INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS

INTERNATIONAL ATOMIC ENERGY AGENCY

UNITED NATIONS EDUCATIONAL, SCIENTIFIC AND CULTURAL ORGANIZATION

Group Theory from a Geometrical Viewpoint

26 March – 6 April 1990 ICTP, Trieste, Italy

Editors

E. Ghys
École Normale Supérieure de Lyon
France

A. Haefliger
Université de Genève
Switzerland

A. Verjovsky
ICTP, Trieste
Italy

World Scientific
Singapore • New Jersey • London • Hong Kong

Geodesics and curvature in metric simplicial complexes 373
M. R. Bridson

Hyperbolizations 464
T. Januszkiewicz

Non-positively curved triangles of groups 491
J. R. Stallings

Complexes of groups and orbihedra 504
A. Haefliger

IV. ACTIONS ON TREES

Dendrology and its applications 543
P. B. Shalen

Cocycles bornés et actions de groupes sur les arbres réels 617
J. Barge and E. Ghys

A quick introduction to Burnside's problem 622
V. Sergiescu

V. MISCELLANEOUS

Projective aspects of the Higman-Thompson group 633
P. A. Greenberg

Hyperbolic geometry and the subgroups of the modular group 645
R. S. Kulkarni

Identities among relations of group presentations 687
S. J. Pride

Closures of totally geodesic immersions in manifolds of constant negative curvature 718
N. A. Shah

List of Contributors 733

GEODESICS AND CURVATURE IN METRIC SIMPLICIAL COMPLEXES

A Dissertation
Presented to the Faculty of the Graduate School
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by
Martin Robert Bridson

August, 1991

Geodesics and Curvature in Metric Simplicial Complexes

Martin Robert Bridson, Ph.D.
Cornell University 1991

Groups which act cocompactly on simplicial trees were completely classified by the work of Bass and Serre. The elegance of this theory is such that the prospect of extending it to higher dimensions is an extremely enticing one. As a first step one must identify a suitable higher-dimensional analogue of a tree. A strong candidate for this role is provided by the class of non-positively curved piecewise Euclidean complexes (which are defined below). The study of such complexes is not new, but has been brought to the fore in recent years by Gromov, who made extensive use of these spaces in his remarkable work on hyperbolic groups.

In his seminal article [17] Gromov (pp. 119–120) states several theorems which are important to the understanding of the geometry of these complexes — theorems which concern the existence of geodesics, and the relationship between local and global definitions of non-positive curvature in simply connected spaces. However, until now the validity of these results had only been established for locally finite complexes. In the study of groups acting on trees the spaces under consideration are not required to be locally finite, and such a restriction would be very limiting in the higher-dimensional case. Thus we are presented with a serious technical difficulty.

In the first three chapters of this thesis we remove this difficulty by proving these theorems for a large class of spaces. This class includes any piecewise Euclidean complex which admits a cocompact action by a group of isometries. We also relate Gromov's ideas to earlier work of others, notably Bruhat and Tits, and Alexandrov, and establish the equivalence of a variety of characterisations of non-positive curvature, both local and global.

These results allow us to analyse the structure of groups which act cocompactly on a simplicial complex by using curvature to convert local (combinatorial) information about the complex into global information which relates to the group action. In Sections 4 and 5 we give a number of examples to show how, in the presence of sufficient local information about the complex, one can establish the existence, or non-existence, of a metric of non-positive curvature. In particular, we prove that the Culler-Vogtmann complex does not support an $Out(F_n)$-equivariant metric of non-positive curvature for $n \geq 3$.

Biographical Sketch

The author was born in the Isle of Man in 1964 and lived there until the age of eighteen. He read mathematics at Hertford College Oxford from 1983 to 1986 and was awarded a B.A. degree with first class honours. He entered Cornell University in 1986 and received an M.S. degree in 1988.

To the memory of my father

Gerald P. Bridson

"Cre t'ayd nagh vel oo er gheddyn?"

Acknowledgements

I would like to thank my thesis advisor, Karen Vogtmann, for introducing me to such a rich area of mathematics, and for skilfully guiding my progress through it. I would also like to thank Marshall Cohen for his great generosity and sound advice, both mathematical and personal.

I have learned something from each of the topologists at Cornell, faculty and graduate students, and I am grateful to them all. In particular I would like to thank Ken Brown and Allen Hatcher, who served on my special committee, and Jim West, who guided me through my early days at Cornell.

I also wish to record my thanks to those who provided me with financial support — the mathematics department and the graduate school at Cornell, the Alfred P. Sloan Foundation, and the Isle of Man government.

Finally, my deepest gratitude is reserved for my family, particularly my mother Lilian, for their unfailing love and support.

Table of Contents

List of Figures viii

Introduction 1

Chapter 1: The Existence of Geodesics in Metric Simplicial Complexes 8

1.1. Definitions 9

1.2. The Intrinsic Metric 11

1.3. The Existence of Shortest m-chains 13

1.4. Taut Chains and Local Geodesics 15

1.5. The Growth of Taut Chains 21

1.6. The Completeness of the Metric 30

1.7. The Two Dimensional Case 32

Chapter 2: Non-Positive Curvarture in Piecewise Euclidean Complexes 35

2.1. The Convexity of the Metric 37

2.2. Global Characterisations of Non-Positive Curvature 47

2.3. Local Characterisations of Non-Positive Curvature 54

2.4. Passing From the Local to the Global Situation 60

2.5. Extending Geodesics 72

Chapter 3: $M(\kappa)$–Complexes of Non-Positive Curvature 74

3.1. The Proof of the Main Theorem 74

3.2. Alternative Forms of the Link Condition 80

Chapter 4: Examples 83
4.1. Simplices of Groups 84
4.2. 3–Complexes with Planar Links 87

Chapter 5: The Curvature of the Culler-Vogtmann Complex 95
5.1. The Culler-Vogtmann Complex K_n 96
5.2. The Link of a Rose in K_3 99
5.3. The Curvature of K_3 110
5.4. The Curvature of K_n 120

References 126

List of Figures

1.1. A taut chain which is not a local geodesic 20

1.2. The complex $k(x_i)$ 24

1.3. The image of the developing map 30

1.4. Shortening a chain which does not have property τ 33

2.1. Superlinear divergence of geodesics 41

2.2. Locally divergence is linear 49

2.3. $CAT(0)$ 51

2.4. A comparison figure preserving angles 53

2.5. Constructing a comparison triangle 54

2.6. Radial projection of a geodesic segment 58

5.1. Four equivalent labelled graphs 98

5.2. A 3–simplex in K_3 98

5.3. A 3–simplex in $K_3/Out(F_3)$ 98

5.4. The face A(+) 103

5.5. The cube 104

5.6. The hexagon H(c) 105

5.7. Fitting the hexagons into the cube 106

5.8. The graphs on the edges of the hexagons 106

5.9. S(o,c), twelve of the remaining seventy two 2–cells 107

5.10. Attaching the 2–disc S(o,c) 108
5.11. The torus 109
5.12. Three 2–cells in $K_3/Out(F_3)$ 111
5.13. The vertices considered in Lemma 5.2 114
5.14. The path P_1 of Lemma 5.2 114
5.15. The marked graph (g_0, G_0) 115
5.16. The roses in the link of (g_0, G_0) 115
5.17. The path P_2 of Lemma 5.3 115
5.18. The graph (g_1, G_1) and its link 117
5.19. The vertices considered in Lemma 5.5 119
5.20. The 1–complex L 119

Introduction

Groups which act cocompactly on simplicial trees were completely classified by the work of Bass and Serre [24]. The elegance of this theory is such that the prospect of extending it to higher dimensions is an extremely enticing one. As a first step one must identify a suitable higher-dimensional analogue of a tree. A strong candidate for this role is provided by the class of *non-positively curved piecewise Euclidean complexes* (which are defined below). The study of such complexes is not new, but has been brought to the fore in recent years by Gromov, who made extensive use of these spaces in his remarkable work on hyperbolic groups.

In his seminal article [17] Gromov (pp. 119–120) states several theorems which are important to the understanding of the geometry of these complexes — theorems which concern the existence of geodesics, and the relationship between local and global definitions of non-positive curvature in simply connected spaces. However, until now the validity of these results had only been established for locally finite complexes. In the study of groups acting on trees the spaces under consideration are not required to be locally finite, and such a restriction would be very limiting in the higher-dimensional case. Thus we are presented with a serious technical difficulty.

In Sections 1–3 of this paper we remove this difficulty by proving these theorems for a large class of spaces. This class includes any piecewise Euclidean complex which admits a cocompact action by a group of isometries. We also relate Gromov's ideas to earlier work of others, notably Bruhat and Tits [10], and Alexandrov [2].

These results allow us to analyse the structure of groups which act cocompactly on a simplicial complex by using curvature to convert local (combinatorial) information about the complex into global information which relates to the group action. In Sections 4 and 5 we give a number of examples to show how, in the presence of sufficient local information about the complex, one can establish the existence, or non-existence, of a metric of non-positive curvature.

In order to state our results we need the following definitions. A *geodesic metric space* is a metric space in which every pair of points can be joined by a *geodesic segment* — a topological arc which, with the induced metric, is isometric to a closed interval of the real line. A large class of examples is provided by *piecewise Euclidean complexes:* Given a simplicial complex K one can metrize the cells of the geometric realisation of K so that each cell is isometric to a Euclidean simplex, and if two cells intersect then the induced metrics on their common face agree. There is a well defined notion of piecewise linear (PL) paths in K and a consistent way of measuring their length. The *intrinsic pseudometric* on K is defined by setting the distance between two points equal to the greatest lower bound on the length of PL paths joining them.

If K is connected and locally finite then this pseudometric is actually a metric, and if K is finite then this construction yields a geodesic metric space. Moreover, this construction works equally well if one replaces Euclidean simplices by hyperbolic or spherical ones, and we use the term *metric simplicial complex* to refer to a complex which is piecewise-Euclidean, -spherical, or -hyperbolic.

If the complex K is not locally finite then in general the intrinsic pseudometric is not a metric. A simple example is the 1–complex with two vertices joined by countably many edges (subdivided to make the complex simplicial), the $n-$th of which has length $1/n$. To avoid the type of limiting behaviour present in this example we restrict our attention to metric simplicial complexes with only finitely many isometry types of cells.

Theorem 1.1: *If K is a metric simplicial complex with only finitely many isometry types of cells then K is a complete geodesic metric space.*

We say that K is of *type B* if it satisfies the hypothesis of Theorem 1.1. We do *not* require spaces of type B to be locally compact. We shall also consider (locally finite) metric simplicial complexes all of whose closed bounded subsets are compact. Such complexes are said to be of *type A*.

Any finite-dimensional simplicial complex can be metrized as a piecewise Euclidean complex of type B by metrizing each cell as a regular simplex of unit edge length. It then satisfies the hypothesis of our Theorem 1.1 and hence is a complete geodesic metric space. Any simplicial action on the space becomes an action by isometries, and the potential exists for studying groups which act in this way via the geometry of geodesics in the complex. This technique works particularly well in the presence of *non-positive curvature*.

Our Main Theorem, which is stated below, establishes the equivalence of various characterisations of non-positive curvature for $M(\kappa)-$complexes of type A or B, where $\kappa \leq 0$ denotes the sectional curvature of the simplices in the given complex. For clarity of exposition, we state and prove our results for piecewise Euclidean complexes (Chapter 2) before generalising to the case $\kappa \leq 0$ (Chapter 3).

In Section 2.1 we study complexes for which there is a unique shortest path between every pair of points, and obtain the following convexity result.

Theorem 2.1: *If K is a piecewise Euclidean complex of type A or B which has unique geodesic segments then any geodesic segments α_0 and α_1 with*

$\alpha_0(0) = \alpha_1(0)$ *satisfy*

$$d(\alpha_0(t), \alpha_1(t)) \leq t \cdot d(\alpha_0(1), \alpha_1(1)) \quad \forall t \in [0,1].$$

We then show, in Section 2.2, that K has unique geodesic segments if and only if it has non-positive curvature as defined by the CAT(0) inequality of Gromov [17], the CN inequality of Bruhat-Tits [10], and Alexandrov's condition on the excess of geodesic triangles:

Theorem 2.7: *If K is a piecewise Euclidean complex of type A or B then the following global characterisations of non-positive curvature are equivalent:*

I) *K has unique geodesic segments.*

II) *K satisfies CAT(0).*

III) *K satisfies CN.*

IV) *Every geodesic triangle in K has non-positive excess.*

This leads to a fixed point theorem in the manner of Bruhat-Tits. In dimension 2 this fixed point theorem is due to Gersten [15] and plays a central role in his work with Stallings on triangles of groups [26].

Fixed Point Theorem: *If K satisfies any of the conditions* I *to* IV *given in Theorem 2.7, and a group Γ acts on K by isometries, such that there is a bounded orbit, then the fixed point set of Γ is non-empty and contractible.*

In Section 2.3 we describe the relationship between local definitions of curvature and prove:

Theorem 2.8: *If K is a piecewise Euclidean complex of type A or B then the following local characterisations of non-positive curvature are equivalent:*

I) *K has unique geodesic segments locally.*

II) *The metric on K is convex locally.*

III) *K satisfies CAT(0) locally.*

IV) *K satisfies CN locally.*

V) *Every point of K has a neighbourhood such that any geodesic triangle in that neighbourhood has non-positive excess.*

VI) *K satisfies the link condition.*

Further, if K is of type B then each of the above conditions is equivalent to

VII) *There exists* $\epsilon_0 > 0$ *such that for all* $x \in K$ *the ball* $B_{\epsilon_0}(x)$ *is geodesically convex and has unique geodesic segments.*

In Section 2.4 we prove the following theorem, which provides the vital link between local and global definitions of non-positive curvature in a simply connected space.

Theorem 2.12: *If K is a piecewise Euclidean complex of type A or B which has unique geodesic segments locally then for every pair of points* $x, y \in K$ *there is a unique shortest path in each homotopy class of paths from* x *to* y *in* K.

This result, for a restricted class of locally-finite complexes, is due to Stone [27]. The strategy of our argument is modelled on the proof of a similar result for smooth manifolds, which can be found in Milnor's book on Morse theory [22].

In Chapter 3 we generalise the preceding results to $M(\kappa)-$simplicial complexes and obtain:

Main Theorem: *If K is a simply connected* $M(\kappa)-$*simplicial complex of type A or B and* $\kappa \leq 0$ *then the following 13 conditions are equivalent:*

Global conditions:

I) *K has unique geodesic segments.*

II) *K satisfies $CAT(\kappa)$ globally.*

III) *K satisfies $CAT(\chi)$ globally, for some χ.*

IV) *K satisfies CN globally.*

V) *The metric on K is convex.*

VI) *Every geodesic triangle in K has non-positive excess.*

Local conditions:

VII) *K has unique geodesic segments locally.*

VIII) *K satisfies $CAT(\kappa)$ locally.*

IX) *K satisfies $CAT(\chi)$ locally, for some χ.*

X) *K satisfies CN locally.*

XI) *The metric on K is convex locally.*

XII) *Every point of K has a neighbourhood such that any geodesic triangle in that neighbourhood has non-positive excess.*

XIII) *K satisfies the link condition.*

The *link condition* listed as condition XIII) is essentially a condition on the combinatorics of the links of vertices in K. Thus, if we have sufficient local (combinatorial) information about the space K then we can deduce global (topological and algebraic) information. (e.g. If K is a non-positively curved complex of type A or B then its universal cover is contractible.) In Chapter 4 we give a number of examples where one can verify the link condition directly. In particular we discuss the work of Gersten and Stallings on triangles of groups. We also work through a specific example to show how the Fixed Point Theorem can lead to a classification of finite subgroups in a group which acts by isometries on a complex of non-positive curvature.

In Chapter 5 we describe a technique for deciding whether or not a complex can be given a metric of non-positive curvature which is equivariant with respect to a given group action. Using this technique we prove the following result.

Theorem 5.6: *If $n \geq 3$ then there does not exist an $Out(F_n)$−equivariant piecewise Euclidean (or piecewise hyperbolic) structure of non-positive curvature on the Culler-Vogtmann complex K_n.*

The Culler-Vogtmann complex is closely analogous to the Teichmüller space of a Riemann surface, and Theorem 5.6 provides an interesting analogue of known results about the curvature of Teichmüller space (see [21]).

The study of curvature in locally finite complexes, particularly PL manifolds, is not new. Work in this area was done by Banchoff and Stone in the sixties, and more recently by Gromov and Thurston [18]. Aitchison and Rubinstein [1] have shown that this theory has powerful applications in the geometry of 3–manifolds. We should also note that in the case of locally finite complexes Ballman ([16], Chapter 10) has given a lucid account of the theorems of Gromov to which we referred at the beginning.

The role which non-positive curvature plays in the study of orbihedra is highlighted by results of Gromov ([17], pp. 127–130), which have been explained in detail by Haefliger ([16], Chapter 11). Recently Haefliger [19] has generalised these results using the techniques which we introduce in this thesis.

1. The Existence of Geodesics in Metric Simplicial Complexes

In this chapter we prove the following theorem:

Theorem 1.1: *If K is a metric simplicial complex with only finitely many isometry types of cells then K is a complete geodesic metric space.*

Notice that we do not require the underlying simplicial complex K to be locally finite. In fact, if K is locally finite then we can weaken the condition that K has only finitely many isometry types of cells, and instead require only that in the intrinsic metric every closed bounded subset of K is compact. Then one can prove that K is a geodesic metric space by using the Arzela-Ascoli Theorem to show that the Birkhoff curve-shortening process converges (cf. [16] Chapter 10). Alternatively, one can view the existence of geodesics in this case as a formal consequence of the case where K has finitely many isometry types of cells; since for any two points $x, y \in K$ with $d(x, y) < N$ we need only consider paths between them which lie in the minimal subcomplex of K containing the ball of radius N about x, and this is a finite complex.

This chapter is organised as follows: In Section 1.1 we give a precise definition of a metric simplicial complex K, and define the intrinsic pseudometric d in terms of m*–chains* (which provide a useful method for describing piecewise-geodesic paths combinatorially). In Sections 1.2–1.5 we establish the existence of geodesic segments in K, and in Section 1.6 we prove that the metric d is complete. For 2–dimensional complexes our proof simplifies considerably, and in Section 1.7 we give the details in this restricted setting to illustrate the ideas involved in the proof of the general case.

1.1 Definitions

A *geodesic segment* in a metric space (X, d) is a topological arc which is isometric to a closed interval of the real line. (So in particular the length of a geodesic segment is equal to the distance between its endpoints.) (X, d) is a *geodesic metric space* if every pair of points in X can be connected by a geodesic segment.

An $n-$*simplex* in Euclidean $n-$space $\mathbf{E^n}$ or hyperbolic $n-$space $\mathbf{H^n}$ is the convex hull of $(n+1)$ points in general position. An $n-$*simplex* in spherical $n-$space $\mathbf{S^n}$ is the intersection of $\mathbf{S^n}$ with the positive cone spanned by $n+1$ linearly independent vectors in $\mathbf{R^{n+1}}$.

An **E** (respectively **S, H**) *simplicial complex* consists of the following information:

I) An (abstract) simplicial complex K.

II) A set $Shapes(K)$ of Euclidean (respectively spherical, hyperbolic) simplices $\sigma_i \subset \mathbf{E^{n_i}}$ (respectively $\mathbf{S^{n_i}}$, $\mathbf{H^{n_i}}$).

III) For every closed simplex $B \subset K$ a simplicial isomorphism $f_B : B \to \sigma(B)$ where $\sigma(B) \in Shapes(K)$ and

$$f_B \circ \left(f_C|_{(B \cap C)}\right)^{-1}$$

is an isometry for all simplices $B, C \subset K$.

Similarly, for any value of κ, we can define an $M(\kappa)-$*simplicial complex,* by requiring that $Shapes(K)$ consist of simplices of constant curvature κ. A *metric simplicial complex* is an $M(\kappa)-$simplicial complex, for some κ. We say that K is *connected* if it is connected in the weak topology.

Remark: It is purely for convenience that we work with simplicial complexes rather than more general complexes in which the cells are metrized as convex polyhedra, and since any such polyhedral complex can be made simplicial by subdivision this involves no loss of generality.

A *line segment* in K is the inverse image of a geodesic arc in $f_B(B)$ for some simplex B. By a *PL path* in K we mean a path which is the concatenation of finitely many such line segments. Notice that the maps f_B induce metrics d_B on the individual simplices of K. These metrics agree on faces of intersection and give a well defined notion of length for PL paths in K. To describe these paths combinatorially we use the following definition:

An *m-chain from x to y* is an $(m+1)$–tuple $C = (x_0, x_1, \ldots, x_m)$ of points in K, such that $x = x_0$, $y = x_m$, and for each i there exists a simplex $B(i)$ containing x_i and x_{i+1}. We call m the *size* of C, and define the *length* of C to be:

$$\lambda(C) = \sum_{i=0}^{m-1} d_{B(i)}(x_i, x_{i+1}).$$

Every m–chain determines a PL path in K, given by the concatenation of the line segments $[x_i, x_{i+1}]$. We denote this path $p(C)$.

If K is connected we can define a distance function:

$$d(x,y) = inf\,\{\lambda(C) : C \text{ a chain from } x \text{ to } y\}.$$

We call d the *intrinsic pseudometric* on K. The example given in the introduction shows that in general d does not define a metric. However, in Chapter 2 we show that if K is any connected metric simplicial complex with $Shapes(K)$, the set of isometry classes of cells of K, *finite* then this formula does define a metric on K. We then refer to d as the *intrinsic metric* on K.

For the remainder of Chapter 1 we fix an arbitrary κ and write M in place of $M(\kappa)$. We always assume that K is a connected M–simplicial complex, and that $Shapes(K)$ is a finite set. The letter d always denotes the intrinsic pseudometric on K.

1.2 The Intrinsic Metric

We begin by proving that $d(x, y)$ defines a metric on K. The fact that it is a pseudometric is immediate, the only difficulty is in showing that the distance between distinct points is nonzero.

Notation: For $x \in K$ we denote the open star of x (i.e., the union of the interiors of the cells containing x) by $st(x)$, and denote the closed star of x by $St(x)$.

Let $\epsilon(x) = inf\{d_B(x, B - st(x)) : B \text{ a closed simplex in } K, x \in B\}$.

Lemma 1.2: $\epsilon(x) > 0$ *for all* $x \in K$.

Proof: Fix x and let B_0 denote the unique simplex of K which contains x in its interior. Define an equivalence relation on the closed simplices of K that contain x by $B \sim B'$ if and only if there is an isometry from B to B' which restricts to the identity on B_0. Notice that $d_B(x, B - st(x))$ is well-defined on equivalence classes.

Because $Shapes(K)$ is finite, there are only finitely many equivalence classes under $\sim$. Choose representatives $\{B_1, \ldots, B_m\}$ for these classes. Then

$$\begin{aligned} \epsilon(x) &= inf\{d_B(x, B - st(x)) : B \text{ a closed simplex in } K,\ x \in B\} \\ &= min\{d_{B_i}(x, B_i - st(x)) : i = 1, \ldots, m\} \end{aligned}$$

and hence $\epsilon(x) > 0$. □

The positivity of the metric now follows from Lemma 1.3, which also shows that for every simplex B the metric d agrees with d_B locally. However, it should be noted that in general the metric which d induces on B is *not* equal to the metric d_B, but instead we have the inequality $d_B \geq d$. But this difficulty (which is illustrated by the following example) is a minor one, which can be rectified by taking a suitable barycentric subdivision of the model simplices in $Shapes(K)$ and giving K the induced simplicial structure.

Example: Consider the 1–complex K with vertices $\{x\}, \{y\}, \{z\}$ and edges $\{x, y\}, \{x, z\}, \{y, z\}$ of length 1, 1 and 4 respectively. If we let $B = \{y, z\}$ then $d_B(a, b) > d(a, b)$ whenever $d_B(a, b) > 3$. However, for every simplex B' in the first barycentric subdivision of K we have $d_{B'}(a, b) = d(a, b)$, for all $a, b \in B'$.

Lemma 1.3: *If $x \in K$ and $d(x, y) < \epsilon(x)$ then there exists a simplex B containing both x and y such that $d_B(x, y) = d(x, y)$.*

Proof: It is enough to show that there exists a simplex B such that $x, y \in B$ and $d_B(x, y) \leq d(x, y)$. This follows immediately from the following assertion, which we prove by induction on m, the size of the given chain:

If $C = (x_0, \ldots, x_m)$ is an m–chain in K and $\lambda(C) < \epsilon(x_0)$ then $x_i \in st(x_0)$ for all i, and for some (and hence any) simplex $B(m)$ containing x_0 and x_m the inequality $d_{B(m)}(x_0, x_m) \leq \lambda(C)$ holds.

The case $m = 1$ follows immediately from the definition of $\epsilon(x_0)$. For $m > 1$: Because C is a chain there is a simplex $B(m)$ containing both x_{m-1} and x_m. Applying our inductive hypothesis to $C' = (x_0, \ldots, x_{m-1})$, we may assume that $x_{m-1} \in st(x_0)$. Hence $x_0 \in B(m)$, and

$$\begin{aligned} d_{B(m)}(x_0, x_m) &\leq d_{B(m)}(x_0, x_{m-1}) + d_{B(m)}(x_{m-1}, x_m) \\ &\leq \lambda(C') + d_{B(m)}(x_{m-1}, x_m) \\ &= \lambda(C) \\ &< \epsilon(x_0). \end{aligned}$$

It follows from the definition of $\epsilon(x_0)$ that $x_m \in st(x_0)$. This completes the induction. □

Note: The metric topology which we have constructed on K coincides with the topology given by barycentric coordinates, and hence is strictly smaller than the *weak topology except when K is locally finite. However, the identity map from K equipped with the weak topology to K equipped with the metric topology is a homotopy equivalence* [14]. *This latter observation is important because it implies that if K is simply connected and has non-positive curvature then it is contractible in both the metric topology and the weak topology.*

1.3 The Existence of Shortest m-chains

Suppose that $C = (x_0, x_1, \ldots, x_m)$ is an $m-$chain from x to y in K satisfying $\lambda(C) = d(x, y)$. Then the PL path obtained by concatenation of the line segments $[x_i, x_{i+1}]$ is a geodesic segment. Therefore, to prove that K is a geodesic metric space it is enough to show that for all $x, y \in K$ the infimum in the definition of $d(x, y)$ is attained. This we do in two stages. Firstly, in this section, we show that if for a fixed integer m there exists some $m-$chain from x to y in K then there exists an m-chain of minimal length from x to y in K. Then, in Section 1.5, we show that there exists a linear function f such that for every pair of points $x, y \in K$

$$inf\left\{\lambda\left(C\right) : C \text{ an } m\text{−chain from } x \text{ to } y,\ m > 0\right\}$$
$$=inf\left\{\lambda\left(C\right) : C \text{ an } m\text{−chain from } x \text{ to } y,\ m \leq f\left(d\left(x,y\right)\right)\right\}.$$

It follows that the infimum in the definition of $d(x, y)$ is attained, and hence K is a geodesic metric space.

Lemma 1.4: *(Moussong) If L is a finite M-simplicial complex, and two points x and y can be joined by an m-chain in L, where m is a fixed integer, then there is a shortest m-chain from x to y in L.*

Proof: Let $X \subset L^{m+1}$ denote the set of m-chains from x to y in L. We show that X is closed and hence compact. The length function λ on m-chains is continuous on X, and therefore attains a minimum if X is closed.

Notice that for any $w, z \in L$, the set $st(w) \cap st(z)$ is empty if and only if there is no closed cell of L containing both w and z. Thus, if **z** $= (z_0, z_1, \ldots, z_m)$ is *not* an m-chain then for some i between 0 and $m-1$ we have $st(z_i) \cap st(z_{i+1}) = \emptyset$. Hence **z** has a neighbourhood disjoint from X, namely $(L \times \ldots L \times st(z_i) \times st(z_{i+1}) \times L \times \ldots \times L)$. So X is closed. □

Given K with $Shapes(K)$ finite, and a positive integer m, one can build a finite set of *"models"*. That is, connected complexes obtained by taking at most m (not necessarily distinct) simplices from $Shapes(K)$ and identifiying faces by isometries. Any subcomplex K_0 of K which can be expressed as the union of at most m closed cells must be isometrically isomorphic to one of these models. The existence of this finite set of models allows us to pass from the case of compact complexes (Lemma 1.4) to the case of interest, complexes with $Shapes(K)$ finite (Lemma 1.5).

Notice that K_0 with the induced metric from K is not in general isometric to the model with its intrinsic metric. However, since the length of a chain is defined in terms of the local metrics d_B, a given m-chain in K_0 and the corresponding m-chain in the model have the same length. This is the key to the following lemma.

Lemma 1.5: *If K is an M-simplicial complex with $Shapes(K)$ finite, and two points x and y can be joined by an m-chain in K, where m is a fixed integer, then there is a shortest m-chain from x to y in K.*

Proof: For any fixed pair of elements x and y there are only finitely many bipointed models, $(K'; x', y')$, for $(K_0; x, y)$ as K_0 runs over all subcomplexes of K which contain both x and y and can be expressed as the union of at most m closed cells. Thus, any m-chain from x to y in K corresponds to an m-chain of the same length from x' to y' in one of the finitely many models under consideration, and vice versa. The present lemma now follows by application of Lemma 1.4 to each of these models. □

Remark: Suppose that one were to weaken the condition that $Shapes(K)$ is finite, and require instead that there be uniform bounds on how small and thin the simplices in $Shapes(K)$ may be. Then a straightforward adaptation of the arguments given in Section 1.2 shows that this weaker assumption is sufficient to ensure that the intrinsic pseudometric on K is indeed a metric. However, under this weaker hypothesis Lemma 1.5 is no longer true, and in general if K is locally inifinite then it is *not* a geodesic metric space, as the following example shows.

Example: We construct a Euclidean 2–complex L as follows: L has a subcomplex L' consisting of three vertices $\{x\}, \{y\}$ and $\{z\}$, and two edges $e_y = \{x, y\}$ and $e_z = \{x, z\}$, both of which are metrized to have unit length. For every integer $n \geq 2$ we metrize a barycentrically subdivided 2–simplex $\sigma(n)$ as an isosceles triangle with two sides of unit length meeting at an angle $(\frac{\pi}{3} + \frac{\pi}{n})$. We then attach each $\sigma(n)$ to L' by identifying one of its edges of unit length with e_y, and the other with e_z. The result is a Euclidean simplicial complex L which is not a geodesic metric space. To see this notice that $d(y, z) = 1$ whereas every m–chain from y to z has length strictly larger than 1. In fact, there exist 2–chains C from y to z with $\lambda(C)$ arbitrarily close to $1 = d(y, z)$, so the conclusion of Lemma 1.5 fails to hold in this case.

1.4 Taut Chains and Local Geodesics

In this section we introduce the notion of a *taut chain*. Intuitively speaking, this is an m–chain whose length cannot be shortened by perturbation unless one allows the integer m to increase. In Lemma 1.6 we show that the minimising chains yielded by Lemma 1.5 are taut. Then in Section 1.5 we formalise the idea that the size of a taut chain is directly related to its length, and this leads to the desired bound on the size of chains which one must consider when seeking a geodesic segment between a given pair of points in K.

In the present section we also define what it means for an $m-$chain to be a local geodesic. Local geodesics do not play an essential role in the proof of Theorem 1.1, but they do (as the name suggests) describe the local behaviour of geodesic segments, and the understanding which this description provides proves to be important in later chapters. In fact, in Chapter 2 we show that in simply-connected complexes of non-positive curvature there is a one-to-one correspondence between geodesic segments and taut local geodesics. This observation provides a link between the results presented here and the work of Gersten and Stallings on non-positively curved 2–complexes [26], where they proved the existence of geodesics using a local criterion.

Before defining what it means for an $m-$chain to be taut, we must make some observations about small subcomplexes of K. Suppose that B and B' are closed simplices in K which have non-empty intersection, and consider $L = B \cup B'$ equipped with its intrinsic metric. (Notice that in general the intrinsic metric on L is not equal to restriction of the intrinsic metric on K.) Suppose that x and y lie in the same simplex of L. Then it follows immediately from the triangle inequalities for d_B and $d_{B'}$ that the line segment $[x, y]$ is a geodesic segment in L. (Here we are using the fact that K is simplicial, and thus if $x, y \in B \cap B'$ then the line segments joining x to y in B and B' coincide.) From this it follows that if $x \in B$ and $y \notin B$ then $d(x, y) = inf\{d_B(x, z) + d_{B'}(z, y) : z \in B \cap B'\}$. And since $z \mapsto d_B(x, z) + d_{B'}(z, y)$ is a continuous function on the compact set $B \cap B'$

it attains a minimum. Therefore L is a geodesic metric space, and the minimal $m-$chain associated to any geodesic segment has size at most 2.

