A General Small Cancellation Theory

Jonathan P. McCammond

September 16, 1999

Abstract

In this article a generalized version of small cancellation theory is developed which is applicable to specific types of high-dimensional simplicial complexes. The usual results on small cancellation groups are then shown to hold in this new setting with only slight modifications. For example, arbitrary dimensional versions of the Poincaré construction and the Cayley complex are described.

0.1 Main Theorems

In this article a generalized version of traditional small cancellation theory is developed which is applicable to specific types of high-dimensional simplicial complexes. The usual results on small cancellation groups are then shown to hold in this new setting with only slight modifications. The main results derived for this general small cancellation theory are summarized below in Theorem A. The notions of general relators, Cayley categories, and general small cancellation presentations are being introduced here, and will be defined in the course of the article.

Theorem A *If $G = \langle A | \mathcal{R} \rangle$ is a general small cancellation presentation with $\alpha \leq \frac{1}{12}$, then the word and conjugacy problems for G are decidable, the Cayley graph is constructible, the Cayley category of the presentation is contractible, and G is the direct limit of hyperbolic groups. If in addition, the presentation satisfies the hypotheses of Lemma 14.17, then every finite subgroup of G is a subgroup of the automorphism group of some general relator in $\mathcal{R}$.*

In a separate article, the general small cancellation theory developed here will be applied to the Burnside groups of sufficiently large exponent. The results obtained for the Burnside groups are similar in nature to the recent results of Ivanov ([8]) and Lysionok ([10]), although it appears that techniques described here will cover several additional cases. See [12] for details.

0.2 Key Concepts

Before beginning the full development of the theory, it seems advisable to provide a brief sketch of the key concepts used in the proof. To this end, the presentation below will avoid precise definitions in favor of rough descriptions which appeal to the intuition.

General Relators: The results obtained in this article are dependent on a type of structure called a general relator, which is introduced here for the first time. In traditional small cancellation theory the cyclically reduced relators can be viewed as a finite partition of the unit circle, with each edge labeled by a generator of the group or, alternatively, as the boundary of a unit disk with these properties. The particular generalization of small cancellation theory developed in this article is based on the idea of using "relators" whose boundaries are homotopically equivalent to the unit circle, and the general relator itself is a topological cone over its boundary. These general relators, which are perhaps best viewed as topological cones over solid tubes, contain a 1-skeleton which can be significantly more complicated than in the traditional theory, but whose local structure is less important than its global topology.

Representatives: Of particular importance in this context is the notion of the winding number of a loop in the boundary of a general relator. A winding number is definable in this situation because of the homotopic equivalence to the unit circle. A loop in the boundary of a general relator with winding number 1 is called a representative of the general relator. Using representatives, it is possible to define more or less traditional van Kampen diagrams over collections of general relators by requiring that the label of every 2-cell in the planar van Kampen diagram be the label of a representative loop in the boundary of some general relator.

General Presentations, Poincaré Constructions, and Cayley Categories: The group corresponding to a set of general relators is defined by forming a variation of the Poincaré construction and setting the group of the presentation equal to the fundamental group of the construction. The universal cover of this construction contains a 1-skeleton which is the Cayley graph of this group. As in the traditional theory, the universal cover of the Poincaré construction will contain multiple copies of a general relator attached to the Cayley graph by functions which agree on the 1-skeletons in their boundaries. The number of such multiplicities are governed by the size of the automorphism group of the particular general relator involved. If these multiple copies are eliminated through a suitable identification, then the resulting structure is called the Cayley category of the presentation.

General Small Cancellation Axioms: By placing sufficient restrictions on the general relators involved it is possible to mimic the traditional proof schemes of small cancellation theory. As an example, the general relators used are required to be thin in the sense that, given any point in a relator and an arbitrary representative of the relator, the distance from the point to the representative is small compared with the minimum length of a representative. Because of this fact, it is possible to speak in a loose way of the distance traveled around the boundary of a general relator. In terms of this distance-like function it is assumed that if one general relator contains more than a specified fraction of the boundary of the other general relator then the boundary of the first relator actually contains the entire second relator.

0.3 Overview

The article is divided into five parts. The first part develops a theory of structures iteratively built out of cones. These structures are a conical version of CW complexes. The second part considers only those conical structures which topologically resemble the 2-cells traditionally used to create Poincaré constructions from group presentations. These are called general relators. The third part of the article investigates constructions which are built out of general relators. The topics include extended versions of presentations, Poincaré constructions, and covering spaces. In Part IV, the focus is narrowed once again to consider the effect that suitably generalized small cancellation conditions have on the groups presented using general relators. A list of axioms for a general small cancellation theory is presented. Finally, in Part V, the axioms are used to prove the remaining results listed in Theorem A.

0.4 Acknowledgements

The author would like to thank John Rhodes and John Stallings for providing an opportunity to lecture on these results in their seminars at the University of California at Berkeley. Professor John Rhodes deserves a separate acknowledgement for the almost weekly conversations held over the past few years on the various issues surrounding Burnside groups and semigroups. Roger Alperin, Paul Brown, John Stallings and other members of Stallings' seminar helped initiate the author into the world of geometric group theory. Ken Brown provided a useful reference at a crucial point in the proof. Professor Ol'shanskii provided enormous help by pointing out several locations which needed additional argumentation and clarification. Lastly, the author would like to acknowledge the influence of the works of M. Gromov, S. Ivanov, A. Yu. Ol'shanskii, D. Quillen, and W. Thurston on the development of the ideas contained in this article.

Contents

List of Figures

Part I
Cones

In this part the theory of structures iteratively built out of cones is developed. The definitions of a conical CW complex, a cone complex, a circular complex, etc. are, as far as I am aware, being introduced here for the first time. The idea of representing these structures by categories and functors was specifically inspired by the work of Quillen ([17]). Many of these structures mimic the more traditional structures of algebraic topology in great detail, so that many readers may find it sufficient to read the definitions and the statements of the lemmas throughout part I. The details have been provided for the sake of completeness.

The contents of this part are divided into three sections. Section 1 briefly reviews the necessary topological preliminaries, before describing the construction of conical CW complexes. Sections 2 and 3 are dedicated to the special cases of these conical CW complexes which can be described by a partially ordered set or a category, respectively. These cases are the ones which will be used in the remaining parts of the article.

1 Conical CW Complexes

Traditional CW complexes are constructed inductively by attaching $(n+1)$-dimensional balls along their spherical boundaries to an already constructed n-skeleton. Conical CW complexes are a variation on this procedure in which topological cones are attached along their bases to an already constructed n-skeleton. As in the traditional case there is both an inductive construction and an internal description of the completed construction which complements the inductive version. The section concludes with some results regarding deformation retractions and coverings of such spaces.

1.1 Topological Preliminaries

The necessary results and definitions from algebraic topology are given below.

Quotient Spaces If X is a topological space, $f : X \to Y$ is a surjection and Y is a set, then there is a unique largest topology on Y for which f is continuous. This topology, called the quotient topology on Y induced by f, is defined by letting a subset $V \subset Y$ be an open set iff $f^{-1}(V)$ is open in X. When the set Y is given this topology the continuous function f is called a quotient map, and Y is called a quotient space. The word 'map' is used in this article both to refer to a continuous function between topological spaces as is done here, and to refer to a particular planar construction used in traditional small cancellation theory. The usage intended should be clear from context. Let X and Y be spaces with $A \subset X$ a closed subset. If f is a map from A to Y, then the space $Y \cup_f X$ is defined as the quotient space of the disjoint union of X and Y by

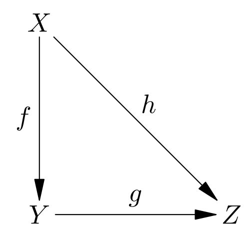

Figure 1: Quotient maps

the equivalence relation $\sim$ generated by the relation $a \sim f(a)$ for all $a \in A$. Lemmas 1.1, 1.2, and 1.3 are standard. Proofs can be found in reference works such as [2]. Lemma 1.1 is illustrated in Figure 1.

Lemma 1.1 *An onto map $f : X \to Y$ is a quotient map if and only if it is true that for all functions $g : Y \to Z$ (the composite $h = g \circ f$ is continuous $\Leftrightarrow$ the function g is continuous).*

Lemma 1.2 *Let $f : X \to Y$ be a quotient map. If Z is a space and $h : X \to Z$ is a map that is constant on each set $f^{-1}(y)$, for all $y \in Y$ then g induces a map $g : Y \to Z$ such that $g \circ f = h$.*

Lemma 1.3 *The canonical map $Y \to Y \cup_f X$ is an embedding onto a closed subspace. The canonical map $X \setminus A \to Y \cup_f X$ is an embedding onto an open subspace.*

Topological Cones A cone over a topological space X is the quotient space of $X \times [0, 1]$ obtained by identifying the subspace $X \times \{0\}$ to a point, together with the quotient map. The cone over X will be denoted $\text{Cone}(X)$. The quotient map $f : X \times [0, 1] \to \text{Cone}(X)$ is contained in the definition because the second coordinate provides a useful partitioning of $\text{Cone}(X)$ into layers. The point $f(x, 0)$ is called the vertex of the cone, the subset $f(X \times [0, 1))$ is called the interior of the cone and the subset $f(X \times \{1\})$ is called the base or the boundary of the cone. Since by Lemma 1.3 the subset $X \times (0, 1]$ is embedded in $\text{Cone}(X)$ by the map f, the base of the cone is canonically homeomorphic to X. When the original space X is viewed as identical with the base of $\text{Cone}(X)$, the cone over X is seen to be an extension of X. If a topological space C is given the structure of a topological cone $\text{Cone}(X)$ for some X then the base of C will be denoted ∂C. The Lemma 1.4 is a special case of Lemma 1.2.

Lemma 1.4 *Let X be a topological space, and let $f : X \times [0, 1] \to Cone(X)$ be the quotient map described above. If $g : X \times [0, 1] \to Y$ is a continuous map, and $g(x, 0) = g(x', 0)$ for all $x, x' \in X$, then there is a unique map $h : Cone(X) \to Y$ for which $h \circ f = g$.*

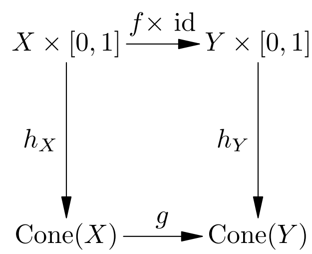

Figure 2: Extending maps over cones

Lemma 1.5 *If X and Y are topological spaces and $f : X \to Y$ is a continuous map between them, then there is a canonical continuous map g : Cone(X) $\to$ Cone(Y) which is identical to f when restricted to the base of each cone. Moreover, the map g is injective iff f is injective, g is surjective iff f is surjective, and g is a homeomorphism iff f is a homeomorphism.*

Proof: Let $h_X : X \times [0,1] \to \text{Cone}(X)$ and $h_Y : Y \times [0,1] \to \text{Cone}(Y)$ be the respective quotient maps. By Lemma 1.4, there is a unique map $g : \text{Cone}(X) \to \text{Cone}(Y)$ such that $h_Y(f(x),t) = g(h_X(x,t))$ for all $x \in X$ and $t \in [0,1]$. See Figure 2. The correspondences between the properties of f and g are straightforward and are left to the reader. □

The particular cones of interest in this article are cones over compact subsets of a Euclidean space such as $\mathbf{R}^n$. If X is a compact subset of $\mathbf{R}^n$ and $\mathbf{R}^n$ is embedded in $\mathbf{R}^n \times \mathbf{R}$, then there is a standard construction of Cone(X) given by the set of points $(tx, 1-t)$ for $t \in [0,1]$ and $x \in X$. The notation tx is meant to indicate the scalar multiplication of the coordinates of the point x in $\mathbf{R}^n$ by the real number t. These points form a subspace of $\mathbf{R}^{n+1}$ which is homeomorphic to Cone(X). Indeed the function which sends (x,t) to $(tx, 1-t)$ is a function from $X \times [0,1]$ to the standard construction of Cone(X) which factors through the quotient and provides the homeomorphism. Alternatively, the standard construction of Cone(X) can be described as the set of all points $(\lambda x, \mu 1)$ with $x \in X$, $\lambda, \mu \geq 0$, and $\lambda + \mu = 1$.

Homotopy Theory and Deformation Retractions If $h : B \times [0,1] \to C$ and $f, g : B \to C$ are maps such that $f(b) = h(b,0)$ and $g(b) = h(b,1)$ for all $b \in B$ then h is called a homotopy between f and g. If A is a subspace of B and $h(a,t) = h(a,t')$ for all $a \in A$ and $t, t' \in [0,1]$ then f and g are said to be homotopic relative to A. Let C be a topological space with B a subspace of C. A deformation retraction from C to B is a homotopy between the identity map on C and a map from C to $B \subset C$ which remains constant on B during the homotopy. More specifically, a deformation retraction is a continuous function from $h : C \times [0,1] \to C$ such that $h(c,0) = c$ for all $c \in C$, $h(b,t) = b$ for all $b \in B$ and $t \in [0,1]$ and $h(c,1) \in B$ for all $c \in C$. A map $f : C \to D$ is called a homotopy equivalence iff there exists a map $g : D \to C$ such that the composition $f \circ g$ is homotopic to the identity map on D and the composition $g \circ f$ is homotopic to the identity map on C. The map f is called a weak homotopy equivalence iff the induced homomorphisms $f_i : \pi_i(C) \to \pi_i(D)$ between the ith homotopy groups are isomorphisms for all $i \geq 0$. The relative space (C, B) is called ∞-connected if the relative homotopy groups $\pi_i(C, B)$ are trivial for all i. For traditional complexes these concepts are equivalent in the following sense.

Lemma 1.6 *If C is a connected CW complex and B is a connected subcomplex, then the following conditions are equivalent:*

1) there is a deformation retraction from C to B
2) the inclusion map $i : B \to C$ is a homotopy equivalence
3) the inclusion map $i : B \to C$ is a weak homotopy equivalence
4) the relative CW complex (C, B) is ∞-connected.

Proof: The implications $1 \Rightarrow 2 \Rightarrow 3$ are well-known to be true in general topological spaces. The equivalence of 3 and 4 follows from the fact that C and B are path-connected and from the long exact sequence of homotopy groups. Finally $4 \Rightarrow 1$ is shown in [20]. □

A space C is called contractible if there is a deformation retraction of C to a point, and it is called weakly contractible if all of its homotopy groups are trivial. The next two lemmas are easy consequences of Lemma 1.6. A useful connection between homotopy theory and cones is given in Lemma 1.9.

Lemma 1.7 *A CW complex is contractible iff it is weakly contractible.*

Lemma 1.8 *If B is a subcomplex of a contractible CW complex C then B is contractible iff there is a deformation retraction of C onto B.*

Lemma 1.9 *Let $f : B \to C$ be a continuous map. The map f is homotopic to a constant map iff there is a map from Cone(B) to C whose restriction to its boundary B is the map f. As a special case, notice that a topological loop in a topological space C is homotopic to a point iff there is a map from the unit disk into C so that the restriction of this map to the unit circle is the topological loop.*

Proof: Let $h : B \times [0, 1] \to C$ be a homotopy showing that the map f is homotopic to a constant map, and let $g : B \times [0, 1] \to \text{Cone}(B)$ be the quotient in the definition of Cone(B). By Lemma 1.2 there exists a map $h' : \text{Cone}(B) \to C$ with $h = h'g$, and the map h' satisfies the lemma. Conversely, given a map $h' : \text{Cone}(B) \to C$ as described, the composition $h = h'g$ is a homotopy proving that the restriction of h' to its boundary is homotopic to a constant map. □

The following lemmas are being recorded for later use.

Lemma 1.10 *If C is a topological space, then Cone(C) is contractible.*

Proof: Let $f : C \times [0,1] \to \text{Cone}(C)$ be the quotient map. The function $h(f(c,t),s) = f(c,ts)$ with $c \in C$, and $s,t \in [0,1]$ provides a well-defined deformation retraction from the topological cone Cone(C) to its vertex. □

Lemma 1.11 *If there is a deformation retraction of a CW complex C onto a subcomplex B then there is a deformation retraction of Cone(C) onto Cone(B) which extends this deformation.*

Proof: Let $f : C \times [0,1] \to \text{Cone}(C)$ be the usual quotient map, and let $h : C \times [0,1] \to C$ be a deformation retraction of C onto B. Then the map $g : \text{Cone}(C) \times [0,1] \to \text{Cone}(C)$ defined by $g(f(c,t),s) = (h(c,s),t)$ is a well-defined deformation retraction of Cone(C) onto Cone(B) which extends h. □

If B is a subspace of C, then let $f : B \to B$ be the identity map from B viewed as the base of Cone(B) to B viewed as a subspace of C. By the notation $C \cup_B \text{Cone}(B)$ is meant the union of the cone Cone(B) and the space C identified along the subspace B, or more specifically, the space $C \cup_f \text{Cone}(B)$. Notice that $C \cup_B \text{Cone}(B)$ is canonically situated as a subspace of Cone(C).

Lemma 1.12 *If B is a subcomplex of a connected CW complex C and there exists a deformation retraction from C to B, then $C \cup_B$ Cone(B) is contractible, and there is a deformation retraction from Cone(C) onto $C \cup_B$ Cone(B).*

Proof: The deformation retraction h from C to B can be combined with the identity deformation on Cone(B) which leaves every point in Cone(B) fixed. The extension is well-defined since the two deformations agree on their overlap, B. Together they yield a deformation retraction from $C \cup_B \text{Cone}(B)$ to Cone(B). Since by Lemma 1.10 Cone(B) is contractible, $C \cup_B \text{Cone}(B)$ is also contractible, and by Lemma 1.8 there exists a deformation retraction from the space Cone(C) to $C \cup_B \text{Cone}(B)$. □

Covering Spaces An onto map $f : C \to B$ between connected CW complexes is called a covering map if for every point $b \in B$ there is a connected open set U containing b such that $f^{-1}(U)$ is a nonempty disjoint union of sets $U_i \subset C$ where f restricted to each U_i is a homeomorphism onto U for all i. If $f : C \to B$ is a covering map with $f(c_0) = b_0$ then the group homomorphism $f_1 : \pi_1(C, c_0) \to \pi_1(B, b_0)$ induced by f is an injection. The topological space B is called the base space and C is the covering space or simply a cover of B. A covering space which is simply connected is called a universal cover. Let $g : D \to B$ and $h : D \to C$ be maps between connected spaces and let $f : C \to B$ be a covering map. If $f \circ h = g$, then h is said to be a lift of g. When $g : D \to B$ is understood from context, and there exists such a map $h : D \to C$, then D is said to lift to C. See Figure 3. Lemmas 1.13 and 1.14 are special cases of the Lifting Theorem in algebraic topology. Recall that the notation f_* refers to the homomorphism between fundamental groups which is induced by the map f.

Lemma 1.13 *Let $f : C \to B$ and $g : D \to B$ be maps between connected CW complexes, where f is a covering map. If D is simply connected and b_0, c_0, and*

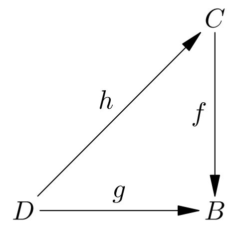

Figure 3: Lifting maps

d_0 are points in these spaces such that $f(c_0) = b_0 = g(d_0)$, then there is a lift $h : D \to C$ of g with $h(d_0) = c_0$. More generally, such a lift exists whenever $f_*(\pi_1(D)) = 0$.

Lemma 1.14 *Let B, C, and D be connected CW complexes, let $f : C \to B$ be a covering map, and let $g : D \to B$ be a continuous map. If h_1 and h_2 are two lifts of g and $h_1(d) = h_2(d)$ for some $d \in D$, then $h_1(d) = h_2(d)$ for all $d \in D$. In other words a lift is uniquely determined by the image of a single point.*

If $f : C \to B$ is a covering map and $g : C \to C$ is a homeomorphism of C with itself such that $f \circ g = f$, then g is called a deck transformation. The set of all deck transformations forms a group. If the group acts transitively on the set $f^{-1}(b_0)$ then C is a regular cover of B and f is a regular covering map.

Lemma 1.15 *Let B be a connected CW complex with fundamental group G. There is a 1 to 1 correspondence between the regular covers of B and the normal subgroups of G. In particular, the regular cover C of B corresponding to the normal subgroup H of G is a regular cover whose fundamental group C is isomorphic to H. Moreover, the group of deck transformations of the covering map $f : C \to B$ is isomorphic to G/H.*

The key fact is that such a regular cover exists for every normal subgroup of G. To illustrate the lemma, notice that the regular cover corresponding to group G is B itself, and the regular cover corresponding to the trivial subgroup is the universal cover of B.

1.2 Inductive Construction

To establish notation, let $\mathbf{N}$ denote the set of natural numbers $0, 1, 2, \ldots$, and let the symbol $\sqcup$ indicate a disjoint union. A set I which is the disjoint union of sets $I_n, n \in \mathbf{N}$ will be called a graded index set. Equivalently, a graded index set can be defined by means of a rank function $r : I \to \mathbf{N}$, with $r(i) = n \iff i \in I_n$. In this case n is called the rank of the element i. If $J = \sqcup_{n \in \mathbf{N}} J_n$ is another graded index set and $f : I \to J$ is a function which preserves the index, meaning that $f(I_n) \subset J_n$ for all $n \in \mathbf{N}$, then f is called a graded index function.

A conical CW complex C is constructed inductively from spaces C_i and maps $\partial\phi_i$ indexed by some graded index set I subject to certain restrictions on the spaces and the maps. The restrictions on the spaces C_i are as follows: the space C_i must be a single point for all i of rank 0, while for all other $i \in I$, C_i must be the topological cone over some nonempty compact subset of a Euclidean space. The base, or boundary, of the cone is called ∂C_i so that $C_i = \text{Cone}(\partial C_i)$ for all $i \in I_n$ with $n \geq 1$. Note that the spaces ∂C_i need not be connected. For convenience, this notation is extended to I_0 by defining the unique point in C_i to be its vertex, defining $\partial C_i = \emptyset$ for each $i \in I_0$, and setting $\text{Cone}(\emptyset) = C_i$, even though technically a point is not a topological cone over the empty set. For all $i \in I_n$, the space C_i is called an n-cone or a cone of rank n, ∂C_i is called its base, and $C_i \setminus \partial C_i$ is called its interior. Again, the terminology is used even in the case $i \in I_0$ although the 0-cones are not technically cones at all. These conventions provide a coherent notation throughout the induction. Finally, for each $n \in \mathbf{N}$ let $X_n = \sqcup_{i \in I_n} C_i$ and $\partial X_n = \sqcup_{i \in I_n} \partial C_i$. Furthermore let $X = \sqcup_{n \in \mathbf{N}} X_n$ and $\partial X = \sqcup_{n \in \mathbf{N}} \partial X_n$.

The construction of a conical CW complex consists of inductively defining spaces $C^{(n)}$, called the n-skeletons of C, with $C^{(n)} \subset C^{(n+1)}$ for all $n \in \mathbf{N}$ and $C = \cup_{n \in \mathbf{N}} C^{(n)}$. The construction begins by defining $C^{(0)} = X_0$. The space $C^{(0)}$ is a discrete set of points in 1 to 1 correspondence with the set I_0 which is given the discrete topology. Next assume that $C^{(n-1)}$ has already been defined for some $n \geq 1$. The restriction on the maps $\partial\phi_i$ mentioned earlier is that for each $i \in I_n$, $\partial\phi_i$ must be a map from ∂C_i to the $(n-1)$-skeleton. If $\{\partial\phi_i : \partial C_i \to C^{(n-1)} | i \in I_n\}$ is a set of such maps, then piecing these together forms a map $\partial\phi_n : \partial X_n \to C^{(n-1)}$ which agrees with $\partial\phi_i$ when ∂X_n is restricted to ∂C_i. The n-skeleton can then be defined as $C^{(n)} = C^{(n-1)} \cup_{\partial\phi_n} X_n$. By Lemma 1.3, $C^{(n-1)}$ is embedded in $C^{(n)}$ as a closed subspace and there is a canonical map $\phi_n : X_n \to C^{(n)}$ which is a homeomorphism when restricted to $X_n \setminus \partial X_n$. More specifically, if the restriction of ϕ_n to the space C_i is called ϕ_i, then the map ϕ_i agrees with $\partial\phi_i$ on ∂C_i and by Lemma 1.3 ϕ_i is a homeomorphism when restricted to $C_i \setminus \partial C_i$. The map ϕ_i is called the characteristic map of the cone C_i and the map $\partial\phi_i$ is called the attaching map of the cone C_i. If C_i is an n-cone then the homeomorphic sets $C_i \setminus \partial C_i$ and $\phi_i(C_i \setminus \partial C_i)$ are called open n-cones. By Lemma 1.3 $\phi_i(C_i \setminus \partial C_i)$ is an open subset of the n-skeleton $C^{(n)}$.

Once the n-skeletons $C^{(n)}$ have been constructed for all n, the space C is defined to be the union of these nested spaces, $C = \cup C^{(n)}$, under the weak topology, the one which specifies that a subset U is open (or equivalently closed) in C iff $U \cap C^{(n)}$ is open (closed) in $C^{(n)}$ for all $n \in \mathbf{N}$. When convenient the characteristic maps ϕ_i and the attaching maps $\partial\phi_i$ will be considered as maps into C instead of into one of the skeletons. Maps $\phi : X \to C$ and $\partial\phi : \partial X \to C$ can be defined by piecing together the maps ϕ_i and $\partial\phi_i$ with $i \in I$, respectively. As is the case with CW complexes, an open cone of height n is only guaranteed to be an open subset of the n-skeleton $C^{(n)}$, and is not necessarily an open subset of the final complex C.

Lemma 1.16 *If C is a conical CW complex then, using the notations estab-*

lished above, C is the disjoint union of its open cones, the n-skeleton is closed in C, and the map $\phi : X \to C$ is a quotient map. Additionally, a subset V of C is open (closed) iff $\phi_i^{-1}(V)$ is open (closed) for all $i \in I$.

Proof: The fact that C is the disjoint union of open cones and the fact that the n-skeleton is closed in C follow inductively from Lemma 1.3 and the definition of the weak topology, respectively. Let V be a subset of C. The weak topology on C guarantees that V is open (closed) in C iff $V \cap C^{(n)}$ is open (closed) for all $n \in \mathbf{N}$. The set $V \cap C^{(n)}$ is open (closed) iff $V \cap C^{(n-1)}$ is open (closed) and the sets $\phi_i^{-1}(V)$ are open (closed) in C_i for all $i \in I_n$ since $C^{(n)}$ is defined by a quotient map. By an easy induction it is clear that $V \cap C^{(n)}$ is open (closed) iff $\phi_i^{-1}(V)$ is open (closed) in C_i for all $i \in \cup_{j=0}^{n} I_j$. Putting this all together, V is open (closed) in C iff $\phi_i^{-1}(V)$ is open (closed) in C_i for all $i \in I$ iff $\phi^{-1}(V)$ is open (closed) in X. Thus by virtue of the definition, ϕ is a quotient map. □

As a consequence of Lemma 1.16, every point c in C can be assigned a rank based on the unique i such that c is in $\phi_i(C_i \setminus \partial C_i)$. Many of the properties and proofs for traditional CW complexes carry over unchanged to conical CW complexes. The following lemma is an example. The proof given in [2] carries over verbatim.

Lemma 1.17 *Let C be a conical CW complex with the notation defined above. If U is a subset of C which has no two points in the same open cone, then U is closed and discrete. If V is a compact subset of C then V is contained in a finite union of open cones. As a consequence, $\phi_i(C_i)$, and indeed every compact subset of C, is contained in a finite subcomplex of C.*

Just as open and closed subsets of conical CW complexes can be determined 'locally', continuous maps from conical CW complexes can also be described 'locally'. This fact is recorded in the following lemma.

Lemma 1.18 *Let C be a conical CW complex with characteristic maps $\phi_i : C_i \to C$, with $i \in I$, and let D be an arbitrary topological space. A function $f : C \to D$ is continuous iff all of the composites $f \circ \phi_i$ are continuous for all $i \in I$. More specifically, if f is a bijection such that f restricted to $\phi_i(C_i)$ is a homeomorphism, and subsets V of D are open in D iff $V \cap f(\phi_i(C_i))$ is open in $f(\phi_i(C_i))$ for all $i \in I$, then f is a homeomorphism.*

Proof: By Lemma 1.16, the map $\phi : X \to C$ is a quotient map. The first statement is then an immediate consequence of Lemma 1.1. The second assertion follows since $V \subset D$ is open iff $V \cap f(\phi_i(C_i))$ is open in $f(\phi_i(C_i))$ for all $i \in I$ iff $f^{-1}(V \cap f(\phi_i(C_i)))$ is open in $\phi_i(C_i)$ for all $i \in I$ iff $f^{-1}(V)$ is open in C by Lemma 1.16. □

1.3 Internal Description

In order to completely describe a conical CW complex it is sufficient to give the cones C_i, their bases ∂C_i, the topological space C, and the characteristic maps

$\phi_i : C_i \to C$, all indexed by some graded index set I. The above list is sufficient since the attaching maps $\partial\phi_i$ and the n-skeletons $C^{(n)}$ can be reconstructed from the given information. Conversely, given a set of topological spaces C_i with subspaces ∂C_i, a topological space C, and a set of maps $\phi_i : C_i \to C$, all indexed by some graded index set I, and satisfying certain minimal conditions, it is possible to show that the space C is homeomorphic to a space arrived at following the inductive procedure described above. This static description of a conical CW complex is called an internal description to distinguish it from the constructive approach taken above. Lemma 1.18 will then be used below to show the equivalence of the constructive and the internal descriptions of conical CW complexes.

Let C_i be a collection of topological spaces with subsets ∂C_i, and let $\phi_i : C_i \to C$ be a set of maps to a topological space C, all indexed by some graded index set I. Assume that C_i is a single point and $\partial C_i = \emptyset$ for all $i \in I_0$ and that $C_i = \text{Cone}(\partial C_i)$ with ∂C_i a nonempty compact subset of some finite-dimensional Euclidean space for all other $i \in I$.

Assume also that C is the disjoint union of the sets $\phi_i(C_i \backslash \partial C_i)$, that $\phi_i(C_i) \subset \partial\phi_j(\partial C_j)$ for some $i \in I_m$ and $j \in I_n$ implies $m < n$, and that C has the weak topology with respect to the maps ϕ_i, $i \in I$. This latter condition means that a subset U is open in C iff $\phi_i^{-1}(U \cap \phi(C_i))$ is open in C_i for all $i \in I$. When all of these conditions are satisfied, the spaces $C_i, \partial C_i$, and the maps ϕ_i are said to provide C with the structure of a conical CW complex. When these spaces and maps are understood from context, C is said to have a conical CW complex structure.

Lemma 1.19 *A topological space C has a conical CW complex structure iff it is homeomorphic to a space arrived at by following the inductive construction described above.*

Proof: If C is a topological space which has a conical CW complex structure, then the spaces C_i, ∂C_i, and the maps ϕ_i restricted to the bases ∂C_i satisfy all of the conditions needed for them to be used to inductively construct a conical CW complex. Lemma 1.18 is then sufficient to show that the constructed complex is homeomorphic with the original topological space C in the obvious way. The converse is immediate from the inductive construction. □

A cone map between two conical CW complexes is a map which, when composed with a characteristic map of the domain, yields a characteristic map of the range. More specifically, let B be a conical CW complex with characteristic maps $\phi_i : B_i \to B$ indexed by a graded index set I, and let C be another conical CW complex with characteristic maps $\phi_j : C_j \to C$ indexed by a graded index set J. A cone map between B and C is a map $f : B \to C$ together with a graded index function $p : I \to J$ and a set of homeomorphisms $h_i : B_i \to C_{p(i)}$ such that $f \circ \phi_i = \phi_{p(i)} \circ h_i$, for all $i \in I$. If $X = \sqcup B_i$, and $Y = \sqcup C_j$, then the maps ϕ_i, ϕ_j, and h_i can be pieced together to form maps ϕ_X, ϕ_Y and h respectively. Using this notation the cone map condition becomes $f \circ \phi_X = \phi_Y \circ h$. See Figure 4. Notice that since p is a graded index function, the homeomorphisms

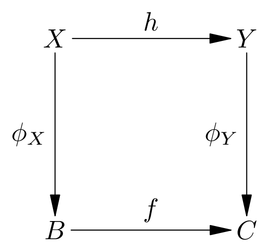

Figure 4: A cone map between conical CW complexes

h_i must send cones of rank n in B to cones of rank n in C, and similarly the map f must send points of rank n in B to points of rank n in C.

Lemma 1.20 *The composition of cone maps is a cone map, and bijective cone maps are homeomorphisms.*

Proof: The first assertion is immediate from the definition, since each of the necessary conditions is well-behaved under composition. For the second assertion let $f : B \to C$ be a bijective cone map between conical CW complexes B and C. Since attaching maps are homeomorphisms when restricted to open cones, f is a bijective cone map, and the interior of every cone B_i is nonempty, the function between the graded index sets must be injective. Also, since f is onto and C is the disjoint union of the nonempty sets $C_i \setminus \partial C_i$, the function between the graded index sets must be surjective, and thus bijective. Thus, without loss of generality, the same set I can represent the graded index set of both B and C.

Next let X, ϕ_X, Y, ϕ_Y, and h be as defined in the inductive construction. Since the index sets are the same, the homeomorphisms $h_i : B_i \to C_i$ combine to form a single homeomorphism $h : X \to Y$. Since $\phi_X \circ h^{-1} = f^{-1} \circ \phi_Y$ as functions, ϕ_Y is a quotient map by Lemma 1.16, and $\phi_X \circ h^{-1}$ is continuous, the function f^{-1} must be continuous, by Lemma 1.1. Thus both f and f^{-1} are continuous, and f is a homeomorphism. □

1.4 Coverings of Conical CW Complexes

In the next lemma, the structure of conical CW complexes is shown to lift through covering maps. Although the proof mimics the proof of the result for traditional CW complexes, there are enough differences to warrant a complete proof.

Lemma 1.21 *Let B be a topological space which has the structure of a conical CW complex, and suppose that all of the cones B_i in B are topologically equivalent to traditional CW complexes. If $f : C \to B$ is a covering map, then f induces a conical CW complex structure on C which is canonically defined.*

Proof: Let B_i, ∂B_i, and ϕ_i be a set of topological spaces, subspaces, and characteristic maps indexed by a graded index set I which provide B with a conical CW complex structure. In particular, this means that these data satisfy the conditions listed immediately prior to Lemma 1.19. After establishing suitable notations, a graded set of spaces, subspaces, and maps for C will be defined from the given data for B, and it will be shown that the new data provide C with a conical CW complex structure.

For every $i \in I$ let v_i be the point in B which is the image under ϕ_i of the vertex of B_i. Next let V_B be the set of all points v_i for $i \in I$ and let $V_C = f^{-1}(V_B)$. By Lemma 1.17 the set V_B is a closed and discrete subset of B since it contains exactly one point in every open cone. Notice that the set V_C is also closed since f is continuous, and the set is discrete since f is a covering map. Let J be a set in 1 to 1 correspondence with V_C, and let v_j be the point in V_C which corresponds to $j \in J$. For every $j \in J$ there is a unique $i \in I$ with $f(v_j) \in \phi_i(B_i \setminus \partial B_i)$, since B is the disjoint union of the images of the interiors of the cones B_i. This serves to define a function $p : J \to I$, and if J_n is defined as $p^{-1}(I_n)$ then J becomes a graded index set and p becomes a graded function. If for a particular $j \in J$, $p(j) = i$, then define $C_j = B_i$, define $\partial C_j = \partial B_i$, and define $\phi_j : C_j \to C$ by lifting $\phi_i : B_i \to B$ through f to C in such a way that the image of the vertex of C_j is v_j. Since by Lemma 1.10 B_i is contractible and B_i has the topology of a traditional CW complex by assumption, such a lift is always possible by Lemma 1.13, and by Lemma 1.14 it is unique.

The spaces C_j, the subspaces ∂C_j, and the maps ϕ_j indexed by the graded set J will provide C with a conical CW complex structure. First, notice that the spaces C_j are either points or cones, and the subspaces ∂C_j are either empty or compact subspaces of finite-dimensional Euclidean spaces as appropriate, since each C_j and each ∂C_j is equal to a space B_i or ∂B_i, and the corresponding spaces have the same rank. Next, if there are indices $j \in J_n$ and $j' \in J_m$ such that $\partial\phi_j(C_j) \cap \phi_{j'}(C_{j'}) \neq \emptyset$, then letting $i = p(j) \in I_n$ and $i' = p(j') \in I_m$, and recalling that $f \circ \phi_j = \phi_{p(j)}$, it is also true that $\partial\phi_i(\partial B_i) \cap \phi_{i'}(B_{i'}) \neq \emptyset$. Thus $n > m$, and consequently the attaching map sends the boundary of the cone C_i to the $(n-1)$-skeleton.

Let $c \in C$ be an arbitrary point, and let i be the unique index such that $f(c) \in \phi_i(B_i \setminus \partial B_i)$. Next let b be a point in $B_i \setminus \partial B_i$ such that $\phi_i(b) = f(c)$. Using the contraction of B_i to its vertex, there is a path in $B_i \setminus \partial B_i$ from b to the vertex. When this path is sent to B by ϕ_i, it becomes a path from $f(c)$ to v_i, and the unique lift of ϕ_i which lifts $f(c)$ to c must lift this path to a path in C which starts at c and ends at a lift of v_i. The lift of v_i must be v_j for some j by definition, and this shows that c is in $\phi_j(C_j \setminus \partial C_j)$ for this j. Since lifts are unique once a single point has been lifted, the j described above is the unique $j \in J$ for which the condition is true, and this shows that C is the disjoint union of the images of the interiors of the cones C_j as required. In symbols, $C = \sqcup_{j \in J} \phi_j(C_j \setminus \partial C_j)$.

It only remains to show that the covering space C has the weak topology with respect to the maps ϕ_j. If U is an open subset of C it is clear that $\phi_j^{-1}(U)$

is open in C_j since each ϕ_j is continuous. So assume instead that $\phi_j^{-1}(U)$ is open in C_j for all $j \in J$, and for the moment, assume that U is contained in an open subset of C which is mapped homeomorphically onto its image under f. The set $\phi_i^{-1}(f(U)) = \cap\{\phi_j^{-1}(U)|p(j) = i\}$ for each $i \in I$ since, if $b \in \phi_i^{-1}(f(U))$ then $\phi_i(b) = f(u)$ for some $u \in U$, there is a lift of ϕ_i to a map ϕ_j for some $j \in J$ which extends the lift of $\phi_i(b)$ to u, and thus b is in $\phi_j^{-1}(U)$ for this j. Conversely, if b is in $\phi_j^{-1}(U)$ then $\phi_j(b) = u$ is in U, $f(\phi_j(b)) = \phi_{p(j)}(b) = \phi_i(b)$ is in $f(U)$, and b is in $\phi_i^{-1}(f(U))$. Since by assumption the sets $\phi_j^{-1}(U)$ are open for all $j \in J$, and each set $\phi_i^{-1}(f(U))$ is the union of such sets, the sets $\phi_i^{-1}(f(U))$ must be open sets for all $i \in I$. By Lemma 1.16 and the fact that B has the structure of a conical CW complex, the set $f(U)$ is open in B. Since by assumption U is contained in an open subset which is mapped homeomorphically onto its image under f, U is open in C.

Let U' be an arbitrary subset of C with $\phi_j^{-1}(U')$ open for all $j \in J$, let u be a point in U', and let U be the intersection of U' with an open set of C containing U which is mapped homeomorphically onto its image by f. Such a subset exists since f is a covering map. Since $\phi_j^{-1}(U)$ is the intersection of $\phi_j^{-1}(U')$ with an open set for each $j \in J$, U is a set satisfying the conditions described above, and thus U is an open subset of C. The union of all such sets U, one for each $u \in U'$, shows that U' is also an open set. □

Lemma 1.22 *Let B be a topological space which has the structure of a conical CW complex, and suppose that all of the cones B_i in B are topologically equivalent to traditional CW complexes. If $f : C \to B$ is a covering map, and C is given the conical CW structure induced by B through f, then f is a cone map. In particular, all of the deck transformations of C relative to f are also cone maps.*

Proof: Since in the proof of Lemma 1.21 the characteristic maps of C were created by lifting the characteristic maps of B through f, it is immediate that the composition of a characteristic map of C with f yields a characteristic map of B. Thus f is a cone map between the conical CW complexes C and B. Since all deck transformations are trivial covering maps of C in their own right, they are cone maps by the above reasoning. □

2 Cone Complexes

Under certain conditions, all of the information needed to construct a conical CW complex can be summarized in a single algebraic structure. In particular, any partially ordered set in which all principal ideals are finite can be used to construct a unique conical CW complex directly from the poset. Moreover, simplicial complexes and cell complexes can be viewed as special cases of this construction. At the end of the section it is shown that all of the conical CW complexes created in this way are in fact polyhedra.

2.1 Posets and Cone Complexes

A partially ordered set, or poset for short, is a set P with a reflexive, antisymmetric, and transitive binary operation $\leq$. If P is a poset with $p \in P$, then define $[p, q] = \{r \in P | p \leq r \leq q\}$, and $(p, q) = \{r \in P | p < r < q\}$, etc. A principal ideal P/p is defined as $P/p = (-, p] = \{q \in P | q \leq p\}$. Principal ideals are denoted P/p because of their close connections with slice categories as described below. An ideal Q of P is a subset of P such that $p \in Q$ implies $P/p \subset Q$. The height of an element p is given by $\text{height}(p) = \sup\{n | \exists p_0 < p_1 < p_2 \ldots < p_n = p\}$. The elements $p_0 < p_1 < p_2 \ldots < p_n$ are called a chain of length n, and a subset of the p_i's form a subchain. Notice that according to this definition minimal elements have height 0, not 1. The height of an ideal, principal ideal or poset P is the supremum of the heights of its elements.

A poset C in which all of the principal ideals are finite is called a cone complex. The fact that the principal ideals are finite guarantees that every element in C has a well-defined and finite height. In particular, every cone complex C can be assumed to be equipped with a rank function r from the elements of C to $\mathbf{N}$ such that $r(p) = 0 \Leftrightarrow p$ is a minimal element, and $p > q \Rightarrow r(p) > r(q)$ for all p, q in C without loss of generality, since the height function shows that such a function always exists. The elements of a cone complex are called open cones, the principal ideals are called closed cones, if $b < c$ then b is called a face of c, and the poset of all faces of c, $\partial c = \{b \in C | b < c\}$, is called the boundary of c. Notice that the boundary of c is the difference between the closed cone c and the open cone c.

The rank of a cone complex or of an ideal of a cone complex is defined as the supremum of the ranks of its elements. A cone of rank n is called an n-cone, and the ideal of all i-cones with $i \leq n$ is called the n-skeleton. The 0-skeleton is also called the vertex set and the 0-cones are called vertices. Notice that the vertices are the only elements which are both open and closed and consequently have no boundary. A subcomplex of a cone complex is any ideal of the poset. A morphism between two cone complexes is an order-preserving function which is an isomorphism when restricted to corresponding closed cones. More explicitly, a morphism between cone complexes is a poset morphism $f : B \to C$ and a set of isomorphisms $h_b : B/b \to C/f(b)$ for all $b \in B$ such that the diagram in Figure 5 commutes. That is, $f \circ \phi_b = \phi_{f(b)} \circ h_b$, where ϕ_b and $\phi_{f(b)}$ represent the natural inclusions.

Lemma 2.1 *If $f : B \to C$ is a morphism between cone complexes B and C and B and C are also closed cones, then the rank of B is less than or equal to the rank of C, and the ranks are equal if and only if f is an isomorphism of cone complexes.*

Proof: Let b and c be elements of B and C respectively such that $B = B/b$ and $C = C/c$. Since posets are antisymmetric, the elements b and c are unique. Moreover, they clearly have the maximum rank in B and C, respectively, so that $\text{rank}(B) = \text{rank}(b)$ and $\text{rank}(C) = \text{rank}(c)$. Next, since f is an isomorphism whenever the domain and range are restricted to corresponding closed cones

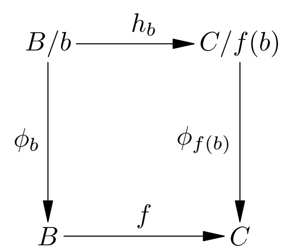

Figure 5: A morphism between cone complexes

and since B itself is a closed cone, $f(B)$ is an isomorphic copy of B situated as a subcomplex of C. It follows that, $\text{rank}(B) = \text{rank}(b) = \text{rank}(f(b)) \leq \text{rank}(c) = \text{rank}(C)$. If, in addition, $\text{rank}(f(b)) = \text{rank}(c)$ then $f(b) = c$ since c is the unique element of maximum rank in C. Thus $B = B/b$ is isomorphic to $C/f(b) = C/c = C$ since f is a morphism between cone complexes. Conversely, if f is an isomorphism then B and C clearly have the same rank. □

Let C be a cone complex and let I be its graded index set. More specifically, let c_i be the element of C which corresponds to $i \in I$, and let the rank of c_i equal the rank of i. Over the course of the next several lemmas a conical CW complex corresponding to C will be defined and will be shown to be unique. Initially, the resemblance will be assumed to be only slight. A geometric realization of the cone complex C is a conical CW complex B with the same index set I, such that the characteristic map ϕ_i is a homeomorphism, and $\partial\phi_i(\partial B_i) = \cup\{\phi_j(B_j) | c_j \in \partial c_i\}$ for all $i \in I$. If B is a geometric realization of the cone complex C, then the space $\phi_i(B_i \setminus \partial B_i)$ in B will be called the topological open cone c_i to distinguish it from the open cone c_i which is an element in the poset C. Once the equivalence of the two notions has been shown, the distinction will be less important.

Lemma 2.2 *If B is a geometric realization of a cone complex C, then the union of a set of topological open cones forms a closed subspace of B iff the set of corresponding elements in C forms an ideal.*

Proof: Let D be an ideal of C, and let B_i be one of the cones of B. The image of B_i is the subspace C/c_i. Since the intersection of the ideals D and C/c_i is a finite ideal, the intersection of the subspaces D and C/c_i is the finite union of open cones in B. Because $D \cap C/c_i$ is an ideal, $\partial\phi_j(\partial B_j)$ is contained whenever $\phi_j(B_j \setminus \partial B_j)$ is contained, so that $D \cap C/c_i$ is the finite union of sets of the form $\phi_j(B_j)$. Thus $D \cap C/c_i$ is closed, $\phi_i^{-1}(D \cap C/c_i) = \phi_i^{-1}(D)$ is closed for all $i \in I$, and by Lemma 1.16 D is closed in B.

Conversely, let D be the union of a set of open cones in B and assume that D is closed. By Lemma 1.16, the sets $\phi_i^{-1}(D)$ are closed for all $i \in I$. Since every point in the boundary of a topological cone such as B_i is the limit of points in its interior, any closed set which contains the interior $B_i \setminus \partial B_i$ must contain the boundary ∂B_i as well. Thus if D contains the topological open cone c_i it must also contain all of the topological open cones c_j where $c_j < c_i$ in C. This implies that the subset of C corresponding to the space D is an ideal. □

If B is a geometric realization of a cone complex C and D is an ideal of C, then the union of the topological open cones in B corresponding to the elements of D is called the subspace D. Thus, it is possible to speak of the subspace C/c or the subspace ∂c of B for each element $c \in C$.

Lemma 2.3 *Every cone complex has a geometric realization as a conical CW complex.*

Proof: Let C be an arbitrary cone complex and let I be the graded index set determined by the rank function of C. A conical CW complex B which is a geometric realization of C will be constructed inductively. Let $B^{(0)}$ be a discrete set of points in 1 to 1 correspondence with the elements in I_0, let B_i be a single point, let ∂B_i be the empty set, and the let $\phi_i : B_i \to B^{(0)}$ be the map which sends the point in B_i to the point in $B^{(0)}$ corresponding to the element i, for all $i \in I_0$. Notice that $B^{(0)}$ is a geometric realization of the 0-skeleton of C. Next suppose that $B^{(n)}$ has already been constructed and is a geometric realization of the n-skeleton of C. For every element $c_i \in C$ of rank $n+1$, the fact that $B^{(n)}$ is a geometric realization of the n-skeleton means that by Lemma 2.2 there is a well-defined closed subspace ∂c_i in $B^{(n)}$. Since ϕ_i is supposed to be a homeomorphism, the space ∂B_i can be defined to be this closed subspace with $B_i = \text{Cone}(\partial B_i)$ and $\partial \phi_i : \partial B_i \to B^{(n)}$ the natural inclusion map. These attaching maps are sufficient to construct the $(n+1)$-skeleton, and more generally, the spaces B_i and the attaching maps $\partial \phi_i : \partial B_i \to B^{(n)}$ indexed by the set I completely determine the inductive construction of a conical CW complex B which satisfies the lemma. □

For every cone complex C there is a standard geometric realization of C which can be described as a subspace of a real vector space whose basis is in 1 to 1 correspondence with the elements of C. Notice that even though the result may not be finite-dimensional, every finitely generated subspace is a Euclidean space. Let v_c be the basis element which corresponds with $c \in C$. The construction proceeds by induction on the skeleta of C. To begin the minimal elements $c \in C$ are placed at the points v_c and this is trivially a geometric realization of the 0-skeleton. Next, suppose that the k-skeleton has been realized geometrically. For every element c of rank $k+1$ in C a new vertex v_c is added, and then the straight line segments connecting v_c with each of the points in the subspace ∂c are added. The points added by the construction of this cone, namely, the point v_c together with the interiors of all of the line segments, form the subset which can be called the topological open cone c since a routine check shows that this construction satisfies the necessary conditions, and that the result is a geometric realization of the $k+1$-skeleton. Continuing in this way yields a geometric realization of the cone complex C in which every space C/c is a standard cone over its base in the Euclidean space spanned by the basis elements corresponding to the elements of C/c. The next lemma shows that every geometric realization of a cone complex is homeomorphic to the standard realization.

Lemma 2.4 *The geometric realization of a cone complex as a conical CW complex is unique up to homeomorphism. In particular, if B and B' are both geometric realizations of a cone complex C, then there is a homeomorphism $f : B \to B'$ which is also a cone map which sends the topological open cone c in B to the topological open cone c in B' for all $c \in C$.*

Proof: The proof is by induction through the skeleta. Any two realizations of a cone complex of rank 0 consist solely of isolated points and these realizations clearly possess the requisite homeomorphism. Next suppose that every cone complex of rank at most n satisfies the lemma, and let C be a cone complex of rank at least $n+1$. If B and B' are two geometric realizations of C, then by induction there is a homeomorphism between their n-skeletons which satisfies the lemma. If c_i is an element of C of rank $n+1$, then the homeomorphism of the n-skeletons restricts to a homeomorphism between the subspaces of B and B' called ∂c_i. Since the characteristic maps g_i are embeddings, this induces a homeomorphism between the subspaces ∂B_i and $\partial B'_i$. By Lemma 1.5 there is a canonical homeomorphism h_i between B_i and B'_i which extends the homeomorphism between the bases. If $X = \sqcup\{B_i | i \in I_{n+1}\}$ and $Y = \sqcup\{B'_i | i \in I_{n+1}\}$, then piecing together the homeomorphisms h_i and the homeomorphism between the n-skeletons yields a homeomorphism $h : B^{(n)} \sqcup X \to B'^{(n)} \sqcup Y$. Since the composition of h with the quotient map from $B'^{(n)} \sqcup Y$ to $B'^{(n+1)}$ is continuous, by Lemma 1.2 there is a continuous map from $B^{(n+1)}$ to $B'^{(n+1)}$. Similarly, composing h^{-1} with the quotient map from $B^{(n)} \sqcup X$ to $B^{(n+1)}$ induces a continuous map in the other direction. It is easy to check that these two continuous maps are inverses of each other, so that they provide the homeomorphism described in the lemma. Thus, by induction, the n-skeleton of B is homeomorphic with the n-skeleton of B' for all n. Since by construction these homeomorphisms agree on their overlaps the union of the homeomorphisms is a bijective function between B and B'. By Lemma 1.16 this function is a homeomorphism. $\Box$

Because of the strong correlation between the algebraic properties of a cone complex and the topological properties of its geometric realization, the two constructions will henceforth be denoted by the same letter. Thus the geometric realization of a cone complex C will be called simply the space C.

Lemma 2.5 *If $f : B \to C$ is a morphism between cone complexes B and C, then there is a cone map from the space B to the space C such that the map restricted to an open cone b is a homeomorphism onto the open cone $f(b)$, for all $b \in B$. Moreover, the cone map corresponding to an isomorphism of cone complexes is a homeomorphism between conical CW complexes.*

Proof: The lemma is true for the standard realizations of B and C since the unique $\mathbf{R}$-linear map from $\oplus_{b\in B}\mathbf{R}$ to $\oplus_{c\in C}\mathbf{R}$ which sends the basis element v_b

to the basis element $v_{f(b)}$ restricts to a cone map from the standard realization of B to the standard realization of C, which satisfies the conditions of the lemma. The details of the verification are left to the reader. If B and C are given arbitrary geometric realizations, then there are homeomorphisms, which are also cone maps, between these other geometric realizations and the standard realization by Lemma 2.4. The composition of these homeomorphic cone maps with the cone map between the standard realizations described above provides the desired cone map. The final statement follows from Lemma 1.20. □

2.2 Simplicial Complexes and Cell Complexes

A simplicial complex can be completely reconstructed from its poset of faces. A description of simplicial complexes in terms of posets will be given below. This description is equivalent to the usual definition in content, as will be shown in Lemma 2.7. Besides providing a familiar example of a cone complex, simplicial complexes and their properties, such as the Simplicial Approximation Theorem, will be needed at various times throughout the article.

Let V be a set and let P be the poset of all non-trivial finite subsets of V ordered by subset. The elements of V can be identified with the minimal elements of P. If C is an ideal of P which contains all of V, then C is called a simplicial complex over V. Since every principal ideal of C is finite, C is also a cone complex. Without loss of generality assume that C comes equipped with a rank function $r : C \to \mathbf{N}$ such that r is order-preserving, and $r(p) = 0 \Leftrightarrow p$ is a minimal element in C. This is possible since the height function shows that such a function always exists. Once a simplicial complex C is viewed as a cone complex, the notions of face, boundary, rank of an element, rank of an ideal, n-skeleton, vertex set, vertices, subcomplex and geometric realization are already defined for C. Some terminology becomes more specialized. For instance, an open cone of a simplicial complex is called an open simplex, a closed cone is called a closed simplex, a cone of rank n is called an n-simplex, and a geometric open cone is called a geometric open simplex.

Notice that morphisms between simplicial complexes, as defined here, must preserve the rank of the simplices. Such a definition is much more restrictive than the usual notion of a simplicial map. Finally, notice that the principal ideals of C are finite Boolean lattices minus the 0 element and that each principal ideal is uniquely determined by the vertices it contains. Conversely, any cone complex whose principal ideals satisfy these two conditions is a simplicial complex.

Lemma 2.6 *A cone complex is a simplicial complex iff all principal ideals are closed simplices and every principal ideal is uniquely determined by the vertices it contains.*

Proof: The forward direction is immediate from the definition, so let C be a cone complex satisfying the two conditions above, let V be the set of minimal elements of C, and let P be the poset of all non-trivial finite subsets of V

ordered by subset. Because every principal ideal in C is uniquely determined by its minimal elements, the order-preserving morphism from C to P which sends every element to the set of vertices of its principal ideal is injective. Thus C can be considered a substructure of P. Since, in addition, the principal ideals of C are closed simplices, the image of C must actually be an ideal of P. By definition, C is a simplicial complex. $\Box$

By Lemma 2.3 every simplicial complex C has a geometric realization. The topological space C is called a polyhedron, and more generally, a polyhedron is any topological space which is homeomorphic to the space of a simplicial complex. This discussion of simplicial complexes concludes with a sketch of a proof that simplicial complexes as defined here correspond with the traditional definition.

Lemma 2.7 *The geometric realization of a simplicial complex as defined here is homeomorphic to the traditional definition of a simplicial complex.*

Proof: (Sketch) If the construction of the standard realization is modified so that the points corresponding to the elements c of rank 1 or more are placed at the average of the coordinates of the vertices in the subspace ∂c, then it can be shown that, because of the special restrictions placed on simplicial complexes, the conditions necessary for the construction to qualify as a geometric realization are still true, and in addition the geometric closed cone c is the convex hull of the vertices it contains. Since these vertices are in general position, the convex hull is the traditional notion of a geometric closed n-simplex. By Lemma 2.4 this geometric construction of C is the geometric realization of C up to homeomorphism, and it corresponds to the traditional definition of a geometric simplicial complex built out of geometric n-simplices. $\Box$

A cell complex is a type of cone complex which is more flexible than a simplicial complex and for that reason it is frequently used in topology as a means of efficiently constructing piecewise linear manifolds. A geometric closed cell is the convex hull of a finite set of points in a Euclidean space $\mathbf{R}^n$. The interior of a closed cell is called an open cell, and the boundary of a cell is the difference between the two. It can be shown that every open (closed) cell is homeomorphic to an open (closed) n-dimensional ball for some unique n called the dimension of the cell. An n-dimensional cell is called an n-cell. If c is the interior of the convex hull of a set of points and b is the interior of the convex hull of a subset of these points which is wholly contained in the boundary of c, then b is called a face of c, and we write $b < c$. A cell complex C is a collection of disjoint open cells, c, in a generalized Euclidean space, such that if $c \in C$ and $b < c$ then $b \in C$, and if $b, c \in C$ then $b \cap c$ is a face of both b and c. The n-skeleton is the cellular subcomplex consisting of the cells with dimension less than or equal to n. If c is an open cell of dimension $n + 1$, then let ∂c be the set of all faces of the cell c. The set ∂c is always finite and, if c is an n-cell, the union of the open cells in ∂c is homeomorphic to an $(n - 1)$-sphere.

