \frac{n-1}{n}$ we see that f bounds a disk diagram D_f which is the union of two disk diagrams of boundary length less than n. Hence f bounds a disk of area at most $2\Delta(n-1)$. By induction

$$2\Delta(n-1) \leq 2 \cdot 2^{n-1} = 2^n$$

and so we have $\Delta(n) \leq 2^n$.

Suppose Γ is strongly shortcut. Choose $L \in \mathbb{N}$ with $L > 3$. Let $\theta \geq \frac{L}{1-\xi}$ bound the lengths of the ξ-almost isometric cycles of Γ. We will prove that the Dehn function of $F_\theta(\Gamma)$ satisfies $\Delta(n) \leq n^{\log_b(2)}$ for $b = \frac{2L}{(L-3)\xi+L+3}$. Note that $b > 1$ and that b tends to $\frac{2}{1+\xi}$ as L goes to infinity. The argument proceeds as in the shortcut case but in the inductive step f bounds a disk diagram which is the union of two disk diagrams of boundary length strictly less than

$$\xi\frac{n}{2} + \left\lceil\frac{n}{2}\right\rceil + 1 \leq \xi\frac{n}{2} + \frac{n}{2} + \frac{1}{2} + 1 = \frac{1}{2}\Big(\xi + 1 + \frac{3}{n}\Big)n \leq \frac{1}{2}\Big(\xi + 1 + \frac{3}{\theta}\Big)n$$

so, since $\theta \geq \frac{L}{1-\xi}$ we have a disk diagram for f of area at most

$$2\Delta\Big(\Big\lfloor\frac{1}{2}\Big(\xi+1+\frac{3(1-\xi)}{L}\Big)n\Big\rfloor\Big) = 2\Delta\Big(\Big\lfloor\frac{1}{2L}\Big((L-3)\xi+L+3\Big)n\Big\rfloor\Big) = 2\Delta\Big(\Big\lfloor\frac{1}{b}n\Big\rfloor\Big)$$

and so by induction we have

$$2\Delta\Big(\Big\lfloor\frac{1}{b}n\Big\rfloor\Big) \leq 2\Big(\Big\lfloor\frac{1}{b}n\Big\rfloor\Big)^{\log_b(2)} \leq 2\Big(\frac{1}{b}n\Big)^{\log_b(2)} = 2\Big(\frac{1}{2}n^{\log_b(2)}\Big) = n^{\log_b(2)}$$

and so we have that $\Delta(n) \leq n^{\log_b(2)}$. $\square$

Corollary 11.2.2. *Let G be a group. If G is shortcut then it has an exponential isoperimetric function. If G is strongly shortcut then it has a polynomial isoperimetric function.*

Corollary 11.2.3. *Let G be a shortcut group. Then G has a decidable word problem.*

11.3 Isodiametric function

Theorem 11.3.1. *Let Γ be a graph. If Γ is shortcut then, for θ large enough, the filling $F_\theta(\Gamma)$ has an exponential isodiametric function. If Γ is strongly shortcut then, for θ large enough, the filling $F_\theta(\Gamma)$ has a polynomial isodiametric function.*

Proof. For a cycle $f\colon C \to \Gamma$ let $\operatorname{diam}(f)$ denote the minimum diameter of a disk diagram for f. Observe that in Construction 11.0.2 $\operatorname{diam}(f) \leq$

$\mathrm{diam}(h|_{P\cup R}) + \mathrm{diam}(h|_{Q\cup R})$. Indeed we may glue together minimal diameter disk diagrams of $h|_{P\cup R}$ and $h|_{Q\cup R}$ along R to obtain a disk diagram for f. Using this observation, the proof follows virtually identically to that of Theorem 11.2.1. $\square$

Corollary 11.3.2. *Let G be a group. If G is shortcut then it has an exponential isodiametric function. If G is strongly shortcut then it has a polynomial isodiametric function.*

Chapter 12

Combinations

In this section we show that (strongly) shortcut graphs and groups are closed under products and that a finite graph of (strongly) shortcut groups with finite edge groups is (strongly) shortcut.

12.1 Products

Let Γ_1 and Γ_2 be simplicial graphs. The *product graph* $\Gamma_1 \times \Gamma_2$ of Γ_1 and Γ_2 is the 1-skeleton of the CW complex product of Γ_1 and Γ_2. The vertex set of $\Gamma_1 \times \Gamma_2$ is $\Gamma_1^0 \times \Gamma_2^0$ and the edges of $\Gamma_1 \times \Gamma_2$ are given by $(u_1, u_2) \sim (v_1, v_2)$ whenever

$$u_1 = v_1 \text{ and } u_2 \sim v_2$$

or

$$u_1 \sim v_1 \text{ and } u_2 = v_2$$

where $\sim$ is the edge relation.

Theorem 12.1.1. *Let Γ_1 and Γ_2 be (strongly) shortcut graphs. Then $\Gamma_1 \times \Gamma_2$ is (strongly) shortcut.*

Proof. Let Γ_1 and Γ_2 be (strongly) shortcut and let θ bound the lengths of the (ξ-almost) isometric cycles of the Γ_i. Let $\Gamma = \Gamma_1 \times \Gamma_2$ and let $f\colon C \to \Gamma$ be a cycle of length $|C| \geq 2\theta$. We combine the shortcut and strongly shortcut cases as follows. If the Γ_i are strongly shortcut then we have θ and ξ as given. Otherwise, by Remark 11.0.1, the Γ_i satisfy $(*)$ for θ as given, for $N = |C|$ and for some ξ depending on N. We will show that for some antipodal pair of points $p, q \in C$ we have $d_\Gamma\big(f(p), f(q)\big) < \big(\frac{1+\xi}{2}\big)\frac{|C|}{2}$.

Each edge of C projects nondegenerately onto exactly one of Γ_1 or Γ_2. Call those edges that project nondegenerately onto Γ_1 *horizontal edges* and those that project nondegenerately onto Γ_2 *vertical edges*. Without loss of generality there are more horizontal edges than vertical edges. Let $f_1\colon C_1 \to \Gamma_1$ be the cycle obtained from $C \to \Gamma_1$ by contracting the vertical edges of C. Then $N \geq |C_1| \geq \frac{|C|}{2} \geq \theta$ so we have $d_{\Gamma_1}\big(f_1(p_1), f_1(q_1)\big) < \xi d_{C_1}(p_1, q_1)$ for some antipodal pair of points $p_1, q_1 \in C_1$. Let $p', q' \in C$ map to p_1 and q_1 under the contraction map $C \to C_1$. We may choose p' and q' so that they are not contained in the interior of any vertical edge. Let ℓ be the number of vertical edges in a geodesic segment of C between p' and q'. Then $d_C(p', q') = \frac{|C_1|}{2} + \ell$ while $d_\Gamma\big(f(p'), f(q')\big) < \xi\frac{|C_1|}{2} + \ell$. Let $P \subset C$ be a geodesic segment of length $\frac{|C|}{2}$ containing p' and q' and let p and q be the endpoints of P with p nearest to p' and q nearest to q'. Then $d_C(p, q) = \frac{|C|}{2}$ while

$$\begin{aligned}
d_\Gamma\big(f(p), f(q)\big) &\leq d_\Gamma\big(f(p), f(p')\big) + d_\Gamma\big(f(p'), f(q')\big) + d_\Gamma\big(f(q'), f(q)\big) \\
&\leq d_C(p, p') + d_\Gamma\big(f(p'), f(q')\big) + d_C(q', q) \\
&= d_C(p, q) - d_C(p', q') + d_\Gamma\big(f(p'), f(q')\big) \\
&< \frac{|C|}{2} - \Big(\frac{|C_1|}{2} + \ell\Big) + \xi\frac{|C_1|}{2} + \ell \\
&= \frac{|C|}{2} - (1 - \xi)\frac{|C_1|}{2} \\
&= \Big(\frac{1 + \xi}{2}\Big)\frac{|C|}{2}
\end{aligned}$$

and so $d_\Gamma\big(f(p), f(q)\big) < \big(\frac{1+\xi}{2}\big)\frac{|C|}{2}$. $\square$

Corollary 12.1.2. *Let G_1 and G_2 be (strongly) shortcut groups. Then $G_1 \times G_2$ is (strongly) shortcut.*

12.2 Trees of shortcut graphs

Let T be a tree. An *arc decomposition* of T is a partition of the edges of T into paths, called *arcs*, whose interior vertices all have degree two. The endpoints of the arcs of an arc decomposition are called *nodes*. Every tree comes equipped with a default arc decomposition whose arcs are simply its edges. A *tree of graphs with discrete edge graphs* is a surjective (possibly degenerate) combinatorial map $\Gamma \to T$ from a graph Γ to a tree T with an arc decomposition such that the preimage of each node v is a connected subgraph and the preimage of the interior of each arc P is a disjoint union of open paths of the same length. The closure of a component $\tilde{P}$ of this disjoint union is called a *lift* of P. We call the preimage of a node v the *vertex graph* Γ_v at v of $\Gamma \to T$. We call the preimage of the midpoint of an arc P of T the *edge graph* Γ_P at P of $\Gamma \to T$. For each endnode v of an arc P of T there is a function $\Gamma_{v,P} \colon \Gamma_P \to \Gamma_v$ sending $p \in \Gamma_e$ to the unique vertex contained in $\tilde{P} \cap \Gamma_v$ where $\tilde{P}$ is the lift of P that contains p. The $\Gamma_{v,P}$ are called the *attaching maps* of $\Gamma \to T$.

For the purpose of discussing trees of graphs, it is convenient to consider combinatorial maps $C \to \Gamma$, with C homeomorphic to S^1, which are not necessarily nondegenerate. Call such a map a *cycles with degeneracies*.

The following lemma is a variant of Jordan's separator theorem for trees [30].

Lemma 12.2.1. *Let T be a tree. Let $f \colon C \to T$ be a cycle with degeneracies of length $|C| \geq 3$. Then for some vertex $w \in f(C^0)$, the metric subspace $f^{-1}(w)$ has diameter at least $\frac{|C|}{3}$.*

Proof. For w in the image of f, consider the metric subspace $f^{-1}(w) \subset C$. Choose w in the image of f such that $f^{-1}(w)$ has the largest possible diameter. Let the vertices $u, v \in f^{-1}(w)$ achieve the diameter of $f^{-1}(w)$. Suppose, for the sake of finding a contradiction, that $d_C(u,v) < \frac{|C|}{3}$. Then the segment P of length $|P| \geq \frac{2|C|}{3}$ between u and v in C intersects $f^{-1}(w)$ only at u and v. Indeed any point $p \in f^{-1}$ must be at distance less than $\frac{|C|}{3}$ to both u and

v and there is no such point of P. But then the first and last edges of P map nondegenerately to some common edge ww' of T. But then $f^{-1}(w')$ has diameter larger than the diameter of $f^{-1}(w)$ contradicting our choice of w. Hence $d_C(u,v) \geq \frac{|C|}{3}$. □

Lemma 12.2.2. *Let T be a tree. Let $f\colon C \to T$ be a cycle with degeneracies of length $|C| \geq 3$. Let $L < \frac{|C|}{3} + 1$ and suppose that for each edge $e \subset f(C)$, the distance between the midpoints of any two edges of $f^{-1}(e)$ is at most L. Then for some vertex $w \in T^0$, any segment $P \subset C$ whose interior is disjoint from $f^{-1}(w)$ has length $|P| \leq L+1$.*

Proof. Let w be as in Lemma 12.2.1 and let P be the closure of a component of $C \setminus f^{-1}(w)$. We need to show that $|P| \leq L+1$. The initial and terminal edges e_1 and e_2 of P map to the same edge of T and so either $|P| \leq L+1$ or $|P| \geq |C| - L + 1$. But $|C| - L + 1 > \frac{2|C|}{3}$ while, by our choice of w, we have $|P| \leq \frac{2|C|}{3}$. Hence $|P| \leq L+1$ as required. □

A cycle with degeneracies is ξ-*almost isometric* if

$$d_\Gamma\big(f(p), f(q)\big) \geq \xi\frac{|C|}{2}$$

for any antipodal pair of points $p, q \in C$.

Lemma 12.2.3. *Let Γ be strongly shortcut with θ bounding the lengths of the ξ-almost isometric cycles of Γ. Then there exist $\xi' \in (0,1)$ and $\theta' \in \mathbb{N}$ depending only on ξ and θ such that θ' bounds the lengths of the ξ'-almost isometric cycles with degeneracies of Γ.*

Proof. Let θ bound the lengths of the ξ-almost isometrically embedded cycles of Γ. Let $f\colon C \to \Gamma$ be a cycle with degeneracies. Define $S \subset C$ as the union of all edges of C that map to vertices under f. We may assume that $S \neq C$ since, otherwise f is the constant map and so satisfies the statement of the theorem for any ξ'. Let $(P_i)_{i=1}^{\ell}$ be the sequence of components of S in the order they are visited in some traversal of C. We begin by showing, for $\theta'' \geq \max\left\{\theta, \frac{8}{1-\xi}\right\}$ and $\xi'' = \frac{1+\xi}{2}$, that if $|C| > \theta''$ and $|P_i|$ is even for each i then there exist antipodal $p, q \in C$ such that $d_\Gamma\big(f(p), f(q)\big) < \xi'' \frac{|C|}{2}$. Call the condition that the $|P_i|$ are even the *parity condition.* Later we will use this result to prove the statement of the theorem, for $\theta' \geq \max\left\{\theta'', \frac{2(2+\xi'')}{1-\xi''}\right\}$ and $\xi' = \frac{1+\xi''}{2}$, with no assumption on the parities of the $|P_i|$.