Definition: An m-chain $C = (x_0, x_1, \ldots, x_m)$ in K is *taut* if it satisfies the following two conditions for $i = 1, \ldots, m-1$: Firstly, there is no simplex containing $\{x_{i-1}, x_i, x_{i+1}\}$. Secondly, if $x_{i-1}, x_i \in B(i)$ and $x_i, x_{i+1} \in B(i+1)$ then the concatenation of the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ is a geodesic segment in $L = B(i) \cup B(i+1)$.

Notice that if a chain is taut then only its first and last entries can lie in the interior of a top dimensional simplex of K.

Lemma 1.6: *For a fixed integer m, if C is an m-chain from x to y in K of minimal length then $p(C)$, the path in K determined by C, is the path determined by some taut n-chain with $n \leq m$.*

Proof: Let $C = (x_0, x_1, \ldots, x_m)$. Suppose that for some i there exists a simplex B containing x_{i-1}, x_i and x_{i+1}, and let C^i denote the $(m-1)-$chain obtained from C by deleting the entry x_i. The triangle inequality for d_B gives $d_B(x_{i-1}, x_{i+1}) \leq d_B(x_{i-1}, x_i) + d_B(x_i, x_{i+1})$, with equality if and only if x_i lies on the line segment $[x_{i-1}, x_{i+1}]$. Thus $\lambda(C^i) \leq \lambda(C)$. But $\lambda(C)$ is minimal, so in fact $\lambda(C^i) = \lambda(C)$, and hence x_i must lie on the line segment $[x_{i-1}, x_{i+1}]$. This implies that C^i determines the same path as C. We can repeat this procedure until no simplex of K contains three successive entries of the resulting chain — the first condition for tautness.

We now show that C satisfies the second condition for tautness. Let B and B' be any cells containing $\{x_{i-1}, x_i\}$ and $\{x_i, x_{i+1}\}$ respectively. Every geodesic segment in the complex $L = B \cup B'$ can be expressed as the concatenation of at most two line segments. So if the concatenation of the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ were not a geodesic segment in L, then we could replace x_i by some $x'_i \in B \cap B'$ to obtain an m−chain from x_0 to x_m which would be strictly shorter than C, contradicting the minimality of $\lambda(C)$. □

As an immediate consequence of Lemma 1.6 we obtain:

Corollary 1.7: $d(x,y) = inf\{\lambda(C) : C$ *a taut chain from* x *to* $y\}$.

We now turn our attention to local geodesics. In order to describe the local behaviour of PL paths we need a precise notion of angle between line segments in K. For this we introduce a spherical metric on the link of each point in K.

Let B be an n−simplex in M^n (the unique simply-connected, complete n−manifold of constant sectional curvature κ). Given $x \in B$, the tangent cone of B at x is defined to be $T_xB = \{v \in T_xM^n : exp_x(\epsilon v) \in B, some\ \epsilon > 0\}$. We identify $\mathbf{S^{n\text{-}1}}$ with the unit sphere in T_xM^n, and define $LK(x, B)$ to be $T_xB \cap \mathbf{S^{n\text{-}1}}$. This is a simplex in $\mathbf{S^{n\text{-}1}}$ if x is a vertex of B, the whole of $\mathbf{S^{n\text{-}1}}$ if x is an interior point of B, and otherwise a closed convex subset of a hemisphere in $\mathbf{S^{n\text{-}1}}$.

Now suppose $x \in B \subseteq K$. We define T_xB to be $T_{f_B(x)}f_B(B)$ and $LK(x, B)$ to be $LK(f_B(x), f_B(B))$.

Definition: The *geometric link* $LK(x, K)$ of x in K is defined to be the disjoint union of the cells $\{LK(x, B) : x \in B \subseteq K\}$, modulo the natural identifications. Specifically, $u \in LK(x, B)$ is identified with $v \in LK(x, B')$ if and only if the differential of $f_{B'}f_B^{-1}$ maps u to v. T_xK is defined similarly.

$LK(x, K)$ is a spherical complex, which can be made simplicial by subdivision. Further, since $Shapes(K)$ is finite, so too is $Shapes(LK(x, K))$. In particular we have an intrinsic metric on connected components of $LK(x, K)$.

Definition: An m-chain $C = (x_0, x_1, \ldots, x_m)$ in K is a *local geodesic* if $x_i \neq x_{i+1}$ for all i, and for $i = 1, \ldots, m-1$ the distance between the points of $LK(x_i, K)$ determined by the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ is at least

π. (If these points lie in different path components of $LK(x, K)$ then we say that they are an infinite distance apart.)

We now prove that if the path associated to an $m-$chain is a geodesic segment then the chain is a local geodesic. Intuitively this is clear, since the definition of a local geodesic is just a precise statement of the fact that the path determined by the given chain has no sharp corners. If such a sharp corner did exist then one could shorten the path by "cutting the corner". The following proof, which is due to Moussong [23], is simply a restatement of this fact in more precise language.

Lemma 1.8: *Let $C = (x_0, x_1, \ldots, x_m)$ be an $m-$chain in K, and suppose that $x_i \neq x_{i+1}$ for all i. If $p(C)$, the path determined by C, is a geodesic segment then C is a local geodesic.*

Proof: Let u and v denote the points of the spherical complex $LK(x_i, K)$ which are determined by the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ respectively.

If $d(u, v) < \pi$ in $LK(x_i, K)$ then there exists a sector S of the unit disc in $\mathbf{E}^2$, and a local isometry $\phi : S \to T_{x_i}L$ which maps the circular arc component of the boundary of S, which we call σ, to a piecewise geodesic path of length $< \pi$ joining u to v in $LK(x_i, K)$.

Consider the line in S which joins the endpoints of σ. This has length $2 - 3\delta$ for some positive δ. The map ϕ sends this line to a PL path from u to v in $T_{x_i}K$. The endpoints of the line segments in this path determine an $m-$chain $(a_0, a_1, \ldots, a_n)$ in $T_{x_i}L$.

On each cell of $T_{x_i}K$ we have the exponential map to the corresponding cell of K, and where two cells meet this is well defined on their common face. So for sufficiently small ϵ we have the following chain joining x_0 to x_m in K: $C' = (x_0, \ldots, x_i, exp_x(\epsilon a_0), \ldots, exp_x(\epsilon a_n), x_{i+1}, \ldots, x_m)$. The map ϕ is a local isometry, and we can choose ϵ so that the exponential map restricted to the ϵ–neighbourhood of the origin increases distances by at most a factor of $1/(1-\delta)$. Hence $\lambda(C') < \lambda(C) - \delta$, contradicting the fact that $p(C)$ is a geodesic segment. □

Remark: Suppose that $C = (x_0, \ldots, x_m)$ is taut. Lemma 1.6, together with the second condition for tautness, implies that for $i = 1, \ldots, m-1$ the distance between the points of $LK(x_i, B(i) \cup B(i+1))$ determined by the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ is at least π. In fact, if x_i is not a vertex of $B(i) \cap B(i+1)$ then this distance is exactly π, and it follows that the image of the concatenation of $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ under any local isometry into $M(\kappa)^n$ is a geodesic arc.

Example: Figure 1.1 illustrates the fact that a taut chain is not necessarily a local geodesic. Here K is a planar 2–complex with three 2–simplices, and the chain (a, b, c) is taut but not a local geodesic.

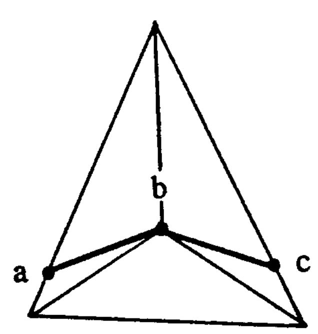

Figure 1.1: A taut chain which is not a local geodesic

1.5 The Growth of Taut Chains

In this section we prove that for large m the length of taut m–chains grows at least linearly with respect to m (Theorem 1.11). Combining this with Corollary 1.7, we see that the infimum in the definition of $d(x, y)$ need only be taken over m–chains with m less than a certain linear function of $d(x, y)$. This, as we explained at the beginning of Section 1.3, completes the proof that K is a geodesic metric space.

We begin by noting a simple fact about spherical geometry. We shall use the term "spherical" to describe complexes whose cells are modelled on a sphere of some fixed radius, which may not be 1. The usage should be clear from the context.

Claim: If P is the vertex of a (2–dimensional) geodesic triangle $\triangle$ in $\mathbf{S^n}$ and S is the boundary of an ϵ–neighbourhood of P in the usual (arclength) metric on $\mathbf{S^n}$, where ϵ is suitably small, then $\triangle \cap S$ is a geodesic arc on S.

Proof: The idea of the proof is to reduce to the case $n = 2$ where the result is clear. Consider $\mathbf{S^n}$ as the unit sphere in $\mathbf{R^{n+1}}$ centred at the origin O. Then S is an $(n-1)$-sphere of Euclidean radius $sin(\epsilon)$ whose centre, which we denote by Q, lies on the line OP. Let V be the three dimensional subspace of $\mathbf{R^{n+1}}$ determined by the vertices of $\triangle$. The intersection of $\mathbf{S^n}$ with V is a 2–sphere of unit radius centred at O. The line OP lies in V, so in particular Q does, and $S \cap V$ consists of those points on the unit 2–sphere in V which are a Euclidean distance $sin(\epsilon)$ from Q. Thus $S \cap V$ is a great circle on S and since $\triangle \cap S$ is an arc of this circle we are done. $\square$

Note that if in the above Claim we replace $\mathbf{S^n}$ by $M(\kappa)^n$ (the unique complete simply-connected n–manifold of constant curvature κ) then the conclusion still holds. In the hyperbolic case this can be seen most easily by thinking of P as the centre of the Poincaré disc model.

Lemma 1.9: *Let K be an M-simplicial complex with $Shapes(K)$ finite. Then there exists a constant ϵ, depending only on $Shapes(K)$, such that for every vertex P of K the set $S(P) = \{x \in K : d(x, P) = \epsilon\}$ is an* **S**-*simplicial complex, whose dimension is one less than that of K.*

Proof: We define ϵ to be the minimum of the following finite set of positive numbers:
$\{d_\sigma(v, \tau) : \sigma \in Shapes(K), v$ a vertex of σ, and τ a face of σ with $v \notin \tau\}$.

With this choice of ϵ it follows from Lemma 1.3 that for every simplex $B \subseteq K$ the metric d_B agrees with d, the intrinsic metric on K, in the ϵ–neighbourhood of each vertex of B. So $S(P)$ is the union of the sets $\{x \in B : d_B(P, x) = \epsilon, P \in B\}$. And by the preceding claim these are spherical simplices. □

Remark: Notice that there is a natural identification of $LK(P, K)$ with $S(P)$.

To complete the inductive step in our proof of Theorem 1.11 we shall need the following notion of radial projection from a vertex in K. The crucial property of this map is that it takes taut chains to taut chains (Lemma 1.10).

Suppose that x is a vertex of a simplex σ in M^n, and that $\epsilon < d(x, \sigma - st(x))$. Then there is a well defined notion of radial projection from x, taking $(\sigma - x)$ onto the intersection of σ with the boundary of the ϵ–ball about x.

Let P and ϵ be as in Lemma 1.9. For every simplex $B \ni P$ we let $f_B(B)$ and $f_B(P)$ play the roles of σ and x in the previous paragraph. We can then pull back the radial projection map, by means of the map f_B, to obtain a map from B onto $S(P) \cap B$. The pull-backs obtained in this way agree on common faces, and so combine to give a *radial projection map* from $St(P)$ to $S(P)$. We denote this map by ρ.

Lemma 1.10: *Let K and P be as in Lemma 1.9. If C is a taut chain in $st(P)$ which does not pass through P then the image of C under radial projection ρ from P to $S(P)$ is a taut chain.*

Proof: Let $C = (x_0, \ldots, x_m)$, and let $C' = (\rho(x_0), \ldots, \rho(x_m))$.

For every $k > 0$ the map ρ gives a 1–1 correspondence between the $k-$simplices of $S(P)$ and those $(k+1)-$simplices in $St(P)$ which meet P. If three sucessive entries $\rho(x_{i-1}), \rho(x_i), \rho(x_{i+1})$ of C' were to lie in some simplex of $S(P)$ then the set $\{x_{i-1}, x_i, x_{i+1}\}$ would be contained in the corresponding simplex of $St(P)$, contradicting the fact that C is taut. Thus C' satisfies the first condition for tautness.

To see that C' satisfies the second condition, consider simplices $\rho(B(i))$ and $\rho(B(i+1))$ which contain the line segments $[\rho(x_{i-1}), \rho(x_i)]$ and $[\rho(x_i), \rho(x_{i+1})]$ respectively. Then the simplices $B(i)$ and $B(i+1)$ contain the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ respectively. We can extend the line segments $[P, x_{i-1}]$ and $[P, x_i]$ until they meet $(B(i) - st(P))$, and the resulting line segments form two sides of a unique geodesic triangle in $B(i)$. (Here we mean a 2–dimensional triangle which is totally geodesic with respect to the metric $d_{B(i)}$.) We denote this triangle by Δ_i. In the same way, the line segments $[P, x_i]$ and $[P, x_{i+1}]$ determine a unique geodesic triangle Δ_{i+1} in $B(i+1)$. These triangles have a common edge, and we can map them into $M(\kappa)^2$ by a local isometry ϕ to obtain the planar 2–complex $k(x_i)$ shown in Figure 1.2.

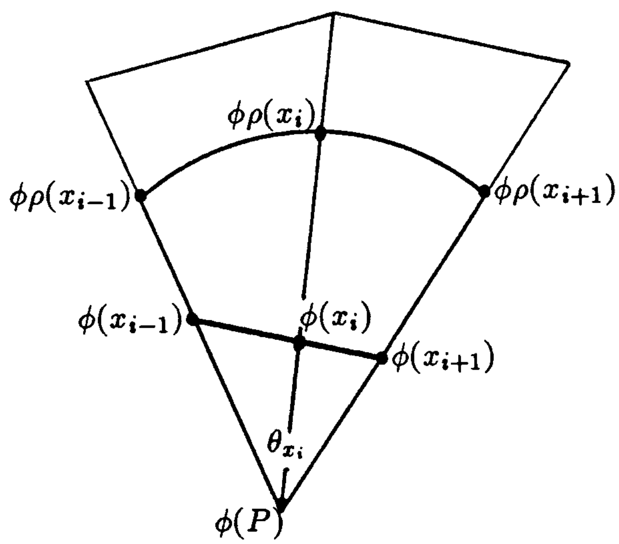

Figure 1.2: The complex $k(x_i)$

Because C is taut, the path in $L = B(i) \cup B(i+1)$ determined by the 3–chain $C_i = (x_{i-1}, x_i, x_{i+1})$ is a geodesic segment, so its image under ϕ is a geodesic arc in $M(\kappa)^2$. Moreover, the angle θ_{x_i} which this arc subtends at $\phi(P)$ must have length less than π, because otherwise the path in $k(x_i)$ determined by the 3–chain $(\phi(x_{i-1}, \phi(P), \phi(x_{i+1}))$ would pull back to a path from x_{i-1} to x_{i+1} in L of length less than $\lambda(C_i)$, contradicting the fact that the path determined by C_i is a geodesic segment.

To prove that C satisfies the second condition for tautness we must show that the path which the 3–chain $C_i' = (\rho(x_{i-1}), \rho(x_i), \rho(x_{i+1}))$ determines in the spherical complex $L' = \rho(B(i)) \cup \rho(B(i+1))$ is a geodesic segment. The map ϕ, which we defined in the previous paragraph, sends this path isometrically onto an arc of the circle of radius ϵ about $\phi(P)$, allowing us to express $\lambda(C_i')$ as a monotone function of the angle θ_{x_i}. If C_i' were not a geodesic segment in L' then there would be a shorter 3–chain $(\rho(x_{i-1}), \rho(y), \rho(x_{i+1}))$ from $\rho(x_{i-1})$

to $\rho(x_{i+1})$ in L'. But then the vertex angle θ_y in the corresponding planar 2–complex $k(y)$ would be strictly smaller than θ_{x_i}, which in turn would imply that the geodesic segment joining the images of x_{i-1} and x_{i+1} in $k(y)$ would pull back to a path from x_{i-1} to x_{i+1} in L of length less than $\lambda(C_i)$. The existence of such a path would contradict the fact that C_i is a geodesic segment in L. Hence C_i' must be a geodesic segment in L'. □

Remark: Using simliar methods one can show that the property of being a local geodesic is also preserved under the radial projection map ρ.

We are now in a position to prove Theorem 1.11. The idea of the proof is to use induction to show that a large taut chain cannot be contained in the neighbourhood of any point. We first use a colouring argument to localise the problem to the star of a single vertex, then we radially project from the vertex to obtain a taut chain in a complex of lower dimension, and use our inductive hypothesis to complete the proof.

Theorem 1.11: *Let K be an M-simplicial complex with $Shapes(K)$ finite. Then there exist constants N and α, depending only on $Shapes(K)$, such that every taut chain of size at least N has length at least α.*

Proof: We proceed by induction on the dimension of K. It is important to the induction that the constants N and α depend only on $Shapes(K)$ and not on the global structure of K itself. For 1–dimensional complexes the result is clear, and the two dimensional case is included as an example in Section 1.7. Assume that the result holds for dimension $n-1$, and that K has dimension n.

In Lemma 1.9 we defined a constant ϵ, depending only on $Shapes(K)$, such that the ϵ–neighbourhood of any vertex $P \in K$ is contained in $st(P)$. We let $\eta_0 = \epsilon/3$ and paint the η_0–neighbourhood of each vertex of K with colour γ_0. Then, for $i = 1, 2, \ldots, n-2$ we paint the η_i–neighbourhood of each i–simplex in K with colour γ_i (except that we do *not* paint over points already coloured with a colour of lower index). The constants η_i (which we define below) depend only on $Shapes(K)$.

We define the constants η_i inductively. At each stage we require that η_i is small enough to ensure that the γ_i–regions corresponding to distinct i–cells are disjoint. It then follows that each point $x \in K$ which is painted with colour γ_i lies in the η_i–neighbourhood of a *unique* i–simplex, and we make the additional requirement that the distance from x to the link of the barycentre of this unique i–simplex must be at least η_i. In the previous paragraph we defined η_0 so that these conditions hold in the case $i = 0$. Suppose now that we have defined η_{i-1} with the desired properties. Consider a simplex $\sigma \in Shapes(K)$ of dimension at least $i + 1$, and let τ^i be an i–dimensional face of σ with barycentre $b_0(\tau^i)$. The complement in τ^i of the open η_{i-1}–neighbourhood of its $(i-1)$–skeleton is compact, and we have chosen η_{i-1} so that this set is non-empty. Let η_{τ^i} denote the distance from this set to the link of $b_0(\tau^i)$ in σ. Then $\eta_i = \frac{1}{3}min\{\eta_{\tau^i} : \tau^i \subseteq \sigma, \sigma \in Shapes(K)\}$ has the desired properties.

Notice that $\eta_{i+1} \leq \eta_i$ for all i. Thus if $x \in K$ is painted with colour γ_i, and b_0 is the barycentre of the unique i–simplex responsible for this colouring, then η_{n-2} is a lower bound on the distance from x to $(St(b_0) - st(b_0))$, the link of b_0 in K.

After we have painted a neighbourhood of the $(n-2)$–skeleton of K in this way, we paint the remainder of K white, which we call colour γ_∞.

Given an m–chain $C = (x_0, x_1, \ldots, x_m)$ we shall refer to the open line segments (x_i, x_{i+1}) as the intervals of C. If $i \neq 0$ or $(m-1)$ then we call (x_i, x_{i+1}) an interior interval. The colouring on K induces a colouring on the intervals of any chain in K by the rule: paint each interval with the colour γ_j

where $j = min\{k :$ some point of $[x_i, x_{i+1}]$ is coloured $\gamma_k\}$.

Every interval of C has the same length as a geodesic arc in some model simplex $\sigma \in Shapes(K)$. If C is taut then only its initial and terminal entries can lie in the open star of an adjacent entry, so for an interior interval the corresponding geodesic segment in σ must have its endpoints in distinct faces of $\partial\sigma$. Moreover, if the interval is painted white then the corresponding geodesic segment must be contained in the complement of the $\eta_{n-2}-$neighbourhood of the $(n-2)-$skeleton of σ. There is a lower bound on the length of such geodesic segments, and taking the minimum of these lower bounds as σ ranges over $Shapes(K)$ we obtain a lower bound, which we call ℓ, on the length of interior intervals of C which are painted white.

If a taut m-chain $(m \geq 3)$ is not entirely contained in the open star of the barycentre of any simplex in K, then either it must contain an interior interval which is painted white, or else it must contain a coloured interval and a point which is not in the star of the barycentre of the cell responsible for that colouring. Thus if we set $\alpha = min\{\ell, \eta_{n-2}, |2\pi\kappa|\}$ then Theorem 1.11 reduces to the following Claim.

Claim: Let κ denote the curvature of the model simplices for K. Suppose that b_0 is the barycentre of a simplex in K, and that C is a taut $m-$chain contained in $st(b_0)$. If $\kappa \leq 0$, or $\kappa > 0$ and $\lambda(C) < 2\pi\kappa$, then $m \leq N$, where N is a constant depending only on $Shapes(K)$.

Proof of claim: We wish to reduce to the case where b_0 is a vertex of K. If b_0 is not a vertex then we change the underlying simplicial structure of K by adding a vertex at b_0 and forming the simplicial join of b_0 with its link. Then for each simplex $B \subseteq st(b_0)$ we must add to $Shapes(K)$ the simplices resulting from the corresponding subdivision of $\sigma(B)$. (However, we do not delete $\sigma(B)$

from $Shapes(K)$, since in general there will exist simplices $B' \subseteq (K - st(B))$ with $\sigma(B) = \sigma(B')$, and these simplices have not been subdivided.) Let $K'(b_0)$ denote the resulting metric simplicial complex. Notice that as b_0 varies there are only finitely many possibilities for $Shapes(K'(b_0))$, and the set of these possibilities, which we denote $\{S_1, \ldots, S_r\}$, depends only on $Shapes(K)$.

Fix a barycentre b_0, and suppose that $Shapes(K'(b_0)) = S_j$. Let $C = (x_0, x_1, \ldots, x_m)$ be a taut $m-$chain in K satisfying the hypotheses of the Claim. Each of the line segments $[x_i, x_{i+1}] \subseteq K$ will be a PL path in $K'(b_0)$. Let $C^i = (x_i, x_i^1, \ldots, x_i^{n_i}, x_{i+1})$ denote the unique chain of minimal length representing this path, and let $\tilde{C} = (x_0, x_0^1, \ldots, x_0^{n_0}, x_1, \ldots, x_{m-1}^{n_{m-1}}, x_m)$. This a taut $m'-$chain in $K'(b_0)$ with $\lambda(\tilde{C}) = \lambda(C)$ and $m' \geq m$. Suppose that the Claim were true in the case where the barycentre in the statement is a vertex. Then there would exist a constant N_j, depending only on $S_j = Shapes(K'(b_0))$, such that $m' \leq N_j$. We would then have a constant $N = max\{N_1, \ldots, N_r\}$ depending only on $Shapes(K)$ such that $m \leq m' \leq N_j \leq N$ as required. So it is enough to consider the case where b_0 is a vertex of K.

Assume that this is so, and let C be as in the statement of the Claim. If $p(C)$, the path determined by C, were to pass through b_0 then because b_0 is a vertex it would have to occur as an entry in C. The first condition for tautness would then imply that C had size at most two. We are only interested in large chains, so we assume $C = (x_0, x_1, \ldots, x_m)$ with $m > 2$ and hence $p(C)$ does not pass through b_0. If we radially project from b_0 onto the spherical complex $S(b_0)$ (which we defined in Lemma 1.8) then, according to Lemma 1.9, the image of C under this map is a taut m-chain. We denote this $m-$chain in $S(b_0)$ by C'.

Let $B(i)$ be a simplex containing $[x_i, x_{i+1}]$. We can extend the line segments

$[b_0, x_i]$ and $[b_0, x_{i+1}]$ until they meet $(B(i) - st(b_0))$, and the resulting line segments form two sides of a unique geodesic triangle in $B(i)$. (Here we mean a 2–dimensional triangle which is totally geodesic with respect to the metric $d_{B(i)}$.) We denote this triangle by $\triangle_i$. For each i the triangles $\triangle_i$ and $\triangle_{i+1}$ have a common edge and we can map these triangles isometrically into $M(\kappa)^2$ to obtain the complex shown in Figure 1.3. Notice that the in general the triangles $\triangle_i$ are not 2–simplices in the simplicial structure on K, and we do not have a bound on the vertex angles.

Because the chain C is taut the image of $p(C)$ under this map is a geodesic arc in $M(\kappa)^2$ of length $\lambda(C)$. Further, this map takes C' isometrically onto an arc of the circle of radius ϵ about the image of b_0. We are assuming that if $\kappa > 0$ then $\lambda(C) < 2\pi\kappa$, so the image of $p(C)$ in $M(\kappa)^2$ cannot be a closed geodesic, and hence it subtends an angle $< 2\pi$ at the image of b_0. Thus the image of C' in $M(\kappa)^2$, and hence C' itself, has length less than $L_\kappa(S_\epsilon)$, the length of a circle of radius ϵ in the plane of constant curvature κ. By induction on $n = dim(K)$ there exists an integer $N(b_0)$, depending only on $Shapes(S(b_0))$, such that every taut chain in the $(n-1)$–dimensional complex $S(b_0)$ of size at least $N(b_0)$ has length at least $L_\kappa(S_\epsilon)$. Hence $size(C) = size(C') < N(b_0)$.

There are only finitely many possibilities for $Shapes(S(b_0))$, and the set of these possibilities depends only on $Shapes(K)$. So setting N equal to the minimum of the corresponding integers $N(b_0)$ finishes the proof. □

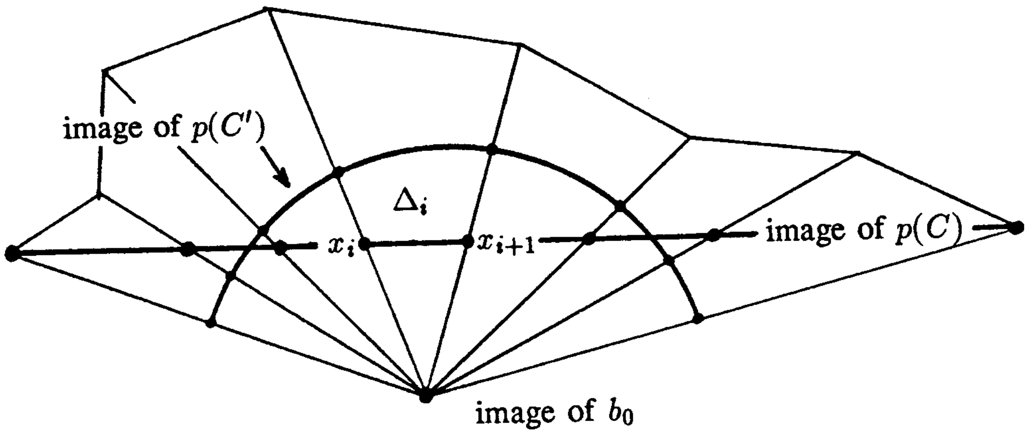

Figure 1.3. The image of the developing map

1.6 The Completeness of the Metric

Theorem 1.12: *If $Shapes(K)$ is finite then d, the intrinsic metric on K, is complete.*

Proof: Let (x_n) be a Cauchy sequence in (K, d), and $\epsilon > 0$. Fix an integer R such that $d(x_n, x_m) < \epsilon$ for all $m, n \geq R$.

By Theorem 1.11, there is an upper bound, N say, on the size of taut chains from x_R to $x_m (m > R)$. It follows that any geodesic segment joining x_R to x_m in K lies is some subcomplex $K_0(m)$ which can be expressed as the union of at most N closed simplices.

Because $Shapes(K)$ is finite there are only finitely many pointed models for $(K_0(m), x_R)$. Let (K_0', x_R') be a model which occurs for infinitely many m, and consider the image in K_0' of the sequence (x_m). Because the model is compact this has a convergent subsequence, which we denote (x'_{n_i}), with limit

x'. We shall prove that (x_{n_i}), the corresponding subsequence in K, converges to a point in the preimage of x'.

We first describe the preimage of x' in K. In general K will have infinitely many bipointed subcomplexes $(K_\lambda; x_R, x_\lambda)$ giving rise to the bipointed model $(K_0'; x_R', x')$. However, the pair (B_0, x_λ), where B_0 is the unique simplex of K which contains x_λ in its interior, is well defined up to isometry type. Suppose $B \in Shapes(K)$ has a face isometric to B_0. For each choice ϕ of isometry from B_0 to a face of B we can measure the distance from $\phi(x_\lambda)$ to $(B - st(\phi(x_\lambda))$. (So if B_0 is a top dimensional cell this is equal to $d(x_\lambda, \partial B_0)$.) Let a_1 be the minimum of these distances over all choices of B and ϕ, and let a_2 be the minimum distance between distinct points in the orbit of x_λ under the action of the isometry group of B_0.

If we set $\epsilon = \frac{1}{3} min\{a_1, a_2\}$ then the ϵ–balls about the x_λ are disjoint. In fact the distance between distinct balls is at least ϵ. Further, for every λ the closed ball $B_\epsilon(x_\lambda)$ is contained in $st(x_\lambda)$, so by Lemma 1.3 the line segment joining x_λ to any point in $B_\epsilon(x_\lambda)$ is a geodesic segment in K. The definition of x' implies that the sequence (x_{n_i}) must lie in the union of the $B_\epsilon(x_\lambda)$ after some finite stage, and since this sequence is Cauchy it must eventually be contained in a single ball. Let x denote the x_λ corresponding to this ball. For large n we have:

$$
\begin{aligned}
d(x, x_{n_i}) &= length\ of\ the\ line\ segment\ [x, x_{n_i}] \\
&= length\ of\ the\ line\ segment\ \left[x', x'_{n_i}\right] \\
&\to 0 \quad as \quad n_i \to \infty.
\end{aligned}
$$

Hence the subsequence (x_{n_i}) converges to x, and since the sequence (x_m) is Cauchy it too converges to x. $\square$

1.7 The Two Dimensional Case

To illustrate the ideas involved in the proof of Theorem 1.11 we now restrict our attention to the 2–dimensional case where many of the details are considerably easier. For simplicity we shall only consider Euclidean complexes — the general case requires only that one consider pictures drawn in $M(\kappa)^2$ rather than $\mathbf{E}^2$.

Throughout this section K denotes a connected 2–dimensional **E**-simplicial complex with $Shapes(K)$ finite and d denotes its intrinsic metric.

We say that an m–chain $C = (x_0, \ldots, x_m)$ in K has *property* τ if it satisfies the following conditions: Firstly, there is no simplex in K containing three successive entries of C. Secondly, if B and B' are 2–simplices in K which contain the line segments $[x_{i-1}, x_i]$ and $[x_i, x_{i+1}]$ respectively, and x_i lies in the interior of an edge common to B and B', then the images of x_{i-1}, x_i and x_{i+1} under any local isometry from $B \cup B'$ into $\mathbf{E}^2$ are colinear.

Property τ is weaker than the property of being taut (which we definied in Section 1.4).

In Section 1.3 we showed that for a fixed integer m there exists a shortest m-chain joining any two points in K. Let C be such a chain. We can delete entries of C until no three sucessive entries lie in the same simplex. The resulting chain must then satisfy property τ, since otherwise we could perform simple moves to produce a shorter chain of the same size and with the same endpoints, as illustrated below.

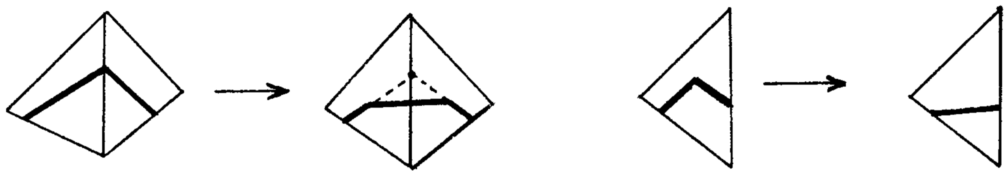

Figure 1.4: Shortening a chain which does not have property τ.

The existence of geodesic segments in the 2–dimensional case now follows from the following lemma.

Lemma 1.13: *There exist constants N and α such that every chain in K which has property τ and size greater than N has length at least α.*

Proof: Choose ϵ small enough to ensure that the ϵ–neighbourhoods of distinct vertices of K are disjoint. Colour these neighbourhoods black and the remainder of K white. Note that there is a lower bound ℓ_1 on the distance from any black region to the link of the corresponding vertex, and a lower bound ℓ_2 on the length of line segments which lie entirely in the white region of some 2–simplex, and have endpoints in the boundary of that simplex.

Let $C = (x_0, \ldots, x_m)$ be an m–chain in K with propery τ. Suppose $m > 3$ and that $p(C)$, the PL path determined by C, is not contained in the open star of any vertex of K. Then either there exists some $i \in \{1, \ldots, m-2\}$ such that the line segment $[x_i, x_{i+1}]$ lies entirely in the white region of K, or else $p(C)$ must contain a point in a black region and a point in the link of the corresponding vertex. In either case C must have length at least $\alpha = min\{\ell_1, \ell_2\}$. So it is enough to show that if m is sufficiently large then $p(C)$ cannot be contained in the open star of any vertex.