Lemma 2.8 *Geometric cell complexes are recoverable up to homeomorphism from their poset of open cells ordered by faces.*

Proof: Once it is realized that the convex hull of a set of points in $\mathbf{R}^n$ can be viewed as a cone over the boundary of the closed cell with the vertex of the cone located at an arbitrary point in the interior, the original geometric cell complex is seen to be a geometric realization of the poset of its faces viewed as a cone complex. By Lemma 2.4, this is the unique geometric realization of the cone complex, and the proof is complete. □

2.3 Quillen's Construction on Posets

A general construction for turning an arbitrary poset into a geometric object was described by D. Quillen in [17]. Starting from a poset P, a new poset called Chain(P) is constructed whose elements are the nonempty finite chains of elements in P with one chain being less in the ordering than another chain iff the former is a subchain of the latter. The resulting poset is a simplicial complex which can be realized geometrically either in the traditional manner or by the cone construction of Lemma 2.3. If the poset is a cone complex, then the construction of Lemma 2.3 can be used directly on the original poset P. The relationship between the Quillen procedure and the cone construction described above is detailed in the following lemmas.

Lemma 2.9 *If P is an arbitrary poset then $C = Chain(P)$ is a simplicial complex in which the vertices of C correspond to the elements of P.*

Proof: The nonempty chains of minimal length are clearly in 1 to 1 correspondence with the elements of P. Since it is also clear that every chain in a poset is uniquely determined by its elements and that the nonempty subchains of a given chain form a Boolean lattice without the 0 element, it follows from Lemma 2.6 that C is a simplicial complex. □

Lemma 2.10 *If P is a poset, $C = Chain(P)$, and P' is a poset obtained from P by forcibly adding a new maximum element called p, then $C' = Chain(P')$ is a simplicial complex containing C as a subcomplex and geometrically C' is the simplicial cone over the topological space C. The vertex of the cone is the 1-element chain p, and the additional open simplices are precisely the chains in P' which contain p as a largest element.*

Proof: Immediate. □

Lemma 2.11 *If C is a cone complex and $C' = Chain(C)$ then there is a homeomorphism f between the conical CW complexes C and C' such that for every open cone c in C, $f(c)$ is the topological union of all of the open simplices in C' which contain an element c as its largest member. In this sense, the simplicial structure of C' is a refinement of the cone structure of C.*

Proof: The proof proceeds by induction on the height of the cone complex C. The lemma is trivially true for cone complexes of height 0 since the spaces involved are collections of isolated points of the same cardinality. If it is assumed that the lemma is true for all cone complexes of height at most k, then Lemma 2.10 guarantees that the homeomorphism f can be extended to the additional $(k+1)$-cones in such a way that the statement of the lemma is preserved. By induction, the lemma holds for all cone complexes of finite height. Since general cone complexes are constructed inductively from their k-skeletons, this is sufficient to prove the lemma in general. □

In the particular case where P is already a simplicial complex, Chain(P) is its barycentric subdivision. In general, given a cone complex P the existence of a simplicial subdivision such as Chain(P) shows that the space P is a polyhedron.

Corollary 2.12 *The geometric realization of a cone complex is a polyhedron.*

Since the geometric realizations of simplicial complexes are a type of CW complex, all of the lemmas about CW complexes given in section 1 are applicable to the spaces of cone complexes, cell complexes, and circular complexes. Although the cone construction and the Quillen construction are homeomorphic topologically, they create distinct partitions of the space and there are advantages to each. The Quillen construction produces a traditional simplicial complex which allows for the utilization of the tremendous number of results already developed for such spaces. The cone construction on the other hand divides the topological space into skeleta in a significant way. In the case of an originally geometric simplicial or cell complex, the cone construction is able to reconstruct the original complex from its poset of faces, whereas the Quillen construction produces either a barycentric subdivision or a simplicial subdivision of the original. Of more importance in the current context is the fact that the skeleta of the space of a particular type of cone complex called a circular complex correspond nicely to algebraic properties of groups defined using these complexes. As a consequence, the partitioning of the underlying space into skeleta by the cone construction is used later in the article to efficiently organize many of the statements and proofs of the lemmas.

3 Cone Categories

The construction of the previous section can be made more flexible by changing the underlying posets to categories. The resulting structure is called a cone category. In this section cone categories are introduced and their properties investigated in a manner which parallels the previous section.

3.1 Categories and Cone Categories

Familiarity with the basic definitions of category theory, such as that of a category and a functor, is assumed. Let C be a category and let c be an object of C. The slice category, C/c, of all objects over c is given by

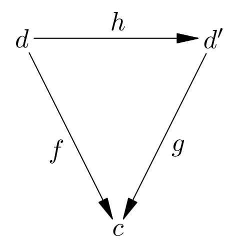

Figure 6: Commuting triangle over c

$$\text{Objects}(C/c) = \{\text{arrows with codomain } c\}$$

$$\text{Arrows}(C/c) = \{\text{arrows which form commuting triangles}\}$$

For example, if $f : d \to c, f' : d' \to c$, and $h : d \to d'$ are arrows in C with $hg = f$ then f and f' are objects in C/c and h is an arrow in C/c from f to f' (see Figure 6). The functor $\phi_c : C/c \to C$ which sends each object $f : d \to c$ to its domain d is called a characteristic functor of C/c on analogy with the functions used in the construction of a CW complex. The connection with CW complexes is examined in more detail below. See Lemma 3.6. Notice that the slice category C/c has the identity arrow of c as a terminal object.

An important fact in this context is that every poset can be viewed as a category by thinking of the elements of the poset as the objects of the category and imagining a unique arrow $p \to q$ iff $p \leq q$. The notion of a slice category in a category generalizes the notion of a principal ideal in a poset in the sense that if a poset P is viewed as a category, then the slice category P/p is isomorphic to the ideal (p) viewed as a category. The attaching map in this case is injective on both objects and arrows. Notice that the usual descriptions of posets and categories use different spatial metaphors. In particular, the ideal P/p can at the same time be described as the elements of the poset 'below' p and, when viewed as a category, as the objects of the category 'over' p.

Lemma 3.1 *A category C is the category corresponding to a poset iff C has singleton hom-sets and contains no non-trivial cycles.*

Proof: The forward direction is immediate, and the other direction is almost immediate. The condition that the hom-sets be singleton sets guarantees that C is the category of a preorder, namely the preorder given by taking the objects with $c \leq d$ iff there is an arrow from c to d. The restriction that there are no non-trivial cycles implies that the preorder is anti-symmetric. □

A useful result about categories in general is the following.

Lemma 3.2 *If $f : B \to C$ and $g : C \to B$ are functors between categories B and C such that fg is the identity functor on C, and gf is the identity functor on B, then f and g are category isomorphisms.*

Proof: Since fg is the identity functor on C, the functor f is onto, and the functor g is 1 to 1 on both objects and arrows. Similarly, since gf is the identity functor on B, g is onto and f is 1 to 1 on both objects and arrows. Thus both f and g are bijections on objects and arrows. $\Box$

A category C in which every slice category C/c is the category of a finite poset is called a cone category. The fact that the slice categories are finite posets guarantees that every object c in C has a well-defined and finite height based on the height of the poset associated with C/c. In particular, every cone category C can be assumed to be equipped with a rank function r from the objects of C to $\mathbf{N}$ such that $r(c) < r(d)$ whenever there exists a non-identity arrow from c to d, and $r(c) = 0 \Leftrightarrow c$ is not the codomain for any non-identity arrows, without loss of generality since the height function shows that such a function always exists. An object c of a cone category C is called an open cone and the slice category C/c with the characteristic functor $\phi_c : C/c \to C$ implicitly understood is called a closed cone. The full subcategory of C/c on the set of all objects other than its terminal object is called the boundary of c and it is denoted ∂c. The restriction of ϕ_c to the subcategory ∂c is called the attaching functor $\partial\phi_c : \partial c \to C$ of the closed cone C/c. The terminology is justified by the fact, shown below, that every cone category can be realized geometrically as a conical CW complex, in which the functions corresponding to the characteristic functors and the attaching functors are the characteristic maps and the attaching maps of the closed topological cones and their boundaries.

Lemma 3.3 *If C is a cone category then C contains no non-trivial cycles and in the slice category C/c the unique terminal object corresponds to the identity arrow at c.*

Proof: If C contains a non-trivial cycle of arrows and c is an object in the cycle then it is easy to check that C/c also contains a non-trivial cycle. This, however, contradicts the fact that C/c is the category of a poset by Lemma 3.1. For the second statement, simply note that closed cones in a cone complex have unique maximum elements, so that the categories of these posets possess unique terminal objects. Since every arrow in C with codomain c factors trivially through the identity arrow id_c at c, id_c is a terminal object in C/c and thus the unique terminal object. $\Box$

Every cone category C determines a graded index set I given by setting I_n in 1 to 1 correspondence with the objects of C which have rank n. Let c_i be the object of C corresponding to $i \in I$. In the same way that the poset ∂c is an ideal in the poset C/c for all c in a cone complex C, the category of the poset ∂c is a full subcategory of the category C/c for all objects c in a cone category C. The restriction of the characteristic functor ϕ_c to the category boundary ∂c, denoted $\partial\phi_c : \partial c \to C$, is called the attaching functor of the closed cone C/c. The n-skeleton is the full subcategory of C on the objects of rank less than or equal to n. A cone functor is a rank-preserving functor $f : B \to C$ between cone categories which induces an isomorphism on slices. Specifically, for every object $b \in B$ there must be a category isomorphism h_b such that the diagram in Figure 5 commutes, and in addition, the rank of b in B must equal the rank of $f(b)$ in C. Notice that the diagram has been reinterpreted in terms of categories and functors, instead of posets and order-preserving maps.

Many of these definitions are identical to the definitions in the previous section. The precise relationship between cone categories and cone complexes is shown in the following lemma.

Lemma 3.4 *A cone category is the category of a cone complex iff all the characteristic functors are injective on objects.*

Proof: The forward direction is immediate. If on the other hand all of the attaching functors are injective on vertices, then all of the hom-sets of C must be singletons since the slice category over the codomain of a non-trivial hom-set produces an attaching map which is not injective on objects. By Lemma 3.3 C contains no non-trivial cycles and thus by Lemma 3.1 C is the category of a poset. In this case the characteristic functors correspond to the embeddings of the principal ideals C/c into the poset C. Since C is a cone category and all of its slices are finite, the poset of C possesses finite principal ideals, making it a cone complex. □

Lemma 3.5 *Let $f : c \to d$ be an arrow in a cone category C, and let C/f be the full subcategory of C/d corresponding to the principal ideal of f in the poset of C/d. There exists a functor $g : C/c \to C/f$ such that $\phi_c = \phi_d \circ g$. The functor g is an category isomorphism which embeds C/c as a subcategory of C/d and sends the terminal object of C/c to the object f in C/d.*

Proof: An object of C/c is an arrow in C with codomain c. Its product with f is then an arrow in C with codomain d and thus an object in C/d. Since the product as an arrow in C factors through f it is also an object in the subcategory C/f. Notice that the identity arrow of c is sent to f under this procedure. Since all of the hom-sets in C/c and C/f are singleton and the ordering of the underlying posets is preserved, there is a unique functor $g : C/c \to C/f$ which extends the function on the objects described above.

The fact that the hom-sets in C/d are singletons implies that if an arrow in C with codomain d factors through f, and thus is an object in C/f, then it must factor uniquely. Since the objects of C/f are precisely the arrows of C with codomain d which factor through f, there is a well-defined function from the objects of C/f to the objects of C/c which is the inverse of the function described above. This second function also preserves the ordering of the underlying posets, and it determines a functor from C/f to C/c. It is easy to check that the second functor is the inverse of the functor g, that g is thus a category isomorphism (by Lemma 3.2), and that $\phi_c = \phi_d \circ g$. □

Recall from the section on conical CW complexes that the topological cones C_i, their bases ∂C_i, and the attaching maps $\partial \phi_i$ from ∂C_i to the appropriate skeleton over some graded index set I are sufficient to reconstruct an entire conical CW complex.

Lemma 3.6 *Every cone category C has a geometric realization as a conical CW complex.*

Proof: Let I be the graded index set determined by the rank function on the objects of the cone category C. For all $i \in I$, let C_i be the geometric realization of the poset of the slice category C/c_i, and let ∂C_i be the conical CW subcomplex of C_i which corresponds to the ideal ∂c_i. These geometric realizations C_i exist by Lemma 2.3. It only remains to describe the attaching maps $\phi_i : \partial C_i \to C^{(n-1)}$ for $i \in I_n$. If $n = 0$ then C_i is a point, ∂C_i is empty, and the attaching maps do not need to be described. So assume that $n > 1$, that the $(n-1)$-skeleton has already been constructed, and that $i \in I_n$ is fixed. In this case the characteristic maps $\phi_j : C_j \to C^{(n-1)}$ for all $j \in I_m$ with $m < n$ have already been defined.

By Lemma 3.5 every proper principal ideal in C/c_i is isomorphic with a slice category over an element c_j of rank strictly less than n. Since this category isomorphism also provides an isomorphism between the cone complexes, it can be geometrically realized as a homeomorphic cone map between the conical CW complex C_j and the appropriate subcomplex of the conical CW complex C_i by Lemma 2.5. The composition of this homeomorphism with the already defined characteristic map on C_j provides a portion of the attaching map for C_i. If this procedure is followed for all proper principal ideals in C/c_i and the resulting portions of the attaching map are pieced together, a complete attaching map for C_i is formed. Another application of Lemma 3.5 insures that the definitions of the portions of the attaching maps derived from different principal ideals agree on their overlaps, so that the resulting map is well-defined. Since this completes the list of information necessary to create a conical CW complex, the proof is complete. □

Lemma 3.7 *Given any cone functor $f : B \to C$ between cone categories there is a cone map between geometric realizations of B and C as conical CW complexes such that the cone map restricted to the topological open cone b is a homeomorphism onto the open cone $f(b)$ for all objects $b \in B$. Moreover, the cone map corresponding to an isomorphism of cone categories is a homeomorphism between conical CW complexes.*

Proof: Let I (J) be the graded index sets for the cone category B (C), let B_i (C_j) be the geometric realization of the slice categories B/b_i (C/c_j), let X (Y) be the disjoint union of the topological open cones B_i (C_j), and let ϕ_X (ϕ_Y) be the quotient map described in Lemma 1.16. If $f(b_i) = c_j$, then by the definition of a cone functor there is a category isomorphism between the slice categories B/b_i and C/c_j and by Lemma 2.4 there is a canonical homeomorphism between the conical CW complexes B_i and C_j. These homeomorphisms combine to form a map $h : X \to Y$.

Define a map f between the geometric realizations of B and C as follows. By Lemma 1.16 each point in the geometric realization of B is contained in a unique open cone b_i, and since the characteristic map ϕ_i restricted to the interior of B_i is a homeomorphism, it lifts uniquely to B_i in X. Then map it by h and then ϕ_Y to a point in the geometric realization of C. This creates a function f such that $f \circ \phi_X = \phi_Y \circ h$ as functions when restricted to the interiors of the cones B_i in X. See Figure 4. By the construction of the attaching maps given above in Lemma 3.6, these functions are equal on the boundaries of the cones as well. Moreover, since $\phi_Y \circ h$ is continuous and ϕ_X is a quotient map, by Lemma 1.1 f is continuous. Clearly this function is a homeomorphism when f is restricted to an open cone in the domain. Finally, an isomorphism between cone categories will produce a map f which is bijective. Thus by Lemma 1.20 f will be a homeomorphism. □

Corollary 3.8 *The geometric realization of a cone category is unique up to homeomorphism.*

3.2 Simplicial Categories and Cell Categories

A simplicial category is a category in which every slice is a closed simplex. That is, for all $c \in \text{Objects}(C)$, the category C/c is isomorphic to the category of faces of an n-dimensional simplex under inclusion for some n. Note that the dimension of the simplex need not be the rank of the object c unless the rank function corresponds to the height function in this instance. An object of C is called an open simplex, and the attaching map $\phi_c : C/c \to C$ is called a closed simplex. For simplicity, C/c will be referred to as a closed simplex with the attaching functor understood. The rank of an open simplex c is given by the rank function on C, and the n-skeleton is the full subcategory on all objects with rank less than or equal to n. The boundary of c, ∂c, is the poset C/c minus its maximal element together with its attaching map. Every simplicial category has a geometric realization which is constructed inductively via its skeleta. The functors $\phi_c : C/c \to C$ are the categorical equivalents of attaching maps, which describe how to attach the boundary of a closed n-simplex into the $(n-1)$-skeleton, hence the name. A simplicial functor between simplicial categories is a rank-preserving functor which is isomorphic on slices. That is, $f : C \to B$ is a simplicial functor iff for all $c \in \text{Objects}(C)$ there is an isomorphism between C/c and $B/f(c)$, with the rank of c in C equal to the rank of $f(c)$ in B.

Simplicial categories illustrate the type of flexibility which is added by the change from poset to categories. The boundary of a simplex in a simplicial category can be collapsed in some way, and distinct simplices can have identical boundaries without being identical.

Lemma 3.9 *A category is a simplicial complex iff all slice categories are categories of closed simplices, all of the attaching maps are injective on objects and every slice is determined by the image of its initial objects.*

Proof: The forward direction is immediate from the definition. Assume that C is a category which satisfies the latter conditions. If a hom-set in C contained two or more arrows then the attaching map of the slice category over the terminal object would not be injective. Thus C has singleton hom-sets, and C is the category of a preorder. If C contained a non-trivial cycle of arrows or an

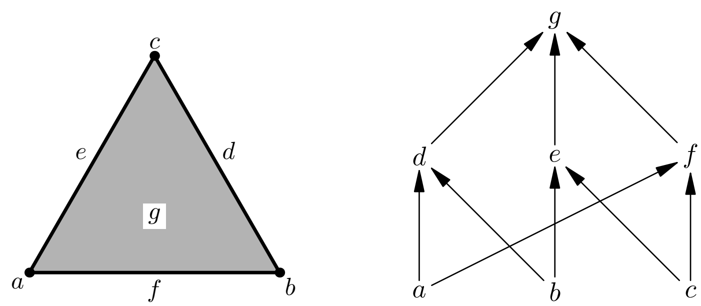

Figure 7: A 2-simplex

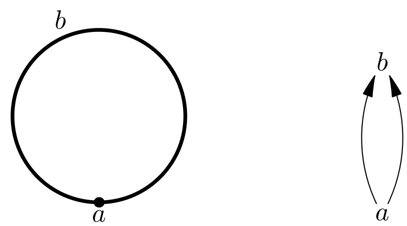

Figure 8: A loop edge

infinite chain of arrows with no first arrow in the chain, then the same condition would hold in one of the slice categories, contradicting the fact they are each categories of closed simplices. Thus C is the category of a poset. Once this is known the remaining conditions become simply those listed in Lemma 2.6 as necessary and sufficient to show that a poset is a simplicial complex. □

To illustrate the differences between a simplicial complex and a simplicial category, three examples are given below. The first is a 2-dimensional simplicial complex and thus also a simplicial category, while the other two are examples of simplicial categories which are not complexes.

Example 1 In Figure 7 there are three vertices, three edges, and one triangle. These correspond to elements and simplices of height 0,1, and 2 respectively.

Example 2 The simplicial category in Figure 8 contains a single vertex and a loop edge, that is, an edge whose endpoints are identical. It is clearly not a simplicial complex since the category used to define it is not a poset. Notice that the geometric realization of the slice category of objects over b is a straight line segment; it is the attaching map which identifies the two endpoints.

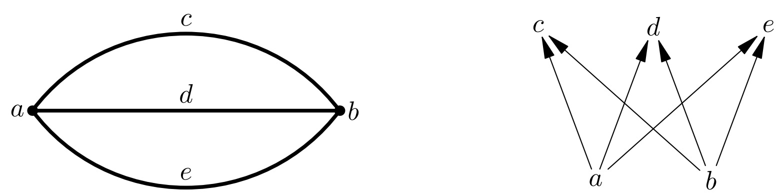

Figure 9: Multiple edges

Example 3 In the simplicial category in Figure 9, the three edges all have isomorphic slice categories and are attached to the same vertices, but the attaching maps have distinct images. This is not possible in a simplicial complex since the simplices are determined by the vertices they contain.

When the type of flexibility described above is incorporated into a cell complex the result is called a cell category, or as it is known elsewhere in the literature, a regular CW complex. A cell category is a category in which the geometric realization of every slice category is homeomorphic with a closed n-cell, for some n. That is, for all $c \in \text{Objects}(C)$, the category C/c is isomorphic to the category of faces of an n-cell under inclusion, for some n. The definitions of open cells, closed cells, cell boundaries, characteristic functors, the dimension of an object, the n-skeleton, and cell functors are analogous with those used above for simplicial categories. The Poincaré construction associated with a group presentation is a non-trivial example of a 2-dimensional cell category. It will be shown by Lemma 3.12 that every cell category can be subdivided to form a simplicial complex, and thus topologically they are polyhedra.

3.3 Quillen's Construction on Categories

Let C be a cone category with distinct objects c_i, $i = 0, 1, \ldots, n$, and arrows $f_i : c_{i-1} \to c_i$. The sequence of composable arrows

$$c_0 \xrightarrow{f_1} c_1 \xrightarrow{f_2} c_2 \cdots c_{n-1} \xrightarrow{f_n} c_n$$

is called a chain of length n. A chain of length 0 is simply an object of C with no arrows. Since the objects are distinct, the arrows f_i are all non-identity arrows. Moreover, since C is a cone category and thus has no non-trivial cycle of arrows, the product of any composable subsequence of the arrows is also not an identity arrow. A subchain of a chain of arrows in C is an arbitrary subset of the objects, arranged in order, with the appropriate products of arrows from the chain placed between the objects. For any arbitrary cone category C, the poset Chain(C) is defined as the set of all finite chains ordered by subchains.

Lemma 3.10 *If C is a cone category, the category of the poset Chain(C) is a simplicial category in which the attaching functors are injective on objects.*

Additionally, the minimal elements of Chain(C) are in 1 to 1 correspondence with the objects of C.

Proof: The principal ideals in the poset Chain(C) are clearly closed simplices since the subchains of a chain are indexed by the objects of the chain which it contains. When the poset is converted into a category, this shows that the resulting category is a simplicial category whose attaching functors are injective on objects. The final statement is a reformulation of the fact that a chain of length 0 is simply an object of C. □

In the original paper by Quillen[17], no restrictions are placed on the objects and arrows. In fact, such restrictions are only possible in categories such as cone categories where there are no non-trivial cycles of arrows. The difference is inconsequential topologically since the resulting constructions are homotopically equivalent, but the construction without the identity arrows is closer in spirit to the poset construction, and the result is cleaner to work with.

Lemma 3.11 *If C is a cone category and C' is the poset Chain(C) then the topological space of C is homeomorphic with that of C' and the homeomorphism carries an open cone c onto the union of the chains (open simplices) of C' which contain c as a terminal object. In this sense, C' is a simplicial subdivision of C.*

Proof: The proof proceeds by induction on the height of the cone category C. The lemma is trivially true for cone categories of height 0 since the spaces involved are collections of isolated points of the same cardinality. If it is assumed that the lemma is true for all cone categories of height at most k, then Lemma 2.10 guarantees that the homeomorphism f can be extended to the additional $(k+1)$-cones in such a way that the statement of the lemma is preserved. By induction, the lemma holds for all cone catgories of finite height. Since general cone caetgories are constructed inductively from their k-skeletons, this is sufficient to prove the lemma in general. □

Lemma 3.12 *If C is a cone category then $C'' =$ Chain(Chain(C)) is a simplicial complex whose topological space is homeomorphic to that of C. Thus every cone category is a polyhedron.*

Proof: The result is a combination of Lemma 3.11 and Lemma 2.11. □

Let $f : B \to C$ be a cone functor between cone categories. Since cone functors preserve the height of the objects, the image of a sequence of arrows which form a chain in B will also form a chain in C. This defines a function called Chain(f) from Chain(B) to Chain(C).

Lemma 3.13 *If $f : B \to C$ is a cone functor between cone categories, then Chain(f) : Chain(B) $\to$ Chain(C) is a simplicial functor between simplicial categories.*

Proof: Since subchains are clearly sent to subchains under Chain(f), it is an order-preserving function, and since the principal ideals of Chain(B) and Chain(C) are the poset of all subchains of a given chain, the restriction of Chain(f) to a principal ideal is clearly an isomorphism. □

3.4 Coverings of Cone Categories

The final result in this section shows that the covers of the geometric realizations of cone categories also have the structure of a cone category. Recall that the same letter is used to denote a cone category, technically a category rather than a topological space, and also to denote its geometric realization as a conical CW complex.

Lemma 3.14 *If B is the geometric realization of a cone category, and $f : C \to B$ is a covering map between a topological space C and the conical CW complex B, then f is the geometric realization of a cone functor $f : C \to B$ between two cone categories. Similarly, all of the deck transformations of C relative to f are geometric realizations of cone functors from the cone category C to itself.*

Proof: Clearly the lemma is true if B has rank 0, so assume that the lemma is true whenever B has rank less than n for some n at least 1. By Lemma 3.12 the conical CW complex B and all of its closed cones B_i are polyhedra and thus are topologically equivalent to traditional CW complexes. Thus Lemma 1.21 can be applied giving C the structure of a conical CW complex and making f a cone map. It only remains to show that the conical CW complex C is also the geometric realization of a cone category. Since the k-skeleton of C is a cover of the k-skeleton of B for all $k < n$, the $(n-1)$-skeleton of C is the geometric realization of a cone category by the inductive hypothesis. Also, since f is a cone map, the closed cones of rank n in C are homeomorphic with closed cones in B, and thus are also geometric realizations of cone categories.

If the category underlying the $(n-1)$-skeleton of C and the categories underlying the n-cones of C are joined together according to the attaching maps the result is a new category of rank n. It is easy to check that the multiplication on this new structure is well-defined, that it is associative, and that the slice categories of the new category correspond to the categories of the closed n-cones of C. Moreover, since the attaching maps of the closed n-cones of C agree with those which would be constructed following the procedure in Lemma 3.6, the n-skeleton of C is the geometric realization of a cone category.

By induction, the k-skeleton of C is the geometric realization of a cone category for all k, and since these categories are compatible, their union is a category whose geometric realization is C. Finally, the second statement of the lemma is an immediate consequence of the first one. □

Part II
Relators

In this part the focus of attention is narrowed to those cone complexes and cone categories in which the boundaries of the cones are homotopically equivalent to S^1. These 'circular' categories are introduced in Section 4. In Section 5, the key concept of a general relator is defined and the technical properties of these relators are investigated.

4 Circular Categories

In this section circular categories are introduced, and notations for words, cycles, paths and loops in labeled circular categories are established. The section concludes with the proof of a 1-dimensional version of the simplicial approximation theorem for circular categories. In particular, it is shown that in a circular category all topological paths which start and end at vertices are homotopic relative to their endpoints to a combinatorial path in the 1-skeleton.

4.1 Circular Complexes and Circular Categories

A circular complex is a special type of cone complex which is tailor-made to work as a building block for a general small cancellation theory. A cone complex whose 1-skeleton is a simplical complex and where the subspace ∂c is homotopically equivalent to a circle for all elements c of rank at least 2 is called a circular complex. By virtue of the fact that every circular complex C is a cone complex, many concepts have already been defined for circular complexes. The elements $c \in C$ are called open circular cones, and the principal ideals C/c are closed circular cones. A cone map between circular complexes is called a circular map.

A circular category, C, is a cone category in which every slice is the category of a closed circular cone. An object of the category is called an open circular cone, the slice category C/c with the attaching functor, ϕ_c, implicitly understood is called a closed circular cone and ∂c, which is equal to C/c minus its maximum element, is the boundary of c. A circular functor is a cone functor between circular categories. The definitions of rank, the n-skeleton, and geometric realization are the same as for cone categories. Notice that the 1-skeleton of a circular category is a graph so that it is possible to speak of edges and vertices in C. Circular categories, merging the specificity of circular complexes with the flexibility of cone categories, will be the constructions of primary interest in the remainder of the article. In the same way that cell categories are an economical way to construct and prove results about manifolds, circular categories are an economical way to construct and prove results about groups.

4.2 Inversions and Words

The language of inversions presented below is an extension of the usage in [5] and is used here to provide a unified way to describe the fact that most functions between invertible objects (such as letters, words, edges, paths, cycles, etc.) commute with inversion. The ease with which a labeled graph can be defined illustrates the advantages of this approach.

Let A be a set and let $\iota : A \to A$ be a function from A to itself such that $\iota(\iota(a)) = a$ for all $a \in A$. The function ι is called an inversion of A and, when ι is understood, A is called an invertible set. The element $\iota(a)$ is the formal inverse of a and is usually written a^{-1}. The set $\{a, \iota(a)\}$ is called an orbit. An inversion homomorphism between invertible sets A and B is a function $f : A \to B$ such that $f(\iota(a)) = \iota(f(a))$ for all $a \in A$. Invertible sets are closed under direct product in the obvious way. If ι_A is the inversion of A and ι_B is the inversion of B, then the function $\iota_{A\times B}$ defined by $\iota_{A\times B}(a, b) = (\iota_A(a), \iota_B(b))$ is an inversion on the product $A \times B$.

An orientation of an invertible set A is a distinguished subset $A^+ \subset A$ such that $\iota(A^+)$ is disjoint from A^+ and $A^+ \cup \iota(A^+) = A$. The elements of A^+ are called the positive elements of A, and those in $\iota(A^+)$ are negative elements. When A^+ is clear from context we say that A is oriented. A necessary and sufficient condition for an orientation to be possible is that the inversion, ι, must be fixed-point-free. If $A^+ \subset A$ is an orientation of A then $\iota(A^+) \subset A$ is another orientation of A called the opposite orientation. If f is an inversion homomorphism between oriented sets A and B, it is called orientation-preserving if $f(A^+) \subset B^+$ and orientation-reversing if $f(A^+) \subset \iota(B^+)$.

An alphabet is a set. All alphabets considered in this article will be, in addition, finite, oriented, and invertible. The free monoid over an alphabet A is denoted A^*, and $|W|$ represents the length of a word $W \in A^*$. If W, X, Y, and Z are possibly empty words such that $W = XYZ$ in A^*, then X is an initial segment of W, Y is a subword of W, and Z is a final segment of W. If $W = XY$ and $Z = YX$ then Z is a cyclic conjugate of W. A cyclic conjugate is proper if both X and Y are nonempty words. The equivalence class of all cyclic conjugates of a word W is called the cycle of W or simply the cycle W. The cycle of W can be thought of as a string of letters in A which are 'written in a circle'. The circle can be then be broken at any point to create a particular cyclic conjugate of W. The inverse of a word, denoted W^{-1}, is the formal inverse of each letter listed in the opposite order, and the inverse of the cycle of the word W is the cycle of W^{-1}. A word, W, is reduced when $\iota(a)a$ does not occur as a subword for any letter $a \in A$, and cyclically reduced if all cyclic conjugates of W are reduced. Define W^0 to be the empty word, $W^{n+1} = W^nW$ for all positive integers and $W^{-n} = (W^{-1})^n$ when $-n$ is a negative integer. If $W = V^n$ for some word V and integer $n > 1$, then W is a power of V. If a reduced word W is not a power of any word V, then W is simple. Equivalently, W is simple iff it is not equal to any of its proper cyclic conjugates. If W is a subword of V^n then W is V-periodic and V is a period of W. The following lemma is easy to prove and well-known in the literature. See for example [13].

Lemma 4.1 *Let X and Y be simple non-conjugate words. If W is both X-periodic and Y-periodic, then $|W| < |X| + |Y|$. As a consequence, whenever X^i is Y-periodic for some integer $i > 1$, $|X| < \frac{1}{i-1}|Y|$.*

A word equivalent to the identity in a free group is called a Dyck word. These words have numerous well-known properties such as the fact that they can be reduced to the empty word by repeatedly removing subwords of the form aa^{-1} for some $a \in A$. Other easy properties are contained in the next two lemmas.

Lemma 4.2 *The set of Dyck words is closed under products, conjugates, and the removal of subwords which are also Dyck words. In addition, every non-empty Dyck word contains a subword of the form aa^{-1} for some $a \in A$, and the cycle of a Dyck word always contains at least two subwords of this form. All Dyck words have an even length. And finally, if W and W^{-1} are words in A^* which are conjugate to each other in the free group over A, then W is a Dyck word and thus equivalent to the identity element.*

4.3 Graphs and Cayley Graphs

A graph is a 1-dimensional simplicial category. Since 0-simplices and 1-simplices are indistinguishable from 0-cells and 1-cells, a graph could also be defined as a 1-dimensional cell category. The 0-simplices are called vertices, and the open 1-simplices are called edges. If a graph is also a simplicial complex then it is called a simplicial graph, in which case the attaching maps are injective and every pair of vertices determines at most 1 edge.

The slice category of an edge contains exactly 2 vertices called the endpoints of the edge, and the two possible orderings of these vertices yield two orientations of the edge. Although the endpoints are distinct in the slice category, they may be identified under the attaching map. An oriented edge is conventionally represented by an arrow going from one endpoint to the other. These vertices are called its initial and terminal vertices, respectively. Note that the arrow used to represent an oriented edge is very different from an arrow in the simplicial category of the graph. If Γ is a graph then the edge set of Γ, written Edges(Γ), is the set of oriented edges of Γ. The edge set of a graph is invertible and orientable since reversing the orientation of an edge generates a fixed-point-free involution of Edges(Γ). An oriented graph is a graph with an orientation of its edge set. To describe an oriented graph it is enough to describe the positive edges.

A graph morphism is a simplicial functor between graphs viewed as simplicial categories. Every graph morphism induces an inversion homomorphism between the edge sets. A graph morphism between oriented graphs is called orientation-preserving or -reversing according to the status of the inversion homomorphism between the edge sets. A morphism from a graph to itself is an automorphism and the set of all automorphisms forms a group called the automorphism group of the graph.

A graph Γ is labeled by an invertible alphabet A if there is an inversion homomorphism f : Edges(Γ) $\to A$. A graph which is labeled by A is called an A-graph or, if A is understood, simply a labeled graph. A labeled graph is deterministic if every pair of distinct oriented edges leaving the same initial vertex has distinct labels. A labeled circular category is one whose 1-skeleton is a labeled graph.

Let p be an inversion homomorphism from an alphabet A to a group G. When the unique extension of p to a monoid homomorphism $p^* : A^* \to G$ is onto, then it is said that A (via p) generates G, the letters in A are generators, and G is called an A-group. If the function p is injective, then the generators can be viewed as elements of G. The size of a generating set is the number of orbits in A. If G and H are A-groups via $p : A \to G$ and $q : A \to H$ respectively, and $f : G \to H$ is a group homomorphism such that $fp = q$, then f is called an A-group homomorphism. When X is a word in A^* sent to an element g in G, then X is called a representative of g.

Lemma 4.3 *If G and H are A-groups, and $f : G \to H$ and $g : H \to G$ are A-group homomorphisms, then f and g are isomorphisms.*

Proof: If $p : A \to G$ and $q : A \to H$ are extended to the monoid homomorphisms $p^* : A^* \to G$ and $q^* : A^* \to H$, then $fp^* = q^*$ and $gq^* = p^*$. Since q^* is onto, so is f. Next, let $p^*(X)$ and $p^*(Y)$ be arbitrary elements of G. If $f(p^*(X)) = f(p^*(Y))$ then $q^*(X) = q^*(Y)$, so that $g(q^*(X)) = g(q^*(Y))$ and $p^*(X) = p^*(Y)$. Thus f is 1 to 1 as well as onto, so that f is an isomorphism. By symmetry, g is an isomorphism. □

The Cayley graph of an A-group G, written $\mathcal{C}(G, A)$, is an A-graph constructed as follows: the vertices correspond to the elements of G, and the edge set corresponds to $G \times A$. The oriented edge associated with (g, a) starts at the vertex g, ends at the vertex ga (or, to be more exact, the vertex labeled by g multiplied on the right by $p(a)$), and is labeled by the letter a. The inverse edge which starts at ga and ends at g corresponds to (ga, a^{-1}) and is labeled a^{-1}. The vertex corresponding to the identity of the group is considered a distinguished vertex. Since Cayley graphs are deterministic and every vertex is the initial vertex of an edge labeled by an arbitrary letter, given any word W and any vertex v in the Cayley graph, there is always a reading of a word W starting at v, and it is unique. In particular, for every word W there is a unique reading of W starting at the base point and the endpoint u of this reading is the vertex corresponding to the element g of G which is equal to the product of the letters in the word W in the group G. This endpoint in turn determines a unique automorphism of the Cayley graph, namely, the one which sends the distinguished vertex to u. Several important and elementary properties of Cayley graphs are recorded in the following lemmas without proof.

Lemma 4.4 *If G is an A-group then its Cayley graph is the unique, connected, deterministic A-graph with a distinguished vertex whose label-preserving automorphism group acts transitively on the vertex set and is canonically isomorphic to G in such a way that the automorphism corresponding to the generator a moves the base point to a vertex connected to the base point by an oriented*

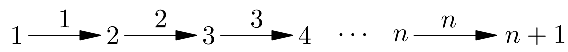

Figure 10: The abstract path, P_n

edge labeled by a, for all $a \in A$. In particular, if Γ is a connected, deterministic A-graph on which the group of label-preserving automorphisms acts transitively on its vertex set, then Γ is the Cayley graph of its automorphism group.

Lemma 4.5 *Let G be an A-group and let C be its Cayley graph. Every word $W \in A^*$ determines a unique automorphism of C. Thus, the automorphism corresponding to W is the identity automorphism iff the word W is equivalent to the identity element in the group, and consequently a path in C is a loop iff the word read by the path is equivalent to 1 in the group.*

Lemma 4.6 *If G is an A-group and H is a normal subgroup of G, then there is a unique label-preserving graph morphism from $\mathcal{C}(G, A)$ to $\mathcal{C}(G/H, A)$ which sends the distinguished vertex to the distinguished vertex. Moreover, one way of constructing such a function is to quotient the graph $\mathcal{C}(G, A)$ by the action of the subgroup of the automorphism group corresponding to H under the canonical isomorphism.*

4.4 Paths, Lines, and Loops

The three examples listed below define types of graphs which will be needed later. To fix notation let $\mathbf{Z}$ be the integers, let $[n] = \{1, 2, \ldots, n\}$ with $[0] = \emptyset$, let $\mathbf{Z}_n = \mathbf{Z}/n\mathbf{Z}$ be the cyclic group of order n, let $\mathbf{D}_{2n}$ be the dihedral group of order $2n$, and let $\mathbf{D}_\infty$ be the infinite dihedral group, which is the relatively free group generated by two involutions.

Example 1 An abstract path, P_n, $n \geq 0$, is an oriented graph with vertices $\cong [n+1]$ and positively oriented edges $\cong [n]$ such that the positive edge i goes from the vertex i to the vertex $i+1$ (see Figure 10). The automorphism group of P_n, $n > 0$, is the cyclic group $\mathbf{Z}_2$, since the only non-trivial automorphism is the morphism which sends the vertex i to $n+2-i$ for all $i \in [n+1]$.

Example 2 An abstract line, L, is an oriented graph with vertices $\cong \mathbf{Z}$ and positively oriented edges $\cong \mathbf{Z}$ such that the positive edge i goes from the vertex i to the vertex $i+1$ (see Figure 11). The automorphisms of L form the group $\mathbf{D}_\infty$, and they can be divided into several types: translations of the graph, reflections which fix a vertex, reflections which fix an edge, and the identity.

Example 3 An abstract loop, L_n, $n \geq 1$, is an oriented graph with vertices $\cong \mathbf{Z}_n$ and positively oriented edges $\cong \mathbf{Z}_n$ such that the positive edge i goes from the vertex i to the vertex $i+1$ (see Figure 12). The automorphisms of

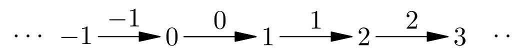

Figure 11: The abstract line, L

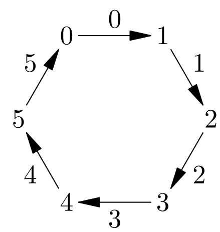

Figure 12: The abstract loop, L_6

L_n form the group $\mathbf{D}_{2n}$, and they can be divided into several types: rotations of the graph, reflections which fix a vertex, reflections which fix an edge, and the identity. If n is odd there is no difference between vertex- and edge-fixing automorphisms; if n is even then there is a difference. In particular notice that a reflection which fixes a vertex or an edge must also fix the graphical element which is directly opposite it. If n is even then for every vertex there is an opposite vertex and for every edge an opposite edge. If n is odd for every vertex there is an opposite edge and vice versa.

An abstract path has two endpoints, namely the vertices 1 and $n+1$, which are called the start and the end of the path respectively, since they are at the start and end of the sequence of positively oriented edges in P_n. The vertex 1 in L and L_n is called its base. All other vertices in P_n, L or L_n are called interior. Let C be a circular category. A finite path is a circular functor from P_n to C or equivalently a graph morphism into the 1-skeleton of C, a loop is a circular functor from L_n to C, and a line or infinite path is a circular functor from L to C. The start (end) of a path is simply the image of the start (end) of P_n and the base of a line or a loop is the image of the base of L or L_n.

The inverse of a path is the same functor from P_n to C but precomposed with the unique non-trivial automorphism of P_n. The inverse of a line or a loop is the same functor but precomposed with the unique non-trivial automorphism which fixes the base. A path, line, or loop is called reduced if the positively oriented edges i and $i+1$ are never sent to edges which are inverses of each other. The degree of a vertex v, written $d(v)$, is the number of oriented edges with v as an initial vertex. A path, loop, or line is called simple iff the graph map into C is injective on vertices. An arc is a simple loop or a maximal simple path in which all interior vertices have degree 2. The length of a path or loop is the number of positive edges in its domain. A path from u to v of shortest length is called a geodesic. A line is called a geodesic if every path it contains is a geodesic. A graph is connected if every pair of vertices can be the endpoints of a path. Every connected graph Γ comes equipped with a metric called the graph metric in which the distance between the vertices u and v is the length of a geodesic from them. If C is a circular category, v is a vertex in C and n is a positive integer, then the ball $\mathrm{Ball}(v,n)$ is the connected subgraph of the 1-skeleton of C consisting of v, all the vertices u whose distance to v in the graph metric is less than n, and all of the edges of C whose endpoints are in $\mathrm{Ball}(v,n)$. The vertex v is called the center of the ball, and n is the radius. Notice that $\mathrm{Ball}(v,0) = \mathrm{Ball}(v,1)$ is the vertex v alone. Sometimes it will be convenient to include in $\mathrm{Ball}(v,n)$ all of the open cones whose entire 1-skeleton is already included under the above definition.

4.5 Labeled Paths and Loops

If the graph P_n is labeled by A then the labels of the positive edges of P_n form a word in A^*, and conversely, every word of length n corresponds to a unique labeling of the positive edges of P_n. Using this correspondence the word W will refer both to a particular finite string of letters and to the associated labeled abstract path. The form intended will be clear from context. The word W is reduced iff its graph is deterministic. If W is a reduced word then its graph is called $\mathrm{str}_1(W)$, the rank 1 straightline construction on W, and if W is not reduced then $\mathrm{str}_1(W)$ is the rank 1 straightline construction of its reduction in the free group.

Similarly, if the graph L_n is labeled by A then the labels of the positive edges form a cycle, and every cycle of length n corresponds to a labeling of the positive edges of L_n. If labelings of L_n which differ by an orientation-preserving automorphism are considered equivalent, then the equivalence classes of labelings of L_n correspond exactly to cycles of length n. Orientation-preserving automorphisms of L_n only serve to change the base and thus produce cyclic conjugates of the word read along the positive edges from the base. The cycle W will also refer to the associated labeled abstract loop. A word is cyclically reduced iff its cycle is reduced iff the corresponding labeling of L_n is deterministic. If the cycle W is reduced then the labeled graph L_n is called $\mathrm{cir}_1(W)$, the rank 1 circular construction on W, and if the cycle W is not reduced (and doesn't freely reduce to the identity), then $\mathrm{cir}_1(W)$ is the rank 1 circular construction of its cyclic reduction in the free group.

Every path in a labeled graph Γ induces a labeling of P_n and thus corresponds to a word $W \in A^*$. The word W is called the word read by the path, the path reads W, and W is readable in Γ. Similarly, a loop in a labeled graph induces a labeling of the abstract loop which corresponds to a cycle, say W. The cycle W is called the cycle read by the loop, the loop reads W and the cycle W is readable in Γ. A path which reads the word W is called the path W, and a loop which reads the cycle W is called the loop W. At this point four distinct situations have been defined which can be indicated by the same finite string of letters, $W \in A^*$, namely, the word W, the cycle W, the path W, and the loop W. To summarize the distinctions: words and cycles refer to labeled abstract paths and loops, respectively, while paths and loops refer

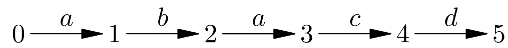

Figure 13: A rank 1 straightline construction

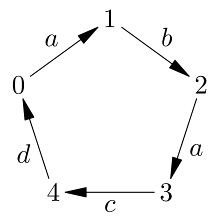

Figure 14: A rank 1 circular construction

to implicitly understood, label-preserving functors from a word or cycle into a labeled complex of some kind.

Example 4 In Figure 13 the rank 1 straightline construction str_1 is illustrated for the word $W = abacd$. Notice that it is also the rank 1 straightline construction of the word $abee^{-1}acd$.

Example 5 In Figure 14 the rank 1 circular construction cir_1 is illustrated for the word $W = abacd$. Notice that it is also the rank 1 circular construction of the word $eabacde^{-1}$.

An automaton is a labeled graph with a specified start state and one or more specified end states. In this article all automata considered will have exactly one end state. A word W is accepted by an automaton Γ if it is read by a path in Γ which starts at the start state and ends at one of the end states. The inverse of a word W, denoted W^{-1}, corresponds to the same labeled graph but with the opposite orientation. The inverse of a cycle, a path, or a loop is defined analogously. If a word is its own inverse then the non-trivial automorphism of P_n is label-preserving. If n is odd then the middle edge is fixed, and the alphabet is non-oriented, and if n is even then the middle two edges show that the word is not reduced. Thus no reduced word over an oriented alphabet is its own inverse. If W is a word read by path from u to v, then W^{-1} is read from v to u. Let X be a path from u to v and let Y be a loop based at v. The loop Y can be conjugated by X to produce a loop based at u. Specifically, the conjugation of Y by X is a loop which reads the word XYX^{-1}.

4.6 Approximations by Paths

A path, in the topological sense of the word, is simply a continuous function from the unit interval $[0, 1]$ to a topological space. The Simplicial Approximation

Theorem on simplicial complexes can be used to show that every topological path between vertices of a circular category is homotopic relative to its endpoints to a path as defined above.

Lemma 4.7 *If C is a circular category and $f : [0, 1] \to C$ is a topological path which starts and ends at vertices in the space C, then f is homotopic relative to its endpoints to a simplicial path in the 1-skeleton of C.*

Proof: The crucial properties of circular cones are that, like simplices and cells, they are contractible spaces with path-connected boundaries when the rank is 2 or more. By Lemma 2.12 the circular category C can be subdivided so that the result is a simplicial complex. The traditional simplicial approximation theorem can then be applied to the subdivided complex. The finite number of vertices in the new path are each contained in a unique open cone of C by Lemma 1.16. Let k be the maximum rank of an open cone of C containing a vertex of the path. This implies that the new path is in the k-skeleton of the original category C. If $k \geq 2$, then the path can be shown to be homotopic to a path in the $(k-1)$-skeleton as follows.

Let c be an open cone of rank k which contains a vertex v of the new path. The edges on either side of v must also be contained in c. By looking far enough in each direction there is a unique vertex in the path which is farthest from v, in the boundary of c, and with all edges and vertices in the path between it and v in the open cone c. Let v_1 and v_2 be these vertices on either side of v. The portion of the path between v_1 and v_2 is in the closed cone C/c and it starts and ends in the boundary of c. Since the closed cone c has a path-connected boundary and is contractible and thus simply connected, there is a path between v_1 and v_2 which is contained in ∂c and is homotopic to the other path, keeping the endpoints fixed. By again applying the simplicial approximation theorem to the new path between v_1 and v_2 in the lower skeleton, the new path is also simplicial.

Technically speaking, the image of the closed cone C/c under the attaching functor may not be either simply connected or contractible. The homotopic path is actually to be found by lifting the portion between v_1 and v_2 into the closed cone C/c, choosing a path in the boundary, applying the Simplicial Approximation Theorem to the path in the boundary, finding a homotopy to the new path, and then pushing the homotopy, via the attaching map, back down into the construction. If this procedure is carried out for each of the open cones of rank k, then the maximum rank of a vertex in the path will be $k-1$, and the path is now in the simplicial subdivision of the $(k-1)$-skeleton of C.

Continuing in this way, the path will eventually be a path in the 1-skeleton. At this point, a similar procedure is followed. For all vertices v in the path which are not vertices of C, there is an open edge c containing v and vertices v_1 and v_2 in the boundary ∂c as defined above. If the lift of the path between v_1 and v_2 to the closed edge C/c lifts v_1 and v_2 to the same endpoint, then the path between v_1 and v_2 is homotopic to the constant path, and thus the original path is homotopic to the path obtained by removing the vertices and edges between v_1 and v_2 and connecting the two sections. If, on the other hand, v_1 and v_2 are at opposite ends of the edge when lifted, then the portion of the path between v_1 and v_2 is homotopic to a nice linear function between the subinterval of $[0, 1]$ and the edge c. When this is done for all vertices in the path which are not vertices in the original category C, the result is a function from unit interval $[0, 1]$ to the 1-skeleton of C which by appropriately partitioning $[0, 1]$ is a circular functor into the 1-skeleton of C. $\Box$

Lemma 4.8 *A circular category is path-connected as a topological space iff its 1-skeleton is a connected graph.*

Proof: Lemma 4.7 shows that being topologically connected implies that the 1-skeleton is connected. To see the converse, notice that by Lemma 1.16 every point is contained in an open cone and thus contained in the image of a closed cone. Moreover, since a closed cone is path-connected and contains a vertex, every point in the space of the circular category is connected by a path to a vertex. These paths combined with the connectedness of the graph prove that the topological space is path-connected. $\Box$

A circular category which satisfies either condition of Lemma 4.8 will be called a connected circular category.

5 General Relators

In this section a generalization of the traditional notion of a relator is defined and some of the basic properties of these general relators are examined. More specifically, the lemmas in this section show that all paths in the boundary of a general relator beyond a certain minimum length can be assigned an orientation relative to the boundary, and that under relatively mild conditions, the shortest loops with winding number 1 must be simple. The section concludes with specialized results on cycles which are piecewise geodesic, relator metrics, and the structure of the automorphism groups of relators.

5.1 Length and Representatives

A generator is an edge labeled by A. A relator in the usual sense can be thought of variously as a cyclically reduced word $W \in A^*$, as a cyclically reduced word $W \in A^*$ together with all cyclic conjugates of W and its inverse, as a cycle, as a deterministically labeled loop, or as a 2-dimensional cell with a deterministic labeling by A. The last description is the most geometric and the one which is used as the basis of the generalization which follows.

A general relator over A is defined to be a circular cone of rank $n > 1$ with a deterministic labeling by A. Note that a general relator is, by definition, a circular complex; it cannot be a circular category which is not also a circular complex. This extends the usual notion of a relator in a group. The restriction that a general relator be a circular complex corresponds to the usual restriction that relators be cyclically reduced. If C is a labeled circular category, and c is an object in C of rank at least 2, then $R = C/c$ is a general relator which inherits its labeling from the attaching map $\phi_c : R \to C$. The boundary of R, written ∂R, is defined to be ∂c, the boundary of c.

Since ∂R is homotopically equivalent to a circle, this implies, among other things, that $\pi_1(\partial R, v) \cong \mathbf{Z}$ for every vertex v in ∂R. Since there are two possible isomorphisms of $\mathbf{Z}$ with itself, there are two possible isomorphisms between $\pi_1(\partial R, v)$ and $\mathbf{Z}$ which are called the two orientations of the relator, R. If a particular isomorphism from $\pi_1(\partial R, v)$ to $\mathbf{Z}$ is chosen then an integer can be assigned to every oriented loop based at v and by path conjugation to every oriented loop in ∂R. Since every loop in R is contained in its 1-skeleton and thus in ∂R, all loops in R are assigned an integer. If the other isomorphism with $\mathbf{Z}$ is chosen then the integer assigned to a particular loop will be opposite in sign. Thus the absolute value of this integer is a characteristic of the unoriented loop, which will be called the winding number of the loop. By the definition of π_1 it is known that a loop in a circle is contractible or homotopic to a point iff it has a winding number of 0. If both the loop and the relator are oriented then a plus or a minus sign is added to the winding number according to whether it is orientation-preserving or -reversing. A loop in R is called a representative of R if it has a winding number of 1, regardless of sign.

Lemma 5.1 *Every general relator has a representative. Specifically, if R is a general relator, then R contains a loop in its 1-skeleton which has winding number 1.*

Proof: Let $f : [0, 1] \to \partial R$ be a topological path which starts and ends at the same point and has winding number 1 as a topological loop. Since the boundary of R is topologically connected f is conjugate to a loop which starts and ends at a vertex in ∂R. Thus, without loss of generality, assume this is true of f. By Lemma 4.7, f is homotopic relative to its endpoints to a path in the 1-skeleton of ∂R. Since the endpoints remain fixed, the loop formed in the 1-skeleton is homotopic relative to its basepoint to the loop formed by f. Because winding numbers are invariant under such homotopies, the proof is complete. □

Lemma 5.2 *For every general relator R there exists a simple loop U with some winding number, say n, such that for all loops V with winding number m,*

$$\frac{|U|}{n} \leq \frac{|V|}{m}$$

Proof: First of all, by Lemma 5.1 there exists at least one loop with a non-trivial winding number. If this loop is not simple then it can be decomposed into two loops of shorter length, one of which again has a non-zero winding number. Continuing in this way a simple loop with non-zero winding number. The process must stop since the loops are getting shorter. Since the relator R is finite, there are only a finite number of simple loops and thus there exists a simple loop U such that the length of U divided by its winding number is the smallest possible among all simple loops with non-zero winding number in R. The ratio of the length of the simple loop U to its winding number is actually minimal among all loops with non-zero winding number in R. The argument proceeds by contradiction.

Suppose that there exists a loop V such that the length of V divided by its winding number is strictly less than the length of U over n. Suppose further that V is selected so that the length of V is the shortest possible among all counterexamples. Since U yields the smallest value among all simple loops, V must be non-simple. Let V_1 and V_2 be loops read in R so that the loop V_1 followed by the loop V_2 is the loop V. If the orientation of the loops V_1 and V_2 are identical then since $\frac{a+c}{b+d}$ is between $\frac{a}{b}$ and $\frac{c}{d}$ for all positive numbers a, b, c, and d, the ratio of length to winding number for at least one of the two smaller loops must be at least as small as that of V, contradicting the minimality of the counterexample V. If, on the other hand, the orientations of the two loops are opposite in sign, then one of the loops has a shorter length and a winding number at least as large, so that its length divided by its winding number is smaller than the ratio for V, again contradicting the minimality of V as a counterexample. Since all possibilities lead to contradictions, it must be that no counterexample exists. □

The length of the general relator R, written $|R|$, is defined to be $\frac{|U|}{n}$, where U and n are as described in Lemma 5.2. A simple loop U with winding number n for which $|U| = n|R|$ is called a standard geodesic of R.

Corollary 5.3 *If V is a loop in a general relator R and m is its winding number then $|V| \geq m|R|$.*

The universal cover of an abstract labeled loop W is called W^∞ since it is an abstract labeled line which reads the bi-infinite word

$$\ldots WWWW \ldots$$

By analogy, if R is a general relator, then the universal cover of ∂R is called R^∞. A path W in ∂R which lifts to a geodesic path in R^∞ will, by an abuse of notation, be called a geodesic. This is equivalent to requiring that the path W be the shortest path in its homotopy class relative to its endpoints. Similarly, a loop W in ∂R is called a geodesic in ∂R if there does not exist a shorter loop in ∂R which is homotopic to W. Since by Corollary 5.3 U is the shortest loop of winding number n, and since loops in circles are homotopic iff they have the same winding number, U certainly deserves to be called a geodesic loop. It is important to realize that geodesic loops in the boundary of a general relator R do not necessarily lift to geodesic loops or lines in the regular covers of ∂R. See Example 1 in Section 6. The standard geodesic U, however, does have this property.

Lemma 5.4 *If U is a standard geodesic of a general relator R, then the natural map from U^∞ to ∂R lifts a line U^∞ in R^∞ which is a geodesic.*

Proof: A path between points in the line U^∞ whose length is strictly less than the portion of U^∞ demarcated by these points can be used to create a loop in ∂R which violates the minimality of the loop U as guaranteed by Lemma 5.2. □

5.2 Ends and Orders

Let R be a general relator, let R^∞ be the universal cover of its boundary, and let $f : R^\infty \to \partial R$ be the covering map. Since $\pi_1(\partial R, v) \cong \mathbf{Z}$ by definition, there is a 1 to 1 correspondence between the vertices of $f^{-1}(v)$ and the integers $\mathbf{Z}$. In particular, if W is an oriented representative of R based at v and v_0 is a vertex with $f(v_0) = v$, then define $v_n \in f^{-1}(v)$ to be the vertex which is the endpoint of the lift of W^n to R^∞ starting at v_0. The following lemma is immediate.

Lemma 5.5 *If R is a general relator, $f : R^\infty \to \partial R$ is the covering map, W is an oriented representative of R based at a vertex v, and v_0 is a vertex in R^∞ with $f(v_0) = v$, then for all integers n the lift of the path W starting at v_n ends at v_{n+1}.*

In order to facilitate the proofs in this subsection it is convenient to assign a unique subscript to every vertex (and open cone) in R^∞. This can be done efficiently by choosing a spanning tree in R. Let T be a fixed spanning tree of the 1-skeleton of R. More generally, T can be extended to a spanning tree T' of the poset which defines the cone complex ∂R. Since T' is a tree and thus simply connected, it follows by Lemma 1.13 and Lemma 1.14 that there are uniquely defined lifts $g_i : T' \to R^\infty$ which send v to v_i. The same lemmas can be used to show that the images of these lifts partition the open cones in R^∞ in the following sense.

