SINGULAR METRICS OF NONPOSITIVE CURVATURE ON BRANCHED COVERS OF RIEMANNIAN MANIFOLDS

By Ruth Charney and Michael Davis

Introduction. According to the well-known Comparison Theorems of Alexandrov and Topogonov, the statement that the sectional curvature of a smooth Riemannian manifold M is bounded above by a real number χ is equivalent to a statement concerning small geodesic triangles in M. One such statement, the so-called "$CAT(\chi)$-inequality", compares distances between points in such a triangle with the corresponding distances in a "comparison triangle" in the complete, simply connected, 2-manifold of constant curvature χ. The advantage of this characterization of bounded curvature is that it makes sense in a context much more general than that of Riemannian metrics; one needs only the concept of a "geodesic space", that is, a metric space in which any two points can be connected by a geodesic segment. Such a geodesic space then has "curvature $\leq \chi$" if, locally, geodesic triangles satisfy the $CAT(\chi)$-inequality. Thus begins one thread in the spectacular web of ideas introduced by Gromov in his 1987 paper [G]. The goal of the present paper is to further develop this thread.

The curvature bound of most interest to us is $\chi = 0$. The fundamental result about complete nonpositively curved Riemannian manifolds is the Cartan-Hadamard Theorem. It implies, in particular, that such a manifold is aspherical. Nonpositive curvature also has interesting group theoretic implications for the fundamental group of a closed manifold, e.g., the group is combable, has solvable word problem, satisfies a quadratic isoperimetric inequality (cf. [AB]), and every solvable subgroup is virtually abelian (a theorem of Lawson-Yau and Gromoll-Wolf, cf. [CE]). Moreover, if the curvature upper bound χ satisfies $\chi < 0$, then the group is hyperbolic in the sense of Gromov [G] and, as such, satisfies a wealth of other group theoretic properties. The more general notion of nonpositive (or negative) curvature for a geodesic metric space (defined via the $CAT(\chi)$-inequality) is sufficiently strong for all of the above results to hold (see [G], [GH] and [Br]).

The current paper was motivated by the following question. Suppose that M is a smooth Riemannian manifold and that $\pi : \tilde{M} \to M$ is a branched covering of M by a manifold $\tilde{M}$. The metric on M can be pulled back, via π, to a (singular)

Manuscript received December 7, 1990; revised May 18, 1992.
Both authors are partially supported by NSF Grant #DMS 8905378.
American Journal of Mathematics 115 (1993), 929–1009.

metric on $\tilde{M}$. When does this induced metric on $\tilde{M}$ have curvature bounded from above? (Or, more generally, when does an orbifold structure on M have curvature bounded from above?) Gromov [G] observes that if the sectional curvature of M is $\leq \chi$ and if each component of the branch set is a totally geodesic submanifold of codimension two in M, then the curvature of $\tilde{M}$ is $\leq \chi$. We wish to understand what happens when the local structure of the branch set is more complicated. (Some observations in this direction are made in [G].)

It is easy to show that two conditions are necessary for $\tilde{M}$ to have curvature $\leq \chi$:

(1) the sectional curvature of M is $\leq \chi$,

(2) locally, the closure of each component of each stratum of the branch set is a convex subset of M.

For each $x \in M$ there is an induced finite sheeted branched covering $\tilde{S}_x \to S_x$, where S_x denotes the unit sphere in the tangent space T_xM. In view of (2), $\tilde{S}_x$ is a piecewise spherical polyhedron. It turns out (cf., Theorem 5.3) that together with (1) and (2) the following condition is necessary and sufficient for $\tilde{M}$ to have curvature bounded from above:

(3) all triangles in $\tilde{S}_x$ satisfy the $CAT(1)$-inequality

Moreover, it follows from a result of Gromov (cf., Theorem 3.1) that a piecewise spherical polyhedron, such as $\tilde{S}_x$, satisfies $CAT(1)$ if and only if it has curvature ≤ 1 (i.e., it *locally* satisfies $CAT(1)$) and every closed geodesic has length at least 2π. Thus, the question of when a branched cover of a smooth Riemannian manifold is nonpositively curved ultimately comes down to a problem in "polyhedral geometry".

In this paper we focus on the following two questions:

(I) Which local structures admit metrics of nonpositive curvature?

(II) What local geometric conditions on the branch set are necessary for condition (3) to hold?

The answer to Question II is particularly interesting. The picture which emerges from our analysis is that the situations in which condition (3) hold are very rigid. For example, suppose the dimension of M is three and for some $x \in M$, $\tilde{S}_x \to S_x$ looks like the quotient map of a 2-sphere by a group G generated by three rotations of orders p, q, and r. (G is an index two subgroup of a reflection group acting on $\mathbb{R}^3$.) The metric d on $S_x(= S^2)$ is the standard spherical metric and the branch set in S_x consists of three points x_p, x_q, and x_r. In this case we discover the following.

THEOREM 8.1. *Let $\tilde{S}_x$ be as above. Then the $CAT(1)$-inequality holds in $\tilde{S}_x$ if and only if*

(i) the singular points x_p, x_q, x_r lie on a great circle in S_x but are not contained in any semicircle, and

(ii) $d(x_i, x_j) \geq \frac{\pi}{k}$ for all $\{i,j,k\} = \{p,q,r\}$.

In other words, near x, the branch locus in M must consist of three coplanar rays meeting at x (like "tinkertoys"), with angles between the rays bounded below by condition (ii).

Similarly, suppose M is a 4-manifold and $\tilde{S}_x \to S_x$ looks like the quotient map of a 3-sphere by the complex reflection group $G(p,q,r)$ generated by three "complex reflections" of orders p, q, and r (as defined in Section 6). Then the branch set in $S_x(=S^3)$ consists of three disjoint great circles P_p, P_q, and P_r and we find the following.

THEOREM 9.1. *Let $\tilde{S}_x$ be as above, then the $CAT(1)$-inequality holds in $\tilde{S}_x$ if and only if P_p, P_q, and P_r are fibers of the Hopf fibration $H : S^3 \to S^2$ and their images x_p, x_q, x_r in S^2 satisfy (i) and (ii) of Theorem 8.1.*

The outline of the paper is as follows. In Chapter I (Sections 1-7) we develop the necessary background material on bounded curvature, polyhedra of piecewise constant curvature, orbifolds, etc, and then (in Section 6) we address question I. Chapter I ends with an application to branched covers of flat tori. Chapter II (Sections 8-12) is devoted to answering Question II. It contains, among others, the two results mentioned just above. The paper ends with an appendix discussing joins of piecewise spherical complexes. In particular, it is shown that the join $X * Y$ of two such complexes X and Y satisfies $CAT(1)$ if and only if both X and Y satisfy $CAT(1)$.

Contents

Chapter I: Orbifolds of Nonpositive Curvature. In this chapter we review some ideas of Gromov. Basic references are [GH] and [G]. In particular, a good bit of Sections 1 and 3 comes from Ballman's article in [GH, Chap. 10], much of Section 4 comes from Haefliger's article [GH, Chap. 11], and Section 5 uses the expository article of Troyanov [GH, Chap. 3].

1. The $CAT(\chi)$-Inequality: Curvature Bounded from Above. Let X be a metric space. A path $c : [a,b] \to X$ is *locally isometric* if for each $t \in (a,b)$ there is a subinterval J containing t such that $c \mid J$ is an isometric map. A *geodesic path* is a constant speed reparametrization of a locally isometric path. Such a geodesic path $c : [a,b] \to$ is *minimal* if c is an isometric map. A *geodesic segment* is the image of a minimal geodesic path together with an orientation. A *closed geodesic* in X is the image of an isometric map from a circle into X. A *triangle* in X consists of three points and geodesic segments connecting them. The metric space X is a *geodesic space* if any two points in X can be joined by a geodesic segment.

For each real number χ and integer $n \geq 2$, let M^n_χ denote the complete, simply connected, Riemannian n-manifold of constant sectional curvature χ. (If $\chi > 0$, then M^n_χ is a sphere of radius $1/\sqrt{\chi}$; if $\chi = 0$, M^n_χ is euclidean n-space; if $\chi < 0$, then M^n_χ is a hyperbolic n-space.)

Let T be a triangle in X. A *comparison triangle for T in M^2_χ* is a triangle $\overline{T}$ in M^2_χ with the same edge lengths as T. It is clear that $\overline{T}$ is unique up to an isometry of M^2_χ and that such a $\overline{T}$ exists if $\chi \leq 0$, or if $\chi > 0$ and the perimeter of T is $\leq 2\pi/\sqrt{\chi}$. (Convention: the condition "the perimeter of T is $\leq 2\pi/\sqrt{\chi}$" is vacuous if $\chi \leq 0$.) There is a unique map $T \to \overline{T}$ which takes each edge of T isometrically onto the corresponding edge of $\overline{T}$. For each $x \in T$ let $\overline{x}$ denote its image in $\overline{T}$ under this map.

A triangle T in X of perimeter $\leq 2\pi/\sqrt{\chi}$ *satisfies the $CAT(\chi)$-inequality* if for any vertex $v \in T$ and point p on the side opposite to v we have:

$$d(v,p) \leq d_\chi(\overline{v},\overline{p}),$$

where $\overline{v}$ and $\overline{p}$ are the corresponding points of a comparison triangle in M^2_χ. (Here d denotes distance in X, while d_χ is distance in M^2_χ.) If $\chi > 0$, then any triangle of perimeter $> 2\pi/\sqrt{\chi}$ is also said to satisfy the $CAT(\chi)$-inequality. A geodesic metric space X is said to *satisfy* $CAT(\chi)$ if the $CAT(\chi)$-inequality holds for every triangle in X. The space X has *curvature* $\leq \chi$, if it satisfies $CAT(\chi)$ locally.

The question of whether or not the $CAT(\chi)$-inequality holds in a geodesic space X is closely connected to the question of uniqueness of geodesic segments.

Indeed, suppose we have two points in X of distance $< \pi/\sqrt{\chi}$ and two distinct geodesic segments connecting them. Such a configuration is called a *digon.* By introducing a third vertex in the interior of an edge of such a digon, one obtains a triangle for which the $CAT(\chi)$-inequality fails. Similarly, if γ is a closed geodesic in X of length $< 2\pi/\sqrt{\chi}$ then by picking any three points of γ as vertices (so that the triangle inequalities hold) we again obtain a triangle for which the $CAT(\chi)$-inequality fails. Thus, if a geodesic space X satisfies $CAT(\chi)$, then any two points of distance $< \pi/\sqrt{\chi}$ are connected by a unique geodesic segment and there are no closed geodesics of length $< 2\pi/\sqrt{\chi}$.

For the remainder of this section we shall assume that X is a locally compact, complete, geodesic metric space.

If, in addition, X has curvature $\leq \chi$, then for any $p \in X$ there is an $\varepsilon > 0$ such that the ε-ball centered at p is convex (cf., Ballman's article in [GH, Chap. 10, Remark 2, p. 182]). In particular, X is locally contractible and hence, has a universal cover. When $\chi \leq 0$, one has the following generalization of the Cartan-Hadamard Theorem.

THEOREM 1.1. (Cartan-Hadamard-Alexandrov-Gromov-Ballman). *If X has curvature ≤ 0, then its universal cover is contractible.*

This result is stated in [G, p. 119] where it is attributed to Cartan, Hadamard and Alexandrov. It is proved in [GH, Chap. 10, Theorem 4, p. 187].

In [GH, Chap. 10, Theorem 7] it is proved that if X has curvature $\leq \chi$ and any two points of X of distance $< \pi/\sqrt{\chi}$ are connected by a unique geodesic segment, then $CAT(\chi)$ holds for X. Ballman's proof actually gives the following result.

LEMMA 1.2. (Ballman [GH, Chap. 10]) *Suppose that X has curvature $\leq \chi$. Let T be a triangle in X such that any two points of T of distance $< \pi/\chi$ are connected by a unique geodesic. Then the $CAT(\chi)$-inequality holds for T.*

We define three numbers. The *systole* of X, denoted by sys(X), is the greatest lower bound of the lengths of closed geodesics in X. The *injectivity radius* of X, denoted by $ir(X)$, is the greatest lower bound of $d(p,q)$, where (p,q) is a pair of points for which the geodesic segment joining p to q is not unique. In other words, $ir(X)$ is one-half the greatest lower bound of the perimeters of all digons in X. Finally, $C_\chi(X)$ is defined to be the greatest lower bound of the perimeters of triangles T in X such that $CAT(\chi)$-inequality fails for T.

LEMMA 1.3. *Suppose that $\chi > 0$ and that X is compact of curvature $\leq \chi$. If $CAT(\chi)$ fails for X, then $sys(X) = 2(ir(X)) = C_\chi(X)$. Hence, $sys(X) \geq 2\pi/\sqrt{\chi}$ if and only if $CAT(\chi)$ holds for X.*

Proof. The $CAT^*(\chi)$-inequality for a triangle T in X asserts that for any two points p and q in T, $d(p,q)$ is $\leq$ the distance between corresponding points in the comparison triangle. It is easy to see that if the $CAT(\chi)$-inequality holds locally, then so does the $CAT^*(\chi)$-inequality. Ballman proves a "Gluing Lemma" for $CAT^*(\chi)$ ([GH, Chap. 10, Lemma 4, pp. 182–183]). Suppose that T is a triangle in X of perimeter $< 2\pi/\sqrt{\chi}$ and that we "chop" T into two new triangles T_1 and T_2 by introducing a geodesic segment connecting a vertex of T to a point on the opposite edge. The Gluing Lemma asserts that if both T_1 and T_2 have perimeter $< 2\pi/\sqrt{\chi}$ and satisfy the $CAT^*(\chi)$-inequality then so does T.

Let C_1 denote the systole of X, let C_2 be twice its injectivity radius, and let $C_3 = C_\chi(X)$.

Claim 1. Either $C_2 \geq 2\pi/\sqrt{\chi}$ or else there is a digon of minimum perimeter C_2.

To see this, suppose that $C_2 < 2\pi/\sqrt{\chi}$. Since X is compact, we can find a sequence of digons the edges of which converge uniformly to geodesic segments of length $\frac{1}{2}C_2$ between points p and q in X. *A priori* these two segments might coincide. Since the $CAT^*(\chi)$-inequality holds locally, the two edges of each digon in our sequence cannot remain close together over their entire length; for otherwise, we could chop such a digon into small triangles and then use the Gluing Lemma to contradict the fact that the $CAT(\chi)$-inequality fails for any triangle obtained by introducing a third vertex on a digon of perimeter $< 2\pi/\sqrt{\chi}$. Hence, our sequence of digons converges to an actual digon of minimum perimeter.

For the remainder of the proof we suppose that the $CAT(\chi)$-inequality fails for some triangle in X, i.e., $C_3 < 2\pi/\sqrt{\chi}$.

Claim 2. $C_1 \geq C_2 \geq C_3$.

Since we can subdivide a digon into a triangle by introducing a third vertex into the interior of an edge, we have that $C_2 \geq C_3$. A pair of points on a closed geodesic of length ℓ are called *opposite* if their distance is $\frac{1}{2}\ell$. Similarly, if D is a digon with vertices p, q, then two points x and y on different edges of D are *opposite* if $d(p,x) + d(p,y) = d(p,q)$. Since a closed geodesic can be subdivided into a digon by introducing a pair of opposite vertices, we have that $C_1 \geq C_2$.

Claim 3. $C_2 = C_3$.

To see this, suppose that T is a triangle of perimeter $< C_2$. Since the injectivity radius of X is $\frac{1}{2}C_2$, any two points of T are connected by a unique geodesic segment; hence, by Lemma 1.2, the $CAT(\chi)$-inequality holds for T. Thus, $C_3 = C_2$.

Claim 4. $C_1 = C_2$.

Since the $CAT(\chi)$-inequality holds for all triangles of perimeter $< C_2$, an easy argument shows that the $CAT^*(\chi)$-inequality also holds for such triangles. By Claim 1, there is a digon D of minimum perimeter C_2. We shall show that D must actually be a closed geodesic. Let x, y be two opposite points in D. If $d(x,y) < \frac{1}{2}C_2$ then by introducing the geodesic segment from x to y we can chop D into two triangles T_1 and T_2 each of perimeter $< C_2$. Since the $CAT^*(\chi)$-inequality holds for such triangles, the Gluing Lemma would imply that it holds for D, a contradiction. Hence, any pair of opposite points of D must be of distance $\frac{1}{2}C_2$. This easily implies that D is the isometric image of a circle of length C_2, i.e., D is a closed geodesic and $C_1 = C_2$. □

2. Polyhedra of Piecewise Constant Curvature. Recall that M^n_χ denotes the simply connected, complete, Riemannian n-manifold of constant curvature χ. A *cell* in M^n_χ (= "convex polyhedral cell") is a compact, nonempty intersection of a finite number of half-spaces.

Suppose that K is a polyhedron. An M_χ*-cell-structure* on K is a cell structure on K together with a collection of maps $\{f_B \mid B \text{ is a cell in } K\}$ where f_B is a homeomorphism from B onto a cell in M^n_χ. If B and C are cells in K with $B \cap C \neq \emptyset$, then we further require that:

(1) $B \cap C$ is a common face of B and of C and

(2) $f_C \circ (f_B)^{-1}$ restricts to an isometry from $f_B(B \cap C)$ to $f_C(B \cap C)$.

Suppose that K is a polyhedron with an M_χ-cell structure. The *intrinsic pseudometric* on K is defined as follows: if $x, y \in K$, then $d(x,y)$ is the greatest lower bound of the length of all paths connecting x to y. ("Pseudometric" means that the possibility $d(x,y) = 0$ or ∞ is not ruled out.) If K is connected and locally finite, then d is obviously a metric.

We assume, from now on, that K is a connected, locally finite polyhedron with an M_χ-cell structure. The metric space K (equipped with the intrinsic metric) is called a *polyhedron of piecewise constant curvature* χ. If $\chi = 0$, K is *piecewise Euclidean*; if $\chi = 1$, it is *piecewise spherical.*

Remark 2.1. It makes sense to speak of an M-cell-structure, where M is some other Riemannian manifold for which there are notations of "half-space" and "convex polyhedral cell". For example, M could be the Cartesian product $M = M^{n_1}_{\chi_1} \times \cdots \times M^{n_k}_{\chi_i}$.

Piecewise spherical polyhedra play a special role in what follows, in that they arise naturally as "links". There are two types of "links" which we wish to define for a piecewise constant curvature polyhedron K: 1) the link of a point in K and 2) the link of a cell in some M_χ-cell-structure on K.

Suppose that E is a cell in M^n_χ and that $x \in E$. Then $\mathrm{Link}(x, E)$ is defined to be the subset of the unit sphere in $T_x(M^n_\chi)$ consisting of all tangent vectors to unit speed geodesics which originate at x and go into E. Obviously, $\mathrm{Link}(x, E)$ is the intersection of a finite number of "hyperplanes" and "half-spaces" in S^{n-1} (the unit sphere in $T_x(M^n_\chi)$). For example, if x is an interior point of E, then $\mathrm{Link}(x,E)$ is a round sphere of dimension one less than the dimension of E.

Now suppose that B is a cell in K, $x \in B$, and that $f_B : B \to M^n_\chi$ is a homeomorphism onto a convex cell in M^n_χ. Then put

$$\begin{aligned}\mathrm{Link}(x,B) &= \mathrm{Link}(f_B(x),f_B(B))\\ \mathrm{Link}(x,K) &= \cup\mathrm{Link}(x,B)\end{aligned}$$

where the union is taken over all cells B containing x and whenever C is a face of B, we identify $\mathrm{Link}(x,C)$ with corresponding face of $\mathrm{Link}(x,B)$. Then $\mathrm{Link}(x,K)$ is a piecewise spherical polyhedron, called the *link of x in K*.

Similarly, if E is a convex cell in M^n_χ and F is a face of E and if x is a point in the relative interior of F, then the set of tangent vectors of unit speed geodesics which originate at x, go into E, and are normal to F form a subpolyhedron of $\mathrm{Link}(x,E)$ which we denote by $\mathrm{Link}(F,E)$. This is a subset of S^{n-1} of dimension equal to $\dim E - \dim F - 1$. It is independent of the choice of x up to an isometry of S^{n-1}. Suppose that K is a polyhedron with M_χ-cell-structure, that B is a cell of K and C a face of B, then, as before, put

$$\begin{aligned}\mathrm{Link}(C,B) &= \mathrm{Link}(f_B(C),f_B(B))\\ \mathrm{Link}(C,K) &= \cup\mathrm{Link}(C,B),\end{aligned}$$

where the union is taken over all cells B containing C and whenever $C < B < A$, we identify $\mathrm{Link}(C,B)$ with the corresponding face of $\mathrm{Link}(C,A)$. Then $\mathrm{Link}(C,K)$ obviously has the structure of a piecewise spherical complex.

Definition 2.2. Let K be a polyhedron of piecewise constant curvature. A path $c : [a,b] \to K$ is a *broken geodesic path* if there exist numbers $t_0,\dots,t_n$ with $a = t_0 < t_1 < \cdots < t_n = b$ so that for each i, $0 \le i < n$, $c \mid [t_i,t_{i+1}]$ is a geodesic path with image lying entirely in some closed cell of K. By a *broken geodesic* we shall mean the image of a broken geodesic path together with an orientation. If a broken geodesic is the image of a geodesic path (that is, a locally isometric path) then it is called a *local geodesic*.

Suppose that x_0, x_1 are two points in some closed cell E of K and that α is a geodesic segment in E from x_0 to x_1. Then α determines a unit tangent vector in $T_{x_0}E$ and hence, a point in $\mathrm{Link}(x_0,K)$ called the *outgoing tangent vector of α at x_0* and denoted by $\alpha'_{out}(x_0)$. Similarly, α determines an *incoming tangent vector* $\alpha'_{in}(x_1) \in \mathrm{Link}(x_1,K)$.

Suppose that for $1 \le i \le n$, α_i is a geodesic segment in some cell of K from x_{i-1} to x_i. The α_i's can be glued together to give a broken geodesic α from x_0 to x_n. We shall use the notation $\alpha = (\alpha_1,\dots,\alpha_n)$. The incoming and outgoing vectors of the broken geodesic α make an "angle" θ_i at each point x_i, defined by,

$$\theta_i = d((\alpha_i)'_{in}(x_i), (\alpha_{i+1})'_{out}(x_i))$$

where d denotes distance in Link(x_i, K). The following lemma is not hard to prove.

LEMMA 2.3. ([M; Lemmas 4.1 and 4.2]). *With notation as above, the broken geodesic $\alpha = (\alpha_1, \dots \alpha_n)$ is a local geodesic if and only if $\theta_i \geq \pi$ for $1 \leq i \leq n$.*

Definition 2.4. Suppose that $E \subset S^k$ is a spherical k-cell and that $F \subset S^\ell$ is a spherical ℓ-cell. Identify S^k and S^ℓ with the unit spheres in subspaces of $\mathbb{R}^{k+\ell+2}$ which are orthogonal complements. Then $E * F$, the *orthogonal join of E and F*, is the $(k + \ell + 1)$-cell in $S^{k+\ell+1}$ defined as the convex hull of E and F (i.e., $E * F$ is the union of all geodesic segments in $S^{k+\ell+1}$ which begin in E and end in F). In an obvious fashion we can extend this to define the notion of the *orthogonal join* $L_1 * L_2$ of two piecewise spherical polyhedra L_1 and L_2. It is a piecewise spherical polyhedron of dimension equal to $\dim L_1 + \dim L_2 + 1$. In particular if L_1 is the round sphere S^{k-1}, then $S^{k-1} * L_2$ is called the *k-fold suspension of L_2*.

The relationship between the two types of links is explained in the following lemma (which is obvious).

LEMMA 2.5. *Suppose that K is a piecewise constant curvature polyhedron, that B is a k-cell in K and that x is a point in the relative interior of B. Then Link(x, K) is isometric to S^{k-1} * Link(B, K), the k-fold suspension of Link(B, K).*

3. The Large Link Condition. In this section we assume that K is a finite dimensional, locally finite polyhedron of piecewise constant curvature χ (equipped with its intrinsic metric). We further assume that there is an $\varepsilon > 0$ so that every closed ball in K of radius ε is compact. (This follows, for example, if K is a finite complex or if it admits a cell structure with only finitely many isometry types of cells.) It follows from this assumption that K is a complete geodesic space (cf. [M, Cor. 4.7]).

The following result is essentially stated by Gromov [G, pp. 119–120]; the nontrivial part of the proof is given in [GH, Chap. 10, Theorem 15].

THEOREM 3.1. (Gromov) *Let K be a polyhedron of piecewise constant curvature χ, as above, equipped with some M_χ-cell-structure. The following conditions are equivalent.*

(1) K has curvature $\leq \chi$.

(2) For each point x in K, the piecewise spherical polyhedron Link(x, K) satisfies the CAT(1)-inequality.

(3) *For each point x in K, the systole of Link(x,K) is $\geq 2\pi$.*

(4) *For each cell B in K, Link(B,K) satisfies the $CAT(1)$-inequality.*

(5) *For each cell B in K, the systole of Link(B,K) is $\geq 2\pi$.*

Proof. The equivalence of statements (1) and (2) is proved in [GH, Chap. 10, Thm. 15]

(2) $\Leftrightarrow$ (4) and (3) $\Leftrightarrow$ (5). Suppose that x is a point in the relative interior of some k-cell B. By Lemma 2.8, $\mathrm{Link}(x,K)$ is isometric to the k-fold suspension of $\mathrm{Link}(B,K)$. The equivalences (2) $\Leftrightarrow$ (4) and (3) $\Leftrightarrow$ (5) follow from Theorems A9 and A10 in the Appendix.

(4) $\Leftrightarrow$ (5). The implication (4) $\Rightarrow$ (5) is trivial. The proof that (5) $\Rightarrow$ (4) is by induction on the dimension of $\mathrm{Link}(B,K)$. Suppose (5) holds. In dimension one, the statement that the systole of a piecewise spherical complex is $\geq 2\pi$ means precisely that it satisfies $CAT(1)$. Let B be a cell in K and suppose by induction that $CAT(1)$ holds for all cells C with $\dim \mathrm{Link}(C,K) < \dim \mathrm{Link}(B,K)$. Put $L = \mathrm{Link}(B,K)$. If C is a cell in K with $B \subset C$, then it defines a cell σ_C in L with $\mathrm{Link}(\sigma_C, L) = \mathrm{Link}(C,K)$. Hence, (4) holds for any cell in L. Since (4) $\Rightarrow$ (2) $\Rightarrow$ (1), L has curvature ≤ 1; since (5) holds, $sys(L) \geq 2\pi$; hence, by Lemma 1.3, $CAT(1)$ holds for L i.e., (5) $\Rightarrow$ (4). This completes the proof. $\square$

If a piecewise constant curvature polyhedron k satisfies any of the conditions (2) through (5) of the Theorem, we say that K has "*large links*".

4. Orbifolds and Orbispaces. Suppose that Γ is a discrete group acting properly and freely on a manifold M. The orbit space M/Γ is then also a manifold and the natural projection is a covering map. If we drop the requirement that Γ act freely (but maintain the requirement that it act properly), then the orbit space naturally has the structure of an "orbifold".

Roughly speaking, an "orbifold" is a topological space which is locally modelled on quotients of $\mathbb{R}^n$ by finite group actions. The notion was introduced by Satake in [S] under the name "V-manifold". A precise definition can be found in [Th], where the terminology "orbifold" is introduced. A more sophisticated viewpoint towards the definition of "orbifold" is taken by Haefliger in [H]. According to Haefliger, an orbifold is a pseudo-group. (A *pseudo-group of local homeomorphisms* consists of a space X together with a collection of "local homeomorphisms" of X, i.e., homeomorphisms between certain open subsets of X. This collection of homeomorphisms is required to be closed under restriction to smaller open subsets; moreover, compositions and inverses are in the collection whenever they are defined.)

An *orbifold* Q consists of a space Q_0 (called the *underlying space* of Q) together with a pseudo-group of local homeomorphisms. The space X for the pseudo-group is a disjoint union of a collection of "uniformizing charts" of coordinate neighborhoods in Q_0. That is to say, $X = \coprod \widetilde{U}_i$ where $\{U_i\}$ is an open cover of Q_0 and where each $\widetilde{U}_i$ is an open subset of Euclidean space stable under some finite group G_i of homeomorphisms and where we have projections $\pi_i : \widetilde{U}_i \to U_i$, inducing homeomorphisms $\widetilde{U}_i/G_i \cong U_i$. Put $\pi = \coprod \pi_i : \coprod \widetilde{U}_i \to Q_0$. The orbifold is then the pseudo-group of all possible local homeomorphisms of $\coprod \widetilde{U}_i$ which commute with π. In particular, if $g \in G_i$ and $\widetilde{V} \subset \widetilde{U}_i$, then $g \mid \widetilde{V} : \widetilde{V} \to g(\widetilde{V})$ is an element of the pseudo-group.

If $x \in U_i$ and $\tilde{x} \in \widetilde{U}_i$ is a point lying over it, then the isotropy subgroup of G_i at $\tilde{x}$ depends only on x, up to isomorphism. It is called the *local group* at x An orbifold in which all local groups are trivial is precisely the same thing as a manifold.

Suppose that H is a finite group of germs of local homeomorphisms of $\mathbb{R}^n$ (taking 0 to 0). A *stratum of type* H in an orbifold Q is the subset of Q_0 consisting of all points x such that the local group at x is isomorphic to H. The *singular set* of Q is the union of all strata of nontrivial type.

If the orbifold is smooth and H is the type of a nonempty stratum, then H is isomorphic to a finite group of linear transformations of $\mathbb{R}^n$. Hence, in this case it follows that the stratum of type H is a manifold.

In [Th], Thurston demonstrates that covering space theory goes over to orbifolds. That is to say, there is an appropriate notion of an orbifold covering, the prototypical example of which is the projection $M/\Gamma' \to M/\Gamma$ where Γ acts properly on M and Γ' is a subgroup of Γ. Each orbifold Q has a universal orbifold cover. Its group of covering transformations is called the *orbifold fundamental group* and denoted by $\pi_1^{\mathrm{orb}}(Q)$.

A new feature which arises at this point in the discussion in [Th] is the notion of a "good orbifold". The universal cover of an orbifold need not be a manifold. If it is, the orbifold is called *good*. Thus, an orbifold is good if and only if it arises as the orbit space of a proper group action on a manifold. Equivalently, an orbifold is good if the natural map from each local group into the orbifold fundamental group is injective [H]. A basic example of a bad orbifold is the "teardrop". Its underlying space is S^2 and there is exactly one singular point, the local group at which is cyclic of order n, $n > 1$. (The teardrop is its own universal cover.)