For this one simply observes that if $p(C)$ were contained in $st(b_0)$ then we could join each line segment $[x_i, x_{i+1}]$ to b_0 and map the resulting triangles isometrically into $\mathbf{E}^2$ to obtain a complex of the type shown in Figure 1.3. Because $Shapes(K)$ is finite there is a lower bound, θ say, on the angles subtended at the vertex, so taking $N = \pi/\theta + 2$ we are done. □

2. Non-Positive Curvature in Piecewise Euclidean Complexes

In this chapter we prove the following special case of the Main Theorem:

Theorem : *If K is a simply connected piecewise Euclidean complex of type A or B then the following 11 conditions are equivalent:*

Global conditions:

I) *K has unique geodesic segments.*

II) *K satisfies CAT(0) globally.*

III) *K satisfies CN globally.*

IV) *The metric on K is convex.*

V) *Every geodesic triangle in K has non-positive excess.*

Local conditions:

VI) *K has unique geodesic segments locally.*

VII) *K satisfies CAT(0) locally.*

VIII) *K satisfies CN locally.*

IX) *The metric on K is convex locally.*

X) *Every point of K has a neighbourhood such that any geodesic triangle contained in that neighbourhood has non-positive excess.*

XI) *K satisfies the link condition.*

The proof of this theorem involves a number of auxillary results about the geometry of piecewise Euclidean complexes, which are of interest in their own right. These include not only the theorems stated in the introduction, but some of the supporting lemmas aswell.

The most difficult step in the proof is in passing from the local to the global situation. This we do in Section 2.4, modelling our argument on that used by Milnor in his book on Morse theory [22] to prove the Cartan-Hadamard Theorem, and subsequently by Stone [27] to prove the corresponding result for PL manifolds. Although the same strategy of proof works here, the details of our proof have little in common with those of the original, due to the absence of both local compactness and a smooth structure.

As a postscript we show in Section 2.5 that geodesic segments in a piecewise Euclidean complex of non-positive curvature can be extended indefinitely, provided that during the extension the endpoint of the extended geodesic does not have a contractible link.

Terminology

We recall the definition of the regularity conditions A and B:

A) *Every subset of K which is closed and bounded in the intrinsic pseudometric is compact.*

B) $Shapes(K)$ *is finite.*

We do not require complexes of type A to be finite dimensional, nor do we require complexes of type B to be locally compact. (In the literature spaces satisfying condition A are sometimes called proper.)

As we explained in the introduction, our main interest is in complexes of type B, but some of our proofs require knowledge of the corresponding result for spaces of type A, and including spaces of type A into our development requires only a minimal amount of extra work.

For the remainder of Chapter 2 the letter K, without further qualification, shall denote a piecewise Euclidean complex of type A or B.

In Chapter 1 we descibed paths combinatorially using $m-$chains. In this chapter we adopt a more analytic approach, and to do so we use the following definitions.

Definition: A *geodesic segment* joining x to y in K is a PL path $\alpha : [0,1] \to K$ with $\alpha(0) = x, \alpha(1) = y$ which is parameterised proportional to arc length, and has length $L(\alpha) = d(x,y)$.

Definition: A PL path $\beta : [0,1] \to K$ is a *local geodesic* if it is parameterised proportional to arc length and every $t \in [0,1]$ has a neighbourhood $[s,s']$ such that $d_K(\beta(t'), \beta(t'')) = L(\beta)|t' - t''|$ for all $t', t'' \in [s,s']$.

Our insistence that all paths under consideration are defined on $[0,1]$ and parameterised proportional to arc length can prove inconvenient, but it allows us to describe convexity properties of the metric by means of inequalities involving geodesic segments.

Any PL path $\alpha : [0,1] \to K$ has a unique minimal expression as the concatenation of line segments in K, and the endpoints of these line segments form an m–chain. It is easy to check that α is a local geodesic if and only if this m–chain is a local geodesic (as defined in Section 1.3). If α is a geodesic segment then this chain is also taut.

2.1 The Convexity of the Metric in Spaces with Unique Geodesic Segments.

In this section we prove that if K has unique geodesic segments (i.e. every pair of points can be joined by a unique geodesic segment in K) then every local geodesic in K is a geodesic segment, and the intrinsic metric is convex in the following sense: If α and β are geodesic segments in K then the function $f_{\alpha,\beta} : [0,1] \to K$ given by $f(t) = d(\alpha(t), \beta(t))$ is convex. The latter assertion follows easily from the following result, the proof of which occupies the remainder of this section.

Theorem 2.1: *If K has unique geodesic segments, and α_0 and α_1 are geodesic segments in K with $\alpha_0(0) = \alpha_1(0)$ then*

$$d(\alpha_0(t), \alpha_1(t)) \leq t \cdot d(\alpha_0(1), \alpha_1(1)) \quad \forall t \in [0,1].$$

The first step in the proof of Theorem 2.1 is to show that geodesic segments in K vary continuously with their endpoints (this statement will be made precise in Lemma 2.3). To prove this fact for complexes of type B we introduce the notion of a *corridor.*

Definition: An $M(\kappa)$*–corridor* is an $M(\kappa)$–simplicial complex which can be expressed as the union of finitely many closed simplices $\{\sigma_1, \ldots, \sigma_m\}$ such that the following conditions hold for $i = 1, \ldots, m-1$: Firstly, $\tau_i = \sigma_i \cap \sigma_{i+1}$ is a non-empty proper subface of both σ_i and σ_{i+1}. Secondly $\tau_i \cap \tau_{i+1}$ (which may be empty) is properly contained in both τ_i and τ_{i+1}. Thirdly, $\sigma_i \cap \sigma_{i+r} = \bigcap_{j=0}^{r-1} \tau_{i+j}$ for all $r > 0$. We will denote such a corridor by $\Gamma = (\sigma_1, \ldots, \sigma_m)$.

This third condition ensures that any path in Γ which has one of its endpoints in each of σ_1 and σ_m must intersect all of the closed simplices σ_i. Notice that Γ is compact and hence a geodesic metric space of type A.

We shall be concerned with corridors which arise in the following way: Suppose that $C = (x_0, \ldots, x_m)$ is an m–chain in an $M(\kappa)$–simplicial complex K, and that no three sucessive entries of C lie in any closed simplex of K (we shall call an m–chain with this property *minimal*). We let $B(i)$ denote the unique closed simplex of smallest dimension which contains the line segment $[x_{i-1}, x_i]$, and consider the disjoint union of the m (not necessarily distinct) simplices $\sigma(B(i)) \in Shapes(K)$, modulo the equivalence relation generated by: $f_{B(i)}(x) \sim f_{B(i+1)}(x)$ whenever $x \in B(i) \cap B(i+1)$. (The maps f_B are those defined in the definition of a metric simplicial complex in Section 1.1.) It is easy to check that because C is minimal this is an $M(\kappa)$–corridor. We denote this corridor Γ_C. If C is the unique taut chain associated to a geodesic segment α then we denote this corridor by $\Gamma(\alpha)$.

Definition: An $M(\kappa)-$corridor Γ is said to *occur in* K if $\Gamma = \Gamma_C$ for some minimal $m-$chain C in K.

Remark: In general the map from Γ_C into K induced by $f^{-1}_{B(i)} : \sigma(B(i)) \to B(i)$ is *not* an injection.

If K is of type B then for every integer m there are only finitely many isometry types of corridors of length m which can occur in K. As in Section 1.3, we represent these isometry types by *model* corridors, and according to Theorem 1.11 any geodesic segment in K whose length is bounded by a certain linear function of m is the image of a geodesic segment in one of these models.

If $\kappa \leq 0$ then the corridor Γ has unique geodesic segments. This follows by application of the following lemma to Γ and its sub-corridors.

Lemma 2.2: *Suppose* $\Gamma = (\sigma_1, \ldots \sigma_m)$ *is an* $M(\kappa)-$*corridor where* $\kappa \leq 0$. *If* $x \in \sigma_1$ *and* $y \in \sigma_m$ *then there is a unique geodesic segment from* x *to* y *in* Γ.

Proof: In Theorem 1.1 we proved that there exists a geodesic segment from x to y in Γ, so the only question is that of uniqueness. Suppose that there are two geodesic segments from x to y, given by distinct m-chains $C_1 = (a_0, \ldots, a_m)$ and $C_2 = (b_0, \ldots, b_m)$, where $a_0 = b_0 = x, a_m = b_m = y$ and $a_i, b_i \in \sigma_i \cap \sigma_{i+1}$. Let c_i be the midpoint of the line segment $[a_i, b_i]$. The simplex σ_i is isometric to a convex subset of $M(\kappa)^n$ for some n. Hence

$$d_{\sigma_i}(c_i, c_{i+1}) \leq \frac{1}{2}(d_{\sigma_i}(a_i, a_{i+1}) + d_{\sigma_i}(b_i, b_{i+1}))$$

with strict inequality in the case $i = min\{j : a_j \neq b_j\}$. Summing over i shows that $C = (c_0, \ldots, c_m)$ is an $m-$chain from x to y in Γ whose length satisfies the inequality:

$$\lambda(C) < \frac{1}{2}(\lambda(C_1) + \lambda(C_2)) = d_\Gamma(x,y)$$

which is a contradiction. □

We now prove that if K has unique geodesic segments then they vary continuously with their endpoints. For paths $\alpha, \beta : [0,1] \to K$ we use the notation $\|\alpha - \beta\| = sup\{d(\alpha(t), \beta(t)) : t \in [0,1]\}$.

Lemma 2.3: *Suppose that K has unique geodesic segments. Given $x, y \in K$, and sequences of points $x_i \to x$ and $y_i \to y$, let α denote the geodesic segment from x to y and let α_i denote the geodesic segment from x_i to y_i. Then $\|\alpha - \alpha_i\| \to 0$ as $i \to \infty$.*

Proof: First we prove the weaker statement that $\alpha_i(t) \to \alpha(t)$ as $i \to \infty$ for all t, an easy compactness argument then shows that the convergence is uniform.

Case 1: Suppose that K is of type A. Then the closed ball of radius $2d(x,y)$ about x is compact. After some finite stage all the α_i must be contained in this set. Suppose, for contradiction, that for some t it is not true that $d(\alpha_i(t), \alpha(t)) \to 0$ as $i \to \infty$. Then by compactness $(\alpha_i(t))$ has a convergent subsequence $(\alpha_j(t))$ with limit $z \neq \alpha(t)$. So we have

$$d(x_j, z) - d(x_j, \alpha_j(t)) \to 0$$
$$d(y_j, z) - d(y_j, \alpha_j(t)) \to 0 \qquad as\ j \to \infty.$$

But on the other hand,

$$d(x_j, \alpha_j(t)) = t\, d(x_j, y_j) \to t\, d(x,y)$$
$$d(y_j, \alpha_j(t)) = (1-t)\ d(x_j, y_j) \to (1-t)\ d(x,y) \qquad as\ j \to \infty.$$

So z satisfies

$$d(x,z) = \lim_{j\to\infty} d(x_j, \alpha_j(t)) = t\ d(x,y) \quad \text{and}$$
$$d(y,z) = \lim_{j\to\infty} d(y_j, \alpha_j(t)) = (1-t)\ d(x,y).$$

But $\alpha(t)$ is the unique point of K which satisfies these equations. Hence $z = \alpha(t)$, contrary to hypothesis.

Case 2: Now suppose that K is of type B. After some finite stage $x_i \in st(x)$ and $y_i \in st(y)$. Hence there is an upper bound on the distances $d(x_i, y_i)$, and so by Theorem 1.11 there exists an integer N such that each α_i is the path determined by a taut chain of size at most N. Thus α_i is the image of a geodesic segment $\tilde{\alpha}_i$ in at least one of the finitely many bipointed model corridors $(\tilde{x}, \tilde{y}; \sigma_0, \ldots, \sigma_N)$, where $\sigma_i \in Shapes(K)$ and $\tilde{w}$ denotes the image of a point $w \in K$ under the identification of a corridor in K with the corresponding model corridor.

Suppose that μ is a model which contains infinitely many of the model geodesic segments $\tilde{\alpha}_i$, and let (α_j) denote corresponding subsequence of (α_i). We proved in Lemma 2.2 that μ satisfies the hypotheses of Case 1, so the $\tilde{\alpha}_j$ converge pointwise to the unique geodesic segment $\tilde{\alpha}^\mu$ from $\tilde{x}$ to $\tilde{y}$ in μ. Hence

$$d_K(x_j, y_j) = L(\alpha_j) = L(\tilde{\alpha}_j) = d_\mu(\tilde{x}_j, \tilde{y}_j)$$

$$\to d_\mu(\tilde{x}, \tilde{y}) = L(\tilde{\alpha}^\mu) \quad as\ j \to \infty.$$

But $d_K(x_j, y_j) \to d_K(x, y)$ as $j \to \infty$, so $d_\mu(\tilde{x}, \tilde{y}) = d_K(x, y)$ and hence $\tilde{\alpha}^\mu$ models the unique geodesic segment α from x to y in K. In particular, because α is unique, any occurence of the bipointed model μ in K must contain α. Also notice that the image in K of the geodesic segment from $\tilde{\alpha}_j(t)$ to $\tilde{\alpha}^\mu(t)$ in μ is a path of the same length from $\alpha_j(t)$ to $\alpha(t)$ in K. Hence $d_\mu(\tilde{\alpha}^\mu(t), \tilde{\alpha}_j(t)) \geq d_K(\alpha(t), \alpha_j(t))$ for all t.

There are only finitely many models μ, so for sufficiently large integers R we can decompose $(\alpha_i)_{i>R}$ into finitely many infinite subsequences, each of which consists of geodesic segments which can be modelled in some fixed μ. For each such subsequence (α_j^μ) and for every $t \in [0, 1]$

$$d_K\left(\alpha(t), \alpha_j^\mu(t)\right) \leq d_\mu\left(\tilde{\alpha}^\mu(t), \tilde{\alpha}_j^\mu(t)\right) \to 0 \quad as\ j \to \infty.$$

Hence the sequence (α_i) converges pointwise to α.

Uniform Convergence: Assume that i is sufficiently large so that $L(\alpha_i) < L(\alpha) + 1$. Given ϵ and $t \in [0,1]$, we have proved the existence of an integer $N(t)$ such that $d(\alpha(t), \alpha_i(t)) < \epsilon$ whenever $i > N(t)$. Hence

$$\begin{aligned} & d(\alpha(t+\delta), \alpha_i(t+\delta)) \\ & \quad \leq d(\alpha(t+\delta), \alpha(t)) + d(\alpha(t), \alpha_i(t)) + d(\alpha_i(t), \alpha_i(t+\delta)) \\ & \quad < |\delta|\ L(\alpha) + \epsilon + |\delta|\ L(\alpha_i) \\ & \quad < 2\epsilon \qquad\qquad \textit{whenever} \quad |\delta| < \frac{\epsilon}{2L(\alpha)+1} \end{aligned}$$

The result now follows from the compactness of $[0,1]$. □

Remark: More generally, the conclusion of Lemma 2.3 holds in any open subset of K with the property that any two points in that subset can be joined by a unique geodesic segment, and that this geodesic segment lies entirely within the given subset. The same remark applies to Lemma 2.4. This importance of this observation is that it allows us to appeal to Lemmas 2.3 and 2.4 when dealing with complexes which have unique geodesic segments locally. (See Section 2.3.)

We shall use Lemma 2.3 to reduce the proof of Theorem 2.1 to the case where the geodesic segments under consideration are uniformly close. We now give the proof in this restricted setting.

Lemma 2.4: *If K has unique geodesic segments and α is a geodesic segment in K then there exists $\epsilon > 0$ such that any geodesic segment β with $\beta(0) = \alpha(0)$ and $\|\alpha - \beta\| < \epsilon$ satisfies*

$$d(\alpha(t), \beta(t)) \leq t \cdot d(\alpha(1), \beta(1)) \quad \forall t \in [0,1].$$

Proof: Represent α and β by the taut chains $(a_0, \ldots, a_m)$ and $(b_0, \ldots, b_r)$ respectively. Choose η sufficiently small so that $B_{4\eta}(a_i) \subset st(a_i)$ and the open balls $B_{4\eta}(a_i)$ are disjoint. Then choose $\epsilon < \eta$ so that for every i the ϵ–neighbourhood of each compact arc $[a_i, a_{i+1}] - (B_{4\eta}(a_i) \cup B_{4\eta}(a_{i+1}))$ is contained in $st(a_i) \cap st(a_{i+1})$. Decompose $\{b_0, \ldots, b_r\}$ into sets

$$S_0 = \{b_0, b_1, \ldots, b_{n_0}\}, S_1 = \{b_{n_0}, b_{n_0+1}, \ldots, b_{n_1}\}, \ldots, S_m = \{b_{n_{m-1}}, \ldots, b_r\}$$

such that $S_i \subseteq st(a_i)$. (With our choice of η and ϵ this can be done for any chain which determines a path in K uniformly ϵ–close to α.)

Let B_j be a closed simplex of K which contains the line segment $[b_j, b_{j+1}]$. If $b_j \in st(a_i)$ then $B_j \cap st(b_j) \subseteq B_j \cap st(a_i)$ is non-empty. In particular if $j = n_i$ then there is a unique geodesic triangle (with respect to the metric d_{B_j}) in B_j with vertices $\{a_i, a_{i+1}, b_{n_i}\}$. We can map this (possibly degenerate) triangle isometrically into $\mathbf{E}^2$ to obtain the geodesic triangle Δ_i shown in Figure 2.1. Similarly, for every $b_j \in S_i$ we have a unique geodesic triangle $\Delta(b_j, b_{j+1}, a_i)$ in B_j. We map each of these triangles isometrically into $\mathbf{E}^2$, as indicated in Figure 2.1. Here x' denotes the image in $\mathbf{E}^2$ of $x \in K$.

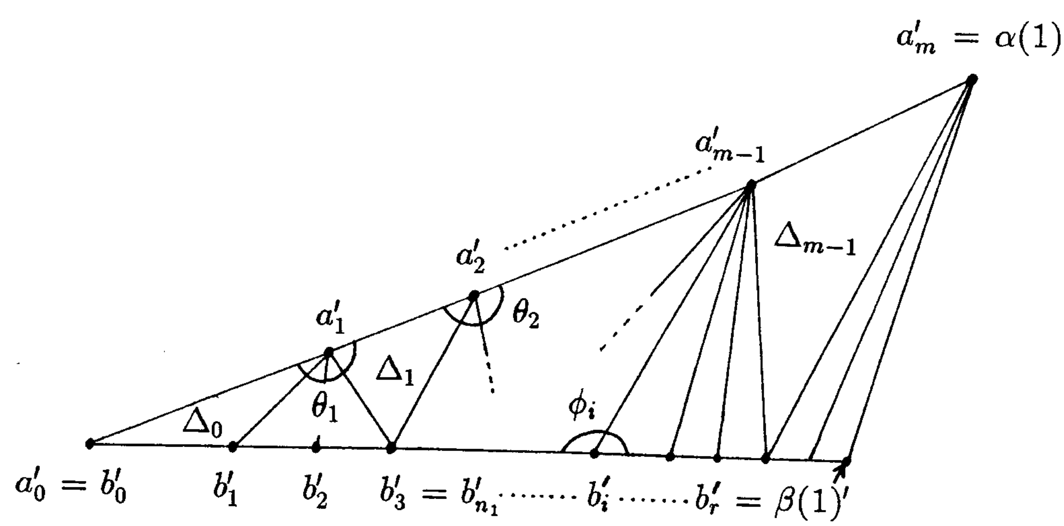

Figure 2.1: Superlinear divergence of geodesics

By barycentrically subdividing the simplices in $Shapes(K)$ and giving K the induced simplicial structure, we may assume that the restriction of the intrinsic metric on K to each simplex B agrees with the local metric d_B. Hence a line segment joining x' to y' in $\mathbf{E}^2$ lifts to a geodesic segment in K if x' and y' lie in the same triangle, while in general the line segment joining x' to y' in $\mathbf{E}^2$ lifts to a path from x to y in K of length $\geq d_K(x, y)$. Thus

$$\begin{aligned} &d_{E^2}\left(\alpha(1)', \beta(1)'\right) = d_K\left(\alpha(1), \beta(1)\right), \quad \text{and} \\ &d_{E^2}\left(\alpha(t)', \beta(t)'\right) \geq d_K\left(\alpha(t), \beta(t)\right) \;\; \forall t \in [0,1]. \end{aligned} \tag{1}$$

The angles θ_i and ϕ_i shown in Figure 2.1 are all $\geq \pi$, because α and β are geodesic segments. Therefore the derivative (which is defined almost everywhere) of the continuous piecewise linear function $\psi : t \mapsto d_{E^2}(\alpha(t)', \beta(t)')$ is non-decreasing. It is also non-negative, so ψ is convex and since $\psi(0) = 0$ we have the following inequality

$$d_{E^2}\left(\alpha(t)', \beta(t)'\right) \leq t \cdot d_{E^2}\left(\alpha(1)', \beta(1)'\right) \;\; \forall t \in [0,1]. \tag{2}$$

Combining (1) and (2) completes the proof. □

Examining the proof of Lemma 2.4 we see that the essential use of the fact that α and β are geodesic segments is to infer that $a_{i-1} \notin st(a_i)$ for $i \neq 1$, and that the angles θ_i and ϕ_i are all $\geq \pi$. To ensure that these conditions hold, one need only assume that α and β are local geodesics. Thus we obtain the following strengthening of Lemma 2.4.

Lemma 2.4*: *If K has unique geodesic segments and α is a local geodesic in K then there exists $\epsilon > 0$ such that any local geodesic β with $\|\alpha - \beta\| < \epsilon$ and $\beta(0) = \alpha(0)$ satisfies*

$$d\left(\alpha(t), \beta(t)\right) \leq t\, d\left(\alpha(1), \beta(1)\right) \qquad \forall t \in [0,1].$$

An important consequence of this generalisation is the following result, which is needed in the proof of Theorem 2.7.

Corollary : *If K has unique geodesic segments then every local geodesic in K is a geodesic segment. Hence there is a unique local geodesic joining any two points in K.*

Proof: Consider a local geodesic $\gamma : [0,1] \to K$ and let ϵ be as in Lemma 2.4. We will prove that $S = \{s : d(\gamma(0), \gamma(s)) = \text{length of } \gamma|_{[0,s]}\}$ is the whole of $[0,1]$. The set S is defined by a closed condition and contains 0, so it suffices to prove that it is open. Fix $s \in S$. Lemma 2.3 implies that for sufficiently small δ the geodesic segment from $\gamma(0)$ to $\gamma(s+\delta)$ is uniformly ϵ–close to the local geodesic $\gamma'(t) = \gamma(t(s+\delta))$. So by Lemma 2.4* these paths must coincide. □

We now turn to the proof of Theorem 2.1, and then conclude this section by giving two easy corollaries.

Theorem 2.1: *If K has unique geodesic segments, and α_0 and α_1 are geodesic segments in K with $\alpha_0(0) = \alpha_1(0)$ then*

$$d(\alpha_0(t), \alpha_1(t)) \leq t \cdot d(\alpha_0(1), \alpha_1(1)) \quad \forall t \in [0,1].$$

*Proof:*Let σ denote the unique geodesic segment from $\alpha_0(1)$ to $\alpha_1(1)$ and let α_s denote the geodesic segment from $\alpha_0(0)$ to $\sigma(s)$. We will prove that the set $\Sigma = \{s : d(\alpha_0(t), \alpha_s(t)) \leq td(\alpha_0(1), \alpha_s(1)) \forall t \in [0,1]\}$ is the whole of $[0,1]$.

Σ contains 0 and is closed by Lemma 2.3, because if for some t the inequality $d(\alpha_0(t), \alpha_s(t)) > td(\alpha_0(1), \alpha_s(1))$ holds at s then it holds in a neighbourhood of s. To see that Σ is open: Fix $s \in \Sigma$. Then by Lemma 2.4 there exists $\epsilon > 0$ such that if $\|\alpha_s - \alpha_{s+\delta}\| \leq \epsilon$ then

$$d(\alpha_s(t), \alpha_{s+\delta}(t)) \leq t \cdot d(\alpha_s(1), \alpha_{s+\delta}(1)) \quad \forall t \in [0,1].$$

Lemma 2.3 implies that $\|\alpha_s - \alpha_{s+\delta}\| \leq \epsilon$ for sufficiently small δ. For such δ we have the following inequality, which shows that $(s+\delta) \in \Sigma$.

$$\begin{aligned} d(\alpha_0(t), \alpha_{s+\delta}(t)) &\le d(\alpha_0(t), \alpha_s(t)) + d(\alpha_s(t), \alpha_{s+\delta}(t)) \\ &\le t \cdot d(\alpha_0(1), \alpha_s(1)) + t \cdot d(\alpha_s(1), \alpha_{s+\delta}(1)) \\ &= t \cdot d(\alpha_0(1), \alpha_{s+\delta}(1)) \qquad \forall t \in [0,1]. \end{aligned}$$

Hence Σ is open. □

Corollary : *If K has unique geodesic segments then for any geodesic segments α and β in K the function $f_{\alpha,\beta}(t) = d(\alpha(t), \beta(t))$ is convex.*

Proof: We show that for any geodesic segments α and β in K

$$d(\alpha(t), \beta(t)) \le t\, d(\alpha(1), \beta(1)) + (1-t)\, d(\alpha(0), \beta(0)) \quad \forall t \in [0,1].$$

The result then follows by applying this inequality to the initial segments of α and β. Let γ denote the geodesic segment from $\alpha(0)$ to $\beta(1)$. Applying Theorem 2.1 twice, once using α, γ and once using γ^*, β^* (where the star denotes reverse orientation), yields

$$\begin{aligned} d(\alpha(t), \beta(t)) &\le d(\alpha(t), \gamma(t)) + d(\gamma(t), \beta(t)) \\ &\le t\, d(\alpha(1), \gamma(1)) + (1-t)\, d(\gamma(0), \beta(0)) \\ &= t\, d(\alpha(1), \beta(1)) + (1-t)\, d(\alpha(0), \beta(0)) \quad \forall t \in [0,1], \end{aligned}$$

as required. □

Another immediate consequence of Theorem 2.1 is the following.

Corollary : *If K has unique geodesic segments and α_x denotes the geodesic segment from $x \in K$ to a fixed point $x_o \in K$, then $F : K \times I \to K$ given by $F(x,t) = \alpha_x(t)$ is a Lipschitz contraction of K to x_0.*

Proof: We have $d(F(x,t), F(y,t)) = d(\alpha_x(t), \alpha_y(t)) \le t d(x,y)$ for all $t \in [0,1]$, hence result with Lipschitz constant 1. □

2.2 Global Characterisations of Non-Positive Curvature, and the Fixed Point Theorem.

In this section we define three global criteria for non-positive curvature and show that if K is a piecewise Euclidean complex of type A or B then each of these criteria is equivalent to requiring that K has unique geodesic segments. One of these criteria is the CN-inequality of Bruhat and Tits, and this leads to the Fixed Point Theorem stated in the introduction. The other descriptions of non-positive curvature which we consider in this section are the $CAT(0)$ inequality, as defined by Gromov, and Alexandrov's condition on the excess of geodesic triangles. We begin by describing this last condition and proving that it is equivalent to the uniqueness of geodesic segments in K.

Definition: Given geodesic segments α and β in K with common initial point x, the *angle* between α and β at x is the distance between the points determined by α and β in the spherical complex $LK(x, K)$.

Definition: A *geodesic triangle* T in K consists of three points in K (the vertices of T) and a choice of geodesic segment between each pair of vertices (the edges of T).

We use the notation $T = \Delta(x_0, x_1, x_2)$, where x_0, x_1, x_2 are the vertices of T. But it should be noted that T is not uniquely determined by its vertices unless geodesic segments are unique in K.

Definition: Let $T = \Delta(x_0, x_1, x_2)$ be a geodesic triangle in K. A triangle $T' = \Delta(x'_0, x'_1, x'_2)$ in $\mathbf{E}^2$ is called a *comparison triangle* for T if $d_{E^2}(x'_i, x'_j) = d_K(x_i, x_j)$ for $i, j \in \{0, 1, 2\}$.

Alexandrov [2] defined curvature via the notion of *the excess of a triangle*, relating the sum of the *sup angles* of a geodesic triangle in the space under consideration to the sum of the angles in the comparison triangle. If K is a complex of type A or B and α and β are geodesic segments in K with a common initial point x then the sup angle between α and β at x is the lesser of π and the angle between α and β as defined above. Given a geodesic triangle Δ in K, let $|\Delta|$ denote the sum of the sup angles at the vertices.

Definition: The *excess* of a geodesic triangle Δ in K is $|\Delta| - \pi$.

Alexandrov defined non-positive curvature in geodesic metric spaces by the condition that every geodesic triangle in the given space has non-positive excess. In order to prove that K satisfies this condition if and only if it has unique geodesic segments we must first introduce some notation.

Suppose that K has unique geodesic segments and consider geodesic segments α and β in K with $\alpha(0) = \beta(0) = x, \alpha(1) = y$, and $\beta(1) = z$. Let Θ denote the angle between α and β at x, and suppose that $\Theta \leq \pi$. For sufficiently small t the geodesic segment from $\alpha(t)$ to $\beta(t)$ is contained in $st(x)$. We can form the join of this geodesic segment with x and map the resulting 2–complex isometrically into $\mathbf{E}^2$, as shown in Figure 2.2. Let x', y', z' be points in $\mathbf{E}^2$ with the property that $d_K(x, y) = d_{E^2}(x', y'), d_K(x, z) = d_{E^2}(x', z')$, and the angle between the line segments $[x', y']$ and $[x', z']$ (which we call α' and β' respectively) is Θ.

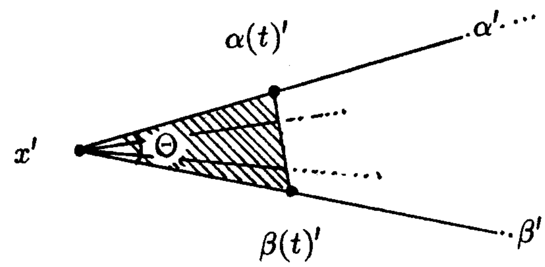

Figure 2.2: Locally divergence is linear

Lemma 2.5: *If y, z, y' and z' are as above then $d_K(y,z) \geq d_{E^2}(y',z')$.*

Proof: If ϵ is sufficiently small then $d_K(\alpha(\epsilon),\beta(\epsilon)) = d_{E^2}(\alpha'(\epsilon),\beta'(\epsilon))$, because Figure 2.2 is isometric to its preimage in K. By elementary Euclidean geometry

$$d_{E^2}(\alpha'(\epsilon),\beta'(\epsilon)) = \epsilon\, d_{E^2}(\alpha'(1),\beta'(1)) = \epsilon\, d_{E^2}(y',z').$$

And by Theorem 2.1

$$d_K(\alpha(\epsilon),\beta(\epsilon)) \leq \epsilon\, d_K(\alpha(1),\beta(1)) = \epsilon\, d_K(y,z).$$

Combining these inequalities finishes the proof. □

Remark: It is not difficult to show that Lemma 2.5 is equivalent to Theorem 2.1.

Gersten calls Lemma 2.5 the Topogonov inequality (cf. Theorem 3.1), and in [15] he derives the CN-inequality directly from this result. He was concerned only with the 2–dimensional case and used a local definition of geodesic, but with the benefit of the corollary to Lemma 2.4* his argument would work equally well in the present setting.

Proposition 2.6: *K has unique geodesic segments if and only every geodesic triangle in K has non-positive excess.*

Proof: Suppose that every geodesic triangle in K has non-positive excess, and fix $x, y \in K$. Consider a geodesic triangle formed by taking two geodesic segments α and β from x to y and adding a third vertex z at the midpoint of α. Since α is a geodesic segment, the angle at z is $\geq \pi$. Hence the angles at x and y are zero and consequently α and β coincide near their endpoints. The same argument shows that at no point can α and β diverge, and hence they coincide.

Conversely, suppose that K has unique geodesic segments and that T is a geodesic triangle in K. If one of the edges of T is the concatenation of the other two sides then the sup angles of T are $\pi, 0$ and 0. If not then, because every local geodesic in K is a geodesic segment, each vertex angle must be less than π. It then follows from Lemma 2.5 that each vertex angle is no greater than the corresponding angle in a comparison triangle for T, and hence the sum of these angles is no greater than π. □

Gromov [17] defines a geodesic metric space X to be non-positively curved if it satisfies $CAT(0)$, the comparison axiom of Alexandrov-Topogonov.

CAT(0): Let $T = \Delta(x_0, x_1, x_2)$ be a geodesic triangle in X, and let y a point on the side of T which has endpoints x_1 and x_2. Choose a comparison triangle $T' = \Delta(x_0', x_1', x_2')$ in $\mathbf{E}^2$ and let y' denote the unique point on the line segment $[x_1', x_2']$ such that $d_{E^2}(x_i', y') = d_K(x_i, y)$ for $i = 1, 2$. Then $d_{E^2}(x_0', y') \geq d_K(x_0, y)$.