Lemma 5.6 *If R is a general relator, and f, g_i, and T' are as described above, then the open cones in the images $g_i(T')$ are disjoint in R^∞ and their union contains every open cone in R^∞.*

For each open cone u in R, define $u_i = g_i(u)$. Since the images $g_i(T')$ partition the open cones in R^∞, this definition assigns a unique subscript to each open cone in R^∞.

Lemma 5.7 *Let R be a general relator, and let f, v, W and v_i, $i \in \mathbf{Z}$ be as described above. If a finite portion of the 1-skeleton of R^∞ is removed, then there is a large positive integer N, based on n and R, such that the open cones u_i are in the same connected component of the result for all $i \geq N$ and for all open cones u in R. Similarly, there is a large negative integer N' such that the open cones u_i are also in a connected component of the result for all $i \leq N'$ and for all u in R.*

Proof: Let j be the largest subscript assigned to a open cone in the ball $\text{Ball}(v_0, n)$. Next, consider the path which starts at v_0 and reads the word W, and let k be the absolute value of the smallest subscript assigned to a vertex or edge in this path. The statement of the lemma will be true for all N strictly bigger than $j+k$. To see this, notice that $g_i(T')$ is connected and all of the open cones involved has the same subscript so that for all $i > j$ and for all $u \in R$, u_i is in the same connected component as v_i. And finally, v_i is connected to v_{i+1} by a path reading W and all of the vertices and edges involved have subscripts of at least $i-k$. An analogous argument works for large negative subscripts. □

When a finite portion of the 1-skeleton of R^∞ is removed, then the connected component of the result which contains the vertices v_i for arbitrarily large positive integers i will be called the positive end, and it will be denoted $R^{+\infty}$. Similarly, the connected component of the result which contains the vertices v_i for arbitrarily large negative integers i will be called the negative end, and it will be denoted $R^{-\infty}$. Notice that these do not necessarily describe distinct connected components. When there is some finite portion which when removed causes these components to become distinct, then these ends are said to be disconnected. The labeling of these components as ends conforms to the use of this term by John Stallings ([19]). The following Corolllary is recorded for later use. It states that the positive and the negative end are the only possible ends.

Corollary 5.8 *If R is general relator and a finite number of open cones are removed from R^∞ then there are at most two connected components which contain an infinite number of open cones. In particular the only connected components with an infinite number of open cones are the (possibly identical) ends $R^{+\infty}$ and $R^{-\infty}$.*

Proof: This follows immediately from Lemma 5.7 once it is noticed that there are only a finite number of open cones which are assigned to each subscript. □

Now that the ends of R^∞ have been defined, the next several lemmas will investigate the ordering induced on disjoint, connected pieces of R^∞ which disconnect the ends. Each finite B_i under consideration will be the finite union of open cones. It may not be a subcomplex since the boundary of an open cone in B_i need not be included in B_i. Also, notice that B_i is required to be connected in the topological sense, but that the portion of the 1-skeleton in B_i need not be connected. Such a B_i is said to be connected to $R^{+\infty}$ if there are (topological) paths which start at open cones in B_i and end at vertices of arbitrarily large positive subscript. Similarly, B_i is connected to $R^{-\infty}$ if there are paths which start at open cones in B_i and end at vertices of arbitrarily large negative subscript. When no restrictions are placed on the connecting paths it is clear that all B_i are connected to both $R^{+\infty}$ and $R^{-\infty}$, but once restrictions are introduced, this is no longer the case.

Lemma 5.9 *For $i = 1, 2$, let B_i be a finite union of open cones in R^∞ which is topologically connected and whose removal disconnects the ends $R^{+\infty}$ and $R^{-\infty}$. If B_1 and B_2 are disjoint, then the following statements are equivalent:*

(1) B_1 can be connected to $R^{-\infty}$ by paths which is disjoint from B_2.

(2) B_1 cannot be connected to $R^{+\infty}$ by paths which is disjoint from B_2.
(3) B_2 can be connected to $R^{+\infty}$ by paths which is disjoint from B_1.
(4) B_2 cannot be connected to $R^{-\infty}$ by paths which is disjoint from B_1.

Proof: If (1) is true, and (2) is false, then the B_2 does not disconnect the ends, contradicting the hypothesis of the lemma. Thus $(1 \Rightarrow 2)$. A similar argument shows that $(3 \Rightarrow 4)$. Next, let U^∞ be the lift to R^∞ of the standard geodesic of R, and consider $B_i \cap U^\infty$. The points in the intersection are at most a bounded distance apart, bounded by the finite diameter of B_i. Thus there is a last vertex or edge in U^∞ which is contained in B_i. Notice that there must be at least one vertex or edge in the intersection since B_i disconnects the ends and U^∞ connects them. Moreover, the last vertices/edges in U^∞ of B_1 and B_2 are distinct since B_1 and B_2 are disjoint. Finally the B_i which contains the later vertex/edge is connected to the end $R^{+\infty}$ by the remaining portion of U^∞. This shows that at least one of the B_i is connected to R^∞ by paths which are disjoint from the other, and as a consequence $(2 \Rightarrow 3)$. An analogous argument shows $(4 \Rightarrow 1)$. □

If any of the four equivalent conditions are satisfied then write $B_1 < B_2$. Notice that $B_1 < B_2$ implies (assumes) that they are disjoint.

Corollary 5.10 *For $i = 1, 2$, let B_i be a finite union of open cones in R^∞ which is topologically connected and whose removal disconnects the ends $R^{+\infty}$ and $R^{-\infty}$. If $B_1 < B_2$, then $B_2 \not< B_1$. Moreover, if B_1 and B_2 are disjoint, then either $B_1 < B_2$ or $B_2 < B_1$.*

Proof: The first statement is immediate since condition (2) for the inequality $B_1 < B_2$ and condition (3) for the ineqality $B_2 < B_1$ directly contradict each other. The second statement is immediate from the proof of Lemma 5.9. □

Lemma 5.11 *For $i = 1, 2$, let B_i be a finite union of open cones in R^∞ which is topologically connected and whose removal disconnects the ends $R^{+\infty}$ and $R^{-\infty}$. If $B_1 < B_2$ and $B_2 < B_3$, then B_2 is 'between' B_1 and B_3 in the sense that any path connecting a point in B_1 to a point in B_3 must contain a point in B_2. In particular B_1 and B_3 are disjoint, and moreover, $B_1 < B_3$.*

Proof: By the definition of $<$, B_1 is connected to $R^{-\infty}$, and B_3 is connected to $R^{+\infty}$ by paths which are disjoint from B_2. If B_1 were connected to B_3 by a path disjoint from B_2, then B_2 would not disconnect the ends. This contradiction proves the first statement. If B_1 and B_3 are not disjoint, the the empty path connects the common point to itself. Finally, notice that by the first part of the lemma, the paths connecting B_1 to $R^{-\infty}$ which are disjoint from B_2, must also be disjoint from B_3. □

For $i = 1, 2$, let B_i For $i = 1, 2$, let B_i be a finite union of open cones in R^∞ which is topologically connected and whose removal disconnects the ends $R^{+\infty}$ and $R^{-\infty}$.

Corollary 5.12 *Let $F = \{B_i | i \in I\}$ be a collection where each B_i is a finite union of open cones in R^∞ which is topologically connected and whose removal disconnects the ends $R^{+\infty}$ and $R^{-\infty}$. If the elements of F are pairwise disjoint, then $<$ defines a discrete total linear order on F which is order isomorphic to the integers, the positive integers, the negative integers, or the integers from* 1 *to n.*

Proof: The ordering is total and linear as an immediate consequence of Corollary 5.10 and Lemma 5.11. To see that it is discrete, let B_1 and B_2 be arbitrary members of F with $B_1 < B_2$. Since $B_1 \cup B_2$ contains only a finite number of open cones, by Corollary 5.8 there are only a finite number of open cones which are not in either B_1, B_2, $R^{+\infty}$ or $R^{-\infty}$. In particular, there are only a finite number of members of F which contain one of these open cones, and thus only a finite number of C in F with $B_1 < C < B_2$. Since F is totally ordered, there is a well-defined member of F which comes after B_1 and a well-defined member of F which comes before B_2. The final comment simply lists the possible discrete total orders. □

5.3 Width and Orientation

The width of a general relator R, written ω_R, can be described as the smallest non-negative integer such that the removal of a ball of radius ω_R centered at any vertex in R^∞ will disconnect the ends. Let u be a vertex in R, u_0 a lift of u to R^∞, W an oriented representative, and n a positive integer. If the removal of a ball of radius n centered at u_0 disconnects the ends of R^∞, then the ends will also be disconnected by the removal of a ball of radius n centered at any other lift of u. Also notice that the orientation of W and even the choice of the representative W is irrelevant since the reversal of the orientation merely switches the labels on the ends, and another representative would leave the lifts of u identical. By the above reasoning, it is sufficient to check that the ends are disconnected for a single lift of each of the finite vertices in R instead of the infinite number of vertices in R^∞. The following lemma shows that such a number always exists.

Lemma 5.13 *Let R be a general relator, and let f, v, W and v_i, $i \in \mathbf{Z}$ be as described above. There is a unique smallest integer ω_R such that for all vertices u in R^∞, the removal of a ball of radius ω_R centered at u will always disconnect the ends.*

Proof: The existence of such a unique minimum non-negative integer will be immediate once it is shown that there is at least one which satisfies the conditions. To begin let T be the fixed maximal spanning tree described in the previous section and let k be the maximum distance between vertices in the tree T, so that u_i and w_i are within k units of each other for all u and w in R, even when the possible paths are restricted to those which only use vertices with subscript i. Next, consider the edges of R. If a particular edge is lifted to R^∞ then its endpoints have subscripts, and the absolute value of the difference between these subscripts is a value which is independent of the lift of the edge. This is because any two lifts differ by an automorphism of R^∞ and the automorphism which sends v_0 to v_i also sends u_j to u_{i+j} for all $u \in R$ and for all $j \in \mathbf{Z}$. Since R is finite, there are only a finite number of edges in R and thus a maximum absolute change in the value of a subscript when traversing a single edge. Call this largest value m.

Finally, if n is chosen so that n is greater than $2k + m|W|$, then the ball $\mathrm{Ball}(u_i, n)$ contains all vertices with subscripts in the range $[i - m, i + m]$. This can be seen by tracing a path from u_i to v_i to the v_j with the correct subscript, and then to the final vertex with subscript j. These can be chosen so that their lengths are at most k, $|i-j| \cdot |W|$, and k, respectively. This must disconnect the vertices with large negative subscripts from those with large positive subscripts. For suppose otherwise: let X be a path from any vertex with a subscript less than $i - m$ with a vertex with subscript greater than $i + m$. Since it cannot contain any vertices with a subscript in $[i - m, i + m]$, it must contain a single edge which starts at a vertex with subscript less than $i-m$ and ends at a vertex with subscript greater than $i+m$, but this contradicts the definition of m. Since u and i are arbitrary, this shows that there is at least one n which works, and this completes the proof. □

If $4\omega_R \leq |R|$ then the general relator R is called thin. Let U be a path in a general relator R from v to u and let W be an oriented representative of R based at v, and let $v_i, i \in \mathbf{Z}$ be defined as above using the word W. If the path U in R^∞ which starts at v_0 ends at a vertex in the same connected component as the vertices v_i with large positive subscripts once the ball $\mathrm{Ball}(v_0, \omega_R)$ has been removed from the 1-skeleton of R^∞, then U is said to have the same orientation as W, or rather that U is positively oriented relative to W. If U ends at a vertex in the same connected component as the vertices v_i with large negative subscripts then U is said to have the opposite orientation as W, or that U is negatively oriented relative to W. If neither is the case, and this possibility does occur, then U is called unoriented relative to W.

Lemma 5.14 *Let R be a general relator and let U be a path in R whose lift to R^∞ has endpoints which are a distance of at least $2\omega_R$ apart. If U starts at a vertex v in R, and W is an oriented representative of R which is based at the same vertex v then U is oriented either positively or negatively relative to W.*

Proof: Let the vertices v_i, $i \in \mathbf{Z}$ be defined as usual from W and a particular lift v_0 of v. By the definition of ω_R it is known that when the ball $\mathrm{Ball}(v_0, \omega_R)$ is removed from the 1-skeleton of R^∞, the v_i with large negative and the v_j with large positive subscripts are contained in distinct connected components, but it is not known whether there are other connected components. The proof will show that all of the vertices in the other connected components must be very close to v_0. Let u be the endpoint of the path U lifted to R^∞ so that it starts at v_0, and assume that u is in a connected component other than the two expected ones.

Consider a path in R^∞ from a vertex v_i with a large negative subscript to a vertex v_j with a large positive subscript. If this path contains a vertex in a component other than the two expected ones and the ball $\text{Ball}(v_0, \omega_R)$, then there must be a subpath which starts and ends in $\text{Ball}(v_0, \omega_R)$ which contains the vertex in the other component. Since $\text{Ball}(v_0, \omega_R)$ is connected (through paths to and from v_0), this subpath can be removed and replaced by a subpath completely in $\text{Ball}(v_0, \omega_R)$. Continuing in this way, the path from v_i to v_j will eventually contain only vertices in the components of v_i and v_j, and the ball $\text{Ball}(v_0, \omega_R)$.

The assumption about the path U implies that the ball $\text{Ball}(u, \omega_R)$ does not contain any of the vertices in $\text{Ball}(v_0, \omega_R)$. But this in turn means that $\text{Ball}(u, \omega_R)$ does not contain a vertex of the path described above, and so does not disconnect the component of v_i from the component of v_j, contradicting the definition of ω_R. Thus U must end in one of the two standard components. □

Up until this point, the orientation of a path U has only been defined relative to an oriented representative of R. The orientation of a path U relative to an unoriented representative W will now be defined, provided that U is a subword of the cycle W. The definition goes as follows. Let W' be the cyclic conjugate of the cycle W or W^{-1} such that U is an initial segment of the word W', that is, so that $W' = UV$ for some word V. The orientation of U relative to the unoriented representative W is then the orientation of U relative to the oriented representative W'. Using this definition, it is possible to have a representative loop $W = XUYU^{-1}$ and to have both instances of the word U be positively oriented relative to W. Intuitively, the orientation of U is positively oriented relative to W if the subword U, as it is situated in W, helps to complete the loop, and it is negative if, as it is situated in W, goes in the opposite direction.

5.4 Simple Representatives

Although by Lemma 5.1 every general relator contains a representative loop, it is not true that this representative must be simple or even that a simple representative must exist at all. A presentation of $\mathbf{Z}$ is given in Section 6 which will provide a counterexample. The lemmas below show that if, however, a general relator is thin, then there does have to exist a simple representative, which can then be used to create a deformation retraction.

Lemma 5.15 *For all geodesic paths X read in R, there exist paths Y, P and Q in R such that the path Y is a geodesic read in a standard geodesic loop U, the lengths of P and Q are each less than ω_R, and the cycle $XPYQ$ is readable in R as a contractible loop. Since both X and Y are geodesic, $||X| - |Y|| \leq |P| + |Q| < 2\omega_R$.*

Proof: Lift the path X to R^∞ and consider its relation to the line U^∞, the lift of a standard geodesic loop U. For convenience the lift of X will also be called X. Since the balls of radius ω_R disconnect R^∞, there is a vertex or edge of U^∞ contained in the ball centered at the start of X. And since edges are included in an open ball only if its endpoints are, there must be a vertex. If Q is the geodesic path from the vertex in U^∞ to the start of X, then the length of Q is strictly less than ω_R. Similarly, there is a path P from the endpoint of X to a vertex of U^∞ whose length is strictly less than ω_R. By Lemma 5.4 the portion of U^∞ marked off by the end of P and the start of Q is a geodesic path which will be called Y. Since $XPYQ$ is a loop in the universal cover, it is contractible. The homotopy which proves the contractibility of $XPYQ$ in R^∞, when composed with the projection down into the boundary of R, provides a contraction of the image of the loop $XPYQ$ in the boundary of R. Finally, X geodesic implies that $|X| \leq |Y| + |P| + |Q|$ and Y geodesic implies that $|Y| \leq |X| + |P| + |Q|$. These inequalities can be rearranged to show that $||X| - |Y|| \leq |P| + |Q| < 2\omega_R$. □

The next lemma is a technical result which will be used to prove Lemma 5.17.

Lemma 5.16 *Let $X = X_1X_2$ be a geodesic loop in a general relator R with $|X_i| \geq 2\omega_R$ for $i = 1, 2$. Let x_0 be a lift of x to R^∞, let x_i be the endpoint of the path reading X^i which starts at x_0, let x_i' be the endpoint of the path reading X_1 starting at x_i, and let B_i and B_i' be the open balls of radius ω_R centered at x_i and x_i' respectively. The following four implications are true:*

(1) If $B_i < B_i'$ then $B_{i-1}' < B_i < B_i' < B_{i+1}$
(2) If $B_i' < B_i$ then $B_{i+1} < B_i' < B_i < B_{i-1}'$
(3) If $B_i' < B_{i+1}$ then $B_i < B_i' < B_{i+1} < B_{i+1}'$
(4) If $B_{i+1} < B_i'$ then $B_{i+1}' < B_{i+1} < B_i' < B_i$

In particular, $B_0 < B_0'$ implies that $B_i < B_j$ for all integers $i < j$, and $B_0' < B_0$ implies that $B_j < B_i$ for all integers $i < j$.

Proof: Since all of these implications have essentially identical proofs, it is sufficient to show that $B_i < B_i'$ implies $B_i' < B_{i+1}$. First of all, each of these three balls is a finite subgraph of the 1-skeleton of R^∞ whose removal disconnects the infinite ends by the definition of ω_R and by Lemma 5.13. Next, B_i' and B_{i+1} are disjoint since X_2 is the shortest path between their centers and it was to have a length of at least $2\omega_R$. Similarly B_i and B_i' are disjoint and B_i and B_{i+1} are disjoint, since the shortest distances between their centers are $|X| \geq m|R| \geq |R| \geq 4\omega_R$ and $|X_1| \geq 2\omega_R$, respectively. Therefore, by Corollary 5.12, B_i, B_i' and B_{i+1} are linearly ordered. Without loss of generality, assume that $B_i < B_{i+1} < B_i'$. By Lemma 5.11, the path X_1 from x_i to x_i' must contain a point in B_{i+1}. Let $X_1 = X_{11}X_{12}$ be a partition of X_1 so that the endpoint of X_{11} lies in B_{i+1}. The fact that $X_{12}X2$ is a subword of X guarantees that it is a geodesic, its least is at least that of X_2 which in turn is at least $2\omega_R$, but since it starts and ends in B_{i+1}, there is a path connecting the startpoint to the endpoint of length strictly less than $2\omega_R$, contradiction. Thus the assumption that $B_{i+1} < B_i'$ must have been false. The final statements follow immediately by iterating the other implications. □

Lemma 5.17 *For all geodesic loops X in a thin relator R, $||X| - m|R|| < 2\omega_R$ where m is the winding number of the loop X, and moreover the winding number* *of X is uniquely determined by this inequality. In particular, $\frac{|X|}{|R|}$ rounded to the nearest integer is the winding number of X.*

Proof: Let U be a standard geodesic loop read in R and let n be its winding number. Orient the loops U and X so that they induce the same orientation on R. Next, let x be a point in the loop X, let P be a path of minimal length from x to any point u in U, and without loss of generality assume that X and U are the words determined by the vertices x and u. Since the loops X and U are geodesic loops, by definition X and U are geodesic words. Since the loops X^n and U^m both have the same winding number, namely, nm, it follows that $X^n P U^{-m} P^{-1}$ is a loop of winding number 0.

Let x_0 be any point in universal cover R^∞ which is sent to x by the covering map f. Next, let u_0 be the endpoint of the path P, lifted to R^∞ so that it starts at x_0. Then define x_i and u_i as the endpoints of the paths X^i and U^i lifted to R^∞ to start at x_0 and u_0, respectively. Since the loop $X^n P U^{-m} P^{-1}$ has winding number 0 in ∂R, it lifts to a loop based at x_0 in R^∞. Notice that this shows that the path P lifts so that it starts at x_n and ends at u_m. If the loop $X^n P U^{-m} P^{-1}$ is lifted so that it starts at x_n then we can conclude that the path P lifted to start at x_{2n} will end at u_{2m}. In general, the vertex x_{jn} is connected to the vertex u_{jm} by a path whose length is at most the length of P, for all integers j. The above steps could be repeated with a loop Y in ∂R of winding number 1, based at a vertex y and with a path Q starting at y and ending at x. The same argument would show that y_{jmn} is connected to x_{jn} by a path whose length is at most the length of Q, and it is connected to the vertex u_{jm} by a path whose length is at most the length of QP. This will now be used to show that the vertices x_i for large positive i are contained in the positive end $R^{+\infty}$.

Consider the removal of the ball of radius ω_R centered at x_0. By the definition of ω_R this ball disconnects the ends of R^∞. Since u_{jm} for large positive values of j is within $|QP|$ of y_{jmn}, and since by Lemma 5.4 U^∞ is an infinite geodesic, it follows that j can be chosen large enough so that the path QP from y_{jmn} to u_{jm} is disjoint from the ball around x_0. If this were false, then $|QP| + \omega_R$ would bound $|U^{jm}|$ for all j, contradiction. Similarly, there is a large negative k such that the path QP from y_{kmn} to u_{km} is disjoint from the ball centered at x_0. This shows specifically that the path $U^{(j-k)m}$ connecting u_{km} to u_{jm} in R^∞ is a path connecting the two disconnected ends $R^{+\infty}$ and $R^{-\infty}$, and must therefore pass through the ball centered at x_0. In particular, there is a point in U^∞ which is connected to x_0 by a path of length less than ω_R. Since the path P was chosen to have minimal length, this shows that $|P| < \omega$.

Next, by repeating the above arguments with x_i replacing x_0 as the center of the ball to be removed, there always exists some vertex in U^∞ which is connected to x_i by a path of length less than ω_R. Define v_i to be the last vertex or edge in $\mathrm{Ball}(x_i, \omega_R) \cap U^\infty$. Since edges are contained in open balls only when their endpoints are as well, the last vertex or edge is in fact a vertex. Let P_i denote the shortest path from x_i to v_i. Notice that $|P_i| < \omega_R$ for all i. Also note that by symmetry v_0 and v_n are connected by a cyclic conjugate of the word

U^m whose length is exactly $|U^m| = m|U| = mn|R|$. The second equals sign follows from the definition of $|R|$ as the length of its standard geodesic divided by its winding number. Since U^∞ is a geodesic in R^∞, it follows that $mn|R|$ is the minimum distance in R^∞ from v_0 to v_n.

Because R is thin and X is a geodesic loop, $|X| \geq m|R| \geq |R| \geq 4\omega_R$. In particular, there is a partition of $X = X_1X_2$ such that $|X_i| \geq 2omega_R$ for $i = 1, 2$. Let B_i be the open ball $\text{Ball}(x_i, \omega_R)$. Since it is known that $B_0 < B_m$, it follows by Lemma 5.16 that $B_i < B_j$ for all $i < j$. In particular, the vertices v_i occur in order within U^∞, and since U^∞ is a geodesic the length of the minimal path from v_0 to v_n is exactly the minimal distance from v_0 to v_1 plus the minimal distance from v_1 to v_2 plus $\ldots$ plus the minimal distance from v_{n-1} to v_n. This in turn means that there must be some subscript i $(0 \leq i < n)$ such that the distance from v_i to v_{i+1} is at most $m|R|$. For this i identified above, there is a loop $XP_{i+1}VP_i^{-1}$ which goes from x_i to x_{i+1} to v_{i+1} to v_i and back to x_i. The word V (or its inverse) is a subword of U^∞. Since X is a geodesic loop in the general relator R, the path X lifts to a geodesic in R^∞ (since any shorter path in R^∞ is mapped under the covering map f to a shorter path in R, contradiction). Thus, $|X| \leq |P_i| + |P_{i+1}| + |V| \leq 2\omega_R + m|R|$ by the above estimates. On the other hand $|X| \geq m|R|$ by Corollary 5.3. Finally, $m|R| \leq |X| < m|R| + 2\omega_R$, or $||X| - m|R|| \leq 2\omega_R$, as was to be shown. If in addition R is a thin relator, then by definition $|R| \geq 4\omega_R$, and in particular the above inequalities become $m|R| \leq |X| < m|R| + 2\omega_R \leq m|R| + \frac{1}{2}|R|$, or $m \leq \frac{|X|}{|R|} < m + \frac{1}{2}$, which completes the proof. □

Lemma 5.18 *If X and Y are loops in a thin relator R which are based at the same vertex, and X followed by Y is a geodesic loop in R, then X and Y are identically oriented. In particular, the winding number of XY is that of X plus that of Y.*

Proof: By Lemma 5.17 the length of the geodesic loop XY is strictly within $2\omega_R$ of the winding number of XY times $|R|$. Since XY is a geodesic, Y does not have winding number 0, or else it could be removed without changing the homotopy class of the loop. Consequently, $|Y| \geq |R|$ and $|X| \leq |XY| - |R|$. Using the fact that R is thin in combination with Lemma 5.17, the winding number of X is seen to be strictly less than that of XY. A similar argument can be used to show the winding number of Y is less than that of XY. Since the winding number of XY is a combination of the winding number of X and the winding number of Y with the appropriate signs attached, the appropriate signs must both be positive in this case. □

Although the content of Lemma 5.18 has great intuitive appeal, the assumption that the relator is thin is necessary for the conclusion to follow, as Example 1 in Section 6 will show. The same counterexample will also show that the conclusion of the following lemma is unwarranted when the word 'thin' is removed.

Lemma 5.19 *If X is a geodesic representative of a thin general relator R, then X is simple.*

Proof: Let X be a geodesic representative which is not simple. In this case X can be written as the sum of two distinct loops, X_1 and X_2, based at the same vertex. By Lemma 5.18, the sum of the winding numbers of X_1 and X_2 equals the winding number of X, which is 1. Thus one of the loops is of winding number 0 and can be removed without changing the homotopy class of the loop. This would imply that the origin loop X is not a geodesic, contradiction, so the assumption that X is not simple must be false. □

Lemma 5.20 *If X is a geodesic representative of a thin general relator R, then there is a deformation retraction from ∂R to X.*

Proof: By Lemma 5.19, the representative X is simple, and since X is simple the functor f from L_n to R is injective and can be viewed as an inclusion map. Because the winding number of X is 1, it induces an isomorphism on fundamental groups. All of the higher homotopy groups are trivial in both spaces which means that f is actually a weak homotopy equivalence, and by Lemma 1.6 there is a deformation retraction from ∂R to X. □

5.5 Geodesic n-gons

If a loop $V = V_1V_2 \dots V_n$ is read in a labeled circular category C and each V_i is a geodesic in the homotopy class relative its endpoints in C, then the loop V is called a geodesic n-gon in C. Let V be a geodesic n-gon in the boundary of a general relator R and let the endpoints of V_i be called v_{i-1} and v_i respectively. If the loop V is a representative of R and for each i, $\mathrm{Ball}(v_i, \omega_R)$ does not contain any of the other endpoints or any of the edges of the geodesic paths other than the ones which contain v_i as an endpoint, then V is called a geodesic n-gon representing R. This condition is precisely what is needed to prove the following lemma.

Lemma 5.21 *If $V = V_1V_2 \dots V_n$ is a geodesic n-gon representing a general relator R, then $|V| < |R| + 2n\omega_R$.*

Proof: Let U be the standard geodesic and suppose that it has a winding number of k. Lift V^k to R^∞ and consider a particular reading of the infinite geodesic U^∞. Let v_0 be the initial vertex of V^k and let v_i be the terminal vertex of the initial segment V^i for all $i = 1, \dots, k$. By the definition of ω_R there is a vertex u_i in U^∞ which passes strictly within ω_R units of each v_i. Moreover, since U has winding number k the vertex u_k can be chosen so that there are paths from v_0 to u_0 and from v_k to u_k which have the same label. In particular, pushing the path from v_0 to u_0 into R and then lifting the result so that it starts at v_k will select the vertex u_k having this property. One consequence of this choice is that the length of the subword of U^∞ marked off by u_0 and u_k is exactly $k|R|$, and thus there is an integer i such that u_i and u_{i+1} mark off a subword of U^∞ of length at most $|R|$. Without loss of generality, assume that $i = 0$ is an integer with this property. Attention now shifts to the word V read from v_0 to v_1.

The terminal vertex of each V_j is within ω_R units of a vertex in U^∞, for $j = 1, \ldots n-1$. Let U_j be the portion of U^∞ marked off by the vertices of U^∞ selected to correspond to the endpoints of V_j in V for all $j = 1, \ldots, n$. Since the definition of a geodesic n-gon representing R implies that the selected vertices in U^∞ occur between u_0 and u_1 in order, the sum of the lengths of the U_j, $j = 1, \ldots, n$ is equal to the length of the path in U^∞ between u_0 and u_1 and thus is at most $|R|$. Since each V_j is a geodesic, $|V_j| < |U_j| + 2\omega_R$ for all $j = 1, \ldots, n$. Combining these inequalities yields the result. □

By the diameter of a general relator R is meant the largest distance between two vertices in R, and it is denoted Diameter(R).

Lemma 5.22 *If R is a thin general relator then Diameter(R) $\leq \frac{1}{2}|R| + 2\omega_R$.*

Proof: Let u and v be vertices in R and let U be a geodesic path between them. If $|U| < 2\omega_R$ then the lemma holds for this pair of vertices, so suppose that $|U| \geq 2\omega_R$. Since U is a geodesic, the lift of the vertex u to R^∞ is strictly between two distinct lifts of the vertex v. The lift of the path U connects the lift of u to one of the lifts of v. Let V be a geodesic path in R^∞ from the other neighboring lift of v to the lift of u. The image of VU in R is now a geodesic 2-gon representing R. By Lemma 5.21 $|R| + 4\omega > |U| + |V| \geq 2|U|$ since U is a geodesic in R. Thus $|U|$ is less than $\frac{1}{2}|R| + 2\omega_R$, completing the proof. □

Suppose that $V = V_1 V_2 \ldots V_n$ is a geodesic n-gon in the boundary of a general relator R, but that it is not a geodesic n-gon representing R because there is an i such that Ball(v_i, ω_R) contains the vertex v_{i-1}. If V' is a geodesic path v_{i-1} to v_{i+1} which is homotopic to $V_i V_{i+1}$ then the path $V_i V_{i+1}$ can be replaced by the path V', in a process called the consolidation of a pair of vertices. The new loop is a geodesic $(n-1)$-gon in the boundary of R, and its length has been shortened by less than $2\omega_R$ units. To see this notice that since V_{i+1} is a geodesic, $|V_{i+1}| \leq |V'| + |V_i|$, and this can be rewritten as $|V_{i+1}| + |V_i| - |V'| \leq 2|V_i| < 2\omega_R$. When these observations are combined with Lemma 5.21, the result is the following.

Lemma 5.23 *If $V = V_1 V_2 \ldots V_n$ is a geodesic n-gon in the boundary of a general relator R, and there exists a sequence of consolidations of pairs of vertices which results in a geodesic k-gon which represents R, then $|V| < |R| + 2n\omega_R$.*

Proof: The comments above show that the $n - k$ consolidations shorten the length less than $2(n-k)\omega_R$, and by Lemma 5.21, the resulting k-gon has length less than $|R| + 2k\omega_R$. □

Lemma 5.24 *Let $V = V_1 V_2 \ldots V_n$ be a geodesic n-gon of winding number 1 in the boundary of a general relator R. If every subword of the form $V_i V_{i+1} \cdots V_j$ in the cycle V which has an orientation is positively oriented with respect to V, then $|V| < |R| + 2n\omega_R$.*

Proof: The properties of V listed in the lemma are inherited under consolidation of pairs of vertices. Moreover, once no more consolidations are possible, it is clear that the result is a geodesic n-gon representing R. The result then follows by Lemma 5.23. □

5.6 Relator Metrics

In order to generalize the small cancellation hypotheses to the context of general relators it is necessary to introduce appropriate functions which can be used to measure the extent to which a particular path wraps around the boundary of a general relator. These functions will be called relator metrics.

Let R be a general relator. A relator metric on R, denoted d_R, is any metric on R^∞, the universal cover of the boundary of R, which satisfies two additional properties: (1) it must be invariant under label-preserving automorphisms of the universal cover, and (2) the distance between the endpoints of a path which forms a loop when pushed into R must be at least the winding number of the loop. A relator metric on R can be used to measure the length of a path in R, by defining its length as the length of a lift of the path to the universal cover. By property (1) this definition is well-defined. By an abuse of notation, d_R will also represent the function which assigns a length to every path in R, so that the length of a path U will be written $d_R(U)$. Alternatively, the function assigning a length to every path in R can be defined directly without using the universal cover. Specifically, a function d_R which assigns a non-negative real number to every path in R is called a relator metric if it satisfies the following six properties:

1. $d_R(U) = d_R(V)$ whenever UV^{-1} is a contractible loop in ∂R
2. $d_R(U) \geq 0$ and $d_R(U) = 0$ iff U is a contractible loop in ∂R
3. $d_R(U) = d_R(U^{-1})$
4. $d_R(UV) \leq d_R(U) + d_R(V)$
5. if U is a path which forms a loop in ∂R then $d_R(U) \geq$ wind.num.(U)
6. if U and V are paths in R which differ by an automorphism of R then $d_R(U) = d_R(V)$.

Property 1 guarantees that the function is only dependent on the homotopy class of the path, which implies that it defines a function on the endpoints of the lift of the path to the universal cover. Properties 2,3, and 4 are the usual axioms for a metric written in terms of paths pushed down to R. Finally, the fifth and sixth properties correspond to the extra conditions placed on relator metrics as described above. A general relator R together with a relator metric on R is called a measured relator. A set of general relators $\mathcal{R}$ together with the set of corresponding relator metrics d_R, one for each $R \in \mathcal{R}$, is called a set of measured relators.

In traditional small cancellation theory, the relators are cyclically reduced so that the loops formed by these cycles are deterministic and have a width of 0. The natural measure of distance in such a setting is the graph metric on R^∞. In particular if U is a path in R then the length of U can be defined as the geodesic distance between the endpoints of U when it is lifted to R^∞. For traditional relators this is equivalent to the length of the reduction of U in the free group. Another metric, which is a variation on the graph metric, is what will be called the normalized graph metric. The length of a path U in the normalized graph metric is given by dividing the length of U in the graph metric by the length of the relator in which it is read. The normalized graph metric of a path U in a relator R will be denoted $|U|_R$. Note that if U is a path in R which lifts to a geodesic in R^∞ then $|U|_R = \frac{|U|}{|R|}$. Also, property 5 is always an equality for the normalized graph metric. Finally, since a length $|R|$ was defined for all general relators R, the normalized graph metric is well-defined for any general relator.

Lemma 5.25 *The sum of the lengths of the arcs of a representative cell of a measured relator is at least 1. That is, if $W = U_1U_2 \ldots U_j$ is a representative of a measured relator R then $\sum_{i=1}^{j} d_R(U_i) \geq 1$. Consequently, if $UXVY$ is a representative of R in which both X and Y measure at most $1 - 3\alpha$, and both U and V measure less than α, then both X and Y measure at least α.*

Proof: By properties 4 and 6, $\sum_{i=1}^{j} d_R(U_i) \geq d_R(W) \geq 1$. For the second statement, the measures of the four pieces add up to at least 1, but since $d_R(X) \leq 1 - 3\alpha$, $d_R(U) < \alpha$, and $d_R(V) < \alpha$, it follows that $d_R(Y) > \alpha$. The proof for X is analogous. □

General relators are more complicated than traditional relators because they typically have a non-zero width. As a consequence, distance can be measured in many different ways, using possibly distinct metrics. This gives rise to a variety of ways in which a word or a cycle may be 'reduced'. The different types of reductions are described below as a prelude to the general theory.

Let $\mathcal{R}$ be a set of measured relators, and let μ be a real number with $\frac{1}{2} \leq \mu < 1$. A word W is called μ-reduced if it is reduced in the free group and there do not exist words U and V and a general relator $R \in \mathcal{R}$ such that U is a subword of W, UV is a representative of R, and $|U| > \mu|UV|$. Alternatively, the inequality can be written as $|U| > \frac{\mu}{1-\mu}|V|$. In this form it is clear that over a set of standard relators, a word is $\frac{1}{2}$-reduced iff it is Dehn-reduced in the usual sense, that is, iff it does not contain strictly more than one-half of a representative of a relator. Over a set of general relators, the notion of being one-half of a relator loses its sense, but being μ-reduced is still well-defined. As a result, even in the general case, a word will be called Dehn-reduced iff it is $\frac{1}{2}$-reduced with respect to $\mathcal{R}$. A cycle W will be called μ-reduced iff all of the cyclic conjugates of the word W are μ-reduced. Cycles which are $\frac{1}{2}$-reduced are called Dehn-reduced cycles.

The other types of 'reduced' words and cycles will be defined using relator metrics. Let $\mathcal{R}$ be a set of measured relators, and let μ be a non-negative real number. A word W is called μ-free if it is reduced in the free group and does not contain strictly more than μ of a relator in $\mathcal{R}$ as measured by the relator metric. Specifically, there cannot exist a word U and a general relator R such that U is a subword of W which is readable in R with $d_R(U) > \mu$. The cycle of W is called μ-free iff all of the cyclic conjugates of the word W are μ-free. Next, let $\mathcal{R}$ be a set of measured relators, and let μ be a real number strictly between 0 and 1. A word W is called μ-complement-free if it is reduced in the free group and there does not exist a subword of W with a complement in a particular relator measuring less than μ. Specifically, there must not exist words U and V and a general relator R such that U is a subword of W, UV is a representative of R, and $d_R(V) < \mu$. A cycle is called μ-complement-free iff all of its cyclic conjugates are μ-complement-free. Notice that by the properties of relator metrics, $d_R(V) < \mu$ means that $d_R(U) > 1 - \mu$. Thus a word or cycle W which is not μ-complement-free is not $(1 - \mu)$-free. Or, said differently, if W is $(1 - \mu)$-free then it is also μ-complement-free.

Both the normalized graph metric and relator metrics will be retained since graph metrics are necessary to complete the inductive step in Dehn's algorithm, but relator metrics are more convenient for various constructions.

5.7 Automorphism Groups

Under certain conditions, the automorphism groups of general relators are easy to describe. One such condition involves what are called crucial cones. Let R be a general relator, and consider an open cone in ∂R which is not contained in the boundary of any other open cone in ∂R. Let B be this open cone together with all of the open cones in its boundary which are not contained in any open cones in R except those which are themselves contained in this particular open cone. Topologically speaking, B is the complement of the closure of the complement of a maximal closed cone in R. In the case of a traditional relator, this is a description of the open edges in the boundary of the relator. Notice that this collection of open cones B contains an open cone of maximum rank and that all of B is contained in the closure of this open cone. Moreover, the closure of B is itself either an edge or a general relator S contained in ∂R. Let B_0 be a lift of this B to R^∞, and let S_0 be the lift of S which contains B_0. If the removal of B_0 disconnects the ends of R^∞ then define B^+ to be the intersection of S_0 with the positive end $R^{+\infty}$. Similarly, let B^- be the intersection of S_0 with the negative end $R^{-\infty}$. If the removal of B_0 not only disconntects the ends of R^∞ but it is also true that B^+ and B^- do not contain any loops which have nontrivial winding number in ∂S, then B_0 is called a crucial cone in R^∞ By extension, B will be called a crucial cone in ∂R. In a traditional relator, all of the open edges of the boundary of the 2-cell are crucial. The terminology derives from the fact that B is crucial for the completion of loops in ∂R, in the sense that every representative of R must contain at least one vertex or edge in B.

A slight generalization of this goes as follows: let C be a subcomplex of ∂R such that $g(C) = C$ for all $g \in \mathrm{Aut}(R)$, and let $C^\infty = f^{-1}(C)$ where

$f : R^\infty \to \partial R$ is the usual covering map. So long as C is connected and C contains a representative of R, it can be proved that C^∞ has at most two ends, denoted $C^{+\infty}$ and $C^{-\infty}$, using the same proof as before. Such a subcomplex C will be called a core of the general relator R. As above, let B be the complement of the closure of the complement of a maximal open cone in C, let S be the closure of B, let B_0 be a lift of B to C^∞, let S_0 be the lift of S to C^∞ which contains B_0, and let B^+ and B^- be the intersections of S_0 with $C^{+\infty}$ and $C^{-\infty}$ respectively. If B_0 disconnects the ends of C^∞, and B^+ and B^- do not contain any loops with nontrivial winding number when viewed as paths in S_0, then B_0 will be called a crucial cone in C^∞ and B will be called a crucial cone in C. In the following, R will be said to have a crucial cone in its boundary if there exists some core C and some B for which B is a crucial cone of C.

Now suppose that R is a general relator which has a crucial cone and fix a core C which contains one. Since the crucial cones in C^∞ disconnect the ends of C^∞, and since by construction they are pairwise disjoint, by Lemma 5.12 they are totally ordered. The proofs of the lemmas stated earlier in this section are unchanged if C and C^∞ are used instead of ∂R and R^∞. One way to see this is to notice that since C is a subcomplex, the cone over C is a general relator in its own right. Next, since all of the lifts of a crucial cone in C are crucial cones in C^∞, there is no first or last crucial cone in C^∞ and the total ordering of the crucial cones in C^∞ is order isomorphic to $(\mathbf{Z}, <)$. If this ordering is imposed on the crucial cones in C, there is a well-defined notion of a next and a previous crucial cone along the boundary ∂R in a particular direction, although it is no longer, strictly speaking, an ordering. If there is only one crucial cone in ∂R, then this next crucial cone in ∂R may turn out to the original crucial cone itself.

It should also be noted that the property of being a crucial cone is preserved under automorphisms of the relator R, and it is also preserved under automorphisms of R^∞. By one of the defining properties of C, an automorphism of R is also an automorphism of C. In addition, the total ordering of the crucial cones in C^∞ is either preserved or reversed under an automorphism depending on whether the automorphism preserves or reverses the ends of C^∞. This is immediate from the definition of the ordering. Thus, if R contains a crucial cone in its boundary there is a group homomorphism from $\mathrm{Aut}(R^\infty)$ to $\mathrm{Aut}(\mathbf{Z}, <)$ which is the infinite dihedral group $\mathbf{D}_\infty$. The homomorphism is given by observing the action of the automorphism on the linearly ordered crucial cones in C^∞. Similarly, if C contains exactly r crucial cones, then there is a group homomorphism from $\mathrm{Aut}(R)$ to the dihedral group of order $2r$, which is denoted $\mathbf{D}_{2r}$.

The kernels of these homomorphisms contain those automorphisms of R^∞ and ∂R which fix all of the crucial cones in C^∞ and C respectively. Notice that the morphism from $\mathrm{Aut}(R^\infty)$ to $\mathbf{D}_\infty$ also shows that the only elements in $\mathrm{Aut}(R^\infty)$ of finite order are either orientation-reversing or else they are automorphisms in the kernel of this map. The following lemmas will show that the kernels of these homomorphisms are 2-groups. Lemma 5.26 and Lemma 5.27 will be proved by simultaneous induction on the rank n of the general relator R under consideration. As a final note, a set of general relators is called closed under subcones if whenever S is a general relator contained in $R \in \mathcal{R}$, S is also in $\mathcal{R}$.

Lemma 5.26 *Let $\mathcal{R}$ be a set of general relators closed under subcones and suppose that all general relators in $\mathcal{R}$ have at least one crucial cone in their boundary. If H is a subgroup of the automorphism group of a general relator $R \in \mathcal{R}$, then for some r there is a group homomorphism from H to $\mathbf{D}_{2r}$, whose kernel is a 2-group. In particular, H is isomorphic to a cyclic or a dihedral group extended by a 2-group.*

Proof: Let C be a core of R and let r be the number of crucial cones in C. Since the presence of crucial cones in C already implies the existence of a group homomorphism $f : \mathrm{Aut}(R) \to \mathbf{D}_{2r}$, it only remains to show that the kernel of this map is a 2-group. Let B be a crucial cone in C, and let B_0 be a lift of B to C^∞. If g is an automorphism of R which fixes all crucial cones in C, then there is an automorphism g' of R^∞ which fixes B_0 and is equal to g when composed with the covering map from C^∞ to C.

Let S be the general relator which results from the closure of B_0 Define B^+ and B^- be the intersection of S with $C^{+\infty}$ and $C^{-\infty}$ respectively. Since B_0 is a crucial cone in C^∞, B_0 disconnects the two infinite ends of C^∞, and B^+ and B^- are thus disjoint. Moreover, since the action of g' either preserves the ends or switches them, it follows that either $g'(B^+) = B^+$ and $g'(B^-) = B^-$, or else $g'(B^+) = B^-$ and $g'(B^-) = B^+$. By definition, B^+ and B^- do not contain the lift loop in S with a nontrivial winding number. Thus the general relator S with subcomplexes B^+ and B^- satisfies all of the hypotheses of Lemma 5.27. Moreover, since the rank of S is strictly less than of R, Lemma 5.27 can be applied to show that g' is an automorphism whose order is a power of 2. This in turn implies that the order of g itself is a power of 2, and that the kernel of the homomorphism f is a 2-group. □

Lemma 5.27 *Let $\mathcal{R}$ be a set of general relators closed under subcones and suppose that all general relators in $\mathcal{R}$ have at least one crucial cone in their boundary. If R is a general relator in $\mathcal{R}$, B_1 and B_2 are disjoint subcomplexes of ∂R which do not contains any loops with a nontrivial winding number, and H is the subgroup of $Aut(R)$ of automorphisms which either sends B_1 to B_2 and vice versa or else which fixes B_1 and B_2 separately, then H is a 2-group.*

Proof: Let C be a core of R and let r be the number of crucial cones in C. The presence of crucial cones in C already implies the existence of a group homomorphism $f : \mathrm{Aut}(R) \to \mathbf{D}_{2r}$. Let g be an automorphism in H, and consider g^2. By the definition of H, g^2 must fix both B_1 and B_2, and in addition, it is orientation-preserving. Since B_1 does not contain any loop with a nontrivial winding number, by Lemma 1.13 it is possible to lift B_1 to R^∞. Call this lift B_0. Corresponding to the automorphism g^2 of C, there is an automorphism of R^∞ (and C^∞) which fixes B_0. Since this automorphism of C^∞ is also orientation-preserving, but of finite order, it must fix all of the crucial cones in C^∞. As a consequence, g^2 must fix all of the crucial cones in C, and thus g^2 is a member of the kernel of the map from $\mathrm{Aut}(R)$ to $\mathbf{D}_{2r}$. By Lemma 5.26, the order of g^2 must be a power of 2. This shows that the order of g itself is a power of 2, and since g was chosen at random, the group H is a 2-group. □

Part III
Constructions

In Part III general presentations and various constructions over general presentations are investigated. In Section 6 $\mathcal{R}$-categories are defined and their properties are studied. Included in this section is a definition of a general version of the traditional Poincaré construction, which is used to define the group of a general presentation. Section 7 extends the combinatorial group theory notion of a planar map in the context of general relators. More flexible structures which correspond to maps on surfaces such as spheres are also described. Finally, in Section 8 the Cayley category of a general presentation is defined and its properties are discussed.

6 $\mathcal{R}$-Categories

This section contains a detailed study of $\mathcal{R}$-categories. After general presentations and $\mathcal{R}$-categories are defined, the notion of the collapse of an $\mathcal{R}$-category is introduced. This leads to a generalization of the usual Poincaré construction. The section concludes with a detailed description of the relationship between general presentations and traditional presentations.

6.1 General Presentations

Let $\mathcal{R}$ be a set of relators. A presentation is a set of generators A together with a set $\mathcal{R}$ of relators labeled by A. The group G defined by the presentation is the quotient of the free group on the alphabet A by the normal closure of the relators in $\mathcal{R}$ thought of as words in A^*. Presentations are written $G = \langle A|\mathcal{R}\rangle$. The Cayley graph of a presentation $G = \langle A|\mathcal{R}\rangle$ is $\mathcal{C}(G, A)$. The Cayley graph of a presentation is of interest both because of the information it carries and because it is constructible if and only if the word problem for the presentation is decidable.

The familiar Poincaré construction of a presentation is a 2-dimensional cell category whose fundamental group is the group described by the presentation. Specifically, it is a 2-dimensional cell category with 1 object of dimension 0, objects of dimension 1 corresponding to the orbits of A, and objects of dimension 2 corresponding to the relators in $\mathcal{R}$. One way of producing the Poincaré construction of a presentation is to take the generators A, viewed as labeled edges, and the relators in $\mathcal{R}$, viewed as labeled 2-cells, and first identify all of the vertices and then collapse the resulting complex until a collapsed 2-dimensional cell category labeled by A emerges. The operation of collapsing a category is precisely defined below.

Example 1 The presentation $G = \langle u, v|u^2 = v^3, uv = vu\rangle$ is a presentation of the integers. The Cayley graph of G is shown in Figure 15. The edges labeled u

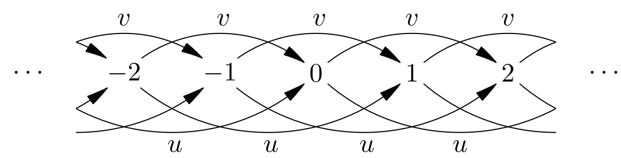

Figure 15: A presentation of **Z**

and v are oriented so that they start at the lesser number and end at the greater number. The Poincaré construction of the presentation is a 2-dimensional cell category homotopically equivalent to S^1. Let X be the Poincaré construction of the presentation and let $R = \text{Cone}(X)$. Even though the cone category R is not a general relator since X is not a cone complex, X can, if desired, be altered to produce a general relator which provides the promised counterexamples, by first subdividing the 1-skeleton into thirds, adding a new vertex in the center of each 2-cell, connecting each new vertex to each of the old vertices in its boundary, and then subdividing each of the recently added edges into a large number of pieces, say one hundred pieces each. With such an alteration it is clear that any geodesic loop in the new construction will be contained in the old 1-skeleton. For convenience the original construction X will be used in the description below.

In R the shortest path of winding number 1 is uv^{-1} which is of length 2 and non-simple. The edge u is itself a loop which has winding number 3 and the edge v is a loop with winding number 2. The length of R is $\frac{1}{3}$, the edge u is the standard geodesic, and its width, defined below, is 2 since the removal of a vertex and all of the vertices located one edge away does not disconnect the universal cover of the boundary of the relator (see Figure 15). The construction R, however, is not thin, and it does not satisify the conclusions of Lemma 5.18 and Lemma 5.19.

A general presentation, denoted $\langle A|\mathcal{R}\rangle$, is defined as an alphabet A of invertible generators and a set $\mathcal{R}$ of general relators labeled by A which is closed under subcones. That is, if R is in $\mathcal{R}$, and S is a general relator contained in R, then S is in $\mathcal{R}$. The group associated with a general presentation will not be defined until after the introduction of the general Poincaré construction later in this section. A rank function on a general presentation is a function from $A \cup \mathcal{R}$ to $\mathbf{N}$ such that the rank of each generator in A is 1, and all of the general relators $R \in \mathcal{R}$ are assigned a rank of at least 2 under the restriction that whenever there is a label-preserving circular functor $f : R \to S$ between two general relators $R, S \in \mathcal{R}$, then $\text{rank}(R) \leq \text{rank}(S)$ with equality iff the f is an isomorphism. Every general presentation can be assumed to be equipped with such a function, since the height function proves that such a function always exists.

The set of all general relators of rank k is denoted $\mathcal{R}_k$, and the set of all general relators of rank at most k is written $\mathcal{R}(k)$. Clearly, $\mathcal{R}$ is partitioned into a disjoint union of the sets $\mathcal{R}_1$, $\mathcal{R}_2$, etc., and $\mathcal{R}(k) = \mathcal{R}(k-1) \cup \mathcal{R}_k = \cup_{i=1}^{k} \mathcal{R}_i$.

The condition on the rank function guarantees that if R is a general relator in $\mathcal{R}_k$, then all of the relators in its boundary are contained in $\mathcal{R}(k-1)$. Notice that $\mathcal{R}_1 = \mathcal{R}(1) = \emptyset$. A general presentation $G = \langle A|\mathcal{R}\rangle$ where $\mathcal{R}$ is a set of measured relators is called a measured presentation.

Let $\langle A|\mathcal{R}\rangle$ be a general presentation. A labeled circular category in which every slice category is either a vertex, a generator from A or a general relator in $\mathcal{R}$ is called a circular category over $\langle A|\mathcal{R}\rangle$ or, more simply, an $\mathcal{R}$-category with the set of labels implicitly understood. If the underlying circular category is also a circular complex, then the result could more accurately be labeled an $\mathcal{R}$-complex. Notice that the rank function of the general presentation induces a rank function on every $\mathcal{R}$-category which assigns vertices a rank of 0, labeled edges a rank of 1, and all other objects c a rank based on the rank of the general relator isomorphic to its slice category. A label-preserving circular functor between $\mathcal{R}$-categories is called an $\mathcal{R}$-functor. Notice that such a functor automatically preserves the induced ranks on its objects. An $\mathcal{R}$-functor $f : C \to B$ is called onto if every object $b \in B$ is the image of some object $c \in C$. An automorphism of an $\mathcal{R}$-category C is an $\mathcal{R}$-functor from C to itself which is a isomorphism. The automorphisms of an $\mathcal{R}$-category C form a group called the automorphism group of C.

6.2 Collapse of an $\mathcal{R}$-Category

An $\mathcal{R}$-category C is called collapsed if whenever there exist $\mathcal{R}$-functors $f, g : B \to C$ where B is a connected $\mathcal{R}$-category and $f(v) = g(v)$ for a vertex v in B, then f and g must be equal. Equivalently C is called collapsed if its 1-skeleton is deterministic and distinct circular cones of dimension at least 2 have distinct boundaries. To say that two open cones c_1 and c_2 have the same boundaries means that C/c_1 and C/c_2 are isomorphic as labeled circular complexes and that the attaching functors agree on corresponding parts of the boundaries. Specifically, $h : C/c_1 \to C/c_2$ is an isomorphism which shows that c_1 and c_2 have the same boundaries, if the attaching functor ϕ_{c_1} is equal to the composition of h with the attaching functor ϕ_{c_2}. Clearly, the first definition implies the second, while Lemma 6.3 shows that the two definitions are equivalent.

A geometric description of an identification of the closed cone C/c_1 with the closed cone C/c_2 goes as follows. Because h is an isomorphism between categories, its geometric realization is a homeomorphism between the polyhedra C/c_1 and C/c_2. Since attaching maps such as ϕ_{c_1} and ϕ_{c_2} are also homeomorphisms once restricted to their open cones c_1 and c_2, there is a homeomorphism $\phi_{c_2} h \phi_{c_1}^{-1}$ between points in the open cone c_1 and the open cone c_2 in the geometric circular category C. The collapse of c_1 and c_2 is the quotient of C obtained by identifying the points in the open cones c_1 and c_2 under the above homeomorphism. The quotient map from C to the quotient is a continuous map.

A categorical description of the same operation is to identify the objects c_1 and c_2 in C, and to identify an arrow $f : d \to c_1$ with an arrow $g : d \to c_2$ iff h sends f to g. Notice that the letters f and g on the righthand side represent the objects in the slice categories which correspond to the arrows f and g in C. This relation determines a unique quotient category in the sense of MacLane [11], and a unique functor from the original category onto the quotient category. These correspond to the quotient space and the quotient map described above.

The collapse of a finite $\mathcal{R}$-category C can be obtained by identifying the endpoints of oriented edges which have the same initial vertex and the same label, identifying open circular cones which have the same boundaries and then repeating. Since the category is finite, and the number of objects is decreasing, the process must stop. And since the collapsing operations are confluent the result, called the collapse of C, is well-defined. For more general $\mathcal{R}$-categories, the usual limit construction suffices: define an equivalence relation on the objects and arrows of an $\mathcal{R}$-category C where two objects (arrows) are equivalent iff there is a finite sequence of collapses after which they are equal. It only remains to check that the result of quotienting by this relation is a well-defined, collapsed $\mathcal{R}$-category $\overline{C}$ and that the obvious functor $f : C \to \overline{C}$ is a label-preserving circular functor. The category $\overline{C}$ is called the collapse of C and the functor f is called the collapsing functor. The collapsing functor has the additional property that every other functor from C to a collapsed category B factors through f. This proves the following lemma.

Lemma 6.1 *If $\mathcal{R}$ is a set of general relators, and C is an $\mathcal{R}$-category, then there exist a collapsed $\mathcal{R}$-category $\overline{C}$ and an onto $\mathcal{R}$-functor $f : C \to \overline{C}$ which are universal in the sense that, given any collapsed $\mathcal{R}$-category B and $\mathcal{R}$-functor $g : C \to B$, there is a unique functor $h : \overline{C} \to B$ such that $hf = g$.*

The reversibility of the order in which collapses take place makes it possible to continue to collapse the edges of C until a deterministic 1-skeleton emerges, and then to collapse all 2-cones until a collapsed 2-skeleton emerges, etc. Because the collapsing of higher-dimensional cones does not affect lower-dimensional skeleta, the deterministic 1-skeleton which emerges is the 1-skeleton of the final completely collapsed version, similarly for the collapsed 2-skeleton, the collasped 3-skeleton, etc. This result is summarized in the following lemma.

Lemma 6.2 *Let $\mathcal{R}$ be a set of general relators, and let C be an $\mathcal{R}$-category. For every positive integer k there is an $\mathcal{R}$-category C_k with a collapsed k-skeleton, and an $\mathcal{R}$-functor $f : C \to C_k$ which is universal in the sense that, given any $\mathcal{R}$-category B with a collapsed k-skeleton and an $\mathcal{R}$-functor $g : C \to B$, there is a unique functor $h : C_k \to B$ such that f followed by h is equal to g. Moreover, in the particular case where B is the collapse of C, the functor h is an isomorphism when restricted to the k-skeletons of each category.*

If $\langle A|\mathcal{R}\rangle$ is a general presentation and every general relator $R \in \mathcal{R}$ is a collapsed $\mathcal{R}$-category in its own right, then $\mathcal{R}$ is called a collapsed set of general relators. The fact that an $\mathcal{R}$-category is collapsed restricts the number of possible $\mathcal{R}$-functors available. This is true when the collapsed category is either the domain or the range of the functor.