It is suggested in [G, §4.5] that the notion of "orbifold" can be broadened to "orbispace". Roughly, an orbispace should be a topological space which is locally modelled on orbit spaces of certain finite group actions. The precise definition is accomplished in Haefliger's article [GH, Chap. 11] along the same lines as the definition for orbifold given above; basically, an orbispace is a pseudogroup.

In order to accomplish this broadening of the definition another new concept is needed: the notion of "rigidity". An orbispace is *rigid* if in each local model only the identity element of the group can fix (pointwise) a nonempty open set. It follows from Newman's Theorem (cf. [B, p. 157]) that every orbifold is rigid.

An orbispace is an *orbihedron* if each local model is a polyhedron and if each element of the pseudogroup is the restriction of a *PL*-homeomorphism.

The local group at a point in the underlying space of an orbispace is defined as before. An orbispace is *nonsingular* if each local group is trivial; an orbispace is *good* if it is covered by a nonsingular orbispace.

As suggested in [G, §4.5], covering space theory goes through, at least for rigid orbispaces. In particular, a rigid orbispace admits a universal orbispace cover.

Suppose that Q is an orbispace and that $\coprod \widetilde{U}_i$ is a disjoint collection of uniformizing charts occurring in the definition of the pseudogroup. (There is a natural notion of equivalence of pseudo-groups, which allows a change in the space $\coprod \widetilde{U}_i$; equivalent pseudogroups are regarded as defining the same orbispace.) One can impose geometric conditions on an orbispace by imposing them on each uniformizing chart $\widetilde{U}_i$ and requiring the pseudo-group to act by isometries. For example, we have the following:

(1) a *metric* on an orbispace means that each $\widetilde{U}_i$ admits a metric and that each element of the pseudogroup is an isometry,

(2) an orbihedron with metric is of *piecewise constant curvature* χ if each $\widetilde{U}_i$ is an open subset of a polyhedron of piecewise constant curvature χ,

(3) an orbifold admits a *Riemannian metric* if there is a Riemannian metric on each $\widetilde{U}_i$ (so that each element of the pseudogroup is an isometry).

Definition 4.1. A metric on an orbispace Q has *curvature* $\leq \chi$ if (up to an equivalence) each $\widetilde{U}_i$ is a geodesic space satisfying the $CAT(\chi)$-inequality.

The following generalization of the Cartan-Hadamard Theorem is stated in [G, p. 128] and proved by Haefliger in [GH, Chap. 11, Thm. 8].

THEOREM 4.2. (Gromov) *Let Q be a rigid orbispace of curvature ≤ 0. Then Q is good. (It then follows from Theorem 1.1 that the universal orbispace cover of Q is nonsingular and contractible.)*

Theorem 4.2 has recently been extended to nonrigid orbihedra by Spieler in [Sp].

Suppose that Q is an orbihedron of piecewise constant curvature. Let $x \in \widetilde{U}$ be a point over x. Then the link of x in Q_0 is naturally a piecewise spherical orbihedron which we denote by $\mathrm{Link}(x, Q)$. In fact, $\mathrm{Link}(x, Q)$ is the quotient space of the piecewise spherical polyhedron $\widetilde{\mathrm{Link}}(x, Q) = \mathrm{Link}(\tilde{x}, \widetilde{U})$ by a finite group G_x of isometries. If Q is cellulated so that each closed stratum is a subcomplex, then similar remarks hold for the link of each cell. In particular, for each cell B, $\mathrm{Link}(B, Q)$ is a good piecewise spherical orbihedron finitely covered by the nonsingular piecewise spherical polyhedron $\widetilde{\mathrm{Link}}(B, Q)$. The proof of Theorem 3.1 gives the following result.

PROPOSITION 4.3. *Let Q be an orbihedron of piecewise constant curvature χ. The following conditions are equivalent.*

(1) *Q has curvature $\leq \chi$.*

(2) *For each point $x \in Q$, $\widetilde{Link}(x, Q)$ satisfies the $CAT(1)$-inequality.*

(3) *For each cell B, $\widetilde{Link}(B, Q)$ satisfies the $CAT(1)$-inequality.*

5. Branched Covers of Riemannian Manifolds.

5.1. The Basic Example. Let m be an integer ≥ 2 and let $\pi_m : \mathbb{C} \to \mathbb{C}$ be the map $\pi_m(z) = z^m$. The cyclic group of m^{th} roots of unity acts on $\mathbb{C}$ by multiplication and we may regard π_m as the orbit map. Identify $\mathbb{C}$ with $\mathbb{R}^2$. The map $\pi_m : \mathbb{R}^2 \to \mathbb{R}^2$ is the prototypical example of an "m-fold branched cover"; "the branch set" is the origin. Put some smooth Riemannian metric on $\mathbb{R}^2$ (regarded as the range of π_m). Denote the corresponding distance function by d. Using π_m, one can pull back d to a metric $\tilde{d}$ on $\mathbb{R}^2$ (regarded as the domain of π_m). Denote $\mathbb{R}^2$ with the metric $\tilde{d}$ by $\widetilde{\mathbb{R}}^2$. The metric $\tilde{d}$ is smooth on the complement of 0. At 0, where there is, so to speak, "a cone angle of $2\pi m$", it is singular. The metric on $\widetilde{\mathbb{R}}^2$ can be understood as follows. Let R be a geodesic ray from the origin in $\mathbb{R}^2$. "Slit open" $\mathbb{R}^2$ along R. The resulting space H is homeomorphic to the half-plane. Its boundary consists of two copies of R (joined at their endpoints) which we denote by R_0 and R_1. To construct the branched cover $\widetilde{\mathbb{R}}^2$ one takes $H \times \{0, 1, \ldots, m-1\}$ and identifies $R_0 \times \{i\}$ with $R_1 \times \{i+1\}$, where the integers are taken modulo m. Each copy of H in $\widetilde{\mathbb{R}}^2$ is isometric to a slit-open version of $\mathbb{R}^2$. (If $\mathbb{R}^2$ is given the Riemannian metric of constant curvature χ, $\chi \leq 0$, then it follows from the above description that $\widetilde{\mathbb{R}}^2$ has the structure of a polyhedron of piecewise constant curvature χ.)

How does one describe the above example in the language of orbifolds? The underlying space of the quotient orbifold $\widetilde{\mathbb{R}}^2/\mathbb{Z}_m$ is naturally identified with $\mathbb{R}^2$. Thus, the underlying space of $\widetilde{\mathbb{R}}^2/\mathbb{Z}_m$ is naturally a smooth manifold and the orbifold metric on $\widetilde{\mathbb{R}}^2/\mathbb{Z}_m$ (read: the metric on $\widetilde{\mathbb{R}}^2$) is induced from a smooth Riemannian metric on the underlying space.

In a similar fashion it makes sense to speak of a metric orbifold Q such that its underlying space Q_0 is an n-dimensional manifold and such that the orbifold metric on Q is induced from a smooth Riemannian metric on Q_0. The goal of this paper is to analyze curvature properties of such orbifolds.

As we shall see in the next section a smooth structure on Q_0 does not necessarily induce a smooth orbifold structure on Q. Also, a singular stratum need not be a locally flat submanifold of Q_0. To see this point, consider the following example (which is not germane to the rest of the paper). The binary icosahedral group I acts freely on S^3. The orbit manifold H^3 is Poincaré's homology 3-sphere. Hence, its double suspension $\Sigma^2 H^3$ is an orbifold. (It is the quotient of a linear I-action on S^5.) On the other hand, a celebrated theorem of Cannon [Ca] and Edwards states that $\Sigma^2 H^3$ is homeomorphic to S^5, hence, it is possible to put a smooth structure and a smooth Riemannian metric on $\Sigma^2 H^3$. The singular set in $\Sigma^2 H^3$ consists of a single stratum—the suspension circle. Note that there is no smooth structure on $\Sigma^2 H^3$ in which this circle is smoothly embedded. (The complement of the circle in $\Sigma^2 H^3$ ($\cong S^5$) is not simply connected.)

The following lemma shows that such a pathological behavior cannot happen if we require the orbifold Q to have curvature bounded from above.

LEMMA 5.2. *Suppose that an orbifold Q has a metric induced from a smooth Riemannian metric on its underlying space Q_0. The following two conditions are necessary for the orbifold metric on Q to have curvature $\leq \chi$ (for some real number χ).*

(1) The sectional curvature of Q_0 is $\leq \chi$.

(2) Locally, the closure of each path component of each stratum must be convex.

Proof. (1) If Q satisfies $CAT(\chi)$ locally, then each nonsingular point of Q_0 must have a neighborhood in which the sectional curvature is bounded from above by χ. Since the set of nonsingular points is open and dense in Q_0, the necessity of (1) follows.

(2) Suppose that the local group at $x \in Q_0$ is G. Then there is a coordinate neighborhood U of x uniformized by $\widetilde{U}$ such that G acts on $\widetilde{U}$ with $\widetilde{U}/G \cong U$ and such that $\widetilde{U}$ satisfies $CAT(\chi)$. For H a subgroup of G, let U_H denote a component of a stratum of type H in U and $\overline{U}_H$ its closure. Let y_0 and y_1 be points of $\overline{U}_H$ and γ a geodesic segment connecting them. Then γ is homotopic rel endpoints to a path in $\overline{U}_H$. It follows that some lift $\tilde{\gamma}$ of γ to $\widetilde{U}$ is a geodesic segment connecting points $\tilde{y}_0$, $\tilde{y}_1$ in $\widetilde{U}^H$ lying over y_0 and y_1. ($\widetilde{U}^H$ denotes the fixed point set of H on $\widetilde{U}$.) If for some $h \in H$, $h\tilde{\gamma} \neq \tilde{\gamma}$, then $h\tilde{\gamma}$ and $\tilde{\gamma}$ fit together to give a geodesic digon from $\tilde{y}_0$ to $\tilde{y}_1$ contradicting $CAT(\chi)$. (If $\chi > 0$, assume $d(y_0, y_1) < \pi/\sqrt{\chi}$.) Hence, $\tilde{\gamma}$ lies in $\widetilde{U}^H$, i.e., condition (2) is necessary. □

We retain the notation in the preceding proof: Q is an orbifold of curvature $\leq \chi$, $x \in Q_0$, U is a coordinate neighborhood of x, G is the local group at x, and $\widetilde{U}$ is a uniformization of U. Denote by T_x the tangent space of Q_0 at x, by S_x the unit sphere in T_x, and by $B_x(r)$ the open ball of radius r about the origin in T_x. Let $\exp : T_x \to Q_0$ denote the exponential map. Choose r small enough so that the restriction of exp to $B_x(r)$ is a diffeomorphism onto its image and so that $\exp(B_x(r)) \subset U$. Using exp one pulls back the orbifold structure on U to an orbifold structure on $B_x(r)$. By Lemma 5.2, (2) each component of each stratum of $B_x(r)$ is a convex subset of some linear subspace of T_x. Let $\lambda : [0, r) \to [0, \infty)$ be a diffeomorphism, linear near 0. The radial expansion $x \to \lambda(|x|)x/|x|$ is a diffeomorphism of $B_x(r)$ with T_x. Using it the orbifold structure on $B_x(r)$ induces one on T_x. As before, the closure of each component of each stratum of T_x is a convex subset containing 0. The orbifold fundamental group of T_x is G. Let $\widetilde{T}_x \to T_x$ be the universal orbifold cover and $\widetilde{B}_x(r)$ and $\widetilde{S}_x$ the corresponding covers of $B_x(r)$ and S_x, respectively. The map $\exp : B_x(r) \to U$ is covered by a G-equivariant embedding $\widetilde{\exp} : \widetilde{B}_x(r) \to \widetilde{U}$.

The Riemannian metric on Q_0 gives an inner product on T_x. Since the singular set of T_x is a finite union of convex polyhedral cones, it follows that $\widetilde{B}_x$ and $\widetilde{B}_x(r)$ are naturally piecewise Euclidean polyhedra. Similarly, $\widetilde{S}_x$ is a piecewise spherical polyhedron.

Using polar coordinates one can identify T_x with constant curvature space M^n_χ, provided $\chi \leq 0$, in such a fashion that distance from the origin is preserved. If $\chi > 0$, then one can, in a similar fashion, identify $B_x(r)$ with an open ball in M^n_χ provided $r \leq \pi/\sqrt{\chi}$.

The main result of this section is the following.

THEOREM 5.3. *Let Q be an orbifold with orbifold metric induced from a smooth Riemannian metric on its underlying space Q_0. The following conditions are necessary and sufficient for the curvature of Q to be $\leq \chi$.*

(1) *The sectional curvature of Q_0 is $\leq \chi$.*

(2) *Locally, the closure of each path component of each stratum is convex in Q_0.*

(3) *For each $x \in Q_0$ the piecewise spherical polyhedron $\widetilde{S}_x$ satisfies $CAT(1)$.*

Before proving this theorem, let us recall Alexandrov's version of the $CAT(\chi)$ inequalities in terms of "angles" (cf. [GH, Chap. 3]).

A *geodesic hinge* in a geodesic space X consists of two geodesic segments $g : [0, a] \to X$ and $h : [0, b] \to X$, parametrized by arc length, with the same initial point x (i.e., $g(0) = h(0) = x$). For any $(s, t) \in [0, a] \times [0, b]$ we have a triangle $\triangle$ in X with vertices $x, g(s)$ and $h(t)$. Let $\triangle^*$ be a comparison triangle in M^2_χ and let $\alpha(s, t)$ be the angle at the vertex of $\triangle^*$ corresponding to x. Alexandrov's condition $A(\chi)$ is the following:

$A(\chi)$: For any geodesic hinge, the function $\alpha(s, t)$ is a monotone nondecreasing function of s and t.

It is proved in [GH, Chap. 3, §1, Thm. 4] that the conditions $A(\chi)$ and $CAT(\chi)$ are equivalent.

We now return to the situation of Theorem 5.3: Q is an orbifold with underlying space Q_0 a Riemannian manifold, U is a neighborhood of x in Q_0, $\widetilde{U}$ is a uniformization of U, $\tilde{x}$ is the lift of x in $\widetilde{U}$, G is the group of deck transformation of $\widetilde{U} \to U$, B_x is a small ball about the origin in T_x with $\exp(B_x) \subset U$ and $\widetilde{B}_x$ is its inverse image in $\widetilde{T}_x$. Also, T_x is given the metric of constant curvature χ and $\widetilde{B}_x$ and $\widetilde{T}_x$ have the induced piecewise constant curvature metrics $\tilde{d}$.

Let $u : [0,a] \to \widetilde{U}$ and $v : [0,b] \to \widetilde{U}$ be a geodesic hinge in $\widetilde{U}$ based at $\tilde{x}$. This geodesic hinge has two well-defined "unit tangent vectors" u' and v' in $\widetilde{S}_x$. The distance θ from u' to v' in $\widetilde{S}_x$ is the *angle at* $\tilde{x}$ of the geodesic hinge. Let $u'(s)$ and $v'(t)$ denote the geodesic hinge in $\widetilde{T}_x$ in directions u' and v'. Recalling the construction in Alexandrov's condition, the geodesic hinge $u(s), v(t)$ in $\widetilde{U}$ gives us a family of angles $\alpha(s,t)$ for the comparison triangles in M^2_χ. Similarly, $u'(s)$ and $v'(t)$ in $\widetilde{T}_x$ gives us a family of angles $\alpha'(s,t)$. It is clear that as $(s,t) \to (0,0)$ both $\alpha(s,t)$ and $\alpha'(s,t)$ approach θ.

Proof of Necessity in Theorem 5.3. The necessity of conditions (1) and (2) was established in Lemma 5.2. Suppose that condition (3) fails. Then there is a triangle $\triangle'$ in $\widetilde{S}_x$ for which $CAT(1)$ fails. Let $\triangle^*$ be the comparison triangle in S^2. Let u', v', w' be the vertices of $\triangle'$ and u^*, v^*, w^* the corresponding vertices of $\triangle^*$. The failure of $CAT(1)$ entails the existence of a point z' on $\triangle'$, say on the segment from v' to w', such that the distance to the opposite vertex is too large, i.e.,

$$d(u', z') > d(u^*, z^*) \tag{1}$$

where z^* is the point on $\triangle^*$ corresponding to z'. Put $\theta' = d(u', z')$ and $\theta^* = d(u^*, z^*)$ so that (1) becomes:

$$\theta' > \theta^*. \tag{1$'$}$$

Denote the geodesic ray from the origin in $\widetilde{T}_x$ in direction u' parametrized by arc length by $t \to u'(t)$. Similarly, let $t \to v'(t)$ and $t \to w'(t)$ be geodesic rays in directions v' and w'. Regarding S^2 as the unit sphere in the tangent space to M^3_χ at a point x^*, we obtain, in a similar fashion, geodesic rays $u^*(t)$, $v^*(t)$, $w^*(t)$ in M^3_χ. Let $z'(t)$ be the point on the segment from $v'(t)$ to $w'(t)$ such that

$$\frac{\tilde{d}(v'(t), z'(t))}{\tilde{d}(v'(t), w'(t))} = \frac{d(v', z')}{d(v', w')} \tag{2}$$

Let $z^*(t)$ be the point in M^3_χ on the segment from $v^*(t)$ to $w^*(t)$ in direction z^*. It follows from (1)$'$ that the rays $u'(t)$ and $z'(t)$ make a larger angle at $0 \in \widetilde{T}_x$ then the rays $u^*(t)$ and $z^*(t)$ make at $x^* \in M^3_\chi$. Thus,

$$\tilde{d}(u'(t), z'(t)) > d(u^*(t), z^*(t))$$

and in fact

$$\lim_{t \to 0} \frac{\tilde{d}(u'(t), z'(t))}{d(u^*(t), z^*(t))} = 1 + c \tag{3}$$

where c is some positive constant (which could be solved for explicitly in terms of the "angles" θ' and θ^*). Let U be a coordinate neighborhood of x, uniformized by $\widetilde{U}$ and let B_x be a ball about the origin in T_x of radius small enough so that $\exp(B_x) \subset U$. Let $u(t), v(t), w(t)$ be the images of $u'(t)$, $v'(t)$, $w'(t)$ under $\widetilde{\exp} : \widetilde{B}_x \to \widetilde{U}$. Let $z(t)$ be the point on the segment from $v(t)$ to $w(t)$ so that

$$\frac{d(v(t), z(t))}{d(v(t), w(t))} = \frac{d(v', z')}{d(v', w')} \tag{4}$$

Let $\tilde{x}$ be the lift of x to $\widetilde{U}$ and let $\triangle_1, \triangle_2, \triangle_3$ be triangles in $\widetilde{U}$ with vertex sets $\{u(t), v(t), \tilde{x}\}$, $\{u(t), w(t), \tilde{x}\}$ and $\{v(t), w(t), \tilde{x}\}$, respectively. The comparison triangles for $\triangle_1, \triangle_2, \triangle_3$ can be glued together along their corresponding sides to form the embedded 1-skeleton of a tetrahedron in M^3_χ. Let $\hat{u}(t)$, $\hat{v}(t)$, and $\hat{w}(t)$ be the points in M^3_χ corresponding to $u(t), v(t)$, and $w(t)$, respectively, and $\hat{\triangle} \subset M^3_\chi$ and $\triangle \subset \widetilde{U}$, the corresponding triangles. Thus, $\hat{\triangle}$ is a comparison triangle for $\triangle$. Let $\hat{z}(t)$ be the point on $\hat{\triangle}$ corresponding to $z(t)$. Up to first order, distance in U is the same as in T_x. The same statement therefore holds for the induced metrics in $\widetilde{U}$ and $\widetilde{T}_x$. Thus, up to first order, the distance from $u'(t)$ to $z'(t)$ depends only on the angle that these rays make at $0 \in \widetilde{T}_x$. This angle is θ'. As $t \to 0$, the angle between the rays $u(t)$ and $z(t)$ at $\tilde{x}$ also goes to θ'. Hence,

$$\lim_{t \to 0} \frac{d(u'(t), z'(t))}{d(u(t), z(t))} = 1 \tag{5}$$

In M^3_x, the angle between the rays $u^*(t)$ and $z^*(t)$ is θ^*; while as $t \to 0$ the angle between $\hat{u}(t)$ and $\hat{z}(t)$ also approaches θ^*. Hence,

$$\lim_{t \to 0} \frac{d(\hat{u}(A), \hat{z}(t))}{d(u^*(t), z^*(t))} = 1. \tag{6}$$

Combining (3), (5) and (6) we get

$$\lim_{t \to 0} \frac{d(u(t), z(t))}{d(\hat{u}(t), \hat{z}(t))} = 1 + c \tag{7}$$

That is to say, for arbitrarily small values of t, the $CAT(\chi)$ inequality fails for $\triangle$. This establishes the necessity of condition (3) in Theorem 5.3. □

Our proof of sufficiency in Theorem 5.3 is a modification of the arguments in [GH, Chap. 3 §2]. The proof is organized as a sequence of lemmas, 5.4 to 5.7, below.

LEMMA 5.4. *Suppose that Q satisfies conditions (1), (2) and (3) in Theorem 5.3. Then $\widetilde{T}_x$ is a polyhedron of piecewise constant curvature χ and it satisfies $CAT(\chi)$.*

Proof. This follows immediately from the basic result of Gromov, Theorem 3.1. The point is that condition (3) is precisely the condition that $\widetilde{T}_x$ has "large links". □

LEMMA 5.5. *Suppose that Q satisfies conditions (1) and (2) in Theorem 5.3. Let $u : [0, a] \to \widetilde{U}$ and $v : [0, b] \to \widetilde{U}$ be a geodesic hinge in $\widetilde{U}$ based at $\tilde{x}$. Let $u'(t)$ and $v'(s)$ be the corresponding rays in $\widetilde{T}_x$, i.e., $u(t) = \widetilde{exp}\, u'(t)$ and $v(s) = \widetilde{exp}\, v'(s)$. Then*

$$d(u(t), v(s)) \geq \tilde{d}(u'(t), v'(s)).$$

(In other words, "$\widetilde{exp}$ spreads geodesic hinges".)

Proof. Since the sectional curvature of Q_0 is $\leq \chi$, the map $\exp : B_x \to U$ spreads geodesic hinges (cf. [Gv, p. 197]). The lemma follows easily from this and the fact that geodesic segments in $\widetilde{U}$ project to piecewise geodesic segments in U. □

Recalling the construction in Alexandrov's condition, the geodesic hinge $u(s)$, $v(t)$ in $\widetilde{U}$ gives us a family of angles $\alpha(s, t)$ for the comparison triangles in M^2_χ. Similarly, the geodesic hinge $u'(s)$, $v'(t)$ in $\widetilde{T}_x$ gives a family of angles $\alpha'(s, t)$. Both geodesic hinges make the same angle, call it θ.

LEMMA 5.6. *Suppose that Q satisfies conditions (1), (2) and (3) in Theorem 5.3. Suppose further, as above, that $u(s)$, $v(t)$ is a geodesic hinge in $\widetilde{U}$ and that $u'(s)$, $v'(t)$ is the corresponding geodesic hinge in $\widetilde{T}_x$. Then*

$$\alpha(s, t) \geq \alpha'(s, t) \geq \theta.$$

Proof. It follows immediately from Lemma 5.5 that $\alpha(s, t) \geq \alpha'(s, t)$. By Lemma 5.4, $\widetilde{T}_x$ satisfies $CAT(\chi)$, so by Alexandrov's condition $A(\chi)$, $\alpha'(s, t)$ is a nondecreasing function of s and t. Thus, $\alpha'(s, t) \geq \theta$. □

LEMMA 5.7. *Let $\triangle$ be a triangle in $\widetilde{U}$ and $\triangle^*$ a comparison triangle in M^2_χ. For $i = 1, 2, 3$, let θ_i be the interior angles at the vertices of $\triangle$ and θ^*_i the corresponding angles of $\triangle^*$. Then*

$$\theta_i \leq \theta^*_i.$$

Proof. This follows immediately from the previous lemma. □

Proof of Sufficiency in Theorem 5.3. It follows from Lemma 5.7 that any geodesic hinge in $\widetilde{U}$ satisfies Alexandrov's condition $A(\chi)$. The argument is exactly the same as the proof of Théoremè 9 on pages 55 and 56 of [GH; Chap 3 §2]. Since $A(\chi)$ and $CAT(\chi)$ are equivalent, Theorem 5.3 is proved. □

6. Local Models: Orbifold Structures on Euclidean Space. In the previous section we focused on orbifolds with metric induced from a smooth Riemannian metric on the underlying space and with curvature bounded from above. The local model for such an orbifold at a point x in its underlying space is the tangent space at x with its induced orbifold structure. Such a local model is an orbihedron V satisfying the following conditions.

(A) The underlying space of V is identified with a real n-dimensional vector space V_0 equipped with an inner product.

(B) The closure of each component of each stratum of V is a convex subset of V_0 containing the origin.

(C) V is good. Its orbihedral fundamental group G is finite and is identified with the local group at the origin.

(D) The universal cover $\widetilde{V}$ (which by (B) has a natural PL-structure) is PL homeomorphic to $\mathbb{R}^n$.

(E) The piecewise Euclidean polyhedron $\widetilde{V}$ satisfies $CAT(0)$.

Remark 6.1. If $S(V)$ denotes the restriction of V to the unit sphere in V_0 and $\widetilde{S}(V)$ denotes its universal cover, then (D) and (E) are equivalent to the conditions that $\widetilde{S}(V)$ be PL-homeomorphic to S^{n-1} and that $\widetilde{S}(V)$ satisfy $CAT(1)$.

Of course, these conditions are motivated by Theorem 5.3. In particular, (B) is implied by condition (2) in 5.3 and (E) is implied by condition (3) in 5.3. Condition (D) is forced on us if we are interested in orbifolds (rather than just orbihedra).

The purpose of this section is to develop some examples which satisfy conditions (A) through (E) above.

Products. The Cartesian product $V \times V'$ of two orbifolds, defined in the obvious manner, is again an orbifold. The underlying space of $V \times V'$ is $V_0 \times V_0'$ and its orbifold fundamental group is $G \times G'$ (where G and G' denote the orbifold fundamental groups of V and V', respectively).

PROPOSITION 6.2. *Suppose that V and V' are orbifolds satisfying conditions (A) through (E) above. Let the inner product on $V_0 \times V_0'$ be the direct product of the inner products on V_0 and V_0'. Then $V \times V'$ also satisfies conditions (A) through (E).*

Proof. It is obvious that $V \times V'$ satisfies (A) through (D). Since V and V' satisfy (E), the vector spaces V_0 and V_0' have inner products so that the induced metrics on $\widetilde{V}$ and $\widetilde{V}'$ satisfy $CAT(0)$. The direct product of these inner products induces the product metric on $\widetilde{V} \times \widetilde{V}'$. To prove that $V \times V'$ satisfies (E) it suffices to show that $\widetilde{V} \times \widetilde{V}'$ satisfies $CAT(0)$. This is immediate from the following lemma. □

LEMMA 6.3. *Suppose that P_1 and P_2 are piecewise Euclidean polyhedra of curvature ≤ 0. Then $P_1 \times P_2$ with the product metric is a piecewise Euclidean polyhedron and its curvature is ≤ 0.*

Proof. The link of a point $(x_1,x_2) \in P_1 \times P_2$ is the "orthogonal join" of $L_{x_1} = \mathrm{Link}(x_1,P_1)$ and $L_{x_2} = \mathrm{Link}(x_2,P_2)$. By Theorem 3.1, L_{x_1} and L_{x_2} satisfy $CAT(1)$; hence, by Theorem A10 in the Appendix, so does their orthogonal join. The product $P_1 \times P_2$ has a natural piecewise Euclidean structure. We have just proved that $P_1 \times P_2$ has "large links". Hence, the result of Gromov, Theorem 3.1, implies that $P_1 \times P_2$ has curvature ≤ 0. □

Next we want to state a partial converse to Proposition 6.2 (namely Proposition 6.5 below). To state the converse we need to introduce a new condition. Suppose that V satisfies conditions (A) through (E). Let $S_1, \ldots, S_k$ denote the components of the codimension two strata of V and let H_i be the linear subspace of codimension two in V_0 spanned by S_i. Put $H = \cap H_i$. Consider the following condition:

(F) $H = \{0\}$.

If V is the product orbifold $\mathbb{R}^m \times V'$, $m > 0$, where $\mathbb{R}^m$ has empty singular set, then we shall say that V has a *trivial factor*. If $\mathbb{R}^m$ is a trivial factor of V, then $\mathbb{R}^m \subset H$. Hence, if V satisfies (F), then it has no trivial factor.

Remark 6.4. Suppose that the orbifolds W, V, and V' satisfy (A) through (D) and that V and V' have no trivial factors. Suppose further that there is a strata-preserving homeomorphism $V_0 \times V_0' \to W_0$. Then the image of (nonsingular set of V_0) $\times 0$ in W_0 is a component of a stratum. By (D), its closure $V_0 \times 0$ is convex; hence, a linear subspace. Similarly, the image of $0 \times V_0'$ is also a linear subspace of W_0.

PROPOSITION 6.5. *Suppose V and V' are orbifolds satisfying (A) through (F). Let W be the product orbifold, $W = V \times V'$, equipped with an inner product $\langle \ , \rangle$ on $W_0 = V_0 \times V_0'$ which restricts to the given inner products on V_0 and V_0'. Then W satisfies (A) through (D) and (F). Moreover, it satisfies (E) if and only if $V_0 \times 0$ and $0 \times V_0'$ are complementary orthogonal subspaces in W_0.*

The proof of Proposition 6.5 depends on the following lemma.

LEMMA 6.6. *Suppose Q is an orbifold with underlying space $Q_0 = S^3$, singular set two disjoint great circles C_1 and C_2, and local groups $\mathbb{Z}/p$ on C_1 and $\mathbb{Z}/q$ on C_2. Then the universal cover $\widetilde{Q}$ of Q satisfies $CAT(1)$ if and only if C_1 and C_2 are orthogonal.*

Proof. If C_1 and C_2 are orthogonal, then Q_0 is the orthogonal join of C_1 and C_2 and $\widetilde{Q}$ is the orthogonal join of $\widetilde{C}_1$ and $\widetilde{C}_2$, where $\widetilde{C}_1, \widetilde{C}_2$ are, respectively, the q-fold and p-fold covers of C_1 and C_2. Clearly $\widetilde{C}_1$ and $\widetilde{C}_2$ satisfy $CAT(1)$, so by Theorem A10 of the Appendix, so does $\widetilde{Q}$.