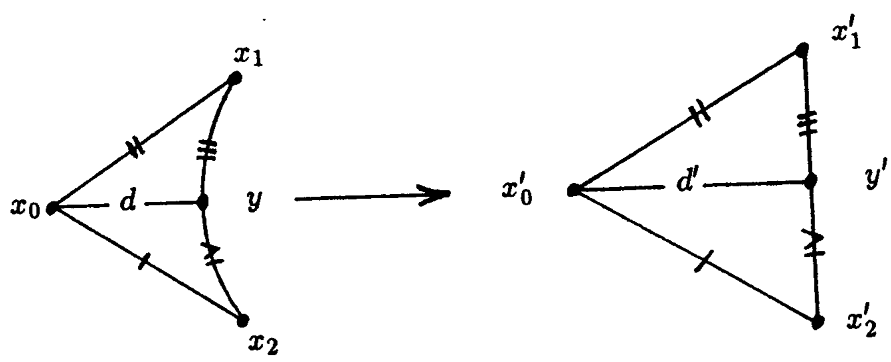

Figure 2.3: *CAT(0)*

CN: A metric space X satisfies the CN-condition of Bruhat-Tits [10] if for every pair of elements $x_1, x_2 \in X$ there exists a point m with $d(x_1, m) = d(x_2, m)$ such that for all $x_0 \in X$

$$d(x_1, x_0)^2 + d(x_2, x_0)^2 \geq 2d(m, x_0)^2 + \frac{1}{2}d(x_1, x_2)^2 . \qquad \text{(CN)}$$

The main result of this section is the following:

Theorem 2.7: *If K is a piecewise Euclidean complex of type A or B then the following global characterisations of non-positive curvature are equivalent:*

I) *K has unique geodesic segments.*

II) *K satisfies CAT(0).*

III) *K satisfies CN.*

IV) *Every geodesic triangle in K has non-positive excess.*

Proof: The equivalence of I and IV was proved in Proposition 2.6. We will now show II $\Rightarrow$III $\Rightarrow$I $\Rightarrow$II. Most of the work is in the final implication.

II $\Rightarrow$III: Suppose that K satisfies $CAT(0)$. Fix $T = \Delta(x_0, x_1, x_2)$ with comparison triangle $T' = \Delta(x_0', x_1', x_2')$. Let m denote the midpoint of the edge of T which has endpoints x_1 and x_2, and let m' denote its comparison point. It is easy to check that $\mathbf{E}^2$ satisfies the CN condition (in fact, by the parallelogram law, one gets equality in all cases) so the inequality (CN) holds for x_0', x_1', x_2', m'. Removing the primes the value of each term in (CN) stays the same with the exception of the first term on the right hand side, for which by $CAT(0)$ we have $d_{E^2}(x_0', m') \geq d_K(x_0, m)$.

III $\Rightarrow$I: Suppose that there are two geodesic segments from x_1 to x_2 in K. Shortening them if necessary, we may assume that their midpoints, which we denote by m and x_0, are distinct. Then $d(x_1, x_0)^2 + d(x_2, x_0)^2 = \frac{1}{2}d(x_1, x_2)^2$, so (CN) implies that $d(m, x_0) = 0$, contrary to hypothesis.

I $\Rightarrow$II: Let $T = \Delta(x_0, x_1, x_2) \subset K$, and let y be some point on the unique geodesic segment from x_1 to x_2 (for which we adopt the notation $[x_1, x_2]$). Let θ denote the angle between $[y, x_1]$ and $[y, x_0]$ at y, and let ϕ denote the angle between $[y, x_2]$ and $[y, x_0]$ at y.

First suppose that $\theta \geq \pi$, then it follows that the concatenation of $[x_1, y]$ and $[y, x_0]$ is a local geodesic from x_1 to x_0. As a corollary to Lemma 2.4* we showed that every local geodesic in K is a geodesic segment. Hence

$$d_K(x_0, y) = d_K(x_0, x_1) - d_K(x_1, y).$$

The triangle inequality applied to the comparison triangle $T' = \Delta(x_0', x_1', x_2')$ in $\mathbf{E}^2$ yields

$$\begin{aligned} d_{E^2}(x_0', y') &\geq d_{E^2}(x_0', x_1') - d_{E^2}(x_1', y') \\ &= d_K(x_0, x_1) - d_K(x_1, y). \end{aligned}$$

Hence $d_K(x_0, y) \leq d_{E^2}(x_0', y')$, so we are done if $\theta \geq \pi$, and similarly if $\phi \geq \pi$.

So we may assume that θ and ϕ are both less than π. This enables us to construct the planar 1–complex shown in Figure 2.4. Here $d_K(x_1, y) = d_{E^2}(x_1'', y''), d_K(x_2, y) = d_{E^2}(x_2'', y''), d_K(x_0, y) = d_{E^2}(x_0'', y'')$ and the angles θ' and ϕ' are equal to θ and ϕ respectively. Notice that since $[x, y]$ is a geodesic segment $\phi' + \theta' \geq \pi$.

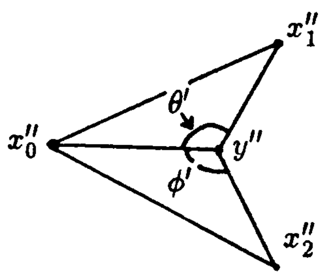

Figure 2.4: A comparison figure preserving angles

We now transform Figure 2.4 into a comparison triangle for T by motions which do not decrease $d_{E^2}(x_0'', y'') = d_K(x_0, y)$, thus completing the proof of the theorem.

By Lemma 2.5, $d_K(x_0, x_1) \geq d_{E^2}(x_0'', x_1'')$ and $d_K(x_0, x_2) \geq d_{E^2}(x_0'', x_2'')$. Also, since $\theta < \pi$ and geodesic segments are unique in K, we have $d_K(x_0, y) + d_K(y, x_1) > d_K(x_0, x_1)$. So without changing the length of the edges incident at y'' we can increase the angles θ' and ϕ' until the distances $d_{E^2}(x_0'', x_1'')$ and $d_{E^2}(x_0'', x_2'')$ reach $d_K(x_0, x_1)$ and $d_K(x_0, x_2)$ respectively. This brings us to the situation shown in Figure 2.5. We now increase the angle ψ until it is equal to π. While doing so we keep all edge lengths constant, except that $d_{E^2}(x_0'', y'')$ is allowed to increase. This completes the construction of the comparison triangle, and with it the proof of Theorem 2.7. □

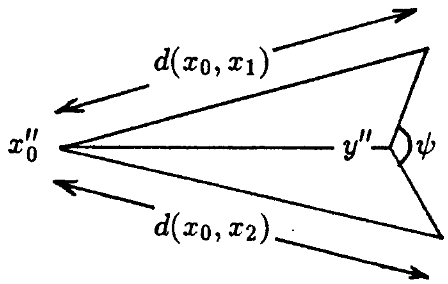

Figure 2.5: Constructing a comparison triangle

Bruhat and Tits showed that if a metric space X is complete and satisfies the CN condition then any group G which acts on X by isometries and has a non-empty bounded orbit stabilises some point of X. An elegant proof of this, due to Serre, can be found in [8] pp.157–158. Moreover, if X is a *geodesic* metric space then there is a unique geodesic segment connecting any two points which are fixed by G. Since G acts by isometries it must fix this geodesic segment pointwise. It follows that the fixed point set of G is contractible. So we have the following consequence of Theorem 2.7.

Fixed Point Theorem: *If K satisfies any of the conditions* I *to* IV *given in Theorem 2.7, and a group Γ acts on K by isometries, such that there is a bounded orbit, then the fixed point set of Γ is non-empty and contractible.*

2.3 Local Characterisations of Non-Positive Curvature.

One would like to define what it means for a Euclidean simplicial complex to have non-positive curvature locally. The obvious way to do this is to select one of the global characterisations of non-positive curvature given in Sections 2.1

and 2.2 and require that every point of the given complex have a neighbourhood in which this condition holds. Theorem 2.8 (which is stated below) shows that this definition does not depend on which global charactersiation we choose.

Theorem 2.8 also relates local curvature to the structure of links in the given complex by means of the link condition. It is this relationship which enables us to determine whether or not certain simplicial complexes, whose local (combinatorial) structure is sufficiently well understood, support a structure of non-positive curvature (see Proposition 4.4 and Theorem 5.6).

We maintain the convention that the letter K, without further qualification, denotes a piecewise Euclidean complex of type A or B.

Definition: K satisfies the *link condition* if for every $x \in K$ two points y', z' in the spherical simplicial complex $LK(x, K)$ can be joined by a *unique* geodesic segment in $LK(x, K)$ whenever $d(y', z') < \pi$.

Definition: An open set U in K is said to be *geodesically convex* if for all $x, y \in U$ *every* geodesic segment joining x to y in K is contained in U.

Definition: K is said to satisfy a property *locally* if K can be covered by geodesically convex open sets each of which satisfies the given property.

Theorem 2.8: *For piecewise Euclidean complexes K of type A or B the following local characterisations of non-positive curvature are equivalent:*

I) *K has unique geodesic segments locally.*

II) *The metric on K is convex locally.*

III) *K satisfies $CAT(0)$ locally.*

IV) *K satisfies CN locally.*

V) *Every point of K has a neighbourhood such that any geodesic triangle contained in that neighbourhood has non-positive excess.*

VI) *K satisfies the link condition.*

Further, if K is of type B then each of the above conditions is equivalent to

VII) *There exists $\epsilon_0 > 0$ such that for all $x \in K$ the ball $B_{\epsilon_0}(x)$ is geodesically convex and has unique geodesic segments.*

If K is of type A, then condition VII is strictly stronger than the other conditions, as the following example shows.

Example: Let L be the 1–complex which has vertices v_n, and which for every integer n has three 1–cells $\{v_{2n-1}, v_{2n}\}, \{v_{2n-1}, v_{2n+1}\}, \{v_{2n}, v_{2n+1}\}$, each of length $1/n$. This is a Euclidean simplicial complex of type A which satisfies conditions I to VI of Theorem 2.8, but does not satisfy condition VII.

Proof of Theorem 2.8: We begin by showing that conditions I to IV are equivalent. Let U be a geodesically convex open subset of K, and suppose that geodesic segments are unique in U. The Remark following the proof of Lemma 2.3 implies that the conclusions of Lemmas 2.3 and 2.4 are valid in U. Theorem 2.1 follows from these results in an entirely formal way, hence the metric on U is convex. Then, a formal translation of the arguments given in Section 2.2 proves that U satisfies $CAT(0), CN$ and Alexandrov's condition on the excess of a geodesic triangle. The argument given in the proof of Proposition 2.6 shows that this last condition implies that geodesic segments are unique in U. Hence conditions I to V are equivalent.

It is clear that VII⇒I, so it suffices to show that I and VI are equivalent and that if K is of type B then VI⇒VII. This follows from the following three results.

First we show that *any* **E**-simplicial complex of type A or B can be covered with geodesically convex open sets. Here, as in Lemma 1.2, $\epsilon(x)$ denotes the distance from x to its link in K.

Lemma 2.9: *Given $x \in K$, the open ball $B_\epsilon(x)$ is geodesically convex whenever $\epsilon < \frac{1}{2}\epsilon(x)$.*

Proof: Fix $y, z \in B_\epsilon(x)$ and notice that if α is a geodesic segment joining y to z in K then its image is contained in $B_{\epsilon(x)} \subseteq st(x)$. If we form the join of α with x and map the resulting 2–complex isometrically into $\mathbf{E}^2$ then we obtain the situation illustrated in Figure 2.6. It follows that

$$d(\alpha(t), x) \leq \max\{d(x,y), d(x,z)\} < \epsilon$$

as required. □

Lemma 2.10: *If K is of type B then there is a constant $\epsilon_0 > 0$ such that for every $x \in K$ there exists y with $B_{\epsilon_0}(x) \subseteq B_{\frac{1}{2}\epsilon(y)}(y)$.*

Proof: The idea of the proof is to construct a constant ϵ_0 such that if $x \in K^{(n)}$ lies in the $2\epsilon_0$–neighbourhood of the $(n-1)$–skeleton of K then there exists $y \in K^{(n-1)}$ such that $B_{\epsilon_0}(x) \subseteq B_{\frac{1}{2}\epsilon(y)}(y)$.

Let Σ denote the disjoint union of $\{\sigma : \sigma \in Shapes(K)\}$, and define $\eta : \Sigma \to (0, \infty)$ by

$$\eta(s) = \min\{d(s,F) : \sigma \in Shapes(K), F \text{ a face of } \sigma,\ s \in (\sigma - F)\}.$$

Note that η is continuous on the interior of simplices.

Let η_0 be $\frac{1}{4}$ of the minimum value attained by η on $\Sigma^{(0)}$, the 0–skeleton of Σ. Then inductively for $n \leq D = dim(K)$ we define η_n to be $\frac{1}{4}$ of the minimum value attained by η on the compact set obtained by deleting the η_{n-1} neighbourhood of $\Sigma^{(n-1)}$ from $\Sigma^{(n)}$. Notice that $4\eta_n \leq \eta_{n-1}$. Set $\epsilon_0 = \eta_D$.

Arguing by induction on n it is easy to see that if $x \in K^{(n)} - K^{(n-1)}$ then either $2\eta_n < \epsilon(x)$ or else there exists $y \in K^{(n-1)}$ with $B_{2\eta_n}(x) \subseteq B_{\frac{1}{2}\epsilon(y)}(y)$. □

Proposition 2.11: *K satisfies the link condition if and only if for every $x \in K$ geodesic segments are unique in $B_\epsilon(x)$ whenever $\epsilon < \epsilon(x)/2$.*

Proof: Fix $x \in K$ and let $\epsilon < \epsilon(x)/2$. The exponential map allows us to identify $LK(x, K)$ with the boundary of the closed ball $B_\epsilon(x)$, by a map which scales the metric by a factor of ϵ. Given $y, z \in B_\epsilon(x)$ let y', z' denote the points in $LK(x, K)$ determined by the line segments $[x, y], [x, z]$. Lemma 1.2 shows that if $y' = z'$ then $[y, z]$ is the unique geodesic segment joining y to z in K. So we may assume that $y' \neq z'$.

Suppose that there exists a geodesic segment α joining y to z in $(B_\epsilon(x) - x)$. In Lemma 1.10 we showed that α', the image of α under radial projection, is represented by a taut chain in $LK(x, K)$, and $L(\alpha') < \pi$ (see Figure 1.3). On the other hand, any piecewise geodesic path α' of length $< \pi$ which joins y' to z' in $LK(x, K)$ is the image of a path α from y to z in $(B_\epsilon(x) - x)$, as illustrated in Figure 2.6.

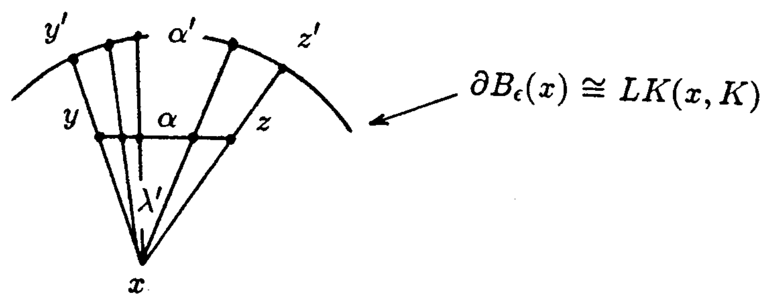

Figure 2.6: Radial projection of a geodesic segment

Let $\lambda = L(\alpha)$ and $\lambda' = L(\alpha') < \pi$. By the cosine rule we have

$$\lambda^2 = d(x,y)^2 + d(x,z)^2 - 2d(x,y)\,d(x,z)\cos\lambda'.$$

Hence minimising λ is equivalent to minimising λ'. So if $d(y',z') < \pi$ then there is a unique geodesic segment joining y' to z' in $LK(x,K)$ if and only if there is a unique geodesic segment joining y to z in $(B_\epsilon(x) - x)$.

To complete the proof we must show that if $d(y',z') \geq \pi$ then there is a unique geodesic segment from y to z in $B_\epsilon(x)$. But this is clear, because by the arguments given above there is no local geodesic from y to z in $(B_\epsilon(x) - x)$, and the line segment $[x,w]$ is the unique geodesic segment from x to w for all $w \in B_\epsilon(x)$. Hence the concatenation of the line segments $[y,x]$ and $[x,z]$ is the unique geodesic segment from y to z in $B_\epsilon(x)$. □

We now return to the proof of Theorem 2.8. If condition I holds, then Lemma 2.9 implies that for every $x \in K$ there exists $\delta(x) > 0$ such that $B_\epsilon(x)$ is geodesically convex and has unique geodesic segments whenever $\epsilon < \delta(x)$. It then follows from the "if" implication of Proposition 2.11 that K satisfies the link condition. Conversely, the "only if" implication of Proposition 2.11 shows that VI⇒I.

In fact, these arguments shows that if condition I holds then geodesic segments are unique in $B_{\epsilon(y)/2}(y)$ for every $y \in K$. Moreover, the uniqueness of geodesic segments implies that the restriction of the metric to each such ball is convex. It follows that if $B_{\epsilon_0}(x) \subseteq B_{\epsilon(y)/2}(y)$ then $B_{\epsilon_0}(x)$ is geodesically convex and has unique geodesic segments. The equivalence of conditions I and VII for complexes of type B now follows from Lemma 2.10. □

2.4 Passing From the Local to the Global Situation.

We now turn to the the most difficult step in the proof of the Main Theorem. Namely, that of relating local characterisations of non-positive curvature to global characterisations.

Theorem 2.12: *If K is a piecewise Euclidean complex of type A or B which has unique geodesic segments locally then for all $x, y \in K$ there is a unique shortest path in each homotopy class of paths from x to y in K.*

Corollary : *If K is simply connected and satisfies any of the local characterisations of non-positive curvature given in Theorem 2.8 then it satisfies all of the global characterisations of non-positive curvature given in Theorems 2.1 and 2.7.*

Thus Theorem 2.12 completes the proof of the special case of the Main Theorem which we stated at the beginning of Chapter 2, and shows that for simply connected simplicial complexes the existence of a piecewise Euclidean metric of non-positive curvature is, as one would expect, a local condition.

Throughout Section 2.4 we assume that K satisfies the hypotheses of Theorem 2.12, and that $p, q \in K$ are fixed.

Let Ω denote the set of PL paths $\alpha : [0,1] \rightarrow K$ such that $\alpha(0) = p$ and $\alpha(1) = q$, equipped with the metric topology given by $\|\alpha - \beta\| = sup\{d(\alpha(t), \beta(t)) : t \in [0,1]\}$. We continue to denote the length of α by $L(\alpha)$, and denote the length of $\alpha|_{[0,t]}$ by $L(\alpha, t)$. Notice that $L(\alpha, t)$ is a piecewise linear function of t and hence is differentiable almost everywhere.

Definition: If $\alpha \in \Omega$ then the *energy* of α is $E(\alpha) = \int_o^1 (\frac{dL}{dt}(\alpha, t))^2 dt$.

The Cauchy-Schwarz inequality implies that $E(\alpha) \geq L(\alpha)^2$, with equality if and only if α is parameterised proportional to arc length.

Although E is not continuous on Ω, it is continuous on the the following subspace of "broken geodesics".

$$\Omega(n) = \left\{\alpha \in \Omega : t \mapsto \alpha\left(\frac{i+t}{n}\right) \text{ is a geodesic segment for } i = 0, \dots, n-1\right\}$$

On this subspace E is given by the formula

$$E(\alpha) = \sum_{i=0}^{n-1} \int_{t_i}^{t_{i+1}} \left(\frac{d}{dt} L(\alpha, t)\right)^2 dt = \sum_{i=0}^{n-1} \frac{d(\alpha(t_{i+1}), \alpha(t_i))^2}{t_{i+1} - t_i}$$
$$= n \sum_{i=0}^{n-1} d(\alpha(t_{i+1}), \alpha(t_i))^2$$

where $t_i = i/n$. To prove Theorem 2.12 we analyse the convexity properties of E on the subspace

$$\Omega(n, e) = \{\alpha \in \Omega(n) : E(\alpha) \leq e\}.$$

Proof of Theorem 2.12

The strategy of our argument is as follows (cf. [22] Theorem 19.2, and [27] Theorem 1). Fix a homotopy class of paths from p to q in K and let α_0 and α_1 be two shortest paths in this class. In Lemma 2.13 we show that there exist constants n and e such that α_0 and α_1 lie in the same path component of $\Omega(n, e)$. Moreover, we can choose e so that every $\alpha \in \Omega(n, e)$ is uniquely determined by the sequence of points $(\alpha(t_i))_{i=1}^{n-1}$. This allows us to embed $\Omega(n, e)$ in K^{n-1}. We then extend E to a continuous function on the whole of K^{n-1} and prove that it has strong convexity properties (Proposition 2.17). This leads to the following result (see Lemmas 2.19 and 2.20):

The set of local minima of $E|_{\Omega(n,e)}$ is discrete, and there is a strong deformation retraction of $\Omega(n, e)$ onto this set.

This will complete the proof of Theorem 2.12, because the fact that α_0 is a shortest path in its homotopy class implies that E restricted to the corresponding component of $\Omega(n,e)$ attains its minimum value at α_0. The same is true of α_1. But E has a unique local minimum in each component of $\Omega(n,e)$, hence α_0 and α_1 must coincide.

Lemma 2.13: *Suppose that K is of type B, and let ϵ_1 be such that for every $x \in K$ the ball $B_{4\epsilon_1}(x)$ is geodesically convex and has unique geodesic segments. If α_0 and α_1 are two shortest paths in the same path component of Ω then there exists a positive integer n and a constant $e > 0$ such that α_0 and α_1 lie in the same path component of $\Omega(n,e)$. Moreover, if we let $t_i = i/n$ then $d(\alpha(t_i), \alpha(t_{i+1})) < \epsilon_1$ for every $\alpha \in \Omega(n,e)$ and $i \in \{1,\ldots,n-1\}$.*

Proof: Let $D : I \times I \to K$ be a homotopy between the maps α_0 and α_1, and denote $D(s,\tau)$ by $\alpha_s(\tau)$. Because D is uniformly continuous, we can find an integer m such that if $\tau_i = i/m$ then $d(\alpha_s(\tau_i), \alpha_s(\tau_{i+1})) < \epsilon_1$ for all $s \in I$ and $i \in \{0,\ldots,m-1\}$. We can then replace each α_s by the path α'_s which is defined to be the concatenation of the *unique* geodesic segments joining $\alpha_s(\tau_i)$ to $\alpha_s(\tau_{i+1})$ in K.

More precisely, we reparameterise the geodesic segment joining $\alpha_s(\tau_i)$ to $\alpha_s(\tau_{i+1})$ to obtain a path $\sigma_s^i : [\tau_i, \tau_{i+1}] \to K$ which is parameterised proportional to arc length, and define $\alpha'_s : [0,1] \to K$ by $\alpha'_s|_{[\tau_i,\tau_{i+1}]} = \sigma_s^i$. Notice that $\alpha_0 = \alpha'_0$ and $\alpha_1 = \alpha'_1$ because α_0 and α_1 are shortest paths in their homotopy class and $B_{\epsilon_1}(x)$ is simply connected (indeed contractible) for all x.

Geodesics vary continuously with their endpoints in $B_{\epsilon_1}(x)$ for every $x \in K$, and hence $s \mapsto \alpha'_s$ is a continuous path from α_0 to α_1 in $\Omega(m)$. The energy function E is continuous along this path and hence attains a maximum value, which we denote by e.

For every α'_s we have the following inequality

$$E(\alpha'_s) = \sum_{i=0}^{m-1} \int_{\tau_i}^{\tau_{i+1}} \left(\frac{d}{dt}L(\alpha'_s, \tau)\right)^2 d\tau = m \sum_{i=0}^{m-1} d(\alpha'_s(\tau_{i+1}), \alpha'_s(\tau_i))^2 \leq m^2 \epsilon_1^2,$$

which implies that $e \leq m^2 \epsilon_1^2$.

Thus if we let $n = m^2$ and $t_i = i/n$ then for every $\alpha \in \Omega(n, e)$ and $i \in \{0, \ldots, m-1\}$ we have

$$\begin{aligned} d(\alpha(t_i), \alpha(t_{i+1}))^2 &= \left(\int_{t_i}^{t_{i+1}} \frac{d}{dt}L(\alpha, t)\, dt\right)^2 \\ &\leq (t_{i+1} - t_i) \int_0^1 \left(\frac{d}{dt}L(\alpha, t)\right)^2 dt \\ &\leq \frac{1}{n} e \leq \epsilon_1^2. \end{aligned}$$

Where the second line comes from the Cauchy-Schwarz inequality in $L^2[0,1]$ applied to $\frac{d}{dt}L(\alpha, t)$ and the characteristic function of $[t_i, t_{i+1}]$. □

Before proceeding we must dispense with an irksome technicality. Namely, if K is of type A then in general the fact that K has unique geodesic segments locally does not imply that there exists a constant ϵ_0 such that geodesic segments are unique in $B_{\epsilon_0}(x)$ for every $x \in K$. However, one can apply the method of Proposition 2.11 to show that for every integer R there exists $\epsilon(R) > 0$ such that if $d(x, p) < R$ then $B_{\epsilon(R)}(x)$ is geodesically convex and has unique geodesic segments. This allows us to modify the proof of Lemma 2.13 to obtain an analogous result if K is of type A, as we now indicate.

For the purposes of this discussion we retain the notation which we introduced in the proof of Lemma 2.13. Fix an integer R so that the image of the homotopy disc D is contained in the ball of radius R about p. We can choose m so that $d(\alpha_s(\tau_i), \alpha_s(\tau_{i+1})) < \epsilon(R)$ for all $s \in [0,1]$ and $\tau = i/m$ where $i \in \{0, \ldots, m-1\}$. As in the proof of Lemma 2.13, we replace each α_s by α'_s, the concatenation of the geodesic segments joining $\alpha_s(\tau_i)$ to $\alpha_s(\tau_{i+1})$. This gives a continuous path from α_0 to α_1 in $\Omega(m)$. Let e denote the maximum value which E attains on this path.

If $\alpha \in \Omega$ and $E(\alpha) < e$ then the image of α is contained in the ball of radius $\sqrt{e}$ about p. Fix an integer R' so that $R' \gg \sqrt{e}$. The final inequality in Lemma 2.13 shows that if we take n sufficiently large then $d(\alpha(t_i), \alpha(t_{i+1})) < \epsilon(R')$ for every $\alpha \in \Omega(n,e)$ and $i \in \{0, \ldots, n-1\}$, where $t_i = i/n$.

Thus, if we write ϵ_1 in place of $\epsilon(R')/4$ then the essential conclusion of Lemma 2.13 remains valid if K is of type A. Namely, a path $\alpha \in \Omega(n,e)$ is determined by the sequence of points $(\alpha(t_i))$, and for every $x \in K$ which we shall need to consider (i.e. those which lies in a small neighbourhood of the set $\{\alpha(t) : \alpha \in \Omega(n,e)\}$) the ball $B_{\epsilon_1}(x)$ is geodesically convex and has unique geodesic segments.

For the remainder of this section we assume that n *and* e *are as in Lemma 2.13 or as in the preceding discussion, according to whether* K *is of type B or of type A. We also retain the notation* t_i *for* i/n.

Lemma 2.14: *The map* $\Phi : \Omega(n,e) \to K^{n-1}$ *given by* $\alpha \mapsto \langle \alpha(t_1), \ldots, \alpha(t_{n-1}) \rangle$ *is an injection.*

Proof: Elements of $\Omega(n,e)$ are geodesic on the subintervals $[t_i, t_{i+1}]$, and we chose ϵ_1 so that if $d(\alpha(t_i), \alpha(t_{i+1})) < \epsilon_1$ then there is a unique geodesic segment from $\alpha(t_i)$ to $\alpha(t_{i+1})$ in K. $\square$

To complete the proof of Theorem 2.12 we study the image of $\Omega(n,e)$ under the map Φ, which we defined in Lemma 2.14. We begin by describing the metric structure of K^{n-1}, whose elements we denote by $X = \langle x_1, \ldots, x_m \rangle$. Each closed cell $\Sigma \subset K^{n-1}$ is the $(n-1)$–fold Cartesian product of simplices in K. The partial metric on $\Sigma = \sigma_1 \times \ldots \times \sigma_{n-1}$ is given by $d(X,Y)^2 = \sum_{i=1}^{n-1} d(x_i, y_i)^2$; and K^{n-1} is a piecewise Euclidean space of type A or B according to the type of K.

Lemma 2.15: *$\alpha(t) = \langle \alpha_1(t), \ldots, \alpha_{n-1}(t) \rangle$ is a geodesic segment in K^{n-1} if and only if α_i is a geodesic segment in K for every i.*

Proof: According to the Cauchy-Schwarz inequality, a path α is a geodesic segment if and only if it has minimal energy among all paths which have the same endpoints as α. Therefore it is enough to consider paths which are parameterised proportional arc length, and show that $E(\alpha) = \Sigma_{i=1}^{n-1} E(\alpha_i)$.

Suppose that α is represented by the m–chain $(\alpha(\tau_0), \ldots, \alpha(\tau_m))$. Then α_i is represented by the m–chain $(\alpha_i(\tau_0), \ldots, \alpha_i(\tau_m))$, and is parameterised proportional to arc length on each subinterval $[\tau_i, \tau_{i+1}]$. Hence

$$E(\alpha) = \sum_{i=0}^{m-1} \int_{\tau_i}^{\tau_{i+1}} \left(\frac{d}{dt} L(\alpha, t) \right)^2 dt = \sum_{i=0}^{m-1} \frac{d(\alpha(\tau_{i+1}), \alpha(\tau_i))^2}{\tau_{i+1} - \tau_i}$$

$$\sum_{j=1}^{n-1} E(\alpha_j) = \sum_{j=1}^{n-1} \int_0^1 \left(\frac{d}{dt} L(\alpha_j, t) \right)^2 dt = \sum_{j=1}^{n-1} \sum_{i=0}^{m-1} \frac{d(\alpha_j(\tau_{i+1}), \alpha_j(\tau_i))^2}{\tau_{i+1} - \tau_i}$$

The last terms in each row are equal because of the definition of the metric on the individual cells of K^{n-1}. □

We have proved that the map $\Phi : \alpha \mapsto \langle \alpha(t_1), \ldots, \alpha(t_m) \rangle$ is injective, and it is easy to see that it is continuous ($d(\Phi(\alpha), \Phi(\beta)) \leq (n-1)^{\frac{1}{2}} \|\alpha - \beta\|$). In fact Φ is a homeomorphism of $\Omega(n,e)$ onto its image, because in the region of K

under consideration geodesic segments vary continuously with their endpoints in balls of radius $4\epsilon_1$, and $L(\alpha|_{[t_i,t_{i+1}]}) < \epsilon_1$ for all $\alpha \in \Omega(n,e)$.

We denote the image of Φ in K^{n-1} by P_e. We extend the map $(E \circ \Phi^{-1}) : P_e \to [0,\infty)$ to a continuous function (which we shall also call E) defined on the whole of K^{n-1}, by the formula

$$E(X) = n\left(d(p,x_1)^2 + \sum_{i=1}^{n-2} d(x_i,x_{i+1})^2 + d(x_{n-1},q)^2\right).$$

We shall prove (Proposition 2.17) that this extended function has a strong convexity property which enables us to strong deformation retract P_e (and hence $\Omega(n,e)$) onto a discrete set of local minima of E, thus completing the proof of Theorem 2.12.

We shall need the following consequence of Theorem 2.8.

Lemma 2.16: *Let $x \in K$, and suppose that $B_\delta(x)$ is geodesically convex and has unique geodesic segments. If $\alpha : [0,1] \to B_\delta(x)$ is a geodesic segment in $B_\delta(x)$ such that $\alpha(0) \neq \alpha(1)$ and $D = d(x,\alpha(0)) = d(x,\alpha(1)) > 0$ then $d(x,\alpha(t)) < D$ for all $t \in (0,1)$.*

Proof: By Theorem 2.8 $CAT(0)$ holds in $B_\delta(x)$, so it is enough to consider the case $K = \mathbf{E}^2$, where the result is clear. □

The proof of the following result is essentially due to Stone [27].

Proposition 2.17: *If $X,Y \in K^{n-1}$ are distinct points in the ϵ_1–neighbourhood of P_e and $d(X,Y) < \epsilon_1$ then there is a unique geodesic segment from X to Y in K^{n-1}, and E is strictly convex along this geodesic segment.*

Proof: Notice that $d(X,Y) < \epsilon_1$ implies that $d(x_i,y_i) < \epsilon_1$ for all i. Hence there exists a unique geodesic segment from x_i to y_i in K, and we denote this by β_i. The uniqueness of the β_i, together with Lemma 2.15, implies that

$\beta(t) = \langle\beta_1(t), \ldots, \beta_{n-1}(t)\rangle$ is the unique geodesic segment from X to Y in K^{n-1}.