Assume f satisfies the parity condition and $|C| > \theta''$. We obtain from f a cycle without degeneracies $f' \colon C \to \Gamma$ by mapping each edge of P_i onto $f(e_i)$ where e_i is the edge that follows P_i in some fixed orientation of C. This is possible since f satisfies the parity condition. Thus f' folds P_i onto $f(e_i)$ in a zig-zag fashion. Then for any point $p \in C$, we have $d_\Gamma\big(f(p), f'(p)\big) \leq 1$. If $|C| > \theta''$ then there is a pair of antipodal points $p, q \in C$ such that $d_\Gamma\big(f'(p), f'(q)\big) \leq \xi \frac{|C|}{2}$. Hence

$$d_\Gamma\big(f(p), f(q)\big) < \xi \frac{|C|}{2} + 2 = \left(\xi + \frac{4}{|C|}\right)\frac{|C|}{2} \leq \left(\xi + \frac{4}{\theta''}\right)\frac{|C|}{2} \leq \xi'' \frac{|C|}{2}$$

as required.

We now consider general $f \colon C \to \Gamma$ which does not necessarily satisfy the parity condition. Assume that $|C| > \theta'$. Let $i_0 < i_1 < \cdots < i_{m-1}$ be the set of indices for which $|P_i|$ is not even and assume $|C| > \theta'$. Obtain a cycle $f' \colon C' \to \Gamma$ from f by contracting an edge in each P_{i_j} with j odd and expanding a vertex v to an edge $e \mapsto f(v)$ in each P_{i_j} with j even. Then f' satisfies the parity condition and we have $|C| \leq |C'| \leq |C| + 1$. There is a relation $R \subset C \times C'$ with pRq if and only if one of the following holds.

1. p was obtained directly from q
2. p is contained in an edge that was contracted to q
3. p is a vertex which was expanded to an edge that contains q

By the alternating nature of the expansions and contractions we see that if pRp' and qRq' then $d_C(p,q) \geq d_{C'}(p',q') - 1$. By the previous paragraph, we have a pair of antipodal points $p', q' \in C'$ such that $d_\Gamma\big(f'(p'), f'(q')\big) < \xi'' \frac{|C'|}{2}$. Take any $p'', q'' \in C$ satisfying $p''Rp'$ and $q''Rq'$. Then $f(p'') = f'(p')$ and $f(q'') = f'(q')$ and so $d_\Gamma\big(f(p''), f(q'')\big) < \xi'' \frac{|C'|}{2} \leq \xi'' \frac{|C|+1}{2}$ and yet $d_C(p'', q'') \geq \frac{|C|}{2} - 1$. Hence, as f is 1-Lipschitz, for some antipodal pair of points $p, q \in C$ we have

$$\begin{aligned} d_\Gamma\big(f(p), f(q)\big) &< \xi'' \frac{|C|+1}{2} + 1 = \left(\xi'' + \frac{\xi''+2}{|C|}\right)\frac{|C|}{2} \\ &< \left(\xi'' + \frac{\xi''+2}{\theta'}\right)\frac{|C|}{2} \\ &\leq \xi' \frac{|C|}{2} \end{aligned}$$

as required. □

Theorem 12.2.4. *Let $\varphi\colon \Gamma \to T$ be a tree of graphs with discrete edge graphs satisfying the following two conditions.*

1. *The vertex graphs Γ_v are uniformly (strongly) shortcut in the sense that there exists $\theta \geq 3$ (and $\xi \in (0,1)$) such that θ bounds the lengths of the (ξ-almost) isometric cycles of every vertex graph.*
2. *For some $M \in \mathbb{N}$, for every attaching map $\Gamma_{v,P}$ of φ, the diameter of $\Gamma_{v,P}(\Gamma_P)$ is at most M.*
3. *Every arc of T has length M.*

Then Γ is (strongly) shortcut.

Proof. We will first consider the case where the vertex graphs are shortcut. Let $f\colon C \to \Gamma$ be an isometric cycle. If f maps entirely into a single vertex graph then $|C| \leq \theta$ by hypothesis. So, suppose the image of f contains some edge in the lift $\tilde{P}$ of an arc P of T. Then, since f is injective, it must traverse all of $\tilde{P}$ and, by consideration of $\varphi \circ f$, it must also traverse some other lift

$\hat{P}$ of P in the opposite direction. Let Q and Q' be the segments of C which map isomorphically to $\tilde{P}$ and $\hat{P}$. Let $u \in Q$ and $u' \in Q'$ be endpoints of Q and Q' mapping to the same vertex graph. Then $d_\Gamma\big(f(u), f(u')\big) \leq M$ and so $d_C(u, u') \leq M$ whereas Q and Q' each have length M. Hence the geodesic segment R of C between u and u' is disjoint from the interiors of Q and Q'. The same goes for the geodesic segment R' between the other pair of endpoints of Q and Q'. Then C is covered by the segments Q, Q', R and R', each of which has length at most M. Hence the lengths of the isometric cycles of Γ are bounded by $\max\{\theta, 4M\}$.

The case where the vertex graphs are strongly shortcut requires a more delicate argument relying on the preceeding lemmas. By Lemma 12.2.3 we can replace θ and ξ so that θ bounds the lengths of the ξ-almost isometric cycles with degeneracies of all the vertex graphs. Let $\xi' = \frac{\xi+2}{3}$ and let $\theta' = \max\big\{\theta, \frac{18M+6}{1-\xi}\big\}$. Let $f\colon C \to \Gamma$ be a ξ'-almost isometric cycle. We will prove that $|C| \leq \theta'$. Since $\frac{18M}{1-\xi} \geq 4$ we may assume that $|C| \geq 4$.

If the image of f is contained entirely in a single vertex graph then $|C| \leq \theta \leq \theta'$ since $\xi' \geq \xi$. So let us assume that f is not confined to a single vertex graph. Then $\varphi \circ f$ maps some pair of distinct edges e and e' of C onto a common edge of some arc P of T. Then f maps e and e' onto a pair of edges in the same relative position in lifts of P. Since P has length M and the attaching maps of φ have diameter bounded by M, we have $d_\Gamma(p_{\bar{e}}, p_{\bar{e}'}) \leq 2M$ where $p_{\bar{e}}$ and $p_{\bar{e}'}$ are the midpoints of $\bar{e}$ and $\bar{e}'$. Then by Proposition 10.0.1 $d_C(p_e, p_{e'}) \leq 2M + (1-\xi')\frac{|C|}{2}$ where p_e and $p_{e'}$ are the midpoints of e and e'.

Let $L = 2M + (1-\xi')\frac{|C|}{2}$. If $L \geq \frac{|C|}{3} + 1$ then we have

$$\begin{aligned} 0 &\leq 2M + \Big(1 - \frac{\xi+2}{3}\Big)\frac{|C|}{2} - \frac{|C|}{3} - 1 \\ &= 2M - 1 + \Big(\frac{1}{2} - \frac{\xi+2}{6} - \frac{1}{3}\Big)|C| \\ &= 2M - 1 - \Big(\frac{\xi+1}{6}\Big)|C| \end{aligned}$$

and so $|C| \leq \frac{6(2M-1)}{\xi+1} \leq \theta'$. So we may assume that $L < \frac{|C|}{3} + 1$ and can apply Lemma 12.2.2 to $\varphi \circ f$ and L to obtain a vertex $w \in T^0$ such that any segment $Q \subset C$ whose interior is disjoint from $(\varphi \circ f)^{-1}(w)$ has length $|Q| \leq L + 1$. Let $v \in (\varphi \circ f)^{-1}(w)$ be a vertex.

Suppose w is an interior vertex of an arc P of T. Let p be the antipode of v and let p' be a point of $(\varphi \circ f)^{-1}(w)$ that is nearest to p. Then $d_C(p, p') \leq \frac{L+1}{2}$ and so $d_C(p', v) \geq \frac{|C|}{2} - \frac{L+1}{2}$. So, since arcs have length M and the images of attaching maps of φ have diameter at most M, we have

$$\begin{aligned} 2M &\geq d_\Gamma\big(f(p'), f(v)\big) \\ &\geq d_C(p', v) - (1-\xi')\frac{|C|}{2} \\ &\geq \frac{|C|}{2} - \frac{L+1}{2} - (1-\xi')\frac{|C|}{2} \\ &= \xi'\frac{|C|}{2} - \frac{L+1}{2} \end{aligned}$$

where the second inequality holds by Proposition 10.0.1. So, recalling that $L = 2M + (1-\xi')\frac{|C|}{2}$ we have

$$2M \geq \xi'\frac{|C|}{2} - M - (1-\xi')\frac{|C|}{4} - \frac{1}{2}$$

which gives $|C| \leq \frac{4M+2}{3\xi'-1} = \frac{4M+2}{\xi+1} \leq \theta'$.

Suppose w is a node of T. Then $(\varphi \circ f)^{-1}(w) = f^{-1}(\Gamma_w)$. Let $(P_i)_i$ be the components of $f^{-1}(\Gamma_w)$ and let $(Q_j)_j$ be the closures of the components of $C \setminus f^{-1}(\Gamma_w)$. Then $|Q_j| \leq L + 1$ for each j and f maps each P_i into Γ_w

and maps each Q_j into the closure of the complement of Γ_w. We will define a cycle with degeneracies $f'\colon C \to \Gamma_w$ that agrees with f on the P_i and that maps each Q_j onto a geodesic. To see that this is possible, we need only to show that the endpoints of each Q_j map to a distance of at most $|Q_j|$ in Γ_w. The endpoints of Q_j map to a distance of at most M since M is bound on the diameters of the attaching maps of φ. So we need only consider the case where $|Q_j| < M$. But then Q_j is not long enough for $f|_{Q_j}$ to traverse the lift of an arc of T since the arcs have length M. Hence the endpoints of Q_j map to the same vertex of Γ_w. So we are able to define $f'\colon C \to \Gamma_w$. Then, for a point $p \in C$, we have $d_\Gamma\big(f(p), f'(p)\big) \le \frac{M}{2} + \frac{L+1}{2}$. So if $p, q \in C$ are any antipodal pair then

$$\begin{aligned} d_{\Gamma_w}\big(f'(p), f'(q)\big) &\ge d_\Gamma\big(f'(p), f'(q)\big) \\ &\ge d_\Gamma\big(f(p), f(q)\big) - M - L - 1 \\ &\ge \xi' \frac{|C|}{2} - M - L - 1 \end{aligned}$$

and so $d_{\Gamma_w}\big(f'(p), f'(q)\big) \ge \big(\xi' - \frac{2(M+L+1)}{|C|}\big)\frac{|C|}{2}$. So as long as $\xi' - \frac{2(M+L+1)}{|C|} \ge \xi$ then $|C| \le \theta \le \theta'$. If $\xi' - \frac{2(M+L+1)}{|C|} < \xi$ then we have

$$\begin{aligned} 0 &< \xi - \frac{\xi+2}{3} + \frac{2\Big(M + 2M + \big(1 - \frac{\xi+2}{3}\big)\frac{|C|}{2} + 1\Big)}{|C|} \\ &= \frac{2\xi - 2}{3} + \frac{6M}{|C|} + \left(1 - \frac{\xi+2}{3}\right) + \frac{2}{|C|} \\ &= \frac{\xi - 1}{3} + \frac{6M+2}{|C|} \end{aligned}$$

and so $|C| < \frac{18M+6}{1-\xi} \le \theta'$. $\square$

Corollary 12.2.5. *Let $\mathscr{G}$ be a finite graph of (strongly) shortcut groups with finite edge groups. Then the fundamental group of $\mathscr{G}$ is (strongly) shortcut.*

Proof. Let Γ be the underlying graph of $\mathscr{G}$. We construct a graph of spaces $\mathscr{H}$ on Γ such that the fundamental group functor sends $\mathscr{H}$ to $\mathscr{G}$. See Scott and Wall for this viewpoint on graphs of groups [46]. By Corollary 11.1.2, we can choose the vertex spaces so that their universal covers have (strongly) shortcut 1-skeleton. The 1-skeleton $\tilde{\Gamma}$ of the universal cover of $\mathscr{H}$ has the structure $\tilde{\Gamma} \to T$ of a tree of graphs $\tilde{\Gamma} \to T$ on which $\pi_1(\mathscr{G})$ acts freely and cocompactly. For M large enough, subdividing each edge of T into an arc of length M results in a tree of graphs that satisfies the conditions of Theorem 12.2.4. $\square$

Corollary 12.2.6. *Amalgamations and HNN extensions of (strongly) shortcut groups over finite subgroups are (strongly) shortcut.*

Note that $\mathrm{BS}(1,2)$ is an HNN extension of $\mathbb{Z}$ but is not strongly shortcut. Hence, we see that the condition that the edge groups be finite is essential in the strong shortcut case.