Conversely, suppose C_1 and C_2 are *not* orthogonal. Let α be a geodesic segment from C_1 to C_2 of minimal length. By minimality, α must be perpendicular to C_1 and C_2 at its endpoints, and since C_1 and C_2 are not orthogonal, the length of α is $< \frac{\pi}{2}$. Let $\tilde{\alpha}$ be a lift of α to $\widetilde{Q}$. The fundamental group of Q is $\mathbb{Z}/p \times \mathbb{Z}/q = \langle s \rangle \times \langle t \rangle$, hence the sequence $\gamma = (\tilde{\alpha}, t\tilde{\alpha}^{-1}, st\tilde{\alpha}, s\tilde{\alpha}^{-1})$ of geodesic segments is a closed broken geodesic in $\widetilde{Q}$. (Here $\tilde{\alpha}^{-1}$ denotes $\tilde{\alpha}$ with its orientation reversed.) The length of γ is four times the length of $\tilde{\alpha}$, hence length $(\gamma) < 2\pi$. We claim that γ is a closed *local* geodesic, so the systole of $\widetilde{Q}$ is $< 2\pi$. By Lemma 1.3, this means that $CAT(1)$ fails in $\widetilde{Q}$.

To show that γ is a local geodesic, it suffices, by Lemma 2.6, to show that if two lifts of $\tilde{\alpha}$ meet at $\tilde{x}$, then their tangent vectors have distance $\geq \pi$ in $\widetilde{L}_x = \mathrm{Link}(\tilde{x}, \widetilde{Q})$. Now if $\tilde{x}$ lies over $x \in Q_0$, then $\widetilde{L}_x$ is the universal cover of the orbifold $L_x = \mathrm{Link}(x, Q)$. Say $x \in C_1$. Then L_x has underlying space S^2 and two singular points (the two tangent vectors to C_1 at x) each with local group $\mathbb{Z}/p$. Viewing the singular points as the "north and south poles" of S^2, the tangent vector to α appears as a point on the equator (since α is perpendicular to C_1 at x). The universal cover $\widetilde{L}_x$ of L_x is the suspension of a circle of circumference $2\pi p$. (We may view $\widetilde{L}_x$ as obtained from L_x by slitting the 2-sphere along a geodesic between the north and south poles and gluing together p copies of this slit sphere.) In particular, if $\tilde{\alpha}_1, \ldots, \tilde{\alpha}_p$ are the lifts of α at $\tilde{x}$, then the tangent vectors to $\tilde{\alpha}, \ldots, \tilde{\alpha}_p$ appear as p points in $\widetilde{L}_x$ equally spaced along its equator. It follows from Lemma A6 of the appendix that the distance between any two of these points in $\widetilde{L}_x$ is π. $\square$

Proof of Proposition 6.5. Clearly W satisfies (A) through (D). If $V_0 \times 0$ and $0 \times V_0'$ are orthogonal in W_0, then $\langle\ ,\ \rangle$ is necessarily the direct product of the given inner products on V_0 and V_0'. It follows from Proposition 6.2 that W satisfies (E).

Conversely, assume W satisfies (E). Denote a component of a codimension two stratum of V by S_i and let H_i be the codimension two linear subspace it spans. Let P_i denote the orthogonal 2-plane to H_i in V_0 and let P be the subspace of V_0 spanned by the P_i. Similarly, let S_j' be a component of a codimension two stratum of V' and define subspaces H_j', P_j', and P' of V_0' analogously. Consider a component of a stratum of W of the form $S_i \times S_j'$. The link of a top dimensional cell in $S_i \times S_j'$ has the form $S(P_i \oplus P_j')$. Since W satisfies (E), this link must satisfy

$CAT(1)$. Now the underlying space of $S(P_i \oplus P'_j)$ is a standard 3-sphere and its singular set is the union of two great circles $H_i \cap S(P_i \oplus P')$ and $H'_j \cap S(P_i \oplus P'_j)$. By Lemma 6.6, these circles are orthogonal. Since $P_i = H_i^{\perp}$ and $P'_j = (H'_j)^{\perp}$, it follows that P_i and P'_j must be orthogonal. Hence, P is orthogonal to P'. Condition (F) implies that $P = V_0$ and $P' = V'_0$ and the proposition is proved. □

Real Reflection Groups. An interesting class of groups for which conditions (A) through (D) hold are the orientation-preserving subgroups of real reflection groups. We describe these below.

A linear involution of $\mathbb{R}^n$ is a *real reflection* if the eigenvalue 1 has multiplicity $n-1$. A finite subgroup $\widehat{G}$ of $GL(n, \mathbb{R})$ is a *real reflection group* if it is generated by real reflections. The subgroup G of index two in $\widehat{G}$ defined by $G = \widehat{G} \cap SL(n, \mathbb{R})$ is called the *orientation-preserving subgroup* of $\widehat{G}$.

The theory of such real reflection groups $\widehat{G}$ is extremely well-known, e.g., see [Bo]. In particular, (a) $\widehat{G}$ is a Coxeter group, (b) there is a " fundamental chamber $C \subset \mathbb{R}^n$ such that $\mathbb{R}^n/\widehat{G} \cong C$, (c) if $\widehat{G}$ fixes no point of $\mathbb{R}^n$ other than the origin, then C is a simplicial cone (cf. [Bo, Prop. 11, p. 88]). Condition (c) means that $S^{n-1} \cap C$ is a spherical $(n-1)$-simplex $\triangle^{n-1}$.

We now suppose (c) holds. The underlying space of the orbifold S^{n-1}/G is then obviously homeomorphic to the double of $\triangle^{n-1}$. We can identify the underlying space of S^{n-1}/G with S^{n-1} by identifying one copy of $\triangle^{n-1}$ with a spherical simplex (or hemisphere) σ^{n-1} in S^{n-1} and the other copy of $\triangle^{n-1}$ with the exterior, $S^{n-1} - \operatorname{int}(\sigma^{n-1})$. The underlying space of $\mathbb{R}^n/G$ is then identified with $\mathbb{R}^n$ in such a fashion that (B) holds. We therefore have the following result.

LEMMA 6.7. *Suppose that G is the orientation-preserving subgroup of a real reflection group. Let V be the orbifold $\mathbb{R}^n/G$ and identify its underlying space V_0 with $\mathbb{R}^n$, as above. Then V satisfies conditions (A) through (D).*

In Section 8 we shall decide when such a V satisfies the curvature condition (E). This will depend on the geometry of σ^{n-1} in S^{n-1}.

Complex Reflection Groups. A complex linear automorphism of $\mathbb{C}^n$ is a *complex reflection* if it has finite order and if the eigenvalue 1 has multiplicity $n-1$. (Thus, a "complex reflection" is actually a rotation about a complex hyperplane.) A finite subgroup G of $GL(n, \mathbb{C})$ is a *complex reflection group* if it is generated by complex reflections.

Our basic example (Example 5.1) of $\mathbb{Z}/m$ acting on $\mathbb{C}$ is a complex reflection group, as is the Cartesian product of copies of the basic example.

A well-known theorem of Chevalley [Ch] asserts that for a finite subgroup G of $GL(n, \mathbb{C})$, the following two conditions are equivalent:

(i) G is a complex reflection group,

(ii) $\mathbb{C}^n/G$ is algebraically isomorphic to $\mathbb{C}^n$.

Condition (ii) implies that $\mathbb{C}^n/G$ is actually homeomorphic to $\mathbb{C}^n$. Thus, the orbifold of a complex reflection group G satisfies conditions (A), (C) and (D).

The singular set of $\mathbb{C}^n/G$ ($\cong \mathbb{C}^n$) is the image of the union of the reflecting hyperplanes. The image of such a hyperplane is homeomorphic to $\mathbb{C}^{n-1}$; however, it need not be embedded in a locally flat fashion in the orbit space. In fact, in most cases it is locally knotted. For example, the dihedral group D_m of order $2m$ is a real reflection group on $\mathbb{R}^2$; complexifying, it becomes a complex reflection group on $\mathbb{C}^2$. If m is odd, then the intersection of the singular set in $\mathbb{C}^2/D_n$ with S^3 is a torus knot of type $(m,2)$; hence, if m is odd and ≥ 3 the singular set is locally knotted. We note that for such singular sets condition (B) cannot hold.

In addition to cyclic groups, there is one further class of irreducible complex reflection groups for which (B) holds. These are certain complex reflection groups acting on $\mathbb{C}^2$ which we denote by $G(p,q,r)$. They will be discussed at the end of this section.

Subspaces in General Position. Suppose we are given a collection of linear codimension two subspaces $\{H_1,\dots,H_\ell\}$ in $\mathbb{R}^m$ in general position. (By this we mean that if we take any subset of $\{H_1,\dots,H_\ell\}$, then the orthogonal complements to the H_i in this subset span a subspace of greatest possible dimension in $\mathbb{R}^m$.) Label each H_i by an integer a_i, $a_i \geq 2$. This is the data for a piecewise Euclidean orbihedron V, with underlying space $\mathbb{R}^m$, with singular set $\cup H_i$ and with local group at a generic point of H_i cyclic of order a_i.

By construction, V satisfies (A) and (B). What about (C), (D) and (E)? In the case $2\ell \leq m$, this is answered by the following proposition.

PROPOSITION 6.8. *With notation as above, if $2\ell \leq m$, then V satisfies (A) through (D), and its orbifold fundamental group is $G = \mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_\ell$. Moreover, V satisfies (E) if and only if the two dimensional subspaces $H_1^\perp,\dots,H_\ell^\perp$ are mutually orthogonal.*

Proof. Let $P_i = H_i^\perp$ be the orthogonal complement of H_i in $\mathbb{R}^n$. Let $H = \cap H_i$, $P = H^\perp = P_1 + \cdots + P_\ell$. If $2\ell \leq m$, then the P_i's are linearly independent, hence, $\dim P = 2\ell$ and $\dim H = m - 2\ell$. In this case, the orbifold structure on each P_i is a basic two-dimensional example (as in Example 5.1) and the orbifold structure on P is a product of these basic examples. Hence $V = H \times P = H \times P_1 \times \cdots \times P_\ell$ and the proposition follows from the discussion of products above. $\square$

We turn now to the case $2\ell > m$. Let G denote the orbihedral fundamental group of V. We need to decide when V is good, when G is finite and when the universal cover $\widetilde{V}$ is PL-homeomorphic to $\mathbb{R}^m$. The group G is defined as follows. Let $S = \cup H_i$ be the singular set in V_0. Let γ_i be a small circle in $V_0 - S$ normal to H_i and let N be the normal subgroup of $\pi_1(V_0 - S)$ generated by the $(\gamma_i)^{a_i}$. Then $G = \pi_1(V_0 - S)/N$. Since $H_1(V_0 - S)$ is free abelian on $\gamma_1,\dots,\gamma_\ell$, it follows that the abelianization of G is the product of cyclic groups $\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_\ell$. The unit sphere in $\mathbb{R}^m$ inherits the structure of a piecewise spherical orbihedron $S(V)$. Let $\widetilde{S}^{m-1}$ denote its $(\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_\ell)$-fold branched cover.

Let us first consider the case $m = 2\ell - 1$.

LEMMA 6.9. *Let V be as above with $m = 2\ell - 1$. Let $\widetilde{S}^{2\ell-2} \to S(V)$ denote the branched cover with $\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_\ell$ as group of deck transformations. Then*

(i) the Euler characteristic of $\widetilde{S}^{2\ell-2}$ is given by the formula:

$$\chi(\widetilde{S}^{2\ell-2}) = 2\left[(a_1 + \cdots + a_\ell) - \sum_{i<j} a_i a_j\right],$$

(ii) $\chi(\widetilde{S}^{2\ell-2}) \leq 0$ (and consequently $\widetilde{S}^{2\ell-2}$ is not a rational homology sphere)

(iii) any finite sheeted orbihedral covering of $\widetilde{S}^{2\ell-2}$ is not a rational homology sphere.

Proof. It suffices to prove (i). (Indeed, it is an easy exercise to see that the right hand side of the formula in (i) is ≤ 0, provided each $a_i \geq 2$. A transfer argument shows that (ii) implies (iii).)

To prove (i), let χ_j denote the Euler characteristic of $\widetilde{S}^{2j-2}$, $j \geq 2$, where the branch set corresponds to the intersection of $H_1, \ldots, H_j$ with S^{2j-2}. Put $\chi_1 = 2a_1$. Consider the $\mathbb{Z}/a_{j+1}$-action on $\widetilde{S}^{2j}$. The orbit space S^{2j}_* is a $(\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_j)$-fold branched cover of S^{2j}, thought of as the unit sphere in the product orbifold. Thus, S^{2j}_* is homeomorphic to S^{2j}. The image of the fixed set of $\mathbb{Z}/a_{j+1}$ in S^{2j}_* is $\widetilde{S}^{2j-2}$. Therefore,

$$\begin{aligned}
\chi_{j+1} &= \chi(\widetilde{S}^{2j}) \\
&= a_{j+1}[\chi(S^{2j}_*) - (1 - (a_{j+1})^{-1})\chi(\widetilde{S}^{2j-2})] \\
&= a_{j+1}[2 - (1 - (a_{j+1})^{-1})\chi_j].
\end{aligned}$$

Assume by induction that the formula in (i) holds for $\ell = j$. Substituting this into the above formula for χ_{j+1}, we get that it also holds for $\ell = j + 1$; hence, the result. $\square$

Part (ii) of the above lemma shows that if $m = 2\ell - 1$, then the orbihedron V is not an orbifold. More generally, part (iii) implies that the same is true whenever $2\ell > m$ and m is odd. For in this case, setting $\ell' = \frac{m+1}{2}$, $\widetilde{S}^{m-1}$ is a finite sheeted covering of the $(\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_{\ell'})$-covering of $S(V)$.

Thus, when $2\ell > m$, to find V satisfying (A) through (E), we must assume that the dimension m of V is even, say $m = 2n$. In this case, the arrangement of the codimension two linear subspaces $H_1, \ldots, H_\ell$ in $\mathbb{R}^m$ is topologically equivalent to an arrangement of linear hyperplanes in general position in $\mathbb{C}^n$. Such configurations are studied in the next subsection.

Actions on Brieskorn Complete Intersections. We are now interested in the following situation. The underlying space of V is identified with $\mathbb{C}^n$. The singular set is the union of $n+k$ complex linear hyperplanes $H_1, \ldots, H_{n+k}$ in general position in $\mathbb{C}^n$. The hyperplane H_i is labeled by an integer a_i, $a_i \geq 2$. Such an arrangement of hyperplanes can be obtained by intersecting the coordinate hyperplanes in $\mathbb{C}^{n+k}$ with a generic n-dimensional complex linear subspace.

Given an $(n+k)$-tuple of integers $a = (a_1, \ldots, a_{n+k})$, $a_i \geq 2$, let G_a denote the product of cyclic groups, $G_a = \mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_{n+k}$, and let $\pi_a : \mathbb{C}^{n+k} \to \mathbb{C}^{n+k}$ be the map $(z_1, \ldots, z_{n+k}) \to ((z_1)^{a_1}, \ldots, (z_{n+k})^{a_{n+k}})$. Then π_a is the orbit map for the natural product G_a-action on $\mathbb{C}^{n+k}$. As usual, we denote the domain of π_a by $\widetilde{\mathbb{C}}^{n+k}$ and we continue to denote the range by $\mathbb{C}^{n+k}$.

Let $(w_1, \ldots w_{n+k})$ be linear coordinates in $\mathbb{C}^{n+k}$ so that $w_i = (z_i)^{a_i}$. Consider a generic n-dimensional linear subspace V_0 of $\mathbb{C}^{n+k}$ defined by k linear equations, $L_i(w) = 0$, $i = 1, \ldots, k$, where

$$L_i(w) = \sum_{j=1}^{n+k} \alpha_{ij} w_j$$

and where (α_{ij}) is a $k \times (n+k)$ matrix of rank k. Let $f_i(z) = L_i(\pi_a(z))$, i.e.,

$$f_i(z) = \sum_{j=1}^{n+k} \alpha_{ij} (z_j)^{a_j}$$

The inverse image of V_0 in $\widetilde{\mathbb{C}}^{n+k}$ is the *Brieskorn complete intersection*

$$Y^n(a,k) = \{z \in \mathbb{C}^{n+k} \mid f_1(z) = \cdots = f_k(z) = 0\}.$$

The orbifold V associated to the G_a-action on $Y^n(a,k)$ is of the form described above. For $k = 1$ the varieties $Y^n(a,k)$ were studied by Brieskorn [Bri], Milnor [Mi] and others; for $k > 1$ there are analogous results due to Hamm [Hm].

Put $\Sigma^{2n-1}(a,k) = Y^n(a,k) \cap S^{2n+2k-1}$, where $S^{2n+2k-1}$ denotes the unit sphere in $\widetilde{\mathbb{C}}^{n+k}$. Then $\Sigma^{2n-1}(a,k)$ and $Y^n(a,k)$ are G_a-stable subsets of $\widetilde{\mathbb{C}}^{n+k}$. Moreover, (i) $\Sigma^{2n-1}(a,k)$ is a smooth $(n-2)$-connected closed manifold of dimension $2n-1$, (ii) the spherical structure on the unit sphere in V_0 pulls back to give $\Sigma^{2n-1}(a,k)$ the structure of a piecewise spherical polyhedron, (iii) the variety $Y^n(a,k)$ is homeomorphic to the cone on $\Sigma^{2n-1}(a,k)$, i.e., $Y^n(a,k) \cong \Sigma^{2n-1}(a,k) \times [0,\infty]/\sim$, where $(x,0) \sim (y,0)$ for all $x, y \in \Sigma^{2n-1}(a,k)$, and (iv) $Y^n(a,k)$ is naturally a piecewise Euclidean polyhedron.

It follows from (i) that $\Sigma^{2n-1}(a,k)$ is simply connected provided $n > 2$. Thus, if $n > 2$, then the orbihedral fundamental group of V is the product of cyclic groups G_a, the universal cover of $S(V)$ is $\Sigma^{2n-1}(a,k)$, and the universal cover of V is $Y^n(a,k)$. In this case, therefore, V satisfies (A), (B), and (C). As for (D), conditions on $a = (a_1, \dots, a_{n+k})$ such that $Y^n(a,k)$ is PL-homeomorphic to $\mathbb{R}^n$ are given by Hamm in [Hm]. On the other hand, by the lemma which follows, V never satisfies (E).

LEMMA 6.10. *If $n > 2$, the manifold $Y^n(a,k)$ can never satisfy $CAT(0)$.*

Proof. If $Y^n(a,k)$ satisfies $CAT(0)$, then the link of each stratum must satisfy $CAT(1)$. If $n > 2$, then for any $i \neq j \neq \ell$, $\operatorname{codim}_{\mathbb{R}}(H_i \cap H_j \cap H_\ell) > \operatorname{codim}_{\mathbb{R}}(H_i \cap H_j) = 4$. It follows that $H_i \cap H_j$ contains a stratum (of real codimension 4) which does not intersect any other H_ℓ. The link L of this stratum is $S(H_i^\perp \oplus H_j^\perp)(\cong S^3)$ with singular set two great circles $H_i \cap L$ and $H_j \cap L$. By Lemma 6.6 these circles are orthogonal and hence $H_i^\perp$ is orthogonal to $H_j^\perp$. But this must hold for every i,j, which is impossible (for dimension reasons) since there are more than n of the H_i's. □

Now suppose $n = 2$. The underlying space of $S(V)$ is S^3. Let $\pi : S^3 \to S^2$ be the Hopf map. This gives $S(V)$ the structure of a "Seifert fibered orbifold" (i.e., the singular set is a union of fibers in some Siefert fiber space structure). There are $k+2$ singular fibers labeled $a_1, \dots a_{k+2}$. This induces an orbifold structure on S^2 with $k+2$ branch points, and $\Sigma^3(a,k) \to S^2$ is a Siefert fiber space in the usual sense. The map π induces a short exact sequence

$$1 \to C \to \pi_1^{orb}(S(V)) \to \pi_1^{orb}(S^2) \to 1$$

where C is the center of $\pi_1^{orb}(S(V))$. It follows that if $\pi_1^{orb}(S(V))$ is a finite group, then so is $\pi_1^{orb}(S^2)$. As is well-known, the group $\pi_1^{orb}(S^2)$ is finite if and only if $k = 1$ and the triple of integers (a_1, a_2, a_3) satisfies the inequality: $a_1^{-1}+a_2^{-1}+a_3^{-1} > 1$. It turns out that this condition is also sufficient for the group $\pi_1^{orb}(S(V))$ to be finite. Moreover, it is well-known that in this case the universal cover of $\Sigma^3(a,k)$ (and hence of $S(V)$) is S^3 and that the group $\pi_1^{orb}(S(V))$ can be realized as a finite subgroup of $GL(2,\mathbb{C})$ generated by three complex reflections of order a_1, a_2 and a_3. We denote this complex reflection group by $G(a_1, a_2, a_3)$. The actual instances of (a_1, a_2, a_3) satisfying the inequality are: $(m,2,2)$, $(3,3,2)$, $(4,3,2)$ and $(5,3,2)$; the corresponding group $G(a_1, a_2, a_3)$ has order $4m^2$, 144, 576, and 3600, respectively. (See [C; p. 93].)

We have therefore shown that when $n = 2$, the orbifold V satisfies conditions (A) through (D) if and only if $k = 1$ and (a_1, a_2, a_3) satisfies the above inequality. Much of Chapter II of this paper is devoted to deciding when V satisfies (E) (see Theorem 9.1).

To summarize, if V is an orbifold associated to a collection of codimension two subspaces $H_1, \ldots, H_\ell$ in general position in $\mathbb{R}^m$, and V satisfies (A) through (E), then either $2\ell \leq m$ and $\pi_1^{orb}(V) = \mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_\ell$, or $m = 4$, $\ell = 3$, and $\pi_1^{orb}(V) = G(a_1, a_2, a_3)$ with $a_1^{-1} + a_2^{-1} + a_3^{-1} > 1$. Conversely, if V is one of these types, then V always satisfies (A) through (D) and the geometry of the singular set determines when it satisfies (E).

Conjecture. Suppose that V is an orbifold satisfying conditions (A) through (E). Then we conjecture that V is equivalent to an orthogonal product of irreducible linear orbifolds of the following four types.

(1) $V_0 \cong \mathbb{R}^1$, the singular set is empty and G is the trivial group.

(2) $V_0 \cong \mathbb{C}$, the singular set is the origin, and the group G is cyclic of order m.

(3) $V_0 \cong \mathbb{C}^2$, the singular set consists of three lines through the origin, and the group G is the complex reflection group $G(a_1, a_2, a_3) \subset GL(2, \mathbb{C})$.

(4) $V_0 \cong \mathbb{R}^n$, the singular set is isomorphic to the cone on the $(n-3)$-skeleton of an $(n-1)$-simplex, and the group G is the orientation-preserving subgroup of a real reflection group.

We have already observed that orbifolds of types (1)–(4) satisfy conditions (A) through (D) and that types (1) and (2) satisfy (E). In Chapter II we will determine when orbifolds of types (3) and (4) satisfy (E).

7. An Application: Orbifold Structures on Flat Tori. In this section we apply the results of the previous section to analyze orbifold structures on flat tori. Let Q be an orbifold whose underlying space is an n-dimensional flat torus T^n, and suppose that the singular set $\Sigma \subset T^n$ is a union $\Sigma = S_1 \cup \cdots \cup S_k$ of flat codimension 2 subtori of T^n. If $x \in S_{i_1} \cap \cdots \cap S_{i_\ell}$ is a singular point, and $T_x(\cong \mathbb{R}^n)$ is the tangent space to T^n at x, then the tangent space $T_x S_{i_j}$ to S_{i_j} at x is a linear subspace of codimension 2 in T_x. Let N_{i_j} denote the linear 2-plane orthogonal to $T_x S_{i_j}$. We say that $S_1, \ldots, S_k$ are in *general position* if whenever $x \in S_{i_1} \cap \cdots \cap S_{i_\ell}$, the subspaces $N_{i_1}, \ldots, N_{i_\ell}$ are linearly independent in T_x (or equivalently, if $2\ell \leq n$ and $T_x S_{i_1}, \ldots, T_x S_{i_\ell}$ are in general position in the sense of section 6). We say that $S_1, \ldots, S_k$ *intersect orthogonally* if whenever $x \in S_{i_1} \cap \cdots \cap S_{i_\ell}$, the subspaces $N_{i_1}, \ldots, N_{i_\ell}$ are mutually orthogonal.

The tangent space T_x inherits the structure of an orbifold with singular set $T_x S_{i_1} \cup \cdots \cup T_x S_{i_\ell}$ (where $S_{i_1}, \ldots, S_{i_\ell}$ are all the singular subtori containing x). In particular, if x is a generic point in S_i, then T_x splits (as an orbifold) into an orthogonal product $T_x = T_x S_i \times N_i$ of a trivial orbifold, $T_x S_i$, and a 2-dimensional orbifold on N_i of the type of our basic example 5.1. In particular, the local group G_x is a cyclic group $G_x \cong \mathbb{Z}/a_i$. More generally, if $S_1, \ldots, S_k$ are in general position, then it follows from Proposition 6.8 that the local group at a generic point in $S_{i_1} \cap \cdots \cap S_{i_\ell}$ is the product of cyclic groups $G_x \cong \mathbb{Z}/a_{i_1} \times \cdots \times \mathbb{Z}/a_{i_\ell}$, that the tangent orbifold T_x satisfies (A) through (D) and that T_x satisfies (E) if and only if $N_{i_1}, \dots, N_{i_\ell}$ are mutually orthogonal. As an immediate corollary we obtain the following theorem.

THEOREM 7.1. *Assume $S_1, \dots, S_k$ are in general position. Then Q has nonpositive curvature if and only if $S_1, \dots, S_k$ intersect orthogonally.*

Proof. By Theorem 5.3, Q has curvature ≤ 0 if and only if, for all $x \in Q_0$, the universal cover $\widetilde{S}_x$ of the unit sphere in the tangent orbifold satisfies $CAT(1)$. But this is equivalent to condition (E) for T_x. □

Remark 7.2. It is possible that Q could have curvature ≤ 0 even if $S_1, \dots, S_k$ are not in general position. For example, let $n = 4$. Suppose that the singular set of Q is $S_1 \cup S_2 \cup S_3$ where any two S_i, S_j (and hence all three) intersect in the single point x_0, and suppose that the local group at x_0 is the complex reflection group $G(p,q,r)$ defined in Section 6. As above, Q has nonpositive curvature if and only if $\widetilde{S}_{x_0}$ satisfies $CAT(1)$ (where $\widetilde{S}_{x_0}$ is the orbifold universal cover of the unit sphere $S_{x_0} \subset T_{x_0}$). The manifold $\widetilde{S}_{x_0}$ is a branched cover of S^3 with covering group $G(p,q,r)$. It will be shown in Theorem 9.1 that such a branched cover *can* (under appropriate hypotheses) satisfy $CAT(1)$.

It follows from Theorems 7.1 and 4.1 that if $S_1, \dots, S_k$ intersect orthogonally, then Q is a good, aspherical orbifold. That is, its universal cover is a contractible manifold. On the other hand, "good" and "aspherical" are topological conditions which may hold even if Q does not have nonpositive curvature. In fact, we will show that general position for $S_1, \dots, S_k$ is sufficient to guarantee Q good and a slightly stronger condition, "complete general position", suffices to give Q aspherical. The technique is to embed Q in an orbifold Q' with underlying space T^{n+m} and singular set a union of codimension two subtori intersecting orthogonally. For this, we use the following orthogonalization procedure.

LEMMA 7.3. *Let $\{v_1, \dots, v_k\}$ be a finite set of vectors in $\mathbb{R}^n$. Then $\exists m \geq 0$ and vectors $\{w_1, \dots, w_k\} \in \mathbb{R}^{n+m}$ such that*

(i) $w_1, \dots, w_k$ *are mutually orthogonal,*

(ii) $\pi_n(w_i) = v_i$ *where* $\pi_n : \mathbb{R}^{n+m} \to \mathbb{R}^n$ *is projection on the first n factors, and*

(iii) if $v_1, \dots, v_k$ *are rational (i.e. lie in* $\mathbb{Q}^n$*) then so are* $w_1, \dots, w_k$.

Proof. The proof proceeds by induction on k. For $k = 1$, take $m = 0$, $w_1 = v_1$. Now consider a set $\{v_1, \dots, v_{k+1}\}$ and assume we have found an orthogonal set $\{w_1, \dots, w_k\}$ in $\mathbb{R}^{n+m}$ with $\pi_n(w_i) = v_i$. Let $f_1, \dots, f_k$ denote the last k standard basis vectors in $\mathbb{R}^{n+m+k}$. Set $w'_i = w_i + f_i$, $i = 1, \dots, k$ and $w'_{k+1} = v_{k+1} - \sum_{i=1}^{k} (v_{k+1} \cdot w_i) f_i$. Then $\{w'_1, \dots, w'_{k+1}\}$ satisfies the requirements of the lemma. □

We return to the case of an orbifold structure Q on T^n with singular set $\Sigma = S_1 \cup \cdots \cup S_k$. Let $p : \mathbb{R}^n \to \mathbb{R}^n/\mathbb{Z}^n = T^n$ be the projection. Then for each S_i, there is a unique linear codimension two subspace P_i in $\mathbb{R}^n$, defined over $\mathbb{Q}$, such that S_i is the image of a $\mathbb{Z}^n$-translate of P_i. We say S_i and S_j are *parallel* if $P_i = P_j$. Reordering if necessary, assume that $P_1, \dots, P_\ell$ are the *distinct* P_i's (i.e. $\{S_1, \dots, S_\ell\}$ is a maximal nonparallel subset of $\{S_1, \dots, S_k\}$). We say that $S_1, \dots, S_k$ are in *complete general position* if $2\ell \leq n$ and $P_1, \dots, P_\ell$ are in general position.

Example 7.4. To see that complete general position is stronger that general position, consider the following planes in $\mathbb{R}^4$

$$\begin{aligned} P_1 &= \{(x_1, x_2, x_3, x_4) \mid x_1 = x_2 = 0\} \\ P_2 &= \{(x_1, x_2, x_3, x_4) \mid x_3 = x_4 = 0\} \\ P_3 &= \{(x_1, x_2, x_3, x_4) \mid (x_1 - x_3) = (x_2 - x_4) = 0\}. \end{aligned}$$

Let $S_1, S_2, S_3 \subset T^4$ be the images of P_1, P_2 and $P_3 + (\frac{1}{2}, 0, 0, 0)$, respectively. Then S_1, S_2, S_3 are in general position, but not complete general position.