Lemma 6.3 *Let C be a connected $\mathcal{R}$-category, and let B be an $\mathcal{R}$-category which is collapsed in the sense that it has a deterministic 1-skeleton, and distinct cones of rank at least 2 have distinct boundaries. If $f, g : C \to B$ are $\mathcal{R}$-functors with $f(v) = g(v)$ for some vertex v of C, then f and g are identical. Thus, the two definitions of collapsed categories are equivalent.*

Proof: Since C is connected, the 1-skeleton of B is deterministic, and the functors f and g agree at a point, f and g must agree on the entire 1-skeleton of C. Next assume that f and g agree on the k-skeleton of C. Since the images under f and g of a $(k+1)$-cone agree on the boundary of the cone and C is collapsed the images must be identical. Thus f and g agree on the $(k+1)$-skeleton of C, and by induction they agree completely. □

Lemma 6.4 *If $f : B \to C$ is an $\mathcal{R}$-functor between $\mathcal{R}$-categories, and f is injective on objects, then f is actually an embedding of B as a subcategory of C.*

Proof: If f is not an embedding, then it is not injective on arrows. Since f is injective on objects, however, these arrows must originate in the same hom-set in B. In particular, they have the same terminal object, say b. Next, the arrows terminating at b in B are in 1 to 1 correspondence with the arrows terminating at b in B/b, which are in 1 to 1 correspondence with the arrows terminating at $f(b)$ in $C/f(b)$ since f is an $\mathcal{R}$-functor, and these are in 1 to 1 correspondence with the arrows terminating at $f(b)$ in C, contradiction. Thus f is an embedding. □

Lemma 6.5 *If $f : B \to C$ is an $\mathcal{R}$-functor between $\mathcal{R}$-categories, B is collapsed, and f is injective on vertices, then f is an embedding of B as a subcategory of C.*

Proof: If f is not an embedding, then by Lemma 6.4 it is not injective on objects. It is, by assumption, injective on vertices, so the identification must be between objects with rank at least 1. Notice that the identified objects must have the same rank since rank is preserved by $\mathcal{R}$-functors. If two distinct objects b_1 and b_2 in B are chosen so that $f(b_1) = f(b_2)$, but f is injective when restricted to objects of lesser rank, then the isomorphisms linking B/b_1 and B/b_2 with $C/f(b_1) = C/f(b_2)$ can be combined to show that B/b_1 and B/b_2 are isomorphic. Because of the way in which b_1 and b_2 were chosen, b_1 and b_2 must also have identical boundaries. Since B is collapsed, this is impossible, and the assumption that f is not an embedding must be false. □

6.3 Poincaré Constructions

A collapsed $\mathcal{R}$-category with a single vertex is known as a Poincaré construction. The fact that Poincaré constructions are collapsed places severe restrictions on the number of possible $\mathcal{R}$-functors as is shown above. The restrictions are in fact enough to show that Poincaré constructions are in 1 to 1 correspondence with general presentations.

Lemma 6.6 *Let B and C be collapsed, connected $\mathcal{R}$-categories which have automorphism groups which act transitively on their vertex sets. If, in addition, B is contained in C then all of the automorphisms of B extend uniquely to automorphisms of C.*

Proof: Let u and v be vertices in B. By assumption there are automorphisms f and g of B and C, respectively, which send u to v, and by Lemma 6.3 they are unique. If the range of f is extended to C via the embedding of B, and the domain of g is restricted to B, then both are functors from B to C which send u to v. Lemma 6.3 now shows that they are identical. □

Lemma 6.7 *If B and C are $\mathcal{R}$-categories, $f, g : C \to B$ are $\mathcal{R}$-functors, and B is a Poincaré construction, then f and g are identical.*

Proof: Since by definition B has only one vertex Lemma 6.3 can be applied to each of the connected components of C to show that the functors f and g are identical. □

Lemma 6.8 *If B and C are $\mathcal{R}$-categories, B is a Poincaré construction, and the generators and general relators used to construct C are all used in the construction of B, then there is a unique $\mathcal{R}$-functor from C to B. Consequently, distinct objects in a Poincaré construction have non-isomorphic slice categories.*

Proof: Once it is shown that such an $\mathcal{R}$-functor exists, Lemma 6.7 guarantees that it is unique. The existence will be shown by induction on the rank of C. If C has rank 0 then the functor sending the objects of C to the unique vertex of B is clearly an $\mathcal{R}$-functor. Next suppose that the lemma is true for all $\mathcal{R}$-categories of rank less than k, and let C have rank exactly k. By the induction hypothesis there is a unique functor f from the $(k-1)$-skeleton of C into B. Since the objects of rank k in C are terminal objects, it is enough to extend the functor to each of these objects individually. If c is an object of rank k in c then by assumption there must be an object b in B such that C/c is isomorphic to B/b as $\mathcal{R}$-categories. Moreover, the boundary of C/c mapped into C and then into B must agree with the image of the corresponding boundary of B/b in B by induction. Thus the extension of f so that c is sent to b preserves the fact that f is an $\mathcal{R}$-functor. Doing this for all objects of rank k shows that the lemma is true for $\mathcal{R}$-categories of rank k. Thus by induction the lemma is true for all $\mathcal{R}$-categories of finite rank, and since $\mathcal{R}$-categories of infinite rank are constructed skeleton by skeleton, the lemma is also true in general. The final statement of the lemma follows from the first since distinct objects with isomorphic slice categories would have characteristic functors which would contradict the first statement. □

Lemma 6.9 *A Poincaré construction is uniquely determined by the generators and general relators used in its construction. In particular there is a 1 to 1 correspondence between Poincaré constructions and general presentations.*

Proof: Let B_1 and B_2 be two Poincaré constructions for which the exact same set of generators and general relators occur as slice categories. By Lemma 6.8 there is a unique circular functor $f : B_1 \to B_2$ and again there is a unique circular functor $g : B_2 \to B_1$. By Lemma 6.8 again, the $\mathcal{R}$-functor gf must be equal to the identity map on B_1 and the $\mathcal{R}$-functor fg must be equal to the identity map on B_2. Thus the functors f and g are $\mathcal{R}$-category isomorphisms and B_1 and B_2 are isomorphic as $\mathcal{R}$-categories.

To show the second part simply notice that by starting with a general presentation $\langle A|\mathcal{R}\rangle$ it is possible to take edges labeled by the orbits of A, and all of the general relators in $\mathcal{R}$, and then identify all of the vertices involved, and finally collapse the result to produce an $\mathcal{R}$-category which contains precisely the generators and general relators of the general presentation. By the first part of the lemma, this Poincaré construction is uniquely determined, and the general presentation can clearly be recovered from the construction. □

In light of Lemma 6.9, define the content of a labeled circular category C as the set of generators and general relators which occur in C as slice categories. Lemma 6.9 can then be restated as: Poincaré constructions are uniquely determined by their content. The unique Poincaré construction corresponding to a general presentation $\langle A|\mathcal{R}\rangle$ will be called the Poincaré construction of the presentation and it will sometimes be denoted $\mathcal{P}(A, \mathcal{R})$. The fundamental group of the Poincaré construction as a topological space provides a way of assigning a group to every general presentation. If the fundamental group is G then G is called the group of the presentation and for simplicity this is indicated as $G = \langle A|\mathcal{R}\rangle$.

Notice that a set of generators and general relators which is not closed under subcones still creates a well-defined Poincaré construction which could be used to define the group of such a presentation. If such general presentations are allowed, however, additional care must be taken since not all slice categories of rank at least 2 would be general relators, and the statement of Lemma 6.9 would have to be altered. To avoid these situations, general presentations will always be assumed to be closed under subcones, or else it is at least tacitly understood that a general presentation refers to a set of general relators which have been 'closed' in this way.

6.4 Covers and Retractions

The earlier results on coverings and deformation retractions can be extended to $\mathcal{R}$-categories.

Lemma 6.10 *If $f : C \to B$ is a covering map and B is an $\mathcal{R}$-category, then C can also be given an $\mathcal{R}$-category structure under which the map f becomes an $\mathcal{R}$-functor. As a consequence the deck transformations of f are also $\mathcal{R}$-functors.*

Proof: By Lemma 3.14 C is the geometric realization of a cone category, and f is the topological version of a cone functor between cone categories. Since a circular category is a cone category with restrictions on the types of cones, and since the slices of C are isomorphic to slices of B by the definition of a cone functor, C is also a circular category. Finally, the induced labeling on C turns C into an $\mathcal{R}$-category. The last statement follows since every deck transformation is a trivial covering of C. □

Lemma 6.11 *Let $\mathcal{R}$ be a set of general relators, and let B be an $\mathcal{R}$-category whose fundamental group is G. For every normal subgroup H of G there is an $\mathcal{R}$-category C and an $\mathcal{R}$-functor $f : C \to B$ such that f is a regular covering, the fundamental group of C is H, and the group of deck transformations of f are automorphisms of C which form a group isomorphic to G/H.*

Proof: From Lemma 1.6 it is clear that a regular cover exists which satisfies the conditions on the fundamental group and the group of deck transformations, and by Lemma 6.10, these can be given the structure of an $\mathcal{R}$-category and $\mathcal{R}$-functors. □

Lemma 6.12 *Let B be the Poincaré construction of a general presentation $G = \langle A|\mathcal{R}\rangle$. The regular covers of B correspond to the normal subgroups of its fundamental group G, and the 1-skeleton of the regular cover associated with the normal subgroup H is the Cayley graph $\mathcal{C}(G/H, A)$. In particular, if C is the universal cover of B, then a path in C is a loop, and thus homotopic to a point, iff the word read by the path is equivalent to 1 in the group G.*

Proof: The first assertion is a special case of Lemma 6.11. Since by the definition of a regular cover the 1-skeleton is a connected A-graph on which the label-preserving automorphisms act transitively on the vertices, it must be the Cayley graph of its automorphism group, by Lemma 4.4. Next, a path in C is also a path in its 1-skeleton which is the Cayley graph $\mathcal{C}(G, A)$ by Lemma 6.12. By Lemma 4.5, this path is a loop iff the word read by the path is equivalent to 1 in the group G. Finally, since the universal cover is simply connected by definition, such a loop is homotopic to a point. □

Lemma 6.13 *Let $G = \langle A|\mathcal{R}\rangle$ be a fixed general presentation. There is a 1 to 1 correspondence between the normal subgroups H of G, the Cayley graphs $\mathcal{C}(G/H, A)$, the regular covers of the Poincaré construction $\mathcal{P}(A, \mathcal{R})$, and the collapses of these regular covers.*

Proof: The correspondence between the normal subgroups of G, their Cayley graphs, and the regular covers of the Poincaré construction of the presentation has already been alluded to in Lemma 6.11 and Lemma 6.12. By Lemma 6.1 the regular cover of the Poincaré construction corresponding to the normal subgroup H has a unique collapse. By Lemma 6.2 the 1-skeleton of the collapsed category is identical to that of the regular cover and by Lemma 6.12 this 1-skeleton is the Cayley graph $\mathcal{C}(G/H, A)$. Since the regular cover can be reconstructed from this 1-skeleton, the correspondence between regular covers and their collapses must be 1 to 1. □

A general presentation over a set of thin general relators is closely connected with certain specific standard presentations of the same group. The connection between the two types of presentation is elucidated through the use of deformation retractions. Let $\mathcal{R}$ be a set of general relators, and let $\mathcal{R}'$ be a set of cycles in 1 to 1 correspondence with $\mathcal{R}$ such that given any general relator $R \in \mathcal{R}$, the corresponding cycle in $\mathcal{R}'$ can be read as a simple representative of R. The set $\mathcal{R}'$ is called a set of standard representatives for $\mathcal{R}$. Recall that if the general relators in $\mathcal{R}$ are assumed to be thin, then by Lemma 5.19 there always exists at least one set of standard representatives for $\mathcal{R}$.

Lemma 6.14 *Let $\mathcal{R}$ be a set of general relators and let $\mathcal{R}'$ be a set of standard representatives for $\mathcal{R}$. If C is an $\mathcal{R}$-category then there is a deformation retraction of C onto a 2-dimensional subcomplex of the simplicial subdivision of C which can be viewed as a simplicial subdivision of an $\mathcal{R}'$-category. In addition, the deformation can be chosen so that the 1-skeleton remains fixed throughout.*

Proof: The proof proceeds by induction on the rank of the $\mathcal{R}$-category C. First of all, notice that the lemma is trivially true if C is an $\mathcal{R}$-category of rank 1, since in this case C is simply an A-graph, and thus trivially an $\mathcal{R}'$-category. The deformation retraction mentioned in the lemma is the identity retraction which keeps the 1-skeleton fixed throughout. Suppose that the lemma is true for all sets of general relators $\mathcal{R}$ and $\mathcal{R}$-categories of rank at most k. Let $\mathcal{R}(k)$ be the subset of $\mathcal{R}$ consisting of the general relators in $\mathcal{R}$ of rank at most k. It should be clear that an $\mathcal{R}$-category of rank k is also an $\mathcal{R}(k)$-category. Let $\mathcal{R}'(k)$ be the subset of $\mathcal{R}'$ corresponding to the general relators in $\mathcal{R}(k)$.

Consider a general relator R in $\mathcal{R}$ of rank $k+1$ with a corresponding simple representative U in $\mathcal{R}'$. Since the boundary of R, ∂R, is an $\mathcal{R}(k)$-category, by assumption there is a deformation retraction of ∂R onto a 2-dimensional subcomplex, say S, of the simplicial subdivision of ∂R which can be viewed as a simplicial subdivision of an $\mathcal{R}'(k)$-category. By Lemma 1.11 the deformation retraction of ∂R onto S can be extended to a deformation retraction of R onto $\text{Cone}(S)$. Since the 1-skeleton is fixed throughout, the loop U is a simple loop S as well as in ∂R, and the injection of U into S is a weak homotopy equivalence, and thus by Lemma 1.6 there is a deformation retraction from S to U. By Lemma 1.12 there is a deformation retraction from $\text{Cone}(S)$ onto $S \cup_U \text{Cone}(U)$. Combining these two deformations, it has been shown that for all relators R of rank $k+1$ there is a deformation retraction of R onto $S \cup_U \text{Cone}(U)$ which extends the prior retraction of ∂R onto S. Since the 1-skeleton remains fixed throughout, U is in the 1-skeleton, and $\text{Cone}(U)$ is the 2-cell of a relator in $\mathcal{R}'$, the deformation from R to $S \cup_U \text{Cone}(U)$ satisfies the conditions of the lemma.

Next consider an arbitrary cone category C of rank $k+1$. By assumption there is a deformation retraction of the k-skeleton which satisfies the lemma. Since this retraction of the k-skeleton of C induces a deformation retraction of the boundary of each general relator R of rank $k+1$ which is attached to the k-skeleton via the attaching map ϕ_R. By the argument above, there is a deformation retraction of R onto an $\mathcal{R}'$-category which extends the deformation retraction of its boundary. The image of this deformation retraction of R is consistent with the earlier retraction of the k-skeleton, so that the deformation retractions on each of the relators of rank $k+1$ combined with the deformation on the k-skeleton yield a single deformation on C which satisfies the lemma.

Finally, a deformation retraction on an arbitrary cone category C which satisfies the requirements of the lemma can be constructed inductively via the deformation retractions on the various skeleta. □

Lemma 6.15 *If $\mathcal{R}$ is a set of thin general relators and C is an $\mathcal{R}$-category, then $H_i(C) = 0$ for all $i > 2$.*

Proof: By the previous lemma, every $\mathcal{R}$-category has the homotopy type of a 2-dimensional cell category. Since homology groups are well-defined up to homotopy type, it is sufficient to show that the conclusion holds for 2-dimensional cell categories. But since cell categories are particular types of CW complexes, and the n-th homology groups of a CW complex are generated by the characteristic maps on its n-cells, the homology groups beyond dimension 2 must be trivial. □

Lemma 6.16 *Let $\mathcal{R}$ be a set of general relators, and let $\mathcal{R}'$ be a set of standard representatives for $\mathcal{R}$. The Poincaré construction $\mathcal{P}(A, \mathcal{R}')$ can be viewed as a connected subcomplex of the simplicial subdivision of the Poincaré construction $\mathcal{P}(A, \mathcal{R})$, and the inclusion map is a homotopy equivalence. In particular, the fundamental groups of the two constructions are the same, showing that $G = \langle A|\mathcal{R}'\rangle$ is a standard presentation of the group of the general presentation $\langle A|\mathcal{R}\rangle$.*

Proof: By Lemma 6.14, there is a deformation retraction of the Poincaré construction $\mathcal{P}(A, \mathcal{R})$ onto an $\mathcal{R}'$-category during which the 1-skeleton remains fixed. In particular, the $\mathcal{R}'$-category which results contains only 1 vertex, and there is a unique way of attaching every labeled 2-cell in $\mathcal{R}'$ to the 1-skeleton. Finally, since every general relator $R \in \mathcal{R}$ occurs exactly once in the Poincaré construction, then every standard representative in $\mathcal{R}'$ occurs exactly once, and the resulting deformation retraction must be already collapsed as an $\mathcal{R}'$-category. By Lemma 6.9 it is the Poincaré construction of the presentation $\langle A|\mathcal{R}'\rangle$. Since deformation retractions preserve homotopy type, the fundamental groups must be the same. □

According to Lemma 6.16, if $\mathcal{R}'$ is a set of standard representatives of a set $\mathcal{R}$ of general relators, then $G = \langle A|\mathcal{R}\rangle = \langle A|\mathcal{R}'\rangle$. As a consequence, G is independent of the choice of standard representatives. Given any general presentation $G = \langle A|\mathcal{R}\rangle$, there is a general presentation $G(k) = \langle A|\mathcal{R}(k)\rangle$ for each value of $k \in \mathbf{Z}^+$. If $\mathcal{R}$ has a set of standard representatives $\mathcal{R}'$, then the general presentations of G, $G(1)$, $G(2)$, etc., can be converted into standard presentations. In this case there are obvious group homomorphisms from $G(k)$ onto $G(k+1)$ since $G(k+1)$ is obtained from $G(k)$ by the addition of relations, and the group G can itself be described as the direct limit of the groups $G(k)$. This observation is recorded below. The final lemma is an application of Lemma 6.16 which will be needed later.

Lemma 6.17 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let $\mathcal{R}$ have a set of standard representatives. Then G is the direct limit of the groups $G(k)$.*

Lemma 6.18 *The universal cover of a deformation retraction is a deformation retraction of the universal cover. In particular, if $\mathcal{R}$ is a set of general relators and $\mathcal{R}'$ is a set of standard representatives, then $\mathcal{P}(A, \mathcal{R}')$ can be viewed as a subspace of $\mathcal{P}(A, \mathcal{R})$, and the universal cover of the former can be viewed as a subspace of the universal cover of the latter.*

Proof: Let C be the universal cover of an $\mathcal{R}$-category B with covering map $f : C \to B$, and let $h : B \times I \to B$ be a deformation retraction from B onto a subcomplex. Since C is simply connected, so is $C \times I$, and thus the function $g : C \times I \to B$ defined by $g(c, x) = h(f(c), x)$ must lift through f to C by Lemma 1.13. A routine verification shows that the lift of g is a deformation retraction of C onto the universal cover of the subcomplex of B. $\square$

A similar result holds for regular covers of $\mathcal{R}$-categories, but it will not be needed here. The deformation retraction of a collapse, however, is not necessarily even of the same homotopy type as the collapse of the deformation retraction, as illustrated by Example 2 in Section 8.

7 $\mathcal{R}$-Structures

In this section the focus is on maps (in the sense of small cancellation theory), $\mathcal{R}$-diagrams, and $\mathcal{R}$-spheres. After the key concept of a map is defined, some examples are given, and a few results from Lyndon and Schupp ([9]) are quoted. In the few cases where the proofs in [9] are stronger than the statements of the lemmas to which they are attached, the statements given here have been altered to capture the full import of the proofs. The extra flexibility will be used in the general version of small cancellation theory. The section concludes with a general version of van Kampen's Lemma which is applicable in all $\mathcal{R}$-categories so long as the general relators in $\mathcal{R}$ are thin.

7.1 Maps

A map is an embedding of the geometric realization of a finite 2-dimensional cell category in the Euclidean plane or, equivalently, it is a bounded finite union of disjoint subsets of the plane homeomorphic to points, open intervals and open disks, and subject to those conditions implicit in the first chacterization. The implicit conditions are, namely, that the topological boundary of an open disk be the union of a finite set of points and open intervals, and that the topological boundary of an open interval be a finite set of points. The latter is the definition of a map given by Ol'shanskii in [15].

Let N be a component of the complement of a connected map M. The boundary of N is the unique reduced loop in M for which there is an ϵ-deformation to a loop in N for all arbitrarily small $\epsilon > 0$. The boundary of

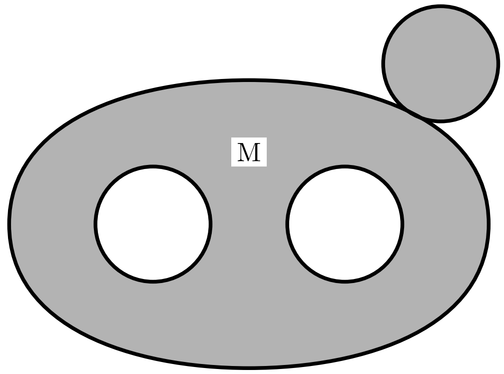

Figure 16: Non-simple boundary loops

a connected map M is the set of boundaries of the components of its complement, and the boundary of a disconnected map M is the union of the boundaries of its connected components. The set of loops which is the boundary of M is denoted ∂M.

Example 1 In Figure 16 a map M is shown schematically. The shaded portion represents the union of the embedded points, lines and cells. The boundary of the map consists of three loops. The two boundary loops arising from the bounded components of the complement happen to be simple loops and the boundary loop defined using the unbounded component is not simple. This example shows how the use of ϵ-deformations in the definition selects the way of traversing the non-simple loop which corresponds to our intuitive notion of a boundary.

If M is connected and simply connected — the most important case — then ∂M is a single loop. If M is an annular map, that is, if M is connected and the complement of M has exactly two components, then ∂M consists of two loops, which may or may not be disjoint.

Example 2 The map sketched in Figure 17 is an example of an annular map in which the two boundary loops are not disjoint.

By a region of M is meant either an open 2-cell D in M or a component of the complement of M. Every edge in M is on the border of exactly two regions. The standard orientation of the plane induces an orientation of every region of M, and thus an orientation of the loops ∂D and an orientation of the loops in ∂M. Every edge of M is given opposite orientations by the two regions it borders. Thus every oriented edge in M is a properly oriented edge of a unique boundary loop of a unique region of M.

Notice that any orientation of the plane induces a clockwise orientation on one of the boundary cycles of an annular diagram and a counter-clockwise ori-

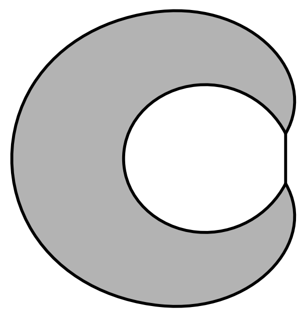

Figure 17: Non-disjoint boundary loops

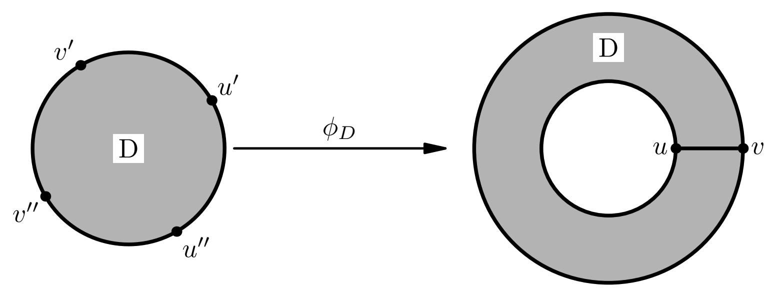

Figure 18: A self-bordering cell with attaching map

entation on the other. Since the standard orientation of the plane is counterclockwise, and boundary cycles are oriented according to the orientation of the component of the complement which they bound, the standard orientation of the boundary cycles of an annular diagram are counterclockwise on the inner cycle and clockwise on the outer cycle.

An arc of the 2-cell D is a path in ∂D whose image is an arc in M. A boundary arc of D is an arc of D whose image is in ∂M. An interior arc of D is an arc of D which is not a boundary arc. The number of arcs in D is called the degree of D and written $d(D)$. The number of interior arcs in D is called the interior degree of D and written $i(D)$. Notice that a single arc in M might have two distinct preimages in ∂D if the two orientations of the edges in the arc both lift to oriented edges in D, in which case it is counted twice in $d(D)$ and $i(D)$. A boundary cell is a 2-cell D whose image in M intersects ∂M, and a cell which is not a boundary cell is an interior cell. A boundary cell D such that $\partial D \cap \partial M$ consists of a single boundary arc is called an exposed cell. The notation $\sum^*$ denotes a summation performed only over the exposed cells. Finally, a $C(p)$-map is a map with no vertices of degree 1 and with $i(D) \geq p$ whenever D contains no boundary arcs.

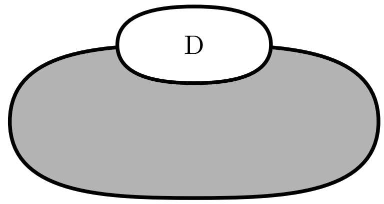

Figure 19: An exposed cell

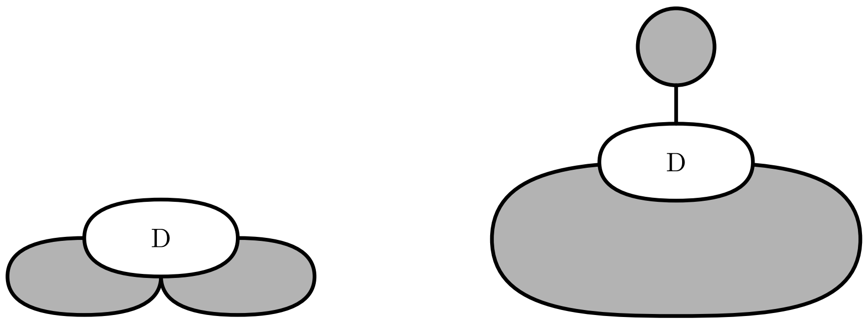

Figure 20: Two non-exposed cells

Example 3 On the righthand side of Figure 18 is a map containing a single 2-cell which is self-bordering. It also contains two vertices and three edges. On the lefthand side is the slice of the 2-cell D. Notice that paths from u' to v' and from u'' to v'' are both sent under the attaching map ϕ_D to the arc between u and v, which causes this arc to be counted twice when calculating the degrees of D. Specifically, the degree of D, $d(D)$, is 4, while the interior degree of D, $i(D)$, is 2, even though the map contains only 3 arcs, and only 1 interior arc. Notice that even though ∂D is a simple loop, the image of ∂D under ϕ_D is no longer simple.

Example 4 A typical exposed cell is shown schematically in Figure 19, while Figure 20 shows two examples of 2-cells which are almost but not quite exposed cells. The first example is not exposed because the intersection of ∂D and ∂M is a single arc plus an isolated point on the other side of the cell. The existence of a single arc of D which contains all of the edges in the overlap is a necessary but not a sufficient condition for D to be exposed. In the second example, the edges in the intersection of ∂D and ∂M lie on a path but D is not exposed since the vertex of degree 3 in the interior of the path means that the path is not a single arc.

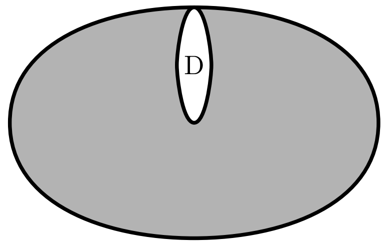

Figure 21: A boundary cell

Example 5 In order for a cell to qualify as a boundary cell it is not necessary for it to contain an edge of ∂M. A vertex from ∂M is sufficient. This situation is illustrated in Figure 21.

Using slightly different notation, Lyndon and Schupp [9] prove the following three lemmas.

Lemma 7.1 *If M is a connected and simply connected $C(6)$-map with more than one 2-cell, then*

$$\sum{}^{*}(4 - i(D)) \geq 6$$

.

Lemma 7.2 (Greedlinger) *If M is a connected and simply connected $C(6)$-map with more than one 2-cell, then there exist exposed cells satisfying one of the following:*
1) $i(D_1), i(D_2) \leq 1$
2) $i(D_1), i(D_2), i(D_3) \leq 2$
3) $i(D_1), i(D_2) \leq 2$ and $i(D_3), i(D_4) \leq 3$
4) $i(D_1) \leq 2$ and $i(D_2), i(D_3), i(D_4), i(D_5) \leq 3$
5) $i(D_1), i(D_2), i(D_3), i(D_4), i(D_5), i(D_6) \leq 3$

Notice that Greedlinger's Lemma is more than just a statement of the consequences of the previous lemma. It states, for example, that if only two exposed cells have internal degrees less than 4 then both must internal degrees ≤ 1, even though $i(D_1) = 0$ and $i(D_2) = 2$ satisfies the earlier inequality.

Lemma 7.3 *Let M be an annular $C(6)$-map and let X and Y be the boundary loops of M. If for every exposed cell D in M an edge of X in ∂D implies $i(D) > 4$, and an edge of Y in ∂D implies $i(D) > 3$, then every 2-cell in M has exactly two boundary arcs, one arc in each boundary loop, and its internal degree is at most 2.*

These lemmas correspond to lemmas V.4.3, V.4.5 and a combination of lemmas V.5.3 and V.5.5 in [9]. In Lemma V.5.3, Lyndon and Schupp make use of

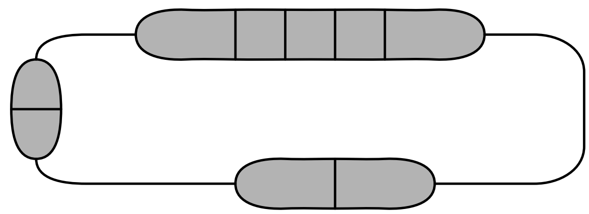

Figure 22: Islands and bridges

the additional assumption that M is a $C(7)$-map, but with a little extra work, this can be reduced to the assumption that M is a $C(6)$-map, given the assumption which has been added in this formulation that $i(D) > 4$ for exposed cells on one of the boundary loops. More specifically, in equation (5.3) the second and third summations are non-positive, and the first sum must be negative under the conditions listed above. This contradiction shows that hypothesis (C) in Lemma V.5.3 is false, and thus hypothesis (C) in Lemma V.5.5 is true.

In a map such as that described in Lemma 7.3, an arc shared by the two boundary loops of M is called a bridge, since it looks like a bridge connecting two islands. An island is a connected component of the annular diagram once the bridges are removed. The boundary of an island consists of a path in each of the boundary cycles such that the two paths have the same endpoints but their interiors are disjoint. An internal arc in such a diagram is called a rung, and boundary arcs which are not bridges are called sides. The terminology arises from thinking of the islands as ladders.

Example 6 The annular map shown in Figure 22 illustrates the terminology above. It contains three islands which contain one, two and five 2-cells respectively. The island with five cells contains four rungs and ten side arcs. The island with only one cell contains two side arcs and no rungs. The three islands are linked together by three bridges.

7.2 $\mathcal{R}$-Diagrams and $\mathcal{R}$-Spheres

A map M is said to be labeled by A if the embedded cell category is labeled by A. To accommodate later usage, the definitions given below will be stated in terms of an arbitrary set of cycles. If $\mathcal{R}$ is a set of standard relators, then $\mathcal{R}$ itself can be viewed as a set of cycles. If $\mathcal{R}$ is a set of general relators, then a set of cycles can be produced by considering all of the representatives of the general relators in $\mathcal{R}$. By Lemma 5.1 there exists at least one representative for each general relator in $\mathcal{R}$. By an abuse of notation, $\mathcal{R}$ will also be used to denote the set of cycles derived from a set of general relators.

Let M be a map labeled by A and let $\mathcal{R}$ be viewed as a set of cycles. If D is a 2-cell of M attached by the map $\phi_D : D \to M$ then the labeling of M induces a labeling on the boundary of D. A van Kampen diagram over $\mathcal{R}$ is a labeled map M such that the cycle read on the boundary of each 2-cell is in $\mathcal{R}$. Diagrams over $\mathcal{R}$ are also called $\mathcal{R}$-diagrams. More generally, if M is a labeled, 2-dimensional cell category, and for every 2-cell D attached to M the cycle read on the boundary of D is contained in $\mathcal{R}$, then M is called an $\mathcal{R}$-structure. If in addition the cell category M is homeomorphic to a 2-sphere, then M is called an $\mathcal{R}$-sphere.

Let M and N be $\mathcal{R}$-diagrams and let $f : N \to M$ be a label-preserving cell functor. The $\mathcal{R}$-diagram N is called a subdiagram of M if the composition of f with the embedding of M into the plane can itself be made into an embedding by an ϵ-deformation of the composition for all $\epsilon > 0$. This definition is formulated so that the simply connected diagram obtained by cutting along a simple path connecting the two boundary loops of an annular diagram is an example of a subdiagram. Similarly, if N is an $\mathcal{R}$-diagram, M is an $\mathcal{R}$-sphere, and $f : N \to M$ is a label-preserving cell functor such that for all $\epsilon > 0$ there is an ϵ-deformation which turns f into an embedding into the underlying 2-sphere, then N is called a subdiagram of the $\mathcal{R}$-sphere M.

If every cycle W is readable as a representative cycle in at most one general relator R in $\mathcal{R}$, then $\mathcal{R}$ is said to have distinct representatives. When a set $\mathcal{R}$ of general relators has distinct representatives, it is possible to assign a number to each 2-cell in an $\mathcal{R}$-structure based on the rank of the general relators it represents. The rank of an $\mathcal{R}$-structure is defined to be the maximum of the ranks of its cells, and the type of an $\mathcal{R}$-structure, denoted Type(M), is the number of 2-cells of each rank, plus the number of edges in the structure. If M is an $\mathcal{R}$-structure of rank k, then the type of M is an ordered k-tuple. Types are ordered lexiconigraphically, first by rank and then by the number of cells in each rank from largest to smallest, and finally by the number of edges.

Specifically, if M and N are $\mathcal{R}$-structures of types Type(M) and Type(N) respectively, then Type(M) $<$ Type(N) if the rank of M is less than that of N, or if they are both of rank k but M has fewer cells of rank k, or if they both have rank k and they both have the same number of rank k cells but M has fewer cells of rank $k - 1$, etc., or if M and N have exactly the same number of cells in every rank and M contains fewer edges. The most important fact about the ordering of types is that, because it is a lexiconigraphic ordering of noetherian orders, it is itself noetherian, meaning that all infinitely decreasing chains must stabilize. In particular, it is possible to use induction on types.

Let α be a fixed positive constant. A path U in a measured relator R is long relative to α iff $d_R(U) \geq \alpha$, otherwise it is said to be short relative to α. Such paths will also be referred to simply as long and short when the constant α is understood from context. An arc U in a diagram Δ is called a long arc if U is in the boundary cycle of a 2-cell D which represents a measured relator R, and U is long under the reading of U in R given by composing the reading of U in ∂D with the reading of ∂D as a representative in R. In small cancellation presentations different readings of the same representative in a relator differ by an automorphism of the relator and thus yield the same length on the arc U by Property 6. In other situations, a particular reading of the boundary cycles of the cells in the measured relators which they represent must be chosen and

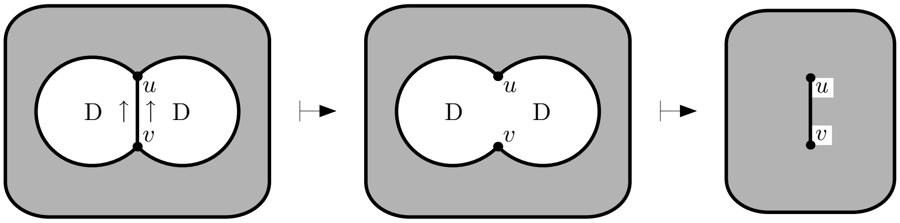

Figure 23: Removing a cancellable pair

fixed as part of the definition of a diagram over such a set of measured relators.

Lemma 7.4 *Let $\mathcal{R}$ be a set of measured relators, and suppose that $\alpha \leq \frac{1}{p}$. If Δ is an $\mathcal{R}$-diagram with no vertices of degree 1 and no internal arcs which are long relative to α, then Δ is a $C(p+1)$-map.*

Proof: If D is a representative cell in Δ with no boundary arcs, then by definition it contains only short internal arcs, each of which has length strictly less than $\alpha \leq \frac{1}{p}$. Since by Lemma 5.25 the lengths must add to at least 1, there must be at least $p+1$ arcs. This shows that for every such D, $d(D) \geq p+1$ and completes the proof. □

Two $\mathcal{R}$-diagrams are called equivalent if they have exactly the same list of boundary cycles, even when orientations and repetitions are taken into consideration. The goal of much of small cancellation theory is to start with an $\mathcal{R}$-diagram and then systematically alter it to obtain an equivalent diagram which has additional, more managable properties. The most familiar type of alteration is the removal of what is known as a cancellable pair. The description of a cancellable pair given below is a slight generalization of the usual definition. Let N be an $\mathcal{R}$-diagram with exactly two closed 2-cells which intersect in exactly one closed edge, and suppose that N has no other 2-cells, edges or vertices. If the 1-skeleton of N can be mapped into a general relator R so that the boundary cycle of each of the 2-cells is sent to a representative of R, and the image of the boundary cycle of N has winding number 0, then N is called a cancellable pair. If Δ is an $\mathcal{R}$-diagram or an $\mathcal{R}$-sphere which contains N as a subdiagram, then intuitively N can be 'cut out' and the hole can be 'sewn up'. Complications can arise because of the creation of new $\mathcal{R}$-spheres. These will be treated in more detail below. An uncomplicated case is illustrated schematically in Figure 23.

It should be noted that the image of N in Δ may contain many more than two boundary arcs and one internal edge. The numbers are derived from a consideration of N as an $\mathcal{R}$-digram in its own right prior to its embedding in Δ. In particular, even if two cells in Δ contain more than one internal edge from which these cells form a cancellable pair, the embedding of the $\mathcal{R}$-diagram effectively selects one of the edges for consideration. Finally, notice that in traditional small cancellation theory, the set $\mathcal{R}$ of possible cycles is reduced, so that the possible cancellable pairs are severely restricted. In particular, if N is a cancellable pair over a traditional set of reduced cycles and v is a vertex on the shared edge of N, the boundary cycle of one 2-cell starting at v, reading counter-clockwise, must be the same as the boundary cycle of the other 2-cell read clockwise starting at v. Thus the notion of a cancellable pair given here is an extension of the traditional concept.

7.3 Boundaries

Let M be an $\mathcal{R}$-diagram, and let $W_1, W_2, \ldots, W_k$ be the properly oriented boundary cycles of M, with possible repetition. If there does not exist an $\mathcal{R}$-diagram whose properly oriented boundary cycles are a non-empty proper subset of this list then these cycles are said to be linked. The concept of an $\mathcal{R}$-diagram with linked boundaries can be reformulated in terms of $\mathcal{R}$-spheres. Let $\mathcal{R}^+$ be the union of the cycles $\mathcal{R}$ and the boundary cycles of an $\mathcal{R}$-diagram M. The embedding of M in the plane combined with the standard embedding of the plane into the 2-sphere creates an embedding of M in the 2-sphere which can then be viewed as a $\mathcal{R}^+$-sphere. From this perspective, one cell has been added to M for each connected component of the complement of M in the 2-sphere. Then these new cells will be called 'phantom' cells.

The boundaries of M are not linked if there is an $\mathcal{R}^+$-sphere whose phantom cells can be put in 1 to 1 correspondence with some, but not all, of these phantom cells. The orientations, labels, and the number of repetitions of the boundary cycles, as well as the absence of any other phantom cells, are clearly important factors in determining whether a particular $\mathcal{R}$-sphere demonstrates that a list of cycles is not linked. Conversely, the boundary cycles of M are linked if no such $\mathcal{R}^+$-sphere exists. By far the most important examples of diagrams with linked boundaries are arbitrary connected and simply connected $\mathcal{R}$-diagrams and annular $\mathcal{R}$-diagrams in which the boundary cycles do not bound connected and simply connected $\mathcal{R}$-diagrams.

The next two lemmas involve identifications of boundaries. The first is a simple topological result, while the second is a variation which takes into account the labeling on the boundary. Despite the simplicity of their proofs, these results will be extremely useful. The difference between Lemma 7.5 and Lemma 7.6 is that in the latter the identifications do not collapse edges to points. The lemmas are stated in terms of one-point products. If B and C are topological spaces, then a one-point product is given by selecting a point in B and a point in C and then identifying them. The one-point product of a finite list of spaces is simply a repetition of this operation a finite number of times. In this definition, the points selected can vary from product to product. Thus, for example, the one-point product of a finite number of edges is a finite tree, and every finite tree is the one-point product of a finite number of edges. In an arbitrary topological space the choice of points to attach is completely arbitrary. In circular categories, however, the choice of points will be restricted to the vertices in the 1-skeletons.

Lemma 7.5 *If M is a connected and simply connected map, then the identification of its entire boundary to a single point yields a structure which is a one-point product of a finite number of spheres.*

Proof: The identification of the boundary to a point will be done in stages. First remove all cut edges by identifying them to a point one at a time. Next proceed by induction on the number of cut vertices which exist in M. If none exist, then either M is a single point or, since there are no cut edges, the boundary loop of M is simple and its identification to a point produces a sphere. If a cut vertex exists then the inductive hypothesis can be applied to each of the connected components formed by its removal and then these structures can be reattached using a one-point product. □

Lemma 7.6 *Let $\mathcal{R}$ be a set of cycles, and let M be a connected and simply connected $\mathcal{R}$-diagram. If the boundary cycle of M is a Dyck word, then there is an identification of the boundary edges so that the result is a one-point product of a finite number of $\mathcal{R}$-spheres and labeled edges.*

Proof: The proof proceeds by induction on the length of the boundary cycle. By Lemma 4.2 the length is even, so the minimal length is 2, and the boundary cycle must be aa^{-1} for some $a \in A$. If the edges labeled a and a^{-1} are distinct then their identification forms a sphere both in the case where the diagram contains two vertices and where it contains only a single vertex. If these edges are not distinct then this edge is in itself the complete diagram. In either case, the lemma is true for diagrams whose boundary cycles have length 2.

Next assume that the result has been shown for all diagrams whose boundary cycles have length at most $2(k-1)$, $k > 1$, and let M be a diagram with a boundary cycle of length $2k$. By Lemma 4.2 there exist a pair of adjacent oriented edges in the boundary cycle labeled aa^{-1} for some $a \in A$. Let e_1 and e_2 be such a pair of adjacent edges with v the vertex between them in the boundary cycle. If the initial vertex of e_1 and the final vertex of e_2 are distinct then the map can be stretched so that the edges e_1 and e_2 and their vertices are identified in a portion of the unbounded component of the complement, thus keeping the map planar. The inductive hypothesis can then be applied to this new diagram.

If, however, the initial vertex of e_1 and the terminal vertex of e_2 are one and the same vertex, say a vertex called u, then the subdiagram bounded by the loop e_1e_2 is connected, simply connected and attached to the rest of M only at u. If this subdiagram is temporarily removed then the inductive hypothesis can be applied to each of the two connected components. Finally, it is clear that the pieces can be reattached at appropriate points using a one-point product. This completes the proof. □

Notice that since in Lemma 7.6 the only identifications are of edges, the number and the rank of the 2-cell is unchanged, and the type of the $\mathcal{R}$-structure remains constant.

7.4 A General Van Kampen Lemma

Lemma 7.7 and Lemma 7.8 will provide the connection between the geometry of $\mathcal{R}$-diagrams and the algebra of presentations which is needed to prove a general version of van Kampen's Lemma. For proofs see [9].

Lemma 7.7 (van Kampen) *Let $G = \langle A|\mathcal{R}\rangle$ be a group presentation. A word $W \in A^*$ is equivalent to 1 in G iff there is a connected and simply connected $\mathcal{R}$-diagram M whose boundary cycle is W.*

Lemma 7.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a group presentation, and let X and Y be words which are not equivalent to the identity in G. The words X and Y are conjugate in G iff there is an annular $\mathcal{R}$-diagram M with boundary cycles X and Y^{-1}.*

These results can be extended to general presentations using the earlier results on deformation retractions.

Lemma 7.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation. If the general relators in $\mathcal{R}$ are thin, then a word $W \in A^*$ is equivalent to 1 in G iff there is a connected and simply connected $\mathcal{R}$-diagram M whose boundary cycle is W.*

Proof: Since the general relators are thin, there is a set of standard representatives $\mathcal{R}'$ for $\mathcal{R}$ by Lemma 5.19. By Lemma 6.16, $G = \langle A|\mathcal{R}'\rangle$ is a standard group presentation for G. Since all $\mathcal{R}'$-diagrams are also $\mathcal{R}$-diagrams, the result follows by Lemma 7.7. □

Lemma 7.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let X and Y be words which are not equivalent to the identity in G. If the general relators in $\mathcal{R}$ are thin, then the words X and Y are conjugate in G iff there is an annular $\mathcal{R}$-diagram M with boundary cycles X and Y^{-1}.*

Proof: The proof is the same as in the previous lemma. Since the general relators are thin, there is a set of standard representatives $\mathcal{R}'$ for $\mathcal{R}$, by Lemma 5.19. By Lemma 6.16 $G = \langle A|\mathcal{R}'\rangle$ is a standard group presentation for G. Since all $\mathcal{R}'$-diagrams are also $\mathcal{R}$-diagrams, the result follows by Lemma 7.8. □

An $\mathcal{R}$-diagram like that described in Lemma 7.9 is called a proof that W equals 1 in G, and the $\mathcal{R}$-diagrams described in Lemma 7.10 are called conjugacy diagrams, or proofs that X and Y are conjugate in G. A more detailed analysis using the Simplicial Approximation Theorem would allow the condition that the relators are thin to be removed from the hypotheses of Lemma 7.9 and Lemma 7.10, but these lemmas will suffice for the general small cancellation theory discussed here, since all of the general relators in a general small cancellation presentation will necessarily be thin. Two useful consequences of Lemma 7.9 are given below.

Lemma 7.11 *Let $\mathcal{R}$ be a set of thin general relators, and let C be an $\mathcal{R}$-category. If W is a contractible loop in C, then there is a connected and simply connected $\mathcal{R}$-diagram whose boundary cycle is W.*

Proof: By Lemma 1.9 there is a map from the unit disk into C where the restriction of the map to the boundary recovers the reading of the loop W. Composing this with the unique map to the Poincaré construction of these general relators (Lemma 6.8) yields an image of a disk which shows that W is contractible in the Poincaré construction. By definition this means that W is equivalent to 1 in the group, and by Lemma 7.9 an $\mathcal{R}$-diagram such as the one described must exist. □

Lemma 7.12 *Let $\mathcal{R}$ be a set of thin general relators, and let R be a general relator in $\mathcal{R}$. If W is a loop in R of winding number 0, then there is a connected and simply connected $\mathcal{R}$-diagram Δ with boundary W which contains 2-cells which represent only general relators in ∂R. In particular, the 2-cells in Δ represent only general relators whose height is strictly less than that of R.*

Proof: Since W has winding number 0 it is contractible in ∂R, and thus Lemma 7.11 can be used to complete the proof. □

7.5 Removing Cancellable Pairs

The next several lemmas show how cancellable pairs can be systematically removed.

Lemma 7.13 *Let $\mathcal{R}$ be a set of graded, thin general relators with distinct representatives. If Δ is a $\mathcal{R}$-sphere which contains a cancellable pair N as a subdiagram, then the cancellable pair N can be removed by a process which creates a one-point product of labeled edges and $\mathcal{R}$-spheres called Δ' such that Type$(\Delta) >$ Type(Δ').*

Proof: If the image of the two open cells of N and the single internal edge are removed from Δ, then the remaining $\mathcal{R}$-structure is connected and simply connected, and since it can be embedded in the plane, it is also an $\mathcal{R}$-diagram, say Δ'. The boundary of Δ' is by definition a cycle W which can be read as a loop of winding number 0 in the boundary of some general relator $R \in \mathcal{R}$. By Lemma 7.12 there is a connected and simply connected $\mathcal{R}$-diagram Δ'' with the boundary W, but whose 2-cells represent general relators only in ∂R. The $\mathcal{R}$-diagram Δ'' can be embedded in the plane upside-down, so that its boundary cycle is W^{-1}, and attached to Δ' at a single point producing a new connected and simply connected $\mathcal{R}$-diagram with boundary cycle WW^{-1}. The boundary cycle is clearly a Dyck word, and the type of this $\mathcal{R}$-diagram is strictly less than that of the $\mathcal{R}$-sphere Δ since two cells representing R were removed but only cells of strictly lower rank were added. Finally, Lemma 7.6 can be used to identify the boundary edges and complete the proof. □

Lemma 7.14 *Let $\mathcal{R}$ be a set of graded, thin general relators with distinct representatives. If Δ is a one-point product of a finite number of labeled edges and $\mathcal{R}$-spheres, then the process of repeatedly removing cancellable pairs can be used to create a one-point product of $\mathcal{R}$-spheres and labeled edges called Δ' such that Type$(\Delta) \geq$ Type(Δ') and Δ' contains no cancellable pairs.*

Proof: The proof proceeds by induction on the type of Δ. If Δ has no cells then it is a finite tree, and the lemma is true. So suppose that Δ is a finite one-point product of labeled edges and $\mathcal{R}$-spheres and that the lemma has been shown for all such $\mathcal{R}$-structures of strictly lower type. If there does not exist a cancellable pair in Δ then Δ itself satisfies the lemma, so let N be a cancellable pair in Δ. Since the two 2-cells in N share an edge, N must actually be a subdiagram of one of the $\mathcal{R}$-spheres in the product which forms Δ. Since the procedure used in Lemma 7.13 to remove a cancellable pair removes two 2-cells and an internal edge but no vertices, any vertices used for one-point attachments remain even after the removal of the cancellable pair. Thus it is possible to form the same one-point products with the other labeled edges and $\mathcal{R}$-spheres after the cancellable pair has been removed. Since this new product has strictly lower type, the induction hypothesis completes the proof. □

Lemma 7.15 *Let $\mathcal{R}$ be a set of graded, thin general relators with distinct representatives. If Δ is an $\mathcal{R}$-diagram whose boundaries are linked, then there exists an equivalent $\mathcal{R}$-diagram Δ' with no cancellable pairs, and Type$(\Delta) \geq$ Type(Δ').*

Proof: The proof again proceeds by induction on the type of Δ. If Δ contains no 2-cells then clearly Δ has no cancellable pairs, and Δ itself satisfies the lemma. So suppose that Δ is an $\mathcal{R}$-diagram with linked boundaries and that the lemma has been shown for all such $\mathcal{R}$-diagrams of strictly lower type. If there does not exist a cancellable pair in Δ then Δ itself satisfies the lemma, so let N be a cancellable pair in Δ.

Adding phantom cells for each of the connected components of the complement creates an $\mathcal{R}^+$-sphere Δ'. The cancellable pair N is a subdiagram of Δ' as well. If the cancellable pair N is removed using the procedure in Lemma 7.13, then the phantom cells persist in the result. Moreover, since they are linked, they must all be contained in the same $\mathcal{R}^+$-sphere. Ignoring the other labeled edges and $\mathcal{R}^+$-spheres in the one-point product and removing the phantom cells leaves an $\mathcal{R}$-diagram with exactly the same boundaries as Δ, but of strictly lower type. The induction hypothesis can now be used to complete the proof. □

8 Cayley Categories

The primary goal of this section is to introduce the Cayley category of a general presentation, and to prove its fundamental properties. Along the way a few results are given on regular covers and automorphism groups of $\mathcal{R}$-categories. The section concludes with a pair of key examples of these concepts.

8.1 Automorphisms

The next several lemmas investigate the automorphism group of regular covers of Poincaré constructions, as well as those of the collapsed versions of these spaces. Of particular interest will be the collapse of the universal cover of a Poincaré construction because of its role in the later developments.

Lemma 8.1 *Let C be a regular cover of the Poincaré construction of a general presentation. There is an isomorphism between two closed cones C/c_1 and C/c_2 iff c_1 and c_2 are lifts of the same object c in the Poincaré construction iff there is a deck transformation of C which sends c_1 to c_2. Moreover, such automorphisms, when they exist, are unique.*

Proof: The first two conditions are equivalent by Lemma 6.8, while the equivalence of the last two follows from the definition of a regular cover. The uniqueness of this automorphism is guaranteed by Lemma 1.14. □

Lemma 8.2 *If C is a regular cover of the Poincaré construction of a general presentation, then the group of label-preserving automorphisms of $\overline{C}$ is isomorphic to that of C. Moreover, both automorphism groups are determined by their action on a single vertex.*

Proof: Since the 1-skeleton of C is a Cayley graph by Lemma 6.12 and Cayley graphs are deterministic, Lemma 6.2 shows that the 1-skeleton of C is the same as that of $\overline{C}$. In addition the construction of $\overline{C}$ is invariant of any label-preserving automorphism, so the deck transformations of C must induce label-preserving automorphisms of $\overline{C}$, and since the 1-skeleton remains the same, distinct automorphisms of C must remain distinct. Thus the label-preserving automorphisms of $\overline{C}$ contain at least the deck transformations of C as a subgroup. Finally, since any such automorphism must send a vertex to a vertex, and an automorphism is trivially a covering map, Lemma 1.14 shows that the two groups are identical, and that their actions are determined by a single vertex. □

Lemma 8.3 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let C be the universal cover of its Poincaré construction. If $\overline{C}$ is the collapse of C, then the 1-skeleton of $\overline{C}$ is the same as the 1-skeleton of C, and thus is the Cayley graph $\mathcal{C}(G, A)$. Moreover, the automorphism group of $\overline{C}$ is transitive on its vertices and isomorphic to G.*

Proof: The result is a special case of Lemma 8.2 using the fact that the fundamental group of C is trivial and the deck transformations form the group G. □

Lemma 8.4 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation and let C be the universal cover of its Poincaré construction. If $\overline{C}$ is the collapse of C then $\overline{C}$ is connected and simply connected.*

Proof: By Lemma 8.3 the 1-skeleton of $\overline{C}$ is a Cayley graph, and thus by Lemma 4.4 the 1-skeleton is connected. By Lemma 4.8 this means that $\overline{C}$ is connected. The proof that $\overline{C}$ is also simply connected will proceed by contradiction. If $\overline{C}$ is not simply connected then it has its own universal cover, say C' with a covering map $h : C' \to \overline{C}$. By Lemma 1.13, there is a lift of $f : C \to \overline{C}$ to $g : C \to C'$. Let v be a vertex in $\overline{C}$ and let $u = g(f^{-1}(v))$ which is well-defined since f is an isomorphism on the 1-skeleton by Lemma 6.2. If u' is another vertex in C' such that $h(u') = h(u) = v$, then choose a path U from u to u'. Under h this creates a loop at v, which becomes a loop U at $f^{-1}(v)$, since f is isomorphic on the 1-skeletons. Thus under g, the path in C' must have started out as a loop, and $u = u'$. This shows that $h^{-1}(v)$ contains only a single vertex. Since by Lemma 1.15 the cardinality of $h^{-1}(v)$ is the order of the fundamental group of $\overline{C}$, $\overline{C}$ is simply connected. □

The following lemma shows in more detail the relationship between the universal cover of a Poincaré construction and its collapse.

Lemma 8.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let $f : C \to \overline{C}$ be the $\mathcal{R}$-functor between the universal cover of the Poincaré construction of the presentation and its collapse. If c is an open cone in $\overline{C}$, then the number of open cones in $f^{-1}(c)$ is the order of the group $Aut(\overline{C}/c)$ divided by the order of the normal subgroup of automorphisms g such that $\phi_c g = \phi_c$. In particular, if the attaching functor ϕ_c is injective on vertices then the objects in $f^{-1}(c)$ are in 1 to 1 correspondence with $Aut(\overline{C}/c)$.*

Proof: First of all, if c is a vertex or an edge, then the lemma is trivially true, so assume that c is an open cone of rank at least 2. Let c_1 and c_2 be two open cones in $f^{-1}(c)$. Since they both map to the same general relator in the Poincaré construction, by Lemma 8.1 there is a unique automorphism of C which takes c_1 to c_2. This induces an isomorphism between C/c_1 and C/c_2 since the automorphism is a circular functor (Lemma 6.11 and the definition of a circular functor). Under f the isomorphism becomes an automorphism of the closed cone $\overline{C}/c$ with itself.

Conversely, given any automorphism g of $\overline{C}/c$ let W be a path read in the 1-skeleton of $\overline{C}/c$ from a vertex v to $g(v)$. Under ϕ_c the word W is also a path read in $\overline{C}$. Since by Lemma 6.2 the 1-skeleton remains unchanged under the collapsing functor f, it is also read in the 1-skeleton of C. By Lemma 8.1 there is a unique automorphism of C which sends $f^{-1}(\phi_c(v))$ to $f^{-1}(\phi_C(g(v)))$. Since the automorphism of $\overline{C}/c$ induced by W fixes the 1-skeleton of $\overline{C}/c$ and the open cone c, the corresponding images in $\overline{C}$ are also fixed. This in turn means that this automorphism of C must send an open cone in $f^{-1}(c)$ such as c_1 to another open cone in this set. By Lemma 4.5 the automorphism of C induced by W is the identity iff the path W forms a loop in C (or $\overline{C}$ iff the word W is equivalent to 1 in G. □

An $\mathcal{R}$-category C is called proper if it is an $\mathcal{R}$-complex. By Lemma 3.4 this is true iff the characteristic functors of the slice categories are injective on vertices. A general presentation is called proper when the universal cover of its Poincaré construction is proper. Since the $\mathcal{R}$-functor from the universal cover of a Poincaré construction is an isomorphism when restricted to the 1-skeleton by Lemma 8.3, the universal cover is proper iff its collapse is as well. Thus Lemma 8.5 has the following corollary.