It remains to show that E is strictly convex along β. We are assuming that there exists $Z \in P_e$ such that $d(Z, X) < \epsilon_1$. Hence

$$d(x_i, x_{i+1}) \leq d(x_i, z_i) + d(z_i, z_{i+1}) + d(z_{i+1}, x_{i+1}) < 3\epsilon_1.$$

It follows that β_i and β_{i+1} are both contained in $B_{4\epsilon_1}(x_i)$, which is geodesically convex and has unique geodesic segments. Further, by the convexity of the metric on $B_{4\epsilon_1}(x_i)$ we have

$$d(p, \beta_1(t)) \leq (1-t)\ d(p, x_1) + t\, d(p, y_1)$$

$$d(\beta_{n-1}(t), q) \leq (1-t)\ d(x_{n-1}, q) + t\, d(y_{n-1}, q)$$

$$d(\beta_i(t), \beta_{i+1}(t)) \leq (1-t)\ d(x_i, x_{i+1}) + t\, d(y_i, y_{i+1}).$$

By Lemma 2.15 the first of these inequalities is strict for $t \in (0,1)$ unless $d(p, x_1) \neq d(p, y_1)$ or $x_1 = y_1$. And if $x_1 = y_1$ then by the same argument the bottom inequality is strict for $i = 1$ and $t \in (0,1)$ unless $d(x_1, x_2) \neq d(y_1, y_2)$ or $x_2 = y_2$. The points $X, Y \in K^{n-1}$ are distinct, so proceeding in this way we see that one of the inequalities given above is strict for all $t \in (0,1)$, or else there is some i for which $d(x_i, x_{i+1}) \neq d(y_i, y_{i+1})$. (Here we have adopted the convention that $z_0 = p$ and $z_n = q$ for all $Z \in K^{n-1}$.)

Using the convexity of the function $f(a) = a^2$ one can show that for any non-negative numbers a, b, c if $a \leq (1-t)b + tc$ then $a^2 \leq (1-t)b^2 + tc^2$, with equality only if the original inequality was actually an equality, and either $t \in \{0,1\}$, or $b = c = a$. Applying this observation to the inequalities given above, and adding the resulting inequalities we obtain

$$E(\beta(t)) \leq (1-t)\,E(X) + t\,E(Y)$$

with strict inequality for all $t \in (0,1)$. $\square$

By subdividing the simplices of K we may assume that every closed cell in K^{n-1} which meets P_e has diameter bounded above by ϵ_1. Let L denote the union of these closed cells, and notice that if two points in L are $4\epsilon_1-$close then they can be joined by a unique geodesic segment in L. It follows from Lemma 2.4* that if these points are elements of the same closed cell in L then the line segment joining them in L (which is *a priori* only a local geodesic) is a geodesic segement. In particular, by Proposition 2.17, E is strictly convex along this line segment. Hence $E|_\Sigma$ attains its minimum value at a unique point $X_\Sigma \in \Sigma$ for every closed cell Σ in L. Inductively, in order of increasing dimension, we star each cell with this point to obtain a (simplicial) subdivision of L.

More precisely, if Σ is a 1–cell and $\partial\Sigma = \{X_0, X_1\}$ then we introduce new edges $\{X_0, X_\Sigma\}$ and $\{X_\Sigma, X_1\}$ into the underlying cell structure of L. Then inductively we assume that we have constructed the desired simplicial subdivision on the $(i-1)-$skeleton of L and consider an $i-$cell Σ. For every $j-$dimensional face τ^j of $\partial\Sigma$ that does not contain X_Σ we introduce the union of the line segments $\{[X_\Sigma, Y] : Y \in \tau^j\}$ as a new $(j+1)-$simplex in the cell structure on L. Notice that we have not changed the metric structure on L, only the underlying combinatorial structure.

Let J denote the union of the closed cells in this subdivision which meet P_e. Notice that if C is a closed cell in J then E is strictly convex along all of the line segments in C and $E|_C$ attains its unique minimum at a vertex.

Proposition 2.18: *The set* $\{E(v) : v$ *is a vertex of* $J\}$ *is finite.*

Proof: If K is of type A then the result is trivial because J is compact. So we may assume that K is of type B. The vertices of J are of two types, those which were vertices of L, and those which we introduced in the construction of J. First we show that $\{E(X) : X \text{ is a vertex of } L\}$ is a finite set.

Every vertex of K^{n-1} is of the form $X = \langle v_1, \ldots, v_{n-1} \rangle$ where each of the coordinates v_i is a vertex of K. If $X \in L$ then (as we showed in the proof of Proposition 2.17) $d(v_i, v_{i+1}) < 3\epsilon_1$ for all i. By Theorem 1.11, there exists an integer N such that any geodesic segment in K of length less than $3\epsilon_1$ can be modelled in one of the finitely many corridors which occur in K and have length less than N. Let Λ denote the set of such corridors. Then for $i \in \{1, \ldots, n-2\}$ we have $d(v_i, v_{i+1}) \in \{d(u,v) : u, v \text{ vertices of some } \mu \in \Lambda\}$. This last set is finite, so $d(v_i, v_{i+1})$ takes on only finitely many values as X varies over the vertices of L. A similar argument shows that $d(p, v_i)$ and $d(q, v_{i+1})$ can only take on finitely many values, and since $E(X)$ depends only on these quantities the set $\{E(X) : X \text{ is a vertex of } L\}$ is finite.

It remains to show that the set of values which E can take at vertices which were introduced in the construction of J is finite. This set can be written as $\{inf E|_\Sigma : \Sigma \text{ a closed cell in } L \text{ and } \Sigma \cap P_e \neq \emptyset\}$, and we will use this description to show that the set is finite. The idea of the proof is to show that $inf E|_\Sigma$ depends only on the metric structure of a set of complexes which we obtain by concatenating $n-$tuples of corridors in Λ. Because Λ is a finite set there are only finitely many possible isometric models for such complexes.

Let $\Sigma = (B_1 \times \ldots \times B_{n-1})$ be a closed cell in L which meets P_e, and fix $X \in \Sigma$. Because $d(x_i, x_{i+1}) < 3\epsilon_1$ and geodesic segments are unique in $B_{4\epsilon_1}(x)$ for all $x \in K$, there is a unique geodesic segment from x_i to x_{i+1} in K for every $i \in \{0, \ldots, n-1\}$ (recall the convention that $x_0 = p, x_n = q$). We parameterise this geodesic segment proportional to arc length to obtain $\gamma_X^i : [t_i, t_{i+1}] \to K$, and denote the concatenation of the γ_X^i by $\gamma_X : [0,1] \to K$. Notice that $\gamma_X \in \Omega(n)$ and $E(X) = E(\gamma_X)$ (which is greater than e if $X \notin P_e$).

Let Γ_i denote the model corridor $\Gamma(\gamma_i^X)$ (here we are using the notation which we introduced with the definition of a corridor at the beginning of Section 2.1), and let $\tilde{\gamma}_X^i$ denote the preimage of γ_X^i in Γ_i. We concatenate the corridors Γ_i by identifying the unique face of Γ_{i+1} which contains $\tilde{\gamma}_X^{i+1}(t_{i+1})$ in its interior with the unique face of Γ_i which contains $\tilde{\gamma}_X^i(t_{i+1})$ in its interior. We denote the resulting complex by $C(X) = \Gamma_0 * \ldots * \Gamma_{n-1}$, and identify Γ_i with its image in $C(X)$. Notice that the defining maps $\Gamma(\gamma_i^X) \to K$ induce a map $\phi : C(X) \to K$ which is length preserving.

The concatenation of the paths $\tilde{\gamma}_X^i$ form a path from $\tilde{\gamma}_X^0(0) = \tilde{p}$ to $\tilde{\gamma}_X^{n-1}(1) = \tilde{q}$ in $C(X)$, and if we denote this path by $\tilde{\gamma}_X$ then $\phi \circ \tilde{\gamma}_X = \gamma_X$. In particular, because ϕ is length preserving $E(\gamma_X) = E(\tilde{\gamma}_X)$.

On the other hand, if $\tilde{\gamma}$ is any path in $C(X)$ with the property $\tilde{\gamma}(0) = \tilde{p}$, $\tilde{\gamma}(1) = \tilde{q}$ and for every i the image of $\tilde{\gamma}|_{[t_i,t_{i+1}]}$ is contained in Γ_i then $E(\tilde{\gamma}) \geq E(Y)$ for some $Y \in \Sigma$. To see this notice that our construction of $C(X)$ was such that $\Gamma_i \cap \Gamma_{i+1}$ is the unique simplex of $C(X)$ which contains $\tilde{\gamma}_X(t_i)$ in its interior. Hence for every i the point $\phi \circ \tilde{\gamma}(t_i)$ lies in the unique simplex of K which contains x_i in its interior, and this implies that $\langle \phi \circ \tilde{\gamma}(t_1), \ldots, \phi \circ \tilde{\gamma}(t_{n-1}) \rangle$ is an element of Σ. Let Y denote this element. If $\phi \circ \tilde{\gamma}|_{[t_i,t_{i+1}]} = \gamma_Y^i$ for every i then $E(\tilde{\gamma}) = E(\phi \circ \tilde{\gamma}) = E(\gamma_Y)$, if not then $E(\tilde{\gamma}) = E(\phi \circ \tilde{\gamma}) > E(\gamma_Y)$.

Thus if we let $e(C(X); \tilde{p}, \tilde{q})$ denote the infimum of $E(\tilde{\gamma})$ taken over all such paths $\tilde{\gamma}$, then the minimum value which E attains on Σ is

$$inf \left\{ e\left(C\left(X \right); \tilde{p}, \tilde{q} \right) \, : X \in \Sigma \right\}.$$

The value of this expression depends only on the metric structure of the bipointed complex $(C(X); \tilde{p}, \tilde{q})$, and since $C(X)$ was obtained by concatenating corridors from the finite set Λ, there are only finitely many possible isometric models for $(C(X); \tilde{p}, \tilde{q})$. $\square$

We regard $V = \{E(v) : v \in P_e \text{ is a vertex of } J\}$ as the set of *critical values* of $E|_J$. Proposition 2.18 implies that this set is finite, and we write $V = \{e_1, \ldots, e_m\}$ where $0 < e_1 < \ldots < e_m \leq e$. Notice that J was constructed so that V contains all the local minima of $E|_{P_e}$.

We consider the following subcomplexes of J:

$$H_i = \bigcup \{\sigma \subset P_e : E(W) \leq e_i \ \forall \, vertices \ W \in \sigma\}, \quad 1 \leq i \leq m.$$

The strict convexity of $E|_J$ along geodesic segments implies that H_i is a *full* subcomplex of J and that $H_m \subseteq \{X \in J : E(X) \leq e\} = P_e$. It also implies that no two adjacent vertices (i.e., vertices that cobound a 1–cell in J) correspond to the same e_i. In particular, H_1 is a discrete set consisting of global minima for $E|_{P_e}$.

The following two lemmas complete the proof of Theorem 2.12.

Lemma 2.19: *P_e strong deformation retracts onto H_m.*

Proof: We constructed J as a simplicial subdivision of a neighbourhood of P_e in L. So any closed simplex σ of J which is not contained in P_e can be written uniquely as a join $(\sigma' \star \sigma'')$ where $\sigma' \subset H_m$ and σ'' is disjoint from P_e. The straight line retraction of $(\sigma - \sigma'')$ onto σ' takes $(P_e \cap \sigma)$ into itself because E is strictly decreasing along each of the line segments from σ'' to σ'. These straight line retractions of individual simplices agree on common faces, and hence combine to give the desired retraction of P_e onto H_m. $\square$

Lemma 2.20: *H_i strong deformation retracts onto the union of H_{i-1} with a discrete set of vertices at which E has a local minimum.*

Proof: Consider a vertex W with $E(W) = e_i$. If W is an isolated point of H_i then the value of E at all of the adjacent vertices of J is greater than e_i and hence E has a local minimum at W.

H_i is a *full* subcomplex of J, so if W is not an isolated point of H_i then $St(W, J) \cap H_i$ is the cone $W \star (link(W, J) \cap H_{i-1})$, i.e. it is the union of the line segments $\{[X, W] : X \in link(W, J) \text{ with } E(X) < e_i\}$. We use this cone structure to define a strong deformation retraction D of $St(W, J) \cap H_i$ onto $link(W, J) \cap H_{i-1}$ as follows:

Fix a point $X_0 \in link(W, J) \cap H_{i-1}$ and let $\alpha : [0, 1] \to St(W, J)$ denote the line segment $[W, X_0]$. We define D by requiring that at time t the map $Y \mapsto D(Y, t)$ sends the line segment $[X, W]$ linearly onto the unique geodesic segment from X to $\alpha(t)$ in J, for every $X \in (link(W; J) \cap H_{i-1})$ and $t \in [0, 1]$. Lemma 2.3 ensures that the homotopy is continuous. And the convexity of $E|_J$ implies that $[X, \alpha(t)]$ is contained in H_i for all $t \in [0, 1]$, and is contained in H_{i-1} for $t = 1$. (In general the image of D may not be contained in $St(W)$, but this is not important.) □

2.5 Extending Geodesics

In Chapter 1 we showed that if K is a complex of type A or B then its intrinsic metric is complete. One would like to deduce from this that if K is simply connected and has non-positive curvature then any isometric embedding of an interval of the real line can be extended to an isometric embedding of the whole line (i.e., geodesic segments in K can be extended indefinitely). This is not true as stated, but to make it true we need only make an exclusion analogous to that which one would make for boundary points in a complete Riemannian manifold of non-positive sectional curvature.

Proposition 2.21: *Suppose that K is a simply connected piecewise Euclidean complex of type A or B which satisfies any of the characterisations of non-positive curvature given in Theorems 2.7 and 2.8. Then any isometric map from an interval of the real line into K can be extended to an isometric embedding of the whole real line, or to an embedding of a closed interval the image of whose endpoints lie in the set*

$$\Sigma = \{x \in K : \; LK(x, K) \;\; is \;\; contractible\}.$$

Proof: K is complete, so any isometric map of an interval of the real line into K can be extended to the closure of that interval. So it is enough to prove that we can extend any isometry $\alpha : I \to K$ where I has a finite endpoint a with $x = \alpha(a) \notin \Sigma$.

We are assuming that K satisfies the link condition, so there must exist $Q \in LK(x, K)$ with $d(P, Q) \geq \pi$. For if this were not the case then there would be a unique geodesic segment from P to every point of $LK(x, K)$, implying that $LK(x, K)$ is contractible, contrary to hypothesis.

Let q be the image of Q under the natural identification of $LK(x, K)$ with the sphere $S_{\epsilon(x)/2}(x)$. The angle between the geodesic segments α and $[x, q]$ at x is $d(P, Q)$, and hence is at least π. So extending α by $[x, q]$ we get a local geodesic. But any local geodesic in K is a geodesic segment. □

3. M(κ)-Complexes of Non-positive Curvature

All of the results which we proved for piecewise Euclidean complexes in Chapter 2 have analogues in the world of piecewise hyperbolic complexes. Indeed, it was only for clarity of exposition that we did not give a unified treatment of these results in the previous chapter. This we do now, by considering $M(\kappa)$–simplicial complexes with $\kappa \leq 0$.

3.1 The Proof of The Main Theorem

The following definitions are needed for the statement of the Main Theorem.

If $\chi > 0$ then we say that a geodesic triangle is χ*–small* if the sum of the lengths of its sides is no greater than $2\pi/\sqrt{\chi}$. (If $\chi \leq 0$ then the condition of being χ*–small* is vacuous.)

Following Gromov [17] we say that a geodesic metric space X has curvature $\leq \chi$ if it satisfies the following condition locally.

CAT(χ): Let $T = \Delta(x_0, x_1, x_2)$ be a χ–small geodesic triangle in X, and let y be a point on the side of T which has endpoints x_1 and x_2. Choose a comparison triangle $T' = \Delta(x_0', x_1', x_2')$ in $M(\kappa)^2$ (the plane of constant curvature κ) and let y' denote the unique point on the geodesic segment $[x_1', x_2']$ such that $d_{M(\kappa)^2}(x_i', y') = d_K(x_i, y)$ for $i = 1, 2$. Then $d_{M(\kappa)^2}(x_0', y') \geq d_K(x_0, y)$.

The Main Theorem: *If K is a simply connected $M(\kappa)-$simplicial complex of type A or B and $\kappa \leq 0$ then the following 13 conditions are equivalent:*

Global conditions:

I) *K has unique geodesic segments.*

II) *K satisfies $CAT(\kappa)$ globally.*

III) *K satisfies $CAT(\chi)$ globally, for some χ.*

IV) *K satisfies CN globally.*

V) *The metric on K is convex.*

VI) *Every geodesic triangle in K has non-positive excess.*

Local conditions:

VII) *K has unique geodesic segments locally.*

VIII) *K satisfies $CAT(\kappa)$ locally.*

IX) *K satisfies $CAT(\chi)$ locally, for some χ.*

X) *K satisfies CN locally.*

XI) *The metric on K is convex locally.*

XII) *Every point of K has a neighbourhood such that any geodesic triangle contained in that neighbourhood has non-positive excess.*

XIII) *K satisfies the link condition.*

In Chapter 2 we proved a special case of the Main Theorem, and for the most part the proofs which we gave there are equally valid (*mutatis mutandis*) in the present setting. When this is the case we do not repeat the details of the argument, or the accompanying motivation.

Remark: At first glance it may seem strange that in this class of spaces

$CAT(\kappa)$ is equivalent to $CAT(\chi)$, where χ is arbitrary. However this merely reflects the fact that if we metrize the simplices of a given complex to have constant negative curvature, then in doing so we redistribute the natural curvature of the complex, and concentrate it in the skeleton of codimension 2. In particular, if the given complex does not support a structure of negative curvature (e.g. if it is not aspherical) then forcing the simplices to be negatively curved concentrates an infinite amount of positive curvature at certain points in the complex, and thus any condition which gives a bound on the local curvature of the space, such as $CAT(\chi)$, fails in the neighbourhood of such a point.

For the remainder of Chapter 3 the letter K, without further qualification denotes an $M(\kappa)$–simplicial complex of type A or B.

Intuitively speaking, a geodesic metric space X has curvature $\leq \kappa$ if geodesic segments in X diverge more quickly than geodesics in $M(\kappa)^2$. To make this notion precise we introduce the following "divergence function".

Given $\theta \in [0, \pi]$ and $a, b \in [0, \infty)$ we can choose geodesic segments $\alpha, \beta : [0, 1] \to M(\kappa)^2$ with lengths $L(\alpha) = a$ and $L(\beta) = b$, which have a common initial point and meet at an angle θ. We then define

$$D(\theta, a, b) = d(\alpha(1), \beta(1)).$$

The homogeneity of $M(\kappa)^2$ ensures that this definition is independent of the choices of α and β.

In the proof of Theorem 3.1 we shall need the fact that for fixed a and b the function $\theta \mapsto D(\theta, a, b)$ is monotone increasing. This follows from the cosine rule in $M(\kappa)^2$. We shall also need the fact that for all $a, b, c > 0$, and all $\theta, \theta', \theta'' \in [0, \pi]$ such that $\theta = \theta' + \theta''$,

$$D(\theta, a, c) \leq D(\theta', a, b) + D(\theta'', b, c).$$

This is simply a restatement of the triangle inequality in $M(\kappa)^2$.

The following comparison theorem is closely analoguous to the Topogonov inequality for manifolds of non-positive curvature [5].

Theorem 3.1: *Suppose that K has unique geodesic segments, and let α_0 and α_1 be geodesic segments in K with common initial point $\alpha_0(0) = \alpha_1(0) = x$. If the angle between α_0 and α_1 at x (which we denote by $\theta_{0,1}$) is less than π then*

$$D\left(\theta_{0,1}, L\left(\alpha_0\right), L\left(\alpha_1\right)\right) \leq d\left(\alpha_0\left(1\right), \alpha_1\left(1\right)\right).$$

Our proof of Theorem 3.1 relies on the following lemma, which can be proved using the construction given in Lemma 2.4. (See Figure 2.1.)

Lemma 3.2: *If K has unique geodesic segments then for every $\eta > 0$ and every local geodesic α in K, there exists a constant $\epsilon > 0$ with the following property: If β is a local geodesic in K for which $\beta(0) = \alpha(0)$ and $\|\alpha - \beta\| < \epsilon$ then Θ, the angle between α and β at $\alpha(0)$, is less than η and*

$$D\left(\Theta, L\left(\alpha\right), L\left(\beta\right)\right) \leq d\left(\alpha\left(1\right), \beta\left(1\right)\right).$$

Corollary : *If K has unique geodesic segments then every local geodesic in K is a geodesic segment. Hence there is a unique local geodesic joining any two points in K.*

Proof of Theorem 3.1: Let σ denote the unique geodesic segment from $\alpha_0(1)$ to $\alpha_1(1)$. We denote the geodesic segment from $\alpha_0(0)$ to $\sigma(s)$ by α_s, and the angle between α_s and α_t at x by $\theta_{s,t}$. It follows from the preceding corollary that $\theta_{s,t} < \pi$ for all $s, t \in [0, 1)$.

We shall prove that the set

$$\Sigma = \left\{s \, : D\left(\theta_{0,s}, L\left(\alpha_0\right), L\left(\alpha_s\right)\right) \leq d_K\left(\alpha_0\left(1\right), \alpha_s\left(1\right)\right)\right\}$$

is the whole of $[0, 1]$.

Lemma 3.2 implies that $\theta_{0,s}$ and $L(\alpha_s)$ vary continuously with s, and hence Σ is closed. To see that it is open, fix $s < 1$ with $s \in \Sigma$. In Lemma 2.3 we proved that geodesic segements vary continuously with their endpoints, so it follows from Lemma 3.2 that if $\delta > 0$ is small enough then $\theta_{s,s+\delta} < \pi - \theta_{0,s}$ and

$$D(\theta_{s,s+\delta}, L(\alpha_s), L(\alpha_{s+\delta})) \leq d(\alpha_s(1), \alpha_{s+\delta}(1)).$$

The following inequality shows that $s + \delta \in \Sigma$.

$$\begin{aligned} d(\alpha_0(1), \alpha_{s+\delta}(1)) &= d(\alpha_0(1), \alpha_s(1)) + d(\alpha_s(1), \alpha_{s+\delta}(1)) \\ &\geq D(\theta_{0,s}, L(\alpha_0), L(\alpha_s)) + D(\theta_{s,s+\delta}, L(\alpha_s), L(\alpha_{s+\delta})) \\ &\geq D(\theta_{0,s} + \theta_{s,s+\delta}, L(\alpha_0), L(\alpha_{s+\delta})) \\ &\geq D(\theta_{0,s+\delta}, L(\alpha_0), L(\alpha_{s+\delta})). \end{aligned}$$

Here we have used the triangle inequality in $LK(x, K)$ to deduce that $\theta_{0,s+\delta} \leq \theta_{0,s} + \theta_{s,s+\delta}$. □

Lemma 3.3: *If a geodesic metric space X satifies $CAT(0)$ then the metric on X is convex.*

Proof: Fix $t \in [0, 1]$. Given geodesic segments α and β in X with $\alpha(0) = \beta(0)$, and $s \in [0, 1]$, we consider the geodesic triangles $\Delta_s \subseteq X$ which have vertices $\{\alpha(0), \alpha(1), \beta(s)\}$. Let $\Delta'_s \subseteq \mathbf{E}^2$ be a comparison triangle for Δ_s. We denote the image of $z \in \Delta_s$ in Δ'_s by z'_s, and the vertex angle at $\alpha(0)'_s$ by ϕ_s.

$CAT(0)$ applied to Δ_1 yields $d(\beta(t), \alpha(1)) \leq d(\beta(t)'_1, \alpha(1)'_1)$. And since $d(\beta(t)'_t, \alpha(1)'_t) = d(\beta(t), \alpha(1))$, this implies that $\phi_t \leq \phi_1$. Hence

$$d(\alpha(t)'_t, \beta(t)'_t) \leq d(\alpha(t)'_1, \beta(t)'_1).$$

$CAT(0)$ applied to Δ_t yields $d(\alpha(t), \beta(t)) \leq d(\beta(t)'_t, \alpha(t)'_t)$. Hence

$$d(\alpha(t), \beta(t)) \leq d(\alpha(t)'_1, \beta(t)'_1).$$

Finally, from elementary Euclidean geometry, we have

$$d\left(\alpha(t)_1', \beta(t)_1'\right) = t\ d\left(\alpha(1)_1', \beta(1)_1'\right) = t\ d(\alpha(1), \beta(1)).$$

Hence

$$d(\alpha(t), \beta(t)) \leq t\ d(\alpha(1), \beta(1)).$$

The fact that the function $t \mapsto d(\alpha(t), \beta(t))$ is a convex function for any geodesics segments α and β in K follows easily from the case $\alpha(0) = \beta(0)$, as we showed in Section 2.1. □

Proof of The Main Theorem:

I⇒VI : Theorem 3.1 implies that if geodesic segments are unique in K then each of the vertex angles of a geodesic triangle in K is no larger than the corresponding angle of a comparison triangle in $M(\kappa)^2$.

I⇒II : In the proof of Theorem 2.7 we used a special case of the Topogonov inequality (Lemma 2.5) to show that a piecewise Euclidean complex with unique geodesic segments satisfies $CAT(0)$. A direct translation of that argument, employing Theorem 3.2 in place of Lemma 2.5, proves the present implication.

II⇒IV : If X is any geodesic metric space which satisfies $CAT(\kappa)$ then it satisfies $CAT(\chi)$ for all $\chi \geq \kappa$. And in the proof of Theorem 2.5 we proved that if X satisfies $CAT(0)$ then it satisfies CN.

II⇒V: This is the content of Lemma 3.3.

IV, V, VI⇒I: These implications follow easily from the definition of the given conditions, as we verified in Section 2.3.

Local versions of the preceding arguments prove the equivalence of conditions VII, VIII, XI, X and XII.

XIII⇔VII: For the case $\kappa = 0$ this is the content of Proposition 2.11. In the proof of that proposition the only point at which we used the hypothesis that K was piecewise Euclidean was to employ the cosine rule when comparing the length of a chain in $st(x, K)$ to the length of the projected chain in $LK(x, K)$. The cosine rule in hyperbolic geometry serves the desired purpose equally well. Hence the conclusion of Proposition 2.11 is valid for piecewise hyperbolic complexes.

II⇒III⇒IX: Trivial.

IX⇒VII: Immediate from the definition.

As in the case $\kappa = 0$ we pass from the local to the global situation by proving VII⇒I (which completes the proof of the theorem). To do so we observe that the results presented in Section 2.4 depend only on the local convexity properties of the metric on K, and hence are equally valid for piecewise hyperbolic complexes. One might be concerned by the fact that in the hyperbolic case the $n-$fold Cartesian product K^n is not itself piecewise hyperbolic, but this is of no consequence since our results concerning the geometry of K^n relied only on the fact that a path in K^n is a geodesic segment if and only if its projection onto each coordinate is a geodesic segment. And this observation remains valid in the piecewise hyperbolic case. □

3.2 Alternative Forms of the Link Condition

The *link condition*, which we defined in Section 2.3, is a condition which bounds the curvature of the spherical simplicial complexes which occur as links in some given complex K. It seems natural to ask whether one can prove a result, analogous to the Main Theorem, which would relate this condition to other characterisations of bounded curvature in spherical simplicial complexes.

Such a result might then yield alternative ways of verifying the link condition for a given complex.

In fact the Main Theorem does not have a direct analogue in the world of spherical simplicial complexes, but in this section we outline the proof of a partial analogue (Lemma 3.4) which gives two alternative characterisations of bounded curvature in the link complexes of an $M(\kappa)$–simplicial complex of type A or B. Each of these conditions is equivalent to the link condition, and in Proposition 4.5 we use one of these characterisations to show that a certain class of 3–complexes are non-positively curved.

Using the techniques of Chapter 2 one can show that if L is a spherical simplicial complex of type A or B then it satisfies Gromov's $CAT(1)$ condition if and only if there is a *unique* geodesic segment from x to y in L whenever $d(x, y) < \pi$. Similarly, one can show that L has unique geodesics segments locally if and only if it satisfies $CAT(1)$ locally. However, if L is simply connected and satisfies $CAT(1)$ locally then it does *not* follow that L satisfies $CAT(1)$ globally, as the following example shows.

Example: Consider a geodesic triangulation of the standard 2–sphere cut open along a 1–simplex of length less than π. Let L denote the spherical 2–complex obtained by identifying two copies of this space along the boundary loop γ. L is (topologically) a 2–sphere, and satisfies $CAT(1)$ locally. However, it does not satisfy $CAT(1)$ globally, because the image of γ in L is a geodesic circle of length strictly less than π.

This example highlights the difficulty which arises when trying to relate local descriptions of bounded curvature in spherical simplicial complexes to global descriptions. Namely, that one must take account of the possible occurence of short geodesic circles (i.e., isometric embeddings of a circle of length $< 2\pi$).

This motivates the following definition. The *systole* of L, which we denote $sys(L)$, is the infimum of the lengths of geodesic circles in L. We wish to relate $sys(L)$ to $inj(L)$, the *injectivity radius* of L, which is defined to be $sup\{r :$ there is a unique geodesic segment from x to y in L whenever $d(x,y) \leq r\}$.

The argument given by Charney and Davis in Lemma 1.3 of [11], shows that if L is a finite spherical simplicial complex with $inj(L) > 0$ then it fails to satisfy $CAT(1)$ if and only if $\frac{1}{2}sys(L) = inj(L) < \pi$. The hypothesis that the complex is finite is only used to ensure that L has the following property: If $inj(L) \in (0, \infty)$ then there is a (non-degenerate) geodesic bigon in L whose sides have length exactly $inj(L)$.

By employing the method of finite models (which we used in Lemmas 1.5 and 2.3) it is easy to show that complexes of type B also have this property. Moreover, if L is of type B then Lemma 2.10 implies that $inj(L) > 0$ if and only if L has unique geodesics locally. Hence we obtain the following result (cf. [17], p.120).

Lemma 3.4: *If K is an $M(\kappa)$–simplicial complex of type A or B (where κ is arbitrary) then the following are equivalent:*

I) *K satisfies the link condition. (i.e., $inj(LK(x,K)) \geq \pi$ for every $x \in K$).*

II) *For every $x \in K$ the complex $LK(x,K)$ has unique geodesic segments locally and $sys(LK(x,K)) \geq 2\pi$.*

III) *For every $x \in K$ the complex $LK(x,K)$ satisfies $CAT(1)$ globally.*

4. Examples

In this chapter we describe some examples of complexes of non-positive curvature. In particular we are interested in complexes which are not locally finite. Interesting examples of this type arise in the work of Gersten and Stallings on triangles of groups, and in Section 3.1 we describe their results in the context of the work presented here. Then in Section 3.2 we discuss non-positively curved 3–complexes which have planar links. A particular example of such a complex arises in the study of the group $GL_n(Z[\omega])$, where ω is a primitive sixth root of unity. We use this example to illustrate how the Fixed Point Theorem proved in Section 2.2 can be used to classify the finite subgroups of a group which acts on a complex of non-positive curvature. Before proceeding to Section 2.1 we briefly mention two other classes of examples.

A particularly rich source of non-positively curved complexes is provided by Euclidean buildings. The theory of buildings is extensive and well-documented (see for example [8], [9], [10], [28], and [29]). So too is the role of non-positive curvature in understanding the geometry of buildings of Euclidean type. We do not attempt to develop anything of this theory here, but merely note that Euclidean buildings provide interesting examples of piecewise Euclidean complexes of type B.

Another class of examples which is well-understood is that of metric simplicial trees. Any metric simplicial tree satisfies $CAT(\kappa)$ for all κ, and hence "a non-positively curved simply connected 1–dimensional Euclidean simplicial complex of type B" is just another name for a metric simplicial tree in which the set of edge lengths is finite. Groups which act on such complexes were completely classified by the work of Bass and Serre [24].

4.1 Simplices of Groups.

The simplest example of a group which acts cocompactly (but not freely) on a tree is an amalgamated free product. So it seems reasonable that as a first step towards generalising Bass-Serre theory to higher dimensions one should look for a 2–dimensional analogue of an amalgamated free product. This is the starting point for the work of Gersten and Stallings on triangles of groups.

Definition: An *n-simplex* of groups is a contravariant functor from the poset of faces of an $n-$simplex ordered by inclusion, into the category of groups and monomorphisms.

We think of such a functor T as a diagram of groups, and refer to the image under T of a vertex as a vertex group, the image of a 1–simplex as an edge group, and so on. We denote the direct limit (or generalised pushout) of this diagram in the category of groups by $\Gamma(T)$.

Definition: An $n-$simplex of groups T is said to be *realisable* if the canonical map from each vertex group of T into $\Gamma(T)$ is an injection.