Chapter 13

Examples

In this section we prove that hyperbolic graphs, 1-skeleta of CAT(0) cube complexes, the standard Cayley graphs of finitely generated Coxeter groups and all Cayley graphs of $\mathbb{Z}$ and $\mathbb{Z}^2$ are strongly shortcut. In particular, hyperbolic groups, cocompactly cubulated groups and finitely generated Coxeter groups are strongly shortcut.

13.1 Hyperbolic graphs

In this section we will prove that hyperbolic graphs are strongly shortcut. To do so we will make use of the following proposition whose proof is given in Bridson and Haefliger [5].

Proposition 13.1.1 (Specialization of Proposition 1.6 of Part III of Bridson and Haefliger [5])**.** *Let Γ be a δ-hyperbolic graph. Let $f\colon P \to \Gamma$ be a 1-Lipchitz map to Γ from a compact interval $P \subset \mathbb{R}$. If $Q \subset \Gamma$ is the image of a geodesic joining the endpoints of f, then*

$$d_\Gamma\big(x, f(P)\big) \leq \delta \max\{0, \log_2 |P|\} + 1$$

for every $x \in Q$.

Theorem 13.1.2. *Let Γ be a hyperbolic graph. Then Γ is strongly shortcut.*

Proof. Let $\delta \geq 1$ be a hyperbolicity constant for Γ. Suppose $f\colon C \to \Gamma$ is a $\frac{3}{4}$-almost isometric cycle of length $|C| \geq 2$. Let $y, y' \in C$ be a pair of antipodal points and let $P_1 \subset C$ and $P_2 \subset C$ be the two segments of C between y and y'. Let Q be the image of a geodesic in Γ from $f(y)$ to $f(y')$ and let x be the midpoint of Q. Then, by Proposition 13.1.1, there are points $p_1 \in P_1$ and $p_2 \in P_2$ such that $f(p_1)$ and $f(p_2)$ are each at distance at most $\delta \log_2 \frac{|C|}{2} + 1$ from x in Γ. Then, since f is $\frac{3}{4}$-almost isometric, we have $|Q| \geq \frac{3|C|}{8}$ and

$$\begin{aligned} \frac{3|C|}{16} \leq \frac{1}{2}|Q| &= d_\Gamma\big(f(y), x\big) \\ &\leq d_\Gamma\big(f(y), f(p_1)\big) + d_\Gamma\big(f(p_1), x\big) \\ &\leq d_C(y, p_1) + \delta \log_2 \frac{|C|}{2} + 1 \end{aligned}$$

and so $d_C(y, p_1) \geq \frac{3|C|}{16} - \delta \log_2 \frac{|C|}{2} - 1$. By the same argument we have the same lower bound for $d_C(y, p_2)$ and $d_C(y', p_1)$ and $d_C(y', p_2)$. Hence

$$d_C(p_1, p_2) \geq \frac{3|C|}{8} - 2\delta \log_2 \frac{|C|}{2} - 2$$

and so, by Proposition 10.0.1, we have

$$d_\Gamma\big(f(p_1), f(p_2)\big) \geq \frac{3|C|}{8} - 2\delta \log_2 \frac{|C|}{2} - 2 - \Big(1 - \frac{3}{4}\Big)\frac{|C|}{2} = \frac{|C|}{4} - 2\delta \log_2 \frac{|C|}{2} - 2$$

but $f(p_1)$ and $f(p_2)$ are both within a distance of $\delta \log_2 \frac{|C|}{2} + 1$ to x and so $d_\Gamma\big(f(p_1), f(p_2)\big) \leq 2\delta \log_2 \frac{|C|}{2} + 2$. Hence we have

$$|C| \leq 16\Big(\delta \log_2 \frac{|C|}{2} + 1\Big)$$

which bounds the length $|C|$ of f. $\square$

Corollary 13.1.3. *Hyperbolic groups are strongly shortcut.*

13.2 CAT(0) cube complexes

In this section we will prove that the 1-skeleton of a finite-dimensional CAT(0) cube complex is strongly shortcut. The proof rests on a theorem about edge colorings of cycles.

Let C be a cycle. An *edge coloring* of C is a function $\alpha\colon C^{(1)} \to W$ from the set $C^{(1)}$ of edges of C to some set W of *colors*. A cycle C along with an edge coloring $\alpha\colon C^{(1)} \to W$ is a *wall cycle* if α is surjective and, for each $w \in W$, the number $|\alpha^{-1}(w)|$ of edges of color w is even. In this case we may refer to the elements of W as *walls*.

Let (C, α) be a wall cycle. A combinatorial segment $P \subset C$ *crosses* a wall $w \in W$ if the number of edges of P colored w is odd. A combinatorial segment $P \subset C$ *begins and ends* with a wall $w \in W$ if the initial and terminal edges of P map to w under α. Two distinct walls $w_1, w_2 \in W$ *cross* if for some combinatorial segment $P \subset C$, we have that P begins and ends with one of the two walls and P crosses the other of the two walls. The *dimension* d of a wall cycle (C, α) is defined as $d = \max\{1, n\}$ where n is the size of the largest set $S \subseteq W$ of pairwise crossing walls. The *wall crossing distance* $d_\alpha(u, v)$ between a pair of vertices $u, v \in C^0$ is defined as the number of walls crossed by a segment $P \subset C$ from u to v. Note that the choice of segment P does not matter since each wall appears an even number of times along C.

Proposition 13.2.1. *Let X be a* CAT(0) *cube complex. Let W be the set of hyperplanes of X and let $\beta\colon X^{(1)} \to W$ map each edge e of X to the hyperplane that e crosses. Then for any cycle $f\colon C \to X^1$, the coloring (C, α) is a wall cycle, where $\alpha(e) = \beta\big(f(e)\big)$ for $e \in C^{(1)}$. Moreover, two crossing walls of (C, α) must cross in X and so the dimension of X is at least the dimension of (C, α). Lastly, the wall crossing distance on (C, α) satisfies $d_\alpha(u, v) = d_{X^1}\big(f(u), f(v)\big)$.*

Proof. That (C, α) is a wall cycle is a consequence of the fact that hyperplanes of a CAT(0) complex are two-sided. That two crossing walls of (C, α) must cross in X is a consequence of the fact that hyperplanes are connected and two-sided. The dimension of a CAT(0) cube complex is equal to the size of the largest set of its pairwise crossing hyperplanes. Finally, the combinatorial distance between two vertices of a CAT(0) cube complex is equal to the number of hyperplanes separating them. □

In light of Proposition 13.2.1, the following theorem implies that the 1-skeleta of d-dimensional CAT(0) cube complexes are strongly shortcut.

Theorem 13.2.2. *Let (C, α) be a d-dimensional wall cycle. If $d_\alpha(u,v) \geq \left(\frac{5d-1}{5d}\right)\frac{|C|}{2}$ for all antipodal pairs of vertices $u, v \in C^0$ then $|C| \leq \frac{50d^2}{5d-1}$.*

Corollary 13.2.3. *The 1-skeleta of finite dimensional* CAT(0) *cube complexes are strongly shortcut.*

Corollary 13.2.4. *Cocompactly cubulated groups are strongly shortcut.*

A group is *cocompactly cubulated* if it acts properly and cocompactly on a CAT(0) cube complex.

The proof of Theorem 13.2.2 relies several lemmas and on the following theorem of Turan.

Theorem 13.2.5 (Turan's Theorem)**.** *Let Γ be a simplicial graph on n vertices. If every complete subgraph of Γ has at most $d \in \mathbb{N}$ vertices then Γ has at most $\left(\frac{d-1}{d}\right)\frac{n^2}{2}$ edges.*

Several proofs of Turan's Theorem are given in Aigner and Ziegler [1].

Lemma 13.2.6. *Let (C, α) be a wall cycle and suppose that for some $\xi \in (0,1)$ we have $d_\alpha(u,v) \geq \xi\frac{|C|}{2}$ for every antipodal pair $u, v \in C^0$. Let $W' = \{w \in W : |\alpha^{-1}(w)| = 2\}$. Then $|W \setminus W'| \leq \frac{1-\xi}{\xi}|W|$.*

Proof. Partition $C^{(1)}$ into two sets S and T such that

$$\left|(\alpha|_S)^{-1}(w)\right| = \left|(\alpha|_T)^{-1}(w)\right|$$

for each $w \in W$. This is always possible since $|\alpha^{-1}(w)|$ is even for each $w \in W$. Viewing the elements of S as colored by W, every color appears in S and those colors in $W \setminus W'$ appear at least twice. Hence $|W| \leq |S| - |W \setminus W'| =$

$\frac{|C|}{2} + |W'| - |W|$. Let u and v be an antipodal pair of vertices. Then we have

$$\xi \frac{|C|}{2} \leq d_\alpha(u, v) \leq |W| \leq \frac{|C|}{2} + |W'| - |W|$$

and so $|W'| \geq |W| - (1-\xi)\frac{|C|}{2} \geq |W| - \frac{1-\xi}{\xi}|W| = \frac{2\xi-1}{\xi}|W|$. Hence we have $|W \setminus W'| = |W| - |W'| \leq \left(1 - \frac{2\xi-1}{\xi}\right)|W| = \frac{1-\xi}{\xi}|W|$. $\square$

Let (C, α) be a wall cycle and let w be a wall of (C, α). Let $X_w \subset C$ denote the set of all midpoints of edges colored w and let $\operatorname{diam} X_w$ denote the diameter of X_w as a metric subspace of (C, d_C). For a pair of vertices $u, v \in C^0$ we say that w *contributes* to $\{u, v\}$ if a geodesic segment from u to v crosses w. Hence $d_\alpha(u, v)$ is equal to the number of walls contributing to $\{u, v\}$.

Lemma 13.2.7. *Let (C, α) be a wall cycle and let $w \in W$ be a wall such that the number of edges colored w is exactly 2. Then w contributes to $\{u, v\}$ for exactly $\operatorname{diam} X_w$ antipodal pairs of vertices $u, v \in C^0$.*

Proof. Let $P \subset C$ be a segment beginning and ending with w of length $|P| = \operatorname{diam} X_w + 1$. Then w contributes to an antipodal pair $\{u, v\}$ if and only if one of u or v is an interior vertex of P and there are exactly $|P| - 1 = \operatorname{diam} X_w$ such pairs. $\square$

Lemma 13.2.8. *Let (C, α) be a wall cycle and suppose that, for some $\xi \in (0, 1)$, we have $d_\alpha(u, v) \geq \xi \frac{|C|}{2}$ for all antipodal pairs of vertices $u, v \in C^0$. Let $w \in W$ be a wall. Then w crosses at least $\operatorname{diam} X_w - 1 - (1-\xi)\frac{|C|}{2}$ walls.*

Proof. Consider first the case where $\operatorname{diam} X_w = \frac{|C|}{2}$. Then there exist a pair of antipodal edges e and e' colored w. Let $u \in e$ and $u' \in e'$ be a pair of antipodal vertices and let $P \subset C$ be a segment with endpoints u and u'. Note that P contains exactly one of e or e'. Without loss of generality P contains e. Then, since $d_\alpha(u, u') \geq \xi \frac{|C|}{2}$, then P must cross at least $\xi \frac{|C|}{2} - 1$ walls aside from w. Then the same must hold for $P \cup e'$ and so w crosses at least $\xi\frac{|C|}{2} - 1 = \operatorname{diam} X_w - 1 - (1-\xi)\frac{|C|}{2}$ walls.

Consider now the case $\operatorname{diam} X_w < \frac{|C|}{2}$. We have a geodesic segment $P \subset C$ beginning and ending with w such that $|P| = \operatorname{diam} X_w + 1$. Let u and v be the endpoints of P, let u' be the antipode of u and let Q be the geodesic segment containing P and having endpoints u and u'. Then we have

$$\begin{aligned}
\xi\frac{|C|}{2} \leq d_\alpha(u,u') &\leq d_\alpha(u,v) + d_\alpha(v,u') \\
&\leq d_\alpha(u,v) + d_C(v,u') \\
&= d_\alpha(u,v) + \frac{|C|}{2} - (\operatorname{diam} X_w + 1)
\end{aligned}$$

and so we have $d_\alpha(u,v) \geq \operatorname{diam} X_w + 1 - (1-\xi)\frac{|C|}{2}$. But w crosses at least $d_\alpha(u,v) - 1$ walls and so we are done. □

Proof of Theorem 13.2.2. Let $\xi = \left(\frac{5d-1}{5d}\right)$. For each vertex pair $\{u,v\}$ and each wall $w \in W$, let $\mathbb{1}_w^{\{u,v\}}$ be defined as follows.

$$\mathbb{1}_w^{\{u,v\}} = \begin{cases} 1 & \text{if } w \text{ contributes to } \{u,v\} \\ 0 & \text{otherwise} \end{cases}$$

Let $W' \subseteq W$ be the set of walls which color exactly two edges of C.