Corollary 8.6 *If $G = \langle A|\mathcal{R}\rangle$ is a proper general presentation, then the universal cover of its Poincaré construction is collapsed iff there are no non-trivial automorphisms of the general relators $R \in \mathcal{R}$.*

8.2 Cayley Categories

A Cayley category is a collapsed, connected, simply connected $\mathcal{R}$-category whose group of label-preserving automorphisms acts transitively on its vertex set. By Lemma 8.3 and Lemma 8.4, the collapse of the universal cover of a Poincaré construction is an example of a Cayley category. The goal of the next few lemmas is to show that these are actually the only examples.

Lemma 8.7 *Let B be a Cayley category and let v be a vertex of B. If R is a general relator used in the construction of B and u is a vertex in R then there exists a unique $\mathcal{R}$-functor from R to B which sends u to v. Similarly, if a is the label on one of the oriented edges of B then there exists a unique oriented edge with v as its initial vertex which is labeled by a.*

Proof: The existence portion of these two statements follows from the facts that they occur somewhere in B and that by definition the automorphism group of B acts transitively on its vertices. Since generators and general relators are connected and B is collapsed, Lemma 6.3 guarantees the uniqueness of each. □

Lemma 8.8 *Let B be a Cayley category, and let Δ be a connected and simply connected $\mathcal{R}$-diagram with boundary cycle W. If B contains all of the generators used and general relators represented in Δ, then the path W is always read as a loop in B.*

Proof: The proof is by induction on the number of 2-cells in Δ. If Δ contains no cells, then it is a finite tree. Start by sending the base of the loop W to a vertex in B. The partial functor into the 1-skeleton of B can be extended one edge at a time. By Lemma 8.7 there is always an edge in B which extends the functor while preserving the label on the edge. The completed functor shows that W is read as a loop in B. If Δ contains exactly one cell, then the boundary of this cell can be read in the unique copy of the general relator it represents attached to the appropriate point in B. This partial map can then be extended to the finite trees which might be attached to the boundary of the single cell one edge at a time, exactly as before. Finally, suppose that the lemma has been shown for all $\mathcal{R}$-diagrams with fewer than k cells and let Δ have exactly k cells, $k > 1$. Using a cell which contains an edge in the boundary, Δ can easily be split into two subdiagrams which have strictly fewer cells and whose boundary cycles once combined yield a word which reduces to W in the free group. Since by induction both of these words label loops in B, so does their concatenation. Moreover, since the 1-skeleton of B is deterministic, the reductions in the free group necessary to produce W can be accomplished in B as well, showing that W is read as a loop. □

Lemma 8.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation in which all of the general relators are thin, and let B be a Cayley category whose content is identical with that of G. If W is a word which is equivalent to 1 in G then the path W always forms a loop when read in B.*

Proof: By Lemma 7.9 there is a connected and simply connected $\mathcal{R}$-diagram whose boundary cycle is W, and then Lemma 8.8 completes the proof. $\Box$

Lemma 8.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation in which all of the general relators are thin. If B_1 and B_2 are Cayley categories whose content is identical with that of G, then they have isomorphic 1-skeletons. In particular the 1-skeleton of a Cayley category which contains generators A and general relators $\mathcal{R}$ is the Cayley graph $\mathcal{C}(G, A)$ where G is the group of the general presentation $G = \langle A|\mathcal{R}\rangle$.*

Proof: Let B be either Cayley category. If a word W is equivalent to 1 in G then by Lemma 8.9 it is read as a loop in B. Conversely let W be read as a loop in B. Since B is simply connected, this loop is also contractible. Thus, by Lemma 7.11 there is a connected and simply connected $\mathcal{R}$-diagram with boundary cycle W, and by Lemma 7.9 W is equivalent to 1 in G. By definition, the 1-skeleton of B is deterministic and the automorphism group acts transitively on its vertices. Since by Lemma 4.8 it is also connected, Lemma 4.4 shows that it is a Cayley graph. Since it is now known that W is read as a loop in B iff W is equivalent to 1 in G, the Cayley graph in question must be $\mathcal{C}(G, A)$. Since all Cayley categories with this content have $\mathcal{C}(G, A)$ as their 1-skeleton, the proof is complete. $\Box$

Lemma 8.11 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let B be a Cayley category whose content is identical with that of G. A path in B forms a loop iff the word read by the path is equivalent to 1 in G. In particular, given words X and Y, XY^{-1} is equivalent to 1 in G iff X and Y are equivalent in G iff X and Y can be read in the Cayley category as paths which start and end at the same vertices.*

Proof: The first statement is a combination of Lemma 8.10 and Lemma 4.5 while the second is an immediate consequence of the first. $\Box$

Lemma 8.12 *All Cayley categories with the same content are isomorphic. In particular Cayley categories are in 1 to 1 correspondence with general presentations and Poincaré constructions, and all Cayley categories are the collapse of the universal cover of a Poincaré construction of some general presentation.*

Proof: Let B_1 and B_2 be two Cayley categories which have the same content. By Lemma 8.10 there is an isomorphism between the 1-skeletons of B_1 and B_2. Suppose that there is an $\mathcal{R}$-isomorphism from the $(k-1)$-skeleton of B_1 to the $(k-1)$-skeleton of B_2. If b is an object of B_1 of rank exactly k, then b is a terminal object and extension of the functor to b will have no effect on the other elements of the k-skeleton, and thus can be chosen independently. By assumption the image of the boundary of b is already set in correspondence with objects and arrows in B_2. In particular, if a vertex in the image of B_1/b is considered, there is a particular vertex in B_2 to which it corresponds. By Lemma 8.7 there exists an $\mathcal{R}$-functor from B_1/b extending this correspondence. Since by Lemma 6.3 this agrees on the boundary of b with the existing identification of the $(k-1)$-skeletons, there is a coherent extension of the functor to the object b. When this is done for all objects of rank k, an $\mathcal{R}$-functor has been produced from the k-skeleton of B_1 to that of B_2. Since by Lemma 8.7 the choices were forced, and since the same procedure could be applied from B_2 to B_1 with the identical correspondence, the $\mathcal{R}$-functor between the k-skeletons must be an isomorphism. Continuing in this way produces the desired isomorphism between B_1 and B_2. As mentioned above, the existence of a Cayley category whose content is the same as that of a given general presentation has already been shown by Lemma 8.3 and Lemma 8.4. This completes the proof. □

In [9] Lyndon and Schupp define a similar construction for traditional 2-dimensional presentations over reduced sets of relators, which they call the Cayley complex of the presentation. The construction defined in [9] is collapsed and it agrees with the definition given here of the Cayley category of a presentation in the cases where both are defined.

8.3 Cyclics and Dihedrals

The two examples given below are good illustrations of the constructions described in the previous sections.

Example 1 Let $\mathbf{Z}_n = \langle a | a^n \rangle$ be the cyclic group of order n. The Cayley graph $\mathcal{C}(\mathbf{Z}_n, a)$ is the abstract loop L_n with all positively oriented edges labeled by a. The cone over this loop is the labeled 2-cell corresponding to the cycle a^n. If a single copy of this 2-cell is attached to the Cayley graph the result is the Cayley category of the presentation. If, on the other hand, n distinct copies of the relator are attached to the loop which is the Cayley graph, then the resulting construction is the universal cover of the Poincaré construction of the presentation. Finally, the Poincaré construction can be obtained by identifying the vertices either of the Cayley category or of the universal cover, and then collapsing the result. The Poincaré construction consists of a single vertex, a single loop edge labeled a, with a single copy of the 2-cell attached.

Example 1 highlights the distinction between the universal cover of the Poincaré construction and the Cayley category of a presentation. The universal cover has many convenient topological properties which are not shared by the Cayley category. However, if the primary intent of an investigation is to prove results about the van Kampen diagrams over the presentation, then the Cayley category is a more appropriate construction to study.

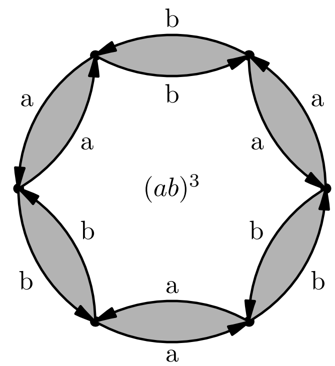

Figure 24: The dihedral group $\mathbf{D}_6$

Example 2 Let $\mathbf{D}_{2n} = \langle a, b | a^2, b^2, (ab)^n \rangle$. The Cayley graph $\mathcal{C}(\mathbf{D}_{2n}, \{a, b\})$ consists of a vertex set $\cong \mathbf{Z}_{2n}$ with two edges between every pair of consecutive vertices. Between the vertices $2i$ and $2i+1$ there are two positively oriented edges, one starting at each vertex and ending at the other, both labeled a. Similarly, between $2i+1$ and $2i+2$ there are two positively oriented edges, one starting at each vertex and ending at the other, both labeled b. In Figure 24 the Cayley graph of $\mathbf{D}_6$ is represented. The arrows are meant to indicate that starting at the vertex 0 and reading clockwise around the outside of the circles produces the word $(ba)^3$, and that reading counter-clockwise around the inside of the circles produces the word $(ab)^3$.

Let S be the Cayley graph $\mathcal{C}(\mathbf{D}_{2n}, A)$ in which between each pair of vertices a 2-cell labeled either a^2 or b^2 is attached depending on the labeled edges available. The Cayley category of the presentation is given by attaching two copies of the 2-cell labeled $(ab)^n$ onto the construction S so that their edges are disjoint. Topologically, the construction is a 2-sphere. It is simply connected but not contractible. The universal cover is given by attaching n distinct copies of the 2-cell labeled $(ab)^n$ to both the top and the bottom, and 2 distinct copies of the 2-cells labeled a^2 and b^2 whenever they occur.

Let $R = \text{Cone}(S)$. The general presentation $\langle A | a^2, b^2, R \rangle$ is an alternative presentation of $\mathbf{D}_{2n}$. Since the relators a^2 and b^2 are contained in R, the Poincaré construction of this presentation is given by identifying the vertices of R and collapsing the result. The universal cover of the Poincaré construction is more easily visualized as $2n$ copies of the general relator R attached to a single copy of the boundary $\partial R = S$. The Cayley category is isomorphic to the general relator R itself.

Since the relators in the standard presentation are simple representatives of the general relators in the general presentation, by Lemma 6.16 there is a deformation retraction of the Poincaré construction of the general presentation onto the Poincaré construction of the standard presentation. By Lemma 6.18, this deformation lifts to a deformation retraction between the universal covers.

That there is no deformation retraction between the Cayley categories is clear since in the case of the general presentation the Cayley category is contractible while in the case of the standard presentation it is not. Finally, for large n the general presentation given here is a general small cancellation group as defined in Section 9.

Because 2-dimensional relators are sufficient to describe any group, geometric group theory is usually done with 2-dimensional cell categories. But by allowing more general, higher-dimensional relators, most of the desirable properties of ordinary relators are preserved, and in addition, more groups fall within the purview of small cancelation theory. In particular, general relators allow the Burnside groups for sufficiently large exponents to be described as generalized twelfth groups. Moreover, in Section 12 it will be shown that the Cayley category of a general small cancellation presentation is contractible even in cases such as the presentation of the dihedral groups as given above, where the Cayley complex of a retracted presentation of the same group is not contractible. The contractibility of this canonical geometric construction is indicative of the benefits, both geometrically and algebraically, to the researcher who uses general relators.

Part IV
Small Cancellation Theory

In Part IV the focus is narrowed once again to examine those general presentations which satisfy a version of small cancellation theory. In Section 9 the general small cancellation theory alluded to in the title of the article is presented along with its most immediate consequences. In Section 10 it is shown how Dehn's algorithm can be applied to $\mathcal{R}$-diagrams over general small cancellation presentations to prove the decidability of the word and conjugacy problems for these groups. Finally, in Section 11 finitely presented general small cancellation groups are shown to be Gromov hyperbolic groups, and arbitrary general small cancellation groups are shown to be the direct limit of hyperbolic groups.

9 General Small Cancellation Theory

This section begins by discussing the traditional small cancellation theory and the ways it can be generalized. Then the axioms used to define general small cancellation groups and presentations are given. After discussing the axioms, a few of the more immediate consequences are shown.

9.1 Traditional Small Cancellation Theory

Let $\mathcal{R}$ be a set of standard relators. In traditional small cancellation theory the relators in $\mathcal{R}$ are thought of as a set of cyclically reduced words closed under inverse and cyclic conjugates. If R and S are distinct words in $\mathcal{R}$, and $R = UV_1$ and $S = UV_2$, then U is called a piece relative to $\mathcal{R}$. Alternatively, a piece can be thought of as a possible label on an internal arc in an $\mathcal{R}$-diagram which has no cancellable pairs. In the first definition it is important to realize that the words R and S may be cyclic conjugates and thus belong to the same cycle even though they are distinct words. Similarly, in the second definition the 2-cells on either side of the internal arc may represent the same relator.

A third definition uses $\mathcal{R}$-functors. Let the relators in $\mathcal{R}$ be viewed as closed cones over labeled abstract loops reading cycles which are reduced in the free group. A word U is not a piece relative to $\mathcal{R}$ if whenever U is read in relators R and S by $\mathcal{R}$-functors f and g respectively, there always exists an $\mathcal{R}$-functor $h : R \to S$ such that $hf = g$. By Lemma 2.1, the functor h must be an isomorphism (since in the traditional theory all relators have the same rank), thus showing that this definition is equivalent to the first two definitions.

The traditional theory focuses on two types of hypotheses called $C(p)$ and $C'(\alpha)$. The first is stated in terms of pieces and the second uses the normalized graph metric. A set of standard relators is said to satisfy $C(p)$ if the relators are cyclically reduced and no word in $\mathcal{R}$ is the product of fewer than p pieces. The notions of a $C(p)$ set of relators and a $C(p)$-map are closely related.

Lemma 9.1 *If $\mathcal{R}$ is a $C(p)$ set of standard relators, and Δ is an $\mathcal{R}$-diagram with cyclically reduced boundaries and no cancellable pairs, then the underlying map of Δ is a $C(p)$-map.*

Proof: To prove the result it is sufficient to show that there are no vertices of degree 1, and that all internal arcs are pieces. Since the boundary cycles of Δ are cyclically reduced, and since $\mathcal{R}$ is a $C(p)$ set of relators, the boundary cycle of every region is cyclically reduced. If Δ contained a vertex v of degree 1, then a small disk around v would be strictly contained in some region and the boundary cycle of this region would not be cyclically reduced, contradiction. If an internal arc is not a piece of a relator, then the label on the arc is the initial segment of a unique word in $\mathcal{R}$ and either the arc is on the border of two distinct 2-cells, or the arc is read in the boundary of a single 2-cell in two distinct ways. In the former case the cells must form a cancellable pair, contradicting the assumption about Δ. And in this latter case, the two readings yield opposite orientations of the boundary, forcing the boundary of the 2-cell to be a word in the free group which is conjugate to its own inverse. By Lemma 4.2 the only word for which this is true is the empty word. Thus every internal arc must be a piece and the proof is complete. □

Lemma 9.2 *Let $\mathcal{R}$ be a $C(p)$ set of standard relators. If Δ is an $\mathcal{R}$-diagram whose boundary cycles are cyclically reduced and linked, then there exists an equivalent $\mathcal{R}$-diagram Δ' over a $C(p)$-map with Type$(\Delta) \geq$ Type(Δ').*

Proof: Since a collection of reduced cycles clearly has distinct representatives and a rank function which assigns every relator the rank of 2, the result follows immediately from Lemma 7.15 and Lemma 9.1 □

A set of cyclically reduced standard relators $\mathcal{R}$ is said to satisfy $C'(\alpha)$ if every piece of every relator has a length strictly less than α in the normalized graph metric. More specifically, let $\mathcal{R}$ be viewed as a set of measured relators by defining $d_R(U)$ as the length of the path U in the normalized graph metric on R for all paths U in R. The condition $C'(\alpha)$ is satisfied by $\mathcal{R}$ if whenever U is a piece relative to $\mathcal{R}$ and U is readable in R, then $d_R(U) < \alpha$. In other words, $\mathcal{R}$ satisfies the condition $C'(\alpha)$ if whenever $d_R(U) \geq \alpha$ for some path U in R, then U is not a piece relative to $\mathcal{R}$. The third description of a piece, given above, can be used to state the condition $C'(\alpha)$ in the form in which it appears in the general theory. A set of standard relators $\mathcal{R}$ is said to satisfy $C'(\alpha)$ if whenever U is read in relators R and S by $\mathcal{R}$-functors f and g respectively and $d_R(U) \geq \alpha$, there exists an $\mathcal{R}$-functor h such that $hf = g$.

The above discussion motivates the following definition. Let $\mathcal{R}$ be a set of measured relators, let C be an $\mathcal{R}$-category, and let μ be a non-negative real. The $\mathcal{R}$-category C is called μ-closed if whenever a word U is read in C by an $\mathcal{R}$-functor g and U is also read in a general relator R by a functor f with $d_R(U) \geq \mu$, there exists a unique functor $h : R \to C$ such that $hf = g$. As an example, Cayley categories are constructed so that every general relator is attached to every vertex in every possible way, so Cayley categories are 0-closed.

Another example is the condition $C'(\alpha)$. A set of standard relators $\mathcal{R}$ satisfies $C'(\alpha)$ iff every relator $R \in \mathcal{R}$ is α-closed with respect to $\mathcal{R}$. The notion of being μ-closed is related to the earlier notions of being μ-reduced, μ-free, or μ-complement-free. If a labeled abstract path or loop is μ-closed, then the word or cycle with which it is labeled is μ-free. Conversely, because of the inclusion or exclusion of equality, abstract paths and loops labeled by μ-free words and cycles are structures which are λ-closed for all $\lambda > \mu$, but not necessarily structures which are μ-closed. For example, in the traditional theory, where the normalized graph metric is the relator metric used, it is clear that a word is μ-reduced iff it is μ-free, and it is also clear that a word can contain exactly one-half of a relator so that it is $\frac{1}{2}$-reduced (Dehn-reduced) but not necessarily $\frac{1}{2}$-closed.

Lemma 9.2 can be combined with the lemmas from Lyndon and Schupp to yield the results below. The results are stated in terms of k-remnants. A word V is called a k-remnant if some word $R \in \mathcal{R}$ has the form $R = U_1 U_2 U_3 \ldots U_k V$ where $U_1, \ldots, U_k$ are pieces.

Lemma 9.3 *Let $\mathcal{R}$ be a set of standard relators satisfying $C(6)$. If W is a cyclically reduced word equivalent to 1 in the group $G = \langle A|\mathcal{R}\rangle$ then either W is equal to 1, W is a relator, or a cyclic conjugate of W has the form $U_1V_1 \ldots U_nV_n$ where each V_k is an $i(V_k)$-remnant such that the number n of the V_k and the numbers $i(V_k)$ satisfy the relation*

$$\sum_{k=1}^{n} (4 - i(V_k)) \geq 6$$

Proof: Consider a reduced proof-diagram whose boundary reads W which must exist by Lemma 7.7 and Lemma 9.2. If W is not equal to 1 or to a relator in $\mathcal{R}$, then the reduced diagram must contain more than one 2-cell. The result follows from Lemma 7.1 by simply letting the words V_k represent the label on the unique boundary arcs of the exposed cells of the diagram. □

Lemma 9.4 *Let $\mathcal{R}$ be a set of standard relators satisfying $C'(\alpha)$ for some $\alpha \leq \frac{1}{8}$. If X and Y are cyclically Dehn-reduced words which are conjugate in the group $G = \langle A|\mathcal{R}\rangle$ but not equivalent to 1 in G, then there exists an annular $\mathcal{R}$-diagram Δ with boundary cycles X and Y^{-1} such that Δ contains no cancellable pairs, and every 2-cell in Δ has internal degree of at most 2. Moreover, every 2-cell contains exactly one arc in each boundary cycle, and the length of the boundary arcs is always more than $\frac{1}{4}$ and no more than $\frac{1}{2}$ of the length of the relator represented by the cell.*

Proof: Since $\mathcal{R}$ is a $C(6)$ set of relators, the diagram Δ is a $C(6)$-map by Lemma 9.1. Let D be an arbitrary exposed cell. Without loss of generality, let the unique boundary arc of D be part of the boundary cycle labeled X. Since X is Dehn-reduced at most half of the boundary of D is contained in X, which means that at least half of the boundary of D is internal. Since the internal arcs are pieces (Lemma 9.1), and by assumption each piece is strictly less than $\frac{1}{8}$ of the boundary, the internal degree must be greater than 4. Thus every exposed cell D in Δ has $i(D) > 4$, and Lemma 7.3 can be applied to complete the proof. The lower bound on the length the boundary arcs follows from the observation that both rungs if they exist are strictly less than $\frac{1}{8}$ of the boundary and a side is at most $\frac{1}{2}$, so the other side must contain more than $\frac{1}{4}$ of the boundary. □

Lemma 9.4 implies in particular that there are cyclic conjugates X_1 of X and Y_1 of Y and a piece U of some relator R such that $UX_1 = Y_1U$. Moreover, more than $\frac{1}{4}$ of the length of R is contained in both X_1 and Y_1, significantly restricting the number of relators which can take on this role. The above lemmas quickly lead to the solution to the word problem and the conjugacy problem for small cancellation groups using a procedure known as Dehn's algorithm. Since the traditional versions are well-known and the general versions will be covered in depth later, no further details will be given here. One application will be given here since it is not as familiar, and since it is also the basis of the construction used in Section 13 to create α-closures of words and cycles. The following lemma is a consequence of the details given in the proof of Lemma 9.4.

Lemma 9.5 *Let $\mathcal{R}$ be a set of standard relators satisfying $C'(\alpha)$ for some $\alpha \leq \frac{1}{8}$. If X is a cyclically Dehn-reduced word which is not equivalent to 1 in the group $G = \langle A|\mathcal{R}\rangle$, then there exists a finite $\mathcal{R}$-structure which contains every cyclically Dehn-reduced word conjugate to X in the group G read as a loop in its 1-skeleton, homotopic to the loop reading X.*

Proof: First find the necessarily finite set of all paths in the abstract loop X which read more than one-fourth of a relator R in $\mathcal{R}$, whose length is at most $|X|$. Next, simply attach all of the appropriate relators R, thought of here as labeled 2-cells, to these paths in the loop X so that the intersection of the loop X and ∂R contains the particular path in question. If the resulting structure is collapsed then a finite $\mathcal{R}$-structure is formed which satisfies the conclusion of the lemma.

To see that this is true, let Y be any non-trivial cyclically Dehn-reduced word conjugate to X in G, and let Δ be a reduced annular diagram proving that X and Y are conjugate. The structure of Δ is given by Lemma 9.4. The result follows from the observation that the cell category of Δ maps into the construction. Begin by sending the boundary cycle of Δ labeled X to the image of the loop X under the collapsing functor. Then extend this map to the cells of Δ. The structure of Δ described in Lemma 9.4 guarantees that all of the cells needed for the extension have already been added to the construction, and the collapsed nature of the construction guarantees that all identifications of boundary edges in Δ are also identified in the construction. Once Δ is mapped into the construction it is clear that the construction must contain a loop reading Y and the image of Δ provides the homotopy to the loop reading X. □

One thing to be aware of is that the $\mathcal{R}$-structure constructed in Corollary 9.5 may not be planar, although for 2-dimensional complexes at worst it is homeomorphic to a Möbius band. Example 1 illustrates this worst-case scenario.

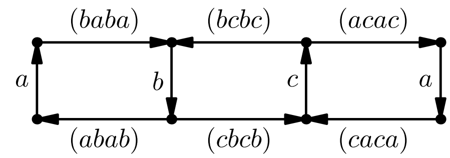

Figure 25: A Möbius strip which contains all reduced conjugates of U

Example 1 Consider the set of relators $\mathcal{R} = \{(ab)^5, (bc)^5, (ca)^5\}$ and the word $X = ababa(bcbc)^{-1}acac$. If the left and right edges labeled a in Figure 25 are identified with the proper orientation, then the resulting non-planar $\mathcal{R}$-structure is the construction described in Corollary 9.5 applied to X. In this instance it produces a Möbius band. In reading the figure, the reader should keep in mind that the parentheses indicate the order in which the label is written (left to right) may conflict with the order in which the letters occur in the arc (indicated by the direction of the arrow). For instance, the arc at the bottom of the leftmost rectangle is labeled $abab$, which means that the letters a, b, a, b are encountered when reading the arc from right to left. From an examination of the construction, it becomes clear that there are exactly 2 cycles, other than X itself, which are non-trivial, Dehn-reduced, and conjugate with X in the group $G = \langle a, b, c | \mathcal{R} \rangle$.

9.2 General Small Cancellation Axioms

Various axioms for a 2-dimensional generalized small cancellation theory have been proposed by Rips ([18]) and Ol'shanskii ([16]). These axiom systems are related but not identical to the system proposed below. The main innovation in the present case is the introduction of the geometry of the general relators and the resulting insight into the structure of the Cayley graph of the group. The description of the axioms of a general small cancellation presentation involves five constants. The constant α is analogous to the constant used in traditional small cancellation theory in that it measures the degree to which one relator can be contained in another without being subsumed. The other constants are more or less specific to general relators. The axioms are as follows:

Axiom 1 *There is a constant α such that every general relator $R \in \mathcal{R}$ is α-closed with respect to $\mathcal{R}$. In particular, if U is a word readable in general relators R and S via $\mathcal{R}$-functors f and g respectively, and $d_R(U) \geq \alpha$, then, since S is α-closed, there exists a unique $\mathcal{R}$-functor $h : R \to S$ such that $hf = g$.*

Axiom 2 *There is a constant β such that whenever a word U is readable in general relators $R, S \in \mathcal{R}$ by $\mathcal{R}$-functors f and g respectively, and either rank$(R) <$ rank(S) or rank$(R) =$ rank(S) but there does not exist an $\mathcal{R}$-functor $h : R \to S$ with $hf = g$, then $d_S(U) < \beta$.*

Axiom 3 *There is a constant* γ *such that for all general relators* $R \in \mathcal{R}$, $\omega_R \leq \gamma|R|$.

Axiom 4 *There is a constant* δ *such that the length of a path* U *in a general relator* $R \in \mathcal{R}$ *in the relator metric* d_R *is within* δ *of its length in the normalized graph metric on the boundary of* R. *Specifically* $|d_R(U) - |U|_R| \leq \delta$.

Axiom 5 *There is a constant* ϵ *such that whenever* U *is the shortest possible path from a vertex in a general relator* $R \in \mathcal{R}$ *to a loop with non-zero winding number in* R, *the length of* U *in the relator metric is at most* ϵ. *That is,* $d_R(U) \leq \epsilon$.

Axiom 6 *If* $W = XUYU^{-1}$ *is a representative of a general relator* R *in which both instances of* U *are properly oriented with respect to* W, *and* $d_R(U) \geq \alpha$, *then there exists a word* V *such that the cycle* $XVYV^{-1}$ *is readable as a contractible loop in* ∂R *extending the reading of* X *given by* W. *The cycle thus bounds a connected and simply connected* $\mathcal{R}$*-diagram* Δ *with rank*$(\Delta) <$ *rank*(R).

Axiom 7 *The constants* α, β, γ, δ, *and* ϵ *satisfy the following constraints:* $\beta \leq \alpha$, *and* $\gamma, \delta, \epsilon < \alpha$, *and* $2\gamma + \delta \leq \alpha \leq \frac{1}{6}$.

A general small cancellation presentation $G = \langle A|\mathcal{R}\rangle$ is a measured presentation which satisfies the axioms listed above. A group which possesses a general small cancellation presentation is called a general small cancellation group.

9.3 Basic Consequences

As stated above, the constant α corresponds to the constant used in the traditional theory. The other four constants provide bounds on various aspects of general relators. The constant β bounds the length of a path in a general relator of lower rank when it is measured by the relator metric of a general relator of higher rank. The constant γ bounds the ratio of the width of the general relator to its length. The constant δ bounds the difference between the normalized graph metric and the relator metric on the boundary of a general relator. And the constant ϵ bounds the length of geodesics going across general relators as measured by their relator metrics.

The general small cancellation theory presented here claims to be an extension of the more traditional theory in the following sense. Every traditional small cancellation presentation which satisfies the condition $C'(\alpha)$ for some $\alpha \leq \frac{1}{6}$ is also a general small cancellation for the same value of α, with $\beta = \alpha$, and with $\gamma = \delta = \epsilon = 0$. The constants γ, δ, and ϵ are zero because traditional relators have width 0 and the metric used is the normalized graph metric. The comparison can also be made more explicit. The hypothesis $C'(\alpha)$ is formulated in terms of pieces. Axiom 1 states that all pieces measure less than α in the relator metric, which in this case is the same as the normalized graph metric on R. To see this, view the relators R and S as closed cones of height 2. Then by Lemma 2.1 the functor mentioned in the axiom must be an isomorphism. Thus, words which contain more than α of the boundary of a relator are readable in only one relator in $\mathcal{R}$. The fact that Axiom 2 is true when $\beta = \alpha$ is true more generally, and this is the content of Lemma 9.6. In the traditional theory, the relators are cyclically reduced, which makes the loops labeled by these cycles deterministic with width 0. Thus Axiom 3 is satisfied. By using the normalized graph metric, the lefthand side of Axiom 4 is always 0, and thus the inequality is true. Again, since the width of the relators is 0, Axiom 5 is immediate. And by Lemma 4.2, Axiom 6 is predicated on a situation which never occurs in the traditional theory. The constraints in Axiom 7 are immediate.

In the general case, Axiom 1 is a rough way of saying that either two general relators have a small overlap or one of them is contained in the boundary of the other, Axiom 3 guarantees the 'niceness' of the relators, and Axiom 4 guarantees the 'niceness' of the relator metrics. Axiom 6 provides a reduction in a situation which does not occur in the 2-dimensional relators. The other two axioms are derivable from those already listed. In Lemma 9.6 it is shown that $\beta = \alpha$ satisfies Axiom 2, and in Lemma 9.7 it is shown that $\epsilon = \gamma + \delta$ satisfies Axiom 5.

Lemma 9.6 *If $G = \langle A|\mathcal{R}\rangle$ is a measured presentation satisfying Axiom 1, then Axiom 2 is also satisfied for the constant $\beta = \alpha$.*

Proof: Suppose that $d_S(U) \geq \alpha$ in one of the situations described by Axiom 2. By Axiom 1 there is a functor $h' : S \to R$ with $h'g = f$. But then Lemma 2.1 combined with the assumptions of Axiom 2 shows that R and S have the same rank and that h' is an isomorphism. If $h : R \to S$ is the inverse of the isomorphism h', then $hf = g$, which is contrary to the initial assumption. $\square$

Lemma 9.7 *If $G = \langle A|\mathcal{R}\rangle$ is a measured presentation satisfying Axiom 3 and Axiom 4, then Axiom 5 is also satisfied for the constant $\epsilon = \gamma + \delta$.*

Proof: If U is the geodesic described in Axiom 5, then by the definition of the width of a general relator and Axiom 3 respectively, $|U| < \omega_R \leq \gamma|R|$. Thus, by Axiom 4, $d_R(U) < \gamma + \delta$. $\square$

The reason for the redundancy is to introduce all five of the constants at the same time and in the same place. The constants are retained because certain constructions, such as the α-closure of an $\mathcal{R}$-category, are only possible if these constants satisfy stricter inequalities than the easy bounds mentioned above. The only constraints which have been included in Axiom 7 are those needed to reproduce the most basic results of small cancellation theory. Other useful restrictions will be added as needed in specific situations. The remaining results in this section will show that general small cancellation presentations satisfy the special conditions needed to apply some of the lemmas developed earlier. Specifically, the general relators in a general small cancellation presentation are thin and collapsed, all long arcs are oriented, and the set $\mathcal{R}$ has distinct representatives.

Lemma 9.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If R is a general relator in $\mathcal{R}_k$, then R is thin and collapsed, and ∂R is α-closed with respect to $\mathcal{R}(k-1)$.*

Proof: By Axiom 7, $\gamma \leq \frac{1}{12} < \frac{1}{4}$. Thus $4\omega_R < |R|$ and R is by definition thin. Next, notice that the 1-skeleton of R is deterministic by the definition of a general relator. Let S be a general relator in R. By Lemma 5.1, the boundary of S includes a representative cycle U, and by one of the properties of relator metrics the word U has $d_S(U) \geq 1$. Thus by Axiom 1 there is a unique functor from S to R which includes this reading of the word U in R. In particular, no two distinct closed cones in R have identical boundaries. Thus R is collapsed. Finally, the fact that ∂R is α-closed with respect to $\mathcal{R}(k-1)$ is an immediate consequence of Axiom 1 and the fact that $\mathcal{R}$-functors from lower-ranked general relators into R must have their image in ∂R. □

Lemma 9.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If U is a path in a general relator R, and U is long in R, then U is also oriented with respect to R. More specifically, whenever $d_R(U) \geq \alpha$, $|U| \geq 2\omega_R$.*

Proof: Since by definition $d_R(U)$ is a function on the endpoints of the path U lifted to R^∞, without loss of generality U can be assumed to be a geodesic in R^∞. Then by Axiom 4, Axiom 7, and Axiom 3 respectively, $|U| \geq (\alpha - \delta)|R| \geq 2\gamma|R| \geq 2\omega_R$. □

Lemma 9.10 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation, and a word W is readable in relators R_1 and R_2 with $d_{R_1}(W) \geq \beta$ and $d_{R_2}(W) \geq \beta$, then R_1 and R_2 are isomorphic as $\mathcal{R}$-categories, and the isomorphism can be chosen so that the paths reading W are identified. In particular, for each word W there is at most one general relator R in which $d_R(W) \geq \beta$.*

Proof: Without loss of generality, assume that $\text{rank}(R_1) \leq \text{rank}(R_2)$. Then by Axiom 2, $\text{rank}(R_1) = \text{rank}(R_2)$ and there is a functor from R_1 to R_2 sending the reading of W in R_1 to the reading of W in R_2. By Lemma 2.1, this functor is an isomorphism. □

Lemma 9.11 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation, and a cycle W is readable as a loop with non-zero winding number in relators R_1 and R_2, then R_1 and R_2 are isomorphic as $\mathcal{R}$-categories, and the isomorphism can be chosen so that the loops reading W are identified. In particular, the set $\mathcal{R}$ has distinct representatives.*

Proof: By the definition of relator metrics, $d_{R_1}(W) \geq 1$ and $d_{R_2}(W) \geq 1$. Thus Lemma 9.10 can be used to complete the proof. □

Lemma 9.12 *Let $G = \langle A|\mathcal{R}\rangle$ be a graded, measured presentation, and let $\alpha, \beta, \gamma, \delta,$ and ϵ be fixed constants. The presentation of G is a general small cancellation presentation iff for all $k \in$ integers $G(k) = \langle A|\mathcal{R}(k)\rangle$ is a general* *small cancellation presentation. In particular, if sets of general relators $\mathcal{R}_k$ are created so that $G(k) = \langle A|\mathcal{R}(k)\rangle$ is a general small cancellation presentation for the same set of constants then $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation.*

Proof: Notice that a counterexample to either Axiom 1 or Axiom 2 involves only two general relators, that a counterexample to Axiom 3, Axiom 4, or Axiom 5 involves only one general relator, that the ranks of the general relators needed to produce a counterexample to Axiom 6 are bounded by the rank of R, and that a counterexample to Axiom 7 does not involve any of the general relators. Thus any counterexample proving that $G = \langle A|\mathcal{R}\rangle$ is not a general small cancellation presentation can also be used to show that the presentation of $G(k)$ is not a general small cancellation presentation for some positive integer k. □

10 Dehn's Algorithm

This section begins with a proof that certain $\mathcal{R}$-diagrams can be altered to produce equivalent $\mathcal{R}$-diagrams over $C(p)$-maps. Once the proof is completed, it will be possible to apply the known results on $C(p)$-maps to diagrams over general small cancellation presentations. These results are then used to prove that the word and conjugacy problems are decidable using a version of Dehn's algorithm. Another use of Dehn's algorithm is to show that the Cayley category of a general small cancellation presentation is proper. This result is given at the end of the section.

10.1 Dehn-reduced Words and Cycles

Recall that a word W is Dehn-reduced with respect to $\mathcal{R}$ if W is reduced in the free group and there do not exist a word U, a word V, and a general relator $R \in \mathcal{R}$ such that U is a subword of W, UV^{-1} is a representative of R and $|U| > |V|$. If such a situation does arise then it is clear that replacing U with V in W produces a strictly shorter word which is equivalent to W in the group $G = \langle A|\mathcal{R}\rangle$.

Lemma 10.1 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation and $\mathcal{R}'$ is a subset of $\mathcal{R}$, then it is decidable whether a word W is Dehn-reduced with respect to $\mathcal{R}'$.*

Proof: If W is not reduced in the free group, then this fact is easy to discover in finite time. If there exists a word U, a word V and a general relator $R \in \mathcal{R}'$ showing that W is not Dehn-reduced, then this can also be discovered in finite time, since $2|W| \geq 2|U| > |U|+|V| \geq |R|$, showing that there is a bound on the length of the general relators which need to be considered. Since by Lemma 9.11 general relators have distinct representatives, and the alphabet A is finite, this bound of the length of the general relator R guarantees that only a finite number of general relators exist which can satisfy this condition. Finally, since the list of general relators of interest is finite and is derived solely from the length of W, it is easy to exhaustively check whether there exist words U and V in each general relator in the list which together satisfy the requirements. □

Lemma 10.2 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. For every word W there exists an effectively constructible word of shorter length which is Dehn-reduced with respect to $\mathcal{R}'$ and equivalent to W in G.*

Proof: First of all, by Lemma 10.1 it is decidable whether a given word W is Dehn-reduced with respect to $\mathcal{R}'$ or not. If W is not Dehn-reduced then it can be reduced either by removing the subword of the form aa^{-1} or by replacing the subword U with the word V to create a strictly shorter word which is equivalent to W in G. Since these reductions shorten the length of the word, only a finite number of such reductions can occur before the word is Dehn-reduced. □

Recall that a word W is called μ-free with respect to a set of general relators $\mathcal{R}$ if W is reduced in the free group and there do not exist a subword U of W and a general relator $R \in \mathcal{R}$ such that $d_R(U) > \mu$.

Lemma 10.3 *If U is a word readable in a general relator R and $d_R(U) > \frac{1}{2} + 2\gamma + \delta$ then there exist a subword U' in U and a word V such that $U'V^{-1}$ is a representative of R and $|U'| > |V|$. In particular, words which are Dehn-reduced with respect to a set of general relators $\mathcal{R}$ must also be $(\frac{1}{2} + 2\gamma + \delta)$-free and thus $(\frac{1}{2} + \alpha)$-free with respect to $\mathcal{R}$.*

Proof: Since $d_R(U) \geq \alpha$, the path U is oriented, by Lemma 9.9. Lift U to a path in R^∞ and call the initial vertex v_0 and the final vertex u. Pick v_1 as the unique vertex such that the path from v_0 to v_1 is a representative and v_1 and u are on the same side of v_0, meaning that both vertices are in the same connected component of R^∞ when $\text{Ball}(v_0, \omega_R)$ is removed. The proof is divided into two cases.

Case 1: If there is a vertex u' of the lifted path U contained in $\text{Ball}(v_1, 2\omega_R)$ then define V to be a geodesic from v_1 to u' and define U' as the initial segment of U which ends at u'. By definition $|V| < 2\omega_R \leq 2\gamma|R|$ by Axiom 3. Also, by construction $U'V^{-1}$ is a representative of R, so that $|U'| + |V| \geq |R|$ by the definition of $|R|$, and $|U'| > (1 - 2\gamma)|R|$. Finally, $1 - 2\gamma > 2\gamma$ by Axiom 7, so $|U'| > (1 - 2\gamma)|R| > 2\gamma|R| > |V|$.

Case 2: If there is not such vertex, then the balls of radius ω_R centered at v_0, v_1, and u are disjoint. More specifically, since these balls disconnect R^∞, the ball centered at u must be between those centered at v_0 and v_1. Define V to be a geodesic path from v_1 to u, and define $U' = U$. The word UV^{-1} is clearly a representative of R. It only remains to show that $|U| > |V|$. If, without loss of generality, it is assumed that U is also geodesic, then since UV^{-1} is a geodesic 2-gon which represents R, Lemma 5.21 shows that $|U| + |V| < |R| + 4\omega_R \leq (1 + 4\gamma)|R|$. Next, since $d_R(U) > \frac{1}{2} + 2\gamma + \delta$, Axiom 4 implies that $|U| > (\frac{1}{2} + 2\gamma)|R|$. Combining these inequalities shows that $(\frac{1}{2} + 2\gamma)|R| > |V|$, and thus $|U| > |V|$. Finally, the fact that Dehn-reduced words are also $(\frac{1}{2} + \alpha)$-free follows from Axiom 7. □

Lemma 10.4 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let B be an $\mathcal{R}$-category, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If B is μ-closed with respect to $\mathcal{R}'$ for some $\mu \geq \frac{1}{2} + 2\gamma + \delta$, then given any path X in B, there is a path Y between the same two vertices which is μ-free with respect to $\mathcal{R}'$, homotopic to X relative to its endpoints, and equivalent to X in $G' = \langle A|\mathcal{R}'\rangle$. In addition $|X| \geq |Y|$.*

Proof: If X is not itself μ-free with respect to $\mathcal{R}'$ then there exist a word U and a general relator $R \in \mathcal{R}'$ such that U is a subword of X and $d_R(U) > \mu$. By Lemma 10.3 there are words U' and V such that U' is a subword of U and thus of X, and the replacement of U' with V creates a shorter word between the same two vertices, homotopic to the original path X. Since the length of the word is strictly descreasing, such replacements can occur only a finite number of times before they stop at a word which satisfies the conditions of the lemma. □

Recall that a cycle W is called Dehn-reduced iff all of the cyclic conjugates of the word W are Dehn-reduced, and similarly, that the cycle W is called μ-free iff all of the cyclic conjugates of the word W are μ-free. Notice that even if a cycle W is Dehn-reduced this does not imply that the cycle W^2 is Dehn-reduced since there may exist a subword U of length slightly longer than W which leads to a reduction. A similar comment applies to μ-free cycles.

Lemma 10.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. For every cycle W there exists an effectively constructible cycle of shorter length which is Dehn-reduced with respect to $\mathcal{R}'$ and conjugate to W in G. Moreover, the resulting Dehn-reduced cycle is $(\frac{1}{2} + 2\gamma + \delta)$-free with respect to $\mathcal{R}'$.*

Proof: By Lemma 10.2 the word W can be Dehn-reduced with respect to $\mathcal{R}'$. If one of the cyclic conjugates of the result is not Dehn-reduced then Dehn-reduce the cyclic conjugate and repeat this process. Since the reductions shorten the length of the cycle and since cyclic conjugation leaves the length unchanged, the process must stop after a finite number of steps. The result is clearly conjugate to W in G, and Dehn-reduced with respect to $\mathcal{R}'$. Finally, since all of the cyclic conjugates are Dehn-reduced with respect to $\mathcal{R}'$, all of the cyclic conjugates are also $(\frac{1}{2} + 2\gamma + \delta)$-free with respect to $\mathcal{R}'$ by Lemma 10.3. Thus the resulting cycle is $(\frac{1}{2} + 2\gamma + \delta)$-free with respect to $\mathcal{R}'$. □

Lemma 10.6 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let B be an $\mathcal{R}$-category, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If B is μ-closed with respect to $\mathcal{R}'$ for some $\mu \geq \frac{1}{2} + \delta + 2\gamma$, then given any loop X in B, there is a loop Y which is μ-free with respect to $\mathcal{R}'$, homotopic to X, and conjugate to X in $G' = \langle A|\mathcal{R}'\rangle$. In addition $|X| \geq |Y|$.*

Proof: The result follows by merely applying Lemma 10.4 repeatedly to the cyclic conjugates of the loop obtained so far. Since the length of the loop is strictly decreasing, the process stops after a finite number of steps at a loop which satisfies the conditions of the lemma. □

Although it is immediate that a word W is Dehn-reduced whenever the cycle W is Dehn-reduced, the reverse is not always true. The lemmas below illustrates some conditions under which it is possible to conclude that a cycle is Dehn-reduced from the fact that a word is Dehn-reduced.

Lemma 10.7 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If the word W^{n+1} is Dehn-reduced with respect to $\mathcal{R}'$, then the cycle W^n is Dehn-reduced with respect to $\mathcal{R}'$.*

Proof: The result follows from the observation that all words readable in the cycle of W^n with a length less than $|W^n|$ are also readable in the word W^{n+1}. Thus, if the cycle W^n is not Dehn-reduced, then the word W^{n+1} is also not Dehn-reduced. □

If a path is read in a collapsed $\mathcal{R}$-category and there is an automorphism of this $\mathcal{R}$-category which sends the start vertex of the path to its end vertex, then this automorphism is called the automorphism represented by the path. Notice that the definition of a collapsed $\mathcal{R}$-category guarantees that this automorphism is unique whenever it exists, thereby justifying the use of the definite article.

Lemma 10.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. In addition, let W be a word, and let R be a general relator which does not contain any power of W as a loop. If U is a W-periodic word which is readable in R with $d_R(U) > 2\beta$, then $|U| < 2|W|$.*

Proof: Assume $|U| \geq 2|W|$, and without loss of generality, assume that W is an initial segment of U, so that $U = WV$ and $|V| \geq |W|$. If $d_R(V) \geq \beta$, then by Axiom 2, the readings of V as initial and final segments of U and thus in R differ by an automorphism of R. Since W is a path in R between the initial vertices of these readings, W represents this automorphism of R. As a consequence, W^n is a loop in R for some power of W, contradicting the assumptions stated in the lemma. Thus $d_R(V)$ must be strictly less than β. Similarly, if $d_R(W) \geq \beta$, then the readings of W as initial segments of U and of V lead to an automorphism of R which can be represented by W, which again leads to power of W being a loop in R. Thus $d_R(W)$ must also be strictly less than β. But in this case, by the properties of relator metrics, $d_R(U) \leq d_R(W) + d_R(V) \leq 2\beta$. This final contraction shows that the initial assumption that $|U| \geq 2|W|$ must itself be false. □

Lemma 10.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If W is a word such that no power of W is ever readable as a loop with nontrivial winding number in any of the general relators in $\mathcal{R}'$, and either the word W^3 or the cycle W^2 is Dehn-reduced with respect to $\mathcal{R}'$, then the cycle W^i is Dehn-reduced with respect to $\mathcal{R}'$ for all $i > 1$.*

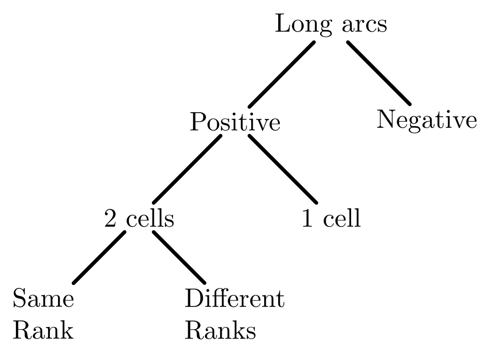

Figure 26: The four varieties of long arcs

Proof: Suppose that the cycle of W^i is not Dehn-reduced with respect to $\mathcal{R}'$. Then there exists a W-periodic word U and a general relator $R \in \mathcal{R}'$ with a path U readable in R such that $|U| \geq \frac{1}{2}|R|$. By Axiom 4, $d_R(U) \geq \frac{1}{2} - \delta > 2\beta$ by Axiom 7. Thus by Lemma 10.8, the length of U is less than twice that of W, so that U is readable in both the word W^3 and the cycle W^2. The path U and the relator R prove that the word W^3 and the cycle W^2 are not Dehn-reduced, contradiction. $\square$

10.2 Long Arcs and Ladders

The first step in applying Dehn's algorithm to diagrams over general small cancellation presentations is to show that the key diagrams can be modified so that the underlying maps are $C(6)$-maps. This is shown in Lemma 10.16. Once this is established, the results quoted from [9] can be applied. The proof proceeds by showing that each of four kinds of long internal arcs in a diagram can be removed in such a way that the type of the diagram is always strictly decreasing. The four kinds of long internal arcs are long negative arcs, long positive arcs between a single cell and itself, long positive arcs between 2 distinct cells of different ranks, and long positive arcs between 2 distinct cells of the same rank. The relations between these kinds of arcs are depicted in Figure 26. In each of the four cases it will be shown that the diagram can be altered to produce an equivalent diagram of lower type. Since diagram types satisfy the descending chain condition, it is clear that every sequence of reductions must terminate in a finite number of steps in a diagram in which none of the above situations occur.

Lemma 10.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If $W = UV$ is a representative of a general relator $R \in \mathcal{R}$ such that $d_R(U) \geq \alpha$ and the orientation of U is the opposite of that of W, then there exist paths V_1, V_2, and U' such that $V = V_1V_2$, the loop $U'V_2$ is a representative of R, and $|U'| < |U|$.*

Proof: Figure 27 schematically illustrates the argument below. First lift W

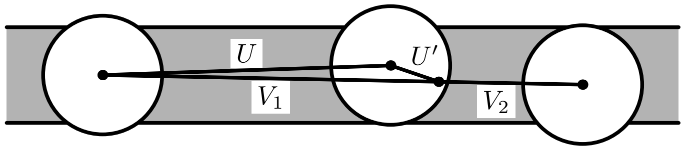

Figure 27: Shortening long negative arcs

from R to R^∞ and let v_0 and v_1 be its initial and terminal vertices, respectively. Furthermore, let u be the terminal vertex of the path U in W. Next, draw balls of radius ω_R around v_0, v_1, and u. Since U and W have opposite orientations, the ball centered at v_0 must be strictly between the balls centered at v_1 and u. In particular, since the path V starts at u and ends at v_1 there must exist a vertex u' in V which lies strictly within ω_R units of v_0. Use the vertex u' to define paths V_1 and V_2 such that $V = V_1V_2$, and define U' to be a geodesic path from v_0 to u'. Clearly, $U'V_2$ is a representative of R, and by Lemma 9.9, $|U| \geq 2\omega_R$ while by construction $|U'| < \omega_R$. □

Lemma 10.11 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let Δ be an $\mathcal{R}$-diagram whose boundary cycles are linked. If D is a 2-cell in Δ and U is an internal arc which is both long and negative in D, then there is an $\mathcal{R}$-diagram Δ' which is equivalent to Δ but with Type$(\Delta) >$ Type(Δ').*

Proof: Let $W = UV$ be the boundary cycle of D. By Lemma 10.10 there are words V_1, V_2, and U' such that $V = V_1V_2$, the loop $U'V_2$ is a representative of R, and $|U'| < |U|$. The proof divides into three cases.

Case 1: If U' is non-empty, then simply remove the edges and vertices in the interior of the arc U along with the interior of the cell or cells on either side of U. Then add a new arc corresponding to the word U' which starts at the initial vertex of U and ends at the terminal vertex of V_1. A new cell or two new cells can then be attached to the new holes. By construction $U'V_2$ is a representative of R. Since the cell on the other side of the arc U' originally contained U in its boundary, by Axiom 1 the relator S which the cell represents contains a complete copy of R. Since $U'V_1^{-1}U^{-1}$ is a loop of winding number 0 in R, the same is true as a loop in S. Thus the replacement of U with $U'(V_1)^{-1}$ creates a new boundary cycle which is also a representative of S. Notice that this argument remains unchanged even in the case where the arc U borders the same cell in two distinct ways. Since the new $\mathcal{R}$-diagram contains exactly the same number of cells in each rank but fewer edges, the type of the new $\mathcal{R}$-diagram is strictly less than that of Δ.

Case 2: If U' is empty and the initial vertex of U is distinct from the terminal vertex of V_1 then the construction proceeds as in Case 1, except that once the interior of U and the interior of the cell(s) on either side of U have been removed, the map can be stretched so that these vertices are identified in the space where the cell(s) used to be. The rest of the proof is as before.

Case 3: If U' is empty, and the initial vertex v of U is identical with the terminal vertex of V_1, then first attach phantom cells to Δ to form an $\mathcal{R}^+$-sphere, and then remove the interior of the arc U along with the interior of the cells on either side of U. The fact that these vertices are identical forces the conclusion that the cells on either side of U are distinct. Let D' be the other cell, and let UV' be its boundary. The result of the removals is a connected and simply connected $\mathcal{R}^+$-diagram with v as a cut point. This $\mathcal{R}^+$-diagram can be viewed as the one-point product of two connected and simply connected $\mathcal{R}^+$-diagrams with boundary cycles $U'V_2 = V_2$ and $U'V_1^{-1}V' = V_1^{-1}V'$. Each of these will be considered separately. Since $U'V_2$ is a representative of R, a cell can be added to form an $\mathcal{R}^+$-sphere. By the same argument as in Case 1, the cycle $U'V_1^{-1}V'$ is a representative of the same general relator as the original boundary of D'. Thus a cell can be added to this $\mathcal{R}^+$-diagram so that it becomes an $\mathcal{R}^+$-sphere as well. Consider the location of the phantom cells. Since the boundary cycles are linked they must all be in one or the other of the $\mathcal{R}^+$-spheres. From this sphere, remove the phantom cells and the result will be an $\mathcal{R}$-diagram of strictly lower type, which is equivalent to Δ. □

Lemma 10.12 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let Δ be an $\mathcal{R}$-diagram whose boundary cycles are linked. If D is a labeled 2-cell in Δ, with boundary cycle W, and U is an internal arc of D in two distinct ways which is long in D, then there is an $\mathcal{R}$-diagram Δ' which is equivalent to Δ but with Type$(\Delta) >$ Type(Δ').*

Proof: If either instance of U in the boundary ∂D is negatively oriented then Lemma 10.11 can be applied to complete the proof, so assume that both instances are positively oriented with respect to ∂D. The cycle read by the boundary of D has the form $XUYU^{-1}$ for some non-empty words X and Y. The words X and Y must be non-empty since both instances of U are positively oriented. If, say, X was empty then since R is deterministic, the two paths labeled U would have the same image in R, showing that they have opposite orientations with respect to ∂D.

If the interior of D and the interior of the arc U are removed, the result is an annular hole with boundary cycles labeled X and Y. Let u be the vertex in the boundary cycle from which X can be read, and let u' be the vertex in the other boundary cycle from which Y can be read. The $\mathcal{R}$-diagram of strictly lower rank with boundary cycle $XVYV^{-1}$, whose existence is guaranteed by Axiom 6, can be placed in the annular hole so that the initial vertex of the first instance of V is attached to u, and the terminal vertex of this instance of V is attached to u'. See Figure 28 for an illustration. The boundary of the hole is a Dyck word, and since all of the cells added have strictly lower rank than the one which was removed the result is an $\mathcal{R}$-diagram of strictly lower type.

At this point the diagram contains all of the original boundary cycles plus a region bounded by a Dyck word. If phantom cells are added for each of the boundary cycles of Δ then the result is a connected and simply connected $\mathcal{R}^+$ diagram whose boundary is a Dyck word. By Lemma 7.6 the boundary

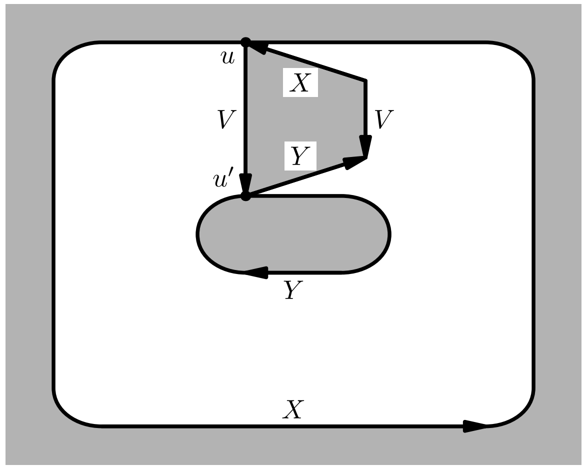

Figure 28: An illustration of Lemma 10.12

edges can be identified and since the boundary cycles of Δ are linked, all of the phantom cells end up in the same $\mathcal{R}$-sphere of the one-point product. Ignoring the other items in the one-point product and removing the phantom cells leaves an $\mathcal{R}$-diagram Δ' which is equivalent to Δ and of strictly lower type. □

Lemma 10.13 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let Δ be an $\mathcal{R}$-diagram whose boundary cycles are linked. If D_1 and D_2 are distinct 2-cells in Δ with the same rank, U is an internal arc which borders both D_1 and D_2, and U is long in D_1, then D_1 and D_2 form a cancellable pair. In consequence there is an $\mathcal{R}$-diagram Δ' which is equivalent to Δ but with Type$(\Delta) >$ Type(Δ').*

Proof: Let UV_1 be the representative of D_1 and UV_2 be the representative of D_2. Since U is long in D_1 and contained in the boundary of D_2, by Axiom 1 the entire relator R_1 maps into R_2. Since by assumption R_1 and R_2 have the same rank, by Lemma 2.1 R_1 and R_2 are isomorphic and will hereafter be referred to as R, and the functor from R_1 to R_2 can then be described as an automorphism of R. Since the arc U is long in D_1, U is long in D_2 as well. By Lemma 10.11 both instances of U can be assumed to be positively oriented. The cycle $V_1(V_2)^{-1}$ is then a loop which is readable in R and its winding number is 0. Thus the 2-cells D_1 and D_2 form a cancellable pair. Since by Lemma 9.11 $\mathcal{R}$ has distinct representatives, Lemma 7.15 yields the desired result. □

Lemma 10.14 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let Δ be an $\mathcal{R}$-diagram whose boundary cycles are linked. If D_1 and D_2 are distinct 2-cells in Δ of different ranks, and U is an internal arc which borders both D_1 and D_2 with U long in D_1, then there exists an $\mathcal{R}$-diagram Δ' which is equivalent to Δ but with Type$(\Delta) >$ Type(Δ').*

Proof: Let UV_1 be the representative of D_1 and UV_2 be the representative of D_2. Since U is long in D_1, and contained in the boundary of D_2, by Axiom 1 the entire relator R_1 maps into R_2. By Lemma 2.1 this means that the rank of R_1 is at most that of R_2. Since the ranks are assumed to be distinct, R_1 must be of strictly lower rank, and the functor from R_1 to R_2 sends R_1 into the boundary of R_2. Thus the loop read by the representative of R_1 is a loop with winding number 0 in R_2, and $V_1(V_2)^{-1}$ is also a representative of R_2. Thus the diagram Δ' obtained by simply erasing the interior of the arc U is also an $\mathcal{R}$-diagram and one of strictly lower type. □

Lemma 10.15 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If Δ is an $\mathcal{R}$-diagram whose boundary cycles are linked, then there exists an effectively constructible $\mathcal{R}$-diagram Δ' which is equivalent to Δ, contains no long internal arcs, and for which Type$(\Delta) \geq$ Type(Δ').*

Proof: If Δ has no long internal arcs then there is nothing to prove, so assume that such an arc exists. Depending on which kind of long internal arc it is, one of the previous four lemmas can be applied to produce an equivalent diagram of strictly lower type. If the new diagram is also not reduced then the same process can be repeated. Since diagram types satisfy the descending chain condition, the process must terminate after a finite number of steps at an $\mathcal{R}$-diagram equivalent to Δ which contains no long internal arcs. □

Lemma 10.16 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{p}$. If Δ is an $\mathcal{R}$-diagram whose boundary cycles are linked and cyclically reduced, then there is an $\mathcal{R}$-diagram Δ' which is equivalent to Δ and such that the underlying map of Δ' is a $C(p{+}1)$-map. In particular, since all general small cancellation presentations satisfy $\alpha \leq \frac{1}{6}$, the $\mathcal{R}$-diagram Δ' is over a $C(7)$-map.*

Proof: The result follows immediately from Lemma 10.15 and Lemma 7.4. □

The next several lemmas apply the results quoted about $C(6)$-maps to diagrams over general small cancellation presentations.