THEOREM 7.5. *Let Q be an orbifold with underlying space T^n and singular set $\Sigma = S_1 \cup \cdots \cup S_k$, a union of flat codimension 2 subtori. If $S_1, \dots, S_k$ are in general position, then Q is good. If $S_1, \dots, S_k$ are in complete general position, then Q is also aspherical.*

Proof. Let $P_1, \dots, P_\ell$ be as above, that is, a collection of distinct codimension 2 subspaces of $\mathbb{R}^n$ such that each S_i is the image of a translate of some P_j under the projection $j : \mathbb{R}^n \to T^n$. Set $N_i = P_i^\perp$ and choose a basis $\{v_{i1}, v_{i2}\}$ for N_i. Apply Lemma 7.3 to the set $\{v_{11}, v_{12}, \dots, v_{\ell 1}, v_{\ell 2}\}$ to get a mutually orthogonal set $\{w_{11}, w_{12}, \dots, w_{\ell 1}, w_{\ell 2}\}$, of vectors in $\mathbb{R}^{n+m}$ such that $\pi_n(w_{ij}) = v_{ij}$. Let N'_i be the subspace generated by w_{i1}, w_{i2} and let P'_i be the orthogonal complement of N'_i in $\mathbb{R}^{n+m}$. By definition, $P'_1, \dots, P'_\ell$ intersect orthogonally. Moreover, viewing $\mathbb{R}^n$ as the first n factors in $\mathbb{R}^{n+m}$ and $\mathbb{R}^m$ as the last m factors, we have

$$P'_i \cap \mathbb{R}^n = (N'_i + \mathbb{R}^m)^\perp = (N_i + \mathbb{R}^m)^\perp = P_i.$$

We now put an orbifold structure on $T^n \times \mathbb{R}^m$ as follows. Let $\overline{p} = p \times id : \mathbb{R}^n \times \mathbb{R}^m \to T^n \times \mathbb{R}^m$. For each S_j, write $S_j = p(P_{i_j} + x_j)$ and set $S'_j = \overline{p}(P'_{ij} + x_j)$. Let Q' be the orbifold structure on $T^n \times \mathbb{R}^m$ with singular set $\Sigma' = S'_1 \cup \cdots \cup S'_k$ and local group on S'_i = local group on $S_i = \mathbb{Z}/a_i$. Restricting Q' to $T^n \times \{0\}$, we get our original orbifold Q. Now by Theorem 7.1, Q' is a good, aspherical orbifold, that is, its universal cover is a contractible manifold. Restricting this cover to Q, we see that Q is also covered by a manifold, hence, Q is good.

Now suppose that $S_1, \dots, S_k$ are in complete general position, or in other words $2\ell \leq n$ and $P_1, \dots, P_\ell$ are in general position. In particular, $P = P_1 \cap \cdots \cap P_\ell$ has dimension $n - 2\ell$. Let $P^\perp$ denote the orthogonal complement of P in $\mathbb{R}^{n+m}$ and define $V = P^\perp \cap P'_1 \cap \cdots \cap P'_\ell$. We claim that (i) projection on $\mathbb{R}^m$ induces an isomorphism $\varphi : V \to \mathbb{R}^m$, and (ii) $P'_i = P_i \oplus V$. To see (i), recall that $P'_i \cap \mathbb{R}^n = P_i$, hence $(P'_1 \cap \cdots \cap P'_\ell) \cap \mathbb{R}^n = P_1 \cap \cdots \cap P_\ell = P$, so $V \cap \mathbb{R}^n = P^\perp \cap P = \{0\}$. Thus $\varphi : V \to \mathbb{R}^m$ is an injection. On the other hand, codim $P^\perp = \dim P = n - 2\ell$ and $\operatorname{codim}(P'_1 \cap \cdots \cap P'_\ell) = \dim(N'_1 + \cdots + N'_\ell) = 2\ell$, so codim $V \leq n$ and $\dim V \geq m$. Combining these observations, we see that $\dim V = m$ and φ is an isomorphism. This proves (i). It follows that $\mathbb{R}^{n+m} = \mathbb{R}^n \oplus V$ and hence $P'_i = (\mathbb{R}^n \cap P'_i) \oplus V = P_i \oplus V$. This proves (ii).

We now show that Q' is isomorphic as an orbifold to $Q \times \mathbb{R}^n$ where $\mathbb{R}^n$ is viewed as an orbifold with all local groups trivial. Define a linear automorphism

$$\phi : \mathbb{R}^n \oplus \mathbb{R}^m \to \mathbb{R}^n \oplus \mathbb{R}^m \qquad (x, y) \mapsto (x, 0) + \varphi^{-1}(y).$$

Clearly ϕ commutes with the $\mathbb{Z}^n$-action on $\mathbb{R}^n$ and hence, descends to homeomorphism

$$\hat{\phi} : T^n \times \mathbb{R}^m \to T^n \times \mathbb{R}^m.$$

Now $\phi(P_i \oplus \mathbb{R}^m) = P_i \oplus V = P'_i$, so $\hat{\phi}(S_i \times \mathbb{R}^m) = S'_i$. Thus, $\hat{\phi}$ induces a strata-preserving homeomorphism from the singular set $\Sigma \times \mathbb{R}^m$ of $Q \times \mathbb{R}^m$ to the singular set Σ' of Q'. Moreover, by definition, the local groups agree under this homeomorphism, hence $\hat{\phi}$ determines an isomorphism of orbifolds $Q \times \mathbb{R}^m \cong Q'$. Now Q' (and hence $Q \times \mathbb{R}^m$) has a contractible universal cover. But the universal cover of $Q \times \mathbb{R}^m$ is $\widetilde{Q} \times \mathbb{R}^m$ where $\widetilde{Q}$ is the universal cover of Q. It follows that $\widetilde{Q}$ is contractible. □

Remark 7.6. The referee has pointed out the following remarks concerning Theorem 7.5.

If $S_1, \dots, S_k$ are in general position, then Q actually has a *finite* covering which is a manifold. To see this, first consider the case where all the S_i's are parallel. In this case, Q is the product of an $(n-2)$-torus and an orbifold O with underlying space T^2. The latter admits a hyperbolic structure and hence has a finite manifold cover determined by a homomorphism of $\pi_1^{\mathrm{orb}}(O)$ into a finite group.

For the general case, we can consider the collection of those S_i's, say $S_1, \dots,$ S_r which are parallel to a fixed plane P_j. Pulling the orbifold structure on Q back to the finite cover described above, gives a finite orbifold cover $Q_j \to Q$ which is nonsingular over $S_1, \dots, S_r$. This cover is determined by a homomorphism $h_j : \pi_1^{\mathrm{orb}}(Q) \to G_j$ into a finite group G_j which is injective on the local groups for $S_1, \dots, S_r$. Doing this for each of the planes $P_1, \dots, P_\ell$, we obtain a homomorphism

$$h = h_1 \times \cdots \times h_\ell : \pi_1^{\mathrm{orb}}(Q) \to G_1 \times \cdots \times G_\ell$$

which is injective on all the local groups of Q (since all local groups are direct products of the local groups for the S_i's). This homomorphism determines a finite cover of Q which is a manifold.

In particular, in dimension $n = 3$, the circles S_i are in general position if and only if they are mutually disjoint and in complete general position if and only if they are disjoint and parallel. In the latter case, Q is the product of a circle and a 2-dimensional hyperbolic orbifold. In the former case, Q can be fibered over the circle with each fiber an orbifold on T^2. The finite manifold cover of Q is then a surface bundle over S^1 and, as such, is aspherical. Thus in dimension 3, general position suffices to imply asphericity.

Finally, we close this section with an example which demonstrates that in dimensions $n > 3$, general position alone is not sufficient to guarantee Q aspherical.

Example 7.7. Let $a_1, \dots, a_n$ be integers ≥ 2. Let $\varphi : \widetilde{\mathbb{C}}^n \to \mathbb{C}^n$ denote the branched cover of $\mathbb{C}^n$ defined by $(z_1, \dots, z_n) \mapsto (z_1^{a_1}, \dots, z_n^{a_n})$ with metric on $\widetilde{\mathbb{C}}^n$ pulled back from the standard (flat) metric on $\mathbb{C}^n$. The branch locus of this covering is the union of the coordinate hyperplanes in $\mathbb{C}^n$ and the covering transformation group is $\mathbb{Z}/a_1 \times \cdots \times \mathbb{Z}/a_n$. Let $w_1, \dots, w_n$ be a system of coordinates in $\mathbb{C}^n$ such that $z_i^{a_i} = w_i$. Let V be the complex hyperplane in $\mathbb{C}^n$ defined by $w_1 + \cdots + w_n = \frac{1}{2}$ and consider the branched cover

$$\widetilde{V} = \varphi^{-1}(V) = \left\{(z_1, \dots, z_n) \in \widetilde{\mathbb{C}}^n \mid z_1^{a_1} + \cdots + z_n^{a_n} = \tfrac{1}{2}\right\}.$$

The branch set in V consists of n linear subspaces, $P_i = \{(w_1, \dots, w_n) \in V \mid w_i = 0\}$, $i = 1, \dots, n$, of (real) codimension 2. It is easy to see that $P_1, \dots, P_n$ are in general position, but not complete general position. Let $\Delta \subset V$ be the geodesic $(n-1)$-simplex in $\mathbb{C}^n$ with vertices $e_1 = (\frac{1}{2}, 0, \dots, 0)$, $e_2 = (0, \frac{1}{2}, \dots, 0) \dots, e_n = (0, \dots, 0, \frac{1}{2})$, and let $\widetilde{\Delta} = \varphi^{-1}(\Delta) \subset \widetilde{V}$. By [Mi], we know that $\widetilde{\Delta}$ is a deformation retract of $\widetilde{V}$ and is homotopy equivalent to a nontrivial wedge of $(n-1)$-spheres. (The number of spheres is $\prod_{i=1}^n (a_i - 1)$.) In particular, $\widetilde{\Delta}$ gives rise to nontrivial elements in $\pi_{n-1}(\widetilde{V})$.

To see that the same phenomenon occurs in branched covers of tori, choose an isometry of V with $\mathbb{R}^{2n-2}$. Then the translation action of $\mathbb{Z}^{2n-2}$ on $\mathbb{R}^{2n-2}$ gives rise to a $\mathbb{Z}^{2n-2}$-action on V so that $V/\mathbb{Z}^{2n-2}$ is a flat torus. Let Q be the orbifold structure on $V/\mathbb{Z}^{2n-2}$ whose singular set Σ is the image of the planes $P_1, \ldots, P_n$, and whose local group at a (generic) point $x \in P_i$ is $\mathbb{Z}/a_i$. Pulling Q back to V gives an orbifold structure Q_V on V whose singular set is the union of all the $\mathbb{Z}^{2n-2}$-translates of $P_1, \ldots, P_n$. Now if $\widetilde{Q}$ is the universal cover of Q, then $\widetilde{Q}$ is a manifold (by Theorem 7.5) and there is a factorization of the covering map

$$\begin{array}{ccccccc} \widetilde{Q} & \longrightarrow & \widetilde{V} & \xrightarrow{\varphi} & V & \longrightarrow & V/\mathbb{Z}^{2n-2}. \\ & & \cup & & \cup & & \\ & & \widetilde{\Delta} & \longrightarrow & \Delta & & \end{array}$$

Since Δ has diameter $\frac{\sqrt{2}}{2}$ and is bounded by $P_1, \ldots, P_n$, it does not intersect any nontrivial $\mathbb{Z}^{2n-2}$-translate of $P_1, \ldots, P_n$. It follows that the inverse image of Δ in $\widetilde{Q}$ is a disjoint union of copies of $\widetilde{\Delta}$. Since $\widetilde{\Delta}$ generates nontrivial elements in $\pi_{n-1}(\widetilde{V})$, its lifts to $\widetilde{Q}$ must generate nontrivial elements in $\pi_{n-1}(\widetilde{Q})$. We conclude that for $n > 2$, Q is not aspherical. In particular, this means that there is no Riemannian metric on the underlying torus $V/\mathbb{Z}^{2n-2}$ so that the induced metric on $\widetilde{Q}$ is nonpositively curved. Odd dimensional examples are obtained by crossing Q with S^1.

Chapter II: Real and Complex Reflection Groups. From Chapter I, we know that if Q is an orbifold of nonpositive curvature, then the local models for Q must be orbifolds V satisfying conditions (A) through (E) of Section 6. As discussed in Section 6, primary examples of V satisfying (A) through (D) are provided by certain real and complex reflection groups acting on $\mathbb{R}^n$. In this chapter we determine when these examples satisfy (E).

Assume conditions (A) through (D). Let G be the orbifold fundamental group of V and let $S(V)$ be the restriction of V to the unit sphere in V_0 ($\cong \mathbb{R}^n$). Then the universal cover $\widetilde{S}(V)$ of $S(V)$ is a branched cover of S^{n-1} with covering group G. Condition (E) is equivalent to the condition that $\widetilde{S}(V)$ satisfy $CAT(1)$. Our goal in this chapter is to analyze branched covers of spheres with covering group an orientation-preserving subgroup of a real reflection group or a complex reflection group of the form $G(p,q,r)$. In particular, we want to determine when the piecewise spherical polyhedron $\widetilde{S}(V)$ satisfies $CAT(1)$. By Lemma 1.3, this is the case if and only if $\widetilde{S}(V)$ has curvature ≤ 1 and systole $\geq 2\pi$. (By Theorem 3.1, $\widetilde{S}(V)$ has curvature ≤ 1 if and only if the link of every point in $\widetilde{S}(V)$ satisfies $CAT(1)$. In most of our examples, the curvature condition will be obvious; the main work will involve the systole condition.)

8. Branched Covers of S^n Associated to Real Reflection Groups. Let G be the orientation-preserving subgroup of a real reflection group of rank $n+1$.

Then, as discussed in Section 6, the action of G on the n-sphere gives rise to a branched cover

$$\pi : \widetilde{S}^n \to \widetilde{S}^n/G = S^n$$

where S^n is given the standard spherical metric d and $\widetilde{S}^n$ has the piecewise spherical metric obtained by pulling back the metric on S^n.

If $n = 1$, G is a cyclic group, $G = \mathbb{Z}/p$, and $\widetilde{S}^1$ is a circle of circumference $2\pi p$. Clearly this satisfies $CAT(1)$. So the first interesting case is the case $n = 2$. The finite real reflection groups (or Coxeter groups) of rank 3 are the triangle groups

$$T(p,q,r) = \langle s_1, s_2, s_3 \mid s_i^2 = (s_1s_2)^p = (s_1s_3)^q = (s_2s_3)^r = 1\rangle$$

with $\frac{1}{p}+\frac{1}{q}+\frac{1}{r} > 1$. The orientation-preserving subgroup $G \subset T(p,q,r)$ is generated by the three rotations, $g_1 = s_1s_2$, $g_2 = s_1s_3$, $g_3 = s_2s_3$, of orders p, q, and r. For such G, the branch locus in S^2 consists of three points x_p, x_q, x_r with local groups $\langle g_1\rangle = \mathbb{Z}/p$, $\langle g_2\rangle = \mathbb{Z}/q$, and $\langle g_3\rangle = \mathbb{Z}/r$, respectively.

THEOREM 8.1. *Let G be the orientation-preserving subgroup of $T(p,q,r)$ and let $\{x_p, x_q, x_r\} \subset S^2$ be the branch set of $\widetilde{S}^2 \to \widetilde{S}^2/G = S^2$. Then $\widetilde{S}^2$ satisfies $CAT(1)$ if and only if*

(i) $d(x_p,x_q) + d(x_q,x_r) + d(x_p,x_r) = 2\pi$, *and*

(ii) $d(x_i,x_j) \geq \frac{\pi}{k}$ *for all* $\{i,j,k\} = \{p,q,r\}$.

Remark 8.2. Condition (i) is equivalent to the conditions that x_p, x_q, x_r all lie on some great circle in S^2 and that they do not all lie on any half of that circle.

Proof of Theorem 8.1. We first observe that $\widetilde{S}^2$ always has curvature ≤ 1. This is true by Theorem 3.1 since for any point $\tilde{x} \in \widetilde{S}^2$, the link of $\tilde{x}$ in $\widetilde{S}^2$ is just a circle of circumference $2\pi s$, where s is the order of the local group at $\tilde{x}$. In particular, these links all satisfy $CAT(1)$. It remains to prove that the systole of $\widetilde{S}^2$ is $\geq 2\pi$ if and only if condition (i) and (ii) of Theorem 8.1 hold.

Let Δ denote the geodesic triangle in S^2 joining x_p, x_p, and x_r. Then Δ divides S^2 into two (closed) regions which we will refer to as "the black and the white regions" of S^2. One of these regions, say the black one, is a spherical triangle with all angles $\leq \pi$. The white region may then be regarded as a spherical triangle with all angles $\geq \pi$. The inverse images of these triangles under the projection $\widetilde{S}^2 \to S^2$ triangulate $\widetilde{S}^2$ into $2|G|$ regions, half of the regions isometric to the black triangle and half to the white one. The vertices of this triangulation are the inverse images of the singular points x_p, x_q, x_r.

We first verify that conditions (i) and (ii) of the theorem are necessary to have $sys(\widetilde{S}^2) \geq 2\pi$. Note that at a vertex $\tilde{x}_p$ over x_p, we have p white triangles and p black triangles meeting alternately:

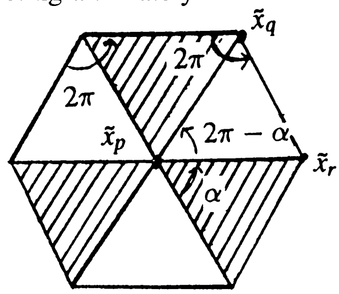

In particular, if γ_1 is the geodesic segment from $\tilde{x}_q$ to $\tilde{x}_p$ and γ_2 the geodesic segment from $\tilde{x}_p$ to $\tilde{x}_r$, then the angle between them is $\geq \pi$. Hence, by Lemma 2.3, $\gamma = (\gamma_1, \gamma_2)$ is a local geodesic at $\tilde{x}_p$. By the same argument, the boundary of any white region is a closed local geodesic in $\widetilde{S}^1$. But the length of this boundary is the perimeter, $\ell(\Delta)$, of the triangle Δ in S^2, namely

$$\ell(\Delta) = d(x_p, x_q) + d(x_q, x_r) + d(x_p, x_r).$$

This shows that condition (i) is necessary. For condition (ii), we note that the simplicial link of the vertex $\tilde{x}_p$ in $\widetilde{S}^2$ consists of $2p$ segments each connecting some $\tilde{x}_q$ to some $\tilde{x}_r$. (By simplicial link we mean the boundary of the region in Figure 1.) In particular, each segment has length $d(x_q, x_r)$. Observing the angles at each vertex, it is clear that this link is a closed local geodesic in $\widetilde{S}^2$. Thus, for $\widetilde{S}^2$ to be large we must have $2p \cdot d(x_q, x_r) \geq 2\pi$. This proves the necessity of condition (ii).

Now assume conditions (i) and (ii) hold and suppose γ is a closed local geodesic in $\widetilde{S}^2$. We first reduce to the case where γ lies in the 1-skeleton of the triangulation of $\widetilde{S}^2$ by black and white regions. As remarked above, condition (i) implies that Δ is a great circle in S^2 and hence, each region of $\widetilde{S}^2$ is isometric to a hemisphere of S^2. Thus, if γ enters the interior of any region, then γ contains a segment γ_1 of length π lying entirely in that region. In this case, the projection of γ to S^2 is a closed broken geodesic which contains a pair of antipodal points. Clearly, any such closed broken geodesic has length $\geq 2\pi$, hence, γ has length $\geq 2\pi$.

Thus, we may assume that $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n)$ is an edge path in the triangulation of $\widetilde{S}^2$. Let $\{i, j, k\} = \{p, q, r\}$. Say a vertex is of *type i* if it lies over x_i and say an edge is of *type i* if its endpoints are of types j and k. Note that any edge of type i has length $d(x_j, x_k)$. If γ contains edges of all three types, then condition (i) implies that length $\gamma \geq 2\pi$. If γ contains exactly two types of edges, then both types of edge must occur at least twice (since γ is closed), so again condition (i) (and Remark 8.2) implies that length $\gamma \geq 2\pi$. Finally, it is easy to see that the shortest closed edge paths containing only type i edges are the simplicial links of type i vertices. These contain exactly $2i$ edges of type i, so condition (ii) implies length $\gamma \geq 2\pi$. □

Next we consider the case of an orientation preserving subgroup G of a real reflection group of rank $n+1 \geq 3$. Our results here will depend on a general theorem on piecewise spherical complexes (Theorem 8.4 below) which is a slight generalization of a result of Gromov ([G, Lemma 4.2.C]).

We preface this theorem with some remarks about "proper" and "almost proper" spherical triangulations. By a *convex n-simplex* in S^n we mean a (nondegenerate) topological n-simplex $\sigma \subset S^n$ such that every codimension 1 face τ of σ lies on a geodesic $(n-1)$-sphere S_τ^{n-1} in S^n and the interior of σ is a single component of

$$S^n \backslash \bigcup_\tau S_\tau^{n-1}.$$

If the spheres S_τ^{n-1} are all distinct, we say that σ is *strictly convex*. For example, a 1-simplex is convex if and only if it has length $\leq \pi$, and it is strictly convex if and only if it has length $< \pi$.

Generally, when considering piecewise spherical complexes, we require our simplices to be strictly convex. This is no real restriction as long as we can pass to subdivisions. However, in the discussion that follows, the size of a simplex is crucial, so subdivision is not always an option. We must, therefore, allow for geodesic simplices which are convex, but not strictly convex.

Definition 8.3. Let L be a geodesic metric space. A triangulation $\mathcal{T}$ of L is an *almost proper* (resp. *proper*) spherical triangulation if every n-simplex of $\mathcal{T}$ is isometric to a convex (resp. strictly convex) n-simplex in S^n.

We assume all of our simplicial complexes are finite dimensional. Recall that a *flag complex* is a simplicial complex $\mathcal{T}$ such that if $\{v_0, \dots, v_k\}$ is a set of vertices in $\mathcal{T}$ which are pairwise joined by edges, then $\{v_0, \dots, v_k\}$ span a k-simplex in $\mathcal{T}$.

THEOREM 8.4. *Let L be a piecewise spherical polyhedron and let $\mathcal{T}$ be a triangulation of L. Assume*

(C1) $\mathcal{T}$ is a flag complex,

(C2) $\mathcal{T}$ is an almost proper spherical triangulation

(C3) if τ_1, τ_2 are disjoint faces of a simplex σ in $\mathcal{T}$, then $d(\tau_1, \tau_2) \geq \frac{\pi}{2}$.

Then L satisfies $CAT(1)$.

Remark. If we assume that $\mathcal{T}$ is *proper*, then (C3) is equivalent to the condition that the length of every 1-simplex in $\mathcal{T}$ is $\geq \frac{\pi}{2}$ (cf. Lemma 8.12 below).

Before proving this theorem we will need the following lemmas.

LEMMA 8.5. *Let $L, \mathcal{T}$ be as above. If $\mathcal{T}$ satisfies (C1), (C2), and (C3), then so does the induced triangulation on $L_\sigma = Link(\sigma, \mathcal{T})$ for every simplex σ of $\mathcal{T}$.*

Proof. Clearly (C1) and (C2) are inherited by links. It remains to consider (C3). We proceed by induction on the dimension of L. If $\dim L = 1$, then links are discrete so there is nothing to prove. Suppose $\dim L = k > 1$. Consider first the case when $\sigma = v$, a vertex of L. Let

$$\begin{aligned} \overline{B}(v) &= \{x \in L \mid d(x, v) \leq \tfrac{\pi}{2}\} \\ st(v) &= \text{ union of simplices containing } v \end{aligned}$$

Then (C3) implies that $\overline{B}(v) \subset st(v)$, hence we may identify L_v, the link of v in L, with the set

$$L_v = \left\{x \in L \mid d(x, v) = \tfrac{\pi}{2}\right\}.$$

In particular, a simplex $\overline{\mu}$ of L_v is of the form

$$\overline{\mu} = \left\{x \in \mu \mid d_\mu(x, v) = \tfrac{\pi}{2}\right\}$$

where μ is a simplex of L containing v. Two faces $\overline{\tau}_1, \overline{\tau}_2$ of $\overline{\mu}$ are disjoint if they correspond to faces τ_1, τ_2 of μ whose intersection is precisely v. Let $x_1 \in \overline{\tau}_1, x_2 \in \overline{\tau}_2$. We must show $d(x_1, x_2) \geq \frac{\pi}{2}$. Consider the spherical triangle $\Delta(v, y_1, y_2)$ in μ

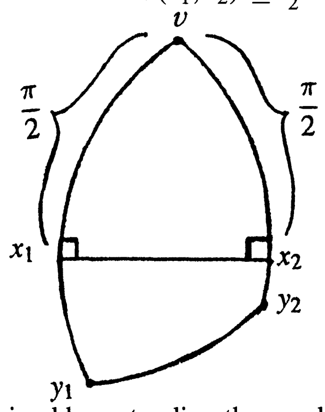

where y_1, y_2 are determined by extending the geodesic segments $\overline{vx_i}$ to the boundary of $st(v)$. Thus, y_i lies in the face τ_i' of τ_i opposite v. Since $\tau_1 \cap \tau_2 = \{v\}$, $\tau_1 \cap \tau_2' = \tau_1' \cap \tau_2 = \emptyset$. Thus, condition (C3) on L implies that $d(y_1, x_2) \geq \frac{\pi}{2}$ and $d(y_2, x_1) \geq \frac{\pi}{2}$. Now $y_1 \neq y_2$ so at most one of y_1, y_2 is antipodal to v. Say $d(v, y_1) < \pi$, so $d(x_1, y_1) < \frac{\pi}{2}$. Then the spherical triangle $\Delta(x_1, y_1, x_2)$ is a right triangle with one leg, $\overline{x_1 y_1}$, of length $< \frac{\pi}{2}$ and hypotenuse, $\overline{y_1 x_2}$, of length $\geq \frac{\pi}{2}$. By an easy exercise in spherical geometry, it follows that $d(x_1, x_2) \geq \frac{\pi}{2}$. This completes the case when σ is a vertex.

If $\dim \sigma > 0$, choose a vertex v of σ. Let $\overline{\sigma} = \mathrm{Link}(v, \sigma) \subset L_v$. Then $\mathrm{Link}(\sigma, L) = \mathrm{Link}(\overline{\sigma}, L_v)$. But we have already shown that the $(k-1)$-dimensional complex L_v satisfies (C3). Hence, by induction, the same holds for $\mathrm{Link}(\overline{\sigma}, L_v)$. □

LEMMA 8.6. *If $L, \mathcal{T}$ satisfy (C1), (C2), and (C3), then so does the induced triangulation on ΣL.*

Proof. It is clear that (C1) and (C2) hold in ΣL. To see that (C3) holds in ΣL, note that every simplex of ΣL lies in the (upper or lower) cone on L, so it suffices to show that $\mathrm{cone}(L) = L * v_0$ satisfies (C3). For a k-simplex σ in L, let $\hat{\sigma}$ denote the corresponding $(k+1)$-simplex, $\hat{\sigma} = \sigma * v_0$, in the cone. We must show that $d_{\hat{\sigma}}(\tau_1, \hat{\tau}_2) \geq \frac{\pi}{2}$ whenever τ_1, τ_2 are disjoint faces of σ. Let $x \in \tau_1$, $y \in \hat{\tau}_2$. If $y = v_0$, then $d(x, y) = \frac{\pi}{2}$ by definition. If $y \neq v_0$, let $z \in \tau_2$ be the projection of y to L and consider the spherical right triangle $\Delta(x, y, z) \subset \hat{\sigma}$.

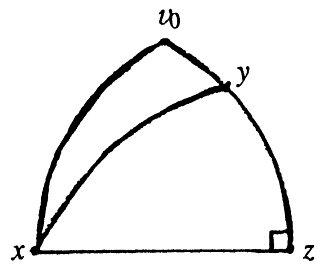

By condition (C3) on L, $d(x, z) \geq \frac{\pi}{2}$ and clearly $d(y, z) < \frac{\pi}{2}$. It follows that $d(x, y) \geq \frac{\pi}{2}$. □

The next lemma is a variation on Lemma 9.8 of [M]. The proof is essentially the same.

LEMMA 8.7. *Suppose $L, \mathcal{T}$ satisfy (C2) and (C3). Let v be a vertex of L and let γ be a local geodesic segment in $st(v)$ which intersects $\partial(st(v))$ precisely in its endpoints. Then γ has length $\geq \pi$.*

Proof. We proceed by induction on $\dim L$. If $\dim L = 1$, then L is made up of edges of length $\geq \frac{\pi}{2}$. In this case the lemma is clear. Let $\dim L = n > 1$, and let v be a vertex of L. Let γ be as in the lemma and write $\gamma = (\gamma_1, \ldots, \gamma_r)$ where γ_i is a segment of γ contained in a single simplex σ_i of $st(v)$. If γ passes through v, then by condition (C3), γ has length $\geq \pi$. So we may assume that γ does not pass through v. Consider the collection of rays in $st(v)$ emanating from v and passing through a point of γ. The union of these rays forms a piecewise spherical 2-complex with 2-simplices $\tau_i = \cup$ (rays from v to γ_i). We can map each τ_i isometrically onto a 2-simplex in S^2, $f : \tau_i \to S^2$ with $f(v)$ at the north pole and the common face of τ_i and τ_{i+1} (containing the endpoint of γ_i) mapping to a common face of $f(\tau_i)$ and $f(\tau_{i+1})$. Let $X = \cup f(\tau_i) \subset S^2$. The image of γ in X

is locally geodesic in X and hence also in S^2 (since $f(\gamma)$ intersects ∂X only at its endpoints). That is, $f(\gamma)$ is a segment of a great circle in S^2.