We have

$$d_\alpha(u,v) = \sum_{w \in W} \mathbb{1}_w^{\{u,v\}}$$

and, by Lemma 13.2.7, for $w \in W'$ we have

$$\operatorname{diam} X_w = \sum_{\{u,v\}} \mathbb{1}_w^{\{u,v\}}$$

where the sum ranges over all antipodal pairs of vertices $\{u,v\}$. Let Γ be the simplicial graph with vertex set W and where two walls are joined by an edge if they cross. Then we have

$$
\begin{aligned}
&|\Gamma^{(1)}| \\
&\geq \frac{1}{2}\sum_{w\in W'}\Big(\deg(w)\Big) \\
&\geq \frac{1}{2}\sum_{w\in W'}\Big(\operatorname{diam} X_w - 1 - (1-\xi)\frac{|C|}{2}\Big) \\
&= \frac{1}{2}\sum_{w\in W'}\Big(\sum_{\{u,v\}} \mathbb{1}_w^{\{u,v\}}\Big) - \frac{1}{2}|W'| - \frac{1}{2}|W'|(1-\xi)\frac{|C|}{2} \\
&= \frac{1}{2}\sum_{\{u,v\}}\Big(\sum_{w\in W} \mathbb{1}_w^{\{u,v\}}\Big) - \frac{1}{2}\sum_{w\in W\setminus W'}\Big(\sum_{\{u,v\}} \mathbb{1}_w^{\{u,v\}}\Big) - \frac{1}{2}|W'| - \frac{1}{2}|W'|(1-\xi)\frac{|C|}{2} \\
&= \frac{1}{2}\sum_{\{u,v\}}\Big(d_\alpha(u,v)\Big) - \frac{1}{2}\sum_{w\in W\setminus W'}\Big(\sum_{\{u,v\}} \mathbb{1}_w^{\{u,v\}}\Big) - \frac{1}{2}|W'| - \frac{1}{2}|W'|(1-\xi)\frac{|C|}{2} \\
&\geq \frac{1}{2}\xi\Big(\frac{|C|}{2}\Big)^2 - \frac{1}{2}|W\setminus W'|\cdot\frac{|C|}{2} - \frac{1}{2}|W'| - \frac{1}{2}|W'|(1-\xi)\frac{|C|}{2} \\
&\geq \frac{1}{2}\xi|W|^2 - \frac{1}{2}\Big(\frac{1-\xi}{\xi}\Big)|W|\cdot\frac{1}{\xi}|W| - \frac{1}{2}|W| - \frac{1}{2}|W|(1-\xi)\frac{1}{\xi}|W| \\
&= \Big(\xi - \frac{1-\xi}{\xi} - \frac{1-\xi}{\xi^2} - \frac{1}{|W|}\Big)\frac{|W|^2}{2} \\
&= \Big(\xi + 1 - \frac{1}{\xi^2} - \frac{1}{|W|}\Big)\frac{|W|^2}{2} \\
&\geq \Big(\xi + 1 - \frac{1}{\xi^2} - \frac{2}{\xi|C|}\Big)\frac{|W|^2}{2}
\end{aligned}
$$

where the second inequality holds by Lemma 13.2.8 and the second to last inequality holds by Lemma 13.2.6. We now verify that $4x-3 \leq x+1-\frac{1}{x^2}$ for $x \in \left[\frac{4}{5},1\right]$, noting that it suffices to check the inequality for $x=\frac{4}{5}$ and $x=1$ since $x \mapsto x+1-\frac{1}{x^2}$ is a concave function. Then, since $\xi = \frac{5d-1}{5d} \in \left[\frac{4}{5},1\right]$ we have

$$
|\Gamma^{(1)}| \geq \Big(4\xi - 3 - \frac{2}{\xi|C|}\Big)\frac{|W|^2}{2} = \Big(\frac{4(5d-1)}{5d} - 3 - \frac{2}{|C|}\cdot\frac{5d}{5d-1}\Big)\frac{|W|^2}{2}
$$

and, since (C,α) is d-dimensional, we have

$$
\frac{4(5d-1)}{5d} - 3 - \frac{2}{|C|}\cdot\frac{5d}{5d-1} \leq \frac{d-1}{d}
$$

by Turan's Theorem (Theorem 13.2.5). After some rearranging and cancellation this inequality becomes $|C| \leq \frac{50d^2}{5d-1}$. □

13.3 Cayley graphs of Coxeter groups

In this section we use the cubulation of Coxeter groups of Niblo and Reeves [38] and our result on CAT(0) cube complexes to prove that Coxeter groups are strongly shortcut.

Let Γ be a simplicial graph on the vertex set $\{v_1, v_2, \ldots, v_n\}$ with every edge labeled by an integer at least 2. If Γ has an edge e from v_i to v_j then let $m_{ij} = m_{ji}$ denote the label of e. The *Coxeter group* C_Γ defined by Γ is given by the following presentation

$$\left\langle v_1, v_2, \ldots, v_n \mid v_k^2 = 1 \text{ for all } k \text{ and } (v_i v_j)^{m_{ij}} = 1 \text{ for all edges } \{v_i, v_j\} \in \Gamma^{(1)} \right\rangle$$

For a Coxeter group C_Γ, Niblo and Reeves [38] construct a finite dimensional CAT(0) cube complex into whose 1-skeleton the Cayley graph of C_Γ with generating set Γ^0 isometrically embeds. Hence, since the 1-skeleta of CAT(0) cube complexes are strongly shortcut, we have the following theorem.

Theorem 13.3.1. *Coxeter groups are strongly shortcut.*

13.4 Cayley graphs of $\mathbb{Z}$ and $\mathbb{Z}^2$

We have shown that the 1-skeleta of CAT(0) cube complexes are strongly shortcut. In particular, the standard Cayley graphs of the finitely generated free abelian groups are strongly shortcut. In this section we will strengthen this result for $\mathbb{Z}$ and $\mathbb{Z}^2$ by showing that all of their Cayley graphs are strongly shortcut.

Lemma 13.4.1. *Let Γ be a graph and suppose there is a continuous (K, M)-quasi-isometric embedding $\iota\colon \Gamma \to \mathbb{R}^2$. Let $\xi \in (0, 1)$ and let $f\colon C \to \Gamma$ be a ξ-almost isometric cycle. Suppose the image of $\iota \circ f$ is contained in the N-neighborhood of a line $L \subset \mathbb{R}$. Then $|C| \leq \frac{2K}{\xi}(M + 2N)$.*

Proof. By continuity, for some pair of antipodal points $p, q \in C$, the points $\iota \circ f(p)$ and $\iota \circ f(q)$ project perpendicularly to the same point of L. Then

$$d_{\mathbb{R}}\big(\iota \circ f(p), \iota \circ f(q)\big) \leq 2N$$

and so we have

$$\xi \frac{|C|}{2} \leq d_{\Gamma}\big(f(p), f(q)\big) \leq K(M + 2N)$$

and so we have $|C| \leq \frac{2K}{\xi}(M + 2N)$. $\square$

Since the inclusion map $\mathbb{Z} \hookrightarrow \mathbb{R} \times \{0\}$ extends to a continuous quasi-isometric embedding from any Cayley graph of $\mathbb{Z}$, we obtain as a corollary of Lemma 13.4.1 the following theorem.

Theorem 13.4.2. *Every Cayley graph of $\mathbb{Z}$ is strongly shortcut.*

In the remainder of this section we will prove that the same holds for Cayley graphs of $\mathbb{Z}^2$. Let S be a generating set of $\mathbb{Z}^2$, let the Γ be the Cayley graph of $(\mathbb{Z}^2, S)$ and let $\iota\colon \Gamma \to \mathbb{R}^2$ be the (K, M)-quasi-isometry obtained by extending the inclusion map $\mathbb{Z}^2 \hookrightarrow \mathbb{R}^2$ to Γ in such a way that the restriction of ι to each edge is a geodesic.

Lemma 13.4.3. *Let $f\colon C \to \Gamma$ be a ξ-almost isometric embedding. For some constants A and B depending only on S, there is a line in $\mathbb{R}^2$ whose $\big((1 - \xi)A|C| + B\big)$-neighborhood contains the image of $\iota \circ f$.*

Proof. For $x \in \mathbb{R}$ let $|x|$ denote the standard Euclidean norm of x. Let $t \in S$ achieve $|t| = \max\big\{|s| : s \in S\big\}$. Let V be the vector subspace of $\mathbb{R}^2$ generated by t. For $s \in S$ let s_t be the perpendicular projection of s onto t and let $\alpha = \max\big\{|s_t| : s \in S \setminus \{\pm t\}\big\}$. Then $\alpha < |t|$.

By continuity, some pair of antipodal points $p, q \in C$ satisfy $\iota \circ f(p) - \iota \circ f(q) \in V$. Pick $u \in \mathbb{Z}^2$ such that $|u - \iota \circ f(p)| \leq 1$. Then for some $r \in \mathbb{Z}$, we have $|u + rt - \iota \circ f(q)| \leq 1 + |t|$. Then $d_{\Gamma}\big(u, f(p)\big) \leq K(M + 1)$ and

$d_\Gamma\big(u + rt, f(q)\big) \le K\big(M + 1 + |t|\big)$. Hence we have

$$\begin{aligned} \xi\frac{|C|}{2} &\le d_\Gamma\big(f(p), f(q)\big) \\ &\le d_\Gamma\big(f(p), u\big) + d_\Gamma(u, u + rt) + d_\Gamma\big(u + rt, f(q)\big) \\ &\le r + K\big(2M + 2 + |t|\big) \end{aligned}$$

and so $r \ge \xi\frac{|C|}{2} - K\big(2M + 2 + |t|\big)$. Let $P \subset C$ be a segment with endpoints p and q. Label each edge of C by the label of the edge it maps to under f. Let T be the union of all t-labeled edges of C and let ℓ be the total length of the segments of $T \cap P$. Consider the projection of the path $\iota \circ f|_P$ onto the line $p + V$. It has arclength at most $\ell|t| + \big(\frac{|C|}{2} - \ell\big)\alpha$. But the endpoints of $\iota \circ f|_P$ are $\iota \circ f(p)$ and $\iota \circ f(q)$, which are of distance at least $(r-1)|t| - 2$ apart and so $(r-1)|t| - 2 \le \ell|t| + \big(\frac{|C|}{2} - \ell\big)\alpha$. Combining this inequality with $r \ge \xi\frac{|C|}{2} - K\big(2M + 2 + |t|\big)$ we have

$$\left(\xi\frac{|C|}{2} - K\big(2M + 2 + |t|\big) - 1\right)|t| - 2 \le \ell|t| + \left(\frac{|C|}{2} - \ell\right)\alpha$$

which, after some manipulation gives

$$\ell \ge \left(\frac{\xi|t| - \alpha}{|t| - \alpha}\right)\frac{|C|}{2} - \frac{K(2M + 3 + |t|)|t| + 2}{|t| - \alpha}$$

and so we have the following inequality.

$$\frac{|C|}{2} - \ell \le \left(\frac{(1-\xi)|t|}{|t| - \alpha}\right)\frac{|C|}{2} + \frac{K(2M + 3 + |t|)|t| + 2}{|t| - \alpha}$$

But then the projection of $\iota \circ \alpha|_P$ to $V^\perp$ must have length at most

$$\left(\frac{(1-\xi)|t|^2}{|t| - \alpha}\right)\frac{|C|}{2} + \frac{K(2M + 3 + |t|)|t|^2 + 2|t|}{|t| - \alpha}$$

and so the image of $\iota \circ \alpha|_P$ must be contained in a neighborhood of radius

$$\left(\frac{(1-\xi)|t|^2}{2(|t| - \alpha)}\right)\frac{|C|}{2} + \frac{K(2M + 3 + |t|)|t|^2 + 2|t|}{2(|t| - \alpha)}$$

about the line $L = p + V$. Then setting

$$A = \Big(\frac{|t|^2}{4(|t| - \alpha)}\Big)$$

and

$$B = \frac{K(2M + 3 + |t|)|t|^2 + 2|t|}{2(|t| - \alpha)}$$

we are done. □

Theorem 13.4.4. *Every Cayley graph of $\mathbb{Z}^2$ is strongly shortcut.*

Proof. Let $f : C \to \Gamma$ be a ξ-almost isometric cycle. By Lemma 13.4.3 we have a line $L \subset \mathbb{R}^2$ whose $\big((1 - \xi)A|C| + B\big)$-neighborhood contains the image of $\iota \circ f$. So, by Lemma 13.4.1, we have

$$|C| \leq \frac{2K}{\xi}\big(M + 2(1 - \xi)A|C| + 2B\big)$$

and so

$$\left(1 - \frac{4K}{\xi}(1 - \xi)A\right)|C| \leq \frac{2K}{\xi}(M + 2B)$$

which gives us a bound on the length of $|C|$ assuming we have $1 - \frac{4K}{\xi}(1 - \xi)A > 0$. But this condition is equivalent to $\xi > \frac{4KA}{1+4KA}$, which we can satisfy. Hence, for $\xi \in \big(\frac{4KA}{1+4KA}, 1\big)$, there is a bound on the length of the ξ-almost isometric cycles of Γ. □

Chapter 14

The Baumslag-Solitar group BS(1, 2)

The Baumslag-Solitar group BS(1, 2) is defined by the following presentation.

$$\langle a, t \mid tat^{-1} = a^2 \rangle$$

In this section we will show that the standard Cayley graph of $G = \mathrm{BS}(1, 2)$ is shortcut but that adding the generator $\tau = t^2$ results in a Cayley graph $\mathrm{Cay}\big(G, \{a, t, \tau\}\big)$ which is not shortcut. Hence we see that there exists a shortcut group with exponential Dehn function [17] and that the shortcut property for a Cayley graph is not invariant under a change of generating set. We also see that there exists a shortcut group which is not strongly shortcut, since strongly shortcut groups have polynomial isoperimetric function, by Corollary 11.2.2.