There is an obvious paradigm in this situation. Namely, if a group acts without inversions on a simplicial complex, and the quotient is a single $n-$simplex then the diagram of stabilisers and inclusions in a fundamental domain is a realisable $n-$simplex of groups. (A group is said to act without inversions if the fixed point set of any element is a simplicial subcomplex.)

Definition: We say that an $n-$simplex of groups is *geometric* if it arises as the diagram of stabilisers and inclusions for the fundamental domain of the action of a group of isometries on a piecewise Euclidean complex of non-positive curvature. (Since the quotient is compact the complex will necessarily be of type B.)

Remark: This use of the term "geometric" is not a standard one.

If T is a 1–simplex of groups then $\Gamma(T)$ is the amalgamated free product of the vertex groups over the edge group. It is well-known that every 1–simplex of groups is realisable and geometric [24]. This is far from true in dimension 2, but a sufficient condition for a triangle of groups to be both realisable and geometric has been given by Gersten and Stallings, using the idea of the angle between subgroups of an arbitrary group.

Consider a group G and subgroups A, B, C with $C \subseteq A \cap B$. Fix a set of coset representatives $\{g_i\}_{i \in I}$ for G/C. Let $L_G(A, B; C)$ be the (unoriented) graph with vertex set $G/A \coprod G/B$ and 1–simplices $\{\{g_iA, g_iB\} \mid i \in I\}$.

Definition: The *angle in G between A and B as measured over C* is $2\pi/n$, where n is the length of the shortest reduced circuit in $L_G(A, B; C)$.

Definition: Given a triangle of groups T one associates to each vertex the angle between the incident edge groups as measured over the 2–cell group. The triangle is said to be *non-spherical* if the sum of the vertex angles is no greater than π.

The following theorem is proved in [26]. (An alternative proof is given in [7].)

Theorem 4.1: *(Gersten-Stallings) Every non-spherical triangle of groups is realisable.*

Gersten and Stallings also proved that a non-spherical triangle of groups is geometric. We shall now outline a proof of this fact. The proof which we give here is somewhat different to that which was originally given by Gersten and Stallings.

Fix a non-spherical triangle of groups T, and let Γ denote the direct limit of the corresponding diagram of groups. We assume that each of the vertex angles is non-zero. (If one of the vertex angles is zero, then the triangle degenerates to an amalgamated free product.) We denote the edge groups by E_1, E_2, E_3 and the vertex group with incident edge groups E_i and E_j by $V_{i,j}$. The 2–cell group shall be denoted C.

We define an (abstract) simplicial 2–complex $K(T)$ as follows: (For convenience we write all indices mod 3.)

$$K(T)^0 = \coprod_{i<j} \Gamma / V_{i,j}$$

$$K(T)^1 = \coprod_{i=1}^{3} \{ \{ e^i_\lambda V_{i,i+1}, e^i_\lambda V_{i,i-1} \} \mid e^i_\lambda \text{ coset reps for } \Gamma / E_i \}$$

$$K(T)^2 = \{ \{ c_\eta V_{1,2}, c_\eta V_{2,3}, c_\eta V_{3,1} \} \mid c_\eta \text{ coset reps for } \Gamma / C \}.$$

There is a natural action of $\Gamma(T)$ on K, given by left multiplication of cosets. The quotient space for this action is a single 2–simplex, and the cosets containing the identity element form a fundamental domain. Moreover, the pattern of stabilisers in this fundamental domain is precisely the original triangle of groups T. One can also show that the geometric realisation of $K(T)$ is simply connected [7].

Notice that $K(T)$ has a natural labelling, for example a vertex corresponding to a coset of $V_{1,2}$ in Γ is thought of as being labelled $\{1, 2\}$, a 1–simplex corresponding to a coset of E_3 in Γ is thought of as having label $\{3\}$, and so on. If we fix a (hyperbolic or Euclidean) triangle Δ with vertices indexed 1 to 3, and each vertex angle equal to the group theoretic angle at the corresponding vertex of T, then this labelling of $K(T)$ induces a simplicial isomorphism from each 2–simplex in $K(T)$ to Δ. The collection of these maps satisfies the axioms given in Section 1.1, and hence we obtain a metric simplicial complex of type B.

The labelling on $K(T)$ also gives a graph isomorphism from the link of any vertex v which is labelled $\{i, j\}$ to the graph $L_{V_{i,j}}(E_i, E_j; C)$. The metric on the link of a vertex in a 2–dimensional complex is given by the angular measure at the vertex. Thus each edge of the metric graph $LK(v, K(T))$ has length π/n. But n was defined to be the length of the shortest reduced circuit in $L_{V_{i,j}}(E_i, E_j; C)$. Thus for every vertex $v \in K$ the graph $LK(v, K(T))$ contains no geodesic circles of length less than 2π, and hence $K(T)$ satisfies the link condition. Thus we have proved:

Theorem 4.2: *(Gersten-Stallings) Every non-spherical triangle of groups is geometric.*

The following result now follows immediately from the Fixed Point Theorem, which we proved at the end of Section 2.2.

Theorem 4.3: *(Gersten-Stallings) If T is a non-spherical triangle of groups then every bounded (e.g. finite) subgroup of $\Gamma(T)$ is conjugate to a subgroup of one of the vertex groups.*

Recently, Haefliger [19] has shown that any non-positively curved orbihedron which has only finitely many isometry types of cells arises as the quotient of a non-positively curved piecewise Euclidean complex by a group of isometries. This generalises Theorems 4.1 and 4.2.

4.2 3–Complexes With Planar Links.

The link of a vertex in a 2–dimensional metric simplicial complex is a graph, with edge lengths given by the angular measure at the vertex. Thus to verify the link condition it is enough to calculate the lengths of reduced circuits in each such graph. In higher dimensions things are much more delicate, because it is difficult to identify geodesics in the link complexes. However, if one has sufficiently explicit knowledge about the structure of the links then it may still be possible to decide whether a given complex satisfies the link condition. (See for example [23], [11] and Chapter 5 below.) One can also prove more general results for classes of complexes in which the structure of the links is sufficiently simple. For example:

Proposition 4.4: *Suppose that K is an (abstract) simplicial 3–complex such that the link of each vertex is simplicially isomorphic to a triangulation of the plane in which every vertex has valence at least 6. Fix a regular tetrahedron T in $M(\kappa)^3$ (the unique simply connected 3–manifold of constant sectional curvature $\kappa \leq 0$). If we metrize each 3–simplex in K by means of a simplicial isomorphism to T then K satisfies the link condition.*

Our proof of Proposition 4.4 requires the following simple fact from spherical geometry: Suppose that the paths $\alpha, \beta : [0,1] \to S^2$ satisfy $\alpha(0) = \beta(0)$ and $\alpha(1) = \beta(1)$, and that they cobound an embedded disc $D \subset S^2$. We say that D is *good* if β is an arc of a great circle, and α is a piecewise geodesic path of length strictly less than π with the property that the angle between its sucessive geodesic subarcs, as measured in D, is at least π.

Lemma 4.5: *Every good disc in S^2 contains an equilateral spherical triangle of side $\pi/2$.*

Proof: We may assume that $\alpha(0) = \beta(0)$ is the north pole, that the initial segment of β follows the line of longitude $0°W$, and that the longitudinal coordinate of $\alpha(t)$ is $a(t)°W$, where a is an increasing function of t.

Notice that $a(t) \geq 180$ for some $t \in (0,1]$. Hence α does not meet the equator between $0°W$ and $90°W$. For if $\alpha(t)$ did lie on this arc then the paths $\alpha|_{[0,t]}$ and $\alpha|_{[t,1]}$ would both have length at least $\pi/2$, contradicting the fact that the length of α is strictly less than π. Therefore the south-west octant of the sphere is contained in D. □

Corollary : *If Γ is a 2–dimensional* **S**-*corridor whose cells are equilateral spherical triangles of side $l < \pi/2$, then there is a unique local geodesic from x to y for all $x, y \in \Gamma$.*

Proof: By Theorem 1.1 there is a geodesic segment α from x to y in Γ. Any geodesic segment is a local geodesic, so the only issue is that of uniqueness. Suppose that γ is another local geodesic from x to y. Restricting to subpaths if necessary we may assume that α and γ cobound an embedded disc $D \subset \Gamma$. If we orient D then for every $x \in D$ we have a well-defined notion of the clockwise and anticlockwise directions in the spherical complex $LK(x, \Gamma)$ (which is a circle or an arc of a circle).

Because α and γ are local geodesics, the paths in $LK(\alpha(t), \Gamma)$ and $LK(\gamma(t), \Gamma)$ determined by D each have length at least π, for every $t \in [0, 1]$. In fact, we may assume that the path which D determines in $LK(\gamma(t), \Gamma)$ has length exactly π. For if this is not the case then we can replace γ by the unique local geodesic β which has the same initial segment and for which the distance in $LK(\beta(t), \Gamma)$ from the point determined by the forward direction of β to that determined by the backward direction is exactly π, when measured in the anticlockwise direction. We can then replace α by an initial segment so that the disc D' cobounded by α and β is non-singular.

However, if such a disc $D' \subset \Gamma$ were to exist then it would be isometric to a good disc in S^2, and Lemma 4.5 would then give an isometry from an equilateral spherical triangle of side $\pi/2$ into Γ. No such map exists. □

Proof of Proposition 4.4: Fix $x \in K$, and write $L = LK(x, K)$. The 2–cells of L are equilateral spherical triangles with vertex angles strictly between $\pi/3$ and $\pi/2$, so the link of a point $p \in L$ is a circle of length $> 2\pi$ if p is a vertex of L, and of length 2π if p is not a vertex. Thus L satisfies the link condition and has unique geodesic segments locally. So by Lemma 3.4 it is enough to prove that there are no geodesic circles of length less than 2π in L.

We use the phrase *a tesselation line in* L to describe a locally isometric embedding of the real line into L whose image lies in the 1–skeleton. Notice that every edge can be extended to a tesselation line, and that if a geodesic circle in L crosses a tesselation line then it must cross it at least twice.

Suppose that $\sigma \subset L$ is a geodesic circle of length less than 2π. Because L is homeomorphic to the plane, σ bounds an open disc in L. Consider a tesselation line which meets this disc in an arc γ which is outermost among all tesselation lines which intersect the disc. Let α denote the (short) arc of σ which has the same endpoints as γ. Because γ is outermost, there is a corridor in L which contains both α and γ. But these paths are both local geodesics, and the length of α is strictly less than π. This contradicts the preceding corollary. □

Remark: The 3–dimensional nature of Proposition 4.4 is something of an illusion, because any complex K which satisfies the given hypotheses strong deformation retracts onto the following 2–dimensional subcomplex of its first barycentric subdivision.

$$L = K' - \bigcup \{st\,(v, K')\ : v \text{ a vertex of } K\}$$

L can be metrized as a Euclidean simplicial complex of non-positive curvature in such a way that any simplicial isomorphism of L is an isometry. To do this one first observes that for every vertex $v \in K$ the complex $link(v, K')$ is simplicially isomorphic to the first barycentric subdivision of a tesselation of the plane in which every vertex has valence at least 6. Let T_v denote this tesselation. We metrize each 2–simplex in T_v as a Euclidean equilateral triangle of side 1, and give $link(v, K')$ the induced metric. This defines a Euclidean simplicial structure on L which is of both type A and type B, and it is easy to check that it also satisfies the link condition.

Notice that if G is a group which acts on K by simplicial isomorphisms then the retraction of K onto L can be done G–equivariantly, and the induced action is by isometries. In particular, if the vertex stabilisers for this induced action are finite then G is (equi-)semihyperbolic in the sense of [3].

A complex for $\mathbf{GL_2(Z[\omega])}$

An example of a simplicial complex which satisfies the hypotheses of Proposition 4.4 is the following complex which was studied by Roger Alperin [4].

Let R be a ring, and consider the set Λ of free direct summands for R^2. We say that $L_1, L_2 \in \Lambda$ are independent if $L_1 + L_2 = L_1 \oplus L_2 = R^2$. Let $K(R)$ be the (abstract) simplicial complex whose vertices are the elements of Λ and whose q–simplices are those subsets of Λ of the form $\{L_0, \ldots, L_q\}$ for which L_i, L_j are independent for $0 \leq i \neq j \leq q$. The action of $GL_2(R)$ on Λ preserves the relation of independence, and hence gives an induced action on $K(R)$.

Let ω be a primitive sixth root of unity. In the case $R = GL_2(Z[\omega])$ one can show that $K(R)$ is a simply connected 3–complex, and that the link of every vertex is simplicially isomorphic to the standard tesselation of the Euclidean plane by equilateral triangles. It follows from Proposition 4.4 that $K(R)$ can be metrized as a piecewise hyperbolic complex of negative curvature. In particular $K(R)$ is contractible. In [4] Alperin showed that $K(R)$ was contractible by studying a filtration of the space by subcomplexes. He then calculated the Euler charcteristic and homology of $SL_2(Z[\omega])$ by studying its action on the

2–dimensional retract of $K(R)$ which we described in the preceding remark.

According to our Fixed Point Theorem any finite subgroup of $GL_2(Z[\omega])$ stabilises a point of $K(R)$. We shall show that this leads to a classification of the finite subgroups in $GL_2(Z[\omega])$.

The centre of $GL_2(Z[\omega])$ acts trivially on $K(R)$, so we get an induced action of $PGL_2(Z[\omega])$. It is this action which we study. For convenience we write Γ in place of $PGL_2(Z[\omega])$.

If a finite subgroup $H \subset GL_2(Z[\omega])$ fixes a vertex in $K(R)$ then it acts by simplicial isomorphisms on the link of that vertex. If we metrize this link as a Euclidean plane in the standard way then H acts by isometries, and hence has a fixed point. Thus any finite subgroup of Γ stabilises the barycentre $b_0(\sigma)$ of some simplex $\sigma \subset K(R)$ of dimension at least 1. The complex $link(b_0, K)$ is finite, and hence its full group of symmetries, which we denote $sym(b_0)$, is also finite.

It is easy to see that no non-trivial element of Γ fixes a 2–simplex in $K(R)$ pointwise. It follows that if σ is a simplex of dimension at least 1 and $b_0(\sigma)$ is its barycentre then $stab_\Gamma(b_0(\sigma))$ injects into $sym(b_0(\sigma))$.

Γ acts transitively on the set of simplices in $K(R)$ in each dimension. Thus if we choose simplices $\sigma_1, \sigma_2, \sigma_3$ of dimension 1, 2 and 3 respectively, then every finite subgroup of Γ is conjugate to a subgroup of $stab_\Gamma(b_0(\sigma_1)), stab_\Gamma(b_0(\sigma_2))$, or $stab_\Gamma(b_0(\sigma_3))$. We now give explicit descriptions of these stabilisers.

The complex $link(b_0(\sigma_2), K(R))$ is the suspension of the triangle $\partial\sigma_2$. Thus $sym(b_0(\sigma_2))$ is isomorphic to $S_3 \times \mathbf{Z}_2$. Because no non-trivial element of Γ fixes σ_2 pointwise, $stab_\Gamma(b_0(\sigma_2))$ must be contained in the S_3 factor. In fact it is isomorphic to S_3. To see this notice that if we let σ_2 be the 2–simplex with vertices $(1,0),(0,1)$ and $(1,1)$ then $stab_\Gamma(b_0(\sigma_2))$ contains a 3–cycle a and a transposition b given by the following matrices.

$$A = \begin{pmatrix} 0 & 1 \\ -1 & 1 \end{pmatrix} \qquad B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

The complex $link(b_0(\sigma_1), K(R))$ is the suspension of a hexagonal 1–complex, and hence $sym(b_0(\sigma_1))$ is isomorphic to $D_{12} \times \mathbf{Z}_2$. Because no non-trivial element of Γ fixes a 2–simplex in $K(R)$ pointwise, $stab_\Gamma(b_0(\sigma_1))$ does not contain the unique symmetry of the link which interchanges the suspension points and leaves the other vertices fixed. Hence $stab_\Gamma(b_0(\sigma_1))$ is isomorphic to a subgroup of D_{12}, the dihedral group of order 12.

Finally, the complex $link(b_0(\sigma_3), K(R))$ is the boundary of a tetrahedron, and hence $sym(b_0(\sigma_3))$ is isomorphic to S_4. One can show that no $\gamma \in \Gamma$ acts transitively on the vertices of a 3–simplex in $K(R)$, so $stab_\Gamma(b_0(\sigma_3))$ contains no 4–cycles. Hence it is isomorphic to a subgroup of A_4, the alternating group on four letters.

At this stage we have shown that $stab_\Gamma(b_0(\sigma_2)) \cong S_3$, $stab_\Gamma(b_0(\sigma_3)) \hookrightarrow A_4$ and $stab_\Gamma(b_0(\sigma_1)) \hookrightarrow D_{12}$. Thus if we can exhibit subgroups of Γ which are isomorphic to D_{12} and A_4 then these groups must occur as the stabilisers of a 1–simplex and a 3–simplex in $K(R)$ respectively.

Let c and d denote the elements of Γ determined by the matrices C and D (which are given below). These elements generate the subgroup $G_{c,d} = \langle c, d \mid c^6 = b^2 = (cd)^2 = 1 \rangle$, which is a dihedral group of order 12.

$$C = \begin{pmatrix} \omega & 0 \\ 0 & 1 \end{pmatrix} \qquad D = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

On the other hand, the elements $d, e \in \Gamma$ determined by the matrices D and E generate the subgroup $G_{e,f} = \langle e, f \mid e^3 = f^2 = (e^{-1}f)^3 = 1 \rangle$, which is isomorphic to A_4.

$$E = \begin{pmatrix} 0 & 1 \\ -1 & 1 \end{pmatrix} \qquad F = \begin{pmatrix} 0 & \omega \\ \omega^2 & 0 \end{pmatrix}$$

Thus any finite subgroup of $PGL_2(Z[\omega])$ is conjugate to a subgroup of one of the groups $G_{a,b}, G_{c,d}$ or $G_{e,f}$. It follows that any finite subgroup of $GL_2(Z[\omega])$ is conjugate to a subgroup in the preimage of one of these three groups.

5. The Curvature of the Culler-Vogtmann Complex

In this chapter we exemplify a method for deciding if a given complex can be given a piecewise Euclidean structure of non-positive curvature. The idea is to develop a good understanding of the local (combinatorial) structure of the space, and use this to decide whether or not the given complex can satisfy the link condition.

More specifically: let G be a group which acts on a simplicial complex K with finite quotient, and suppose that the quotient space, the local structure of K, and the $G-$stabilisers of the vertices of K can be described explicitly. One would like to know whether or not the space K can be given a $G-$equivariant piecewise Euclidean structure of non-positive curvature. We assume (for the sake of argument) that K/G has been metrized as a piecewise Euclidean complex such that the induced structure on K has non-positive curvature. In Chapter 1 we described the induced spherical simplicial structure on the links of points in K. By analysing the symmetries of these spherical complexes we can identify "conjugate points", i.e., pairs of points which cannot be joined by a unique geodesic segment in the link complex. We are supposing that K satisfies the link condition, so any path joining a pair of conjugate points must have length at least π. Thus from each each choice of such a path we obtain an inequality involving the angles of the cells in the quotient space. The aim is to gather sufficient information to determine these angles, or else to obtain contradictory bounds and hence deduce that K does not support a $G-$equivariant piecewise Euclidean structure of non-positive curvature. The same method can of course be applied when dealing with piecewise hyperbolic structures.

In Section 5.1 we shall describe the action of $Out(F_n)$, the group of outer automorphisms of the free group of rank n, on the Culler-Vogtmann complex K_n. It has been proved ([12], [13], [20]) that K_n is contractible and that the fixed point set of any finite subgroup of $Out(F_n)$ is non-empty and contractible. We saw in Chapter 2 that these properties are indicative of non-positive curvature, so it seems reasonable to ask whether or not K_n supports an $Out(F_n)$–equivariant piecewise Euclidean structure of non-positive curvature. The answer to this question is yes in the case $n = 2$, where it is known that $Out(F_2)$ is isomorphic to $SL_2(\mathbf{Z})$, and that its action on K_2 is the usual action of $SL_2(\mathbf{Z})$ on the Serre tree (see [13]). For $n \geq 3$ we shall prove that K_n does *not* support an $Out(F_n)$–equivariant piecewise Euclidean (or piecewise hyperbolic) structure of non-positive curvature, by using the technique described above.

5.1 The Culler-Vogtmann Complex K_n

In this section we describe the Culler-Vogtmann complex K_n as the geometric realisation of a certain poset consisting of marked graphs. We then describe the natural action of $Out(F_n)$ on this poset. This action is order preserving and hence induces a simplicial action of $Out(F_n)$ on K_n.

Definitions

A *graph* G is a connected 1–dimensional CW-complex. G is said to be *admissible* if it is not homotopy equivalent to any proper subgraph, all of its vertices have valence at least three, and the complement of every edge is connected. Henceforth all graphs are required to be admissible. A graph R is called a *rose* if it has a single vertex, which we denote $v(R)$.

Fix a rose R_0 and identify F_n with $\pi_1(R_0, v(R_0))$. We call this the *standard rose*. A *marking* is a homotopy equivalence $g : R_0 \to G$, where G is an admissible graph. Two markings $g_1 : R_0 \to G_1$ and $g_2 : R_0 \to G_2$ are said to be equivalent if there is a graph isomorphism $i : G_1 \to G_2$ such that the following diagram commutes up to free homotopy.

$$\begin{array}{ccc} R_0 & \xrightarrow{g_1} & G_1 \\ \| & & \downarrow i \\ R_0 & \xrightarrow{g_2} & G_2 \end{array}$$

A *marked graph* is an equivalence class of markings, and the class containing $g : R_0 \to G$ is denoted (g, G).

Let e be an edge of the graph G. We say that (g', G') is obtained from (g, G) by *blowing down the edge* e if there is a cellular homotopy equivalence $d : G \to G'$ which collapses e, is one-to-one on the complement of e, and satisfies $d \circ g \simeq g'$.

Description of K_n

We define a partial ordering on the set of marked graphs by $(g', G') \preceq (g, G)$ if and only if (g', G') is obtained from (g, G) by blowing down finitely many edges. K_n is defined to be the geometric realisation of this partially ordered set. It is easy to prove that an admissible graph whose fundamenal group is free of rank n has at most $(3n - 3)$ edges, hence K_n has dimension $(2n - 3)$. An example of a maximal dimensional cell in the case $n = 3$ is shown in Figure 5.2.

Labelled graphs

We wish to represent a given vertex (g, G) of K_n pictorially, in such a way that both the graph G and the marking g are determined by the picture. To do this we choose a maximal tree T in G and a homotopy inverse to g which sends T to the vertex $v(R_0)$. This map sends each oriented edge of $(G-T)$ to a loop in R_0 based at $v(R_0)$, and we label the edge with the corresponding word in $F_n = \pi_1(R_0, v(R_0))$. It is important to notice that the representation of a given (g, G) by a labelled graph is *not* unique. Different choices of maximal tree and homotopy inverse will give rise to different markings; in particular, altering the labels on a given labelled graph by the action of an inner automorpism of F_n produces a labelled graph representing the same marked graph. Figure 5.1 shows four different labelled graphs representing the same point in K_n.

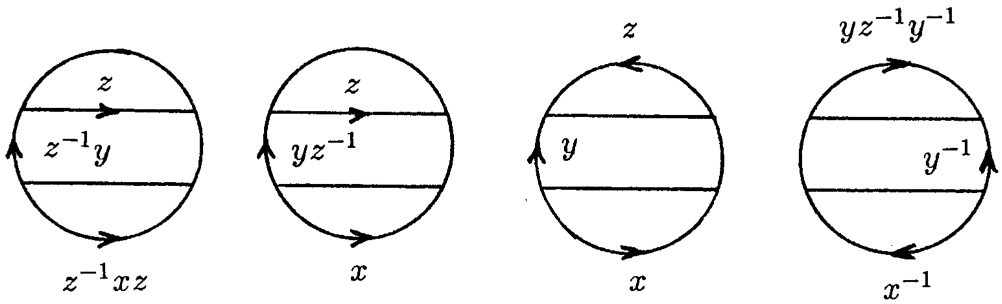

Figure 5.1: Four equivalent labelled graphs

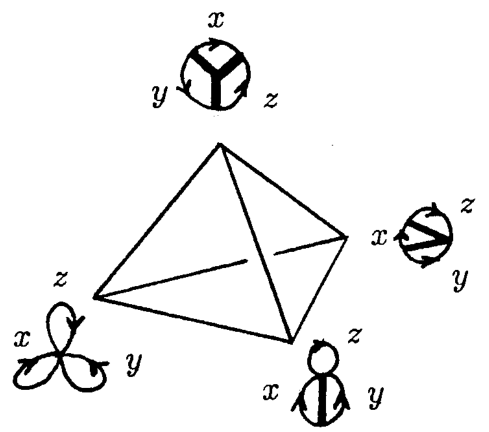

Figure 5.2: A 3–simplex in K_3

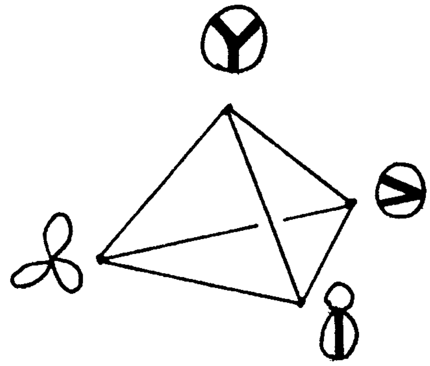

Figure 5.3: A 3–simplex in $K_3/Out(F_3)$

The action of $Out(F_n)$

There is a natural left action of $Out(F_n)$ on the complex K_n, which for our purposes is best described in terms of marked graphs. Given $\phi \in Aut(F_n)$ and a vertex $(g, G) \in K_n$ determined by a labelled graph with edge labels $\{w_1, \ldots, w_n\}$ we define $\phi \centerdot (g, G)$ to be the vertex represented by the same graph with edge labels $\{\phi(w_1), \ldots, \phi(w_n)\}$.

It can be shown that this action is well-defined and preserves the partial ordering $\preceq$, hence it is simplicial. Further, the penultimate sentence of the paragraph on labelled graphs implies that the action of $Inn(F_n)$ is trivial, so we get an induced action of $Out(F_n)$ on K_n.

The $Out(F_n)$–stabiliser of a vertex in K_n is isomorphic to its group of graph automorphisms. This follows from the definition of equivalence for marked graphs, as is proved in [25]. For example, in the case $n = 3$ the isotropy group of a rose is isomorphic to Ω_3, the semi-direct product of $\mathbf{Z}_2 \times \mathbf{Z}_2 \times \mathbf{Z}_2$ with S_3, which occurs naturally as the full group of isometries of a cube in Euclidean 3–space.

We shall represent cells in the quotient complex $K_n/Out(F_n)$ by omitting the edge labels but continuing to draw the edges which are to be blown down in boldface. For example the image of the cell shown in Figure 5.2 will be drawn as in Figure 5.3.

5.2 The Link of a Rose in K_3

In order to carry out the programme of proof outlined in the introduction to this chapter it is essential that we have explicit knowledge of the links of vertices in K_3. The most complicated link which we need to consider is that of a rose. The Gertrude Stein lemma of [13] shows that the links of any two roses are isomorphic via the action of some element of $Out(F_n)$. Moreover, since we are assuming that $Out(F_n)$ acts by isometries, it follows that they are isometric as spherical simplicial complexes. For notational convenience we shall study the link of the standard rose $\rho_0 = (id, R_0)$ and represent all vertices of the link by graphs with labels $\{a, b, c\}$, where $\{a, b, c\}$ is the standard basis for $\pi_1(R_0, v(R_0))$.

Note: In this section we are only concerned with the *combinatorial* structure of the link, not the metric structure.

As we remarked earlier, the stabiliser of a rose is Ω_3 the full group of symmetries of a cube in Euclidean 3–space. Our explicit description of $lk(R_0)$ distinguishes a triangulated 2–sphere embedded in the link. This can be viewed as a cube in a natural way, and the action of $\Omega_3 = stab(\rho_0)$ on $lk(\rho_0)$ is entirely determined by its usual action on this cube. This description of the action of $stab(\rho_0)$ exhibits the symmetries of $lk(\rho_0)$ with sufficient clarity for us to identify conjugate points, i.e. pairs of points in $LK(\rho_0)$ which cannot be joined by a unique geodesic segment. If K_n satisfies the link condition then any path joining a pair of conjugate points must have length at least π. We use this fact to obtain the inequalities given in Section 5.3.

Notation: The ordered triple (a', b', c') denotes the automorphism of F_3 given by $(a, b, c) \mapsto (a', b', c')$, and $[a', b', c']$ denotes the corresponding outer automorphism class.

A vertex of K_3 lies in $st(\rho_0)$ if and only if it can be represented by a labelled graph with labels $\{a, b, c\}$. To list all such vertices one could first list those which have six (the maximal number) of edges and then blow down edges to obtain all the other vertices. Alternatively, one could start at ρ_0 and "blow up" edges (see [13] for details). In any case, the process of assembling and validating such a list is an exhausting one, and for the sake of brevity is omitted. Instead, we simply describe all the 2–cells in $lk(\rho_0)$ (of which there are 408), grouped into convenient subcomplexes, and assemble $lk(\rho_0)$ from these subcomplexes in a suggestive way.

Figure 5.4 shows 32 two-cells which fit together to form a subcomplex of $lk(\rho_0)$ which we call *"the face* $A(+)$*"*. Changing the labels on all the vertex graphs of $A(+)$ by the action of the automorphims $(a^{-1}, b, c), (b, a, c), (b^{-1}, a, c), (c, b, a)$ and (a^{-1}, b, c) respectively, we obtain 5 more "face" subcomplexes $A(-), B(+), B(-), C(+)$ and $C(-)$, each consisiting of 32 two-cells. Making all possible identifications between these faces we obtain the subcomplex shown in Figure 5.5. Here we see the cube referred to earlier, and the action of $stab(\rho_0)$ in permuting the labels $\{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}$ on the graphs corresponds to the usual action of Ω_3 on the cube.

Figure 5.6 shows a further subcomplex consisting of 36 two-cells, which we call *"the hexagon* $H(c)$*"*. The letter c in this notation denotes the fact that this is the label on the edge that is "against the flow" on the central graph. Permuting the labels on all graphs by the automorphisms $(c, b, a), (a, c, b)$ and (a, b, c^{-1}) respectively we obtain "hexagonal" subcomplexes $H(a), H(b)$ and $H(0)$ (in $H(0)$ the arrows associated to the edges labelled a, b, c on the central graph are all confluent).

Notice that there are no edge identifications between the hexagons, and that their boundary edges correspond to all marked graphs in $lk(\rho_0)$ which are of the type shown in Figure 5.7. These edges also occur as cells in the 1–skeleton of the cube shown in Figure 5.5, and making the necessary identifications defines gluing maps from the boundary of each hexagon into the cube. This fits the hexagons into our previous picture beautifully, as shown in Figure 5.8.

The remaining 72 two-cells in $lk(\rho_0)$ can be grouped together so as to form six discs, corresponding to the six possible pairings of distinct hexagons from $\{H(c), H(a), H(b), H(0)\}$. One of these discs, $S(0, c)$, is shown in Figure 5.9. The other five discs are obtained from $S(0, c)$ by permuting the labels on all of the graphs by the action of the automorphisms $(a^{-1}, b, c), (c, b, a), (a, c, b), (a^{-1}, c, b)$, and (c^{-1}, b, a). Figure 5.10 illustrates how these discs fit into the subcomplex of $lk(\rho_0)$ which we have so far constructed. This completes the description of the link of a rose in K_3.

Remark: Our analysis of the link differs from that of Culler and Vogtmann, who described it as "a torus with ten discs attached, having the homotopy type of a wedge of eleven 2–spheres" (see [6] for details). In our construction we obtained $lk(\rho_0)$ (topologically) by taking a 2–sphere and attaching ten 2–discs by injective maps on their boundaries, so the above description of the homotopy type is clear. The torus desribed in [13] is shown in Figure 5.11.