Lemma 10.17 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If W is a cyclically reduced word which is equivalent to 1 in G but which is not a representative of any general relator in $\mathcal{R}$, and Δ is a connected and simply connected $\mathcal{R}$-diagram with boundary cycle W over a $C(7)$-map, then Δ contains at least two exposed cells with internal degree at most 3. Moreover, the labels on the boundary arcs of these exposed cells show that the cycle W contains at least two disjoint subwords U for which there is a general relator R such that $d_R(U) > 1 - 3\alpha$.*

Proof: The combination of Lemma 10.16 and Lemma 7.2 produces the two 2-cells, while Lemma 5.25 furnishes the measurement on the boundary arc. □

Lemma 10.18 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If a non-trivial word W is equivalent to 1 in G, then either W is not reduced in the free group, or there is a subword U in W and a general relator $R \in \mathcal{R}$ such that $d_R(U) > 1 - 3\alpha$. As a consequence, if $3\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then the only Dehn-reduced word which is equivalent to 1 in G is the identity itself. In particular, these results are true when $\alpha \leq \frac{1}{6}$ and $\gamma = \delta = 0$ or whenever $\alpha \leq \frac{1}{8}$.*

Proof: Assume that W is equivalent to 1 in G and reduced in the free group. The only way that cycle W could not also be cyclically reduced in the free group is if the first and the last letters of W are inverses of each other. Removing these letters and repeating the process if necessary shows that there are words V and W' such that $W = VW'V^{-1}$ and W' is cyclically reduced in the free group. If W' is the empty word, then W itself is not reduced, contradicting the initial assumption. Thus W' is non-trivial and by Lemma 10.17 there is a connected and simply connected $\mathcal{R}$-diagram Δ over a $C(7)$-map with boundary cycle W' which contains at least two exposed cells with internal degree at most 3. The $\mathcal{R}$-diagram Δ can be altered to form a new $\mathcal{R}$-diagram Δ' with boundary W by adding a path reading V attached to Δ at the appropriate vertex. This addition potentially occurs in the middle of the unique boundary arc of one of the exposed cells, but this leaves at least one exposed cell D with its boundary arc uninterrupted. The unique boundary arc U of D and the general relator R which it represents complete the proof. □

Lemma 10.19 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let X and Y be cycles which are conjugate but not equivalent to the identity in G. If X is 4α-complement-free and Y is 3α-complement-free, then there exists an annular $\mathcal{R}$-diagram Δ with boundary cycles X and Y^{-1} in which every 2-cell D has exactly two boundary arcs, one in the loop X and the other in Y, and with internal degree at most 2. Moreover, both boundary arcs are long relative to D. In particular, if $4\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then any two Dehn-reduced cycles in G are the boundary cycles of an annular $\mathcal{R}$-diagram such as Δ. This is true, more specifically, when $\alpha \leq \frac{1}{8}$ and $\gamma = \delta = 0$, and whenever $\alpha \leq \frac{1}{10}$.*

Proof: Since X and Y are conjugate and non-trivial, by Lemma 7.10 there is an annular $\mathcal{R}$-diagram Δ with oriented boundary cycles X and the inverse of Y. By Lemma 10.16, it can be assumed that Δ is over a $C(7)$-map. An application of Lemma 7.3 shows that Δ has the desired structure. The fact that both boundary arcs are long in D follows from Lemma 5.25. Finally, the inequality listed in the statement of the lemma was chosen to guarantee that Dehn-reduced cycles, which are $(\frac{1}{2} + 2\gamma + \delta)$-free by Lemma 10.5, would also be $(1 - 4\alpha)$-free, and thus 4α-complement-free. The specific examples clearly satisfy this inequality because of the restriction on the choices of the constants. □

Lemma 10.20 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let X and Y be words which are equivalent but not equivalent to the identity in G. If X is 4α-complement-free and Y is 3α-complement-free, then there exists* *a connected and simply connected $\mathcal{R}$-diagram Δ with boundary cycle XY^{-1} in which every 2-cell D has either exactly two boundary arcs, one in the loop X and the other in Y, and with internal degree at most 2, or one boundary arc which contains an endpoint of X in its interior, and with internal degree at most 1. Moreover, both boundary arcs are long relative to D. In particular, if $4\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then any two equivalent Dehn-reduced words in G form a boundary cycle of an $\mathcal{R}$-diagram such as Δ. This is true, more specifically, when $\alpha \leq \frac{1}{8}$ and $\gamma = \delta = 0$, and whenever $\alpha \leq \frac{1}{10}$.*

Proof: If X and Y have the same first letter or the same last letter, then remove the common letter from both words, and repeat this process until it no longer occurs. The removed letters will be added in afterwards. If either the new X or the new Y is equivalent to 1 in G then they both are, and since they are Dehn-reduced, Lemma 10.18 shows that they are empty words. If letters have been removed which need to be added back in, then the original X and Y were equal and the loop X satisfies the conditions of the lemma. On the other hand if the new X and the new Y are not equivalent to 1 in G, then XY^{-1} is the boundary of a connected and simply connected $\mathcal{R}$-diagram Δ by Lemma 7.9. By Lemma 10.16, it can be assumed that Δ is over a $C(7)$-map.

Since X and Y are 4α-complement-free and 3α-complement-free words, every exposed cell D of Δ either contains an edge of X and has $i(D) > 4$, or it contains an edge of Y and has $i(D) > 3$, or it contains one of the endpoints of X in the interior of its unique boundary arc. Since by Lemma 7.9 the endpoints of the path X are distinct, the map can be stretched and these points identified to form an annular $\mathcal{R}$-diagram Δ' with boundary cycles X and Y^{-1}. If letters have been removed from the original words then a path reading these removed letters can be added instead between the endpoints of X. The map remains $C(7)$ under this procedure, and all exposed cells D now satisfy the conditions necessary to apply Lemma 7.3. This shows that the original $\mathcal{R}$-diagram Δ has the desired structure. The fact that all boundary arcs are long in X and in Y follows from Lemma 5.25. Finally, the inequality listed in the statement of the lemma was chosen to guarantee that Dehn-reduced words, which are $(\frac{1}{2} + 2\gamma + \delta)$-free by Lemma 10.2, would also be $(1 - 4\alpha)$-free, and thus 4α-complement-free. The specific examples clearly satisfy this inequality because of the restriction on the choices of the constants.□

10.3 The Word and Conjugacy Problems

The following lemma shows that under mild additional restrictions on the choice of constants the word problem for general small cancellation groups is always decidable.

Lemma 10.21 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $3\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then the group G has a decidable word problem. In particular, these results are true when $\alpha \leq \frac{1}{6}$ and $\gamma = \delta = 0$ or whenever $\alpha \leq \frac{1}{8}$.*

Proof: Since words X and Y are equivalent in G iff the word $W = XY^{-1}$ is equivalent to 1, it is sufficient to show that it is decidable whether words are equal to 1. By Lemma 10.2 given any word W there is a word Z which is effectively constructible, Dehn-reduced with respect to $\mathcal{R}$, and equivalent to W in the group G. If Z is the empty word, then by construction W is equivalent to 1 in G. Conversely, if Z is not the empty word, then W cannot be equivalent to 1 in G, since the existence of such a word Z would contradict Lemma 10.18. □

One of the consequences of a decidable word problem is that the Cayley graph of the group is effectively constructible.

Lemma 10.22 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $3\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then the Cayley graph $\mathcal{C}(G, A)$ is effectively constructible.*

Proof: Since by Lemma 10.21 it is decidable whether any two words are equal, it is possible to test all pairs of words of length less than n, and from this information to construct the ball $\text{Ball}(v_0, n)$ in the Cayley graph. This can be done successively for larger and larger n, to construct as much of the Cayley graph as is desired. □

Under a slightly stronger restriction on the choice of constants, the conjugacy problem is also decidable.

Lemma 10.23 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $4\alpha + 2\gamma + \delta \leq \frac{1}{2}$, then the group G has a decidable conjugacy problem. In particular, this is true when $\alpha \leq \frac{1}{8}$ and $\gamma = \delta = 0$ or whenever $\alpha \leq \frac{1}{10}$.*

Proof: Let X and Y be words, and by Lemma 10.5 assume without loss of generality that the cycles of X and Y are Dehn-reduced. A word is conjugate to 1 in G iff it is equivalent to 1 iff its Dehn reduction is the empty word by Lemma 10.18 and Lemma 10.5. Thus it is decidable whether a word is conjugate to 1 in G. If neither word is conjugate to 1 but they are conjugate to each other then there exists an $\mathcal{R}$-diagram Δ which satisfies the conclusion of Lemma 10.19. The crux of the problem is to decide in a finite amount of time whether such a diagram Δ exists.

First, start with the Dehn-reduced abstract loop X, and find all subwords U in X and general relators R with $|U| \leq |X|$, U readable in R, and $d_R(U) \geq \alpha$. Since by Axiom 4 $|X| \geq |U| \geq (\alpha - \delta)|R| > 0$, there is an a priori bound on the length of the relator R. By Lemma 9.11 this in turn shows that only a finite number of possible general relators need to be considered in the search for an $\mathcal{R}$-diagram having the properties of Δ. Also, since every cell in Δ would have to contain an edge of X, there is a bound on the number of cells of potential diagrams. And from the bounds on the number of relators and their lengths comes a bound on the number of edges in Δ as well. At this point it is possible simply to exhaustively search among the possible diagrams with the right structure; the outcome of the search decides the conjugacy problem for the words X and Y. □

11 Gromov Hyperbolic Groups

This section derives a number of results on geodesic words which are then used to demonstrate how close general small cancellation presentations are to hyperbolic groups in the sense of M. Gromov [7]. The main results are that finitely presented general small cancellation groups are hyperbolic, and more generally, general small cancellation groups are the direct limit of the hyperbolic groups obtained by choosing finite subsets of the general relators.

11.1 Geodesics

Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If Y is a geodesic in the Cayley category of the presentation then Y is called a geodesic with respect to $\mathcal{R}$. If Y is a geodesic with respect to $\mathcal{R}$ which is equivalent to X in G, then Y is called a geodesic of X in G. Let B be an $\mathcal{R}$-category. Suppose that whenever X is a path in B and Y is a geodesic of the word X in G, there is another path in B reading Y which is homotopic to X relative to its endpoints. In this case B is said to be geodesic-closed with respect to $\mathcal{R}$. If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation and B is an $\mathcal{R}$-category, then under certain conditions it is possible to show that every path in B is homotopic to a path which reads a geodesic in G.

Lemma 11.1 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let B be an $\mathcal{R}$-category which is $(\frac{1}{2} - \delta)$-closed with respect to $\mathcal{R}'$. If X is a path in B, and Y is a Dehn reduction of X with respect to $\mathcal{R}'$, then there is a path reading Y in B between the same two vertices which is homotopic to X relative to its endpoints. Similarly, if X is a loop in B, and Y is a Dehn reduction of X with respect to $\mathcal{R}'$, then there is a loop reading Y in B which is homotopic to the loop X. These results are true in particular whenever B is α-closed with respect to $\mathcal{R}'$.*

Proof: Let Z be a word or cycle obtained from X by a single Dehn reduction. In particular, suppose that there exist words U and V and a general relator R such that UV^{-1} is a representative of R, U is a subword of X, $|V| < |U|$, and Z is the result of replacing the subword U in X with V. Since $2|U| > |U|+|V| \geq |UV|$ and since $|UV| \geq |R|$ by the definition of $|R|$, it follows that $|U| > \frac{1}{2}|R|$. Thus by Axiom 4, $d_R(U) > \frac{1}{2} - \delta$. Because B is $(\frac{1}{2} - \delta)$-closed, there is a functor from R to B which sends the reading of U in R to the reading of U in B. In particular this shows that the word V is readable in B between the same two vertices. Thus Z is readable in B, and in the case where X and Z are words they have the same endpoints. Moreover, since R is contractible, the readings of U and V are homotopic in R relative to their endpoints, and the homotopy thus constructed is sent under the functor into B to a homotopy between the readings of U and V in B relative to their endpoints. And finally, this shows that X and Z are homotopic in B. If X and Z are words, then the homotopy can be performed relative to the endpoints of X. Since Y is obtained from X by a finite sequence of such reductions, the results follow. □

Notice that Y in the statement of the lemma is a Dehn reduction of X and not an arbitrary Dehn-reduced word equivalent to X in G. There is no guarantee, and in fact it is not true, that arbitrary Dehn-reduced equivalent words can always be obtained from a given word by a sequence of Dehn reductions. The next two lemmas address the existence of such paths. The inclusion of such paths in an arbitrary category B requires a strong hypothesis, such as that B is α-closed, not merely that it is $(\frac{1}{2} - \delta)$-closed.

Lemma 11.2 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If $XUYV$ is a representative of a general relator $R \in \mathcal{R}$ and X is $(1 - 3\alpha)$-free, while U and V are short with respect to R, then $d_R(Y) > \alpha$. As a consequence, if $\alpha \leq \frac{1}{8}$ and X is Dehn-reduced then Y is long.*

Proof: Since U and V are short, $d_R(U) < \alpha$, and $d_R(V) < \alpha$. Combining these inequalities with the properties defining relator metrics yields $1 \leq d_R(XUYV) \leq d_R(X) + d_R(U) + d_R(Y) + d_R(V) < d_R(Y) + 1 - \alpha$. Thus $d_R(Y) > \alpha$. Finally, if X is Dehn-reduced with respect to R and $\alpha \leq \frac{1}{8}$, then by Lemma 10.3, $d_R(X) < \frac{1}{2} + \alpha \leq 1 - 3\alpha$. □

Lemma 11.3 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let B be an $\mathcal{R}$-category which is collapsed and α-closed with respect to $\mathcal{R}'$. If X is a path in B and Y is a word which is 3α-complement-free with respect to $\mathcal{R}'$ and equivalent to X in $G' = \langle A|\mathcal{R}'\rangle$, then there is a path reading Y in B which is homotopic to the path X relative to its endpoints. Similarly, if X is a loop in B and Y is a cycle which is 3α-complement-free with respect to $\mathcal{R}'$ and conjugate to X in G', then there is a loop reading Y in B which is homotopic to the loop X.*

Proof: By Lemma 11.1 it is enough to show the result when X is Dehn-reduced. In this case Lemma 10.20 and Lemma 10.19 construct $\mathcal{R}$-diagrams which have very particular structures. Because each cell in the diagram has a boundary arc in Y which is 3α-complement-free and rungs which are short, it follows immediately that the boundary arc in X has length at least α relative to the relator metric of the general relator which the cell represents. Since B is α-closed, the entire boundary of the cell is readable in B as a loop extending the reading of the portion of the boundary in the path X. Next since B is collapsed, the labeling of the 1-skeleton is deterministic so that the word or cycle Y is readable as a path or loop in B. Clearly it is homotopic to the path or loop X. In the case where X and Y are words, the path Y is read between the same vertices as X, and the homotopy can be performed relative to the endpoints of X. □

Lemma 11.4 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let B be an $\mathcal{R}$-category which is collapsed and α-closed with respect to $\mathcal{R}'$. If X^n is a loop in B and Y is a cycle which is 3α-complement-free with respect to $\mathcal{R}'$ and conjugate to X in G', then there is a loop reading Y^n in B which is homotopic to the loop X^n. In particular, if X^n*

is a loop in a general relator $R \in \mathcal{R}_k$, and Y is a cycle which is 3α-complement-free with respect to $\mathcal{R}(k-1)$ and conjugate to X in $G(k-1)$, then there is a loop reading Y^n in R which is homotopic to the loop X^n.

Proof: The proof is nearly identical to the proof of Lemma 11.3 except that after creating a conjugacy diagram between the cycles X and Y it is the n-fold cover of this diagram which is used to create the loop Y^n in B. By Lemma 9.8, the above reasoning is applicable to the situation mentioned in the final sentence. □

Lemma 11.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If B is an $\mathcal{R}$-category which is collapsed and α-closed with respect to $\mathcal{R}'$, then B is geodesic-closed with respect to $\mathcal{R}'$.*

Proof: Since words which are geodesic with respect to $\mathcal{R}'$ are clearly Dehn-reduced with respect to $\mathcal{R}'$, by Lemma 10.3 they are also $(\frac{1}{2}+\alpha)$-free with respect to $\mathcal{R}'$. Because $\alpha \leq \frac{1}{8}$, they are also $(1-3\alpha)$-free with respect to $\mathcal{R}'$, and the result follows from Lemma 11.3. □

11.2 Hyperbolic Groups

Let XYZ be a geodesic triangle in the Cayley graph of a group G. If every vertex in Z is within τ units of a vertex of either X or Y as measured by the graph metric, then it is called τ-thin with respect to the side Z. If it is τ-thin with respect to each of the three sides then the triangle itself is said to be τ-thin. If all geodesic triangles in the Cayley graph of G are τ-thin for some fixed non-negative real constant τ then the group is called τ-hyperbolic, and a group which is τ-hyperbolic for some non-negative τ is simply called a hyperbolic group, or sometimes a Gromov hyperbolic group. The next several lemmas show that subject to certain restrictions on the constants, general small cancellation presentations satisfy a variety of the thin triangle condition.

Lemma 11.6 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let C be its Cayley category. If X, Y, and Z are paths such that the cycle XYZ is readable as a loop in C, then there exist words U, V, and W, such that UV^{-1}, VW^{-1}, and WU^{-1} are $\frac{2}{3}$-reduced and are equivalent to X, Y and Z respectively. In addition, U, V, and W can be chosen so that $|X|+|Y| \geq |U|+|V|+|W|$.*

Proof: Start with any set of words U, V, W, such that UV^{-1}, VW^{-1}, and WU^{-1} are equivalent to X, Y, and Z, such as $U = X$, $W = Y^{-1}$, and $V = \emptyset$. The word $WU^{-1} = Y^{-1}X^{-1}$ is equivalent to Z in G because XYZ is equivalent to 1 in G by Lemma 8.11. To complete the proof, the words U, V, and W will be altered so that the sum $|U|+|V|+|W|$, will always be decreasing.

First, replace U, V, and W with words which are geodesic in $\mathcal{R}$. Next, if the word, say UV^{-1}, contains more than two-thirds of a representative, then since U and V are geodesics, the subword must contain the vertex between U and V^{-1}. Let $U = U_1U_2$ and $V = V_1V_2$ so that $U_2V_2^{-1}$ is the subword in question. The

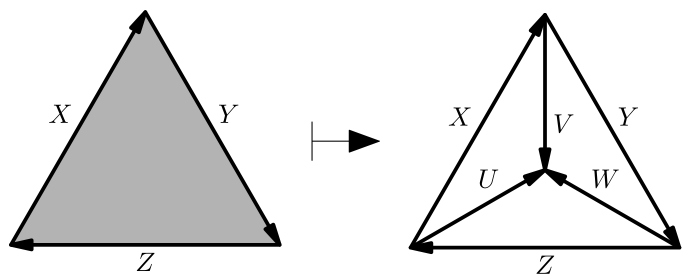

Figure 29: An illustration of Lemma 11.6

remaining portion of the representative must be a word P from the initial vertex of U_2 to the initial vertex of V_2, and it must be strictly shorter than either U_2 or V_2. Without loss of generality, assume it is shorter than U_2. Then U_2 can be replaced with P and the words U, V, and W relabeled accordingly. Specifically, the new U is U_1P, the new V is V_1, and the new W is WV_2^{-1}. Notice that the sum $|U| + |V| + |W|$ strictly decreases under this procedure. Since the sum is a positive integer, this process can be carried out only a finite number of times before it ceases to be possible. When the process stops, the words U, V, and W satisfy the conditions of the lemma. The final statement follows from the observations that $|X| + |Y|$ is the initial value of the sum $|U| + |V| + |W|$, and that this sum only decreases throughout the procedure. □

Lemma 11.7 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $3\alpha + \delta \leq \frac{1}{3}$, and let C be its Cayley category. If X, Y, and Z are $(1 - 4\alpha)$-free paths in C which form a triangle, then every vertex v in Z is connected to a vertex in either X or Y by a path PQ where P and Q are each readable in some general relator in $\mathcal{R}$.*

Proof: Let U, V, and W be the paths created by Lemma 11.6, and let Δ be a connected and simply connected $\mathcal{R}$-diagram with boundary cycle XVU^{-1}. By Lemma 10.15 Δ can be assumed to have only short internal arcs. Suppose there is an exposed cell D with $i(D) \leq 3$, which represents a general relator R, and whose boundary arc is contained in the subword VU^{-1}. Then by Axiom 4, the geodesic connecting the adjoining internal arcs of D has a length of less than $(3\alpha + \delta)|R|$. Since by assumption this is less than $\frac{1}{3}|R|$, then the word read by the boundary arc must be more than two-thirds of a representative of R, contradicting the choice of U, V, and W. Thus Δ has no such exposed cells. On the other hand, since X is $(1 - 4\alpha)$-free, Δ satisfies all of the conditions necessary to apply Lemma 10.20. And as a consequence, every vertex of X is joined to a vertex in UV^{-1} by a path which is readable in a general relator in $\mathcal{R}$, and vice versa. The same is true of Y and Z as well. Thus every vertex in Z is connected to a vertex in either U or W by a path Q which is readable in

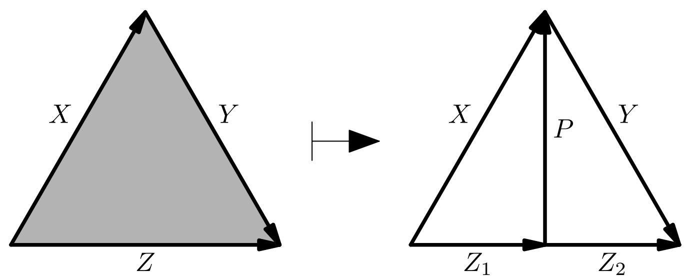

Figure 30: An illustration of Lemma 11.8

some general relator in $\mathcal{R}$. If the vertex is in U then this vertex is connected to a vertex in X by a path readable in a general relator, and if the vertex is in W, then it is connected to a vertex in Y by a path readable in a general relator. □

Lemma 11.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $3\alpha + \delta \leq \frac{1}{3}$ and $\alpha \leq \frac{1}{10}$, and let X, Y, and Z be $(1-4\alpha)$-free paths in its Cayley category which form a triangle. If B and C are α-closed categories in the Cayley category which contain the paths X and Y respectively, then there exist words Z_1, Z_2 and P such that Z_1P is read in B between the same vertices as X, $P^{-1}Z_2$ is read in C between the same vertices as Y, and $Z_1Z_2 = Z$. In particular Z is contained in the union of B and C.*

Proof: By Lemma 11.6 there exist words U, V and W satisfying the conditions of that lemma. Since UV^{-1} is equivalent to X in G and $(1-3\alpha)$-free, by Lemma 11.3 it is contained in B. Similarly, there is a path reading VW^{-1} in C between the same vertices as the path Y. Since the Cayley category is deterministic, the reading of V in both instances is the same. Thus the vertex v which is the terminal vertex of V is contained in both B and C. Let Δ be an $\mathcal{R}$-diagram with no long internal arcs, which shows that WU^{-1} is equivalent to Z in G. Since Z is $(1-4\alpha)$-free, by Lemma 11.2 every cell in Δ is long in WU^{-1}. If it is long in U then since U is contained in B, the entire general relator represented by the cell is contained in B. Similarly, if the cell is long in W then since W is contained in C the entire general relator represented by the cell is contained in C. The only cell in question is the cell D which contains the vertex v in its boundary. Since $\alpha \leq \frac{1}{10}$, $\frac{1}{2} - 3\alpha$ is at least 2α. Thus either the boundary of D in U is long or the boundary of D in W is long. In either case, Z is contained in the union of B and C and if P is a path in Δ to the crossover point, then Z is divided into Z_1 and Z_2 which satisfies the lemma. □

Lemma 11.9 *If $G = \langle A|\mathcal{R}\rangle$ is a finitely presented general small cancellation presentation with $3\alpha+\delta \leq \frac{1}{3}$ and $\alpha \leq \frac{1}{10}$, then the group G is Gromov-hyperbolic. In particular, the result is true when $\alpha = \frac{1}{10}$ and $\gamma = \delta = 0$, or whenever $\alpha \leq \frac{1}{12}$.*

Proof: Since $\mathcal{R}$ is finite, there is a finite bound on the diameters of all of the general relators in $\mathcal{R}$. Call this bound τ. By Lemma 11.7 every geodesic triangle in the Cayley graph is 2τ-thin. Since the choice of τ is independent of the triangle involved, the group G is 2τ-hyperbolic, and G is therefore a hyperbolic group. □

Since every finitely presented general small cancellation group is hyperbolic, the results of this well-developed theory are available for application. In the other direction, the axioms of general small cancellation theory provide an efficient way to construct examples which are necessarily hyperbolic. Potential applications of hyperbolic group theory to specific presentations hinge on the ability to show that the presentation is indeed hyperbolic, using one of the many equivalent definitions of hyperbolicity such as the thin triangle condition. It is precisely at this point where most of the practical difficulties arise and where general small cancellation theory is of the most use. General small cancellation groups provide an array of accessible examples of hyperbolic groups, and the hyperbolicity in this instance is easy to check since all that is needed is to check the axioms of general small cancellation theory.

Lemma 11.10 *If $G = \langle A|\mathcal{R}\rangle$ is a finitely generated general small cancellation presentation with $3\alpha+\delta \leq \frac{1}{3}$ and $\alpha \leq \frac{1}{10}$, then the group G is the direct limit of hyperbolic groups. In particular, the result is true when $\alpha = \frac{1}{10}$ and $\gamma = \delta = 0$, or whenever $\alpha \leq \frac{1}{12}$.*

Proof: Recall that by definition general small cancellation groups are finitely generated. Since A is finite, the number of relators must be countable. Using the previous lemma, add the relators in one at a time, together with any others which need to be added in order to add in that particular relator. The number of relators so far will always be finite, so that by Lemma 11.9 the group defined so far will always be hyperbolic. Also, since the relators were arranged so that all of them are used in the limit, the direct limit of the groups must satisfy all of these relators. Finally, since the group itself satisfies no other relations, the group G must be equal to the limit. □

Since it is known that the Burnside groups are not automatic and thus not hyperbolic (see [5]), one application of Lemma 11.9 is that the Burnside groups cannot be finitely presented using general small cancellation theory. It will be shown that the Burnside groups of sufficiently large exponent do, however, possess a general small cancellation presentation. It follows by Lemma 11.10 that the Burnside groups of sufficiently large exponent are the limit of hyperbolic groups.

Part V
Consequences

Part V is devoted to proving various consequences of the axioms for general small cancellation presentations. Section 12 shows that the Cayley category of a general small cancellation presentation is contractible. Section 13 represents a detour into an area which is useful in the construction of general relators and general small cancellation presentations. In this section it is shown that a given $\mathcal{R}$-category which is $(1 - 3\alpha - 2\epsilon)$-closed can be embedded in another $\mathcal{R}$-category which is α-closed in a unique and minimal way. Finally, in Section 14 the finite subgroups of a general small cancellation presentation are shown to be contained within the automorphism groups of the general relators.

12 Cayley Categories Revisited

In this section the Cayley category of a general small cancellation is shown to be proper and contractible. Then using these results the torsion elements in the group are shown to occur only as automorphisms of general relators. At the end of the section the conditions under which the Poincaré construction of a general small cancellation presentation is an Eilenberg-MacLane space are described.

12.1 Cayley Categories are Proper

The next several lemmas show that Cayley categories are proper, and they provide conditions under which $\mathcal{R}$-categories can be embedded into Cayley categories.

Lemma 12.1 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, and let C be its Cayley category. If B is an $\mathcal{R}$-category which is connected, collapsed and $(1 - 3\alpha)$-closed with respect to $\mathcal{R}$, and $f : B \to C$ is an $\mathcal{R}$-functor, then f is an embedding of B in C and B is simply connected.*

Proof: Let W be a path in B between vertices u and v. The restriction on α guarantees that $1 - 3\alpha \geq \frac{1}{2} + \alpha$. Since B is $(1 - 3\alpha)$-closed, by Lemma 10.4 W can be assumed to be $(1-3\alpha)$-free. If u and v are sent to the same vertex in C, then by Lemma 8.11 W is equivalent to 1 in G. Since W is also $(1 - 3\alpha)$-free, by Lemma 10.18 W is the empty word and u and v are identical in B. Thus f is injective on vertices, and by Lemma 6.5, it is an embedding.

Since by Lemma 4.7 every topological loop based at a vertex is homotopic to a loop in the 1-skeleton, to show B is simply connected it is sufficient to show that each loop in the 1-skeleton of B is contractible to a point. Let W be a loop in B based at v. Since B is embedded in C, W is a loop in C, and by Lemma 8.11 it is equivalent to 1. Thus by Lemma 10.4 and Lemma 10.18, the loop W is homotopic in B to the empty loop based at v. This shows that B is simply connected. □

Lemma 12.2 *Let $G = \langle A|\mathcal{R}\rangle$ be a general presentation, and let C be its Cayley category. If B is a connected, collapsed $\mathcal{R}$-category which is $(1-3\alpha)$-closed with respect to $\mathcal{R}$, then the universal cover of B can be embedded in C. As a special case, note that when W is a word which is $(1-3\alpha)$-free with respect to $\mathcal{R}$, then it is read in C as a simple path.*

Proof: Let B' be the universal cover of B. Since B' is an $\mathcal{R}$-category, by Lemma 6.8 there is a functor from B' to the Poincaré construction of the presentation, and since by definition it is simply connected, this functor lifts to a functor $f : B' \to C$ by Lemma 1.13. Since B' is connected, simply connected, collapsed, and $(1-3\alpha)$-closed, by Lemma 12.1 the functor f is an embedding. The final statement is immediate. □

Lemma 12.3 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. The Cayley category of the presentation is proper, meaning that the underlying circular category is also a circular complex.*

Proof: To prove that the Cayley category is proper, it is sufficient to show that every characteristic functor is injective on vertices by Lemma 3.1. Let C/c be a particular closed cone of C which is isomorphic with a general relator $R \in \mathcal{R}$, and let $\phi_c : C/c \to C$ be its characteristic functor. Notice that C/c is collapsed, connected and α-closed by Lemma 9.8, by the definition of circular cones, and by Axiom 1, respectively. Since $\alpha < 1 - 3\alpha$ for all general small cancellation presentations, Lemma 12.1 is applicable, and it shows that ϕ_c is an embedding and is consequently injective on vertices. □

12.2 Torsion Elements

The next several lemmas show that torsion elements correspond to automorphisms of general relators.

Lemma 12.4 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation and a general relator $R \in \mathcal{R}$ has a non-trivial automorphism, then G contains a torsion element.*

Proof: Suppose R is a general relator which has a non-trivial automorphism $f : R \to R$. If u is a vertex in R and $f(u) = v$, then u and v must be distinct by Lemma 6.3. Let W be a path in R from u to v. Since f is an automorphism, W^i is readable in R starting at u for all natural numbers i. Since R is finite, there is an n such that W^n is a loop in R. Using an attaching functor for R, W^n is also a loop in the Cayley category of the presentation, and thus by Lemma 8.10 and Lemma 4.5 it must be equivalent to 1 in G. However, since the Cayley category is proper by Lemma 12.3, W itself is not a loop, and by Lemma 8.10 and Lemma 4.5 again, W is not equivalent to 1. Thus W is a torsion element of G. □

Lemma 12.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If X^n is readable as a loop with non-trivial winding number in a general relator*

$R \in \mathcal{R}$ for some $n > 1$, but X itself is not readable as a loop in R, then the path X induces a non-trivial automorphism of R.

Proof: Since X^n is a loop with non-trivial winding number, by one of the defining properties of a relator metric, $d_R(X^n) \geq 1 \geq \alpha$. Thus by Axiom 1 and Lemma 2.1 there is an automorphism of R which sends the initial vertex of the path X to its terminal vertex. □

If X is a path in R such as is described in the statement of the lemma, then X is called a rotation of R, and the automorphism represented by X is said to rotate R.

Lemma 12.6 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation. If Y is a word of order n, $1 < n < \infty$ in G, which is cyclically reduced in the free group and whose cycle is $(1-4\alpha)$-free with respect to $\mathcal{R}$, then there exists a general relator $R \in \mathcal{R}$ which is rotated by Y.*

Proof: Since $Y^n = 1$ in G and cyclically reduced, by Lemma 10.15 there exists a connected and simply connected $\mathcal{R}$-diagram Δ with boundary cycle Y^n over a $C(6)$-map. Let Δ be chosen so that it is a diagram of minimal rank having these properties. By Lemma 10.17 Δ contains an exposed 2-cell D representing a general relator R with internal degree at most 3. Let U be the path which reads the unique boundary arc of D. Since each of the internal arcs has length less than α, $d_R(U) > 1-3\alpha$ by Lemma 5.18. Because the cycle of Y is $(1-4\alpha)$-free by assumption, $|U| > |Y|$. Let $U = U_1U_2$ where U_1 is the initial segment of U of length $|Y|$. Since $d_R(U) \leq d_R(U_1) + d_R(U_2)$ by the definition of the relator metric d_R, and since $d_R(U_1) < 1-4\alpha$ by assumption, it must be the case that $d_R(U_2) > \alpha$.

If U is viewed as a path in the loop Y^n, then U_2 can be viewed as the overlap between this path and another path reading U obtained by shifting the start vertex of the path $|Y|$ edges. Since each path individually is readable in R and the overlap U_2 has $d_R(U_2) > \alpha$, by Axiom 1 and Lemma 2.1 the union of the two paths is readable in R. When the path is shifted $|Y|$ more edges, the overlap is again long relative to R, and the reading of the word Y^n can be continued. In this way an arbitrarily long word periodic in Y can be read in the relator R. Since R is finite, Y^i must be a loop for some integer i, and in particular, since R embeds into the Cayley category of the presentation by Lemma 12.3, Y^n must be a loop in R. If the winding number of Y^n is 0 then by Lemma 7.12 there is a connected and simply connected $\mathcal{R}$-diagram with boundary Y^n of rank less than the rank of R, contradicting the choice of Δ. Thus the winding number of Y^n is non-zero. □

Lemma 12.7 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{10}$. If X is a word of order n, $1 < n < \infty$, then X is conjugate in G to a word Y such that a power of Y, namely Y^n, is readable as a loop with non-zero winding number in one of the general relators $R \in \mathcal{R}$, and such that Y is cyclically reduced in the free group and the cycle of Y is $(\frac{1}{2}+\alpha)$-free with respect* *to $\mathcal{R}$. Moreover, the general relator R is uniquely determined by the word X, and the word Y represents a non-trivial automorphism of R.*

Proof: By Lemma 10.5 the cycle X is conjugate in G to a cycle Y which is cyclically reduced in the free group and $(\frac{1}{2}+\alpha)$-free with respect to $\mathcal{R}$. Because X and Y are conjugate in G, they have the same order. Also, since $\alpha \leq \frac{1}{10}$, the cycle Y is also $(1-4\alpha)$-free. Lemma 12.6 can now be applied to show that Y^n is readable as a loop with non-zero winding number in a general relator R. Next, by Lemma 9.11, the general relator R for which this is true is unique. And finally, by Lemma 12.5, Y represents a non-trivial automorphism of R. □

Since Lemma 12.7 guarantees that the general relator rotated by an element of G is unique, any rank function which turns $G = \langle A|\mathcal{R}\rangle$ into a graded presentation also provides a rank function which assigns a finite number to every automorphism of finite order based on the rank of the general relator which it rotates.

Lemma 12.8 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $\alpha \leq \frac{1}{10}$, then G is torsion-free iff all of the general relators in $\mathcal{R}$ have no non-trivial automorphisms.*

Proof: The result is a combination of Lemma 12.4 and Lemma 12.7. □

12.3 Cayley Categories are Contractible

The lemmas below show that the Cayley category of a general small cancellation presentation is always contractible as a topological space. The heart of the proof is contained in the rather long proof of Lemma 12.11. A quick sketch of the argument goes as follows. Given an arbitrary continuous function from a 2-sphere to the geometric realization of an $\mathcal{R}$-category C, the Simplicial Approximation Theorem can be applied to a subdivision of C to yield a simplicial map. After a sequence of technical intermediary steps, the focus devolves upon $\mathcal{R}$-functors from $\mathcal{R}$-spheres to C, at which point it is possible to use general small cancellation theory to show that this continuous map is contractible to a point, the original map is contractible to a point, and thus the second homotopy group of C is trivial. Lemma 12.10 establishes this important special case. More generally, the same results hold when the Cayley category is replaced by an arbitrary connected and simply connected subcategory of the Cayley category.

Lemma 12.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let Δ be an $\mathcal{R}$-sphere. If the boundary cycle of each 2-cell in Δ is cyclically reduced, then Δ must contain an arc which is long with respect to one of the 2-cells which it borders.*

Proof: Let D be one of the 2-cells in Δ. When the interior of D is removed the result is a planar, connected and simply connected $\mathcal{R}$-diagram. If this planar diagram is a $C(7)$-map then by Lemma 10.17 it contains an exposed cell D'

whose unique boundary arc measures more than $1-3\alpha$ in its relator metric. In the latter case, the unique boundary arc of D' forms a long arc between D and D' in the original $\mathcal{R}$-sphere. If, on the other hand, it is not a $C(7)$-map, then since the boundary cycles of the 2-cells are cyclically reduced, by Lemma 7.4 there must be an internal arc which is long relative to one of the 2-cells which it borders. □

Lemma 12.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let C be the Cayley category, let Δ be an $\mathcal{R}$-sphere, and let f be a continuous function from Δ to C. If f is a label-preserving map when restricted to the 1-skeleton of the $\mathcal{R}$-sphere, and f sends every 2-cell into an instance of the general relator which the cell represents, then f is homotopic to a constant map.*

Proof: Suppose that a subword of the form aa^{-1} can be read in the boundary of some 2-cell D in Δ. If the vertex in the middle of the subword has degree 2 in Δ then the entire subword can be identified to a point in Δ. If the vertex in the middle has a degree greater than 2 and the start and end vertices of the subword are distinct in Δ then the two edges can be identified so that only the boundary cycle of D is altered. If the vertex in the middle of the subword has a degree greater than 2 and the start and end vertices of the subword are not distinct then the identification of the two edges creates two spheres which are identified along an edge. Up to homotopy type this can be modified so that the spheres are identified at a point. In all three cases the result is a modification of the original $\mathcal{R}$-sphere and of the original map f. In addition, it is clear that the new map is homotopic to a constant map iff the old map is homotopic to a constant map. Since the new configuration also has a lower type, these operations can be repeated until no such subwords occur. Thus the boundary cycles of the 2-cells in Δ can be assumed to be cyclically reduced without loss of generality. Lemma 12.9 can now be applied to show that there exists a 2-cell D in Δ which must contain a long arc.

The remainder of the proof is by induction on the type of Δ. To begin, consider a sphere with at most two 2-cells. Since the cases of spheres with only one 2-cell, or two 2-cells with different ranks, are easily seen to be impossible, assume that Δ contains two 2-cells representing relators of the same rank. Since the ranks are the same, they form a cancellable pair and they represent the same general relator R. In this case, the assumptions about f guarantee that the image of Δ is contained in an instance of R in C. The general relator R is embedded in C by Lemma 12.3. Since R is contractible, the result follows.

Next suppose that the lemma has been shown for all $\mathcal{R}$-spheres of strictly lower type, and that Δ contains at least three 2-cells. In this case, there must be a 2-cell in Δ which does not border the long arc. If the interior of this 2-cell is removed then the remainder of the sphere can be viewed as a connected and simply connected $\mathcal{R}$-diagram in the plane. The modifications used in Section 10 to remove long interior arcs can then be mimicked in the $\mathcal{R}$-sphere Δ. When the arc in question is negative relative to the boundary cycle of D, for example, then the process described for removing it in Lemma 10.11 was divided into three separate cases. The first two cases create new $\mathcal{R}$-sphere structures of strictly lower type. In the third case a one-point product of two lower-type $\mathcal{R}$-spheres is created. In each case it is easy to see how to create modified functions to go with the modified configurations, so that the new functions are homotopic to a constant map iff the old function is homotopic to a constant map as well.

If the long arc is between two 2-cells of different ranks, then the arc can simply be erased. This creates a new $\mathcal{R}$-sphere structure of strictly lower type, but without changing the function f. If the long arc is between two 2-cells of the same rank then these form a cancellable pair in Δ. The removal of the cancellable pair creates a one-point product of $\mathcal{R}$-spheres of lower type. Once again it is easy to see how to create modified functions to go with the modified configurations, so that the new functions are homotopic to a constant map iff the old function is homotopic to a constant map as well.

Finally, if the long arc borders D in two distinct ways, then because both instances are sent by f to the same path in the general relator which D represents, the orientations of the two instances must be opposite with respect to the boundary of D. In particular, this shows that no long arcs bordering the same cell in two distinct ways with both orientations positive can occur in this context. Thus, a long arc which borders the same cell in two distinct ways can always be dealt with by focusing on the long negative arc. Since all four types of long arcs lead to modifications which result in configurations of strictly lower type, the inductive step completes the proof. □

Lemma 12.11 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation and C is its Cayley category, then $\pi_2(C) = 0$.*

Proof: To show that $\pi_2(C) = 0$ it is enough to show that all maps from the unit 2-sphere are homotopic to constant maps. Due to its length the proof will be divided into steps.

Step 1: Let $f_1 : S^2 \to C$ be an arbitrary continuous function from the unit 2-sphere to the Cayley category C. Since by Lemma 9.8 the general relators in $\mathcal{R}$ are thin, there is a set $\mathcal{R}'$ of standard representatives for $\mathcal{R}$ by Lemma 5.19. Then by Lemma 6.14 there is a deformation retraction of C onto a 2-dimensional $\mathcal{R}'$-category C' such that the 1-skeleton remains fixed throughout. The deformation retraction creates a homotopy of f_1 with a map f_2 from the unit 2-sphere whose image is in C'.

Step 2: Since C is a circular complex, by Lemma 12.3, and the 1-skeletons of C and C' are the same, C' has a simplicial graph as a 1-skeleton. Moreover, because of the procedure used in Lemma 6.14 there is a 1 to 1 correspondence between the 2-cells in C' and the circular cones of height at least 2 in C. In particular, if D_1 and D_2 are 2-cells in C' with the same boundary, then the corresponding closed circular cones C/c_1 and C/c_2 in C share a representative. By Lemma 9.11, Axiom 1, and the fact that C is collapsed, $c_1 = c_2$ and $D_1 = D_2$. Combined with the above, this shows that C' is a circular complex. Thus by Lemma 2.9 Chain(C') is a simplicial complex.

Step 3: Using the Simplicial Approximation Theorem, there is a simplicial subdivision Δ_1 of S^2 and a map $g_1 : \Delta_1 \to \text{Chain}(C')$, such that f_2 and g_1 are homotopic maps, and g_1 is simplicial in the broad sense. The map g_1 cannot yet be assumed to be simplicial in the strict sense adopted in this article since g_1 may be degenerate on simplices; g_1 may send triangles to edges or vertices, or send edges to vertices.

Step 4: Consider the largest connected subcomplex of Δ which is sent to the same point in Chain(C'). If this subcomplex is all of Δ then g_1 factors through a complex in which the entire sphere is identified to a point. If not, then each of the components of the complement of this subcomplex are simply connected, and g_1 factors through a complex in which the boundary of the component has been identified to a point. Lemma 7.5 can be used on the component, and on Δ minus the component, to show that the result is homeomorphic to the one-point product of spheres. This procedure need only be repeated a finite number of times since the total number of simplices is strictly decreasing. When the process stops a new structure Δ_2 and a map $g_2 : \Delta_2 \to \text{Chain}(C')$ have been created, such that g_2 is a simplicial map which is a factor of g_1 and Δ_2 is a one-point product of spheres. In particular g_2 is a factor of a map homotopic with f_1. As a result of the identifications, however, the structure Δ_2 is not a simplicial complex, because multiple edges can have the same endpoints and distinct triangles can have the same boundaries. These multiplicities will be considered one sphere at a time.

Step 5: Two distinct triangles with the same boundaries form a sphere which is sent to a triangle in Chain(C'). Thus the sphere in Δ_2 can be collapsed to a triangle and the map g_2 factors through this new structure. Similarly, if there are two edges with the same vertices, they are necessarily part of the same sphere. If this sphere contains nothing else, then the entire sphere can be collapsed to an edge, and the map g_2 factors through this new structure. If one of the disks bounded by the two edges contains no other vertices or edges, but the other disk does, then this one side can be collapsed to an edge, leaving a sphere formed by the other side. Finally, if both disks bounded by the edges contain other vertices or edges then the identification of the edges creates two spheres joined along an edge. Up to homotopy, these spheres can be moved so that they are joined only at one of the endpoints of the edge. If these reductions are carried out until no such situations exist, the result is a simplicial complex Δ_3 and a simplicial map g_3 such that Δ_3 is the one point product of edges, triangles, and simplicial spheres, and g_3 is a factor of g_2. In particular g_3 is a factor of a map homotopic to f_1.

Step 6: Every vertex in Δ_3 is sent by g_3 to a vertex in Chain(C') which corresponds to an object in C'. If every vertex in Δ_3 is labeled by the height of the corresponding object in C', then all triangles in Δ_3 contain vertices of heights 0, 1, and 2 since they must be distinct, and these are the only three heights of objects in C'. By selecting a finite tree embedded in a sphere and then peeling it off one edge at a time, it is possible to assume, up to homotopy, that every sphere is attached to the other elements in the one-point product at only one vertex, and that it is a vertex of height 0. Call this new one-point product Δ_4 and the corresponding map g_4. By construction g_4 is a factor of a map homotopic to f_1.

Step 7: Let v be a vertex of height 1 in one of the simplicial spheres in Δ_4, and let E be the set of open edges connecting v to vertices of height 0 in the same simplicial sphere. Since $g_4(v)$ is an element in Chain(C') of height 1, it corresponds to an edge e in C'. If the ends opposite v of the selected edges in E all correspond to the same endpoint of the edge e in C', or if E contains more than two edges, then cut the sphere from v along each of the edges in E, and fill in the hole with a 2-cell whose boundary cycle has length at least 4. This can be done up to homotopy by extending the map to the 2-cell by sending it all to the simplicial subdivision of the edge e. Because an edge in a simplicial complex is contractible, there is such an extension. Then repeat this process for every vertex of height 1 in the simplicial sphere in Δ_4 which satisfies the conditions. The result is a sphere which is simplicial except for a number of 2-cells which have been added. Moreover, if every edge in this sphere is considered labeled by the edge in Chain(C') (as opposed to the edge in C') to which it is mapped, then the boundary cycles of the 2-cells are Dyck words over this alphabet.

Step 8: Pick one of the new 2-cells, say D, to remove first. Since the order in which the reducible subwords are removed plays no role in the proof of Lemma 7.6, first identify those pairs of adjacent edges in ∂D where the vertex between the edges has height 1. If the only region on the other side of the edges is one of the new 2-cells then the identification creates a vertex of degree 1 which can be removed along with the edge to which it is attached. If one of the edges borders a simplicial triangle then simply identify the edges. If both of the edges border simplicial triangles then these triangles already shared the edge formed by the vertices of height 1 and 2 by the construction so far. When the edges formed by the vertices of height 0 and 1 are identified, the two triangles share two edges and all three vertices. Up to homotopy, both triangles and their common edges can be removed and the remaining edges identified.

Notice that once all of the identifications of this type have occurred the subcomplex determined by the vertices of height 0 and 1 forms a simplicial subdivision of a graph. As the added 2-cells are removed using the procedure in Lemma 7.6, some care will be taken to preserve this property. In particular, at this point the remaining boundary of the selected 2-cell can be viewed as a simplicial subdivision of a cycle sent to the edge e in C'. Thus the remaining boundary is not only a Dyck word when edges are labeled by the edges in Chain(C') to which they are sent, but it is also the simplicial subdivision of a Dyck word labeled by edges of C'. The rest of the identifications are done in pairs according to the identifications which would take place when the boundary is viewed as a Dyck word in this second sense. Because the identified boundary is contractible, the 2-cell can be viewed as collapsing into the identified boundary. When this procedure has been used to remove all of the new 2-cells, the result is a simplicial complex Δ_5 which is the one-point product of edges, triangles, and simplicial spheres. In addition, each of the simplicial spheres is the simplicial subdivision of a 2-dimensional cell category. There is also a simplicial map g_5 which is a factor of a map homotopic to f_1. Finally, by perhaps adding a finite number of edges it is possible to assume that all of the one-point products in Δ_5 involving spheres occur at vertices of height 0. In this case, the spheres can now be replaced with the 2-dimensional cell categories of which they are subdivisions without affecting the fact that they are in a one-point product attached at vertices.

Step 9: Consider one of the spheres in the domain of g_5, viewed as a 2-dimensional cell category. Let D be a cell in one of these spheres, and let v be the unique vertex contained in the interior of D when the sphere is viewed as a simplicial subdivision of this sphere. If R' is the particular 2-cell in C' determined by $g_5(v)$ then the boundary of D is sent to a cycle in the boundary of R'. Since there are no constraints on the winding numbers of the cycles, the 2-dimensional cell category is not necessarily an $\mathcal{R}'$-sphere, but the boundary cycle of D is at the least equivalent in the free group to an integral power of the word read by the boundary of R' . Alterations will now be made to Δ_5 so that the spheres remaining in the one-point product are all $\mathcal{R}'$-spheres.

Let D and R' be as above, and assume that the winding number n of the image of ∂D in $\partial R'$ is not 1. Pick a representative U of R' which is based at a vertex in the image of ∂D. If the interior of D is removed then the remainder of the sphere can be embedded in the plane so that the boundary of the connected and simply connected map which results from this action is the same as the original boundary of D. By attaching n copies of 2-cells labeled by U (or U^{-1}) to the appropriate vertex, the resulting planar diagram has a boundary cycle U' which has winding number 0 when it is pushed into R'. In particular, the resulting boundary is a Dyck word. By Lemma 7.6 the boundary edges can be identified so that the result is a one-point product of spheres and edges. Because R' is contractible, it is clear that there is a continuous map from Δ_5 to the resulting space (although the map is not cellular), and that there is a cellular map from this space to C' which is a factor of a map homotopic to f_1. Since the net result of the process described above is to remove a 2-cell whose boundary cycle has a winding number other than 1 without introducing any other such cells, the process can be repeated a finite number of times to produce a one-point product of edges, triangles, and $\mathcal{R}'$-spheres called Δ_6 for which there is a map g_6 such that g_6 is a factor of a map homotopic to f_1, and such that g_6 restricted to one of the $\mathcal{R}'$-spheres is a cellular map.

Step 10: Finally, since the cycles in $\mathcal{R}'$ are cyclically reduced by the definition of standard representatives, it follows that each of the $\mathcal{R}'$-spheres mapped into C' corresponds to an $\mathcal{R}$-sphere mapped into C satisfying the hypothesis of Lemma 12.10. Thus by Lemma 12.10 the $\mathcal{R}'$-spheres mapped into C' are homotopic to constant maps. And more broadly, since each of the elements in the one-point product is homotopic to a constant map, so is g_6, and as a result, the same is true of f_1. □

Lemma 12.12 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation, and C is the Cayley category of the presentation, then C is contractible.*

Proof: From the definition of a Cayley category of a general presentation it is immediate that C is connected and simply connected. Thus $\pi_1(C) = 0$, and by Lemma 12.11, $\pi_2(C) = 0$ as well. Using Hurewicz's Theorem, which says that the first non-trivial homotopy group is isomorphic to the first non-trivial homology group, these facts imply that $H_i(C) = 0$ for $i = 1, 2$. Since by Lemma 6.15 $H_i(C) = 0$ for all $i > 2$, all of the homology groups are trivial. Using Hurewicz's Theorem in the other direction, all of the homotopy groups must also be trivial, making C a weakly contractible space. Finally, by Lemma 1.7, C is also contractible. □

Lemma 12.13 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let C be the Cayley category of the presentation. If B is an arbitrary subcategory of C then $\pi_2(B) = 0$. If B is an arbitrary connected and simply connected subcategory of C then B is contractible.*

Proof: To prove the first statement it only needs to be noted that the proofs of Lemma 12.10 and Lemma 12.11 do not make use of the fact that C is the full Cayley category of the general small cancellation group G. The second statement follows from a similar observation about Lemma 12.12; the only properties which are needed in its proof are that the subcategory under consideration be connected and simply connected. □

12.4 Eilenberg-MacLane Spaces

A topological space is called an Eilenberg-MacLane space iff it has exactly one non-trivial homotopy group, say $\pi_n(C, v) = G \neq 1$, in which case the space is called a $\mathrm{K}(G, n)$-space. It is known that with the restriction that G be abelian if n is greater than 1, the space $\mathrm{K}(G, n)$ always exists and is unique up to homotopy equivalence. It is also well-known that a polyhedron is a $\mathrm{K}(G, 1)$-space iff its fundamental group is G and its universal cover is contractible. The next lemma will extend the results of Lemma 12.8 to include Eilenberg-MacLane spaces. Notice that the group G is assumed to be finitely presented.

Lemma 12.14 *If $G = \langle A|\mathcal{R}\rangle$ is a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{10}$, then the following five conditions are equivalent.*
(1) the group G is torsion-free
(2) all of the general relators in $\mathcal{R}$ have no non-trivial automorphisms
(3) the universal cover of the Poincaré construction is collapsed
(4) the universal cover of the Poincaré construction is contractible
(5) the Poincaré construction is a $K(G, 1)$-space.

Proof: ($1 \Leftrightarrow 2 \Leftrightarrow 3$) The first three conditions are equivalent by Lemma 12.8 and Corollary 8.6. ($3 \Rightarrow 4$) The third condition implies that the universal cover of the Poincaré construction is the same as the Cayley category of G. Thus by Lemma 12.12 the fourth condition is also satisfied. ($4 \Leftrightarrow 5$) The last two are equivalent, since the fundamental group of the Poincaré construction is the group G by definition. ($5 \Rightarrow 1$) Finally, notice that the Poincaré construction of the presentation can be subdivided to yield a finite-dimensional simplicial complex. Since it is known that a finite-dimensional $\mathrm{K}(G, 1)$ space implies that the cohomological dimension of G is finite, which in turn implies that G is torsion-free (see [3]), the proof is complete. □

Lemma 12.14 is reminiscent of Lyndon's Theorem which states that the Poincaré construction of a 1-relator group is a K$(G, 1)$-space iff the relator is simple. The above lemma verifies several cases of Lyndon's Theorem and extends the result to the realm of general small cancellation theory. The fact that finitely presented, automorphism-free general small cancellation presentations have finite and immediately constructible Eilenberg-MacLane spaces is also worth emphasizing. The usual argument for the existence of Eilenberg-MacLane spaces involves iteratively killing off the generators of higher and higher homotopy groups and it has a distinctly non-constructive flavor. One consequence of the existence of these finite Eilenberg-MacLane spaces is that the groups constructed in this way have finite cohomological dimension, thus providing a rich source of examples of such groups. In addition, these groups also provide an illustration of a theorem of J. Cannon (see [4]). Cannon proved that a finitely generated group G acts geometrically on some contractible geometry iff G has a finite Eilenberg-MacLane space $K(G, 1)$. In this instance, the contractible geometry alluded to is the Cayley category of the presentation under the natural metric on its simplicial subdivision.

13 Closures of $\mathcal{R}$-Categories

In the first half of this section one generic and two specific constructions are introduced. The generic construction is the μ-closure of an $\mathcal{R}$-category with respect to a given set of general relators. The specific constructions are the straightline construction and the circular construction. These constructions are shown, under certain conditions, to exist, to be collapsed, and, most importantly, to be α-closed. The first half concludes with a few lemmas relating the paths and cycles read in the specific constructions to the constructions themselves. The process described above uses an infinite series of local closures. In the second half of the section the process is refined so that only a finite number of local closures are needed in most cases. The main construction is contained in Lemma 13.25.

13.1 Existence of Minimal Closures

Recall that an $\mathcal{R}$-category C is called μ-closed if whenever a word U is read in C by a functor g, and it is also read in a general relator $R \in \mathcal{R}$ by a functor f with $d_R(U) \geq \mu$, then there exists a functor $h : R \to C$ with $hf = g$. The lemmas below show that every $\mathcal{R}$-category B which is embedded in a μ-closed $\mathcal{R}$-category C is contained in a unique, minimal, μ-closed $\mathcal{R}$-category in C. This $\mathcal{R}$-category will be called the μ-closure of B in C.