Let Γ be the Cayley graph of BS(1, 2) with generating set $\{a, t\}$. Since BS(1, 2) is an HNN extension it has a Bass-Serre tree T. Every vertex of T has two outgoing edges labeled t and one incoming edge labeled t.

Lemma 14.0.1. *Every element of* BS(1, 2) *can be written uniquely in the form* $t^m a^k t^n$ *where* $m, k, n \in \mathbb{Z}$ *and* k *is even only if* $k = m = 0$.

Proof. Given any word representing an element of BS(1, 2) in the standard generators, we may commute positive powers of t to the right and negative powers of t to the left using the relations $t^n a^k = a^{2^n k} t^n$ and $a^k t^{-n} = t^{-n} a^{2^n k}$, with $n \geq 0$, to obtain a representative of the form $t^m a^k t^n$. Then we may apply the relation $a^k = t^n a^{k/2^n} t^{-n}$ if k is a nonzero integer multiple of 2^n, with $n \geq 0$, to obtain a representative of the form $t^m a^k t^n$ where k is even if and only if $k = m = 0$.

To see that this form is unique, let $t^{m'} a^{k'} t^{n'} = t^m a^k t^n$. By consideration of the Bass-Serre tree T we must have $m + n = m' + n'$. Without loss of generality $m \geq m'$ and so we have

$$a^{k'} t^{n'} = t^{m-m'} a^k t^n = a^{2^{m-m'}k} t^{m-m'} t^n = a^{2^{m-m'}k} t^{n'}$$

and so, as the base group embeds in an HNN extension, we have $k' = 2^{m-m'}k$. So, in the case where $k = m = 0$, we have $k' = 0$ and so $m' = 0$, which implies $(m', k', n') = (m, k, n)$. If $k \neq 0$ then $k' \neq 0$ and so k and k' are both odd integers. Hence $2^{m-m'} = 1$, which again implies $(m', k', n') = (m, k, n)$. $\square$

It follows from Lemma 14.0.1 that we have a one-to-one correspondence

$$\begin{aligned} \varphi\colon G &\to \mathbb{Z}\left[\frac{1}{2}\right] \times \mathbb{Z} \\ t^m a^k t^n &\mapsto (2^m k, m + n) \end{aligned}$$

with inverse

$$\begin{aligned} \varphi^{-1}\colon \mathbb{Z}\left[\frac{1}{2}\right] \times \mathbb{Z} &\to G \\ (r, z) &\mapsto \begin{cases} t^{\nu(r)} a^{r/2^{\nu(r)}} t^{z-\nu(r)} & \text{if } r \neq 0 \\ t^z & \text{if } r = 0 \end{cases} \end{aligned}$$

where $\mathbb{Z}\left[\frac{1}{2}\right]$ is the set of dyadic rationals and $\nu(r)$ is defined as follows.

$$\nu(r) = \max\Big\{m \in \mathbb{Z} : r \text{ is an integer multiple of } 2^m\Big\}$$

The *height* of a point (r, z) is z. We use $\mu(g)$ to denote the height of $\varphi(g)$ for $g \in \mathrm{BS}(1, 2)$.

Pushing forward the group operation to $\mathbb{Z}[\frac{1}{2}] \times \mathbb{Z}$ gives the following operation.

$$(r, z) \cdot (r', z') = (r + 2^z r', z + z')$$

Pushing forward the Cayley graph structure gives the following edges.

$$(r, z) \xrightarrow{a} (r + 2^z, z)$$
$$(r, z) \xrightarrow{t} (r, z + 1)$$

The Bass-Serre tree T may be identified with the quotient of this graph which identifies (r, z) and (r', z) if $r - r'$ is an integer multiple of 2^z. This identification preserves height so we may refer to the height $\mu(v)$ of a vertex v of T.

14.1 Geodesics in $\mathrm{BS}(1,2)$

We now prove some lemmas about geodesics in the Cayley graph Γ. We will describe paths in Γ using words in the letters $\{a^{\pm 1}, t^{\pm 1}\}$. We use the notation $w_1 \equiv w_2$ to mean that the words w_1 and w_2 represent the same elements of $\mathrm{BS}(1,2)$. Of course if $w_1 \equiv w_2$ then paths described by w_1 and w_2 and starting at the same initial vertex must have the same final vertex.

Remark 14.1.1. A path in Γ may be projected to a path in T. A backtrack in the projection corresponds to a subword of the form $ta^k t^{-1}$ or $t^{-1}a^{2k}t$. The initial and terminal edges of a path described by twt^{-1} project to the same edge in T if and only if $w \equiv a^k$. The initial and terminal edges of a path described by $t^{-1}wt$ project to the same edge in T if and only if $w \equiv a^{2k}$.

Lemma 14.1.2. *Every finite subtree of T has a unique vertex of minimal height.*

Proof. Let T' be a finite subtree and suppose v and v' are minimal height vertices of T'. Then v and v' correspond to some (r, z) and (r', z) in $\mathbb{Z}[\frac{1}{2}] \times \mathbb{Z}$. There exists a path in T' from v to v' and so we may obtain r' from r by adding and subtracting powers of 2 with exponent at least z. But then $r - r'$ is an integer multiple of 2^z and so $(r, z) = (r', z)$ in T. $\square$

Lemma 14.1.3. *Words of the following forms do not describe geodesic paths.*

1. $ta^{\pm 1}t^{-1}$
2. $t^{-1}a^k$ *and* $a^k t$ *with* $|k| \geq 2$
3. $a^{\varepsilon}t^{-1}a^{-\varepsilon}$ *with* $\varepsilon = \pm 1$.
4. $t^{-1}w_1tw_2t^{-1}$ *with* $w_1 \equiv a^{k_1}$ *and* $w_2 \equiv a^{k_2}$
5. $w \equiv t^h$ *with* $w \neq t^h$.
6. w_1w_2 *with* $w_1w_2 \equiv a^k$ *and where* $\mu(w_1) < 0$

Proof. The following equivalences prove nongeodesicity for (1), (2), (3) and (4).

(1) $ta^{\pm 1}t^{-1} \equiv a^{\pm 2}$

(2) $t^{-1}a^k \equiv at^{-1}a^{k-2}$

(3) $a^{\varepsilon}t^{-1}a^{-\varepsilon} \equiv t^{-1}a^{\varepsilon}$

(4) $t^{-1}w_1tw_2t^{-1} \equiv t^{-1}a^{k_1}a^{2k_2} \equiv t^{-1}a^{2k_2}a^{k_1} \equiv t^{-1}tw_2t^{-1}w_1 \equiv w_2t^{-1}w_1$

(5) Suppose w is geodesic with $w \equiv t^h$. Note that $\nu(w) = h$ so w must contain at least $|h|$ instances of t^{ε} where ε is the sign of h. Hence, as $|w| \leq |t^h| = h$, we see that w cannot contain any instance of $a^{\pm 1}$. But w may not contain any backtracks either and so $w = t^h$.

(6) Suppose $w = w_1w_2$ is geodesic with $w \equiv a^k$ and $\mu(w_1) < 0$. By Lemma 14.1.2, there is a unique vertex v of minimal height of the projection of w to the Bass-Serre tree T and $\mu(v) \leq \mu(w_1) < 0 = \mu(1) = \mu(w)$. Then w must contain a subword of the form $t^{-1}a^kt$ with the a^k part mapping to v under the projection to T. Then (2) implies that $|k| = 1$. So $t^{-1}a^kt = t^{-1}a^{\pm 1}t$ corresponds to a nonbacktracking path in T. It follows, since the projection of w to T is a closed path and v is a cutpoint of T, that w contains another subpath of the form $t^{-1}a^{\pm 1}t$ such that $a^{\pm 1}$ maps to v under the projection to

T. Then w contains a subword of the form $t^{-1}a^{\varepsilon_1}tw't^{-1}a^{\varepsilon_2}t$ with $\varepsilon_i \in \{\pm 1\}$ and $w' \equiv a^{k'}$. But, by (4), we know that $t^{-1}a^{\varepsilon_1}tw't^{-1}$ is not geodesic, which is a contradiction. □

Lemma 14.1.4. *If $\mu(w) = 0$ and every prefix w' of w has $\mu(w') \geq 0$ then $w \equiv a^k$ for some k.*

Proof. Viewing w as a path in $\mathbb{Z}\left[\frac{z}{2}\right]$ starting at $(0,0)$, we see that at each step the first coordinate changes by a positive power of 2. Hence the endpoint of the path is $(k,0)$ for some integer k. □

A word w is *ascending* if it contains only positive powers of t and *descending* if it contains only negative powers of t.

Lemma 14.1.5. *Let w be a geodesic word with $w \equiv t^{-h}a^k$ where $h \geq 0$. Then no prefix w' of w satisfies $\mu(w') < -h$ and we have $w = xy$ where x is ascending and y is descending.*

Proof. Suppose $w = w_1w_2$ where w_1 is the smallest prefix of w with $\mu(w_1) = -h$. Since the vertices of T have indegree 1, any two paths in T of the same negative height have the same endpoint. So the projections of w and w_1 to T have the same endpoint. Hence $w_2 \equiv a^\ell$ for some ℓ. Also, any prefix w' of w with $\mu(w') < -h$ must be longer than w_1 and so $w' = w_1w_2'$. Then w_2' is a prefix of w_2 of height $\mu(w_2') = \mu(w') - \mu(w_1) < 0$. So, by Lemma 14.1.3(6), we have that w_2 is not geodesic, contradicting the geodesicity of w.

We now prove that $w = xy$ such that x is ascending and y is descending. If w has no such decomposition then w has a subword of the form $t^{-1}w't$. An innermost such subword has the form $t^{-1}a^kt$. By Lemma 14.1.3(2), we have $|k| \leq 1$. So $w = w_1t^{-1}a^\varepsilon tw_2$ with $\varepsilon = \pm 1$. We have $\mu(w_1t^{-1}a^\varepsilon) \geq -h$ and so $\mu(w_1t^{-1}a^\varepsilon t) > -h$. So $\mu(w_2) < 0$ and the shortest prefix of w_2 of negative height has the form $w_2't^{-1}$ with $\mu(w_2') = 0$. Then, by Lemma 14.1.4, we have

$w_2' \equiv a^k$ for some k. But then

$$w_1 t^{-1} a^{\varepsilon} t w_2' t^{-1} \equiv w_1 t^{-1} t w_2' t^{-1} a^{\varepsilon} \equiv w_1 w_2' t^{-1} a^{\varepsilon}$$

and so $w_1 t^{-1} a^{\varepsilon} t w_2' t^{-1}$ is a nongeodesic subword of w, which is a contradiction.

□

Lemma 14.1.6. *Let $h \geq 1$, let $k \geq 2^h$ and let $\varepsilon = \pm 1$. Let w be a geodesic word with $w \equiv t^{-h} a^{\varepsilon k}$. Then the first letter of w is not t^{-1}.*

Proof. Suppose the first letter of w is t^{-1}. Then, by Lemma 14.1.5, we see that w is descending. Hence

$$w = t^{-\ell_1} a^{k_1} t^{-\ell_2} a^{k_2} \cdots t^{-\ell_m} a^{k_m}$$

with $\sum_i \ell_i = h$ and $\ell_i > 0$ for all i. So we have

$$w \equiv t^{-\sum_i \ell_i} a^{2^{L_1} k_1 + 2^{L_2} k_2 + \cdots + 2^{L_m} k_m}$$

where $L_j = \sum_{i>j} \ell_i$ and so $\varepsilon k = 2^{L_1} k_1 + 2^{L_2} k_2 + \cdots + 2^{L_m} k_m$. But, by Lemma 14.1.3(2), we have $|k_i| \leq 1$ for all i and so

$$|k| \leq 2^{L_1} + 2^{L_2} + \cdots + 2^{L_m}$$

with

$$0 = L_m < L_{m-1} < \cdots < L_1 < h$$

which implies $|k| \leq \sum_{j=0}^{h-1} 2^j = 2^h - 1$, a contradiction. □

Lemma 14.1.7. *Let $h \geq 1$, let $0 \leq k \leq 2^h$ and let $\varepsilon = \pm 1$. Let w be a geodesic word with $w \equiv t^{-h} a^{\varepsilon k}$. Then w contains only negative powers of t and every prefix w' of w satisfies $w' \equiv t^{-h'} a^{\varepsilon k'}$ where $0 \leq h' \leq h$ and $0 \leq k' \leq 2^{h'}$.*

Proof. Since there is an automorphism of $\mathrm{BS}(1,2)$ fixing t and sending a to a^{-1}, we may assume that $\varepsilon = 1$. The proof is by induction on the length of w.