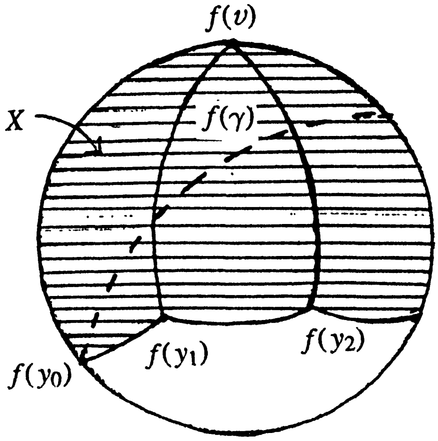

Let $x_0, \ldots, x_r$ denote the endpoints of $\gamma_1, \ldots, \gamma_r$. Let $y_0, \ldots, y_r$ be the radial projections of $x_0, \ldots, x_r$ to $\partial(st(v))$. To show that the length of $f(\gamma)$ (and hence the length of γ) is at least π, it suffices to show that the interior angle of X at $f(y_i)$ is $\geq \pi$, for $i = 1, \ldots, r-1$.

For this, consider $L_i = \operatorname{Link}(y_i, L)$. Let $\alpha = \operatorname{Link}(y_i, \tau_i)$ and $\beta = \operatorname{Link}(y_i, \tau_{i+1})$. Then (α, β) is a local geodesic in L_i and the interior angle of X at $f(y_i)$ is precisely the length of this geodesic. Now let ρ_i be the support of y_i (i.e. ρ_i is the unique simplex of L containing y_i in its interior), and let $k = \dim \rho_i$. Then

$$L_i = \operatorname{Link}(y_i, L) = \Sigma^k \operatorname{Link}(\rho_i, L_i)$$

so the natural triangulation on $\operatorname{Link}(\rho_i, L)$ induces a triangulation on L_i. By Lemmas 8.5 and 8.6, this triangulation satisfies (C2) and (C3). The simplex $v * \rho_i$ spanned by v and ρ_i represents a vertex $\overline{v}$ in $\operatorname{Link}(\rho_i, L)$ and hence also in L_i. (Note that $\overline{v}$ is not, in general, the tangent vector to the geodesic from y_i to v, but rather the orthogonal projection of this vector onto (tangent space to $\rho_i)^{\perp} = \operatorname{Link}(\rho_i, L)$.) The geodesic (α, β) lies in the star of $\overline{v}$ in L_i and has endpoints in $\partial(st(\overline{v}))$. Therefore, our induction hypothesis applies to $(\alpha, \beta) \subset st(\overline{v}) \subset L_i$ to give that length $(\alpha, \beta) \geq \pi$. □

Proof of Theorem 8.4. The proof is by induction on $\dim L$. If $\dim L = 1$, then L is a graph whose edges have length $\geq \frac{\pi}{2}$. The assumption that L is a flag complex means that this graph contains no cycles with less than 4 edges. Thus, L satisfies $CAT(1)$.

Suppose $\dim L = n > 1$. For a vertex v in L write

$$st(v) = \bigcup_{v \in \sigma} \sigma = \text{star of } v$$

$$ost(v) = \bigcup_{v \in \sigma} \text{interior}(\sigma) = \text{ open star of } v.$$

The collection of open stars, $\{ost(v)\}_{v \text{ a vertex}}$, forms an open cover of L and $ost(v) \cap ost(w) \neq \emptyset$ if and only if $w \in st(v)$. Let γ be a closed geodesic in L and let v be a vertex such that $\gamma \cap ost(v) \neq \emptyset$. Then by Lemma 8.7, $ost(v)$ contains (the interior of) a segment of γ of length $\geq \pi$. If γ does not lie entirely in $st(v)$, then it must intersect $ost(w)$ for some vertex $w \notin st(v)$. Again by Lemma 8.7, $ost(w)$ contains (the interior of) a segment of γ of length $\geq \pi$. But $ost(v) \cap ost(w) = \emptyset$ (since $w \notin st(v)$), hence γ has length $\geq 2\pi$.

It remains to consider the case where γ lies entirely in $st(v)$. By condition (C3), we can identify

$$L_v = \mathrm{Link}(v, L) = \left\{x \in st(v) \mid d_L(v, x) = \tfrac{\pi}{2}\right\}.$$

This identification extends to a natural map $f : st(v) \to \Sigma L_v$ which, restricted to any simplex of $st(v)$, is an isometry. (Though f is not, in general, an isometry on the whole space $st(v)$.) By Lemma 8.5, L_v satisfies the hypotheses of the theorem, so by induction, it satisfies $CAT(1)$. By Theorem A.10 of the appendix, ΣL_v also satisfies $CAT(1)$. The image of γ under f is a broken geodesic in ΣL_v with break points occurring only in $f(\partial\, st(v))$. In particular, if $\gamma \subset ost(v)$, $f(\gamma)$ is a closed local geodesic in ΣL_v and so has length $\geq 2\pi$. If γ intersects $\partial\, st(v)$, then by Lemma 8.7, $ost(v)$ contains a segment of γ of length π, so $f(ost(v))$ contains a local geodesic segment of $f(\gamma)$ of length π. But since ΣL_v satisfies $CAT(1)$, any local geodesic of length π is an actual geodesic. In other words, $f(\gamma)$ contains a pair of points x, y of distance π in ΣL_v. It follows that $f(\gamma)$, and hence γ, has length $\geq 2\pi$. □

We now return to the case of a branched cover

$$\widetilde{S}^n \to \widetilde{S}^n / G = S^n$$

in which G is the orientation-preserving subgroup of a real reflection group of rank $n + 1$. The singular set $\Sigma \subset \widetilde{S}^n/G$ is homeomorphic to the $(n-2)$-skeleton of an n-simplex.

THEOREM 8.8. *Let $n \geq 2$ and suppose that $\widetilde{S}^n$ satisfies $CAT(1)$. Then Σ lies on a great $(n-1)$-sphere $S_\Sigma^{n-1} \subset S^n$, and the triangulation $\mathcal{T}_\Sigma$ of S^n whose $(n-2)$-skeleton is Σ and $(n-1)$-skeleton is S_Σ^{n-1} is an almost proper spherical triangulation.*

Proof. As demonstrated in Section 5, $CAT(1)$ for $\widetilde{S}^n$ implies that the faces of Σ are locally geodesic, or in other words, each k-simplex σ of Σ lies on a great k-sphere in S^n. We first show, by induction on n, that Σ itself lies on a great $(n-1)$-sphere. For $n = 2$, this follows from Theorem 8.1.

Suppose $n > 2$. Let v be a vertex of Σ and $\tilde{v} \in \widetilde{S}^n$ a point over v. Then

$CAT(1)$ for $\widetilde{S}^n$ implies $CAT(1)$ for $\widetilde{L}_v = \mathrm{Link}(\tilde{v}, \widetilde{S}^n)$. But $\widetilde{L}_v$ is again a branched cover of the form

$$\widetilde{L}_v \to \widetilde{L}_v / G_v = \mathrm{Link}(v, S^n) = S^{n-1}$$

where G_v is the orientation-preserving subgroup of a rank n reflection group. By induction, therefore, the singular set $\Sigma_v \subset \widetilde{L}_v/G_v = S^{n-1}$ lies on a great $(n-2)$-sphere. It follows that the star of v in Σ lies on a great $(n-1)$-sphere of S^n. But every vertex of Σ is contained in $st(v)$. Hence Σ lies entirely in this $(n-1)$-sphere.

Now let S^{n-1}_Σ denote the $(n-1)$-sphere spanned by Σ. Then Σ determines a triangulation $\mathcal{T}_\Sigma$ of S^n as follows. The n-simplices Δ_1, Δ_2 of $\mathcal{T}_\Sigma$ are the two hemispheres of S^n determined by S^{n-1}_Σ. The $(n-1)$-simplices are the closures of the connected components of $S^{n-1}_\Sigma \backslash \Sigma$. The $(n-2)$-skeleton of $\mathcal{T}_\Sigma$ is Σ. We claim that $\mathcal{T}_\Sigma$ is almost proper. Again we proceed by induction on n. For $n = 2$, Σ consists of three points x_p, x_q, x_r lying on a great circle of S^2. By Theorem 8.1 and Remark 8.2, x_p, x_q, x_r divide this circle into three segments of length $\leq \pi$. It follows that $\mathcal{T}_\Sigma$ is almost proper.

Now suppose $n > 2$. Let σ be a k-simplex of $\mathcal{T}_\Sigma$. We know that the codimension 1 faces $\tau_0, \dots, \tau_k$ of σ lie on great $(k-1)$-spheres. If $k \geq 2$, the convexity of σ will follow if we show that the dihedral angles between the τ_i's are $\leq \pi$. For this, note that if τ_i and τ_j meet along the codimension 2 face $\tau = \tau_i \cap \tau_j$, then the angle between τ_i and τ_j is the length of the edge determined by σ in $\mathrm{Link}(\tau, S^n)$. By induction, the induced triangulation on this link is almost proper so every edge has length $\leq \pi$. We conclude that σ is convex provided $\dim \sigma \geq 2$. It remains to show that every 1-simplex in $\mathcal{T}$ has length $\leq \pi$. Suppose e is a 1-simplex of length $> \pi$. Since $n \geq 3$, e is contained in some 3-simplex σ in $\mathcal{T}_\Sigma$. If τ is a 2-dimensional face of σ, then, by the discussion above, τ is convex. But if e lies in τ, this is possible if and only if τ is a hemisphere with all vertices lying on the great circle containing e. It follows that *all four* vertices of σ lie on this great circle. Indeed, they lie on a segment of this great circle of length $< \pi$. But this is absurd for a nondegenerate 3-simplex. We conclude that the edges of $\mathcal{T}_\Sigma$ must have length $\leq \pi$. □

We are now ready to establish sufficient conditions for $CAT(1)$ to hold in $\widetilde{S}^n$.

THEOREM 8.9. *Let G be the orientation-preserving subgroup of a real reflection group W of rank $n+1$. Assume that the singular set $\Sigma \subset S^n = \widetilde{S}^n/G$ lies on a great $(n-1)$-sphere and that the faces of Σ are totally geodesic. Let $\mathcal{T}_\Sigma$ be the associated triangulation of S^n. If $\mathcal{T}_\Sigma$ satisfies (C2) and (C3), then $CAT(1)$ holds in $\widetilde{S}^n$.*

Proof. If $\mathcal{T}_\Sigma$ satisfies (C2) and (C3), then the same holds for the induced triangulation $\widetilde{\mathcal{T}}_\Sigma$ of $\widetilde{S}^n$. Combinatorially, $\widetilde{\mathcal{T}}_\Sigma$ is the Coxeter complex associated to

W and, as such, is well-known to be a flag complex (see eg. [Bro]; p. 29). It follows from Theorem 8.4 that $\widetilde{S}^n$ satisfies $CAT(1)$. □

Remark 8.10. For G as above, there always exist triangulations $\mathcal{T}_\Sigma$ satisfying (C2) and (C3). For example, take the radial projection of a regular n-simplex in $\mathbb{R}^n$ with center of mass at the origin, and let $\Sigma \subset S^{n-1}$ be the image of the $(n-2)$-skeleton. Then $\mathcal{T}_\Sigma$ satisfies (C2) and (C3).

We have seen (Theorem 8.4) that condition (C3) is not a necessary condition for $CAT(1)$ to hold in $\widetilde{S}^n$, at least in the case $n = 2$. However, there is one situation, namely the case of an abelian reflection group, in which the hypotheses of Theorem 8.9 are necessary and sufficient.

THEOREM 8.11. *Let $n \geq 2$. Suppose G_n is the orientation-preserving subgroup of*

$$W_n = \langle s_1, \ldots, s_{n+1} \mid s_i^2 = (s_i s_j)^2 = 1 \rangle \cong (\mathbb{Z}/2)^{n+1}.$$

Then $\widetilde{S}^n$ satisfies $CAT(1)$ if and only if Σ and the associated triangulation $\mathcal{T}_\Sigma$ satisfy the hypotheses of Theorem 8.9.

LEMMA 8.12. *Let L be a piecewise spherical polyhedron and $\mathcal{T}$ an almost proper triangulation of L. For any simplex σ in $\mathcal{T}$ and any point $x \in \sigma$, there exists a vertex $v \in \sigma$ such that $d(x, v) \leq \frac{\pi}{2}$.*

Proof. The proof is by induction on the dimension of σ. If $\dim \sigma = 1$, then σ is convex if and only if it has length $\leq \pi$. In this case, the lemma is clear. Suppose $\dim \sigma = n \geq 2$. Identify σ with an intersection of half-spaces in S^n. Let $x \in int\, \sigma$. Let y be the point on $\partial\sigma$ nearest to x. Clearly, $d(x, y) \leq \frac{\pi}{2}$. Let τ be the support of y, i.e., τ is the smallest simplex containing y. Then the geodesic from x to y is necessarily perpendicular to τ (since y is the closest point to x on $\partial\sigma$). By induction, there is a vertex $v \in \tau$ such that $d(v, y) \leq \frac{\pi}{2}$. The geodesic triangle with vertices x, y, v is a spherical triangle with a right angle at y. Both legs of this triangle have length $\leq \frac{\pi}{2}$. It follows that the hypotenuse has length $\leq \frac{\pi}{2}$. □

LEMMA 8.13. *Let $L, \mathcal{T}$ be as in Lemma 8.12. Then condition (C3) is equivalent to:*

(C3)′ *every 1-simplex in L and every 1-simplex in L_σ, $\sigma \in \mathcal{T}$, has length $\geq \frac{\pi}{2}$.*

Proof. By definition, (C3) implies that edges in L have length $\geq \frac{\pi}{2}$. By Lemma 8.5, (C3) for L implies ($C3$) for L_σ, so the same holds in L_σ. This shows that (C3) implies (C3)′.

Conversely, assume condition (C3)$'$. Suppose τ and τ^{op} are opposite faces of a simplex $\sigma \in \mathcal{T}$. We will show by induction on $\dim \sigma$ that $d(\tau, \tau^{op}) \geq \frac{\pi}{2}$. If σ is an edge, then τ and τ^{op} are the vertices of σ and $d(\tau, \tau^{op}) =$ length of $\sigma \geq \frac{\pi}{2}$ by (C3)$'$. Suppose $\dim \sigma > 1$. Let γ be a geodesic in σ connecting τ and τ^{op} of minimal length. If γ lies in the boundary of σ, then by the inductive hypothesis it has length $\ell(\gamma) \geq \frac{\pi}{2}$. So we may assume that the endpoints $x \in \tau$, $y \in \tau^{op}$ of γ lie in the interiors of τ and τ^{op}. It follows from the minimality of the length of γ that γ intersects τ and τ^{op} orthogonally. Now either τ or τ^{op} has dimension > 0, say $\dim \tau > 0$. By Lemma 8.12, we can choose a vertex $w \in \tau$ with $d(x, w) \leq \frac{\pi}{2}$. Consider the right spherical triangle spanned by x, y, w with right angle at x. The segment from y to w lies in the boundary of σ, so by induction, $d(y, w) \geq \frac{\pi}{2}$. Now if $d(x, w) < \frac{\pi}{2}$, an easy exercise in spherical geometry shows that $\ell(\gamma) = d(x, y) \geq \frac{\pi}{2}$. If, on the other hand, $d(x, w) = \frac{\pi}{2}$, then $d(y, w) = \frac{\pi}{2}$ and the angle at y is also a right angle. In this case, γ can be viewed as a geodesic segment in $L_w = \mathrm{Link}(w, L)$ connecting opposite faces of $\sigma_w = \mathrm{Link}(w, \sigma)$. By induction, (C3) holds in σ_w, so $\ell(\gamma) \geq \frac{\pi}{2}$. □

Remark 8.14. If an n-simplex σ is strictly convex, then it is contained in some *open* hemisphere of S^n. Thus, if we assume $\mathcal{T}$ is proper, in Lemma 8.12 we can replace all the inequalities by strict inequalities, and in Lemma 8.13 we can assume $d(x, w) < \frac{\pi}{2}$. The proof of Lemma 8.13 then shows that if $\mathcal{T}$ is proper, then (C3) is equivalent to the condition that every 1-simplex in L has length $\geq \frac{\pi}{2}$.

Proof of Theorem 8.11. In light of Theorems 8.8 and 8.9, it remains only to show that if $\widetilde{S}^n$ satisfies $CAT(1)$, then $\mathcal{T}_\Sigma$ satisfies condition (C3). We proceed by induction on n. Let Σ_n be the singular set in S^n and write $\mathcal{T}_n$ for $\mathcal{T}_{\Sigma_n}$. The theorem assumes $n \geq 2$ since for $n = 1$, $\Sigma_1 = \emptyset$, so $\mathcal{T}_1$ is not well-defined. It will be convenient, however, to begin our induction with $n = 1$ since this case occurs as the link of a codimension 2 simplex for larger n. The fact that Σ_n lies on a great $(n-1)$-sphere in S^n means that the dihedral angle between codimension 1 faces in $\mathcal{T}_n$ is always π. Hence, if σ is of codimension 2 in $\mathcal{T}_n$, then the induced triangulation on $L_\sigma = \mathrm{Link}(\sigma, L)$ ($\cong S^1$) consists of two edges of length π. We therefore *define* the triangulation $\mathcal{T}$ of S^1 to consist of two edges of length π meeting at a pair of antipodal points. Clearly this satisfies (C2) and (C3).

Now suppose $n \geq 2$ and that $\widetilde{S}^n$ satisfies $CAT(1)$. In light of Lemma 8.13, it suffices to show that all edges in $\mathcal{T}_n$ have length $\geq \frac{\pi}{2}$. (By induction, the same holds for links in S^n since the induced triangulation on the link L_σ of a codimension k simplex σ is precisely $\mathcal{T}_{k-1}$.) Let Δ_1, Δ_2 denote the two n-simplices of $\mathcal{T}_n$. Recall that $\partial\Delta_1 = \partial\Delta_2$ is a great $(n-1)$-sphere. Let e be an edge of Δ_i and σ be the codimension 2 face opposite e. The branching order at σ is 2. Hence, if $\tilde{\sigma}$ is a lift of σ to $\widetilde{S}^n$, then the star of $\tilde{\sigma}$, $st(\tilde{\sigma})$, consists of two copies of Δ_1 and two copies of Δ_2.

$st(\tilde{\sigma})$:

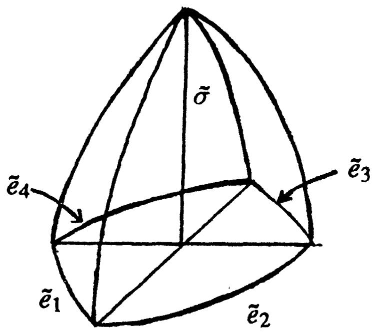

The four lifts of e in $st(\tilde{\sigma})$ form a closed broken geodesic $\gamma = (\tilde{e}_1, \tilde{e}_2, \tilde{e}_3, \tilde{e}_4)$. We claim that γ is a local geodesic in $\widetilde{S}^n$. To see this, let $\tilde{v}$ be the shared vertex of an adjacent pair $(\tilde{e}_i, \tilde{e}_j)$. We must show that the angle between $\tilde{e}_i, \tilde{e}_j$ at $\tilde{v}$ is $\geq \pi$, or in other words, that the tangent vectors x_i, x_j to $\tilde{e}_i, \tilde{e}_j$ at $\tilde{v}$ are at distance $\geq \pi$ in $\widetilde{L}_v = \mathrm{Link}(\tilde{v}, \widetilde{S}^n)$. Now x_i and x_j are vertices in the induced triangulation $\widetilde{\mathcal{T}}_{n-1}$ on $\widetilde{L}_v \cong \widetilde{S}^{n-1}$. Moreover, x_i and x_j do not lie in a common simplex of $\tilde{\tau}_{n-1}$ (since $\tilde{e}_i, \tilde{e}_j$ do not lie in a common simplex of $\tilde{\tau}_n$). By induction, $\mathcal{T}_{n-1}$, and hence $\widetilde{\mathcal{T}}_{n-1}$, satisfies (C3). Clearly (C3) implies that two vertices not contained in a common simplex are of distance $\geq \pi$. This shows that γ is a local geodesic at each vertex $\tilde{v}$. Since $CAT(1)$ holds in $\widetilde{S}^n$, the systole of $\widetilde{S}^n$ is $\geq 2\pi$. Hence $\ell(\gamma) = 4 \cdot \ell(e) \geq 2\pi$, so $\ell(e) \geq \frac{\pi}{2}$. ☐

We end this section with a conjecture for the general case of an orientation-preserving subgroup G of a real reflection group. Recall that if G has rank $n+1$, then the triangulation $\mathcal{T}_\Sigma$ of S^n consists of exactly two n-simplices Δ_1, Δ_2 with their boundaries identified. Thus, if σ is a codimension 2 simplex in $\mathcal{T}_\Sigma$, then there is a unique edge e_σ in $\mathcal{T}_\Sigma$ lying opposite σ in both Δ_1 and Δ_2. Consider the property:

(C4) If σ is a codimension 2 simplex in $\mathcal{T}_\Sigma$, and p is the branching order at σ, then $\ell(e_\sigma) \geq \frac{\pi}{p}$. In addition, the same holds for the induced triangulation $\mathcal{T}_{\Sigma_\tau}$ of $\mathrm{Link}(\tau, S^n)$ for all simplices $\tau \in \mathcal{T}_\Sigma$.

An easy exercise shows that (C4) is equivalent to

(C4)$'$ If σ is a codimension 2 simplex in $\mathcal{T}_\Sigma$, and p is the branching order at σ, then $\ell(e_\sigma) \geq \frac{\pi}{p}$ and for any $\tau \subset \sigma$, the dihedral angle at τ in $\tau * e_\sigma$ is $\geq \frac{\pi}{p}$.

Conjecture. Let G be the orientation-preserving subgroup of a real reflection group of rank $n+1$. Then $\widetilde{S}^n$ satisfies $CAT(1)$ if and only if Σ lies on a great $(n-1)$-sphere, the faces of Σ are totally geodesic, and the associated triangulation $\mathcal{T}_\Sigma$ satisfies (C2) and (C4) (or equivalently (C4)$'$).

Note that this conjecture is supported by Theorems 8.1 and 8.11.

9. Branched covers of S^3 with covering group $G(p,q,r)$. It is conjectured in Section 6 that the only irreducible complex reflection groups which can occur as local groups in an orbifold of nonpositive curvature are the cyclic groups and the groups $G(p,q,r)$ generated by three complex reflections in $\mathbb{C}^2$ of orders $p,q,$ and r with $\frac{1}{p}+\frac{1}{q}+\frac{1}{r}>1$. The action of $G=G(p,q,r)$ on $\mathbb{C}^2$ gives rise to a branched cover $\pi:\widetilde{S}^3\to\widetilde{S}^3/G=S^3$ (where the base, S^3, is given the standard spherical metric and the covering, $\widetilde{S}^3$, is assigned the pullback metric). The branch locus of such a cover consists of three disjoint great circles P_p,P_q,P_r. The local group at $x\in P_m$, $m=p,q,$ or r, is $\mathbb{Z}/m$. The remainder of the paper is devoted to determining which geometric configurations of these 3 great circles give rise to a metric on $\widetilde{S}^3$ satisfying $CAT(1)$.

It will be convenient in the discussion which follows to identify S^3 with the set of unit quaternions and S^2 with the set of totally imaginary unit quaternions (i.e., those $v\in S^3$ such that $v^2=-1$). Note the every unit quaternion $w\in S^3$ can be written in the form $w=e^{v\theta}$ where $v\in S^2$ is totally imaginary. In this terminology, the Hopf map $H:S^3\to S^2$ is given by $H(e^{v\theta})=e^{\overline{v}\theta}ie^{v\theta}=$ the point on S^2 obtained by rotating i about v through an angle of 2θ. The map H is a fibration whose fibers are the circles $e^{i\theta}w$, $0\le\theta<2\pi$, for some $w\in S^3$.

Now let P_1,P_2,P_3 be disjoint great circles in S^3. Say P_1,P_2,P_3 are *Hopf* if there exists an isometry φ of S^3 (possibly reversing orientation) such that $\varphi(P_1)$, $\varphi(P_2)$, $\varphi(P_3)$ are fibers of the Hopf map.

We now state the main theorem of this chapter. Note its similarity with Theorem 8.1.

THEOREM 9.1. *Let $\widetilde{S}^3$ be a branched cover of S^3 with covering group $G(p,q,r)$ and branch locus three disjoint great circles P_p,P_q,P_r of orders $p,q,$ and r respectively. Then $\widetilde{S}^3$ satisfies $CAT(1)$ if and only if P_p,P_q,P_r are Hopf and (up to isometry) their images x_p,x_q,x_r under the Hopf map satisfy the following two conditions:*

(i) $d(x_p,x_q)+d(x_p,x_r)+d(x_q,x_r)=2\pi$

(ii) $d(x_i,x_j)\ge\frac{\pi}{k}$ *where* $\{i,j,k\}=\{p,q,r\}$.

Before embarking on the proof of Theorem 9.1, we offer the following example as an application.

Example 9.2. Let F be a nonpositively curved Riemann surface. Let $M=F\times F$ with the product Riemannian metric. Then M is a smooth 4-manifold of nonpositive curvature. We can construct a variety of interesting orbifold structures of nonpositive curvature on M. The singular sets of these orbifolds are unions of totally geodesic codimension 2 submanifolds, for example, submanifolds of the form

(i) $x\times F$ or $F\times x$, $x\in F$,

(ii) $\gamma_1\times\gamma_2$, γ_1,γ_2 closed geodesics in F,

(iii) the graph of an isometry $\varphi:F\to F$.

The singular submanifolds must intersect in codimension 4 subspaces (i.e. a discrete set of points) and the local group at such a point should be of the form $G_z = \mathbb{Z}/p \times \mathbb{Z}/q$ (if exactly two singular submanifolds meet at z) or $G_z = G(p,q,r)$ (if three singular submanifolds meet at z). Using Theorem 9.1 and the results of Chapter 1, one can then determine exactly which such orbifold structures have nonpositive curvature.

To illustrate, let us consider a particular example. Choose a set of distinct points $x_1, \ldots, x_n$ in F and let $\Delta = \{(x,x) \in F \times F\}$ denote the diagonal in M. Consider an orbifold structure Q on M with singular set consisting of Δ, $x_i \times F$, and $F \times x_i$, $1 \leq i \leq n$. Let (p,q,r) be a 3-tuple of integers ≥ 2 with $\frac{1}{p} + \frac{1}{q} + \frac{1}{r} > 1$, and suppose the local groups G_z of Q are as follows.

(i) For z in the codimension 2 strata

$$G_z = \begin{cases} \mathbb{Z}/p & z = (x_i, x) \\ \mathbb{Z}/q & z = (x, x_i) \\ \mathbb{Z}/r & z = (x, x) \end{cases}$$

where $x \neq x_1, \ldots, x_n$.

(ii) For z in the codimension 4 strata

$$G_z = \begin{cases} \mathbb{Z}/p \times \mathbb{Z}/q & z = (x_i, x_j), \quad i \neq j \\ G(p,q,r) & z = (x_i, x_i) \end{cases}$$

Let T_z denote the tangent space to M at z, viewed as an orbifold, and let S_z denote the unit sphere in T_z. For z in a codimension 2 stratum, T_z is the orthogonal product of a 2-dimensional trivial orbifold and a 2-dimensional cyclic orbifold (as in Example 5.1). For $z = (x_i, x_j)$, $i \neq j$, T_z is the orthogonal product of two 2-dimensional cyclic orbifolds. In both cases, it follows from Proposition 6.2 that $\widetilde{T}_z$ satisfies $CAT(0)$, or equivalently, $\widetilde{S}_z$ satisfies $CAT(1)$. Finally, for $z = (x_i, x_i)$, the singular set of T_z consists of three 2-planes, the tangent planes to $x_i \times F$, $F \times x_i$, and Δ. Thus, S_z is a 3-sphere with singular set three disjoint great circles and covering group $G(p,q,r)$. Since $x_i \times F$ and $F \times x_k$ are orthogonal, we can choose an orthonormal basis v_1, v_2, v_3, v_4 for T_z such that the tangent spaces to $x_i \times F$, $F \times x_i$, and Δ are $\langle v_1, v_2 \rangle$, $\langle v_3, v_4 \rangle$, and $\langle v_1 + v_3, v_2 + v_4 \rangle$, respectively. Identifying v_1, v_2, v_3, v_4 with the basis $1, i, j, k$ for the quarternions, we see that the three singular circles in S_z are $P_p = e^{i\theta}$, $P_q = e^{i\theta} j$, and

$$P_r = e^{i\theta} \left(\frac{1 + j}{\sqrt{2}} \right).$$

Thus P_p, P_q, P_r are Hopf with $H(P_p) = i$, $H(P_q) = -i$, $H(P_r) = k$. Clearly these satisfy conditions (i) and (ii) of Theorem 9.1, hence $\widetilde{S}_z$ satisfies $CAT(1)$. It follows from Theorem 5.3 that Q has curvature ≤ 0, and hence, $\widetilde{Q}$ is a contractible manifold.

Constructions such as the above can be used to provide examples of branched covers $\widetilde{M}$ of Riemannian 4-manifolds M such that $\widetilde{M}$ has nonpositive curvature but is *not* a locally symmetric space (because the ratio of the Euler characteristic to the signature of $\widetilde{M}$ will generally not be that of a locally symmetric 4-manifold). Examples of smooth nonpositively curved manifolds which are not locally symmetric are due to Gromov and Thurston [G-T].

Local Geodesics in $\widetilde{S}^3$. We now begin the proof of Theorem 9.1. Let $\widetilde{S}^3$ be as in Theorem 9.1. Suppose $\tilde{x}$ is a singular point in $\widetilde{S}^3$. Then $\operatorname{Link}(\tilde{x}, \widetilde{S}^3)$ is a branched cover of S^2 with covering group $\mathbb{Z}/m$, $m = p, q$, or r, and branch locus a pair of antipodal points. Or in other words, $\operatorname{Link}(\tilde{x}, \widetilde{S}^3)$ is the suspension of a circle of length $2\pi m$. By Theorem A10 of the appendix, this link satisfies $CAT(1)$, and hence, by Theorem 3.1, $\widetilde{S}^3$ has curvature ≤ 1.