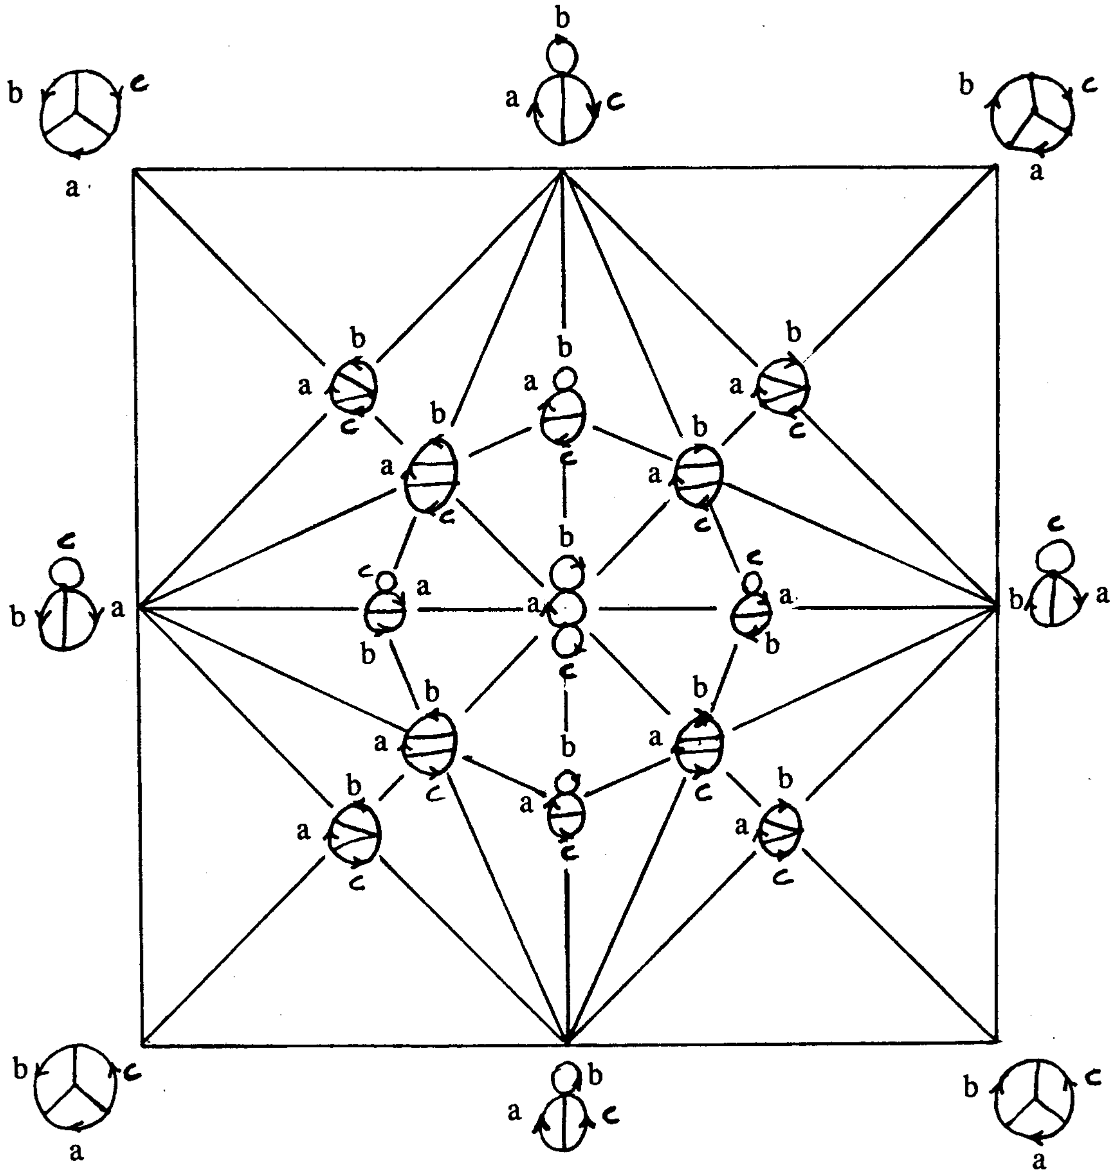

Figure 5.4: The face A(+)

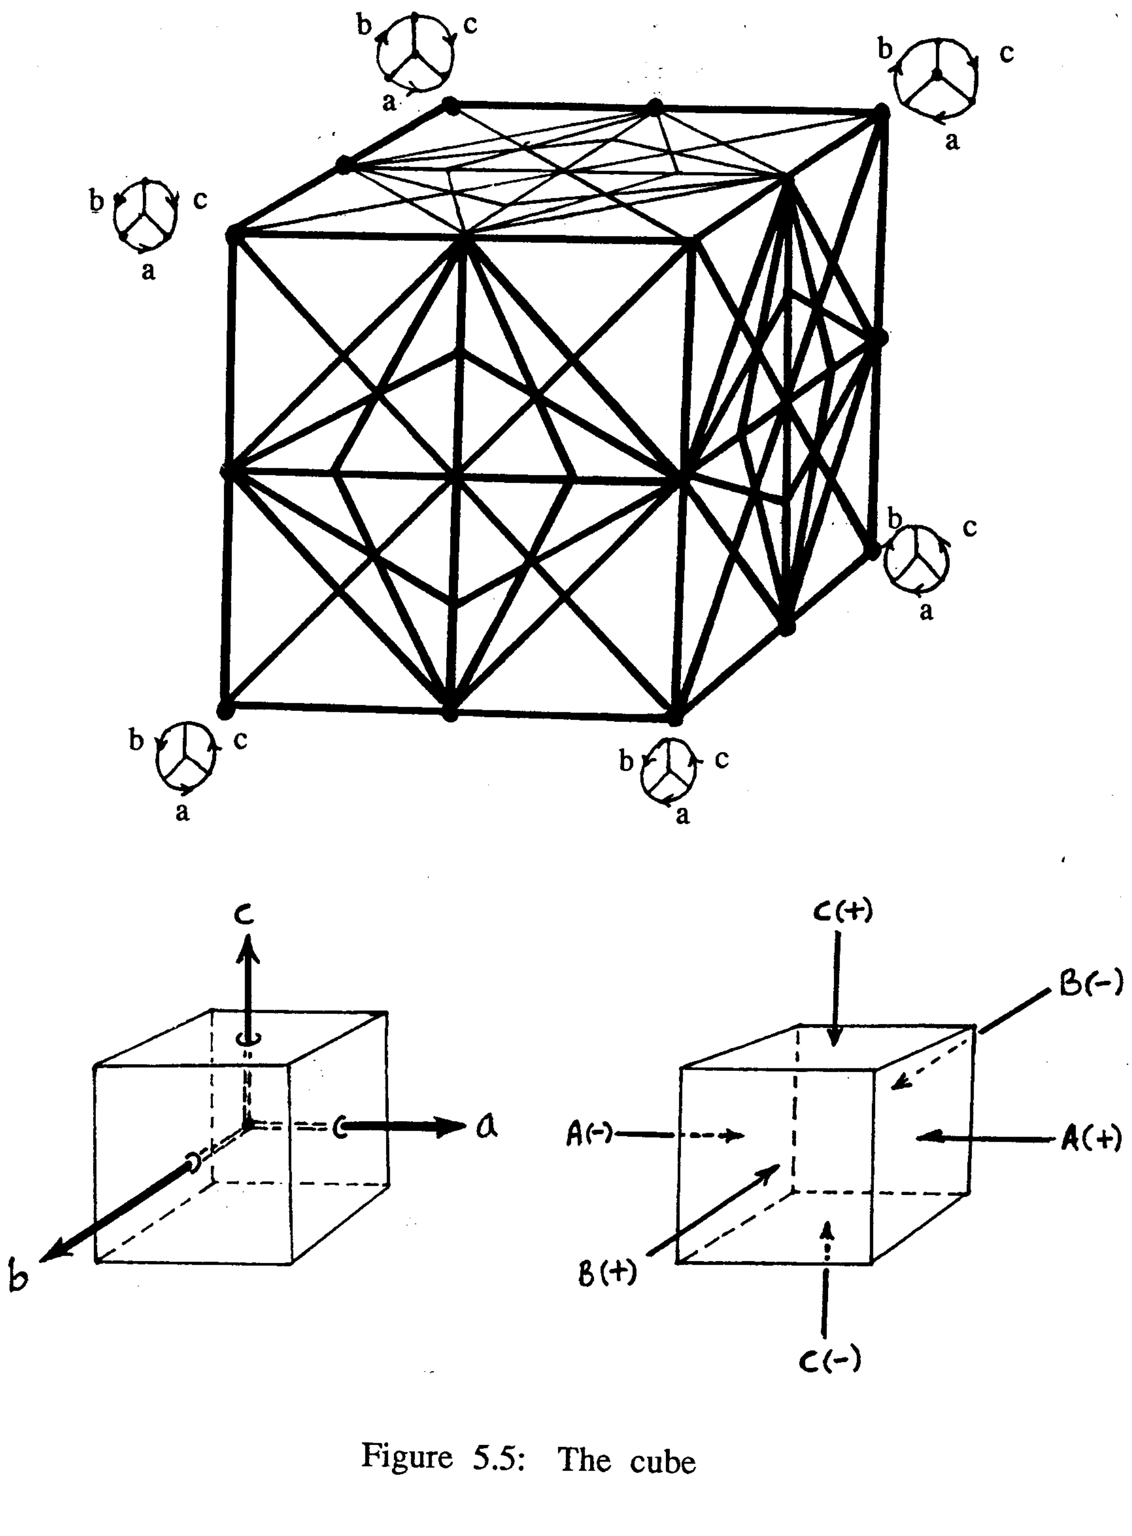

Figure 5.5: The cube

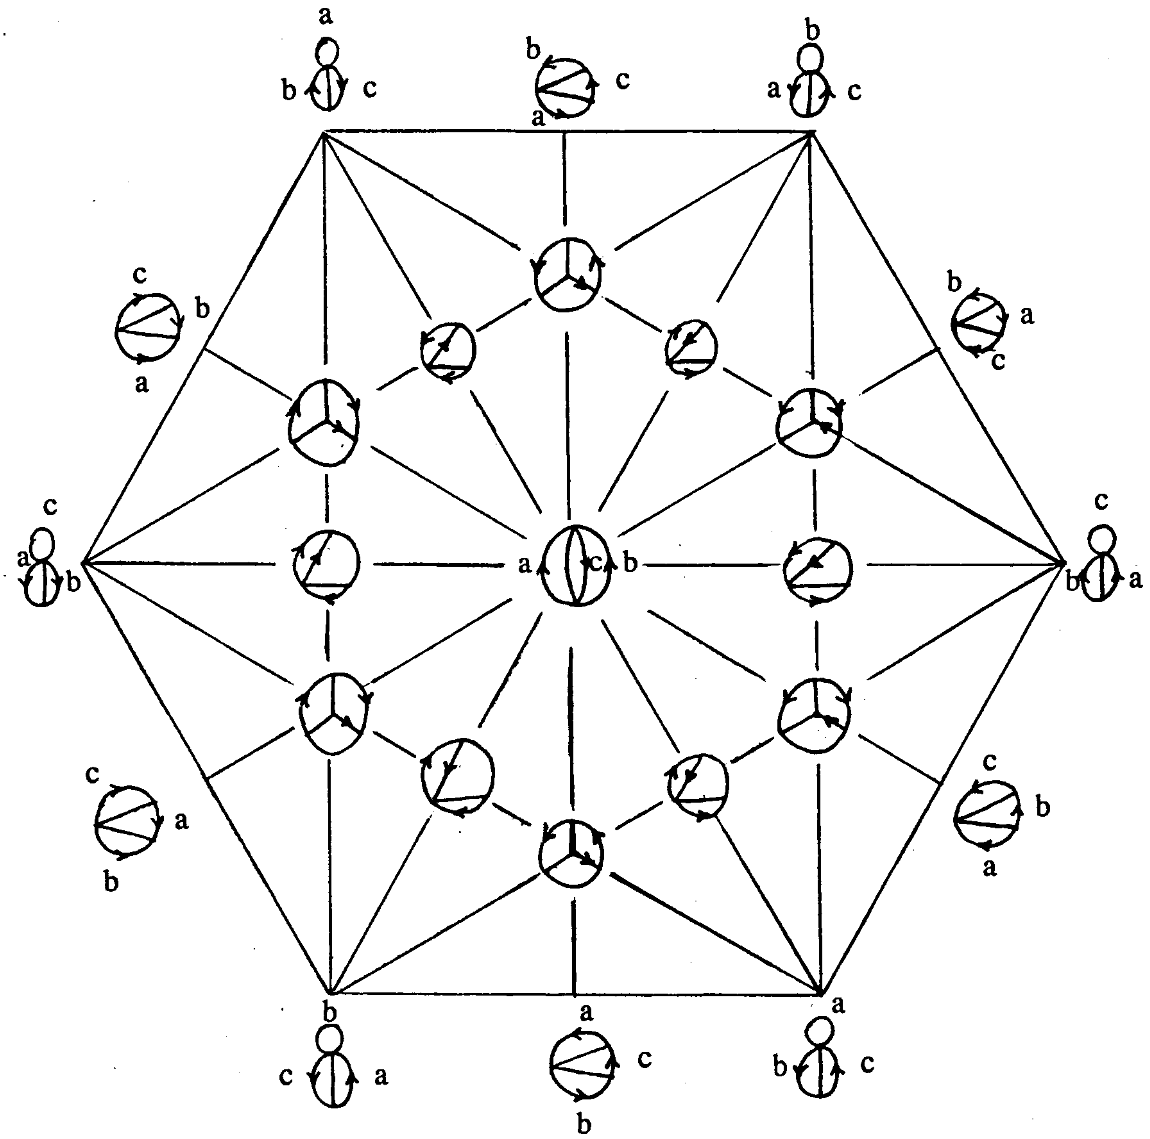

Figure 5.6: The hexagon H(c)

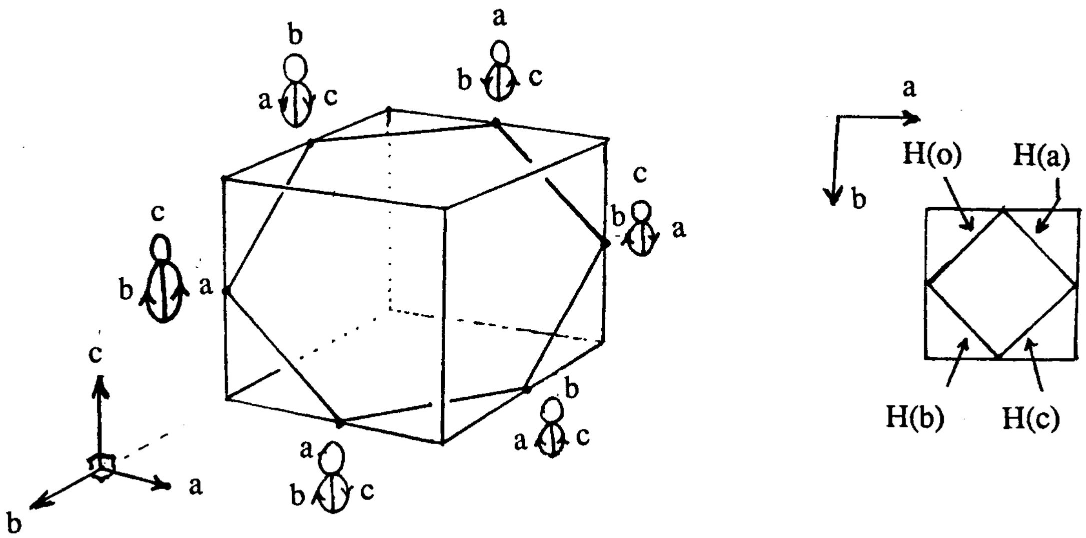

H(0) in the cube

The top face

Figure 5.7: Fitting the hexagons into the cube

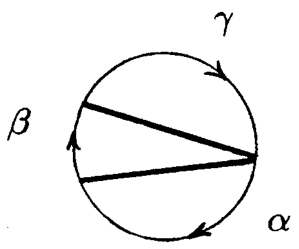

$$\alpha, \beta, \gamma \in \{a^{\pm 1}, b^{\pm 1}, c^{\pm 1}\}$$

Figure 5.8: The graphs in 1–to–1 correspondence with the edges of the hexagons

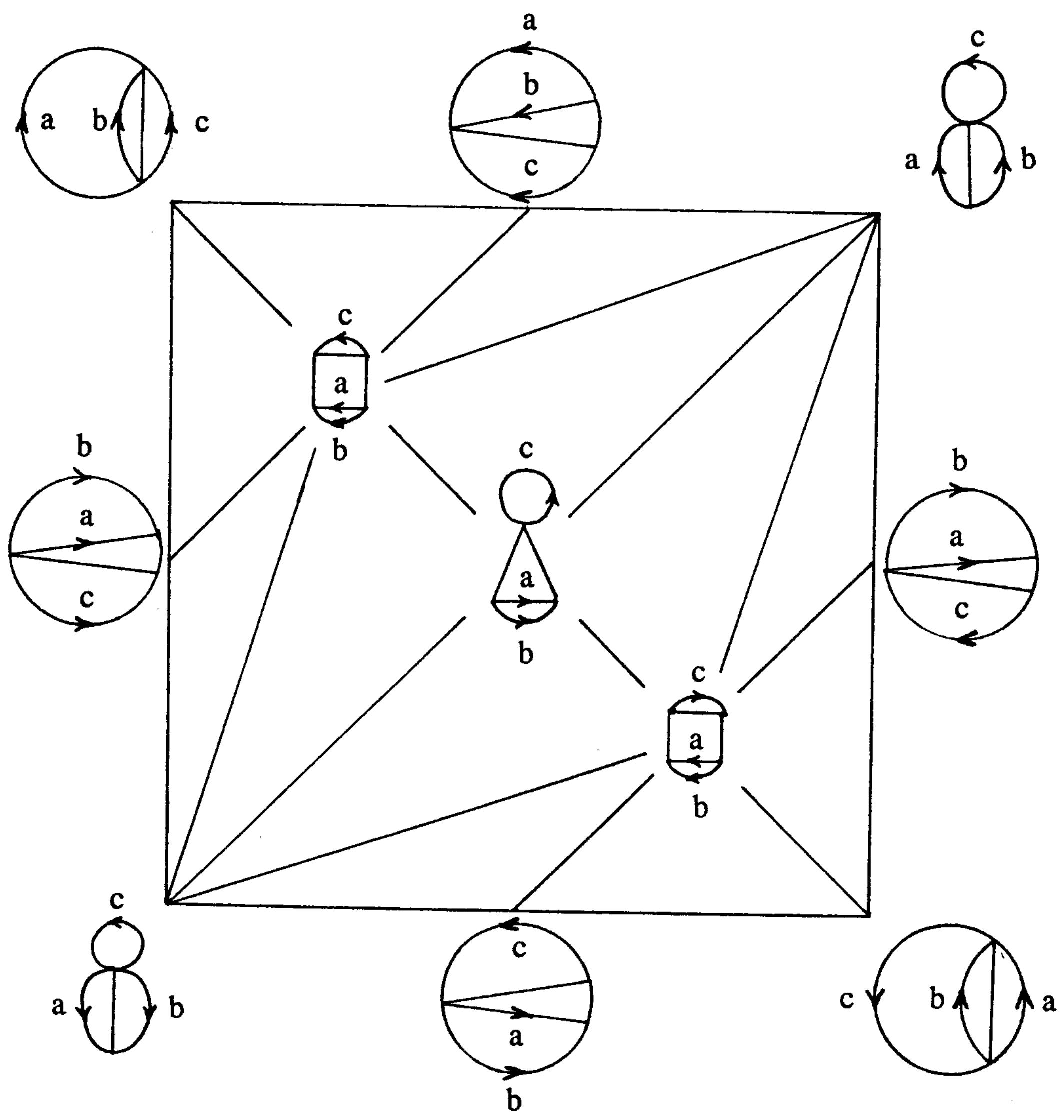

Figure 5.9: S(o,c), twelve of the remaining seventy two 2–cells

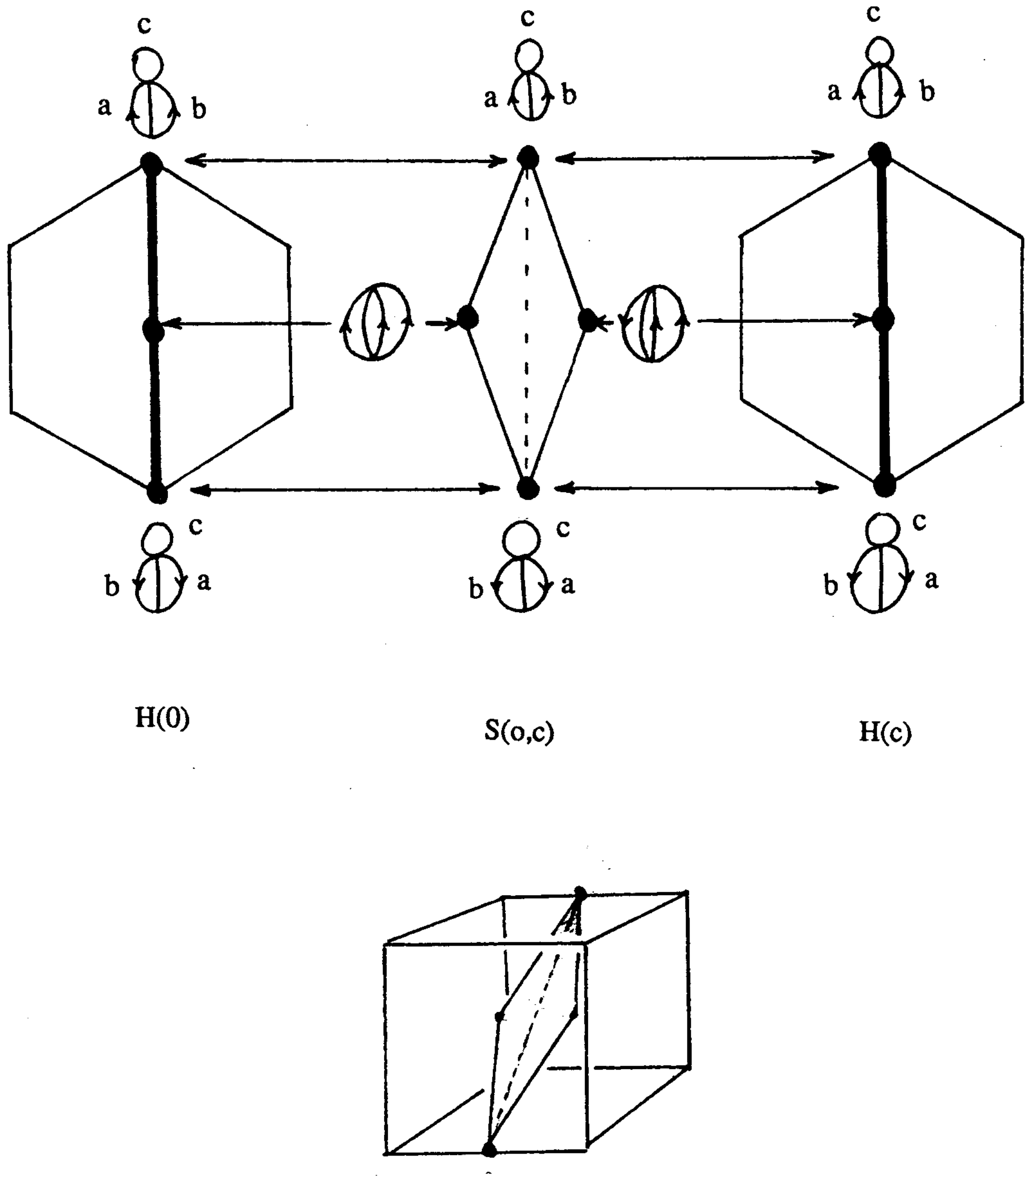

Figure 5.10: Attaching the disc S(o,c)

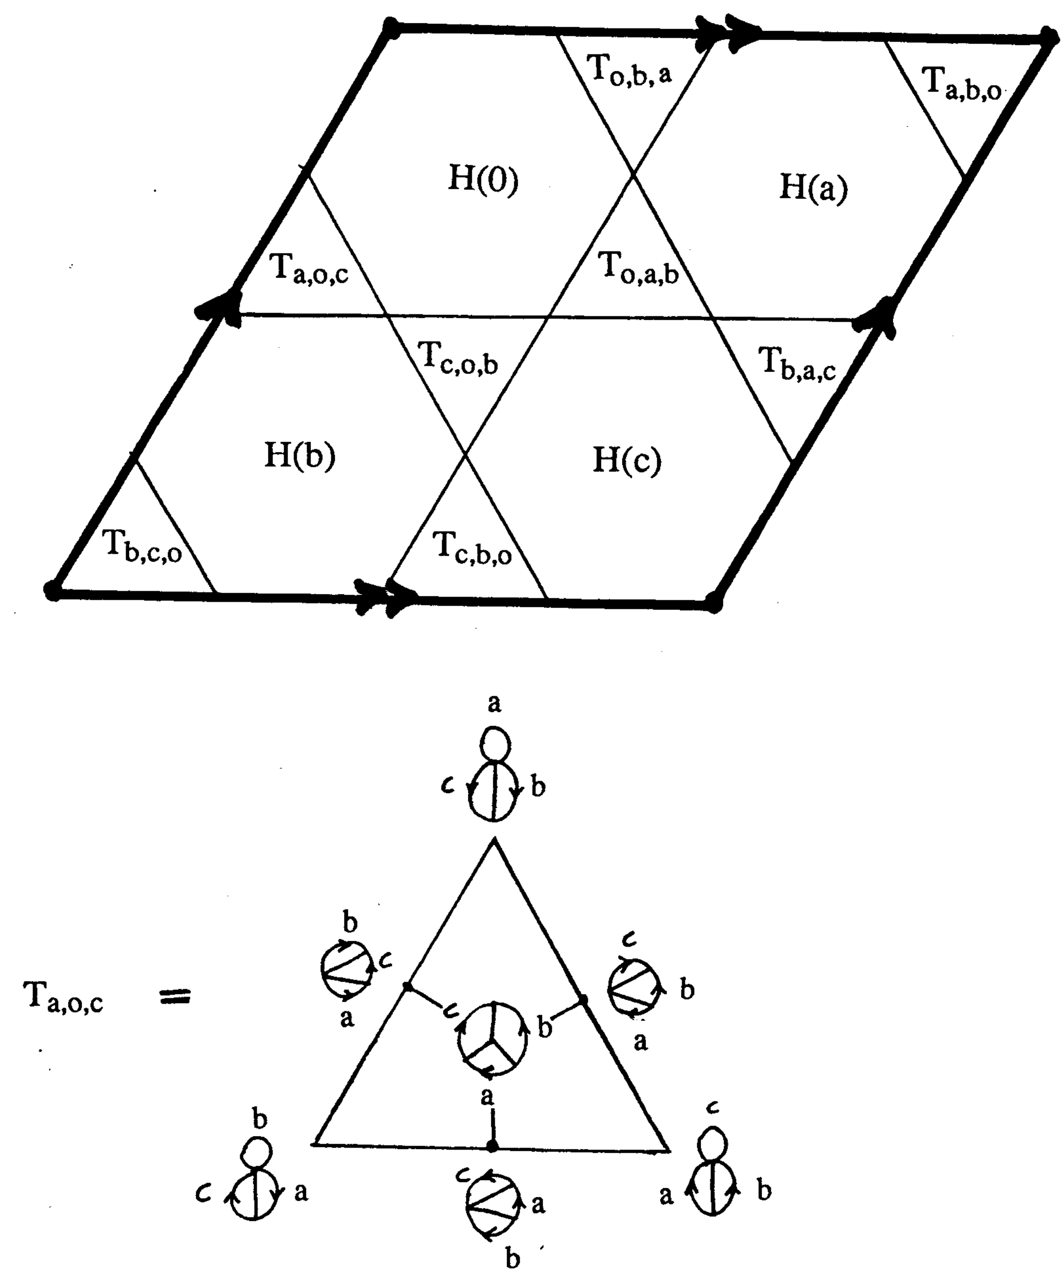

Figure 5.11: The torus

5.3 The Curvature of K_3

Theorem 5.1: *There does not exist an $Out(F_3)$–equivariant piecewise Euclidean (or piecewise hyperbolic) structure of non-positive curvature on K_3.*

The strategy of the proof was described in the introduction to this chapter. *We assume that the cells of $K_3/Out(F_3)$ have been metrized as Euclidean simplices so that the induced piecewise Euclidean structure on K_3 has non-positive curvature.* Then, by studying the links of vertices we obtain contradictory bounds on the angles of the cells in $K_3/Out(F_3)$. Our arguments centre on three 2–cells in $K_3/Out(F_3)$, which are shown in Figure 5.12. We shall prove the following lemmas under the hypothesis that K_3 is non-positively curved (the greek letters denote the angles defined in Figure 5.12).

Lemma 5.2:

$$\frac{\beta}{2} + \gamma + \delta \geq \frac{\pi}{2} .$$

Lemma 5.3:

$$\psi \geq \frac{\pi}{2} .$$

Lemma 5.4:

$$\lambda + \mu + \eta \geq \pi .$$

Lemma 5.5:

$$\theta + \phi \geq \pi .$$

Summing these inequalities we get

$$(\beta + \gamma + \delta) + \psi + (\lambda + \eta + \mu) + (\theta + \phi) \gneqq 3\pi$$

contradicting the fact that these are the angles of three Euclidean triangles. Thus Lemmas 5.2–5.5 together constitute a proof of Theorem 5.1.

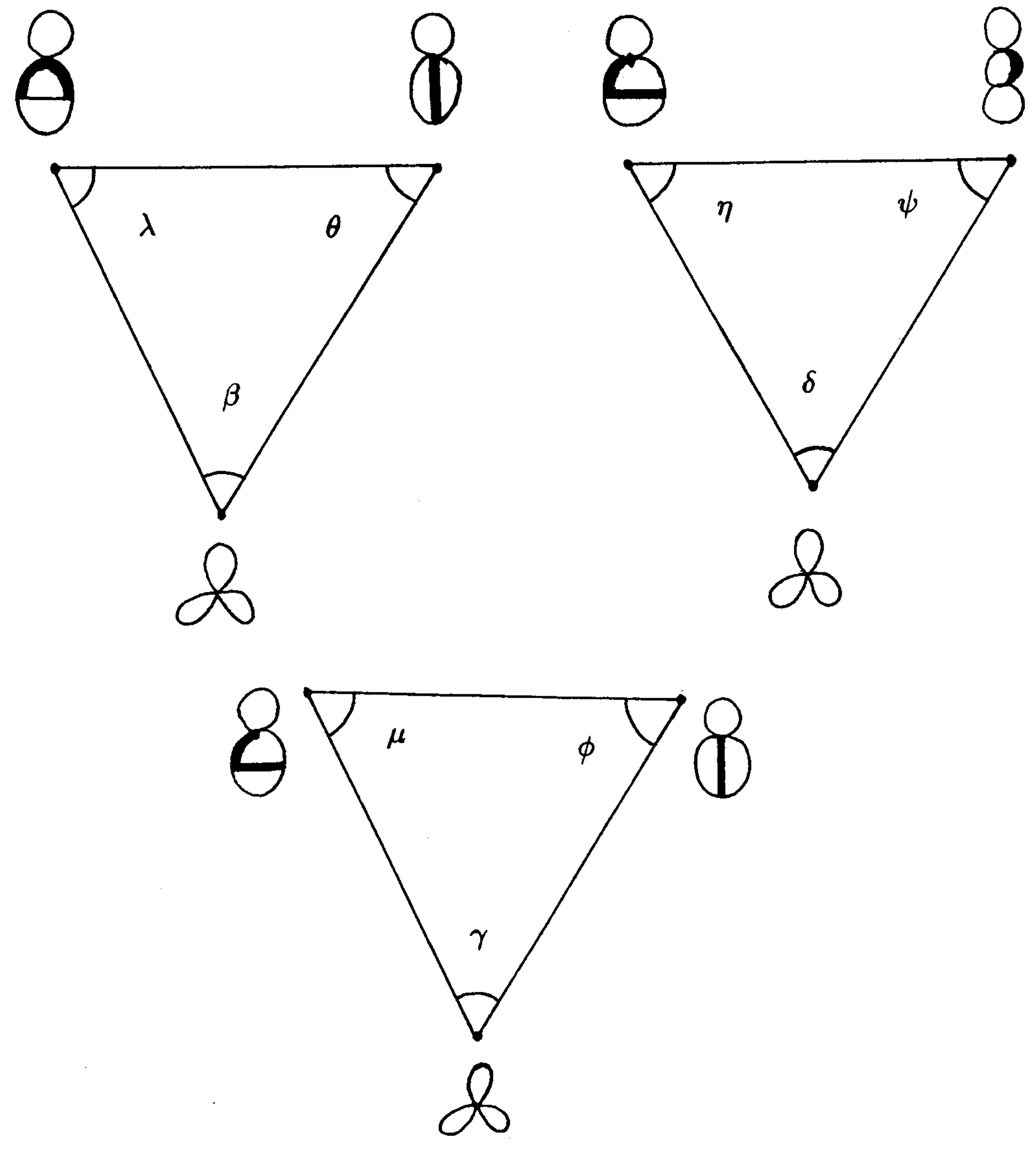

Figure 5.12: Three 2–cells in $K_3/Out(F_3)$

Since all of our arguments involve bounding the length of paths in links of vertices, it is worth recalling that the lengths of the 1–simplices in each of the spherical simplicial complexes $LK((g, G))$ are given by the angles in the 2–simplices of the quotient space $\tilde{K}_3/Out(F_3)$. It is by bounding the lengths of paths in the link complexes $LK((g, G))$ that we obtain bounds on the angles of the 2–simplices in the quotient space.