Lemma 13.1 *Let $\mathcal{R}$ be a set of general relators, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If $\mathcal{F}$ is a non-empty collection of $\mathcal{R}$-categories embedded in a fixed, collapsed $\mathcal{R}$-category C, and all of the categories in $\mathcal{F}$ are μ-closed with respect to $\mathcal{R}'$ for some μ, then the intersection of the $\mathcal{R}$-categories in $\mathcal{F}$ is also μ-closed with respect to $\mathcal{R}'$.*

Proof: Let R be a general relator in $\mathcal{R}'$ and let U be a path in R with $d_R(U) > \mu$. Suppose that U is also readable in an $\mathcal{R}$-category B in $\mathcal{F}$. Since B is μ-closed with respect to $\mathcal{R}'$, there is a characteristic functor from R to B such that the reading of U in R is sent to the reading of U in B. In other words, this characteristic functor on R is a closed cone of B. Since B is embedded in C, there is also a characteristic functor from R to C, and since C is collapsed, by Lemma 6.3 this is the only $\mathcal{R}$-functor from R to C which sends the reading of U in R to the reading of U in C. If U is read in the intersection of all of the $\mathcal{R}$-categories in $\mathcal{F}$ then this particular characteristic functor on R is a closed cone in each of the $\mathcal{R}$-categories, and thus it is a closed cone in the intersection. This shows that the intersection is μ-closed with respect to $\mathcal{R}'$. □

Lemma 13.2 *Let $\mathcal{R}$ be a set of general relators, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If B is an $\mathcal{R}$-category which is embedded in an $\mathcal{R}$-category C and C is μ-closed with respect to $\mathcal{R}'$ for some constant μ, then there is a unique minimum $\mathcal{R}$-category in C which contains B and is μ-closed with respect to $\mathcal{R}'$.*

Proof: Let $\mathcal{F}$ be the set of all $\mathcal{R}$-categories in C which contain B and are μ-closed with respect to $\mathcal{R}$. Since C itself satisfies these conditions, $\mathcal{F}$ is non-empty . Let B' be the intersection of all of the categories in $\mathcal{F}$. The intersection is non-empty since it at least contains B, and by Lemma 13.1, the intersection is also μ-closed with respect to $\mathcal{R}$. Thus B' is itself an element in the set, and it is contained in every other $\mathcal{R}$-category satisfying these conditions. □

The $\mathcal{R}$-category described in the lemma is called the μ-closure of B in C with respect to $\mathcal{R}'$. If C/c is a closed cone in C but not in B which is isomorphic with $R \in \mathcal{R}'$ and U is a path in the intersection of C/c and B with $d_R(U) \geq \mu$, then C/c is certainly a cone which must be included in any μ-closure of B with respect to $\mathcal{R}'$. Such a closed cone is called a cone which must be added to B immediately. The union of B and all of the closed cones in C which must be added to B immediately is called the local μ-closure of B in C with respect to $\mathcal{R}'$. The local μ-closure of B in C respect to $\mathcal{R}'$ will be denoted $\mathrm{L}(B)$ with μ, C, and $\mathcal{R}'$ understood from context. This operation can be iterated. In particular, let $B = \mathrm{L}^0(B)$, let $\mathrm{L}(B) = \mathrm{L}^1(B)$, and let $\mathrm{L}^{k+1}(B) = \mathrm{L}(\mathrm{L}^k(B))$ for all k. Notice that clearly, $\mathrm{L}^i(B) \subset \mathrm{L}^j(B) \subset C$ for all integers $i \leq j$. Finally, let $\mathrm{L}^\infty(B) = \cup_{i=0}^\infty \mathrm{L}^i(B)$. Again it is clear that $\mathrm{L}^\infty(B) \subset C$. The next lemma shows that L^∞ is identical with the μ-closure of B in C with respect to $\mathcal{R}'$.

Lemma 13.3 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let B and C be $\mathcal{R}$-categories, with B embedded in C, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let μ be a constant. If C is itself μ-closed with respect to $\mathcal{R}'$, then $L^\infty(B)$ is the μ-closure of B in C with respect to $\mathcal{R}'$.*

Proof: First note that $\mathrm{L}^\infty(B)$ is itself μ-closed with respect to $\mathcal{R}'$. To see this let U be a path in $\mathrm{L}^\infty(B)$ which is also readable in a general relator $R \in \mathcal{R}'$ with $d_R(U) > \mu$. Since C is μ-closed there is an $\mathcal{R}$-functor $f : R \to C$ which sends the reading of U in R to the path U in $\mathrm{L}^\infty(B)$. Since every edge in the finite path U in $\mathrm{L}^\infty(B)$ is contained in one of the $\mathrm{L}^k(B)$, the $\mathrm{L}^k(B)$ with the largest superscript contains the entire path U, and the closed cone which is the image of R under f is included in $\mathrm{L}^{k+1}(B)$. Thus $\mathrm{L}^\infty(B)$ is μ-closed. If B' denotes the μ-closure of B in C with respect to $\mathcal{R}'$, then by Lemma 13.1 the fact that $\mathrm{L}^\infty(B)$ is μ-closed with respect to $\mathcal{R}'$ implies that $B' \subset \mathrm{L}^\infty(B)$. On the other, by construction it is clear that $\mathrm{L}^k(B)$ must be contained in B' for all k, so that their union, $\mathrm{L}^\infty(B)$, is also contained in B'. Finally, since $B' \subset \mathrm{L}^\infty(B)$ and $\mathrm{L}^\infty(B) \subset B'$, they must be identical. □

Lemma 13.4 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let $\mathcal{R}'$ be a subset of $\mathcal{R}$. For all $\mathcal{R}$-categories B embedded in a Cayley category and for all μ, there exists a μ-closure of B in C with respect to $\mathcal{R}'$. Moreover, the automorphism group of the μ-closure of B contains the automorphism group of B itself.*

Proof: Since B is embedded in the Cayley category, and the Cayley category is 0-closed, and thus μ-closed, the result follows by Lemma 13.2. The second statement follows from Lemma 13.3 and the obeservation that paths which show that a closed cone must be added at one location are sent by automorphisms to paths which show that corresponding closed cones must be added at all of the corresponding locations. □

Lemma 13.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let B be a connected, collapsed $\mathcal{R}$-category. If B is $(1-3\alpha)$-closed with respect to $\mathcal{R}'$, then there is exists an $\mathcal{R}$-category C in which B is embedded, which is μ-closed with respect to $\mathcal{R}'$, and which does not properly contain any other $\mathcal{R}$-category which satisfies these conditions. In addition, the inclusion of B in C induces an isomorphism on their fundamental groups, and both B and C are collapsed.*

Proof: Let H be the fundamental group of B and let B' be its universal cover. By Lemma 12.2 B' embeds in the Cayley category of the presentation, by Lemma 1.15 the group of automorphisms of B' contains a subgroup isomorphic to H, and by Lemma 13.4 it has a μ-closure in the Cayley category. Next, by Lemma 13.3 the automorphism group of the μ-closure contains the automorphism group of B' and thus contains the subgroup isomorphic with H. By Lemma 12.1 the μ-closure of B' is simply connected. If the μ-closure of B' is viewed as a separate $\mathcal{R}$-category not embedded in the Cayley category, and then quotiented by the action of the subgroup of the automorphism group isomorphic with H, then the result is an $\mathcal{R}$-category which satisfies the statements of the lemma. Finally, the fact that B and C are collapsed follows from the fact that their universal covers embed in the Cayley category which is collapsed by definition. □

In the case where either B or its universal cover embeds in the Cayley category of a general small cancellation presentation, it possible to speak unambiguously of the μ-closure of B with respect to $\mathcal{R}'$. Although the above lemma proves the existence of a canonical $\mathcal{R}$-category containing B which is μ-closed with respect to $\mathcal{R}'$ it does not guarantee that the resulting $\mathcal{R}$-category will be finite. The proof that the μ-closure is finite will be postponed until after a more detailed investigation of the structure of μ-closures.

13.2 Closures of Words and Cycles

The next few lemmas describe some of the functorial properties of α-closures, and these are then used to create a well-defined α-closed $\mathcal{R}$ category associated with every word and cycle.

Lemma 13.6 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If $f : B \to C$ is an $\mathcal{R}$-functor between connected, collapsed $\mathcal{R}$-categories, B is $(1-3\alpha)$-closed with respect to $\mathcal{R}'$, and C is μ-closed with respect to $\mathcal{R}$, then a μ-closure of B with respect to $\mathcal{R}'$ exists, and there is a unique $\mathcal{R}$-functor from the μ-closure of B to C which, when restricted to B, is equal to f.*

Proof: Let B' be the μ-closure of B which exists by Lemma 13.5, and let B'/b be a closed cone in B' which is isomorphic to a general relator $R \in \mathcal{R}'$ but which is not a closed cone in B. If U is a path in the intersection of B'/b and B with $d_R(U) > \mu$, then since C is μ-closed there is a unique $\mathcal{R}$-functor from B'/b to C which sends the path U in B'/b to the image of the path U in B under the functor f. Next, by Lemma 1.2 the obvious $\mathcal{R}$-functor from the disjoint union of B'/b and B to C factors through the quotient of the disjoint union obtained by identifying the two paths reading U, which in turn factors through the collapse of this $\mathcal{R}$-category by Lemma 6.1. Repeating this procedure for all of the closed cones which must be added immediately to B shows that there exists a functor from $\mathrm{L}^1(B)$ to C which extends the functor f. More repetitions of the procedure show that there exists a functor from $\mathrm{L}^k(B)$ to C which extends all of the previous functors. Finally, this shows that there is a functor from L^∞ to C which extends f. By Lemma 6.3 such a functor is unique. □

Lemma 13.7 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If B and B' are $(1-3\alpha)$-closed $\mathcal{R}$-categories with B embedded in B', μ and μ' are constants with $1-3\alpha \geq \mu \geq \mu'$, and $\mathcal{R}'$ and $\mathcal{R}''$ are subsets of $\mathcal{R}$ with $\mathcal{R}' \subset \mathcal{R}''$, then the μ-closure of B with respect to $\mathcal{R}'$ is contained in the μ'-closure of B' with respect to $\mathcal{R}''$.*

Proof: Let $f : B \to B'$ be the $\mathcal{R}$-functor which embeds B in B', and apply Lemma 13.6. The conclusion is immediate. □

Lemma 13.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If X and Y are words which are $(1-3\alpha)$-free with respect to $\mathcal{R}'$ and equivalent to each other in the group $G' = \langle A|\mathcal{R}'\rangle$, then there is an isomorphism from the α-closure of the abstract path X with respect to $\mathcal{R}'$ to the α-closure of the abstract path Y with respect to $\mathcal{R}'$ which sends the endpoints of X to the endpoints of Y.*

Proof: Let B and C be the α-closures of X and Y respectively, and notice that both B and C are connected and collapsed. By Lemma 11.3, the path Y is readable in B between the same vertices as X. By Lemma 13.6 this reading of Y extends to a functor $f : C \to B$. Similarly, there is a functor $g : B \to C$ extending the reading of X in C between the same vertices as Y. Since fg is a functor B to itself which fixes a vertex, by Lemma 6.3 it is the identity functor, and similarly, gf is the identity functor on C. Thus, f is the desired isomorphism. □

Lemma 13.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. If X and Y are cycles which are $(1-3\alpha)$-free with respect to $\mathcal{R}'$ and conjugate to each other in the group $G' = \langle A|\mathcal{R}'\rangle$, then there is an isomorphism from the α-closure of the abstract loop X with respect to $\mathcal{R}'$ to the α-closure of the abstract loop Y with respect to $\mathcal{R}'$.*

Proof: Let B and C be the α-closures of X and Y respectively, and let Δ be an annular $\mathcal{R}'$-diagram with no long internal arcs, which shows that X and Y are conjugate in G'. By Lemma 11.3, the loop Y is readable in B, and by Lemma 13.6 this reading of Y extends to a functor $f : C \to B$. Similarly, there is a functor $g : B \to C$ extending the reading of X in C. If the same diagram Δ is used in both applications of Lemma 11.3 to create the readings of the loop X in C and Y in B, then the functor fg will fix the loop X, and the functor gf will fix the loop Y. Then by Lemma 6.3 fg is the identity functor on B, and gf is the identity functor on C. Thus, f is the desired isomorphism. □

Lemma 13.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$. If $\mathcal{R}'$ is a subset of $\mathcal{R}$ and W is an arbitrary word or cycle, then the α-closure of a Dehn reduction of W is a well-defined $\mathcal{R}'$-category.*

Proof: This follows immediately from Lemma 13.8 and Lemma 13.9. □

The α-closure of a Dehn reduction of a word W with respect to $\mathcal{R}'$ is called the straightline construction on W with respect to $\mathcal{R}'$ and it will be denoted $\text{str}(W, \mathcal{R}')$. Similarly, the α-closure of a Dehn reduction of a cycle W with respect to $\mathcal{R}'$ is called the circular construction on W with respect to $\mathcal{R}'$ and it will be denoted $\text{cir}(W, \mathcal{R}')$. The 1-skeletons of $\text{str}(W, \mathcal{R}')$ and $\text{cir}(W, \mathcal{R}')$ can be viewed as automata by specifying the initial and terminal vertices of the original path W as the start and stop states of the automata. In addition an orientation of $\text{cir}(W, \mathcal{R})$ is induced by the orientation of the loop W.

The most important special cases are those in which $\mathcal{R}'$ is compatible with the grading of the general relators, and these cases merit an abbreviated notation. In particular, the constructions $\text{str}(W, \mathcal{R}(k))$ and $\text{cir}(W, \mathcal{R}(k))$ will be denoted as $\text{str}_k(W)$ and $\text{cir}_k(W)$ respectively, and $\text{str}(W, \mathcal{R})$ and $\text{cir}(W, \mathcal{R})$ as $\text{str}(W)$ and $\text{cir}(W)$. Notice that since $\mathcal{R}_1 = \emptyset$, the definitions of $\text{str}_1(W)$ and $\text{cir}_1(W)$ given here coincide with those given in Section 5.

Lemma 13.11 *Let* $G = \langle A|\mathcal{R}\rangle$ *be a general small cancellation presentation with* $\alpha \leq \frac{1}{8}$*, and let* $\mathcal{R}'$ *be a subset of* $\mathcal{R}$*. If* X *is an arbitrary word and* Y *is an arbitrary cycle which is not equivalent to 1 in* $G' = \langle A|\mathcal{R}'\rangle$*, then* $str(X, \mathcal{R}')$ *is simply connected while the fundamental group of* $cir(Y, \mathcal{R}')$ *is* $\mathbf{Z}$*.*

Proof: Without loss of generality assume that X and Y are Dehn-reduced with respect to $\mathcal{R}'$. Then the abstract path X is simply connected, the abstract loop Y has a fundamental group of $\mathbf{Z}$, and by Lemma 13.5 the functor embedding them into their α-closures with respect to $\mathcal{R}'$ induces an isomorphism on the fundamental groups. □

Lemma 13.12 *Let* $G = \langle A|\mathcal{R}\rangle$ *be a general small cancellation presentation with* $\alpha \leq \frac{1}{8}$*, let* C *be its Cayley category, and let* $\mathcal{R}'$ *be a subset of* $\mathcal{R}$*. If* X *is an arbitrary path in* C*, then there is a unique embedding of* $str(X, \mathcal{R}')$ *into* C *which sends the start and stop states of* $str(X, \mathcal{R}')$ *to the start and end vertices of the path* X*.*

Proof: By Lemma 10.2 there exists a Dehn reduction of X with respect to $\mathcal{R}'$. Call this not necessarily unique reduction X'. Since C is μ-closed for all μ, by Lemma 10.4 X' is readable in C between the same vertices as X. By construction $\mathrm{str}(X, \mathcal{R}) = \mathrm{str}(X', \mathcal{R})$, and it is a connected, collapsed, and $(1 - 3\alpha)$-closed $\mathcal{R}'$-category. The latter is true since $1 - 3\alpha \geq \alpha$. Thus by Lemma 12.2 and Lemma 13.11, $\mathrm{str}(X, \mathcal{R}')$ embeds in C. Since by definition the automorphisms of a Cayley category act transitively on its vertices there is an embedding of $\mathrm{str}(X, \mathcal{R}')$ which sends its start state to the start vertex of the path X, and by Lemma 6.3 this embedding is unique. Finally, since $\mathrm{str}(X, \mathcal{R}')$ can be constructed from X', X' is certainly accepted by $\mathrm{str}(X, \mathcal{R}')$ which shows that the stop state of $\mathrm{str}(X, \mathcal{R}')$ has been sent to the end vertex of the path X', which is also the end vertex of the path X. □

Lemma 13.13 *Let* $G = \langle A|\mathcal{R}\rangle$ *be a general small cancellation presentation with* $\alpha \leq \frac{1}{8}$*, let* C *be its Cayley category, and let* B *be an arbitrary subcategory of* C *which is* α*-closed with respect to* $\mathcal{R}$*. If* X *is a path in* B *then the embedding of* $str(X)$ *into* C *is contained in* B*. As a consequence, if* X *is a path in* $str(Y)$*, then there is an embedding of* $str(X)$ *into* $str(Y)$ *which sends the start and stop states of* $str(X)$ *to the start and end vertices of the path* X*.*

Proof: Let X' be a Dehn reduction of X. By Lemma 10.4, X' is readable in B between the same endpoints. By Lemma 13.12, $\mathrm{str}(X) = \mathrm{str}(X')$ can be uniquely embedded in C so that it starts and ends at the same vertices as X and X'. If $\mathrm{str}(X)$ were not contained in B, the intersection of the two would be an α-closed $\mathcal{R}$-category containing X', which contradicts the definition of $\mathrm{str}(X)$ as the minimum such $\mathcal{R}$-category. Thus $\mathrm{str}(X)$ is contained in B. The second assertion follows immediately from the first once $\mathrm{str}(Y)$ is embedded in C by Lemma 13.12. □

13.3 Decidability

A word is said to be accepted by $\mathrm{str}(W, \mathcal{R}')$ if it readable as a path in the construction between the same vertices as the Dehn reduction of W with respect to $\mathcal{R}'$. By analogy, a cycle is said to be accepted by $\mathrm{cir}(W, \mathcal{R}')$ if it is readable as a loop in the construction with a winding number of 1 and with the same orientation as W.

Lemma 13.14 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let X and Y be words which are $(1-3\alpha)$-free with respect to $\mathcal{R}'$. The following statements are equivalent.*
(1) X and Y are equivalent in $G' = \langle A|\mathcal{R}'\rangle$
(2) str$(X, \mathcal{R}')$ and str$(Y, \mathcal{R}')$ are isomorphic as automata
(3) Y is a word accepted by str$(X, \mathcal{R}')$
(4) X is a word accepted by str$(Y, \mathcal{R}')$.

Proof: $(1 \Rightarrow 2 \Rightarrow 3)$ By Lemma 13.8 the first condition implies the second, and the second condition clearly implies the third. $(3 \Rightarrow 1)$ Suppose that Y is accepted by $\mathrm{str}(X, \mathcal{R}')$. By Lemma 13.12 $\mathrm{str}(X, \mathcal{R}')$ embeds in the Cayley category of $G' = \langle A|\mathcal{R}'\rangle$. This shows that X and Y are readable in the Cayley category between the same endpoints, and thus by Lemma 8.11 they are equivalent. $(1 \Leftrightarrow 4)$ Since the first three conditions have been shown to be equivalent, and the first condition is symmetric with respect to X and Y, it follows that the fourth condition is equivalent as well. □

Lemma 13.15 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. Words X and Y are equivalent in $G' = \langle A|\mathcal{R}'\rangle$ iff there is a start- and endpoint-preserving isomorphism between str$(X, \mathcal{R}')$ and str$(Y, \mathcal{R}')$.*

Proof: By Lemma 10.2 there exist Dehn reductions of X and Y with respect to $\mathcal{R}'$. Call these not necessarily unique reductions X' and Y', respectively. By Lemma 10.3 they are $(\frac{1}{2}+\alpha)$-free and thus $(1-3\alpha)$-free with respect to $\mathcal{R}'$ since $\alpha \leq \frac{1}{8}$. Thus Lemma 13.14 shows that X' and Y' are equivalent in G' iff there is a start- and endpoint-preserving isomorphism between $\mathrm{str}(X', \mathcal{R}')$ and $\mathrm{str}(Y', \mathcal{R}')$. Since X and X' are equivalent in G', Y and Y' are equivalent in G', and by definition $\mathrm{str}(X, \mathcal{R}') = \mathrm{str}(X', \mathcal{R}')$ and $\mathrm{str}(Y, \mathcal{R}') = \mathrm{str}(Y', \mathcal{R}')$, the proof is complete. □

Lemma 13.16 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, and let X and Y be cycles which are $(1-3\alpha)$-free with respect to $\mathcal{R}'$. The following statements are equivalent.*
(1) X and Y are conjugate in $G' = \langle A|\mathcal{R}'\rangle$
(2) cir$(X, \mathcal{R}')$ and cir$(Y, \mathcal{R}')$ are isomorphic as oriented $\mathcal{R}'$-categories
(3) Y is a cycle accepted by cir$(X, \mathcal{R}')$
(4) X is a cycle accepted by cir$(Y, \mathcal{R}')$.

Proof: ($1 \Rightarrow 2 \Rightarrow 3$) By Lemma 13.9 the first condition implies the second, and the second condition clearly implies the third. ($3 \Rightarrow 1$) Suppose that Y is accepted by $\text{cir}(X, \mathcal{R}')$. Since the abstract loop X is $(1-3\alpha)$-free with respect to $\mathcal{R}'$, it is also $(1-3\alpha)$-closed. And by Lemma 13.11 its fundamental group is $\mathbf{Z}$ so that it is possible to speak of winding numbers. Let P be a path in $\text{cir}(X, \mathcal{R}')$ from the basepoint of the loop X to the basepoint of the loop Y. Then $PYP^{-1}X^{-1}$ is a loop of winding number 0, and it lifts to a loop in the universal cover of $\text{cir}(X, \mathcal{R}')$. By Lemma 12.2, the universal cover of $\text{cir}(X, \mathcal{R}')$ embeds in the Cayley category of G' so that $PYP^{-1}X^{-1}$ is readable as a loop in the Cayley category. By Lemma 8.11, this word is equivalent to 1 in G', showing that X and Y are conjugate in G'. Thus the first three conditions are equivalent. ($1 \Leftrightarrow 4$) Since both of the first two are symmetric with respect to X and Y, it follows that the fourth condition is equivalent to the others as well. □

Lemma 13.17 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, and let $\mathcal{R}'$ be a subset of $\mathcal{R}$. Cycles X and Y are conjugate in $G' = \langle A|\mathcal{R}'\rangle$ iff there is an orientation-preserving isomorphism between cir$(X, \mathcal{R}')$ and cir$(Y, \mathcal{R}')$.*

Proof: By Lemma 10.5 there exist Dehn reductions of X and Y with respect to $\mathcal{R}'$. Call these not necessarily unique reductions X' and Y', respectively. Also by Lemma 10.5 they are $(\frac{1}{2}+\alpha)$-free and thus $(1-3\alpha)$-free with respect to $\mathcal{R}'$ since $\alpha \leq \frac{1}{8}$. Thus Lemma 13.16 shows that X' and Y' are conjugate in G' iff there is an orientation-preserving isomorphism between $\text{cir}(X', \mathcal{R}')$ and $\text{cir}(Y', \mathcal{R}')$. Since X and X' are conjugate in G', Y and Y' are conjugate in G', and by definition $\text{cir}(X, \mathcal{R}') = \text{cir}(X', \mathcal{R}')$ and $\text{cir}(Y, \mathcal{R}') = \text{cir}(Y', \mathcal{R}')$, the proof is complete. □

13.4 Finite Closures

Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation, and let $\mathcal{R}(k)$ be the general relators in $\mathcal{R}$ of rank at most k. The efficient construction of the α-closure of an $\mathcal{R}$-category with respect to $\mathcal{R}(k)$ which is given below is reminiscent of a construction of Zimin ([21]). For example, if B is an $\mathcal{R}$-category which is sufficiently closed to begin with and all of the general relators in B are in $\mathcal{R}(3)$, then B can be made α-closed by a local closure in rank 1, then rank 2, then rank 1 again, then rank 3, then rank 1, then rank 2, then rank 1. The sequence of ranks has the characteristic Zimin structure, so that like Zimin's construction, the process can be conceptualized as involving three steps. Of these, the first and the third are identical and amount to inductively applying the procedure in the previous rank. After the first step the construction is α-closed in all ranks strictly less than k, but slightly less closed in rank k itself. The second step, which is the only step involving relators of rank k, adds relators of rank k in order to make the result μ-closed in rank k for some small constant μ. This procedure, however, disrupts the closures in the lower ranks. Thus, the third step is needed in which the procedure used to make the lower ranks α-closed is

Rank	Start	After Step 1	After Step 2	After Step 3
1	$1-3\alpha-2\epsilon$	α	3α	α
2	$1-3\alpha-2\epsilon$	α	3α	α
...	...	...	...	...
$k-1$	$1-3\alpha-2\epsilon$	α	3α	α
k	$1-3\alpha-2\epsilon$	$1-3\alpha$	$\alpha-2\epsilon$	α
$(k+1)$	(δ)			$(\delta+2\epsilon)$

Figure 31: Levels of completeness

applied again. The final result is a structure which is α-closed with respect to all of the general relators in $\mathcal{R}(k)$. Because of the back-and-forth nature of the levels of completeness, Figure 31 has been provided to illustrate the process. The final row of the table shows the result of the study of Zimin-type words in the next rank. Once this has been shown, then the entire three-step process is ready to become the first step, or the third step, in the process in the next larger rank.

Because of their connection with the above procedure, the following words will be called Zimin-type words. A rank k Zimin-type word is a word $W = XYZ$ where X and Z are rank $(k-1)$ Zimin-type words and Y is a word readable in a rank k relator. To start the definition, notice that a rank 1 Zimin-type word must be the empty word, since no rank 1 general relators exist. For simplicity, Zimin-type words will be referred to as Zimin words even though this designation is usually reserved for words of a much more restricted form.

Lemma 13.18 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{12}$. If XYZ is a path in a general relator $R \in \mathcal{R}$ with $d_R(XYZ) \geq 1-3\alpha$, and X, Y and Z are $(\frac{1}{2}+\alpha)$-free words with respect to R then at least two of the three words are long in R.*

Proof: Suppose on the contrary that two of the words, say Y and Z, are short in R. Then $1-3\alpha \leq d_R(XYZ) \leq d_R(X)+d_R(Y)+d_R(Z) < d_R(X)+2\alpha$. Thus $d_R(X) > 1-5\alpha \geq \frac{1}{2}+\alpha$, since $\alpha \leq \frac{1}{12}$, contradiction. □

Lemma 13.19 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{10}$, let $\mathcal{R}'$ be a subset of $\mathcal{R}$, let C be a collapsed $\mathcal{R}$-category and let S and T be connected $\mathcal{R}$-categories embedded in C which are each α-closed with respect to $\mathcal{R}'$. Suppose further that S and T intersect and that u and v are vertices in S and T respectively. If X is the shortest path from u to the intersection of S and T, and Y is the shortest path from the endpoint of X to v, then XY is $(1-3\alpha)$-free with respect to $\mathcal{R}'$.*

Proof: First, note that since S and T are α-closed with respect to $\mathcal{R}'$, by Lemma 11.5 they are geodesic-closed with respect to $\mathcal{R}'$ as well, so that X is contained in S, Y is contained in T, and both paths are geodesics with respect to $\mathcal{R}'$. Because X and Y are geodesics with respect to $\mathcal{R}'$, by Lemma 10.3 they are $(\frac{1}{2}+\alpha)$-closed and thus $(1-3\alpha)$-closed with respect to $\mathcal{R}'$. Thus if XY is not $(1-3\alpha)$-closed with respect to $\mathcal{R}'$ it is because there are non-empty words X' and Y' and a general relator $R \in \mathcal{R}'$ such that X' is a final segment of X, Y' is an initial segment of Y, $X'Y'$ is readable in R and $d_R(X'Y') > 1-3\alpha$. Since none of the edges of X are in T, $d_R(Y') < \alpha$, because otherwise R would be included in T and X' would consist of edges in T, contradiction. By one of the properties of relator metrics, $d_R(X') > 1-4\alpha$. But since $\alpha \leq \frac{1}{10}$, $\frac{1}{2}+\alpha \leq 1-4\alpha$, so that by Lemma 10.3, X' is not a geodesic, contradiction. Therefore the assumption that XY is not $(1-3\alpha)$-free is false. □

If U is a subword of a word X and UV^{-1} is a representative of a general relator R, then replacing U with V in X is called a simple substitution.

Lemma 13.20 *Let X be readable in a general relator R and let Y be a word obtained from X by simple substitutions which shorten the length of the word. If none of the simple substitutions involve a representative of R, then $d_R(X) = d_R(Y)$.*

Proof: Assume that Y is obtained from X by a single simple substitution of V for a subword U in X, and that S is the general relator of which UV^{-1} is a representative. Since the substitution shortens the length, $2|U| > |U|+|V| \geq |S|$ so that $|U| > \frac{1}{2}|S|$ and by Axiom 4 and Axiom 7, $d_S(U) > \frac{1}{2} - \delta > \alpha$. Thus by Axiom 1, there is a functor from S to R which extends the reading of U in R. Since by assumption none of the simple substitutions involve a representative of R, it must be that the image of S is contained in ∂R by Lemma 2.1. This shows that V is readable in R between the same vertices as U, and that the new path Y is homotopic to X in ∂R relative to their endpoints. Thus, by one of the properties of relator metrics $d_R(X) = d_R(Y)$. In the general case, Y is formed by a finite number of iterations of this procedure. □

Lemma 13.21 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{8}$, let C be a collapsed $\mathcal{R}$-category and let $\mathcal{F}$ be a family of connected $\mathcal{R}$-categories embedded in C which are each α-closed with respect to $\mathcal{R}(k)$ and whose union is C. If U is a path in C then there exists a path V in C with $|V| \leq |U|$ which is homotopic to U relative to its endpoints, and such that for all $B \in \mathcal{F}$, the connected components of $V \cap B$ are geodesics with respect to $\mathcal{R}(k)$. In addition, if the original path U is readable in a general relator R, then V is also readable in R between the same vertices.*

Proof: If B is a member of $\mathcal{F}$ and U' is a subpath of U in B which is not a geodesic with respect to $\mathcal{R}(k)$, then pick V' to be such a geodesic. Since B is geodesic-closed with respect to $\mathcal{R}(k)$ by Lemma 11.5, there is a path reading V' in B homotopic to U' relative to its endpoints. If U is also read in a general relator R, then since by Lemma 9.8 and Axiom 1, R is collapsed and α-closed, it follows from Lemma 11.5 that the same replacement of U' with V' is permissible in R. After repeating this procedure a finite number of times, the process must stop, since the length of the path is strictly decreasing. The final path which results necessarily satisfies the statement of the lemma. □

Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \geq 3\beta + 2\epsilon$ and $\alpha \leq \frac{1}{12}$, let B be an $\mathcal{R}$-category which is α-closed with respect to $\mathcal{R}(k-1)$ and $(1-3\alpha)$-closed with respect to $\mathcal{R}_k$, let C be the local $(\alpha - 2\beta - 2\epsilon)$-closure of B, and let $\mathcal{F}$ be the set consisting of B and each of the closed cones of rank k in C which have been added to B to form C. Since C is contained in the full $(\alpha - 2\beta - 2\epsilon)$-closure of B with respect to $\mathcal{R}_k$, and by Lemma 13.5, B is embedded in that $\mathcal{R}$-category, it follows that B is also embedded in C. Clearly each of the elements of $\mathcal{F}$ are embedded in C, they are all connected $\mathcal{R}$-categories which are α-closed with respect to $\mathcal{R}(k-1)$, and their union is C. Thus Lemma 13.21 can be applied to C and $\mathcal{F}$ as follows.

Lemma 13.22 *Let G, B, C and $\mathcal{F}$ be as described above. If R, S, and T are distinct elements in $\mathcal{F}$ other than B, and X is a path in S which starts in $R \cap S$, ends in $S \cap T$ and is a geodesic with respect to $\mathcal{R}(k)$, then there is a path Y with the same endpoints as X which is contained in the intersection of $R \cup B \cup T$ with the α-closure of X in S with respect to $\mathcal{R}(k)$.*

Proof: Case 1: If R and T intersect, let v be a vertex in the intersection, let P be the shortest path from the initial vertex of X to v, and let Q be the shortest path from v to the terminal vertex of X. Since R, S, and T are geodesic-closed by Lemma 11.5, P is contained in $R \cap S$, and Q is contained in $S \cap T$. Since P and Q are in the intersection of two distinct closed cones of the same rank in a collapsed $\mathcal{R}$-category, by Axiom 2 $d_S(P) < \beta$ and $d_S(Q) < \beta$. If the winding number of the loop PQX^{-1} in S is n, then $n \leq d_S(PQX^{-1}) < d_S(X) + 2\beta$. If n is not 0 then $d_S(X) > 1 - 2\beta > \frac{1}{2} + \alpha$ and X could not be a geodesic in S by Lemma 10.3. Thus the winding number is 0, and by Lemma 7.12 there is a proof that PQX^{-1} is equivalent to 1 in $G(k-1)$. Since P, Q, and X are $(1-4\alpha)$-free paths with respect to $\mathcal{R}(k-1)$ and they form a triangle in the Cayley graph of $G(k-1)$, by Lemma 11.8 the path X itself is contained in the $R \cup T$.

Case 2: If R and T do not intersect, then let P be the shortest path from the initial vertex of X to a point in $S \cap B$, let Q be the shortest path from a vertex in $S \cap B$ to the terminal vertex of X, and let Y be the shortest path from the terminal vertex of P to the initial vertex of Q. By Lemma 13.19 both PY and YQ are $(1-3\alpha)$-free with respect to $\mathcal{R}(k-1)$. If PYQ is not $(1-3\alpha)$-free with respect to $\mathcal{R}(k-1)$ then there is a subword of PYQ containing a non-empty portion of P, all of Y and a non-empty portion of Q which has a length longer than $1 - 3\alpha$ when measured by the relator metric of some general relator in $\mathcal{R}(k-1)$. By Lemma 13.18 this would mean that either the portion in P or the portion in Q is long in the general relator. If without loss of generality it is assumed that the portion in P is long in the relator, then, since R is α-closed with respect to $\mathcal{R}(k-1)$, the entire relator is contained in R, and R and T would intersect, contradicting the initial assumption. Thus PYQ is $(1-3\alpha)$-free with respect to $\mathcal{R}(k-1)$.

By Axiom 2, $d_S(P) < \beta$, and $d_S(Q) < \beta$. Since Y is a geodesic by Lemma 10.3, $d_S(Y) < \frac{1}{2} + \alpha \leq 1 - 3\alpha$ since $\alpha \leq \frac{1}{12}$. If the loop $PYQX^{-1}$ has a non-zero winding number in S then $1 \leq d_S(PYQX^{-1}) \leq d_S(X) + 1 - 3\alpha + 2\beta \leq d_S(X) + 1 - \alpha$, by the properties of relator metrics and the constraints of the constants listed in Axiom 7. Thus $d_S(X) > \alpha$, and the α-closure of X in S with respect to $\mathcal{R}(k)$ is S itself and thus includes the path PYQ. On the other hand, if $PYQX^{-1}$ has winding number 0, then by Lemma 7.12 and Lemma 8.8 it is a loop in the Cayley category of $G(k-1)$ and thus by Lemma 8.11, X and PYQ are equivalent in $G(k-1)$. By Lemma 11.3 it follows that PYQ is contained in the α-closure of X with respect to $\mathcal{R}(k-1)$ and thus also clearly in the α-closure of X with respect to $\mathcal{R}(k)$. □

Lemma 13.23 *Let G, B, C, and $\mathcal{F}$ be as described above, and let X be a path in C. If S and T are elements of $\mathcal{F}$ which contain the initial and terminal vertices of X respectively, then there is a path with the same endpoints as X which is contained in the intersection of $S \cup B \cup T$ with the α-closure of X with respect to $\mathcal{R}(k)$. In particular, if S and T do not intersect, then the path mentioned above can be chosen to be of the form PYQ where P is the shortest path from the initial vertex of X to a vertex in the intersection of S and B, Q is the shortest path from a vertex in the intersection of B and T to the terminal vertex of X, and Y is the shortest path between the terminal vertex of P and the initial vertex of Q. If S and T do intersect, then the path can be chosen to be of the form PQ, where P is the shortest path from the initial vertex of X to a point in the intersection of S and T, and Q is the shortest path from the terminal vertex of P to the terminal vertex of X.*

Proof: Let $\mathcal{F}'$ be the smallest finite subset of $\mathcal{F}$ such that $\mathcal{F}'$ contains S, B, and T, and such that the union of the elements in $\mathcal{F}'$ contains a path X' which is accepted by the α-closure of X with respect to $\mathcal{R}(k)$. By applying Lemma 13.21 and using Lemma 13.6 it is possible to assume that X' satisfies the conclusion of Lemma 13.21. If $\mathcal{F}'$ contains an element other than S, B, and T, then applying Lemma 13.22 and Lemma 13.6 to the portion of the path X' which reads this element shows that $\mathcal{F}'$ is in fact not minimal, contradiction. Thus $\mathcal{F}'$ contains only S, B, and T. The properties of the particular paths then follow from the proof of Lemma 13.22. □

Lemma 13.24 *If G, B, C, and $\mathcal{F}$ are as described above, then C is 3α-closed with respect to $\mathcal{R}(k-1)$ and $(\alpha - 2\epsilon)$-closed with respect to $\mathcal{R}_k$.*

Proof: Let X be a path readable in C and also in a general relator $R \in \mathcal{R}(k)$ and let S and T be elements in $\mathcal{F}$ which contain the initial and terminal vertices of the path X in C respectively. If there does not exist a functor from R to C which extends the reading of X in C, then the path Y in $S \cup B \cup T$ and readable in R guaranteed by Lemma 13.23 must satisfy $d_R(X) = d_R(Y)$ by Lemma 13.20. Let $Y = Y_1Y_2Y_3$ be a partitioning of Y so that Y_1 lies in S, Y_2 lies in B, and Y_3 lies in T. If S and T intersect so that Y never includes an edge in B, then consider $Y_2 = \emptyset$ and $d_R(Y_2) = 0$. If R is in $\mathcal{R}_k$ then $d_R(Y_1) < \beta$, and $d_R(Y_3) < \beta$

by Axiom 2. If $d_R(Y) \geq \alpha - 2\epsilon$ then $d_R(Y_2) < \alpha - 2\beta - 2\epsilon$ and R is already included in the construction C. Thus C is $(\alpha - 2\epsilon)$-closed with respect to $\mathcal{R}_k$. If R is in $\mathcal{R}(k-1)$ then either $d_R(Y_1) < \alpha$, $d_R(Y_2) < \alpha$, and $d_R(Y_3) < \alpha$, or R is already included in S, B, or T since each of these constructions is α-closed with respect to $\mathcal{R}(k-1)$. Thus C is 3α-closed with respect to $\mathcal{R}(k-1)$. Finally, notice that by construction, every general relator in $\mathcal{F}$ other than B shares a path with B which measures at least $\alpha - 2\beta - 2\epsilon$ in the relator metric of the general relator. □

Lemma 13.25 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{12}$ and $3\beta + 2\epsilon \leq \alpha$, and let ϵ be a constant which satisfies Axiom 5. Suppose that ϵ is also an upper bound on the length of a rank i Zimin word, $i \leq k$, which is read in any general relator R having a rank strictly greater than i, as measured by the relator metric of R. If B is an $\mathcal{R}$-category which is $(1 - 3\alpha - 2\epsilon)$-closed with respect to $\mathcal{R}(k)$, then there is an $\mathcal{R}$-category C containing B which is α-closed with respect to $\mathcal{R}(k)$, and such that every vertex in C is connected to a vertex in B by a path reading a rank k Zimin word. In addition, if B is μ-closed with respect to $\mathcal{R}_j$ with $j > k$, then C is $(\mu+2\epsilon)$-closed with respect to $\mathcal{R}_j$.*

Proof: If $k = 1$ then the result follows immediately from Lemma 13.24, so suppose that the lemma is true for all ranks strictly less than k, and $k > 1$. By applying the lemma in the previous rank, there is a finite $\mathcal{R}$-category B_1 containing B which is α-closed with respect to $\mathcal{R}(k-1)$ and $(1 - 3\alpha)$-closed with respect to $\mathcal{R}_k$. Then by applying Lemma 13.24 there is an $\mathcal{R}$-category B_2 containing B_1 such that B_2 is 3α-closed with respect to $\mathcal{R}(k-1)$ and $(\alpha - 2\epsilon)$-closed with respect to $\mathcal{R}_k$. Finally, by applying this lemma again in the previous rank, there is an $\mathcal{R}$-category C containing B_2 such that C is α-closed with respect to $\mathcal{R}(k)$. By induction every vertex in C is connected to a vertex in B_2 by a rank $(k-1)$ Zimin word. The fact that B_2 is obtained from B_1 by attaching rank k relators to B_1 shows that every vertex in B_2 is connected to a vertex in B_1 by a path readable in a rank k relator. Then, by induction again, every vertex in B_1 is connected to a vertex in B by a rank $(k-1)$ Zimin word. Combining these facts shows that every vertex in C is connected to a vertex in B by a rank k Zimin word.

Finally, let X be a path readable in C and also in a general relator $R \in \mathcal{R}_j$, $j > k$, such that there does not exist a functor from R to C which extends the reading of X in C. By repeatedly using Lemma 13.23 applied to each of the local closures, there is a word PYQ readable in both C and R where both paths are homotopic to X relative to its endpoints, and such that P and Q are rank k Zimin words and Y is a possibly empty path in the $\mathcal{R}$-category B. Moreover, by Lemma 13.20 and the fact that R is not included in C at this point, $d_R(X) = d_R(PYQ)$. By the assumptions stated in the lemma, $d_R(P) < \epsilon$ and $d_R(Q) < \epsilon$. Also, since R is not attached to C and B is μ-closed, $d_R(Y) < \mu$. Using the properties of relator metrics, the combination of these bounds shows that $d_R(X) = d_R(PYQ) < \mu + 2\epsilon$. Thus C is $(\mu + 2\epsilon)$-closed with respect to $\mathcal{R}_j$. □

Under minor restrictions it is possible to prove that the μ-closure of a finite $\mathcal{R}$-category with respect to $\mathcal{R}(k)$ is still finite. The existence of a finite μ-closure follows from the construction of a single finite μ-closed $\mathcal{R}$-category containing the original category. The first lemma shows that if this is the case, the μ-closure is also effectively constructible.

Lemma 13.26 *If B and C are finite $\mathcal{R}$-categories such that B is embedded in C, and C is μ-closed with respect to $\mathcal{R}(k)$, then the unique μ-closure of B in C with respect to $\mathcal{R}(k)$ is effectively constructible.*

Proof: Because of the existence of the $\mathcal{R}$-category C, the finiteness of the μ-closure is immediate. Also, a straightforward but inefficient procedure to find the unique minimum μ-closure of B is to list all of the $\mathcal{R}$-categories in C containing B and then to test to see which of them are μ-closed with respect to $\mathcal{R}(k)$. $\Box$

Lemma 13.27 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{12}$, and $3\beta + 2\epsilon \leq \alpha$. In addition let ϵ be a constant which satisfies Axiom 5 and also such that ϵ is an upper bound on the length of a rank i Zimin word, $i \leq k$, which is read in any general relator R having a rank strictly greater than i, as measured by the relator metric of R. If B is a finite $\mathcal{R}$-category which is $(1-3\alpha-2\epsilon)$-closed with respect to $\mathcal{R}(k)$, C is the α-closure of B with respect to $\mathcal{R}(k)$, and there is a bound on the diameters of the closed cones which are in C but not in B, then C is finite and effectively constructible.*

Proof: The bound on the diameters shows that there is also a bound on the length of rank j Zimin words, $j \leq k$, if each of the component pieces which are readable in a general relator of the appropriate rank is a geodesic in that relator. As shown in the proof of Lemma 13.25, every vertex in C differs from a vertex in B by a path of this type. Since C is collapsed, and the alphabet A is finite, this shows that the number of vertices in C is also finite. Finally, the fact that C is collapsed means that the category C is also finite. $\Box$

Lemma 13.28 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{12}$, and $3\beta + 2\epsilon \leq \alpha$ and $2\beta + \delta + 2\epsilon < \alpha$. In addition let ϵ be a constant which satisfies Axiom 5 and also such that ϵ is an upper bound on the length of a rank i Zimin word, $i \leq k$, which is read in any general relator R having a rank strictly greater than i, as measured by the relator metric of R. If X is a word which is $(1-3\alpha-2\epsilon)$-free with respect to $\mathcal{R}(k)$ then $str_k(X)$ is finite and effectively constructible. As a consequence, when $1-3\alpha-2\epsilon \geq \frac{1}{2}+\alpha$, $str_k(X)$ is finite and effectively constructible for all words X.*

Proof: The abstract path X is an $\mathcal{R}$-category which satisfies Lemma 13.27. The condition of the bounds of the diameters is satisfied because the final result is contained in the $\mathcal{R}$-category produced by a finite number of local closures. In each local closure the added closed cones have paths measuring $\alpha - 2\beta - 2\epsilon$ in the relator metric. Without loss of generality the path can be assumed to be a geodesic and thus by Axiom 4 it has a length of at least $\alpha - 2\beta - 2\epsilon - \delta$ of the length of the relator. Since the construction so far is simply connected and finite, there is a bound on the length of possible geodesics, which when divided by $\alpha - 2\beta - 2\epsilon - \delta > 0$ gives a bound on the length of the newly attached closed cones, and this in turn leads to a finite bound on the diameters of these closed cones. The induction through the local closures shows that the final result has a bound on the diameters of the closed cones in C which are not in B. Lemma 13.27 completes the proof. Finally, in general $\text{str}_k(X)$ is created using a Dehn reduction of the word X, and the inequality listed guarantees that the prior reasoning can be applied to the Dehn reduction. □

14 Finite Subgroups

The strong connection between torsion elements in G and automorphisms of general relators in $\mathcal{R}$, which was examined in Section 12, foreshadows an even stronger connection between the finite subgroups of G and the automorphism groups of the general relators in $\mathcal{R}$, when G is a general small cancellation group. In particular, under relatively mild conditions, every finite subgroup of the group is contained in the automorphism group of some general relator. The converse, that every subgroup of an automorphism group occurs as a finite subgroup, is an easy consequence of the fact that the Cayley category is proper. At the conclusion of the section Theorem A is shown.

14.1 Equivalent Conditions on Finite Subgroups

Lemma 14.4 establishes a set of four equivalent conditions on finite subgroups of general small cancellation groups. Eventually (Lemma 14.17) it will be shown that all of the finite subgroups satisfy these conditions, but for the moment, it will merely be shown that the conditions are equivalent. One of the implications in the proof is proved using a well-known fixed-point lemma for finite groups acting simplicially on finite contractible simplicial complexes, which is given below.

Lemma 14.1 *Let H be a finite group which acts simplicially on a finite contractible simplicial complex B. Then the action of H on B must have a global fixed point iff H is a finite group which is a q-group extended by a cyclic group extended by a p-group, where p and q are primes.*

Proof: See [14]. □

The following technical lemma will be needed:

Lemma 14.2 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{12}$, and let R be a general relator in $\mathcal{R}$. If X is a path in R such that $d_R(X) \geq \frac{1}{2} + \delta$, then either there exists an initial segment U of X such that all paths V in R with the same endpoints as U satisfy $d_R(V) \geq 3\alpha$, or else every edge in R is long.*

Proof: Since $d_R(X) \geq \frac{1}{2} + \delta$, by Axiom 4 $|X|_R \geq \frac{1}{2}$. Since geodesic distance changes at most 1 unit when a word is shortened by a single letter, there must be an initial segment U of X such that $|U|_R$ is at most $\frac{1}{2|R|}$ from $\frac{1}{2}$. Thus $\frac{1}{2} - \frac{1}{2|R|} \leq |U|_R \leq \frac{1}{2} + \frac{1}{2|R|}$. Given any other path V in R between the same endpoints, UV^{-1} forms a loop. If the winding number of this loop is 0, then U and V are homotopic in ∂R and consequently $|V|_R \geq |U|_R$ by definition. If, on the other hand, the winding number is $k \geq 1$, then $|U|_R + |V|_R \geq k \geq 1$. In either case $|V|_R \geq \frac{1}{2} - \frac{1}{2|R|}$. If $|R| \geq 3$, then this implies that $|V|_R \geq \frac{1}{3}$, and that $d_R(V) \geq \frac{1}{3} - \delta \geq 3\alpha$ by Axiom 4 and by assumption. If, on the other hand, $|R| < 3$, then let V be a single edge. By the definition of the normalized graph metric, $|V|_R \geq \frac{1}{|R|} \geq \frac{1}{3}$, so that $d_R(V) \geq \frac{1}{3} - \delta \geq 3\alpha$ by Axiom 4 and by assumption. This shows that every edge in R is long. □

Lemma 14.3 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{12}$ for which str(W) is finite and effectively constructible for all words $W \in A^*$. In addition, let H be a finite subgroup of G, let C be the Cayley category of G, and let B be a finite subcategory of C. If the natural action of H on C restricts to an action of H on B, then there is a finite, contractible, and effectively constructible subcategory B' containing B on which H acts.*

Proof: If B is not connected then the finite number of components can be connected with a finite number of simplicial paths in the 1-skeleton. In order to extend the action of the group H to the resulting $\mathcal{R}$-category, let B_1 be the union of B and the finite number of images of each of the connecting paths under the automorphisms of H. Next, consider all pairs of vertices $u, v \in B_1$. If u and v are viewed as vertices in C then there is a geodesic path X from u to v. By assumption str(X) exists, it is finite, and it is effectively constructible. Moreover, str(X) can be viewed as the α-closure of X in C with respect to $\mathcal{R}$, and by Lemma 13.12 there is a unique embedding of str(X) which sends the startpoint to u and the endpoint to v. The uniqueness of str(X) guarantees that the image of str(X) under $h \in H$ is the α-closure of the image of X with respect to $\mathcal{R}$ as well. Let B_2 be a subcategory of C obtained by adding to B_1 such a straightline construction for every pair of vertices u and v in B_1. Since a str(X) construction is added between u and v iff it is also added between $h(u)$ and $h(v)$ for all $h \in H$, the action of the group H on B_1 extends to B_2.

The key step in the proof involves showing the B_2 is $(1 - 3\alpha)$-closed with respect to $\mathcal{R}$. Consider a path V in B_2 which is readable in a general relator T with $d_T(V) \geq 1 - 3\alpha$. By Lemma 14.2 there is a subword U such that every path in T between the same endpoints as U measures at least 3α. Let u and v be the startpoint and endpoint of U respectively. Since u is in B_2 there must be a path X in B_2 which starts and ends in B_1 such that u is contained in the embedding of str(X) determined by the endpoints of X. Similarly, there must be a path Y in B_2 which starts and ends in B_1 such that v is contained in the embedding of str(Y) determined by the endpoints of Y. Let u' be one of the endpoints of X and let v' be one of the endpoints of Y. Since str(X) is connected and α-closed with respect to $\mathcal{R}$, there is a geodesic X' in str(X) from u to u'. In particular, by Lemma 13.13 $\mathrm{str}(X')$ embeds in $\mathrm{str}(X)$ and thus is contained in B_2. Similarly, since $\mathrm{str}(Y)$ is connected and α-closed with respect to $\mathcal{R}$, there is a geodesic Y' in $\mathrm{str}(Y)$ from v' to v. In particular, by Lemma 13.13 $\mathrm{str}(Y')$ embeds in $\mathrm{str}(Y)$ and thus is contained in B_2. Next, notice that u' and v' are in B_1 so that if Z is a geodesic from u' to v', $\mathrm{str}(Z)$ is contained in B_2 as well. Since the paths X and Y will no longer be needed, relabel X' and Y' as X and Y. The construction so far has produced a path XZY from u to v such that $\mathrm{str}(X)$, $\mathrm{str}(Z)$ and $\mathrm{str}(Y)$ are contained in B_2.

If XZ is not a geodesic in C, then there is a path W between the same endpoints of strictly shorter length. But then by Lemma 11.8, W is contained in the union of $\mathrm{str}(X)$ and $\mathrm{str}(Z)$ and W can be split into X' and Z' such that X' is readable in $\mathrm{str}(X)$ and Z' is readable in $\mathrm{str}(Z')$. By Lemma 13.13, $\mathrm{str}(X')$ and $\mathrm{str}(Z')$ are embedded in $\mathrm{str}(X)$ and $\mathrm{str}(Z)$ and thus in B_2. This produces a strictly shorter word $X'Z'Y$ from u to v such that $\mathrm{str}(X')$, $\mathrm{str}(Z')$ and $\mathrm{str}(Y)$ are in B_2. After this process is repeated at most a finite number of times (on XZ and on ZY alternately), and after suitable relabeling, a path XZY is obtained with $\mathrm{str}(X)$, $\mathrm{str}(Z)$, and $\mathrm{str}(Y)$ contained in B_2, and with the additional property that XZ and ZY are geodesics in C.

If XZY is $(1-3\alpha)$-free then by Lemma 11.3 it is readable in T, and by Lemma 14.2 $d_T(XYZ) \geq 3\alpha$. Thus one of the three subpaths, say X, is long in T and since $\mathrm{str}(X)$ is α-closed, T is contained in $\mathrm{str}(X)$ and thus in B_2. If, however, XZY is not $(1-3\alpha)$-free, then there exists a general relator R and a path P in XZY with $d_R(P) \geq 1-3\alpha$. Since XZ and ZY are geodesics and thus $(\frac{1}{2}+\alpha)$-free, it follows that P contains a portion of X and Y. Thus let $P = X_1ZY_1$. Since $\alpha \leq \frac{1}{12}$, $d_R(X_1) \geq 2\alpha$, and $d_R(Y_1) \geq 2\alpha$. In particular R is contained both in $\mathrm{str}(X)$ and in $\mathrm{str}(Y)$. This in turn implies that the path XZ lies completely in $\mathrm{str}(X)$. If XZ is relabeled as X, then XY is a path from u to v such that $\mathrm{str}(X)$ and $\mathrm{str}(Y)$ are contained in B_2, and both X and Y are geodesics. Applying the reduction argument described above to the path XY shows that without loss of generality, we can assume that XY is a geodesic in C. At this point, by Lemma 11.3 XY is readable in T, and by Lemma 14.2 $d_T(XY) \geq 3\alpha$. Thus one of the two subpaths, say X, is long in T and since $\mathrm{str}(X)$ is α-closed, T is contained in $\mathrm{str}(X)$ and thus in B_2. In either case T must be contained in B_2, so that B_2 is $(1-3\alpha)$-closed with respect to $\mathcal{R}$.

At this point it has been shown that B_2 is connected and that it is $(1-3\alpha)$-closed with respect to $\mathcal{R}$. Since B_2 is a subcategory of C, it is also collapsed. From Lemma 12.1, it follows that B_2 is simply connected, and then by Lemma 12.13, B_2 is contractible. Thus B_2 satisfies the conditions stated in the lemma. □

Lemma 14.4 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{12}$ in which str(W) is finite and effectively constructible for all words $W \in A^*$, and let C be the Cayley category of the presentation. In addition suppose that all general relators in $\mathcal{R}$ contain at least one crucial cone. If H is a finite subgroup of G then the following conditions are equivalent:*

(1) the natural action of H on C has a global fixed point

(2) H is isomorphic to a subgroup of the automorphism group of some general relator $R \in \mathcal{R}$, or the automorphism group of a labeled edge

(3) H is isomorphic to a finite group which is either a cyclic or a dihedral group extended by a 2-group

(4) H is isomorphic to a subgroup of a finite group which is a q-group extended by a cyclic group extended by a p-group, where p and q are primes.

Proof: $(1 \Rightarrow 2)$ Let $C' = \text{Chain}(C)$ be the simplicial subdivision of C. The action of H on C can then be viewed as a simplicial action of H on C' which preserves the heights of the vertices in C'. If the action of H on C' has a global fixed point, then that point must occur in the interior of a unique simplex. Since the vertices of the simplex are ordered by their height, the elements of H cannot permute them. Thus the entire simplex is fixed by H and without loss of generality the global fixed point of H can be chosen to be a vertex of C'. Call this vertex c and consider the slice category C/c. Since the action of H on C preserves height, H must act on the image of the slice category in C and by Lemma 12.3 H acts on C/c itself. Notice that if c has height 0 then H must be trivial, if c has height 1 then C/c is a labeled edge, and if c has a height at least 2 then C/c is a general relator in $\mathcal{R}$. In all three cases, the second condition is satisfied.

$(2 \Rightarrow 3)$ This follows immediately from Lemma 5.26, and the fact that the automorphism group of a labeled edge is at most cyclic order 2. $(3 \Rightarrow 4)$ Since a dihedral group can be viewed as a group of order 2 extended by a cyclic group, this step also follows immediately. $(4 \Rightarrow 1)$ By Lemma 14.3, H acts on a finite contractible subcategory B of C. Thus, by Lemma 14.1 the action of H on B has a global fixed point. Since this global fixed point is also a global fixed point of the action of H on C, the proof is complete. □

Lemma 14.5 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{12}$ in which str(W) is finite and effectively constructible for all words $W \in A^*$, and let H be a finite subgroup of G. In addition suppose that all general relators in $\mathcal{R}$ contain at least one crucial cone. If H is a p-group for some prime p, then H satisfies all of the equivalent conditions listed in Lemma 14.4. In particular this is true when H is a 2-group.*

Proof: The result is immediate since by assumption H satisfies condition (4) of Lemma 14.4. □

Another easy application of Lemma 14.4 uses the solvability of groups of odd order to show that all odd-order subgroups of finitely presented general small cancellation groups are cyclic. Later in this section a more detailed examination of the finite subgroups of general small cancellation groups will provide another proof of this fact, so that the results collected in Theorem A remain independent of the Odd Order Theorem.