If $|w| = 1$ then $w = t^{-1}$ and so satisfies the required conditions. Assume now that $|w| > 1$. Consider the path $f\colon P \to T$ followed by w in the Bass-Serre tree T. Let v_1 and v_2 be the initial and final vertices of this path. The shortest path in T from v_1 to v_2 is labeled t^{-h}. By Lemma 14.1.3(6), the path f may not traverse an edge below v_2. Hence, any instance of t in w corresponds to an edge of T which is ascended by f and later descended. That is, the instance of t is the first letter of a subword $tw't^{-1}$ of w with $w' \equiv a^{k'}$ for some k'. Then, if w has an instance of t then, by Lemma 14.1.3(4), it must occur to the left of any negative power of t. So $w = a^{\eta}tw'$ for some $\eta \in \{-1, 0, 1\}$ and some word w'. But then

$$w' \equiv t^{-1}a^{-\eta}w \equiv t^{-1}a^{-\eta}t^{-h}a^{k} \equiv t^{-(h+1)}a^{-2^{h}\eta+k}$$

and $|-2^{h}\eta + k| \leq 2^{h} + 2^{h} = 2^{h+1}$ and so, by induction, w' must have a prefix of the form t^{-1} or $a^{\varepsilon'}t^{-1}$, where ε' is the sign of $-2^{h}\eta + k$. But then w must contain a subword tt^{-1} or $ta^{\varepsilon'}t^{-1}$, which are not geodesic. So we see that w is descending.

It remains to show that every prefix w' of w satisfies the condition $(*)$ that $w' \equiv t^{-h'}a^{k'}$ where $0 \leq h' \leq h$ and $0 \leq k' \leq 2^{h'}$. That $w' \equiv t^{-h'}a^{\ell}$ with $0 \leq h' \leq h$ holds because w is descending and f does not descend below v_2 in T. We project the path taken by w in Γ to $\mathbb{Z}$ by sending $w' \equiv t^{-h'}a^{\ell}$ to $\pi(w') = 2^{h-h'}\ell$. Then a prefix w' satisfies $(*)$ if and only if its projection satisfies $0 \leq \pi(w') \leq 2^{h}$. Note that $\pi(w')$ is uniquely defined by $w't^{h'-h} \equiv t^{-h}a^{\pi(w')}$ where $w' \equiv t^{-h'}a^{\ell}$. So if two prefixes w' and w'' have the same projection then $(w')^{-1}w'' \equiv t^{h'-h''}$ where h' and w'' are the heights of w' and w''. Hence, by Lemma 14.1.3(5), if $\pi(w') = \pi(w'')$ then $w'' = w't^{h'-h''}$. Let $w_0, w_1, \ldots, w_n$ be the sequence of prefixes of w ordered by length. Then the projected path $\pi(w_0), \pi(w_1), \ldots, \pi(w_n)$ begins at 0 and ends at k with $\pi(w_{i+1}) - \pi(w_i)$ equal to 0 or $\pm 2^{h'}$ with $0 \leq h' \leq h$. Hence if this path leaves the interval $[0, 2^h]$ then there must be some $i < j < \ell$ with $\pi(w_i) = \pi(w_\ell) \in \{0, 2^h\}$ but with $\pi(w_j) \neq \pi(w_i)$, contradicting $w_\ell = w_i t^m$. $\square$

Lemma 14.1.8. *Let $h \geq 1$, let $0 \leq k \leq 2^h$ and let $\varepsilon = \pm 1$. The following statements describe precisely which initial letters a geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ may have.*

1. *If $k < \left(\frac{2}{3}\right) 2^h$ then any geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ has the form $w = t^{-1} w'$.*
2. *If $\left(\frac{2}{3}\right) 2^h < k < \left(\frac{5}{6}\right) 2^h$ then any geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ has the form $w = t^{-1} w'$ or $w = a^\varepsilon w''$ and there exist geodesics of both forms.*
3. *If $k > \left(\frac{5}{6}\right) 2^h$ then any geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ has the form $w = a^\varepsilon w'$.*

Proof. The proof is by induction on h. If $h = 1$ then $k \in \{0, 1, 2\}$. If $k = 0$ then $k < \frac{4}{3} = \left(\frac{2}{3}\right) 2^h$ and the only geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ is t^{-1}. If $k = 1$ then $k < \frac{4}{3} = \left(\frac{2}{3}\right) 2^h$ and the only geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ is $t^{-1} a^\varepsilon$. If $k = 2$ then $k > \frac{5}{3} = \left(\frac{5}{6}\right) 2^h$ and the only geodesic word $w \equiv t^{-h} a^{\varepsilon k}$ is $a^\varepsilon t^{-1}$. So, in all cases, the lemma holds for $h = 1$.

If $h = 2$ then $k \in \{0, 1, 2, 3, 4\}$. By Lemma 14.1.7, we need only consider descending words whose prefixes are equivalent to

$$t^{-h'} a^{k' \varepsilon}$$

for $h' \in \{0, 1, 2\}$ and $k' \in \{0, 1, \ldots, 2^{h'}\}$. We may also exclude words with backtracks and, by Lemma 14.1.3(2), those containing $t^{-1} a^\ell$ with $|\ell| \geq 2$ and, by Lemma 14.1.3(3), those containing $a^\varepsilon t^{-1} a^{-\varepsilon}$. Then the list of all possible geodesics is

$$t^{-2} a^{\varepsilon \ell}$$

$$t^{-1} a^\varepsilon t^{-1} a^{\varepsilon \ell} \equiv t^{-2} a^{\varepsilon (2+\ell)}$$

$$a^\varepsilon t^{-2} a^{-\varepsilon \ell} \equiv t^{-2} a^{\varepsilon (4-\ell)}$$

with $\ell \in \{0,1\}$. Only one pair of the words in this list, namely $a^{\varepsilon}t^{-2}a^{-\varepsilon}$ and $t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}$, are equivalent and they have the same length. Hence the list is exactly the list of all geodesics equivalent to $t^{-2}a^{\varepsilon k}$ with $k \in \{0,1,2,3,4\}$. Now, if $k < \left(\frac{2}{3}\right)2^h$ then $k \in \{0,1,2\}$. All geodesics equivalent to $t^{-2}a^{\varepsilon k}$ with $k \in \{0,1,2\}$ are of the first two forms in the list which have initial letter t^{-1}. The only k with $\left(\frac{2}{3}\right)2^h < k < \left(\frac{5}{6}\right)2^h$ is $k = 3$ and $t^{-2}a^{\varepsilon 3}$ is equivalent, with ℓ set to 1, to both the second form and the third form, which have initial letter t^{-1} and a^{ε}. Finally, if $k > \left(\frac{5}{6}\right)2^h$ then $k = 4$ which is equivalent only to the last form in the list with $\ell = 0$ and this form has initial letter a^{ε}. So we see that the lemma holds for $h = 2$. Going forward we assume that $h > 2$.

Suppose $k < \left(\frac{2}{3}\right)2^h$. To show that a geodesic $w \equiv t^{-h}a^{\varepsilon k}$ has initial letter t^{-1} it suffices, by Lemma 14.1.7, to rule out the possibility that w has the form $a^{\varepsilon}t^{-1}w'$. If that were the case then, by Lemma 14.1.6, the first letter of w' would be either t^{-1} or $a^{-\varepsilon}$ and we would have $w' \equiv ta^{-\varepsilon}t^{-h}a^{\varepsilon k} \equiv t^{-(h-1)}a^{-\varepsilon(2^h-k)}$. If $2^h - k > 2^{h-1}$ then, by Lemma 14.1.6, the first letter of w' is not t^{-1} and so would have to be $a^{-\varepsilon}$. But $a^{\varepsilon}t^{-1}a^{-\varepsilon}$ is not geodesic, by Lemma 14.1.3(3). So we have $2^{h-1} \geq 2^h - k > \left(\frac{2}{3}\right)2^{h-1}$. Then, applying (2) and (3) inductively to $h-1$ and $2^h - k$ and $-\varepsilon$, we see that w' is equivalent to a geodesic of the form $a^{-\varepsilon}w''$. But then $a^{\varepsilon}t^{-1}a^{-\varepsilon}w''$ is geodesic and this cannot be by Lemma 14.1.3(3).

Suppose $k > \left(\frac{5}{6}\right)2^h$. Let w be a geodesic with $w \equiv t^{-h}a^{\varepsilon k}$. Suppose the initial letter of w is not a^{ε}. Then, by Lemma 14.1.7, we have $w = t^{-1}a^{\varepsilon}w'$ for some w'. Hence $w' \equiv a^{-\varepsilon}tt^{-h}a^{\varepsilon k} \equiv t^{-(h-1)}a^{\varepsilon(k-2^{h-1})}$ with $k - 2^{h-1} \leq 2^{h-1}$ and $k - 2^{h-1} > \left(\frac{5}{6}\right)2^h - 2^{h-1} = \left(\frac{2}{3}\right)2^{h-1}$. So, inductively applying (2) and (3), we see that w' can be replaced by a geodesic of the form $a^{\varepsilon}w''$. But then $t^{-1}a^{\varepsilon}a^{\varepsilon}w''$ is a geodesic, which is a contradiction.

Suppose $\left(\frac{2}{3}\right)2^h < k < \left(\frac{5}{6}\right)2^h$. That any geodesic $w \equiv t^{-h}a^{\varepsilon k}$ has the form $w = t^{-1}w'$ or $w = a^{\varepsilon}w''$ follows from Lemma 14.1.7. Consider first the case where we have a geodesic of the form $w = t^{-1}w'$. Then, by Lemma 14.1.7, the initial letter of w' is either t^{-1} or a^{ε}. But $w' \equiv tw \equiv t^{-(h-1)}a^{\varepsilon k}$ with $k > \left(\frac{2}{3}\right)2^h > 2^{h-1}$ and so, by Lemma 14.1.6, the initial letter of w' is a^{ε}. So $w' = a^{\varepsilon}u'$ with $u' \equiv a^{-\varepsilon}w' \equiv t^{-(h-1)}a^{\varepsilon(k-2^{h-1})}$ and $k - 2^{h-1} < \left(\frac{5}{6}\right)2^h - 2^{h-1} = \left(\frac{2}{3}\right)2^{h-1}$. So, by induction, we have $u' = t^{-1}x'$ with $x' \equiv t^{-(h-2)}a^{\varepsilon(k-2^{h-1})}$ and $k - 2^{h-1} > \left(\frac{2}{3}\right)2^h - 2^{h-1} = \left(\frac{2}{3}\right)2^{h-2}$. Then, either by induction if $k \leq 2^{k-2}$ or otherwise by Lemma 14.1.6, we have that x' is equivalent to a geodesic of the form $a^{\varepsilon}y'$. Thus we have a geodesic $t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}y' \equiv t^{-h}a^{\varepsilon k}$. But $a^{\varepsilon}t^{-2}a^{-\varepsilon}$ has the same length as $t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}$ and $a^{\varepsilon}t^{-2}a^{-\varepsilon} \equiv t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}$ and so $a^{\varepsilon}t^{-2}a^{-\varepsilon}y'$ is a geodesic with $a^{\varepsilon}t^{-2}a^{-\varepsilon}y' \equiv t^{-h}a^{\varepsilon k}$. Now, consider the case where we have a geodesic of the form $w = a^{\varepsilon}w''$. By Lemma 14.1.7, the next two letters of w'' are either t^{-2} or $t^{-1}a^{-\varepsilon}$ but $a^{\varepsilon}t^{-1}a^{-\varepsilon}$ is not geodesic, by Lemma 14.1.3(3), so we must have $w = a^{\varepsilon}t^{-2}x''$. Then $x'' \equiv t^2a^{-\varepsilon}t^{-h}a^{\varepsilon k} \equiv t^{-(h-2)}a^{-\varepsilon(2^h-k)}$ with $2^h - k > 2^h - \left(\frac{5}{6}\right)2^h = \left(\frac{2}{3}\right)2^{h-2}$. Then, either by induction if $2^h - k \leq 2^{k-2}$ or otherwise by Lemma 14.1.6, we have that x'' is equivalent to a geodesic of the form $a^{-\varepsilon}y''$. So we have a geodesic $a^{\varepsilon}t^{-2}a^{-\varepsilon}y'' \equiv t^{-h}a^k$ and we may replace $a^{\varepsilon}t^{-2}a^{-\varepsilon}$ with $t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}$ to obtain an equivalent geodesic $t^{-1}a^{\varepsilon}t^{-1}a^{\varepsilon}y''$. $\square$

Lemma 14.1.9. *Let w be a geodesic with $w \equiv a^k$. Then $w = xa^{\ell}y$ such that the following conditions hold.*

1. *x is ascending and does not have terminal letter $a^{\pm 1}$.*
2. *y is descending and does not have initial letter $a^{\pm 1}$.*
3. *ℓ has the same sign as k.*
4. *$-\left(\frac{2}{3}\right)2^h < |k| - 2^h|\ell| < \left(\frac{5}{3}\right)2^h$ where $0 \leq h = \mu(x) = -\mu(y)$.*

Proof. By Lemma 14.1.5, we have $w = xa^{\ell}y$ with x ascending and y descending. Choosing x and y so as to maximize the length of the a^{ℓ} part ensures that the terminal letter of x is t and the initial letter of y_i is t^{-1}. Since $\mu(w) = 0$ and $\mu(a^\ell) = 0$, we have $\mu(x) + \mu(y) = \mu(w) = 0$.