To prove Theorem 9.1, therefore, it suffices to show that $\operatorname{sys}(\widetilde{S}^3) \geq 2\pi$ if and only if P_p, P_q, P_r are Hopf and satisfy conditions (i) and (ii) of the theorem. The condition $\operatorname{sys}(\widetilde{S}^3) \geq 2\pi$ means that every closed local geodesic in $\widetilde{S}^3$ has length $\geq 2\pi$. We begin with a general discussion of local geodesics in $\widetilde{S}^3$. Let γ be a local geodesic (or a closed geodesic) in $\widetilde{S}^3$ and let $\alpha = \pi \circ \gamma$ be its projection to S^3. Then α is a broken geodesic. Moreover, since the restriction of π to the nonsingular part of $\widetilde{S}^3$ is a local isometry, the "breaks" in α can occur only on the singular set. Therefore, we may assume that α is a union of geodesic segments in S^3, $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_k)$, where α_i is a geodesic segment connecting two singular points x_{i-1} to x_i, and α_i contains no singular points in its interior. Likewise, we decompose γ into geodesic segments, and $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_k)$, such that α_i is the image of γ_i. We denote the length of a geodesic segment α_i by $\ell(\alpha_i)$. The length of the broken geodesic α is then

$$\ell(\alpha) = \sum_{i=1}^{k} \ell(\alpha_i).$$

Similarly for γ.

It is possible that some α_i begins and ends on the same singular circle. But in this case, α_i has length π and its endpoints, x_{i-1}, x_i, are antipodal. If α is closed, this implies that α has length $\geq 2\pi$ because any geodesic or broken geodesic from x_i to x_{i-1} will also have length $\geq \pi$. Since we will be concerned only with the existence of closed geodesics of length $< 2\pi$, we may assume that no α_i begins and ends on the same singular circle.

ASSUMPTION 9.3. *Throughout this chapter we will assume that all broken geodesics α in S^3 are of the form $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ where each α_i satisfies*

(i) *α_i contains no singular points in its interior, and*

(ii) *the endpoints of α_i lie on distinct singular circles.*

By "geodesic segment", we will always mean a geodesic segment which satisfies (i) and (ii).

If γ is a local geodesic in $\widetilde{S}^3$, then by Lemma 2.3, for any point $\tilde{x}$ on γ, the incoming and outgoing vectors of γ at $\tilde{x}$ are at distance $\geq \pi$ in the link $\widetilde{L}_x = \mathrm{Link}(\tilde{x}, \widetilde{S}^3)$. Since the isotropy group of $\tilde{x}$ is $\mathbb{Z}/n$, $\widetilde{L}_x$ is the suspension of a circle E of length $2\pi n$ and $\mathbb{Z}/n$ acts on $\widetilde{L}_x$ so that $\widetilde{L}_x/(\mathbb{Z}/n)$ is the standard 2-sphere. Let $\tilde{d}$ denote the distance function in $\widetilde{L}_x$. It follows from Lemmas A3 and A7 of the appendix, that *for any points $\tilde{v}, \tilde{w} \in \widetilde{L}_x$, $\tilde{d}(\tilde{v}, \tilde{w}) \leq \pi$, and $\tilde{d}(\tilde{v}, \tilde{w}) = \pi$ if and only if $\tilde{v}, \tilde{w}$ are the "poles", $\pm z$, of the suspension, or if they lie on opposite "circles of latitude" (i.e. equidistant from the equator) and their projections on the equator E have distance $\geq \pi$.*

$\widetilde{L}_x$:

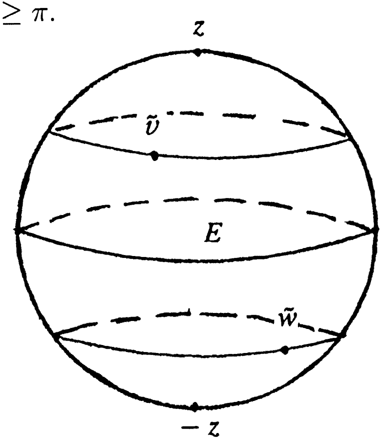

In particular, suppose v, w are non-antipodal points lying on opposite circles of latitude in the standard 2-sphere $\widetilde{L}_x/(\mathbb{Z}/n)$. Given a lift $\tilde{v} \in \widetilde{L}_x$ of v, there is a *unique* lift $\tilde{w}_0 \in L$ of w such that $\tilde{d}(\tilde{v}, \tilde{w}_0) < \pi$. Any other lift $\tilde{w}_1$ of w is a translate of $\tilde{w}_0$ under $\mathbb{Z}/n$ and $\tilde{d}(\tilde{v}, \tilde{w}_1) = \pi$. If v and w are antipodal, then *any* lifts $\tilde{v}, \tilde{w}$ satisfy $\tilde{d}(\tilde{v}, \tilde{w}) = \pi$.

The next two lemmas follow immediately from the preceding discussion.

LEMMA 9.4. *Suppose that $\alpha = (\alpha_1, \dots, \alpha_k)$ is a broken geodesic in S^3 which is the projection of a local geodesic in $\widetilde{S}^3$. Then α satisfies the following "angle condition".*

(Ang) For $i = 1, \dots, k-1$, the incoming vector of α_i at x_i and the outgoing vector of α_{i+1} at x_i make the same angle with the singular circle containing x_i. (Moreover, if α is closed then the same statement is true for α_k and α_1.)

Suppose that $\alpha = (\alpha_1, \dots, \alpha_k)$ is a broken geodesic in S^3 satisfying the angle condition (Ang). The *angle between α_i and α_{i+1}*, denoted $\measuredangle(\alpha_i, \alpha_{i+1})$, is by definition the distance between the incoming vector of α_i at x_i and the outgoing vector of α_{i+1} at x_i (where distance is measured in $\mathrm{Link}(x_i, S^3)$ ($\cong S^2$)), i.e.,

$$\measuredangle(\alpha_i, \alpha_{i+1}) = d((\alpha_i')_{in}(x_i), (\alpha_{i+1}')_{out}(x_i)).$$

Let α_1 be a geodesic segment in S^3 meeting the singular circles only at its endpoints $x_0 \in P_i$ and $x_1 \in P_j$. Let $\tilde{x}_0$ be a point in $\widetilde{S}^3$ lying over x_0 and let n_i denote the order of the isotropy group at $\tilde{x}_0$. Then there are precisely n_i different lifts of α_1 beginning at $\tilde{x}_0$.

LEMMA 9.5. *Suppose that $\alpha = (\alpha_1, \alpha_2)$ is a broken geodesic in S^3 satisfying condition (Ang) at x_1 and that $\measuredangle(\alpha_1, \alpha_2) < \pi$. Let $\tilde{\alpha}_1$ be a lift of α_1 to a geodesic segment in $\widetilde{S}^3$ from $\tilde{x}_0$ to $\tilde{x}_1$. Then there is a unique lift $\tilde{\alpha}_2$ of α_2 such that the initial point of $\tilde{\alpha}_2$ is $\tilde{x}_1$ and such that the angle between $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ is $< \pi$. Thus, if β is any other lift of α_2 beginning at $\tilde{x}_1$, then $(\tilde{\alpha}_1, \beta)$ is a local geodesic in $\widetilde{S}^3$.*

10. Reduction to the Hopf Case. In this section we show that if $\mathrm{sys}(\widetilde{S}^3) \geq 2\pi$ then the singular circles in S^3 must be Hopf. It will be convenient from now on to denote the singular circles by P_1, P_2, P_3 (instead of P_p, P_q, P_r) and set $n_i =$ branching order of P_i.

LEMMA 10.1. *Let P_1, P_2, P_3 be disjoint great circles in S^3 and let $x \in P_1$. Then there exists a unique great circle C in S^3 passing through x and intersecting P_2 and P_3.*

Proof. Consider the great 2-sphere S containing x and P_2. P_3 must intersect this 2-sphere in a pair of antipodal points $\{y, -y\}$. Let C be the circle in S through x and y. □

LEMMA 10.2. *Let P_1, P_2, P_3 be disjoint great circles in S^3. Then the following are equivalent.*

(1) *Every circle intersecting P_1, P_2, and P_3 intersects all three orthogonally.*

(2) *P_1, P_2, P_3 are Hopf and their images in S^2 lie on a great circle.*

Proof. First assume P_1, P_2, P_3 are Hopf. Let C be a circle intersecting P_1, P_2, and P_3. It's projection under the Hopf map H is a circle $H(C)$ in S^2 containing the points $H(P_1)$, $H(P_2)$, and $H(P_3)$. It is easy to verify that C intersects P_i orthogonally if and only if $H(C)$ is a *great* circle.

Conversely, assume every circle intersecting P_1, P_2, and P_3 does so orthogonally. One can show that every circle C in S^3 can be written in the form $C = e^{v\theta}a = ae^{(\bar{a}va)\theta}$, $0 \leq \theta < 2\pi$, for some $a \in S^3 = \{\text{unit quaternions}\}$, $v \in S^2 = \{\text{purely imaginary unit quaternions}\}$, and that two such circles $C = e^{v\theta}a$ and $C' = e^{v'\theta}a'$ intersect orthogonally if and only if $d_{S^2}(v, v') = d_{S^2}(\bar{a}va, \bar{a}'v'a') = \frac{\pi}{2}$. Moreover, circles C and C' are identical if and only if $(v, \bar{a}va) = \pm(v', \bar{a}'v'a')$ in $S^2 \times S^2$; they are of distance $\frac{\pi}{2}$ if and only if $(v, \bar{a}va) = \pm(v', -\bar{a}'v'a')$ in $S^2 \times S^2$.

Up to an isometry of S^3 we may assume that $P_1 = e^{i\theta}$ and that one of the common orthogonal circles is $C_1 = e^{j\theta}$. Let $P_2 = e^{w\theta}a$. Choose another common orthogonal circle $C_2 = e^{v\theta}c$. Since Lemma 10.1 implies that there is a one-parameter family of such C_2, we may assume that $d(C_1, C_2) < \frac{\pi}{2}$. Since P_1 and P_2 intersect both C_1 and C_2 orthogonally, $\langle i, w\rangle$ and $\langle j, v\rangle$ are a pair of orthogonal subspaces of $\mathbb{R}^3$ (= {imaginary quaternions}). (Here $\langle\cdots\rangle$ denotes "the subspace spanned by ...".) Similarly, $\langle i, \bar{a}wa\rangle$ and $\langle j, \bar{c}vc\rangle$ are orthogonal. Since $d(C_1, C_2) < \frac{\pi}{2}$ either $j \neq \pm v$ or $j \neq \pm\bar{c}vc$. Suppose $j \neq \pm v$. Then $\langle j, v\rangle$ is 2-dimensional, so $\langle i, w\rangle$ must be 1-dimensional. Thus, $w = \pm i$ and $P_2 = e^{i\theta}a$. The same argument shows that P_3 is of the form $P_3 = e^{i\theta}b$. Hence, in this case, P_1, P_2, P_3 are Hopf and their images lie in $H(C_1) = H(e^{j\theta})$, a great circle in S^2. Similarly, if $j \neq \pm\bar{c}vc$, then $\bar{a}wa = \pm i$, so $P_2 = ae^{i\theta}$ and by the same argument, $P_3 = be^{i\theta}$. Applying quaternionic conjugation $s \mapsto \bar{s}$ (an isometry of S^3), we get $\overline{P}_1 = e^{i\theta}$, $\overline{P_2} = e^{i\theta}\bar{a}$, $\overline{P}_3 = e^{i\theta}\bar{b}$. Thus, $\overline{P}_1, \overline{P}_2, \overline{P}_3$ are Hopf and their images are contained in the great circle $H(\overline{C}_1)$. The lemma is proved. □

LEMMA 10.3. *Suppose P_1, P_2, P_3 are disjoint great circles in S^3 which do not satisfy (1) and (2) of Lemma 10.2. Then (renumbering the P_i's if necessary) there exists a broken geodesic $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ in S^3 of length $< \pi$ which goes from P_1 to P_2 to P_3 to P_1. Moreover, if α is the shortest such broken geodesic, then α satisfies the angle condition (Ang) and intersects P_1 orthogonally at both ends.*

Proof. Let C be a circle which intersects P_1, P_2 and P_3, and suppose the angle at $C \cap P_1$ is not $\frac{\pi}{2}$. Consider a semicircle of C beginning and ending at P_1. We can view this semicircle as a piecewise geodesic $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ where α_1 is the segment from P_1 to P_2, α_2 the segment from P_2 to P_3, and α_3 the segment from P_3 to P_1. Let $x_1 \in P_2$ be the endpoint of α_1. Since α_1 meets P_1 non-orthogonally, α_1 is *not* the shortest segment from x_2 to P_1. Replacing α_1 by α_1' = orthogonal from x_2 to P_1, we obtain a piecewise geodesic $\alpha' = (\alpha_1', \alpha_2, \alpha_3)$ of length $\ell(\alpha') < \ell(\alpha) = \pi$. This proves the first statement of the lemma.

Now suppose $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ is the shortest piecewise geodesic form P_1 to P_2 to P_3 to P_1. We have already observed that α must be perpendicular to P_1 at both ends. It remains only to show that α satisfies the angle condition. It suffices to prove the following statement: if x_1, x_2 are fixed points in S^3 and P is a circle in S^3 not containing x_1 or x_2, then the shortest piecewise geodesic $\beta = (\beta_1, \beta_2)$

from x_1 to x_2 passing through P satisfies the angle condition. To see this, let φ be an orthogonal rotation about P such that $\varphi(\beta_2)$ lies in the great 2-sphere S determined by P and β_1 in the opposite hemisphere from β_1.

S :

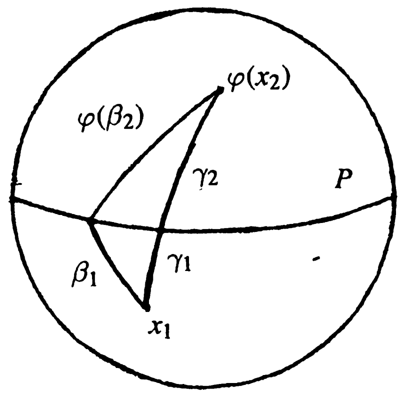

This rotation does not change the angle between P and β_2, so (β_1, β_2) satisfies the angle condition if and only if $(\beta_1, \varphi(\beta_2))$ does. Now if $(\beta_1, \varphi(\beta_2))$ fails to satisfy the angle condition, then the local geodesic $\gamma = (\gamma_1, \gamma_2)$ from x_1 to $\varphi(x_2)$ in S is strictly shorter than $(\beta_1, \varphi(\beta_2))$. Hence, the piecewise geodesic $(\gamma_1, \varphi^{-1}(\gamma_2))$ from x_1 to x_2 is strictly shorter than β. $\square$

The next proposition is the main result of this section. To prove the proposition, we will need a few facts about the group $G = G(p,q,r)$. This group is generated by three elements g_1, g_2, g_3 of orders p, q, and r, respectively, whose product $z = g_1 g_2 g_3$ generates the center of G. If $\widetilde{P}_1, \widetilde{P}_2, \widetilde{P}_3$ are singular circles such that some convex fundamental domain for G intersects all three circles, then we can choose g_1, g_2, g_3 so that g_i fixes $\widetilde{P}_i$. Moreover, since each singular circle is the fixed point set of some subgroup of G, and $z = g_1 g_2 g_3$ is central, it follows that z stabilizes every singular circle.

PROPOSITION 10.4. *Let P_1, P_2, P_3 be the branch circles of $\widetilde{S}^3 \to S^3$. If $sys(\widetilde{S}^3) \geq 2\pi$, then P_1, P_2, P_3 satisfy (1) and (2) of Lemma 10.2.*

Proof. Suppose P_1, P_2, P_3 do *not* satisfy (1) of Lemma 10.2. Let $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ be the minimal piecewise geodesic described in Lemma 10.3. Let

$$\alpha\overline{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \overline{\alpha}_3, \overline{\alpha}_2, \overline{\alpha}_1)$$

where $\overline{\alpha}_i$ is α_i traced in reverse. Then $\alpha\overline{\alpha}$ is a closed, broken geodesic satisfying the angle condition with length $\alpha\overline{\alpha} < 2\pi$. We will show that $\alpha\overline{\alpha}$ can be lifted to a closed local geodesic in $\widetilde{S}^3$, so $\widetilde{S}^3$ does not satisfy $CAT(1)$. In fact, it suffices to show that α can be lifted to a local geodesic γ which begins and ends at the same singular circle $\widetilde{P}_1$. For in this case, choosing a nontrivial element g in $G(p,q,r)$

which fixes $\widetilde{P}_1$ pointwise, $\gamma \cdot \overline{g\gamma}$ is a closed, local geodesic lift of $\alpha\overline{\alpha}$. (Here, we are using the fact that γ must be perpendicular to $\widetilde{P}_1$ at its endpoints.)

To show that such a lift γ exists, let $\{x_1, x_2, x_3, x_4\}$ be the endpoints of the α_i's and consider the tetrahedron R = convex hull of $\{x_1, x_2, x_3, x_4\}$. Assume for the moment that R is nondegenerate, that is, $\{x_1, x_2, x_3, x_4\}$ is not contained in a 2-sphere. We claim that R contains no singular points in its interior. Clearly $\text{int}(R)$ contains no points of P_1 since one edge of R is a segment of P_1. Suppose $\text{int}(R)$ contains a point $r \in P_2$. The geodesic from x_2 to r is a segment of P_2, so P_2 also intersects the face σ of R opposite x_2 in a point r_0. Drop a perpendicular from r_0 to P_1 and let $y \in P_1$ be its endpoint. Consider the piecewise geodesic $\beta = (\beta_1, \beta_2, \alpha_3)$ where β_1 is the geodesic segment from y to r_0 and β_2 the segment from r_0 to x_3. Note that β lies entirely in the face σ.

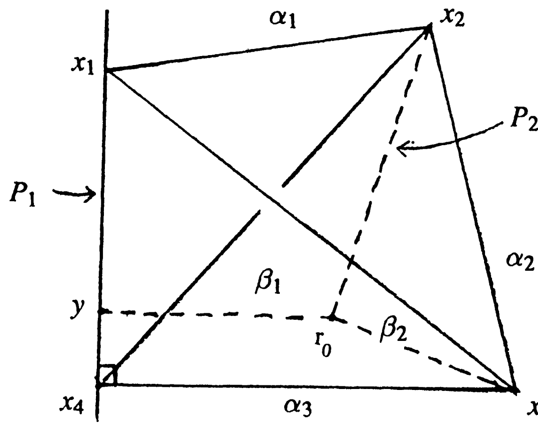

We have

$$
\begin{aligned}
\ell(\beta) &= \ell(\beta_1) + \ell(\beta_2) + \ell(\alpha_3) \\
&\leq \ell(\overline{x_1 x_3}) + \ell(\alpha_3) \\
&< \ell(\alpha_1) + \ell(\alpha_2) + \ell(\alpha_3).
\end{aligned}
$$

But this contradicts the minimality of the length of α, so we conclude that $\text{int}(R)$ contains no points of P_2, and by an analogous argument, it contains no points of P_3. We can therefore lift R to a region $\widetilde{R}$ in $\widetilde{S}^3$ isometric to R. Three of the edges of $\widetilde{R}$ correspond to a lift $\tilde{\alpha} = (\tilde{\alpha}_1, \tilde{\alpha}_2, \tilde{\alpha}_3)$ of α beginning and ending in the same singular circle $\widetilde{P}_1$ and having angles $\leq \pi$. Let $\widetilde{P}_1$ and $\widetilde{P}_2$ be the singular circles containing the endpoints of $\tilde{\alpha}_2$. Choose a set of generators g_1, g_2, g_3 of $G(p,q,r)$ such that g_i fixes $\widetilde{P}_i$ pointwise and $z = g_1 g_2 g_3$ is central. Consider the lift $\gamma = (\tilde{\alpha}_1, g_2\tilde{\alpha}_2, g_2 g_3 \tilde{\alpha}_3)$ of α. By Lemma 9.5, γ is a local geodesic. Moreover, if $\tilde{\alpha}$ begins at $\tilde{x} \in \widetilde{P}_1$ and ends at $\tilde{y} \in \widetilde{P}_1$, then γ begins at $\tilde{x}$ and ends at $g_2 g_3(\tilde{y}) = g_2 g_3 g_1(\tilde{y}) = z(\tilde{y})$. But the center of $G(p,q,r)$ stabilizes each singular circle. Hence $z(\tilde{y}) \in \widetilde{P}_1$.

It remains to consider the case where R is degenerate. That is, where $\{x_1, x_2,$

$x_3, x_4\}$ all lie on some 2-sphere in S^3. In this case we have two possible configurations:

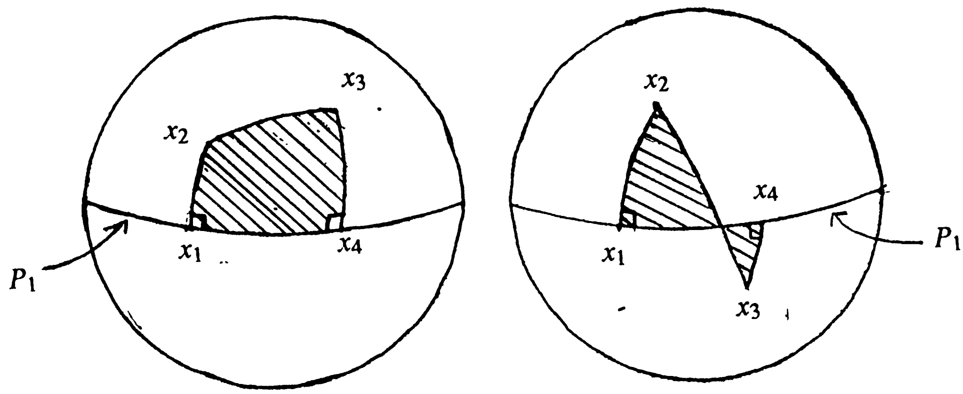

Let D denote the shaded region in either of the diagrams above. Clearly, the minimality of the length of α implies that D contains no singular points in its interior. Hence, as above, we can lift D to an isometric region $\widetilde{D}$ in $\widetilde{S}^3$ and proceed as in the nondegenerate case. $\square$

ASSUMPTION 10.5. *From now on we assume that the singular circles P_1, P_2, P_3 are fibers of the Hopf map and that their images in S^2 lie on a common great circle.*

Under this assumption, we can use the Hopf map to give a more precise description of the action of $G = G(p, q, r)$ on $\widetilde{S}^3$. As remarked above, G is generated by three elements g_1, g_2, g_3 of orders p, q, r whose product $z = g_1 g_2 g_3$ generates the center Z of G. The quotient $\overline{G} = G/Z$ is naturally isomorphic to the orientation-preserving subgroup of $T(p, q, r)$ discussed in Section 8. If $(p, q, r) = (m, 2, 2)$, let $s = m$; otherwise let s be the least common multiple of $\{p, q, r\}$. Then $|Z| = |\overline{G}| = 2s$, hence the order of G is $4s^2$ (see [C; p. 93]).

Now let $H : S^3 \to S^2$ be the Hopf map and let $x_i = H(P_i)$. If we take $\widetilde{S}^2$ to be the branch cover of S^2 with branch locus x_1, x_2, x_3 and covering group $\overline{G}$ (as in Section 8), we get a commutative diagram

$$(10.6)\qquad \begin{array}{ccccc} \widetilde{S}^3 & \xrightarrow{\pi} & S^3 & = & \widetilde{S}^3/G \\ \widetilde{H}\downarrow & & \downarrow H & & \\ \widetilde{S}^2 & \xrightarrow{\pi'} & S^2 & = & \widetilde{S}^2/\overline{G} \end{array}$$

The map $\widetilde{H}$ is also a fibration whose fibers F are circles. The center Z of G acts as a group of translations on F and $\pi(F) = F/Z$ is a Hopf circle in S^3. Thus metrically, F is a circle of length $4s\pi$, and the generator z of Z acts on F as a translation by 2π. The action of Z on $\widetilde{S}^3$ extends to a free, fiberwise action of the full circle group $S^1 = \{e^{i\theta}\}$ with $e^{i\theta}$ acting on each fiber as translation by $2s\theta$. We remark that via this action, an orientation on S^1 gives rise to compatible orientations on all the fibers of $\widetilde{H}$.

As shown in the proof of Theorem 8.1, π' induces a triangulation of $\widetilde{S}^2$ by "black and white" triangles corresponding to the decomposition of S^2 into two regions bounded by the geodesics connecting x_1,x_2,x_3. In terms of this triangulation, generators of $\overline{G}$ are given by rotations $\overline{g}_1,\overline{g}_2,\overline{g}_3$ about the vertices v_1, v_2, v_3 of any (black or white) triangle as indicated below.

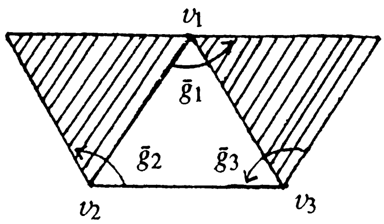

The product $\overline{g}_1\overline{g}_2\overline{g}_3$ is the identity in $\overline{G}$. The generators g_1,g_2,g_3 of G can then be taken as the corresponding rotations about the fibers F_{v_1},F_{v_2},F_{v_3} over v_1, v_2, v_3. Since the stabilizer of v_i is the cyclic group generated by $\overline{g}_i$, and the kernel Z of $G \to \overline{G}$ acts freely on F_{v_i}, it follows that the stabilizer of a point x in F_{v_i} is the cyclic group generated by g_i.

11. Local Geodesics in $\widetilde{S}^3$ in the Hopf Case. We remain under Assumption 10.5: the singular circles P_1,P_2,P_3 are Hopf and their images in S^2 lie on a common great circle. In this section we discuss the relationship between geodesics in $\widetilde{S}^3$ and geodesics in $\widetilde{S}^2$ via the "Hopf map" $\widetilde{H} : \widetilde{S}^3 \to \widetilde{S}^2$, and we prove the "only if" part of Theorem 9.1. Our first goal is to show that any local geodesic α in $\widetilde{S}^2$ lifts to a local geodesic $\tilde{\alpha}$ in $\widetilde{S}^3$ of half the length of α.

By Lemmas 10.1 and 10.2, for any point $x \in P_1$, there is a unique great circle C through x which intersects P_1,P_2,P_3 orthogonally. We call this circle the *common orthogonal circle through* x. We say a geodesic segment β from P_i to P_j is a *standard orthogonal segment* in S^3 if β is a segment of some common orthogonal circle. Likewise any lift of β to $\widetilde{S}^3$ will be called a standard orthogonal segment in $\widetilde{S}^3$. (We remark that if $d(P_i,P_j) = \frac{\pi}{2}$, then every geodesic segment between them intersects P_i and P_j orthogonally. On the other hand, for a given point $x \in P_i$, there is a unique *standard* orthogonal segment from x to P_j.) If β is a standard orthogonal segment from P_i to P_j then its projection under the Hopf map is the unique geodesic segment from $H(P_i)$ to $H(P_j)$ and

$$\ell(\beta) = d(P_i,P_j) = \frac{1}{2}d(H(P_i),H(P_j)) = \frac{1}{2}\ell(H(\beta)).$$

Now recall the commutative diagram (10.6). We want to compare lifts of β to $\widetilde{S}^3$ and lifts of $H(\beta)$ to $\widetilde{S}^2$. Let $x \in P_i$ be the initial point of β and fix a point $\tilde{x} \in \widetilde{S}^3$ over x. If n is the branching order of P_i, then there are n distinct lifts $\widetilde{\beta}_1, \dots, \widetilde{\beta}_n$ of β to $\widetilde{S}^3$ with initial point $\tilde{x}$. These lifts end in distinct singular circles, $\widetilde{P}_j$, $g\widetilde{P}_j, \dots, g^{n-1}\widetilde{P}_j$ where g is a generator of $G_x \cong \mathbb{Z}/n$. It follows that the segments $\widetilde{H}(\widetilde{\beta}_1), \dots, \widetilde{H}(\widetilde{\beta}_n)$ are distinct lifts of $H(\beta)$ with initial point $\widetilde{H}(\tilde{x})$. Since the branching order of $H(x)$ is also n, we see that $\widetilde{H}$ gives a one-to-one correspondence between lifts of β beginning at $\tilde{x}$ and lifts of $H(\beta)$ beginning at $\widetilde{H}(\tilde{x})$. It follows that any broken geodesic $\tilde{\alpha}$ in $\widetilde{S}^2$ beginning at $\widetilde{H}(\tilde{x})$ lifts to a unique broken geodesic $\tilde{\beta}$ in $\widetilde{S}^3$ beginning at $\tilde{x}$ such that each segment of $\tilde{\beta}$ is a standard orthogonal segment.

It remains to show that if $\tilde{\alpha}$ is a local geodesic, then so is its lift $\tilde{\beta}$. Let $\tilde{\beta} = (\beta_1, \beta_2, \dots, \beta_k)$ and $\tilde{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_k)$ where $\alpha_k = \widetilde{H}(\beta_i)$. Then the angle between β_i and β_{i+1} at $\tilde{x}_i$ (= endpoint of β_i) is, by definition, the distance between their tangent vectors in $\mathrm{Link}(\tilde{x}_i, \widetilde{S}^3)$. Now $\mathrm{Link}(\tilde{x}_i, \widetilde{S}^3)$ is isometric to the suspension of $\mathrm{Link}(H(\tilde{x}_i), \widetilde{S}^2)$. Since β_i and β_{i+1} are standard orthogonal segments, their tangent vectors appear on the equator of $\mathrm{Link}(\tilde{x}_i, \widetilde{S}^3)$. Thus, by Proposition A7 of the appendix, the angle between β_i and β_{i+1} is the minimum of π and the angle between α_i and α_{i+1} in $\widetilde{S}^2$. It follows that $\tilde{\beta}$ is a local geodesic if and only if $\tilde{\alpha}$ is. To summarize, we have proved the following.

LEMMA 11.1. *Given a singular point $\tilde{x}$ in $\widetilde{S}^3$ and a broken geodesic $\tilde{\alpha}$ in $\widetilde{S}^2$ beginning at $\widetilde{H}(\tilde{x})$, there exists a unique broken geodesic $\tilde{\beta}$ in $\widetilde{S}^3$ beginning at $\tilde{x}$ such that $\widetilde{H}(\tilde{\beta}) = \tilde{\alpha}$ and such that $\tilde{\beta}$ is made up of standard orthogonal segments. The length of $\tilde{\beta}$ is half that of $\tilde{\alpha}$, and $\tilde{\beta}$ is a local geodesic if and only if $\tilde{\alpha}$ is.*

We are now in a position to prove the "only if" part of Theorem 9.1, which we state as the following.