Lemma 14.6 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $\alpha \leq \frac{1}{12}$ in which str(W) is finite and effectively constructible* *for all words $W \in A^*$, and let H be a finite subgroup of G. In addition suppose that all general relators in $\mathcal{R}$ contain at least one crucial cone. If H is a finite subgroup of G of odd order then the group H is cyclic.*

Proof: By the Odd Order Theorem, H is a solvable group and thus has a composition series whose factors are cyclic groups of prime order. The proof proceeds by induction on the length of the composition series. If the composition series has length 1 then H is cyclic and there is nothing to prove. Next, assume that the result has been shown for all finite subgroups with composition series of length k and let H have a composition series of length $k+1$. By induction, H has a cyclic normal subgroup of prime index. Since H satisfies condition (4) of Lemma 14.4, H is also isomorphic to a subgroup of a finite group which is a dihedral extended by a 2-group. In particular, there is a group homomorphism from H to a dihedral whose kernel is a 2-group. But since H has odd order, the kernel of the map is trivial and H is isomorphic to a subgroup of the dihedral group. Moreover, the absence of elements of order 2 shows that H injects into the cyclic subgroup of the dihedral and H itself is thus cyclic. This completes the induction and the proof. □

14.2 Inequalities and Estimates

The following lemmas are a collection of inequalities and estimates which will be needed in the proof of Lemma 14.12.

Lemma 14.7 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \le \frac{1}{10}$ and let R be a general relator in $\mathcal{R}_k$. If U and V are words which are conjugate in $G(k-1)$ and U^n and V^n are both readable in R as loops with non-zero winding numbers, then $|d_R(U) - d_R(V)| \le 2\gamma + 2\delta < 2\alpha$. In particular, if U and V are arbitrary representative cycles of R then $|d_R(U) - d_R(V)| \le 2\gamma + 2\delta \le \alpha$.*

Proof: Without loss of generality, assume that U is a geodesic in R and let W be a cycle conjugate to U in $G(k-1)$ and of minimal length subject to this restriction. First of all, it is clear that $|U| \ge |W|$ and that the cycle of W is $(\frac{1}{2} + \alpha)$-free with respect to $\mathcal{R}(k-1)$, for otherwise Lemma 10.3 would yield a contradiction to the minimality of W. Next, by Lemma 11.4 W^k is readable as a loop in R with the same winding number as U^k. Finally, notice that the automorphism of R induced by sending the initial vertex of one of the paths U to its endpoint will also send the initial vertex of each of the vertices of each of these paths W to their endpoints.

By the definition of ω_R, there is a path of length at most ω_R from the initial vertex of one of the paths U to a point in the loop W^k. Without loss of generality let W be the conjugate of W determined by the endpoint of this path. The image of this path under the automorphism determined by U sends it to another path between U^k and W^k. Since U is a geodesic it is clear that $|U| \le |W| + 2\omega_R$. Combining these inequalities shows that $|U|$ is within ω_R of $|W| + \omega_R$. Since the word V must also satisfy these restrictions, the length of U and the length of V differ by at most $2\omega_R$, and consequently, $||U|_R - |V|_R| \leq 2\gamma$ by Axiom 3. Using Axiom 4 and Axiom 7, $|d_R(U) - d_R(V)| \leq 2\gamma + 2\delta < 2\alpha$. Finally, let U and V be representative cycles and let P be a path in R between their basepoints. Depending on the relative orientations of U and V either $PUP^{-1}V$ or $PUP^{-1}V^{-1}$ is a loop of winding number 0 and thus by Lemma 7.12 U is conjugate to either V or its inverse. The situation thereby reduces to the one described above. □

Lemma 14.8 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation with $\alpha \leq \frac{1}{12}$. If g is a non-trivial torsion element in G then there exist a word U and a general relator $R \in \mathcal{R}_k$ such that $2\alpha \leq \frac{1}{3} - 2\gamma - 2\delta \leq d_R(U) \leq \frac{1}{2} + \alpha$, the cycle of U is $(1-4\alpha)$-free with respect to $\mathcal{R}$, U represents a power of g as an automorphism of the Cayley category obtained from repeatedly squaring g, and a power of U is a non-trivial loop in R. In addition, all $(1-4\alpha)$-free words which are conjugate to U in $G(k-1)$ and readable in R satisfy the same inequalities as U.*

Proof: Let $n > 1$ be the order of g in G. If u and v are vertices of the Cayley category of G such that the automorphism g sends u to v, then by Lemma 6.3 u and v must be distinct. Let X be a path from u to v. Since g has order n, X^n is readable as a loop in the Cayley category. Next, by Lemma 12.7 there is a word Y and a general relator R such that Y^n is readable in R as a loop with non-zero winding number, with Y cyclically reduced in the free group and $(\frac{1}{2}+\alpha)$-free with respect to $\mathcal{R}$. In particular, $d_R(Y) \leq \frac{1}{2} + \alpha \leq \frac{2}{3}$, since $\alpha \leq \frac{1}{6}$. On the other hand, since Y^n has a non-zero winding number by the properties of relator metrics $d_R(Y^n) \geq 1$, and $d_R(Y^{ni}) \geq i$. Then by properties 4 and 5, $nd_R(Y^i) \geq i$. This shows that by choosing i large enough, it is possible to make $d_R(Y^i)$ as large as desired. In particular, there is an i so that $d_R(Y^i) \geq \frac{1}{3}$. Let j be the smallest non-negative integer such that $d_R(Y^i) \geq \frac{1}{3}$ where $i = 2^j$. If $d_R(Y) \geq \frac{1}{3}$ then $j = 0$. If j is positive, then $i = 2l$ and by the properties of relator metrics $d_R(Y^i) \leq d_R(Y^l) + d_R(Y^l) \leq \frac{2}{3}$. The two instances of a path reading Y^l must have the same measure by property 6, in so far as they measure at least $\frac{1}{6} \geq \alpha$ and thus differ by an automorphism of the general relator R.

Let V be the geodesic in R between the endpoints of Y^i. It will be shown that a power of V is readable in R as a loop with a non-zero winding number. If Y^iV^{-1} is a loop in ∂R of winding number 0, then $d_R(V) = d_R(Y^i)$ by the properties of relator metrics and it is clear that a power of V is readable in R as a loop with non-zero winding number. On the other hand, if $d_R(V) \neq d_R(Y^i)$ then the winding number of Y^iV^{-1} must be at least 1. If it is more than 1 then $d_R(V)$ is at least $\frac{4}{3}$ and $|V| \geq (\frac{4}{3} - \delta)|R| > (\frac{2}{3} + \delta)|R| \geq |Y^{2^i}|$ by Axiom 4. Since this implies that V is not a geodesic, it must be the case that Y^iV^{-1} has winding number exactly 1. Next, let m be the winding number of Y^{ni}. Since $d_R(Y^{ni}) \leq nd_R(Y^i) \leq n\frac{2}{3} < n$, it follows by property 5 of relator metrics that $m < n$. Moreover, by successively removing Y^iV^{-1} it is clear that the winding number of $Y^{ni}(V^{-1})^n$ is n once it is noted that each of these Y^iV^{-1} loops must be similarly oriented. The reason for this is that if successive copies of Y^i have opposite orientations then the symmetry of the situation would require that the winding number of Y^{ni} be 0, contradiction. Finally, since the winding number $Y^{ni}(V^{-1})^n$ is n and the winding number of Y^{ni} is $m < n$, the winding number of V^n is $n - m \neq 0$, which was to be shown.

Since V is a geodesic in R, by Lemma 10.3 $d_R(V) < \frac{1}{2} + \alpha$, and by construction $d_R(V) \geq \frac{1}{3}$ either because $d_R(V) = d_R(Y^i) \geq \frac{1}{3}$ or because $d_R(V) \geq 1 - d_R(Y^i) \geq 1 - \frac{2}{3}$. Let R be a general relator in $\mathcal{R}_k$. By Lemma 10.5, V is conjugate to a word U whose cycle is $(\frac{1}{2} + \alpha)$-free with respect to $\mathcal{R}(k-1)$ and with $|V| \geq |U|$. By Lemma 11.4, U^n is readable as a loop in ∂R and it is homotopic to the loop V^n. In particular it has the same non-trivial winding number. Finally, by Lemma 14.7, $\frac{1}{3} - 2\alpha \leq d_R(U) \leq \frac{1}{2} + 3\alpha$. If $d_R(U)$ is greater than $\frac{1}{2} + \alpha$ then U can be replaced by a geodesic in R between the same endpoints and the arguments of the last two paragraphs can be repeated. Since the lengths of the words under consideration are constantly shortening, a finite number of iterations yields a word U with $\alpha \leq \frac{1}{3} - 2\alpha \leq d_R(U) \leq \frac{1}{2} + \alpha$ which satisfies the requirements of the lemma. The final statement of the lemma follows from the fact that the other words conjugate to U in $G(k-1)$ and readable in R are conjugate to Y^i as well, so that the above argument can be repeated for these words. □

Lemma 14.9 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and suppose that $W = UV$ is a representative of a general relator R in $\mathcal{R}$. If U is long, negatively oriented with respect to W, and a geodesic path in its homotopy class in ∂R, then $|V| \geq |U| + (1-2\gamma)|R|$ and $d_R(V) \geq d_R(U) + 1 - 2\gamma - 2\delta$.*

Proof: Lift the word UV to R^∞, the universal cover of the boundary of R. Since U is long and negative, the ball of radius ω_R centered at the start vertex of U is strictly between the balls of radius ω_R centered at the end of U (which equals the start of V) and at the end of V. Since by the definition of ω_R these balls disconnect R^∞, the path V must contain a vertex within ω_R units of the start of U. Let P be the path from the start of U to such a vertex and let $V = V'V''$ be a split of V determined by this point. Without loss of generality, assume that U is a geodesic in R^∞. Thus $|V'| + |P| \geq |U|$. On the other hand, since PV'' is a representative of R when it is read in ∂R, it follows that $|V''| + |P| \geq |R|$. Thus $|V| + 2|P| > |U| + |R|$, and since $|P| < \gamma|R|$ by Axiom 3, $|V| > |U| + (1-2\gamma)|R|$. Using Axiom 4, $d_R(V) > d_R(U) + 1 - 2\gamma - 2\delta$. □

Lemma 14.10 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and suppose that $W = U_1V_1U_2V_2$ is a representative of a general relator R in $\mathcal{R}$. If U_i is long, negatively oriented with respect to W, and a geodesic path in its homotopy class in ∂R for $i = 1, 2$, then $|V_1| + |V_2| \geq |U_1| + |U_2| + (1-4\gamma)|R|$ and $d_R(V_1) + d_R(V_2) \geq d_R(U_1) + d_R(U_2) + 1 - 4\gamma - 4\delta$.*

Proof: Let X and Y be the paths of shortest length in the homotopy classes of U_1V_1 and U_2V_2 in ∂R, respectively. Without loss of generality, assume that $d_R(X) \geq d_R(Y)$. Since XY is a representative, $2d_R(X) \geq d_R(X) + d_R(Y) \geq 1$, so that $d_R(X) \geq \frac{1}{2} > \alpha$, and X is long. Thus by Lemma 9.9 X is oriented. If

X is negatively oriented with respect to XY, then by Lemma 14.9 $|Y| - |X| \geq (1-2\gamma)|R|$. Since by assumption $d_R(X) \geq d_R(Y)$, by Axiom 4 $2\delta|R| \geq |Y| - |X| \geq (1-2\gamma)|R|$, contradiction. Thus X is positively oriented as well as long.

Next, lift the path W to R^∞, and let u_i, v_i be the starting vertex of U_i, V_i respectively, for $i = 1, 2$. The fact that X is long and positively oriented while U_1 is long and negatively oriented with respect to the cycle W in R^∞ implies that the balls of radius ω_R centered at v_1, u_1 and u_2 (in R^∞) are disjoint and that they occur in that order. Specifically, since by the definition of ω_R, these balls disconnect R^∞, the path V must pass through the ball centered at u_1. Thus there exists a path P from u_1 to a point in V with $|P| \leq \gamma|R|$. Let $V = V'V''$ be the partition of V determined by this endpoint. Since U_1 is a geodesic in R^∞, $|P| + |V'| \geq |U_1|$. Also, since the image of U_2V_2PV'' in ∂R is a representative of R which satisfies the hypotheses of Lemma 14.9 with U_2 playing the role of U, and V_2PV'' playing the role of V, it follows that $|P|+|V''|+|V_2| \geq |U_2|+(1-2\gamma)|R|$. Combining these two inequalities and using the fact that $|P| \leq \gamma|R|$ yields $|V_1| + |V_2| \geq |U_1| + |U_2| + (1-4\gamma)|R|$. Finally, applying Axiom 4 shows that $d_R(V_1)+d_R(V_2) \geq d_R(U_1)+d_R(U_2)+1-4\gamma-4\delta$. □

Lemma 14.11 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and suppose that $W = XUYU^{-1}$ is a representative of a general relator T in $\mathcal{R}$. If both instances of U are properly oriented with respect to W, $d_T(U) \geq \alpha$, and there exists a relator R in $\mathcal{R}$, distinct from T, such that $d_R(X) \geq \alpha$ and X^n is readable as a loop in R with a non-zero winding number, then there exists a word Z such that the cycle $XZX^{-1}Z^{-1}$ is a representative of T and corresponding to the two readings of the word X in this cycle there are two instances of the relator R (say R_1 and R_2) contained in the boundary of T. Both instances of the word X represent the same orientation-reversing automorphism of T which fixes both R_1 and R_2. Moreover, the shortest path between these two instances of R has a length of at least $(\frac{1}{2} - \beta - \delta)|T|$ units.*

Proof: First, by Axiom 6 there is a word V such that $XVYV^{-1}$ is readable as a contractible loop in ∂T. Since the reading of $XVYV^{-1}$ can be chosen so that it extends the reading of Y given by W (because the axiom is symmetric with respect to X and Y), there is another loop in ∂T which reads $XUV^{-1}X^{-1}VU^{-1}$ which is a representative since it is a combination of a loop with winding number 1 and a contractible loop. Choosing $Z = UV^{-1}$ completes the first proof. To distinguish between the two copies of Z and X, let Z_1 be the copy of Z which starts and ends at the start of the readings of X, with Z_2 being the other, and let X_1 be the reading of X which starts at the start of Z_1, with X_2 being the other. Next, by Axiom 1 the readings of X_1 and X_2 can be extended to embeddings of R in T, and since R and T are distinct, the copies of R must be contained in ∂T. Moreover, the extension can be arranged so that the reading of X in T is also in the loop X^n in R, simply by choosing this reading of X in R for use in the application of Axiom 1. Let R_1 and R_2 be the instances of R in T associated with X_1 and X_2 respectively.

Using Axiom 2 and the properties of relator metrics, it is clear that $d_T(X_i) < \beta$ for $i = 1, 2$ and that therefore $d_T(Z_i) > \frac{1}{2} - \beta$ for either $i = 1$ or $i = 2$ and thus for both $i = 1$ and $i = 2$ by another property of relator metrics and the fact that $\frac{1}{2} - \beta > \alpha$. As a technical aside, notice that the two instances of X must have the same measure only because they are long in T. As a consequence there is an automorphism g of T which sends Z_1 to Z_2, which is also the automorphism induced by either X_1 or X_2. Clearly the paths X_1 and X_2 cannot be loops themselves since this would mean that Z_1 and Z_2 correspond to the same reading in T and that at least one of the X loops would have a non-trivial winding number in T. This in turn, by a property of relator metrics and Lemma 9.10, would imply the identity of R and T contrary to the hypothesis in the statement of the lemma. Thus X_1 and X_2 are not loops and Lemma 12.5 can be applied to show that g is also an automorphism of both R_1 and R_2, so that they remain fixed under the action of g.

To prove the final claim, let P be the shortest path from R_1 to R_2, let Q_1 be a path from the start of P to the start of Z_1, and let Q_2 be a path from the end of P to the end of Z_1. Since R_1 and R_2 are stabilized by g, the image of the loop $Q_1^{-1}PQ_2Z_1^{-1}$ is a loop $Q_1^{-1}PQ_2Z_2^{-1}$ where the new Q_1 and the new Q_2 are read in R_1 and R_2 respectively. Combining these loops with the loop $X_1Z_2X_2^{-1}Z_1^{-1}$ produces a loop $Q_1X_1Q_1^{-1}PQ_2X_2^{-1}Q_2^{-1}P^{-1}$ where $Q_iX_iQ_i^{-1}$ is read in R_i, $i = 1, 2$. If this loop has a non-zero winding number then by the same reasoning as above, $d_T(Q_iX_iQ_i^{-1}) < \beta$ for $i = 1, 2$ and consequently $d_T(P) > \frac{1}{2} - \beta$. Thus by Axiom 4 $|P| > (\frac{1}{2} - \beta - \delta)|T|$. Therefore it only remains to show that the winding number of this loop is non-zero.

Suppose that one of the readings of Z, say Z_1, is negatively oriented with respect to the loop $X_1Z_2X_2^{-1}Z_1^{-1}$. Then by Lemma 14.9, $d_T(X_1) + d_T(Z_2) + d_T(X_2) \geq 1 + d_T(Z_1) - 2\gamma - 2\delta$. Since by the properties of relator metrics $d_T(Z_1) = d_T(Z_2)$ and since $d_T(X_i) < \beta, i = 1, 2$, it follows that $2\beta \geq 1 - 2\gamma - 2\delta$ which contradicts the constraints in Axiom 7. Thus both instances of Z must be positively oriented with respect to the loop. Because Z_1 and Z_2 are facing opposite directions within the loop itself, this means that the winding numbers of the loops $Q_1^{-1}PQ_2Z_1^{-1}$ and $Q_1^{-1}PQ_2Z_2^{-1}$ will cancel each other out, so that the loop $Q_1X_1Q_1^{-1}PQ_2X_2^{-1}Q_2^{-1}P^{-1}$ still has a winding number of 1. This completes the proof. □

14.3 Elements of Maximum Rank

The long proof which follows provides the key technical lemma needed to show that all finite subgroups of a general small cancellation group are contained in the automorphism group of some general relator in the presentation.

Lemma 14.12 *Let $G = \langle A|\mathcal{R}\rangle$ be a graded general small cancellation presentation with $2\beta + 2\gamma + \delta \leq \alpha \leq \frac{1}{12}$, let H be a finite subgroup of G, and let C be the Cayley category of the presentation. If $g \in H$ is an element of maximum rank (rank(g) $\geq$ rank(h) for all $h \in H$), and g rotates a general relator $R \in \mathcal{R}$,* *then for all $h \in H$ there is an instance of R in C which is fixed by both g and h.*

Proof: Given non-identity elements $g, h \in H$ of orders n and m respectively, then by Lemma 10.5 there are words X and Y which are cyclically Dehn-reduced with respect to $\mathcal{R}$ and which represent g and h respectively. By Lemma 12.7 the words X^n and Y^m are readable as loops with non-zero winding numbers in uniquely determined general relators R and S respectively. By assumption g has been chosen so that the rank of its associated general relator R is maximal among all of the general relators associated to the elements of H in this way. The proof below will show that each of the elements of H stabilizes an instance of R in the Cayley category which is stabilized by g. The action of g on the Cayley category of G will stabilize some instances of the general relator R in the Cayley category and possibly will not stabilize other instances. Let R_1 be an instance of R in C which is stabilized by g and let S_1 be an instance of S in C which is stabilized by h. The proof is divided into two stages. First it will be shown that R_1 and S_1 can be chosen so that they have a vertex in common. Next it will be shown that the action of h stabilizes the chosen R_1.

Claim 1: R_1 and S_1 can be chosen so that $R_1 \cap S_1 \neq \emptyset$.

Let R_1 and S_1 be chosen so that they are as close to each other in the graph metric as possible, and assume for the moment that $R_1 \cap S_1 = \emptyset$, and assume that P is the shortest possible path in the Cayley category connecting a vertex u' in R_1 to a vertex v' in S_1. By assumption, P is non-empty. Next, by Lemma 14.8 there are powers of g and h, say g^i and h^j, as well as words U and V conjugate to X^i and Y^j respectively, such that these words satisfy the conclusions of Lemma 14.8. In particular, powers of these words are readable as loops in R and S with non-zero winding numbers. Let u and $g^i(u)$ be the endpoints of a path reading U in R_1, and let Q be a path from u to u'. Then $Q^{-1}UQ$ is a path from u' to $g^i(u')$ in R_1. Let U' be a geodesic path in the boundary of R_1 between u' and $g^i(u')$ which lies in the same homotopy class as $Q^{-1}UQ$. Thus together these paths form a loop in R_1 of winding number 0, so that by Lemma 7.12 and Lemma 7.9 U and U' are conjugate in $G(k-1)$. Then, by Lemma 11.4 U^n and $(U')^n$ are both readable in R_1 as loops with the same non-zero winding number, and by Lemma 14.8 U' also satisfies the same conditions as U. Similarly choose V' as a geodesic path in the boundary of S_1 between v' and $h^j(v')$. By the same reasoning V^m and $(V')^m$ are readable in S_1 as loops with the same non-zero winding number, and V' satisfies the conclusions of Lemma 14.8. Thus for convenience U' and V' can be relabeled as U and V since they satisfy the same conditions and the original U and V will no longer be needed in the proof.

Let U, V and the two paths labeled P be called the sides of the cycle $UPV^{-1}P^{-1}$. The first task will be to show that this cycle is 3α-complement-free with respect to $\mathcal{R}$. If this were not true then there would be a word W, a word Z, and a general relator T such that W is a subword of a cyclic conjugate of $UPV^{-1}P^{-1}$, WZ is a representative of T, and $d_T(Z) < 3\alpha$. A case-by-case analysis will show that no such words W and Z can exist. The cases will be divided by the number of sides needed to contain the word W.

Case 1: First of all, the subword W cannot be completely contained in one of the sides labeled P since P is a geodesic in the Cayley category, $1-3\alpha > \frac{1}{2}+\alpha$, and by Lemma 10.3, geodesics are $(\frac{1}{2}+\alpha)$-free with respect to $\mathcal{R}$. Similarly the subword W cannot be completely contained within a side labeled U or V since this would imply that the general relator T is contained in either R or S respectively, and a similar contradiction arises due to the fact that U and V are geodesics in the boundaries of R and S.

Case 2: Next, suppose that W is a subword of two consecutive sides in the cycle $UPV^{-1}P^{-1}$. Without loss of generality assume that W is a subword of UP since the other possibilities can be handled analogously. Let $W = W_1W_2$ where W_1 is the portion of W in U and W_2 is the portion in P. If $d_T(W_1) \geq \alpha$ then by Axiom 1 the instance of the general relator T which contains the path W is contained in the instance of R containing the path U. In particular, this would mean that the path P between S_1 and R_1 would not be disjoint from R_1 except at its endpoint, contradiction. If, on the other hand, $d_T(W_1) < \alpha$ then by the properties of relator metrics, $d_T(W_2) > 1-4\alpha \geq \frac{1}{2}+\alpha$, and this contradicts the assumption that P is a geodesic.

Case 3A: Suppose that W is a subword of three consecutive sides. If the middle of these three sides is labeled P, then without loss of generality assume that W is a subword of UPV^{-1} and write $W = W_1PW_2$ in the obvious way. If either $d_T(W_1) \geq \alpha$ or $d_T(W_2) \geq \alpha$ then a contradiction to the disjointness of P from S_1 and R_1 is obtained as above. If on the other hand both measures are less than α, then $d_T(P) > 1-5\alpha \geq \frac{1}{2}+\alpha$ and P is not a geodesic in the Cayley category.

Case 3B: If the middle of the three sides is labeled either U or V then without loss of generality assume that W is a subword of $P^{-1}UP$ and write $W = W_1UW_2$ in the obvious way. Since W contains all of U and $d_R(U) \geq 2\alpha$, it follows from Axiom 1, Lemma 2.1, and Axiom 2 that R is contained in T, that the rank of R as a cone complex is less than that of T (since equality leads to a contradiction of the disjointness of P and R_1), and that $d_T(U) < \beta$. If either $d_T(W_1) < \alpha$ or $d_T(W_2) < \alpha$ then the measure of the other is more than $1-4\alpha-\beta$ and this contradicts the fact that P is a geodesic. So assume that both ends measure at least α with respect to the general relator T. If W_1 is negatively oriented with respect to the loop W_1UW_2Z, then, by applying Lemma 14.9 to this loop, $d_T(U) + d_T(W_2) + d_T(Z) \geq 1 + d_T(W_1) - 2\gamma - 2\delta$. Since $d_T(U) < \beta$, $d_T(Z) < 3\alpha$, and $d_T(W_1) \geq \alpha$, it follows that $d_T(W_2) > 1-2\alpha-\beta-2\gamma-2\delta \geq 1-4\alpha > \frac{1}{2}+\alpha$, which contradicts the fact that P, and hence W_2, is a geodesic. Thus W_1 can be assumed to be positively oriented with respect to the loop W_1UW_2Z. An analogous argument shows that W_2 must also be positively oriented with respect to the loop. Again without loss of generality assume that $|W_1| \leq |W_2|$, so that $W_2 = W_1^{-1}W_3$ for some possibly empty word W_3. Lemma 14.11 can then be applied to the representative WZ with W_1^{-1} playing the role of U, and U and W_3Z standing in for X and Y respectively.

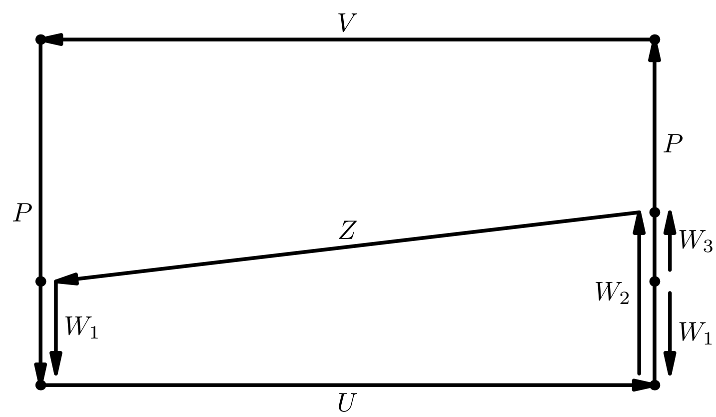

Figure 32: An illustration of Case 3B

Thus there exists another instance of R, called R_2, in T which is fixed by g and which is at least $(\frac{1}{2} - \beta - \delta)|T|$ units away from R_1.

Let u be the start vertex of one of the paths U in the loop U^n in R_2. Since WZ is a representative of T, by Axiom 3 there is a path from u to WZ whose length is at most $\gamma|T|$. Let Q be one such path. If the endpoint of Q lies is W then it must fall within either W_1 or W_2 and at least $(\frac{1}{2} - \beta - \gamma - \delta)|T| > \gamma|T|$ away from U. This shows that S_1 is closer to R_2 than to R_1. Even if the endpoint of Q lies in Z, the length of Z is at most $(3\alpha + \delta)|T|$ by Axiom 4 so that this endpoint lies at most $(\frac{3}{2}\alpha + \frac{1}{2}\delta + \gamma)|T| \leq 2\alpha|T|$ from a point in one of the instances of P. By Lemma 14.11, this point in P must be at least $(\frac{1}{2} - 2\alpha - \beta - \delta)|T| \geq 2\alpha|T|$ so that S_1 is closer to R_2 than to R_1, contradiction. Notice that if the middle of the three sides covered by W had been labeled V instead of U, then W would have been a subword of $P^{-1}UP$, and the conclusion would have been that there is an instance S_2 of S which is stabilized by h and R_1 is closer to S_2 than to S_1.

Case 4: The next possibility is that W is a subword of four consecutive sides. Without loss of generality, assume that W is a subword of $P^{-1}UPV^{-1}$ and write $W = W_1UPW_2$ in the obvious way. Repeating the arguments given above shows that the only possibility which avoids immediate contradictions is the one where $d_T(U) < \beta$ and $d_T(W_2) < \alpha$. If in addition $d_T(W_1) < \alpha$ then $d_T(P) > 1 - 5\alpha - \beta \geq \frac{1}{2} + 2\gamma + \delta$ which contradicts the fact that P is a geodesic. Therefore assume that $d_T(W_1) \geq \alpha$. On the other hand, since there is an initial segment of P which reads W_1^{-1}, let $P = W_1^{-1}W_3$ for some possibly empty word W_3. By property 6 of relator metrics and Axiom 1, this second instance of W_1 is also long in T. If either W_1 is negatively oriented with respect to the loop representing T, then by applying Lemma 14.9 to this loop together with the inequalities listed above, it follows that $d_T(W_3) + d_T(W_1) + \alpha + 3\alpha > 1 + d_T(W_1) - \delta - 2\gamma$. Thus $d_T(W_3) > 1 - 4\alpha - \beta - 2\gamma - \delta \geq \frac{1}{2} + \alpha$ by the

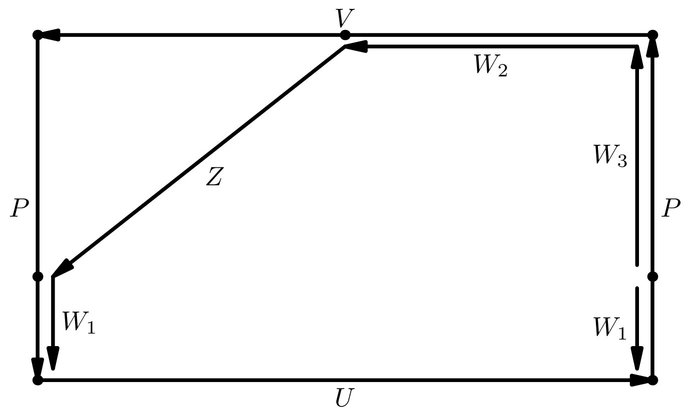

Figure 33: An illustration of Case 4

restrictions assumed on the constants, which contradicts the assumption that P is a geodesic. Both instances of W_1 are thus long and positively oriented with respect to the loop representing T. Lemma 14.11 can now be applied to this situation with W_1^{-1} playing the role of U, and U and W_3W_2Z standing in for X and Y respectively. Thus there exists another instance of R, called R_2, in T which is fixed by g and which is at least $(\frac{1}{2} - \beta - \delta)|T|$ units away from R_1.

Let u be the start vertex of one of the paths U in the loop U^n in R_2. Since WZ is a representative of T, by Axiom 3 there is a path from u to WZ whose length is at most $\gamma|T|$. Let Q be one such path. All possibilities for the endpoint of Q in WZ imply that u lies within $2\alpha|T|$ of W_2 or within $2\alpha|T|$ of one of the instances of P. To see this notice that the endpoint of Q cannot be in U (since by Lemma 14.11 the two instances of R lie at least $(\frac{1}{2} - \beta - \delta)|T|$ apart), that if the endpoint of Q lies in W_1, W_2, or P then the statement is immediate, and that if the endpoint of Q lies in Z then since $|Z| \leq (3\alpha + \delta)|T|$, it follows that u is at most $(\frac{3}{2}\alpha + \frac{1}{2}\delta + \gamma)|T| \leq 2\alpha|T|$ from either W_2 or W_1. Suppose u lies within $2\alpha|T|$ of W_2. From the estimates on the measures of the sides given above, $d_T(P) + d_T(W_1) \geq 1 - 4\alpha - \beta$. Thus $2|P| \geq |P| + |W_1| \geq (1 - 4\alpha - \beta - 2\delta)|T| \geq 6\alpha|T|$, or $|P| \geq 3\alpha|T|$. In either case R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If on the other hand u lies within $2\alpha|T|$ of a point v in one of the instances of P, then since u is at least $(\frac{1}{2} - \beta - \delta)|T|$ from U, it follows that the portion of P between U and v has a length of at least $(\frac{1}{2} - 2\alpha - \beta - \delta)|T| > 2\alpha|T|$. Thus R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If the four sides covered by W had been arranged differently, the conclusion would have been that there is an instance S_2 of S which is stabilized by h and R_1 is closer to S_2 than to S_1.

Case 5A: Suppose that W is a subword of five consecutive sides. This is possible because W might begin and end on the same side without overlapping. If the word W begins and ends on a side labeled U or V then without loss

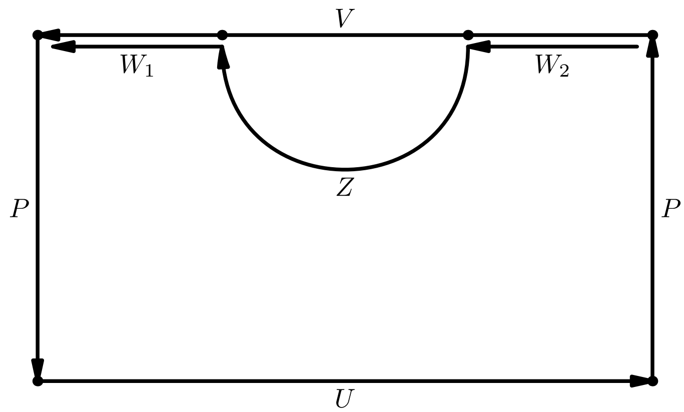

Figure 34: An illustration of Case 5A

of generality assume that W is a subword of $V^{-1}P^{-1}UPV^{-1}$ and write $W = W_1P^{-1}UPW_2$ in the obvious way. As usual it follows fairly immediately that $d_T(W_1) < \alpha$, that $d_T(U) < \beta$, and that $d_T(W_2) < \alpha$. Thus at least one of the instances of P has $d_T(P) \geq \frac{1}{2}(1 - 5\alpha - \beta) > \alpha$. Consequently at least one of the instances of P is long, and by a property of relator metrics, both instances of P are long and have the same measure. If one of the instances is negatively oriented with respect to the loop $W_1P^{-1}UPW_2Z$ then $d_T(P) \geq 1 + d_T(P) - 4\alpha - \beta - 2\gamma - 2\delta$, or $0 \geq 1 - 4\alpha - \beta - 2\gamma - 2\delta$, contradiction. Thus both instances of P are positively oriented with respect to the loop $W_1P^{-1}UPW_2Z$. Lemma 14.11 can now be applied to this situation with P playing the role of U, and U and W_2ZW_1 standing in for X and Y respectively. Thus there exists another instance of R, called R_2, in T which is fixed by g and which is at least $(\frac{1}{2} - \beta - \delta)|T|$ units away from R_1.

As in the earlier cases, let u be the start vertex of one of the paths U in the loop U^n in R_2. There must exist a path from u to WZ whose length is at most $\gamma|T|$. Let Q be one such path. All possibilities for the endpoint of Q in WZ imply that u lies within $2\alpha|T|$ of V or within $\gamma|T|$ of one of the instances of P. To see this notice that the endpoint of Q cannot be in U (since by Lemma 14.11 the two instances of R lie at least $(\frac{1}{2} - \beta - \delta)|T|$ apart), that if the endpoint of Q lies in W_1, W_2 or one of the instances of P then the statement is immediate, and that if the endpoint of Q lies in Z then since $|Z| \leq (3\alpha + \delta)|T|$, it follows that u is at most $(\frac{3}{2}\alpha + \frac{1}{2}\delta + \gamma)|T| \leq 2\alpha|T|$ from either W_2 or W_1. Suppose u lies within $2\alpha|T|$ of W_2. From the estimates on the measures of the sides given above, $d_T(P) + d_T(P) \geq 1 - 5\alpha - \beta$. Thus $2|P| \geq (1 - 5\alpha - \beta - 2\delta)|T| \geq 5\alpha|T|$, or $|P| > 2\alpha|T|$. Thus R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If, on the other hand, u lies within $2\alpha|T|$ of a point v in one of the instances of P, then since u is at least $(\frac{1}{2} - \beta - \delta)|T|$ from U, it follows that the portion of P between U and v has a length of at least $(\frac{1}{2} - 2\alpha - \beta - \delta)|T| > 2\alpha|T|$. Thus

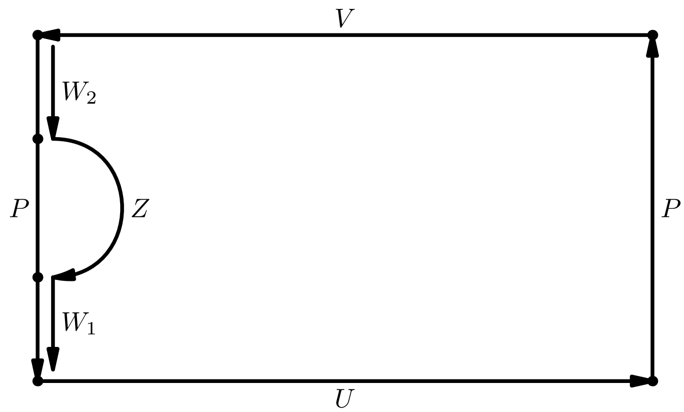

Figure 35: An illustration of Case 5B

in either case R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If the five sides covered by W had been arranged differently, the conclusion would have been that there is an instance S_2 of S which is stabilized by h and R_1 is closer to S_2 than to S_1.

Case 5B: Suppose that W is a subword of five consecutive sides and that W begins and ends on one of the sides labeled P. Then, without loss of generality, assume that W is a subword of $P^{-1}UPV^{-1}P^{-1}$ and write $W = W_1UPV^{-1}W_2$ in the obvious way. Moreover, let $P^{-1} = W_2W_3W_1$ for some possibly empty word W_3. Since R_1 and S_1 are chosen to be as close as possible, it follows that $|Z| \geq |W_3|$, and consequently $d_T(W_3) \leq d_T(Z) + 2\delta \leq 3\alpha + 2\delta$. From the usual argument it can be assumed that $d_T(U) < \beta$ and that $d_T(V) < \beta$. Thus $d_T(P) + d_T(W_1) + d_T(W_2) \geq 1 - 3\alpha - 2\beta$, and $d_T(W_1) + d_T(W_1) + d_T(W_2) + d_T(W_2) \geq 1 - 6\alpha - 2\beta - 2\delta \geq 4\alpha$. Consequently at least one instance of either W_1 or W_2 is long relative to T. By Axiom 1 and a property of relator metrics it follows that the other instance is also long. Without loss of generality, assume that both instances of W_1 are long relative to T.

If one of the instances of W_1 is negatively oriented with respect to the loop WZ then by Lemma 14.9 $d_T(W_1) + d_T(W_2W_3) + d_T(W_2) \geq 1 + d_T(W_1) - 3\alpha - 2\beta - 2\gamma - 2\delta$ and $d_T(W_2) + d_T(W_2) \geq 1 - 6\alpha - 2\beta - 2\gamma - 4\delta \geq 2\alpha$ so that one (and thus both) of the two instances of W_2 are long as well. If both instances of W_1 and W_2 are long and at least one instance of W_1 and one instance of W_2 are negatively oriented, then by Lemma 14.10, $d_T(W_1) + d_T(W_2) + d_T(W_3) \geq d_T(W_1) + d_T(W_2) + 1 - 3\alpha - 2\gamma - 2\delta$. Thus $d_T(W_3) \geq 1 - 3\alpha - 2\gamma - 2\delta \geq 1 - 5\alpha \geq \frac{1}{2} + \alpha$, which contradicts the assumption that P, and thus W_3, is a geodesic. Consequently at least one pair of the long W_i are positively oriented. Without loss of generality, assume that the two instances of W_1 are long and positively oriented with respect to WZ. Lemma 14.11 can now be applied to this situation with W_1^{-1} playing the role of U, and U and $W_3^{-1}W_2^{-1}V^{-1}W_2Z$

standing in for X and Y respectively. Thus there exists another instance of R, called R_2, in T which is fixed by g and which is at least $(\frac{1}{2} - \beta - \delta)|T|$ units away from R_1.

As above, let u be the start vertex of one of the paths U in the loop U^n in R_2. There must exist a path from u to WZ whose length is at most $\gamma|T|$. Let Q be one such path. All possibilities for the endpoint of Q in WZ imply that u lies either within $\gamma|T|$ of V or within $2\alpha|T|$ of one of the instances of P. To see this notice that the endpoint of Q cannot be in U (since by Lemma 14.11 the two instances of R lie at least $(\frac{1}{2}-\beta-\delta)|T|$ apart), that if the endpoint of Q lies in W_1, W_2, V, or the intact instance of P then the statement is immediate, and that if the endpoint of Q lies in Z then since $|Z| \leq (3\alpha + \delta)|T|$, it follows that u is at most $(\frac{3}{2}\alpha + \frac{1}{2}\delta + \gamma)|T| \leq 2\alpha|T|$ from either W_2 or W_1. Suppose u lies within $\gamma|T|$ of V. Then since $|P| \geq |W_1| \geq (\alpha - \delta)|T| > \gamma|T|$, R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If, on the other hand, u lies within $2\alpha|T|$ of a point v in one of the instances of P, then since u is at least $(\frac{1}{2} - \beta - \delta)|T|$ from U, it follows that the portion of P between U and v has a length of at least $(\frac{1}{2} - 2\alpha - \beta - \delta)|T| > 2\alpha|T|$. Thus in either case R_2 is stabilized by g and S_1 is closer to R_2 than to R_1. If the five sides covered by W had been arranged differently, the conclusion would have been that there is an instance S_2 of S which is stabilized by h and R_1 is closer to S_2 than to S_1.

Since these six cases exhaust all of the possibilities, the cycle $UPV^{-1}P^{-1}$ must be 3α-complement-free with respect to $\mathcal{R}$. Let Z be a Dehn-reduced cycle which is conjugate to $UPV^{-1}P^{-1}$ in G. Such a cycle exists by Lemma 10.5. Then by Lemma 7.10, there is a connected $\mathcal{R}$-diagram Δ whose boundary cycles are Z^{-1} and $UPV^{-1}P^{-1}$. Since the cycle Z is Dehn-reduced it follows that it is 4α-complement-free with respect to $\mathcal{R}$. Thus Z^p is readable as a loop with a non-trivial winding number in a general relator T. The importance of the word Z is that it is a word which represents the automorphism g^i followed by h^{-j}, so that T is the unique general relator associated to and stabilized by the automorphism $g^i h^{-j}$. Moreover, since the cycle $UPV^{-1}P^{-1}$ is 3α-complement-free with respect to $\mathcal{R}$, by Lemma 11.4 it follows that a power of this cycle is readable as a loop in T. Since $d_R(U) \geq 2\alpha$, by Axiom 1 there is an $\mathcal{R}$-functor from R to T. Thus the rank of T is at least as great as the rank of R. The rank of T cannot be strictly greater than that of R since this would contradict the initial choice of the automorphism g. The only possibility is for the ranks to be equal, which forces T and R to be identical by Lemma 2.1, and this in turns shows that S_1 and R_1 have a point in common, as was to be shown.

Claim 2: h fixes the instance of R which satisfies Claim 1.

This portion of the proof will be divided into two cases depending on whether the image of R_1 (chosen in Claim 1) under the automorphism h intersects with R_1 or not.

Case 1: Assume for the moment that $R_1 \cap h(R_1) = \emptyset$. This implies, in particular, that $S_1 \cap R_1$ and $S_1 \cap h(R_1)$ are disjoint. Let u be a vertex in $S_1 \cap R_1$, let U be a path in R_1 from $g^{-i}(u)$ to u, and let V be a path in S_1 from

u to $h(u)$. The path $X = UV$ thus represents the automorphism hg^i. Notice that since hg^i is in H and a word representing it is read in an instance of R followed by an instance of S, there is a power of X, say X^n, which is read as a loop in C in an alternating succession of instances of R and S.

If the cycle of X is not 3α-complement-free with respect to $\mathcal{R}$, then there exist a word W, a word Z, and a general relator $T \in \mathcal{R}$, such that WZ is a representative of T, $d_T(Z) \leq 3\alpha$ and W is a subword of a cyclic conjugate of X. If W, when read in X^n, is not contained in a single selected instance of either R or S, then it is still contained in at most three instances since W is a subword of a cyclic conjugate of UV. However, since $d_T(W) \geq 1 - 3\alpha > 3\alpha$, the portion of W in one of these instances is at least α. Thus by Axiom 1, all of T is contained in this instance, but this contradicts the assumption that W is not contained in a single instance of either R or S. If W, when read in X^n, is contained in a single selected instance of R or S, then all of T is contained in the instance by Axiom 1, and thus Z^{-1} can be substituted for W in a conjugate of X to create a shorter word which represents the same automorphism, and whose n-th power is still readable in the same sequence of instances of R and S. In addition, the new cycle X' has a cyclic conjugate $U'V'$ where $U'V'$ is readable as a path in $R_1 \cup S_1$ where U' is a word in R_1 from $R_1 \cap g^{-i}(S_1)$ to $R_1 \cap S_1$, and V' is a word in S_1 from $R_1 \cap S_1$ to $h(R_1) \cap S_1$. Since the length of the representative strictly decreases, such a substitution can occur only a finite number of times before the process stops.

When the process stops, consider the relationship between the words U and U'. Both U and U' are readable as paths in R_1 from $R_1 \cap g^{-i}(S_1)$ to $R_1 \cap S_1$. Since R and S are α-closed, so are $R_1 \cap g^{-i}(S_1)$ and $R_1 \cap S_1$ by Lemma 13.1. Thus by Lemma 11.3 a path between the startpoint of U and that of U' in R_1 (which exists since R_1 is connected), is equivalent to a path in $R_1 \cap g^{-i}(S_1)$. Similarly there is a path in $R_1 \cap S_1$ from the endpoint of U to that of U'. Call these paths P and Q respectively. By Axiom 2, $d_R(P)$ and $d_R(Q)$ are at most β. If the winding number of $PU'Q^{-1}U^{-1}$ is zero, then by the properties of relator metrics, $d_R(U') + d_R(P) + d_R(Q) \geq d_R(PU'Q^{-1}) = d_R(U)$, and by Lemma 14.8, $d_R(U') \geq \frac{1}{3} - 2\beta - 2\gamma - 2\delta$ which is at least α by the restrictions assumed on the constants. If on the other hand the winding number of $PU'Q^{-1}U^{-1}$ is not zero, then the properties of relator metrics show that $d_R(U') + d_R(P) + d_R(Q) + d_R(U) \geq 1$, which by Lemma 14.8 implies that $d_R(U') \geq \frac{1}{2} - \alpha - 2\beta$ which is also greater than α. In either case $d_R(U') \geq \alpha$.

For convenience relabel the word X' as X. At this point the cycle X is 3α-complement-free with respect to $\mathcal{R}$. Let Z be a Dehn-reduced cycle which is conjugate to X in G. Such a cycle exists by Lemma 10.5. Then by Lemma 7.10, there is a connected $\mathcal{R}$-diagram Δ whose boundary cycles are Z^{-1} and X. Since the cycle Z is Dehn-reduced it follows that it is 4α-complement-free with respect to $\mathcal{R}$. Thus Z^p is readable as a loop with a non-trivial winding number in a general relator T. The importance of the word Z is that it is a word which represents the automorphism g^ih (in the sense that it is a word which connects some vertex to its image under the automorphism), so that T is the unique general relator associated to and stabilized by the automorphism g^ih. Moreover, since the cycle X is 3α-complement-free with respect to $\mathcal{R}$, by Lemma 11.4 it follows that a power of this cycle is readable as a loop in T.

Since the subword U' of X^n which is contained in the original R_1 and since the measure of U' is at least α, there exists an $\mathcal{R}$-functor from R to T by Axiom 1. Thus the rank of T is at least as great as the rank of R. The rank of T cannot be strictly greater than that of R since this would contradict the initial choice of the automorphism g. The only possibility is for the ranks to be equal, which forces T and R to be identical by Lemma 2.1 and this in turns shows that $hg^i(R_1) = h(R_1) = R_1$, as was to be shown.

Case 2: Assume that $R_1 \cap h(R_1) \neq \emptyset$. The structure of the proof is nearly identical to that of Case 1 with only minor modifications. Let u be a vertex in $R_1 \cap h(R_1)$, let U be a path in R_1 from $g^{-i}(u)$ to u, and let V be a path in $h(R_1)$ from u to $h(u)$. The path $X = UV$ thus represents the automorphism hg^i. Notice that since hg^i is in H and a word representing it is read in an instance of R, there is a power of X, say X^n, which is read as a loop in C in an succession of instances of R.

If the cycle of X is not 3α-complement-free with respect to $\mathcal{R}$, then there exist a word W, a word Z, and a general relator $T \in \mathcal{R}$ such that WZ is a representative of T, $d_T(Z) \leq 3\alpha$ and W is a subword of a cyclic conjugate of X. If W, when read in X^n, is not contained in a single selected instance of either R, then it is still contained in at most three instances since W is a subword of a cyclic conjugate of UV. However, since $d_T(W) \geq 1 - 3\alpha > 3\alpha$, the portion of W in one of these instances is at least α. Thus by Axiom 1, all of T is contained in this instance, which contradicts the assumption that W is not contained in a single instance of R. If W, when read in X^n, is contained a single selected instance of R, then all of T is contained in the instance by Axiom 1, and thus Z^{-1} can be substituted for W in a conjugate of X to create a shorter word which represents the same automorphism, and whose n-th power is still readable in the same sequence of instances of R_1. In addition, the new cycle X' has a cyclic conjugate $U'V'$ where $U'V'$ is readable as a path in $R_1 \cup h(R_1)$ with U' a word in R_1 from $R_1 \cap g^{-i}h(R_1)$ to $R_1 \cap h(R_1)$. Since the length of the representative strictly decreases, such a substitution can occur only a finite number of times before the process stops.

When the process stops, consider the relationship between the words U and U'. Both U and U' are readable as paths in R_1 from $R_1 \cap g^{-i}h(R_1)$ to $R_1 \cap h(R_1)$. Since R is α-closed, so are $R_1 \cap g^{-i}h(R_1)$ and $R_1 \cap h(R_1)$ by Lemma 13.1. Thus by Lemma 11.3 a path between the startpoint of U and that of U' in R_1 (which exists since R_1 is connected) is equivalent to a path in $R_1 \cap g^{-i}h(R_1)$. Similarly there is a path in $R_1 \cap h(R_1)$ from the endpoint of U to that of U'. Call these paths P and Q respectively. The rest of the proof of this case is word for word the same as in case 1. Thus in either case 1 or case 2, $h(R_1) = R_1$ and the proof is complete. □

Lemma 14.13 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $2\beta + 2\gamma + \delta \leq \alpha \leq \frac{1}{12}$, let C be the Cayley category of the presentation, let H be a finite subgroup of G, and suppose that all general relators* *in $\mathcal{R}$ have at least one crucial cone in their boundary. If $g \in H$ has maximum rank, g rotates a general relator $R \in \mathcal{R}$, and the order of g is not a power of 2, then any instance R_1 of R in C which is fixed by g is also fixed by all of the other $h \in H$.*

Proof: In the proof of Lemma 14.12 the instances of R and S were chosen to be as close to each other as possible. The only places in the proof where the fact that R_1 is as close as possible to S_1 was used were in the later cases of Claim 1. Moreover, each invocation of this condition followed an application of Lemma 14.11. This lemma requires that g^{2^i} be an orientation-reversing automorphism of the higher-rank general relator T. But by Lemma 5.26 this would imply that the order g^{2^i} is a power of 2, and this contradicts the additional assumptions made in the statement of this lemma. Lemma 14.11, therefore, cannot be used under these circumstances, and the fact that R_1 is as close as possible to S_1 is never needed in the proof. Thus the instance of R stabilized by g can be chosen first (and arbitrarily) and then the instance of S stabilized by h can be chosen so that it is as close as possible to this fixed R_1. The proof of Lemma 14.12 then proceeds with the minor alteration described above. As a result, every element $h \in H$ must fix this arbitrarily chosen instance of R stabilized by g. □

Lemma 14.14 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $2\beta + 2\gamma + \delta \leq \alpha \leq \frac{1}{12}$, let H be a finite subgroup of G, and suppose that all general relators in $\mathcal{R}$ have at least one crucial cone in their boundary. If the order of $g \in H$ is not a power of 2, then g has maximum rank in H.*

Proof: By Lemma 14.12, g fixes the relator of maximum order, but by the structure of automorphism groups of relators (Lemma 5.26), g must be a rotation of this general relator. Thus g itself is an element of maximum rank. □

Lemma 14.15 *Let $G = \langle A|\mathcal{R}\rangle$ be a finitely presented general small cancellation presentation with $2\beta + 2\gamma + \delta \leq \alpha \leq \frac{1}{12}$, let H be a finite subgroup of G, and suppose that all general relators in $\mathcal{R}$ have at least one crucial cone in their boundary. If the order of H is not a power of 2, then H is isomorphic to a subgroup of an automorphism group of a general relator in $\mathcal{R}$.*

Proof: Pick an element $g \in H$ whose order is not a power of 2, let R be the unique general relator rotated by g, and let R_1 be an instance of R_1 in C which is fixed by g. By Lemma 14.14, g is an element of maximum rank, and by Lemma 14.13 all of the $h \in H$ stabilize R_1. Thus H is a subgroup of $\text{Aut}(R_1) = \text{Aut}(R)$. □

14.4 Finite Subgroups

The lemmas below complete the task of showing that all of the finite subgroups of a general small cancellation group are contained in the automorphism group of some general relator.

Lemma 14.16 *Let $G = \langle A|\mathcal{R}\rangle$ be a general small cancellation presentation and let C be the Cayley category of G. If H is a finite subgroup of G, then there is a finite set of general relators $\mathcal{R}'$ in $\mathcal{R}$ such that H is also a finite subgroup of the finitely presented general small cancellation presentation $G' = \langle A|\mathcal{R}'\rangle$.*

Proof: Since the general relators in $\mathcal{R}$ are thin by Lemma 9.8, it follows from Lemma 5.19 that there exists a set of standard representatives for the set $\mathcal{R}$. Once G is given a standard presentation, it is a standard exercise in combinatorial group theory to show that there exists a finitely presented subpresentation which also contains the subgroup H. The replacement of the standard representatives by the general relators $R \in \mathcal{R}$ which they represent yields a finite set of general relators $\mathcal{R}'$ which, when closed under subcones, satisfy the requirements of the lemma. □

Lemma 14.17 *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $2\beta + 2\gamma + \delta \leq \alpha \leq \frac{1}{12}$ in which str(W) is finite and effectively constructible for all words $W \in A^*$, and such that all general relators in $\mathcal{R}$ have at least one crucial cone in their boundary, then every finite subgroup of G is a subgroup of the automorphism group of some general relator in $\mathcal{R}$.*

Proof: Let H be a finite subgroup G. By Lemma 14.16 there is a finite subset $\mathcal{R}'$ of $\mathcal{R}$ which is closed under subcones and which also contains the subgroup H. By Lemma 14.5, the lemma is satisfied whenever H is a 2-group, and since by Lemma 13.6, str$(W, \mathcal{R}')$ is finite and effectively constructible, Lemma 14.15 shows that it is satisfied whenever H is not a 2-group. □

14.5 Proof of Theorem A

At this point all of the pieces needed to show Theorem A have been shown separately, so it is merely a matter of collecting the references and of showing that the assumptions listed are enough to satisfy the hypotheses of each of the quoted lemmas.

Theorem A *If $G = \langle A|\mathcal{R}\rangle$ is a general small cancellation presentation with $\alpha \leq \frac{1}{12}$, then the word and conjugacy problems for G are decidable, the Cayley graph is constructible, the Cayley category of the presentation is contractible, and G is the direct limit of hyperbolic groups. If, in addition, the presentation satisfies the hypotheses of Lemma 14.17, then every finite subgroup of G is a subgroup of the automorphism group of some general relator in $\mathcal{R}$.*

Proof: The conclusions follow from Lemma 10.21, Lemma 10.23, Lemma 10.22, Lemma 12.12, Lemma 11.10, and Lemma 14.17, respectively. Any inequality required by one of these lemmas but not listed in the statement above follows immediately from a combination of $\alpha \leq \frac{1}{12}$ and Axiom 7. □

References

[1] T. Bedford, M. Keane, and C. Series, editors. *Ergodic Theory, Symbolic Dynamics and Hyperbolic Spaces.* Oxford University Press, New York, 1991.

[2] G. Bredon. *Topology and Geometry.* Springer-Verlag, New York, 1993.

[3] K. Brown. *Cohomology of Groups.* Springer-Verlag, New York, 1982.

[4] J. Cannon. The theory of negatively curved spaces and groups. In Bedford et al. [1], chapter 11, pages 315–369.

[5] D. Epstein, J. Cannon, D. Holt, S. Levy, M. Paterson, and W. Thurston. *Word Processing in Groups.* Jones and Barlett, Boston, 1992.

[6] S. Gersten, editor. *Essays in Group Theory.* Springer-Verlag, New York, 1987.

[7] M. Gromov. Hyperbolic groups. In Gersten [6], pages 75–263.

[8] S. Ivanov. The free Burnside groups of sufficiently large exponents. *International Journal of Algebra and Computation*, 4:1–308, June 1994.

[9] R. Lyndon and P. Schupp. *Combinatorial Group Theory.* Springer-Verlag, New York, 1977.

[10] I. Lysionok. Infinite Burnside groups of even period. *Izv. Ross. Akad. Nauk. Ser. Mat.*, 60:3–224, 1996.

[11] S. MacLane. *Categories for the Working Mathematician.* Springer-Verlag, New York, 1971.

[12] J. McCammond. Burnside groups via general small cancellation theory. In preparation.

[13] J. McCammond. The solution to the word problem for the relatively free semigroups satisfying $t^a = t^{a+b}$ with $a \geq 6$. *International Journal of Algebra and Computation*, 1:1–32, 1991.

[14] R. Oliver. Fixed-point sets of group actions on finite acyclic complexes. *Comment. Math. Helvetici*, 50:155–177, 1975.

[15] A. Yu. Ol'shanskii. On the Novikov-Adian theorem. *Math. Sbornik*, 118:203–235, 1982.

[16] A. Yu. Ol'shanskii. *Geometry of defining relations in groups.* Kluwer Acad. Pub., 1991.

[17] D. Quillen. *Higher Algebraic K-Theory: I.* Springer-Verlag, New York, 1973.

[18] E. Rips. Generalized small cancellation theory and applications: I. the word problem. *Israel J. Math.*, 41:1–146, 1982.

[19] J. Stallings. On torsion-free groups with infinitely many ends. *Annals of Math. (2)*, 88:312–334, 1968.

[20] G. Whitehead. *Elements of Homotopy Theory.* Springer-Verlag, New York, 1978.

[21] A. I. Zimin. Blocking sets of terms. *Mat. Sb.*, 119:363–375, 1982.

Index