If $h = 0$ then $w = a^k = a^\ell$ and so the remaining conditions hold. So let $h \geq 1$. Then, by Lemma 14.1.3(1), we have $|\ell| \geq 2$. Let ε be the sign of ℓ. The first letter of y is t^{-1}. Hence, by Lemma 14.1.6, we have $y \equiv t^{-h}a^{\varepsilon m}$ for some $|m| \leq 2^h$. Then, by Lemma 14.1.8, we have $|m| < \left(\frac{5}{6}\right)2^h$. Now $a^\varepsilon y$ is a subword of w and $a^\varepsilon y \equiv a^\varepsilon t^{-h} a^{\varepsilon m} \equiv t^{-h}a^{\varepsilon(m+2^h)}$ with $m + 2^h > 2^h - \left(\frac{5}{6}\right)2^h > 0$. Then, by Lemma 14.1.8, we see that $m + 2^h > \left(\frac{2}{3}\right)2^h$. Hence we have $-\left(\frac{1}{3}\right)2^h < m < \left(\frac{5}{6}\right)2^h$. Applying the exact same arguments to the subwords x^{-1} and $a^{-\varepsilon}x^{-1}$ of $w^{-1} = y^{-1}a^{-\ell}x^{-1}$ we see that $x^{-1} \equiv t^{-h}a^{-\varepsilon n}$ with $-\left(\frac{1}{3}\right)2^h < n < \left(\frac{5}{6}\right)2^h$. Hence

$$w \equiv a^{\varepsilon n} t^h a^\ell t^{-h} a^{\varepsilon m} \equiv a^{\varepsilon(n+m+2^h|\ell|)}$$

and so $k = \varepsilon(n + m + 2^h|\ell|)$ which gives

$$-\left(\frac{2}{3}\right)2^h + 2^h|\ell| < \varepsilon k < \left(\frac{5}{3}\right)2^h + 2^h|\ell|$$

and so, as $|\ell| \geq 2$ we see that $\varepsilon k > 0$. Then k has the same sign as ℓ so $\varepsilon k = |k|$ and we have

$$-\left(\frac{2}{3}\right)2^h < |k| - 2^h|\ell| < \left(\frac{5}{3}\right)2^h$$

as required. □

14.2 Isometric cycles in $\mathrm{BS}(1,2)$

Lemma 14.2.1. *Let $f\colon C \to \Gamma$ be an isometrically embedded cycle of length $|C| > 5$. Then f is described by a word of the form*

$$w = a^{\varepsilon_1} w_1 a^{\varepsilon_2} w_2$$

with $\varepsilon_1, \varepsilon_2 \in \{\pm 1\}$ such that w_i satisfies the following properties for each i.

1. *w_i is geodesic with initial letter t and terminal letter t^{-1}.*

2. $w_i \equiv a^{2k_i}$ *with* $|k_1 + k_2| \leq 1$.

Proof. Let $f\colon C \to \Gamma$ be an isometrically embedded cycle. Let $\bar{f}$ be the projection of f to the Bass-Serre tree T. By Lemma 14.1.2, there is a unique vertex v of minimal height of $\bar{f}$. Then f must contain a subpath of the form $t^{-1}a^k t$ with the a^k part mapping to v under $\bar{f}$. Since f is an isometric embedding and $|C| > 5$, Lemma 14.1.3(2) implies that $|k| = 1$. So $t^{-1}a^k t = t^{-1}a^{\pm 1}t$ corresponds to a nonbacktracking path in T. It follows, since f is a closed path and v is a cutpoint of T, that f contains another subpath of the form $t^{-1}a^{\pm 1}t$ such that $\bar{f}$ sends $a^{\pm 1}$ to v. That is f has the form

$$w = a^{\varepsilon_1} t u_1 t^{-1} a^{\varepsilon_2} t u_2 t^{-1}$$

with $\varepsilon_1, \varepsilon_2 \in \{\pm 1\}$, such that $u_i \equiv a^{k_i}$ for some $k_1, k_2 \in Z$. Let $w_i = t u_i t^{-1}$. Then $w_i \equiv a^{2k_i}$ so we have

$$\varepsilon_1 + 2k_1 + \varepsilon_2 + 2k_2 = 0$$

since w is trivial in G, and this implies $|k_1 - k_2| \leq 1$. Also we have

$$a^{\varepsilon_i} w_i a^{\varepsilon_{i+1}} \equiv w_i a^{\varepsilon_i} a^{\varepsilon_{i+1}}$$

which shows that $a^{\varepsilon_i} w_i a^{\varepsilon_{i+1}}$ has the same length and endpoints as $w_i a^{\varepsilon_i} a^{\varepsilon_{i+1}}$. But $w_i a^{\varepsilon_i} a^{\varepsilon_{i+1}}$ cannot be geodesic since it either backtracks or contains $t^{-1}a^k$ with $|k| = 2$ and so $a^{\varepsilon_i} w_i a^{\varepsilon_{i+1}}$ is not geodesic either. Hence, since f is an isometric embedding, the complementary path w_{i+1} is geodesic. □

Theorem 14.2.2. *The standard Cayley graph of of* BS(1, 2) *is shortcut.*

Proof. Let Γ be the standard Cayley graph of BS(1, 2). We will show that there are no isometrically embedded cycles $f\colon C \to \Gamma$ of length $|C| > 5$. For the sake of deriving a contradiction, suppose f is such a cycle. Then, by

Lemma 14.2.1, we have

$$w = a^{\varepsilon_1} w_1 a^{\varepsilon_2} w_2$$

such that w_i is geodesic with initial letter t and terminal letter t^{-1} and $w_i \equiv a^{2k_i}$ with $|k_1 + k_2| \leq 1$. Then, by 14.1.9, we have

$$w_i = x_i a^{\ell_i} y_i$$

where x_i is ascending and has initial and terminal letter t, where y_i is descending and has initial and terminal letter t^{-1}, where ℓ_i has the same sign as k_i and where

$$-\left(\frac{2}{3}\right)2^{h_i} < |k_i| - 2^{h_i}|\ell_i| < \left(\frac{5}{3}\right)2^{h_i}$$

with $0 \leq h_i = \mu(x_i) = -\mu(y_i)$. Since x_i has terminal letter t and y_i has initial letter t^{-1}, we have $h_i \geq 1$ and, by Lemma 14.1.3(1), we have $|\ell_i| \geq 2$.

We may assume that $h_1 \leq h_2$ since otherwise we may replace w with a cyclic permutation. We must have $|k_i| \geq 1$ since otherwise $w_i = 1$ or $w = a^{\pm 1}$, which do not start with t. Hence, as $|k_1 + k_2| \leq 1$ we must have that k_1 and k_2 have opposite signs. Since there is an automorphism of $\mathrm{BS}(1,2)$ fixing t and mapping $a \mapsto a^{-1}$ we may assume that $k_1 > 0$ and $k_2 < 0$. Then $\ell_1 > 0$ and $\ell_2 < 0$.

Let $p_i \in C$ be the midpoint of the subpath a^{ℓ_i} of w. Abusing notation we write the two segments of C between p_1 and p_2 as $a^{\ell_1/2} y_1 a^{\varepsilon_2} x_2 a^{\ell_2/2}$ and $a^{\ell_2/2} y_2 a^{\varepsilon_1} x_1 a^{\ell_1/2}$ which may be thought of as combinatorial paths in the barycentric subdivision of Γ. Since f is an isometric embedding, one of these two paths must be geodesic in Γ. We have

$$y_1 a^{\varepsilon_2} x_2 a^{-1} \equiv y_1 a^{\varepsilon_2} a^{-2^{h_2}} x_2 \equiv a^{-2^{h_2 - h_1}} y_1 a^{\varepsilon_2} x_2$$

$$a^{-1} y_2 a^{\varepsilon_1} x_1 \equiv y_2 a^{-2^{h_2}} a^{\varepsilon_1} x_1 \equiv y_2 a^{\varepsilon_1} x_1 a^{-2^{h_2 - h_1}}$$

and so

$$a^{\ell_1/2}y_1a^{\varepsilon_2}x_2a^{\ell_2/2} \equiv a^{\ell_1/2-2^{h_2-h_1}}y_1a^{\varepsilon_2}x_2a^{\ell_2/2+1}$$

$$a^{\ell_2/2}y_2a^{\varepsilon_1}x_1a^{\ell_1/2} \equiv a^{\ell_2/2+1}y_2a^{\varepsilon_1}x_1a^{\ell_1/2-2^{h_2-h_1}}$$

which, by geodesicity of one of these paths, imply that either

$$|\ell_1/2| + |\ell_2/2| \leq |\ell_1/2 - 2^{h_2-h_1}| + |\ell_2/2+1|$$

or

$$|\ell_2/2| + |\ell_1/2| \leq |\ell_2/2+1| + |\ell_1/2 - 2^{h_2-h_1}|$$

but these two inequalities are equivalent so they must both hold. Then, using $\ell_1 \geq 0$ and $\ell_2 \leq -2$, we obtain $\ell_1 + 2 \leq |\ell_1 - 2^{h_2-h_1+1}|$. Since $\ell_1 \geq 0$, this inequality may only hold if $2^{h_2-h_1+1} > \ell_1$ and so we have $\ell_1 + 1 \leq 2^{h_2-h_1}$. Now $\ell_1 \geq 2$ and so we have $h_2 \geq h_1 + 2$. The following computation makes use of various inequalities which have been established thus far in this proof.

$$\begin{aligned}
\left(\frac{2}{3}\right)2^{h_2-2} + 2^{h_2} &\geq \left(\frac{2}{3}\right)2^{h_1} + 2^{h_2-h_1}\cdot 2^{h_1} \\
&\geq \left(\frac{2}{3}\right)2^{h_1} + (\ell_1+1)2^{h_1} \\
&= \left(\frac{5}{3}\right)2^{h_1} + 2^{h_1}|\ell_1| \\
&> |k_1| \geq |k_2| - 1 > -\left(\frac{2}{3}\right)2^{h_2} + 2^{h_2}|\ell_2| - 1
\end{aligned}$$

So, as $|\ell_2| \geq 2$, we have

$$\left(\frac{2}{3}\right)2^{h_2-2} + 2^{h_2} > -\left(\frac{2}{3}\right)2^{h_2} + 2^{h_2}\cdot 2 - 1$$

which we manipulate to obtain the equivalent inequality $2^{h_2} < 6$. Then $h_2 \leq 2$, which implies that $h_1 = 0$. But this is a contradiction as we established above that $h_1 \geq 1$. $\square$

14.3 A Cayley graph of $\mathrm{BS}(1,2)$ that is not shortcut

We now turn our attention to a different generating set of $\mathrm{BS}(1,2)$. Let Γ be the Cayley graph of $\mathrm{BS}(1,2)$ with the generating set $\{a, t, \tau\}$ where $\tau = t^2$.

Lemma 14.3.1. *Let $k \geq 1$ and let $0 \leq z_{\max} \leq k$. Suppose*

$$\sum_{z=-m}^{z_{\max}} \alpha_z 2^z = 2^k \pm 1$$

where $\alpha_z \in \mathbb{Z}$ and $m \geq 0$.

- *If $z_{\max} = 0$ then $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} - 1$*
- *If $z_{\max} = 1$ then $\sum_z |\alpha_z| \geq 2^{k-z_{\max}}$*
- *If $z_{\max} \geq 2$ then $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} + 1$*

Proof. If $z_{\max} = 0$ then

$$2^k - 1 \leq 2^k \pm 1 = \left|\sum_{z=-m}^{z_{\max}} \alpha_z 2^z\right| \leq \sum_{z=-m}^{z_{\max}} |\alpha_z| 2^{z_{\max}} = \sum_{z=-m}^{z_{\max}} |\alpha_z|$$

and so we have $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} - 1$.

Suppose $z_{\max} = 1$ and $\sum_z |\alpha_z| < 2^{k-z_{\max}}$. Then $\sum_z |\alpha_z| \leq 2^{k-z_{\max}} - 1$ and so

$$\left|\sum_{z=-m}^{z_{\max}} \alpha_z 2^z\right| \leq \sum_{z=-m}^{z_{\max}} |\alpha_z| 2^{z_{\max}} \leq (2^{k-z_{\max}} - 1) \cdot 2^{z_{\max}} = 2^k - 2$$

which is a contradiction. So we have $\sum_z |\alpha_z| \geq 2^{k-z_{\max}}$.

Now, suppose $z_{\max} \geq 2$. Among all $m \geq 0$ and $(\alpha_z)_z$ that satisfy

$$\sum_{z=-m}^{z_{\max}} \alpha_z 2^z = 2^k \pm 1$$

choose an $m \geq 0$ and $(\alpha_z)_z$ that minimizes $\sum_z |\alpha_z|$. We will show that $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} + 1$. We claim that for $z < z_{\max}$, we have $|\alpha_z| \leq 1$. Indeed, if $|\alpha_z| \geq 2$, then we can replace α_z with $\alpha_z - \varepsilon 2$ and α_{z+1} with $\alpha_{z+1} + \varepsilon$, where ε is the sign of α_z. This reduces $\sum_z |\alpha_z|$ while preserving $\sum_z \alpha_z 2^z$ and so contradicts minimality of m and $(\alpha_z)_z$. Since $2^k \pm 1$ is not even, there must be some $\alpha_z \neq 0$ with $z \leq 0$. So if $\alpha_{z_{\max}} \geq 2^{k-z_{\max}}$ then $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} + 1$. So we may assume that $\alpha_{z_{\max}} < 2^{k-z_{\max}}$. Say $\alpha_{z_{\max}} = 2^{k-z_{\max}} - \ell$ with $\ell \geq 1$. Then $2^k \pm 1 - \alpha_{z_{\max}} 2^{z_{\max}} \geq \ell 2^{z_{\max}} - 1$ and so $\sum_{z=-m}^{z_{\max}-1} \alpha_z 2^z \geq \ell 2^{z_{\max}} - 1$. But, since $\alpha_z \leq 1$ for $z < z_{\max}$ we have

$$\sum_{z=-m}^{z_{\max}-1} \alpha_z 2^z \leq \sum_{z=-m}^{z_{\max}-1} 2^z \leq 2^{z_{\max}} - 2^{-m}$$

with equality only if $\alpha_z = 1$ for $-m \leq z < z_{\max}$. It follows that $\ell = 1$ and that $\alpha_z = 1$ for $0 \leq z < z_{\max}$. Hence $\sum_z |\alpha_z| \geq 2^{k-z_{\max}} - 1 + z_{\max} \geq 2^{k-z_{\max}} + 1$ as required. $\square$

Lemma 14.3.2. *Let k and ℓ be nonnegative integers with $\ell \leq k$ and $k \geq 2$. Then the word*

$$w = \tau^\ell a \tau^{-k} a^{\pm 1} \tau^{k-\ell}$$

describes a geodesic in Γ.