PROPOSITION 11.2. *If $sys(\widetilde{S}^3) \geq 2\pi$, then the singular circles P_1, P_2, P_3 are (up to isometry of S^3) Hopf circles and their images $v_i = H(P_i)$ in S^2 satisfy (i) $d(v_1, v_2) + d(v_2, v_3) + d(v_1, v_3) = 2\pi$, and (ii) $d(v_i, v_j) \geq \frac{\pi}{n_k}$ for all $\{i, j, k\} = \{1, 2, 3\}$, where n_k =branching order of P_k.*

Proof. From Proposition 10.4, we know that $\mathrm{sys}(\widetilde{S}^3) \geq 2\pi$ implies that P_1, P_2, P_3 are Hopf and that v_1, v_2, v_3 lie on a great circle in S^2. By Theorem 8.1, we know that conditions (i) and (ii) hold if and only if $\mathrm{sys}(\widetilde{S}^2) \geq 2\pi$, where $\widetilde{S}^2$ is the branched cover of S^2 with branch points v_1, v_2, v_3 of order n_1, n_2, n_3. Thus, it suffices to show that $\mathrm{sys}(\widetilde{S}^2) < 2\pi$ implies $\mathrm{sys}(\widetilde{S}^3) < 2\pi$.

So suppose $\tilde{\alpha}$ is a closed local geodesic in $\widetilde{S}^2$ of length $< 2\pi$. Then by Lemma 11.1, $\tilde{\alpha}$ lifts to a local geodesic $\tilde{\beta}$ in $\widetilde{S}^3$ of length $< \pi$. While $\tilde{\beta}$ is not necessarily closed, it must begin and end at the same singular circle $\widetilde{P}$ (since $\tilde{\alpha} = \widetilde{H}(\tilde{\beta})$ is closed). Let $g \in G(p, q, r)$ be a nontrivial element which fixes $\widetilde{P}$ pointwise. Then $\tilde{\beta} \cdot \overline{g\tilde{\beta}}$ is a closed local geodesic in $\widetilde{S}^3$ of length $2 \cdot \ell(\tilde{\beta}) < 2\pi$. $\square$

Edgepaths. Recall from the proof of Theorem 8.1 that the covering $\pi' : \widetilde{S}^2 \to S^2$ gives rise to a triangulation of $\widetilde{S}^2$ whose vertices are the singular points of $\widetilde{S}^2$. The 1-skeleton of this triangulation maps via π' onto the geodesic triangle Δ connecting the singular points x_1, x_2, x_3 of S^2. Under our current assumptions, Δ is a great circle in S^2.

Let Γ denote the 1-skeleton of this triangulation of $\widetilde{S}^2$. For a singular circle $\widetilde{P}$ in $\widetilde{S}^3$, let $v_{\widetilde{P}} = \widetilde{H}(\widetilde{P})$ denote the corresponding vertex in Γ. An *edgepath* φ in Γ is a broken geodesic in $\widetilde{S}^2$. It will be convenient to describe such an edgepath as a sequence of vertices $\rho = (v_0, v_1, \ldots, v_n)$. To a broken geodesic $\gamma = (\gamma_1, \ldots, \gamma_n)$ in $\widetilde{S}^3$, we associate the edgepath $\rho_\gamma = (v_{\widetilde{P}_0}, v_{\widetilde{P}_1}, \ldots, v_{\widetilde{P}_n})$ where $\widetilde{P}_i$ is the singular circle at which γ_i ends and γ_{i+1} begins. Conversely, given an edgepath $\rho = (v_0, v_1, \ldots, v_n)$ and a point $\tilde{x}_0$ in the singular circle over v_0, Lemma 11.1 implies that there is a unique broken geodesic γ_ρ in $\widetilde{S}^3$ beginning at $\tilde{x}_0$ such that $\widetilde{H}(\gamma_\rho) = \rho$ and γ_ρ is made up of standard orthogonal segments. We call γ_ρ the *orthogonal (broken) geodesic associated to ρ at $\tilde{x}_0$*. Since Δ is a great circle, the angle between any two edges at a vertex in Γ is at least π. Thus, the edgepath ρ is a local geodesic in $\widetilde{S}^2$ if and only if ρ has no "backtracks", i.e., no subsequences $\ldots v_{i-1}, v_i, v_{i+1} \ldots$ with $v_{i-1} = v_{i+1}$. It follows from Lemma 11.1 that γ_ρ is a local geodesic if and only if ρ has no backtracks.

Now choose an orientation on $\widetilde{S}^2$. This restricts to an orientation on each triangle in the triangulation of $\widetilde{S}^2$. By a *fundamental cycle* in Γ we will mean an edgepath $\sigma = (v_0, v_1, v_2, v_0)$ corresponding to the boundary of a positively oriented triangle.

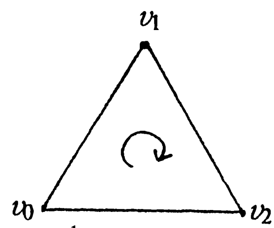

Recall that the circle group S^1 acts as a group of translations on each fiber of the Hopf map, $\widetilde{H} : \widetilde{S}^3 \to \widetilde{S}^2$, with $e^{i\theta}$ acting as translation by $2s\theta$. Denote this action by $x \mapsto x + 2s\theta$.

LEMMA 11.3. *Let σ be a fundamental cycle and let γ_σ be the orthogonal geodesic associated to σ at $\tilde{x}_0$. Then γ_σ ends at $\tilde{x}_0 \pm \pi$. The sign, $\pm$, is independent of σ; it depends only on the orientation of $\widetilde{S}^2$.*

Proof. Since the actions of G on $\widetilde{S}^3$ and $\overline{G}$ on $\widetilde{S}^2$ are orientation-preserving, and since any pair of adjacent triangles in $\widetilde{S}^2$ is a fundamental domain for the action of $\overline{G}$, it suffices to show that there exists a pair σ_1, σ_2 of adjacent fundamental cycles such that if $\gamma_{\sigma_1}, \gamma_{\sigma_2}$ both start at $\tilde{x}_0$, then they both end at $\tilde{x}_0 + \pi$ (or both end at $\tilde{x}_0 - \pi$). Or equivalently, if $\overline{\sigma}_2$ denotes σ_2 with orientation reversed, we must show that γ_{σ_1} ends at $\tilde{x}_0 + \pi$ and $\gamma_{\overline{\sigma}_2}$ ends at $\tilde{x}_0 - \pi$ (or vice versa).

Recall the commutative diagram 10.6.

$$\begin{array}{ccc} \widetilde{S}^3 & \xrightarrow{\pi} & S^3 \\ \widetilde{H}\downarrow & & \downarrow H \\ \widetilde{S}^2 & \xrightarrow{\pi'} & S^2 \end{array}$$

Let $x_0 \in S^3$ be a singular point, say $x_0 \in P_1$. Let C be the common orthogonal circle through x_0. Consider the great 2-sphere $S(C, P_1)$ in S^3 spanned by P_1 and C. Note that the other singular circles P_2 and P_3 intersect $S(C, P_1)$ in pairs of antipodal points $\pm x_2, \pm x_3$ lying on C. Choose a hemisphere of $S(C, P_1)$ bounded by P_1. Then C cuts this hemisphere into two equal digons D_1 and D_2 with vertices $\pm x_0$.

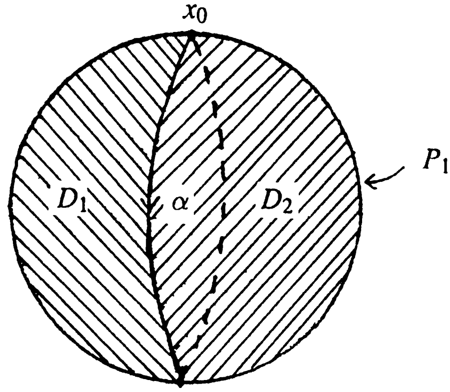

These digons contain no singular points in their interior so they can be lifted to isometric digons $\widetilde{D}_1$ and $\widetilde{D}_2$ in $\widetilde{S}^3$ with one vertex at $\tilde{x}_0$.

Let α be the edge of D_1 and D_2 lying in C (oriented from x_0 to $-x_0$). Then α contains the singular points x_2 and x_3 so the lifts $\tilde{\alpha}_1, \tilde{\alpha}_2$ of α in $\widetilde{D}_1, \widetilde{D}_2$ are broken geodesics beginning at $\tilde{x}_0$. We may, in fact, assume that the first segments of $\tilde{\alpha}_1$ and $\tilde{\alpha}_2$ agree. (If not, replace $\widetilde{D}_2$ by $g \cdot \widetilde{D}_2$ for an appropriate g in the isotropy group of $\tilde{x}_0$.) The projections $\widetilde{H}(\widetilde{D}_1)$, $\widetilde{H}(\widetilde{D}_2)$ of $\widetilde{D}_1, \widetilde{D}_2$ to $\widetilde{S}^2$, are then a pair of adjacent triangles (since $H(D_1), H(D_2)$ are two hemispheres of S^2) with orientations agreeing along the common edge. The boundaries of $\widetilde{H}(\widetilde{D}_1)$ and $\widetilde{H}(\widetilde{D}_2)$ form, respectively, a fundamental cycle σ_1 and the reverse $\overline{\sigma}_2$ of a fundamental cycle σ_2. Moreover, $\tilde{\alpha}_1 = \gamma_{\sigma_1}$ and $\tilde{\alpha}_2 = \gamma_{\overline{\sigma}_2}$ with initial point $\tilde{x}_0$. Now the inverse image of P_1 in $\widetilde{D}_1 \cup \widetilde{D}_2$ is a segment of length 2π in the singular circle $\widetilde{P}_1$ containing $\tilde{x}_0$. In particular, its endpoints are $\tilde{x}_0 + \pi$ and $\tilde{x}_0 - \pi$. It follows that the endpoints of $\tilde{\alpha}_1 (= \gamma_{\sigma_1})$ and $\tilde{\alpha}_2 (= \gamma_{\overline{\sigma}_2})$ are $\tilde{x}_0 + \pi$ and $\tilde{x}_0 - \pi$ (or vice versa). □

In light of the preceding lemma, we may fix an orientation of $\widetilde{S}^2$ so that for every fundamental cycle σ, γ_σ ends at $\tilde{x}_0 + \pi$ (where $\tilde{x}_0$ = initial point of γ_σ).

Now suppose ρ is an arbitrary edgepath in Γ beginning and ending at $v = v_{\widetilde{P}}$. Then there is a unique translation t of $\widetilde{P}$ such that if γ_ρ begins at $\tilde{x} \in \widetilde{P}$, then it ends at $\tilde{x}+t$. Since the fundamental group $\pi_1(\Gamma, v)$ is generated by loops of the form $\rho = \rho_1\sigma\rho_1^{-1}$ where σ is a fundamental cycle, it follows from Lemma 11.3 that t is always a multiple of π. Let T be the group of translations of $\widetilde{P}$ generated by $+\pi$. Since T is abelian (in fact T is cyclic of order $4s$), the homomorphism $\pi_1(\Gamma, v) \to T$, defined by $(\rho_1\sigma\rho_1^{-1}) \mapsto +\pi$, factors through a homomorphism $\tau : H_1(\Gamma) \to T$, defined by $\Sigma\varepsilon_i\sigma_i \mapsto +(\Sigma\varepsilon_i)\pi$, where σ_i is a fundamental cycle. Thus, for any closed edgepath σ, γ_σ is closed if and only if it is homologous to $\Sigma\varepsilon_i\sigma_i$ with $\Sigma\varepsilon_i \equiv 0 \bmod 4s$. We summarize these remarks in the following lemma.

LEMMA 11.4. *Let ρ be a closed edgepath in Γ and let γ_ρ be the orthogonal broken geodesic associated to ρ at $\tilde{x}$. Then*

(i) γ_ρ is a local geodesic if and only if ρ has no backtracks, and

(ii) γ_ρ ends at $\tilde{x}+\tau(\rho)$. In particular, γ_ρ is closed if and only if, in $H_1(\Gamma)$, $\rho \sim \Sigma\varepsilon_i\sigma_i$ where the σ_i's are fundamental cycles and $\Sigma\varepsilon_i \equiv 0$ mod $4s$.

Shifts. The previous lemma gives criteria for determining when a broken geodesic in $\widetilde{S}^3$ consisting of standard orthogonal segments is closed. In this subsection we give similar criteria for an arbitrary broken geodesic in $\widetilde{S}^3$. We first consider geodesic segments in the standard sphere S^3.

Let α be a geodesic segment in S^3 from $x_1 \in P_1$ to $x_2 \in P_2$. Let β be the standard orthogonal segment from x_1 to P_2, and let y be the endpoint of β.

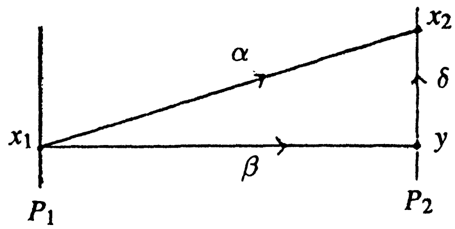

Note that x_2 cannot be antipodal to y since this would imply that α and β constitute half of the common orthogonal circle through x_1. But in that case, either α or β would contain a point of P_3 in its interior, contradicting Assumption 9.3. Thus, there is a unique directed segment δ of P_2 from y to x_2 of length $\ell(\delta) < \pi$. We define the *shift of* α, denoted $s(\alpha)$, to be $\ell(\delta)$ or $-\ell(\delta)$ according to whether the direction of δ agrees or disagrees with the orientation of P_2. Note that, by definition, $0 \leq |s(\alpha)| < \pi$, and that $s(\alpha) = 0$ if and only if α is a standard orthogonal segment. If γ is a geodesic segment in $\widetilde{S}^3$, we define the *shift of* γ, $s(\gamma)$, to be the shift of α where $\alpha = \pi(\gamma)$ is the image of γ in S^3. For a broken geodesic $\alpha = (\alpha_1, \ldots, \alpha_n)$ in S^3, the *total shift of* α is

$$s(\alpha) = \sum_{i=1}^{n} s(\alpha_i)$$

and likewise for a broken geodesic $\gamma = (\gamma_1, \ldots, \gamma_n)$ in $\widetilde{S}^3$.

Suppose γ is a geodesic segment in $\widetilde{S}^3$ from $\tilde{x}_1 \in \widetilde{P}_1$ to $\tilde{x}_2 \in \widetilde{P}_2$ and let $\alpha = \pi(\gamma)$ be its image in S^3. Note that the triangle $T = \Delta(x_1, x_2, y)$ in the figure above contains no singular points in its interior (since the only geodesic through x_1 intersecting all three singular circles is the common orthogonal circle containing β). Hence T lifts to an isometric triangle $\widetilde{T} = \Delta(\tilde{x}_1, \tilde{x}_2, \tilde{y})$ in $\widetilde{S}^3$ with edges γ and $\tilde{\beta}$ where $\tilde{\beta}$ is the standard orthogonal segment form $\tilde{x}_1$ to $\widetilde{P}_2$. The endpoints $\tilde{x}_2$ of γ and $\tilde{y}$ of $\tilde{\beta}$ clearly satisfy $\tilde{x}_2 = \tilde{y} + s(\gamma)$.

More generally, let $\gamma = (\gamma_1, \dots, \gamma_n)$ be a broken geodesic in $\widetilde{S}^3$, ρ the corresponding edgepath, and γ_ρ the orthogonal broken geodesic associated to ρ (with initial point of γ_ρ = initial point of γ). Then γ and γ_ρ end in the same singular circle and, arguing as in the previous paragraph, endpoint (γ) = endpoint$(\gamma_\rho) + s(\gamma)$. In particular, in light of Lemma 11.4, if ρ is closed we obtain the following.

LEMMA 11.5. *Let $\gamma = (\gamma_1, \dots, \gamma_n)$ be a broken geodesic in $\widetilde{S}^3$ beginning at $\tilde{x}$ and suppose the edgepath ρ determined by γ is closed. Then γ ends at $\tilde{x} + \tau(\rho) + s(\gamma)$. In particular, γ is closed if and only if $\tau(\rho) + s(\gamma) \equiv 0$ mod $4\pi s$.*

We close this section with some useful terminology. As usual, P_1, P_2, P_3 denote the singular circles in S^3.

Definition 11.6. Let $\{i, j, k\} = \{1, 2, 3\}$.

(1) A geodesic segment α in S^3 is of *type i* if it connects P_j and P_k.

(2) A singular circle $\widetilde{P}$ in $\widetilde{S}^3$ is of *type i* if its projection to S^3 is P_i.

(3) A geodesic segment γ in $\widetilde{S}^3$ is of *type i* if its projection to S^3 is of type i.

Remark 11.7. In the proof of Theorem 8.1 we defined the "type" of a vertex or edge in Γ. The definitions above are such that $\widetilde{P}$ (resp. γ) in $\widetilde{S}^3$ is of type i if and only if the corresponding vertex (resp. edge) in Γ is of type i. Note also that Assumption 9.3 guarantees that the type of every geodesic segment is well-defined.

Definition 11.8. Let $\gamma = (\gamma_1, \dots, \gamma_n)$ be a closed broken geodesic in S^3 or $\widetilde{S}^3$. We say γ is *even* (resp. *odd*) if γ contains an even (resp. odd) number of segments of each type. Likewise, a closed edgepath ρ in Γ is *even* (resp. *odd*) if it has an even (resp. odd) number of edges of each type. It is easy to see that any closed γ or ρ is necessarily either even or odd. The following lemma follows immediately from the discussion above.

LEMMA 11.9. *Let $\gamma = (\gamma_1, \ldots, \gamma_n)$ be a closed broken geodesic in $\widetilde{S}^3$ and let ρ be the corresponding edgepath. Then the following are equivalent.*

(i) γ *is even (resp. odd)*

(ii) $\rho \sim \Sigma\varepsilon_i\sigma_i$ *with* $\Sigma\varepsilon_i \equiv 0$ *(resp.* $\Sigma\varepsilon_i \equiv 1$*) mod* 2.

(iii) $s(\gamma) = n\pi$ *with* $n \equiv 0$ *(resp.* $n \equiv 1$*) mod* 2.

12. Completion of the proof of Theorem 9.1. To complete the proof of Theorem 9.1, it remains to show that under conditions (i) and (ii) of the theorem, $\widetilde{S}^3$ has systole $\geq 2\pi$. For convenience, let ℓ_i be the length of an edge of type i in $\widetilde{S}^2$. We assume these satisfy conditions (i) and (ii) of Theorem 9.1, namely

(i) $\ell_1 + \ell_2 + \ell_3 = 2\pi$

(ii) $\ell_i \geq \frac{\pi}{n_i}$ where n_i = branching order of P_i.

Note that since each edge has length at most π, (i) implies that $\ell_i + \ell_j \geq \pi$ for any two edge types i and j.

Assuming these conditions, we must show that if $\gamma = (\gamma_1, \ldots, \gamma_n)$ is a closed local geodesic in $\widetilde{S}^3$, then its length is $\geq 2\pi$. We divide the proof into two cases, the orthogonal and non-orthogonal cases (Propositions 12.1 and 12.12). Here we say γ is *orthogonal* if every segment γ_i of γ is perpendicular to the singular circles at its endpoints. Otherwise, we say γ is *non-orthogonal.*

The Orthogonal Case.

PROPOSITION 12.1. *Under conditions (i) and (ii) of Theorem 9.1, if γ is an orthogonal closed local geodesic in $\widetilde{S}^3$, then $\ell(\gamma) \geq 2\pi$.*

Proof. Let $\gamma = (\gamma_1, \ldots, \gamma_n)$ be as in the proposition. We first assume that no pair of singular circles has distance $\frac{\pi}{2}$. In this case, γ is made up of standard orthogonal segments and hence the shift of γ is zero. Moreover, if ρ is the closed edgepath determined by γ then $\gamma = \gamma_\rho$. It follows from Lemma 11.4 that ρ satisfies:

(a) ρ has no backtracks, and

(b) in $H_1(\Gamma)$, $\rho \sim \Sigma\varepsilon_i\sigma_i$ with $\Sigma\varepsilon_i \equiv 0 \bmod 2s$.

Since the length of ρ (viewed as a local geodesic in $\widetilde{S}^2$) is twice that of γ, it suffices to show that any edgepath in Γ satisfying (a) and (b) has length $\geq 4\pi$.

Case 1. ρ contains all 3 types of edges. Since the shift of γ is 0, γ is even. Thus, each type of edge occurs at least twice in ρ, so by (i), $\ell(\rho) \geq 4\pi$.

Case 2. ρ contains only one type of edge, say type 1. Let Γ_1 denote the subgraph of Γ consisting of edges of type 1. Then $H_1(\Gamma_1)$ is generated by cycles of the form $c = \sum_{i=1}^{2n_1} \sigma_i$ where $\sigma_1, \ldots, \sigma_{2n_1}$ are all the fundamental cycles beginning at a vertex v of type 1.

$c =$ 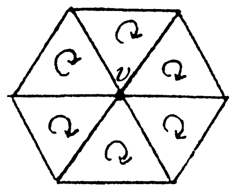

Now $\rho \subset \Gamma_1$, so $\rho \sim \Sigma\varepsilon_j c_j$ in $H_1(\Gamma_1) \subset H_1(\Gamma)$. Clearly the shortest such ρ satisfying (a) and (b) are the edgepaths $\rho = c_1 c_2^{-1}$ where c_1, c_2 share an edge or a vertex. In this case, ρ consists of $4n_1$ edges of type 1, so $\ell(\rho) \geq 4n_1 \cdot \frac{\pi}{n_1} = 4\pi$.

Case 3. ρ contains exactly 2 types of edges, say types 1 and 2. Let $\Gamma_{1,2}$ denote the subgraph of Γ consisting of edges of types 1 and 2. Then $H_1(\Gamma_{1,2})$ is generated by cycles of the form $c = \sigma_1 + \sigma_2$ where σ_1, σ_2 share a common edge of type 3.

$c =$ 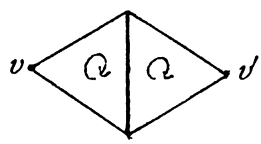

type (v) = type (v') = 3

Since $\rho \subset \Gamma_{1,2}$, $\rho \sim \Sigma\varepsilon_j c_j$ in $H_1(\Gamma)$. Again, it is clear that the shortest such ρ satisfying (a) and (b) is one of the form $\rho = c_1 c_2^{-1}$ where c_1, c_2 share an edge or a vertex. These edgepaths have 4 edges each of types 1 and 2 hence $\ell(\rho) = 4(\ell_1 + \ell_2) \geq 4\pi$.

Finally, we need to consider what happens if some pair of singular circles has distance $\frac{\pi}{2}$, say $d(P_1, P_2) = \frac{\pi}{2}$. This situation differs from the above in that γ can have nonzero shift on segments of type 3. On the other hand, since $d(P_i, P_j) = \frac{1}{2}\ell_k$, condition (i) implies

$$d(P_1, P_3) + d(P_2, P_3) = \pi - d(P_1, P_2) = \tfrac{\pi}{2}.$$

Thus, if γ has length $< 2\pi$, then γ can have at most 2 segments of type 3. (If γ has exactly 3 segments of type 3, then it must also have at least one of types 1 and 2 in order to be closed.) Since each segment has shift of absolute value $< \pi$, the total shift of γ satisfies $0 \leq |s(\gamma)| < 2\pi$. If γ is even, it follows that $s(\gamma) = 0$. In this case, all the arguments in the first part of the proof still hold. The only remaining case is that in which γ is odd and $|s(\gamma)| = \pi$. But if γ is odd and $\ell(\gamma) < 2\pi$, it contains exactly *one* segment of type 3, so $|s(\gamma)| < \pi$. This is a contradiction, so no such γ exists. $\square$

The Non-Orthogonal Case. It remains to analyze the non-orthogonal closed geodesics in $\widetilde{S}^3$. For this we must understand non-orthogonal broken geodesics in S^3. We begin by investigating the link of a singular point x in S^3 from a slightly different viewpoint, that of the "visual sphere" at x. The points in Link(x, S^3) correspond to geodesic rays out of x and we can ask what we "see" as we look along each of these rays.

Let V denote the imaginary quaternions spanned by i,j, and k, that is, $V = \mathbb{R}i \oplus \mathbb{R}j \oplus \mathbb{R}k$. Then S^3 may be identified with the group of unit quaternions, $S^3 = \{e^{v\theta} \mid v \in V, \|v\| = 1\}$, and the Hopf circles in S^3 are the subsets

$$P_a = \{e^{i\theta} \cdot a \mid 0 \leq \theta < 2\pi\}$$

for $a \in S^3$. Every geodesic ray out of $a \in S^3$ is of the form

$$\alpha_v = \{e^{vt} \cdot a \mid t > 0\}$$

for some $v \in V$, $\|v\| = 1$. For such an α_v, the corresponding point in Link(a, S^3) is the unit tangent vector to α_v at a, namely $\alpha'_v = v \cdot a$. Multiplying on the right by a^{-1}, we may therefore identify Link(a, S^3) with $\{v \in V \mid \|v\| = 1\}$, where the ray α_v corresponds to the point v in the link. In particular, the ray along the positively (resp. negatively) oriented direction of P_a corresponds to i (resp. $-i$) in the link. We view these as the "north and south poles" of Link(a, S^3). The "equator" of the link then corresponds to rays α_v with $v \perp i$ and hence $\alpha_v \perp P_a$.

From the "visual sphere" viewpoint, we can label points in Link(a, S^3) according to where the rays α_v lead. For example, suppose that $P_b = \{e^{i\theta} b\}$ is another Hopf circle. We can identify on Link(a, S^3) which rays lead to P_b.

LEMMA 12.2. *The rays from a to P_b form a great circle E in Link(a, S^3) such that the distance d from i to E is equal to $d(P_a, P_b)$. In particular, E is the equator of Link(a, S^3) if and only if P_b is orthogonal to P_a.*

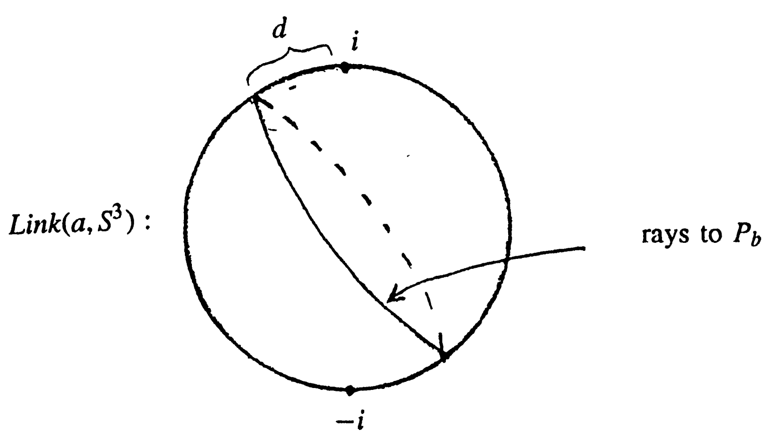

Proof. Up to isometry of S^3, we may assume that $a = 1$, $P_a = \{e^{i\theta}\}$. Let

$H : S^3 \to S^2$ be the Hopf map. Identifying S^2 with the unit sphere in V, H is defined by $H(e^{\nu\theta}) = e^{-\nu\theta} i e^{\nu\theta}$. In other words, $H(e^{\nu\theta})$ is the point obtained by rotating i about v through an angle of 2θ. Let $x_a = i = H(P_a)$, $x_b = H(P_b)$. Consider the ray $\alpha(t) = e^{vt}$, $t > 0$, from a. Projecting to S^2, $H(\alpha(t))$ is a circle (not necessarily a great circle) in S^2 centered at v and passing through x_a.

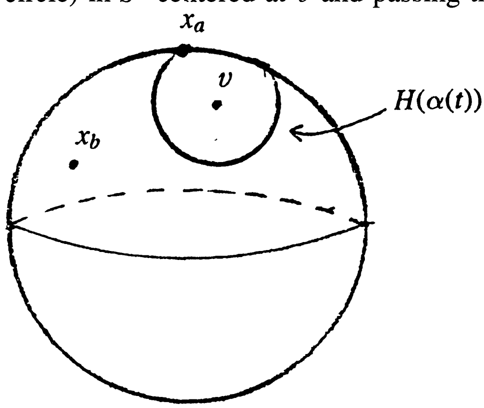

The ray $\alpha(t)$ intersects P_b if and only if $H(\alpha(t))$ passes through x_b, that is, if and only if v lies on the circle E of points equidistant from x_a and x_b. Since the point in Link(a, S^3) corresponding to $\alpha(t) = e^{vt}$ is $\alpha'(0) = v$, we see that the collection of rays from a to P_b correspond to the circle E in $S^2 = \text{Link}(a, S^3)$. Finally, we note that the distance d from the north pole to E satisfies $d = \frac{1}{2}d(x_a, x_b) = d(P_a, P_b)$.□

Now assume P_a, P_b are not orthogonal, i.e. $d(P_a, P_b) < \frac{\pi}{2}$. Let β be the unique orthogonal geodesic from a to P_b of length $\ell(\beta) = d(P_a, P_b)$. Then β corresponds to a point in Link(a, S^3) where the circle E (see the proof above) meets the equator.

LEMMA 12.3. *Assume $d(P_aP_b) < \frac{\pi}{2}$. Then the orientation of P_b is such that it points "upward" from β when viewed in Link(a, S^3):*

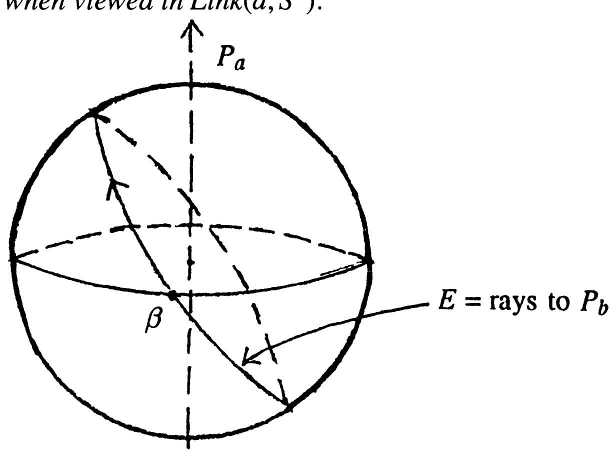

Proof. As before, we may assume $a = 1$. From the proof of the previous lemma, we see that the angle of E from vertical varies with $d(H(P_a), H(P_b)) =$

$2d(P_a, P_b) = 2 \cdot \ell(\beta)$. Now the oriented direction of the Hopf circles varies continuously. Since the direction of P_a points northward from a, for very small β, the direction of P_b must point nearly northward. By continuity, as $\ell(\beta)$ increases, it will continue to point into the northern hemisphere until the circle E reaches the equator, i.e., until $\ell(\beta) = \frac{\pi}{2}$. □

We are now ready to proceed with the analysis of non-orthogonal broken geodesics in S^3. We first show that if $\alpha = (\alpha_1, \ldots, \alpha_n)$ satisfies the angle condition, (Ang), then the angles between α_i and the singular circles at its endpoints are the same for all i. More precisely, we prove the following result.