Before proceeding to the proof of Lemma 5.2 we make a simple (yet crucial) observation about the action of $Out(F_n)$. Adjacent vertices in K_n have underlying graphs which are not homeomorphic. So if an element of $Out(F_n)$ fixes a point in the interior of a simplex in K_n then it fixes the whole simplex pointwise. It follows that if w is a vertex in $lk(v)$, and $\Gamma \subseteq stab(v)$ fixes w but none of the adjacent vertices in $lk(v)$, then Γ acts freely on $lk(w) \cap lk(v)$. In particular there is no Γ–invariant topolgical arc in $lk(v)$ which has an endpoint at w.

Lemma 5.2:

$$\frac{\beta}{2} + \gamma + \delta \geq \frac{\pi}{2}\,.$$

Proof: The proof is based on the structure of $lk(\rho_0)$. Let v_1, v_2 be the vertices of the spherical simplicial complex $LK(\rho_0)$ represented by the graphs shown in Figure 5.13. Each is fixed by the action of both $[a^{-1}, b, c], [a^{-1}, c^{-1}, b^{-1}] \in stab(\rho_0)$. (The action of $[a^{-1}, b, c]$ on the cube shown in Figure 5.5 is by reflection in the (b, c)–plane. The action of $[a^{-1}, c^{-1}, b^{-1}]$ is by rotation through π about the axis (v_1, v_3)). The link of v_2 in $LK(\rho_0)$ has four vertices, represented by the graphs v_4, v_5, v_6, v_7 in Figure 5.13. None of these vertices is fixed by the action of $[a^{-1}, c^{-1}, b^{-1}]$, so neither is any topological arc in $LK(\rho_0)$ with an endpoint at v_2.

If there were a unique geodesic segment from v_1 to v_2 in $LK(\rho_0)$ it would be a topological arc fixed by any automorphism fixing both v_1 and v_2. Therefore such a segment does not exist and, since we are assuming that K_3 satisfies the link axiom, it follows that $d_{LK(\rho_0)}(v_1, v_2) \geq \pi$. Consider P_1, the path shown in Figure 5.14. This joins v_1 to v_2 in $LK(\rho_0)$, and hence must have length at least π. But its actual length is $(\beta + 2\gamma + 2\delta)$. $\square$

Lemma 5.3:

$$\psi \geq \frac{\pi}{2}.$$

Proof: Let (g_0, G_0) be as in Figure 5.15. To prove the lemma we analyse $lk(g_0, G_0)$. (Notice that (g_0, G_0) is the central vertex of the face $B(+)$ of the cube shown in Figure 5.5.)

Given a vertex u in the geometric realisation of any poset, its link can be described as the join of its link in the subcomplex spanned by vertices greater than u (its upper link) and its link in the subcomplex spanned by vertices less than u (its lower link). This makes the calculation of $lk((g_0, G_0))$ a simple matter, since its upper link is a subcomplex of $lk(\rho_0)$, which we have already caclulated, and its lower link consists only of the roses ρ_0 and ρ_1 shown in Figure 5.16.

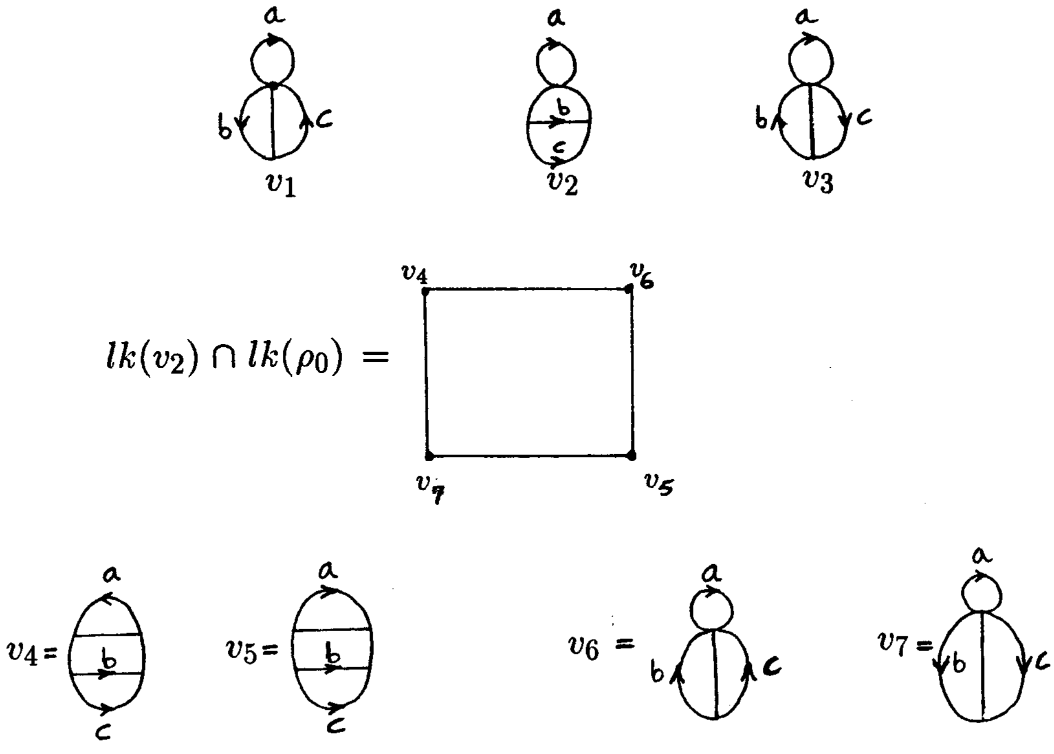

Figure 5.13: The vertices considered in Lemma 5.2

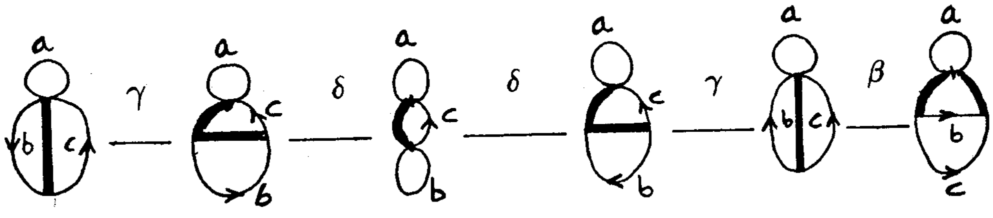

Figure 5.14: The path P_1 of Lemma 5.2

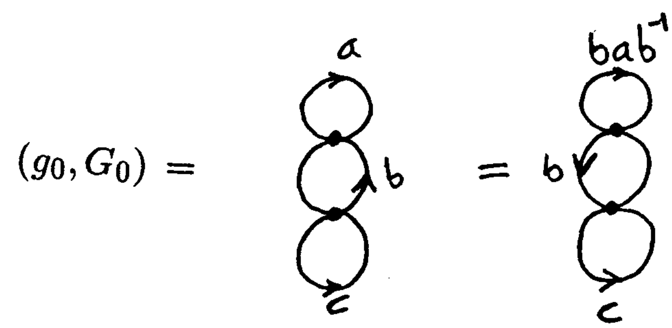

Figure 5.15: The marked graph (g_0,G_0)

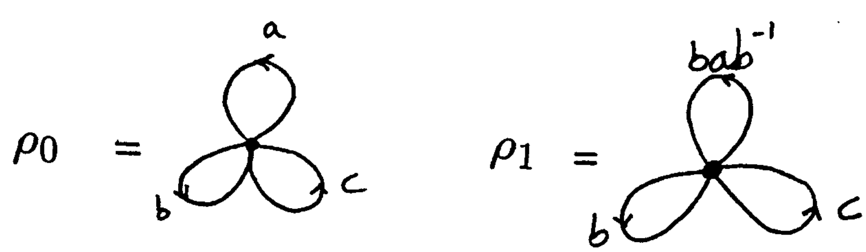

Figure 5.16: The roses in the link of (g_0,G_0)

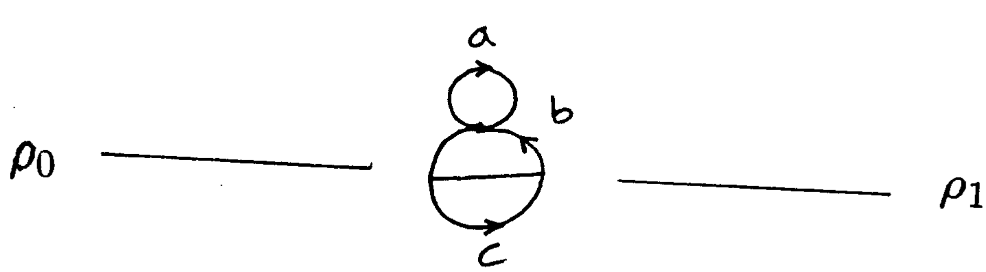

Figure 5.17: The path P_2 of Lemma 5.3

We shall prove that there does not exist a unique geodesic segment from ρ_0 to ρ_1 in $Lk((g_0, G_0))$. The group $stab(\rho_0) \cap stab(\rho_1) \cap stab((g_0, G_0))$ is isomorphic to the dihedral group D_8, which is the stabiliser of the face $B(+)$ of the cube in Figure 5.5 under the action of $\Omega_3 = stab(\rho_0)$. No point in the upper link of (g_0, G_0) $(= lk(\rho_0) \cap lk((g_0, G_0)))$ is fixed by this action. Therefore, as in Lemma 5.2, we deduce that $d_{LK((g_0,G_0))}(\rho_0, \rho_1) \geq \pi$ and hence the path P_2, which is shown in Figure 5.17, is no shorter than π. The actual length of P_2 in $LK((g_0, G_0))$ is 2ψ. □

Remark: The dihedral group referred to in the proof of Lemma 5.3 is:

$$\{[a^{\pm 1}, b, c], [a^{\pm 1}, b, c^{-1}], [c^{\pm 1}, b^{-1}, a], [c^{\pm 1}, b^{-1}, a^{-1}]\}.$$

Lemma 5.4:

$$\lambda + \mu + \eta \geq \pi \, .$$

Proof: For this lemma we need to understand $lk((g_1, G_1))$, which is described in Figure 5.18. We prove that there is not a unique geodesic segment from the vertex ρ_0 to the vertex v' in $LK((g_1, G_1))$.

The automorpism $[a^{-1}, b, c] \in stab((g_1, G_1))$ interchanges the two elements of the upper link of (g_1, G_1) and fixes the lower link pointwise, so if there were a unique geodesic segment from ρ_0 to v' in $LK((g_1, G_1))$ then it would have to be contained in the lower link. But, every path from ρ_0 to v' in the lower link has length at least $\lambda + \mu + \eta$, and there are two paths of precisely this length. Hence the required inequality. □

(g_1, G_1) =

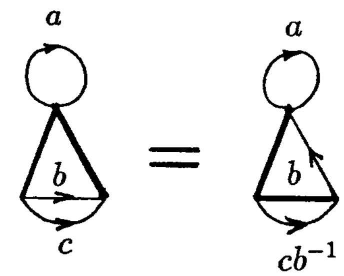

Upper link of (g_1, G_1) =

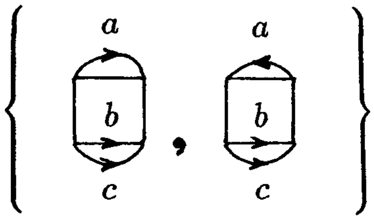

Lower link of (g_1, G_1) :

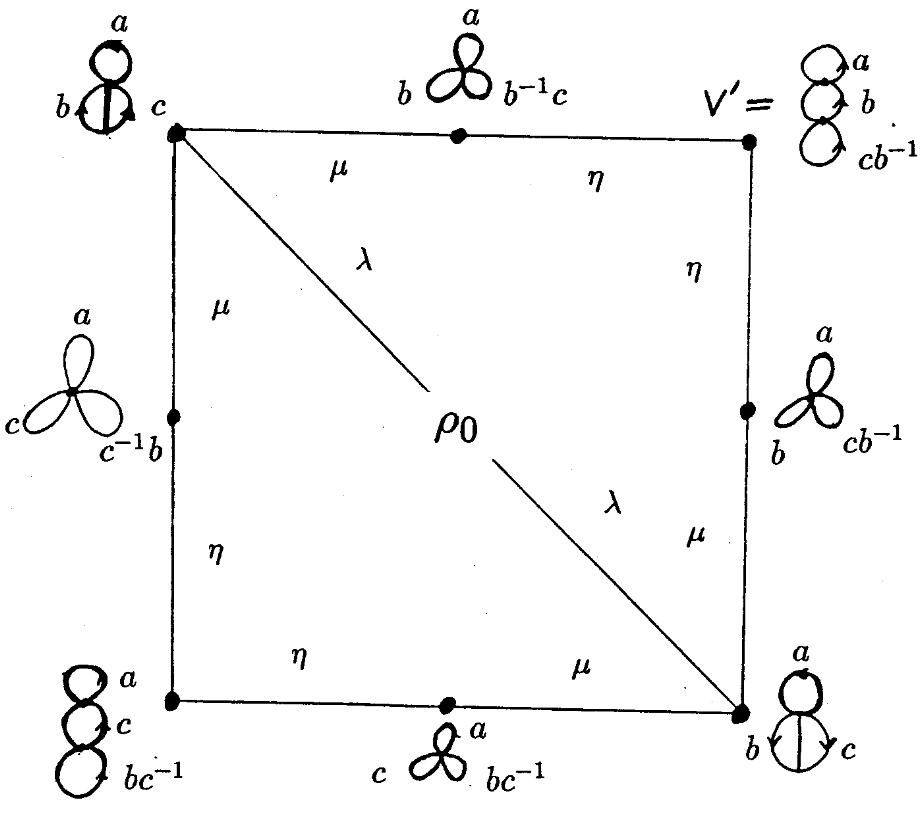

Figure 5.18: The graph (g_1, G_1) and its link

Lemma 5.5:

$$\theta + \phi \geq \pi \ .$$

Proof: Let (g_2, G_2) be as in Figure 5.19. In Section 5.2 we saw this graph as the midpoint of an edge of the cube in Figure 5.6, and calculated its upper link. Its lower link consists of the roses ρ_0, ρ_2, and ρ_3 shown in Figure 5.19. The action of $[a^{-1}, b, c] \in stab((g_2, G_2))$ fixes each of these roses, aswell as three vertices of the upper link of (g_2, G_2). Figure 5.20 shows the 1–dimensional subcomplex of $LK((g_2, G_2))$ fixed by $[a^{-1}, b, c]$. We call this complex L. Each of the edges which is drawn as a solid line has length ϕ and each of the edges which is drawn as a broken line has length θ.

Let p and q be the midpoints of distinct solid edges. If there were a unique geodesic segment from p to q in $LK((g_2, G_2))$ then it would be contained in L. But, all paths from p to q in L have length at least $(\theta + \phi)$ and there are two paths joining them which are of precisely this length. Hence there is not a unique geodesic segment from p to q in $LK((g_2, G_2))$, which imples that $(\theta + \phi) \geq \pi$. $\square$

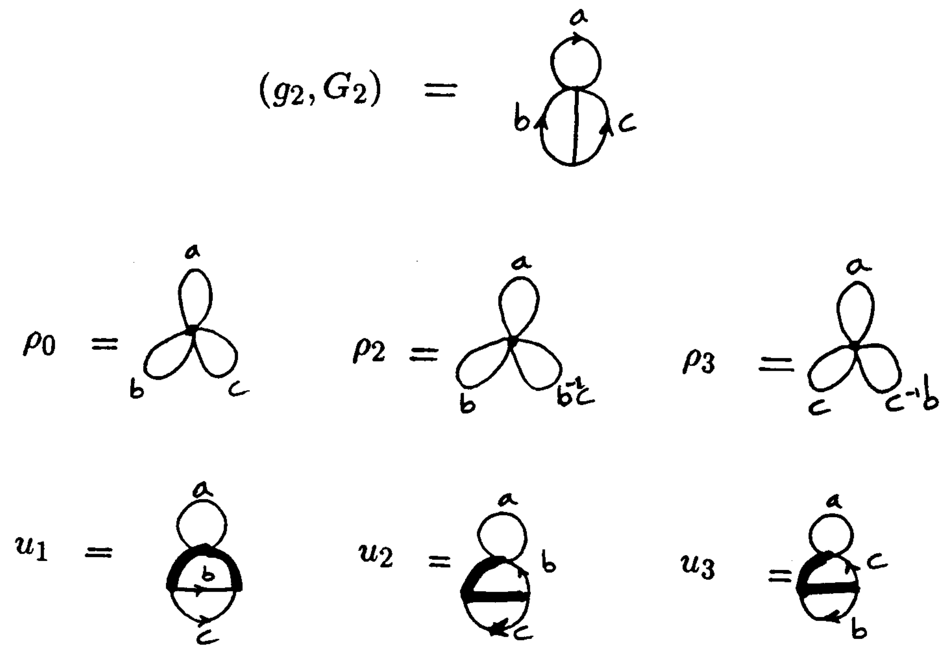

Figure 5.19: The vertices considered in Lemma 5.5

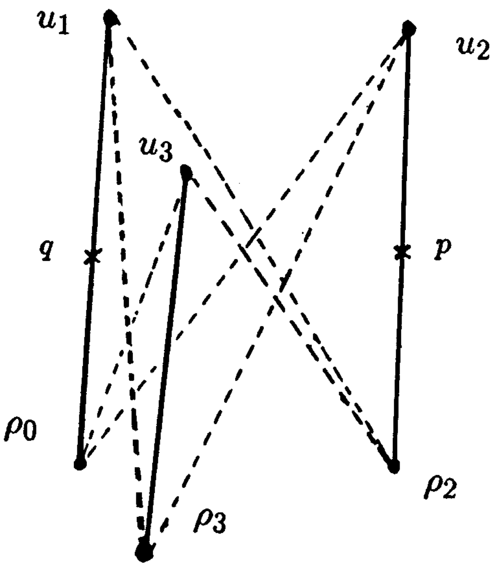

Figure 5.20: The 1–complex L

5.4 The Curvature of K_n

The complexity of the links of vertices in K_3 (particularly that of a rose) makes the prospect of explicitly calculating the links of vertices in K_n for $n > 3$ seem daunting, if not impossible. However, we circumvent this difficulty by restricting our attention to the fixed point set of Γ_n, a particular finite subgroup of $Out(F_n)$. The links of vertices in this set are 2–dimensional, and admit an easy description in terms of the links of vertices in K_3. We use this description to obtain a contradiction in the manner of Section 5.3.

Theorem 5.6: *If $n \geq 3$ then there does not exist an $Out(F_n)$–equivariant piecewise Euclidean (or piecewise hyperbolic) structure of non-positive curvature on K_n.*

Notation: Fix a basis $\{a_1, \ldots, a_{n-2}, b, c\}$ for the free group of rank n.

Let Γ_n denote the subgroup of $Out(F_n)$ generated by the $n-2$ involutions $[a_1^{-1}, a_2, \ldots, a_{n-2}, b, c], \ldots, [a_1, a_2, \ldots, a_{n-2}^{-1}, b, c]$ together with the outer automorphism classes of those automorphisms which fix b and c and act by permutations on the set $\{a_1, \ldots, a_{n-2}\}$. Thus Γ_n is a semidirect product of $\mathbf{Z}_2 \times \ldots \times \mathbf{Z}_2$ ($n-2$ factors) with the symmetric group on $n-2$ letters.

Lemma 5.7 describes the fixed point set of Γ_n. Before stating this lemma we need to make some observations about the geometry of embedded circles in a marked graph and their relationship to cyclic words in the free group.

In the paragraph on labelled graphs we noted that there were certain ambiguities in the choice of labels. However, it is important to notice that the marking g assigns unique labels to reduced circuits in G in the following sense: If we join the endpoints of a labelled edge by the unique topological arc between them in the chosen maximal tree then we obtain an embedded circle in G. This oriented circle corresponds, via the marking g, to a unique conjugacy class of words in F_n (i.e., a cyclic word). Moreover, this conjugacy class, which contains the given edge label, does not depend on the choice of maximal tree. Because the graph under consideration is assumed to have no separating edges it is the union of the circles obtained from the labelled edges by this construction. Further, any two of these circles meet in a (possibly empty) subarc of our chosen maximal tree.

If Φ denotes the isometry of G induced by $\phi \in stab((g, G))$ then the image under Φ of an embedded oriented circle which corresponds to a given cyclic word w is the embedded oriented circle corresponding to the cyclic word $\phi(w)$. In particular, if $\phi(w) = w$ then the embedded circle is mapped to itself by an orientation preserving homeomorphism, and if $\phi(w) = w^{-1}$ then it is mapped to itself by an orientation reversing homeomorphism.

Lemma 5.7: *Suppose that the vertex $(g, G) \in K_n$ can be represented by a labelled graph with the following properties: it has edge labels $\{a_1, \ldots, a_{n-2}, \omega_1, \omega_2\}$ for some words ω_1 and ω_2 involving only the letters b and c, and contains a bouquet of $n-2$ circles with labels $\{a_1, \ldots, a_{n-2}\}$ based at some vertex. Then Γ_n stabilises a vertex $(h, H) \in lk((g, G))$ if and only if (h, H) can be represented by a labelled graph with the same properties.*

Proof: Sufficiency is immediate from the definition of the action of $Out(F_n)$ on K_n. To prove that the given condition on (h, H) is necessary we first observe that *any* vertex in the lower link of (g, G) can be represented by a labelled graph with the desired properties. To see this, consider blowing down an edge of (G, g): The labelling of G by $\{a_1, \ldots, a_{n-2}, \omega_1, \omega_2\}$ arises from choosing a particular maximal tree. If we wish to blow down an edge which is not in this tree, then we must rechoose the tree. Since the new tree must lie in the complement of the loops labelled $\{a_1, \ldots, a_{n-2}\}$ these labels remain unchanged, while the labels ω_1 and ω_2 may be multiplied by words in b and c (see Figure 5.1).

On the other hand, any vertex (h, H) in the upper link of (g, G) can be represented by a labelled graph with edge labels $\{a_1, \ldots, a_{n-2}, \omega_1, \omega_2\}$. Suppose that the automorphism $[a_1^{-1}, a_2, \ldots, a_{n-2}, b, c]$ (which for brevity we call ϕ) stabilises (h, H). If the edge labelled a_1 were not a loop, then the embedded circle obtained by joining its endpoints by the unique arc between them in the chosen maximal tree would intersect the circle determined by some other labelled edge in a non-trivial arc. But this is impossible, because the isometry Φ sends the circle determined by the edge labelled a_1 to itself by an orientation reversing homeomorphism, and sends the circle determined by any other labelled edge to itself by an orientation preserving homeomorphism.

Repeating this argument with a_i in place of a_1 shows that if the graph (h, H) is fixed by the action of Γ_n then the edge labelled a_i is a loop, for all $i \in \{2, \ldots, n-2\}$. The action of the symmetric group S_{n-2}, which occurs in the semi-direct product structure of Γ_n, permutes these loops, hence they must all be based at the same vertex. □

Definition: Suppose $(g, G) \in K_3$ can be represented by a labelled graph with edge labels $\{a, b, c\}$, such that the edge labelled a is a loop. We define $(g, G)(n)$ to be the vertex of K_n represented by the labelled graph obtained from (g, G) by replacing the edge labelled a with a bouquet of circles based at the same vertex, and labelled $a_1, \ldots, a_{n-2}$. (It is easy to check that $(g, G)(n)$ depends only on (g, G) and not on the particular labelling chosen.)

The following is an immediate consequence of Lemma 5.7.

Corollary : *The map* $(h, H) \mapsto (h, H)(n)$ *is a simplicial isomorphism from the fixed point set of* $[a^{-1}, b, c]$ *acting on* $lk((g, G))$ *to the fixed point set of* Γ_n *acting on* $lk((g, G)(n))$.

Proof of Theorem 5.6: The proof centres on the three 2–cells of $K_n/Out(F_n)$ which correspond to the three 2–cells shown in Figure 5.12. More precisely, for each of the 2–cells in Figure 5.12 we choose a 2–cell in K_3 which lies in its preimage and for which each of the vertex graphs has a loop labelled a. We then take the image in K_n of this 2–cell under the map $(H, h) \mapsto (H, h)(n)$, and project this down into $K_n/Out(F_n)$. We retain the names used for the corresponding angles in the rank 3 case.

Suppose that $K_n/Out(F_n)$ can be metrized so that K_n, with the induced structure, is non-positively curved. If this were the case then the sum of the angles shown in Figure 5.12 would be 3π. However, by virtue of the preceding corollary, the arguments given in Section 5.3 apply *(mutatis mutandis)* in the present setting to yield

$$(\beta + \gamma + \delta) + \psi + (\lambda + \eta + \mu) + (\theta + \phi) \gneqq 3\pi.$$

We rephrase Lemmas 5.1–5.4 to clarify this remark.

If $K_n/Out(F_n)$ can be metrized so that K_n, with the induced structure, is non-positively curved then the following assertions hold:

Lemma 5.2 (The case n=3): *Let* v_1 *and* v_2 *be as shown in figure 5.13. These vertices are fixed by both* $[a^{-1}, b, c]$ *and* $[a^{-1}, c^{-1}, b^{-1}]$, *whereas no point of* $lk(v_1) \cap lk(\rho_0)$ *is. Hence*

$$\frac{\beta}{2} + \gamma + \delta \geq \frac{\pi}{2}.$$

Lemma 5.2$'$ (The general case): *Let v_1 and v_2 be as shown in figure 5.13. The vertices $v_1(n)$ and $v_2(n)$ are fixed by the action of $\Gamma_n \times \{1, [a_1, \ldots, a_{n-2}, c^{-1}, b^{-1}]\}$, whereas no point of $lk(v_1(n)) \cap lk(\rho_0(n))$ is. Hence*

$$\frac{\beta}{2} + \gamma + \delta \geq \frac{\pi}{2} \, .$$

Lemma 5.3 (The case n=3): *The only points of $LK((g_0, G_0))$ fixed by both $[a^{-1}, b, c]$ and $[a, b^{-1}, c]$ are the roses ρ_0 and ρ_1 shown in figure 5.16. Hence*

$$\psi \geq \frac{\pi}{2} \, .$$

Lemma 5.3$'$ (The general case): *The only points of $LK((g_0, G_0)(n))$ fixed by the action of $\Gamma_n \times \{1, [a_1, \ldots, a_{n-2}, b^{-1}, c^{-1}]\}$ are the roses $\rho_0(n)$ and $\rho_1(n)$. Where ρ_0 and ρ_1 are as shown in figure 5.16. Hence*

$$\psi \geq \frac{\pi}{2} \, .$$

Lemma 5.4 (The case n=3): *The automorpism $[a^{-1}, b, c] \in stab((g_1, G_1))$ interchanges the two elements of the upper link of (g_1, G_1) and fixes the lower link pointwise. Let v' be as shown in figure 5.18. Every path from ρ_0 to v' in the lower link of (g_1, G_1) has length at least $\lambda + \mu + \eta$, and there are two paths of precisely this length. Hence*

$$\lambda + \mu + \eta \geq \pi \, .$$

Lemma 5.4′ (The general case): *The group* $\Gamma_n \subset stab((g_1, G_1))$ *has no fixed points in the upper link of* $(g_1, G_1)(n)$*, and fixes the lower link pointwise. Let* v' *be as shown in figure 5.18. Every path from* $\rho_0(n)$ *to* $v'(n)$ *in the lower link of* $(g_1, G_1)(n)$ *has length at least* $\lambda + \mu + \eta$*, and there are two paths of precisely this length. Hence*

$$\lambda + \mu + \eta \geq \pi .$$

Lemma 5.5 (The case n=3): *The fixed point set for the action of* $[a^{-1}, b, c]$ *on* $LK((g_2, G_2))$ *is isomorphic to the graph L shown in figure 5.20, with vertices* $\rho_0, \rho_2, \rho_3, u_1, u_2, u_3$ *as defined in figure 5.19. If L is metrized so that the solid edges have length* ϕ *and the broken edges have length* θ*, then any path from p to q in L has length at least* $\theta + \phi$ *and there are two paths of precisely this length. Hence*

$$\theta + \phi \geq \pi .$$

Lemma 5.5′ (The general case): *The fixed point set for the action of* Γ_n *on* $LK((g_2, G_2)(n))$ *is isomorphic to the graph L shown in figure 5.20, with vertices* $\rho_0(n), \rho_2(n), \rho_3(n), u_1(n), u_2(n), u_3(n)$*, where* $\rho_0, \rho_2, \rho_3, u_1, u_2, u_3$ *are as shown in figure 5.19. If L is metrized so that the solid edges have length* ϕ *and the broken edges have length* θ*, then any path from p to q in L has length at least* $\theta + \phi$ *and there are two paths of precisely this length. Hence*

$$\theta + \phi \geq \pi .$$

This concludes the proof of Theorem 5.6. □

REFERENCES

1. I. R. Aitchison and J. H. Rubinstein, *An introduction to polyhedral metrics of non-positive curvature on 3–manifolds*, preprint, University of Melbourne, 1989.

2. A. D. Alexandrov, *A theorem on triangles in a metric space and some applications*, (In Russian), Trudy Mat. Inst. Steks. **38** (1951), 5–23.

3. J. M. Alonso and M. R. Bridson, *Semihyperbolic Groups*, preprint, Cornell University, 1990.

4. R. Alperin, *Homology of* $SL_2(Z[\omega])$, Comm. Math. Helv. **55** (1980), 364–377.

5. W. Ballmann, M. Gromov, and V. Schroeder, *Manifolds of nonpositive curvature,* Progress in mathematics. vol. 61, Birkhäuser, Boston, 1985.

6. T. Brady, *The integral cohomology of* $Out_+(F_3)$, Ph.D Thesis, Cornell University, 1988.

7. M. R. Bridson, *Normal forms in triangles of groups*, (in prepartion).

8. K. S. Brown, *Buildings*, Springer-Verlag, New York, 1989.

9. K. S. Brown, *Five lectures on buildings*, to appear in "Group theory from a geometric viewpoint" (E. Ghys, A. Haefliger, Eds.), Proc. ICTP, Trieste, April 1990.

10. F. Bruhat and J. Tits, *Groupes réductifs sur un corps local I. Données radicielles valuées*, I.H.E.S. Sci. Publ. Math. 41 (1972), 5–251.

11. R. Charney and M. Davis, *Singular metrics of non-positive curvature on branched Riemannian manifolds*, preprint, The Ohio State University, 1990.

12. M. Culler, *Finite groups of outer automorphisms of a free group*, Contributions to group theory, pp. 197–207, Contemp. Math., 33, Amer. Math. Soc., Providence, R.I., 1984.

13. M. Culler and K. Vogtmann, *Moduli of graphs and automorphisms of free groups*, Invent. Math **84** (1986), 91–119.

14. C. H. Dowker, *The topology of metric complexes*, Amer. J. Math. **74** (1952), 555–577.

15. S. M. Gersten, *Convexity properties of non positive curvature*, preprint MSRI, 1989.

16. E. Ghys and P. de la Harpe (Eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov*, Progress in Mathematics, vol. 83, Birkhäuser, Boston 1990.

17. M. Gromov, *Hyperbolic groups*, "Essays in group theory" (S. M. Gersten, Ed.) pp. 75–264, Springer-Verlag, 1987.

18. M. Gromov and W. P. Thurston, *Pinching constants for hyperbolic manifolds*, Invent. Math. **89** (1987), 1–12.

19. A. Haefliger, *Complexes of group and orbihedra*, to appear in "Group theory from a geometric viewpoint" (E. Ghys, A. Haefliger, Eds.), Proc. ICTP, Trieste, April 1990.

20. S. Krstic, K. Vogtmann, (in prepartation).

21. H. Masur, *On a class of geodesics in Teichmüller space*, Ann. of Math. (2) **90** (1969), 1–8.

22. J. Milnor, *Morse theory*, Ann. of Math. Studies, no. 51, Princeton Univ. Press, 1963.

23. G. Moussong, *Hyperbolic coxeter groups*, Ph.D. thesis, The Ohio State University, 1988.

24. J-P. Serre, *Trees*, Springer-Verlag, Berlin, 1980, translation of "Arbres, Amalgames, SL_2", Astérique **46**, 1977.

25. J. Smillie and K. Vogtmann, *A generating function for the Euler charcteristic of Out(F_n)*, J. Pure Appl. Algebra **44** (1987), 329–348.

26. J. R. Stallings, *Non-positively curved triangles of groups*, to appear in "Group theory from a geometric viewpoint" (E. Ghys, A. Haefliger, Eds.), Proc. ICTP, Trieste, April 1990.

27. D. A. Stone, *Geodesics in piecewise linear manifolds*, Trans. Amer. Math. Soc. **215** (1976), 1–44.

28. J. Tits, *On buildings and their applications to group theory*, in "Proceedings of the International Congress of Mathematics", Vancouver 1974, vol. 1, Canad. Math. Congress, Montreal, 1975, pp. 209–220.

29. J. Tits, *Immeubles de type affine*, in "Buildings and the Geometry of Diagrams" (L. A. Rosati, Ed.), Proc. CIME Como 1984, Lecture Notes in Mathematics 1181, Springer-Verlag, Berlin, 1986, pp. 159–190.