Proof. Consider the bijection $\varphi\colon \mathrm{BS}(1,2) \to \mathbb{Z}\left[\frac{1}{2}\right] \times \mathbb{Z}$ described near the beginning of Chapter 14. Then, under this bijection, w describes a path from $(0, k-\ell)$ to $(4^k \pm 1, k-\ell)$. Consider any path $(r_j, z_j)_{j=0}^m$ from $(0, k-\ell)$ to $(4^k \pm 1, k-\ell)$ following edges of the Cayley graph. It will suffice to show that $m \geq 2k+2$, since $2k+2 = |w|$. For each j, either

$$(r_{j+1}, z_{j+1}) - (r_j, z_j) = (\pm 2^{z_j}, 0)$$

which we call a *horizontal step* or

$$(r_{j+1}, z_{j+1}) - (r_j, z_j) = \begin{cases} (0, \pm 1) \\ (0, \pm 2) \end{cases}$$

which we call a *vertical step.* Since $4^k \pm 1$ is not a power of 2, there must be at least two horizontal steps in $(r_j, z_j)_j$. Moreover, since $4^k \pm 1$ is not even, we must have $z_j \leq 0$ for some j. Let $z_{\max} = \max_j z_j$. Extend $(r_j, z_j)_{j=0}^m$ to a continous map $f\colon C \to \mathbb{R}^2$ where $|C| = m$. The projection $C \to \mathbb{R}$ of this map onto the second component captures the vertical behaviour of f. We collapse any edges of C that map to points under $C \to \mathbb{R}$ to obtain a map $\bar{f}\colon \bar{C} \to \mathbb{R}$ where $|\bar{C}|$ is equal to the number of vertical steps of $(r_j, z_j)_j$. The map $\bar{f}$ is 2-Lipschitz and $0, (k-\ell), z_{\max} \in \bar{f}(\bar{C})$ with $z_{\max} \geq k - \ell$ and $0 \leq k - \ell$. Then

$$\left|\bar{f}^{-1}\big([k-\ell, \infty)\big)\right| \geq 2 \cdot \frac{z_{\max} - (k-\ell)}{2}$$

and

$$\left|\bar{f}^{-1}\big((-\infty, k-\ell]\big)\right| \geq 2 \cdot \frac{k-\ell}{2}$$

where $\left|\bar{f}^{-1}(I)\right|$ is the sum of the lengths of all maximal segments of $f^{-1}(I)$. Hence $|\bar{C}| \geq z_{\max}$ and so $(r_j, z_j)_j$ takes at least $z_{\max}$ vertical steps. Then, since there must also be at least two horizontal steps, we see that if $z_{\max} \geq 2k$ then $m \geq 2k + 2$. So we may assume that $z_{\max} < 2k$.

We split into three cases: $z_{\max} = 0$, $z_{\max} = 1$ and $z_{\max} \geq 2$. If $z_{\max} = 0$ then, by Lemma 14.3.1, there are at least $2^{2k} - 1$ horizontal steps. So it suffices to show that $2^{2k} - 1 \geq 2k + 2$ but, since $k \geq 2$, this follows from the fact that $2^x \geq x + 3$ for all $x \geq 4$.

If $z_{\max} = 1$ then, by Lemma 14.3.1, there are at least 2^{2k-1} horizontal steps. There will also be at least $z_{\max} = 1$ vertical step. But 1 vertical step cannot give a closed path and so there are at least 2 vertical steps. So it suffices to show that $2^{2k-1} + 2 \geq 2k + 2$ but, since $k \geq 2$, this follows from the fact that $2^x \geq x + 1$ for all $x \geq 1$.

If $z_{\max} \geq 2$ then, by Lemma 14.3.1, there are at least $2^{2k - z_{\max}} + 1$ horizontal steps. There are also at least $z_{\max}$ vertical steps. So it suffices to show that $2^{2k - z_{\max}} + 1 + z_{\max} \geq 2k + 2$ but, since $z_{\max} < 2k$, this follows from the fact that $2^x \geq x + 1$ for all $x \geq 1$. $\square$

Lemma 14.3.3. *The word*

$$w = a\tau^k a\tau^{-k} a^{-1}\tau^k a^{-1}\tau^{-k}$$

describes an isometric cycle in Γ *for all* $k \geq 1$.

Proof. The word w has length $4k+4$. After possible inversion and/or the application of the automorphism of $\mathrm{BS}(1,2)$ fixing t and sending $a \mapsto a^{-1}$, the cyclic subwords of w of length $2k+2$ are all of the form in Lemma 14.3.2 and so are geodesic. Hence, by Proposition 10.0.3, the word w describes an isometric cycle in Γ. $\square$

Then we have the following theorem and we see that the shortcut property for a Cayley graph is not invariant under a change of generating set.

Theorem 14.3.4. *Let* Γ *be the Cayley graph of* $\mathrm{BS}(1,2)$ *with generating set* $\{a, t, \tau\}$ *where* $\tau = t^2$. *Then* Γ *is not shortcut.*

References

[1] M. Aigner and G. M. Ziegler. *Proofs from THE BOOK*. Springer Berlin Heidelberg, sixth edition, 2018.

[2] S. P. Avann. Metric ternary distributive semi-lattices. *Proc. Amer. Math. Soc.*, 12(3):407–414, 1961.

[3] H.-J. Bandelt. Hereditary modular graphs. *Combinatorica*, 8(2):149–157, 1988.

[4] A. Björner. Posets, regular CW complexes and Bruhat order. *European J. Combin.*, 5(1):7–16, 1984.

[5] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.

[6] J. W. Cannon. Shrinking cell-like decompositions of manifolds. Codimension three. *Ann. of Math. (2)*, 110(1):83–112, 1979.

[7] M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Math.*, 217(1-3):101–113, 2000. Formal power series and algebraic combinatorics (Vienna, 1997).

[8] V. Chepoi. Graphs of some CAT(0) complexes. *Adv. in Appl. Math.*, 24(2):125–179, 2000.

[9] V. Chepoi and D. Osajda. Dismantlability of weakly systolic complexes and applications. *Trans. Amer. Math. Soc.*, 367(2):1247–1272, 2015.

[10] Y. Civan and E. Yalçın. Linear colorings of simplicial complexes and collapsing. *J. Combin. Theory Ser. A*, 114(7):1315–1331, 2007.

[11] M. Dehn. *Papers on group theory and topology*. Springer-Verlag, New York, 1987. Translated from the German and with introductions and an appendix by John Stillwell, With an appendix by Otto Schreier.

[12] T. Elsner. Flats and the flat torus theorem in systolic spaces. *Geom. Topol.*, 13(2):661–698, 2009.

[13] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134(1):90–145, 1998.

[14] M. H. Freedman. The topology of four-dimensional manifolds. *J. Differential Geom.*, 17(3):357–453, 1982.

[15] V. Gerasimov. Fixed-point-free actions on cubings. *Siberian Adv. Math.*, 8(3):26–58, 1998.

[16] S. M. Gersten. Isoperimetric and isodiametric functions of finite presentations. In *Geometric Group Theory: Proceedings of the Symposium held in Sussex, Volume 1*, volume 181 of *London Mathematical Society Lecture Notes Series*, pages 79–96. Cambridge University Press, New York, 1991.

[17] S. M. Gersten. Dehn functions and l_1-norms of finite presentations. In *Algorithms and classification in combinatorial group theory (Berkeley, CA, 1989)*, volume 23 of *Math. Sci. Res. Inst. Publ.*, pages 195–224. Springer, New York, 1992.

[18] M. Greendlinger. Dehn's algorithm for the word problem. *Comm. Pure Appl. Math.*, 13(1):67–83, 1960.

[19] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[20] F. Haglund. Complexes simpliciaux hyperboliques de grande dimension. Preprint, 2003.

[21] F. Haglund and F. Paulin. Simplicité de groupes d'automorphismes d'espaces à courbure négative. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 181–248. Geom. Topol. Publ., Coventry, 1998.

[22] F. Haglund and J. Świątkowski. Separating quasi-convex subgroups in 7-systolic groups. *Groups Geom. Dyn.*, 2(2):223–244, 2008.

[23] F. Haglund and D. T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

[24] A. Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[25] N. Hoda. Quadric complexes. *Preprint, arXiv:1711.05844*, 2017.

[26] G. C. Hruska and D. T. Wise. Finiteness properties of cubulated groups. *Compos. Math.*, 150(3):453–506, 2014.

[27] J. F. P. Hudson. *Piecewise linear topology.* University of Chicago Lecture Notes prepared with the assistance of J. L. Shaneson and J. Lees. W. A. Benjamin, Inc., New York-Amsterdam, 1969.

[28] T. Januszkiewicz and J. Świątkowski. Simplicial nonpositive curvature. *Publ. Math. Inst. Hautes Études Sci.*, 104(1):1–85, 2006.

[29] T. Januszkiewicz and J. Świątkowski. Hyperbolic Coxeter groups of large dimension. *Comment. Math. Helv.*, 78(3):555–583, 2003.

[30] C. Jordan. Sur les assemblages de lignes. *J. Reine Angew. Math.*, 70:185–190, 1869.

[31] G. Kalai. Many triangulated spheres. *Discrete Comput. Geom.*, 3(1):1–14, 1988.

[32] S. Klavžar and H. M. Mulder. Median graphs: characterizations, location theory and related structures. *J. Combin. Math. Combin. Comput.*, 30:103–127, 1999.

[33] J. Matoušek. LC reductions yield isomorphic simplicial complexes. *Contrib. Discrete Math.*, 3(2):37–39, 2008.

[34] L. Nebeský. Median graphs. *Comment. Math. Univ. Carolinae*, 12(2):317–325, 1971.

[35] M. H. A. Newman. The engulfing theorem for topological manifolds. *Ann. of Math. (2)*, 84:555–571, 1966.

[36] G. A. Niblo and L. D. Reeves. Groups acting on CAT(0) cube complexes. *Geom. Topol.*, 1(1):1–7, 1997.

[37] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.

[38] G. A. Niblo and L. D. Reeves. Coxeter groups act on CAT(0) cube complexes. *J. Group Theory*, 6(3):399–413, 2003.

[39] D. Osajda. Connectedness at infinity of systolic complexes and groups. *Groups Geom. Dyn.*, 1(2):183–203, 2007.

[40] D. Osajda and P. Przytycki. Boundaries of systolic groups. *Geom. Topol.*, 13(5):2807–2880, 2009.

[41] G. Perelman. The entropy formula for the Ricci flow and its geometric applications, 2002.

[42] G. Perelman. Finite extinction time for the solutions to the Ricci flow on certain three-manifolds, 2003.

[43] G. Perelman. Ricci flow with surgery on three-manifolds, 2003.

[44] M. Roller. *Poc sets, median algebras and group actions.* Habilitationsschrift, University of Regensburg, 1998.

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

[46] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36 of *London Math. Soc. Lecture Note Ser.*, pages 137–203. Cambridge Univ. Press, Cambridge-New York, 1979.

[47] H. Seifert and W. Threlfall. *Seifert and Threlfall: a textbook of topology*, volume 89 of *Pure and Applied Mathematics*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], New York-London, 1980. Translated from the German edition of 1934 by Michael A. Goldman, With a preface by Joan S. Birman, With "Topology of 3-dimensional fibered spaces" by Seifert, Translated from the German by Wolfgang Heil.

[48] S. Smale. Generalized Poincaré's conjecture in dimensions greater than four. *Ann. of Math. (2)*, 74:391–406, 1961.

[49] V. P. Soltan and V. D. Chepoi. Conditions for invariance of set diameters under d-convexification in a graph. *Kibernetika (Kiev)*, 19(6):14–18, 1983.

[50] J. R. Stallings. Polyhedral homotopy-spheres. *Bull. Amer. Math. Soc.*, 66:485–488, 1960.

[51] D. T. Wise. Sixtolic complexes and their fundamental groups. Preprint, 2003.

[52] D. T. Wise. The structure of groups with a quasiconvex hierarchy, 2011. pp. 1–189. Submitted.

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

[54] E. C. Zeeman. The Poincaré conjecture for $n \geq 5$. In *Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute, 1961)*, pages 198–204. Prentice-Hall, Englewood Cliffs, N.J., 1962.