LEMMA 12.4. *A geodesic segment α in S^3 meets any two Hopf circles at the same angle.*

Proof. Let A be the great circle in S^3 containing α. Let P_1 and P_2 be Hopf circles intersecting A at x_1 and x_2, respectively. Identifying S^3 with the unit quaternions as in Section 9, we have $P_1 = e^{i\theta}x_1$ and $P_2 = e^{i\theta}x_2$, $A = x_1e^{v\theta}$ for some totally imaginary unit quaternion v (cf. proof of Lemma 10.2). In particular, $x_2 = x_1e^{v\rho}$ and multiplication by $e^{v\rho}$ is an isometry of S^3 which preserves A and takes P_1 to P_2. □

It follows from Lemma 12.4 that if $\alpha = (\alpha_1, \ldots, \alpha_n)$ is a broken geodesic satisfying the angle condition (Ang), and if some segment α_i is orthogonal to the singular circles at its endpoints, then the same holds for every segment of α. In particular, in considering non-orthogonal broken geodesics, we may assume that $d(P_i, P_j) < \frac{\pi}{2}$, since if $d(P_i, P_j) = \frac{\pi}{2}$, then any segment between P_i and P_j is orthogonal.

LEMMA 12.5. *Suppose that $\alpha = (\alpha_1, \alpha_2)$ is non-orthogonal and satisfies the angle condition (Ang) at the endpoint x_1 of α_1. Then the shifts of α_1 and α_2 have the same sign.*

Proof. Say the shift of α_1 is positive. Then the incoming vector of α_1 at x_1 is in the southern hemisphere of $\mathrm{Link}(x_1, S^3)$ and hence, by the angle condition, the outgoing vector of α_2 is in the northern hemisphere of $\mathrm{Link}(x_1, S^3)$. Say α_2 begins in P_1 and ends in P_2 and let β be the standard orthogonal segment from $x_1 \in P_1$ to P_2. $\mathrm{Link}(x_1, S^3)$ viewed as the "visual sphere" at x_1, appears as follows.

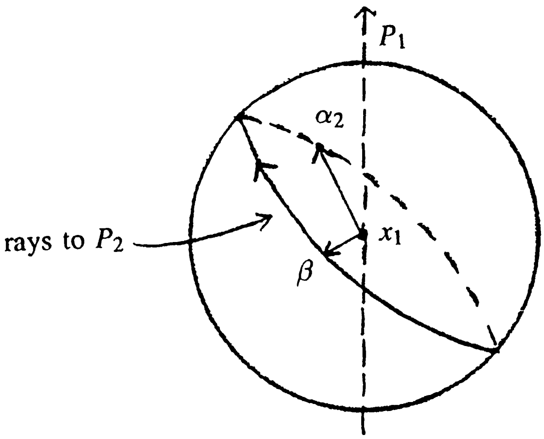

By Lemma 12.3, we know that the orientation of P_2 points upward from β in this visual sphere, hence the shift of α_2 is positive. □

In the next lemma we collect some facts about right spherical triangles. Consider the right spherical triangle

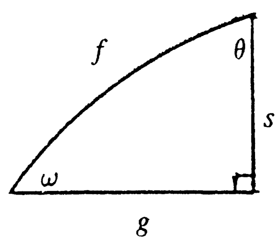

with side lengths f, g, and s as indicated. Clearly, for fixed θ and s, there is a unique such triangle up to isometry. Thus, fixing θ, we may view f, g, and ω as functions of s, $0 < s < \pi$. Write $f = f_\theta(s)$, $g = g_\theta(s)$, and $\omega = \omega_\theta(s)$.

LEMMA 12.6. *Assume* $0 < \theta < \frac{\pi}{2}$. *Then*

(i) $f_\theta(\pi - s) = \pi - f_\theta(s)$

$g_\theta(\pi - s) = g_\theta(s)$

$\omega_\theta(\pi_s) = \pi - \omega_\theta(s)$

(ii) f_θ *and* g_θ *are strictly increasing for* $0 < s < \frac{\pi}{2}$

(iii) *For* $0 < s < \frac{\pi}{2}, f_\theta$ *is concave downward; in particular,* $f_\theta\left(\frac{s}{n}\right) > \frac{1}{n} f_\theta(s)$.

(iv) $f_\theta(s) > s$ *for* $0 < s < \frac{\pi}{2}$

$f_\theta(x) = s$ *for* $s = \frac{\pi}{2}$

$f_\theta(s) < s$ *for* $\frac{\pi}{2} < s < \pi$

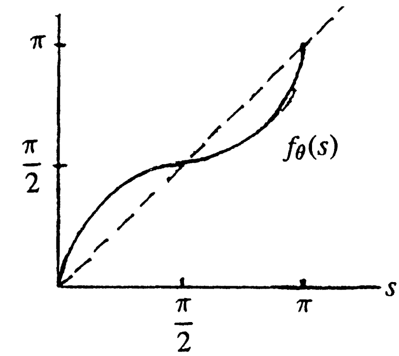

Proof. Consider a wedge of S^2 of angle θ. Cutting the wedge with a geodesic segment orthogonal to the right edge gives rise to two right triangles, each having an angle θ:

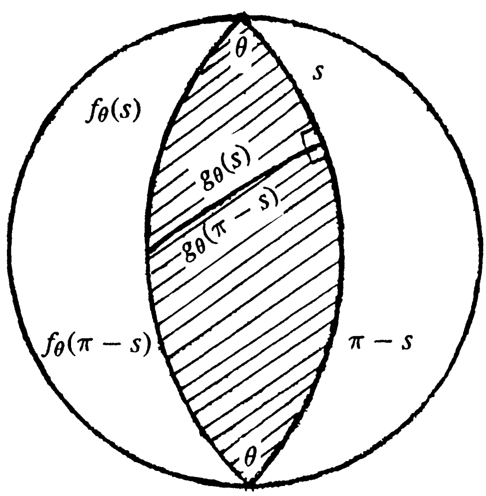

This proves assertion (i) of the lemma.

Now a straightforward exercise in spherical geometry gives the following formulas for f_θ and g_θ on the interval $(0, \frac{\pi}{2})$:

$$\begin{aligned} f_\theta(s) &= \arctan(a \cdot \tan s), \quad a = \sec\theta > 0 \\ g_\theta(s) &= \arctan(b \cdot \sin s), \quad b = \tan\theta > 0. \end{aligned}$$

Direct computation shows that $g'_\theta > 0, f'_\theta > 0$, and $f''_\theta < 0$ on $(0, \frac{\pi}{2})$. Assertions (ii), (iii), and (iv) follow. □

Let $\alpha = (\alpha_1, \alpha_2)$ be a broken geodesic in S^3 satisfying the angle condition. Then α is called a *bounce of type k* if α_1 and α_2 are both of type k. If in addition, the angle between α_1 and α_2 is $< \pi$, then α is a *proper bounce.* Note that if α is a bounce with $\measuredangle(\alpha_1, \alpha_2) = \pi$, then the endpoints of α are antipodal points on some singular circle P_i. Thus, any *closed* broken geodesic containing such a bounce has length $\geq 2\pi$. For our purposes, therefore, we can and will *assume that all bounces are proper.*

LEMMA 12.7. *Let $\alpha = (\alpha_1, \ldots, \alpha_n)$ be a broken geodesic in S^3 satisfying the angle condition. If α_i, α_j are of the same type, then either $\ell(\alpha_i) = \ell(\alpha_j)$ and $s(\alpha_i) = s(\alpha_j)$ or $\ell(\alpha_i) = \pi - \ell(\alpha_j)$ and $s(\alpha_i) = \pi - s(\alpha_j)$. If (α_i, α_{i+1}) is a bounce, then $\ell(\alpha_i) = \ell(\alpha_{i+1})$ and $s(\alpha_i) = s(\alpha_{i+1})$.*

Proof. By Lemma 12.5, we know that the signs of $s(\alpha_i)$ and $s(\alpha_j)$ are the same. Let $s_i = |s(\alpha_i)|$, $s_j = |s(\alpha_j)|$. Consider the right-spherical triangle

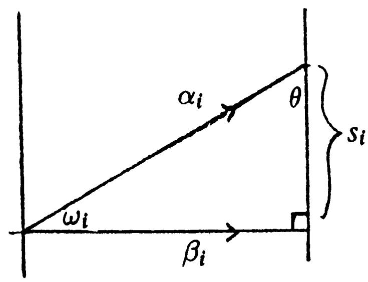

where β_i is a standard orthogonal segment. With notation as in Lemma 12.6, $\ell(\alpha_i) = f_\theta(s_i)$, $\ell(\beta_i) = g_\theta(s_i)$, and $\omega_i = \omega_\theta(s_i)$. Similarly for α_j. But if α_i and α_j are of the same type, say type 3, then $\ell(\beta_i) = \ell(\beta_j) = d(P_1, P_2)$. Thus, $g_\theta(s_i) = g_\theta(s_j)$ and it follows from Lemma 12.6 that either

$$\begin{cases} s_i = s_j \\ f_\theta(s_i) = f_\theta(s_j) \\ \omega_\theta(s_i) = \omega_\theta(s_j) \end{cases} \qquad \text{or} \qquad \begin{cases} s_i = \pi - s_j \\ f_\theta(s_i) = \pi - f_\theta(s_j) \\ \omega_\theta(s_i) = \pi - \omega_\theta(s_j). \end{cases}$$

This proves the first statement of the lemma. Now if (α_i, α_{i+1}) is a bounce, then the angle between α_i and α_{i+1} at x_i is $\omega_i + \omega_{i+1}$, so if the bounce is proper $\omega_i \neq \pi - \omega_{i+1}$. We conclude that $\omega_i = \omega_{i+1}$, so the two triangles are congruent. □

If $\alpha = (\alpha_1, \ldots, \alpha_n)$ is a broken geodesic, we will call a segment α_i a *short edge* if $\ell(\alpha_i) < \frac{\pi}{2}$ and a *long edge* if $\ell(\alpha_i) \geq \frac{\pi}{2}$.

LEMMA 12.8. *Let α be a non-orthogonal geodesic segment. If α is short, then $0 < |s(\alpha)| < \ell(\alpha)$, and if α is long, then $\ell(\alpha) < |s(\alpha)| < \ell(\alpha) + \frac{\pi}{2}$.*

Proof. Let $s = |s(\alpha)|$. This lemma follows immediately from Lemma 12.6, since $\ell(\alpha) = f_\theta(s)$. □

Let ℓ_i, $i = 1, 2, 3$, be as at the beginning of this section. Then, in particular, $\ell_i = 2d(P_j, P_k)$, so any segment in S^3 of type i has length $\geq \frac{\ell_i}{2}$. Assume, as before, that ℓ_1, ℓ_2, ℓ_3 satisfy $\ell_1 + \ell_2 + \ell_3 = 2\pi$ and $\ell_i \geq \frac{\pi}{n_i}$.

LEMMA 12.9. *Suppose* $\alpha = (\alpha_1, \ldots, \alpha_n)$ *is a non-orthogonal closed broken geodesic in* S^3 *satisfying the angle condition (Ang). If* $\ell(\alpha) < 2\pi$, *then* α *satisfies the following.*

(i) All three types of edges occur among the α_i *and two types occur only once.*

(ii) The total shift of α *is* $\pm\pi$.

Remark 12.10. If α is closed, we may choose any singular point on α as our starting point. For α satisfying (i) of the lemma, we may choose the starting point such that the types of α_1 and α_2, say types 1 and 2, occur only once, while all the remaining α_i's are of type 3.

Proof of Lemma 12.9. Suppose α is as in the lemma and that $\ell(\alpha) < 2\pi$. We divide the argument into 3 cases.

Case 1. α has only one type of edge. In this case, α has an even number of edges and every consecutive pair of edges forms a bounce. Since $\ell(\alpha) < 2\pi$, these bounces may be assumed to be proper. Thus, all the edges of α have the same length. If the edges are all short, then by Lemma 12.8, the shift of α satisfies

$$0 < |s(\alpha)| < \ell(\alpha) < 2\pi.$$

If the edges are long, then there are at most 2 edges (since 4 long edges would give $\ell(\alpha) \geq 2\pi$), so again

$$0 < |s(\alpha)| \leq |s(\alpha_1)| + |s(\alpha_2)| < 2\pi.$$

On the other hand, by Lemma 11.9, if α is even then $s(\alpha) \equiv 0 \bmod 2\pi$. Thus, this case cannot occur.

Case 2. α has two types of edges, say types 1 and 2. After possibly choosing a new starting point, α consists of a sequence of (proper) bounces of types 1 and 2. In particular, α is even so $|s(\alpha)| \equiv 0 \bmod 2\pi$. If all the edges of α are short, then, as in case 1,

$$0 < |s(\alpha)| < \ell(\alpha) < 2\pi$$

gives a contradiction.

Hence there must be exactly one bounce, say of type 1, whose edges are long. (If two of the bounces are long, then there are 4 long edges, hence $\ell(\alpha) \geq 2\pi$.) Moreover, we claim that all the other bounces are of type 2. For if (α_i, α_{i+1})

and (α_j, α_{j+1}) are respectively a short and a long bounce of type 1, then by Lemma 12.7

$$\ell(\alpha_i) = \ell(\alpha_{i+1}) = \pi - \ell(\alpha_j) = \pi - \ell(\alpha_{j+1}).$$

So the total length of these four edges is 2π. Thus, we may assume that $\alpha = (\alpha_1, \alpha_2, \dots \alpha_{2n+2})$ where (α_1, α_2) is a long bounce of type 1 and $(\alpha_3, \alpha_4, \dots, \alpha_{2n+2})$ is a series of n short bounces of type 2. Then for $i \geq 3$, $\ell(\alpha_i) = \ell(\alpha_3)$ and $s(\alpha_i) = s(\alpha_3)$, so

$$|s(\alpha)| = 2|s(\alpha_1)| + 2n|s(\alpha_3)| \equiv 0 \mod 2\pi.$$

In fact, since $0 < 2n|s(\alpha_3)| < \ell(\alpha_3, \dots, \alpha_{2n+2}) < \pi$, and $|s(\alpha_1)| < \pi$, we have

$$2|s(\alpha_1)| + 2n|s(\alpha_3)| = 2\pi$$

so

$$|s(\alpha_3)| = \frac{\pi - |s(\alpha_1)|}{n}.$$

Now compare the triangles

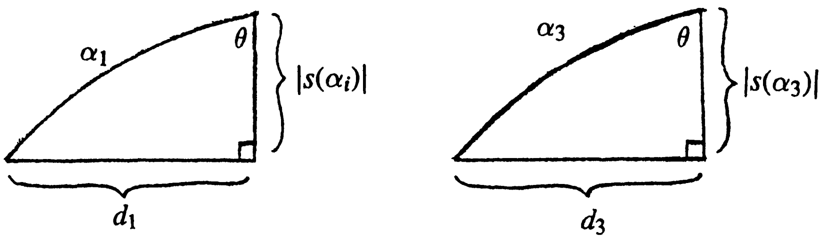

The angles marked θ agree by the angle condition. Moreover, since $d_1, d_2 \leq \frac{\pi}{2}$, we have $\theta \leq \frac{\pi}{2}$. Letting f_θ be as in Lemma 12.6, it follows from that lemma that

$$\begin{aligned}\ell(\alpha_3) &= f_\theta(|s(\alpha_3)|) = f_\theta\left(\frac{\pi - |s(\alpha_1)|}{n}\right) > \frac{1}{n} f_\theta(\pi - |s(\alpha_1)|) \\ &= \frac{1}{n}(\pi - f_\theta(|s(\alpha_1)|)) = \frac{1}{n}(\pi - \ell(\alpha_1)).\end{aligned}$$

Hence,

$$\ell(\alpha) = 2\ell(\alpha_1) + 2n\,\ell(\alpha_3) > 2\pi.$$

We conclude that this case cannot occur.

Case 3. α has three types of edges. Since the length of an edge of type i is $\geq \frac{\ell_i}{2}$ and $\ell_1 + \ell_2 + \ell_3 = 2\pi$, α cannot contain two (or more) edges of each type. Thus α must be odd and contain at most one edge of some type, say type 3. Moreover, $\ell_1 + \ell_2 \geq \pi$, hence if edge types 1 and 2 occur three or more times, we have

$$\ell(\alpha) \geq 3\left(\frac{\ell_1}{2} + \frac{\ell_2}{2}\right) + \frac{\ell_3}{2} = \frac{1}{2}(\ell_1 + \ell_2 + \ell_3) + (\ell_1 + \ell_2) \geq 2\pi.$$

We conclude (since α is odd) that either type 1 or 2, say type 2, occurs only once. This proves assertion (i) of the Lemma.

It remains to prove that $|s(\alpha)| = \pi$. Since π is odd, $s(\alpha) \equiv \pi \bmod 2\pi$, so we need only show that $|s(\alpha)| < 3\pi$. Choosing a new starting point if necessary, we may assume $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_n)$ where α_1, α_2 are of types 2 and 3 and the remaining α_i's are of type 1. In particular, $\alpha_3, \alpha_4, \ldots, \alpha_n$ have the same length and shift. Now if these edges are short, then by Lemma 12.8

$$\begin{aligned} |s(\alpha)| &= (|s(\alpha_1)| + |s(\alpha_2)|) + (|s(\alpha_3)| + \cdots + |s(\alpha_n)|) \\ &< (\ell(\alpha_1) + \ell(\alpha_2) + \pi) + (\ell(\alpha_3) + \cdots + \ell(\alpha_n)) \\ &= \ell(\alpha) + \pi < 3\pi. \end{aligned}$$

On the other hand, if $\alpha_3, \ldots, \alpha_n$ are long, then there is only one such edge (i.e., $n = 3$) since 3 or more such edges would make $\ell(\alpha) \geq 2\pi$. But in this case, since the shift of any one edge has absolute value less than π, we again have

$$|s(\alpha)| = |s(\alpha_1)| + |s(\alpha_2)| + |s(\alpha_2)| < 3\pi.$$

This proves assertion (ii). □

Now suppose $\gamma = (\gamma_1, \ldots, \gamma_n)$ is a closed broken geodesic in $\widetilde{S}^3$ of length $< 2\pi$. Then its image $\alpha = \pi(\gamma)$ in S^3 must be as in Lemma 12.9. To complete the proof of Theorem 9.1, we will show that any such γ contains a pair (γ_i, γ_{i+1}) such that the angle between them, $\measuredangle(\gamma_i, \gamma_{i+1})$, is less than π. In other words, γ cannot be a local geodesic.

We need one final lemma. Recall that if $\alpha = (\alpha_1, \alpha_2)$ is a non-orthogonal broken geodesic in S^3 satisfying the angle condition, and γ_1 is a lift of α_1 to $\widetilde{S}^3$, then by Lemma 9.5 there is a *unique* lift γ_2 of α_2 such that $\measuredangle(\gamma_1, \gamma_2) < \pi$.

LEMMA 12.11. *Let $\alpha = (\alpha_1, \alpha_2)$ be a non-orthogonal broken geodesic in S^3 satisfying the angle condition, and assume type $\alpha_1 \neq$ type α_2. Let $\gamma = (\gamma_1, \gamma_2)$ be a lift of α to a broken geodesic in $\widetilde{S}^3$ and let ρ be the edgepath in Γ corresponding to γ. Suppose α has negative (resp. positive) shift. Then $\measuredangle(\gamma_1, \gamma_2) < \pi$ if and only if ρ corresponds to two positively (resp. negatively) oriented edges of a fundamental cycle in Γ.*

Proof. Replacing α by the reverse path $\overline{\alpha} = (\overline{\alpha}_2, \overline{\alpha}_1)$ if necessary, we may assume that the shift of α is negative. Let x be the common endpoint of α_1 and α_2 and let $\beta = (\beta_1, \beta_2, \beta_3)$ be half of the common, orthogonal circle through x as indicated below.

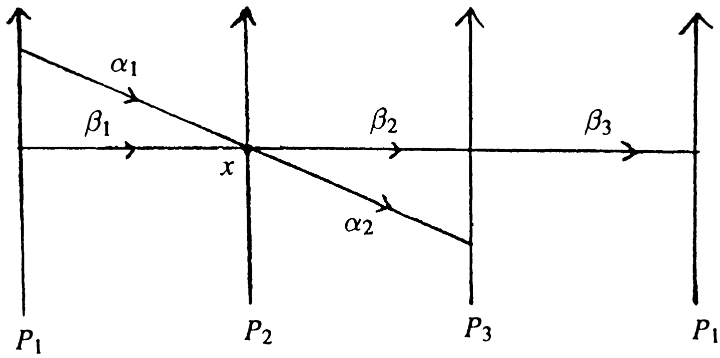

Viewing $\text{Link}(x, S^3)$ as the visual sphere at x (cf. Lemmas 12.2 and 12.3), we have

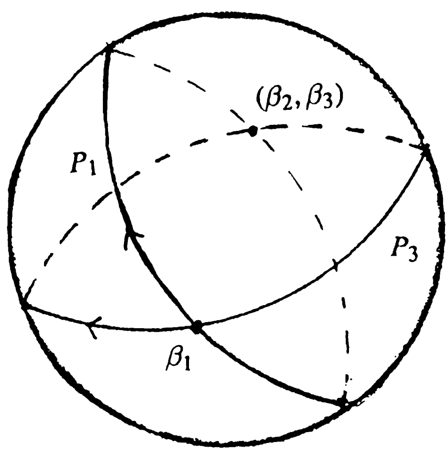

Since β_3 is a continuation of the geodesic β_2, the endpoints of β_2 and β_3 appear as the same point in this link. The great circle through the north pole and β_1 splits the link into two hemispheres, H_1 and H_2.

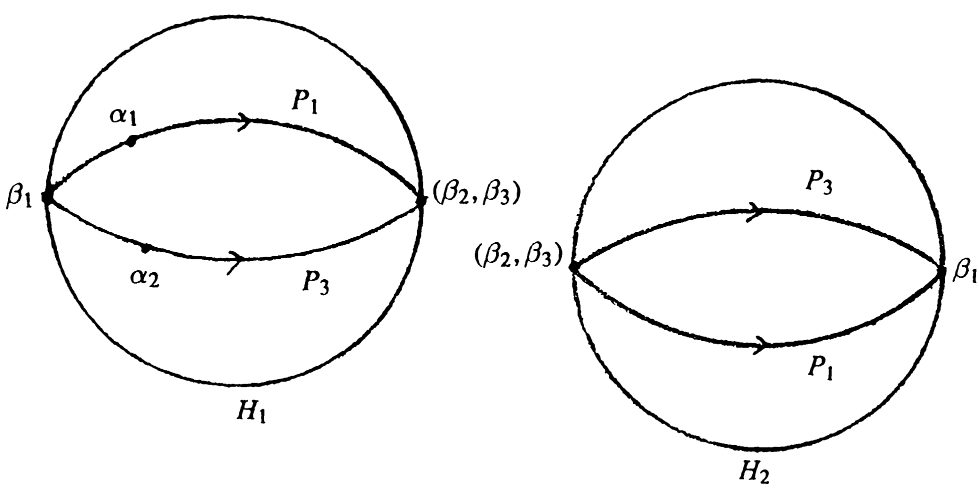

The assumption that the shift of α is negative means that α_1 and α_2 must both appear as points in H_1.

Now consider a lift $\gamma = (\gamma_1, \gamma_2)$ of α with $\tilde{x}$ = endpoint of γ_1 = initial point of γ_2. Let r = branching order of P_2. Then the link of $\tilde{x}$ in $\widetilde{S}^3$ is made up of $2r$ segments isometric, alternately, to H_1 and H_2. The lifts γ_1, γ_2 of α_1, α_2 must appear in segments of type H_1. Clearly, then, $\measuredangle(\gamma_1, \gamma_2) < \pi$ if and only if γ_1 and γ_2 appear in the *same* segment of type H_1. The fact that they appear in the same segment means that the edge path ρ corresponding to γ consists of two edges of an (oriented) triangle σ in Γ.

We claim that this triangle σ is positively oriented. To see this, let $\tilde{\beta} = (\tilde{\beta}_1, \tilde{\beta}_2, \tilde{\beta}_3)$ be the lift of β through $\tilde{x}$ with edgepath σ. By Lemma 11.3, if $\tilde{\beta}$ begins at $\tilde{y}$, then it ends at $\tilde{y} + \pi$ or $\tilde{y} - \pi$ according to whether the triangle σ is positively or negatively oriented. Now the segment of $\mathrm{Link}(\tilde{x}_2, \widetilde{S}^3)$ containing γ and $\tilde{\beta}$ is isometric to its projection H_1 in $\mathrm{Link}(x, S^3)$.

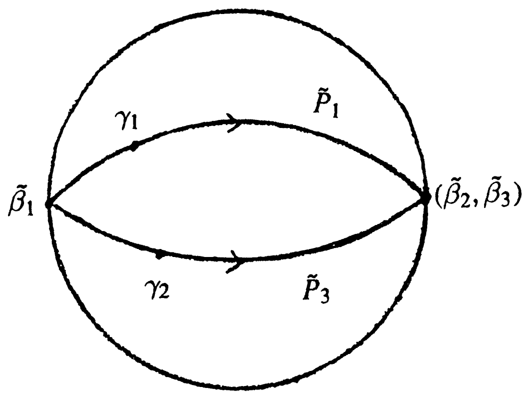

In particular, the endpoint $\tilde{y} \pm \pi$ of $\tilde{\beta}_3$ is a point on $\widetilde{P}_1$ which, viewed from $\tilde{x}_2$ appears the same as the endpoint of $\tilde{\beta}_2$. Noting the orientation of $\widetilde{P}_1$, this point clearly must be $\tilde{y} + \pi$. We conclude that σ is a fundamental cycle. □

PROPOSITION 12.12. *Assuming conditions (i) and (ii) of Theorem 9.1, if γ is a closed, non-orthogonal local geodesic in $\widetilde{S}^3$, then $\ell(\gamma) \geq 2\pi$.*

Proof. Let $\gamma = (\gamma_1, \ldots, \gamma_n)$ be as above and suppose $\ell(\gamma) < 2\pi$. Applying Lemma 12.9 (and Remark 12.10) to the projection of γ in S^3, we may assume that γ_1, γ_2 are of types 1 and 2 while $\gamma_3, \ldots, \gamma_n$ are of type 3, and that $s(\gamma) = \pm\pi$. We will assume $s(\gamma) = -\pi$; the argument for $s(\gamma) = \pi$ is completely analogous.

Let ρ be the closed edgepath in Γ determined by γ, and let $\widetilde{P}$ be the singular circle at which γ begins. Recall from the subsection of Section 11, headed "Edgepaths," the function

$$\tau : H_1(\Gamma) \to \text{ translations of } \widetilde{P}$$

which takes each fundamental cycle to translation by π. By Lemma 11.5, γ is closed if and only if $\tau(\rho) + s(\gamma) \equiv 0 \bmod 4\pi s$, or in other words, $\tau(\rho) = n\pi$ with $n \equiv 1 \bmod 4s$. Now the first two edges e_1, e_2 of ρ are of types 1 and 2. Let v be the vertex (necessarily of type 3) between them. Then e_1 and e_2 are (oriented) edges of some fundamental cycles σ, σ' containing v. Let $d = \sum_{i=1}^{k} \sigma_i$ be the sum of the fundamental cycles $\sigma = \sigma_1, \sigma_2, \ldots, \sigma_k = \sigma'$ between σ and σ' at the vertex v as indicated below,

$d =$ 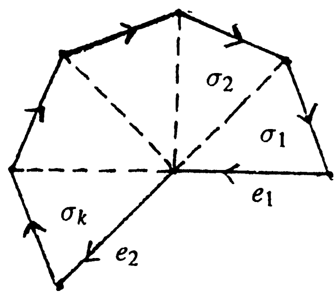

Note that $1 \le k < 2r$, where r = branching order of P_3. In the homology group $H_1(\Gamma)$, $\rho - d$ is homologous to a cycle containing only edges of type 3. Hence,

$$\rho \sim d + \sum_{i=1}^{m} \varepsilon_i c_i,$$

where c_i is the sum of *all* the fundamental cycles about some vertex v_i of type 3. In particular, c_i is the sum of $2r$ fundamental cycles.

It follows that $\tau(\rho) = n\pi$ with

$$n = k + 2r\left(\sum_{i=1}^{m} \varepsilon_i\right).$$

But as noted above, we must have $n \equiv 1 \bmod 4s$. In particular (since r divides s), we have $k \equiv 1 \bmod 2r$. But $1 \le k < 2r$ so $k = 1$. We conclude that d is a single fundamental cycle containing e_1 and e_2 as oriented edges. But the corresponding edges (γ_1, γ_2) of γ have negative shift. It follows from Lemma 12.11 that $\measuredangle(\gamma_1, \gamma_2) < \pi$. This contradicts the assumption that γ is a local geodesic, and hence completes the proof of the proposition. $\square$

Together, Propositions 11.2, 12.1 and 12.12 constitute a proof of Theorem 9.1.

Appendix: Orthogonal Joins of Piecewise Spherical Complexes. Suppose X and Y are finite polyhedra of piecewise constant curvature 1 (cf. Section 2). For convenience, we assume the cellulation on X and Y is simplicial and that no simplex contains two points of distance π. We will define a geodesic metric space $X * Y$, called the orthogonal join of X and Y and prove (Theorem A10) that $X * Y$ satisfies $CAT(1)$ if and only if X and Y each satisfy $CAT(1)$.

To define the orthogonal join, first consider two standard spheres S^n and S^m of radius 1. These can be isometrically embedded orthogonally in the standard sphere S^{n+m+1}. Given any simplices $\sigma^i \subset S^n$, $\tau^j \subset S^m$, the convex hull $\sigma^i * \tau^j$ of σ^i and τ^j in S^{n+m+1} is an $(i+j+1)$-simplex. Now if X and Y are as above, then we can put a metric d on their simplicial join $X * Y$ by assigning the metric described above to $\sigma * \tau$ for each $\sigma \subset X$, $\tau \subset Y$, and letting d be the induced geodesic metric. (That is $d(x,y)$ is the length of the shortest path from x to y.) We call $X * Y$, together with this metric d, the *orthogonal join* of X and Y. We also define the *suspension* and the *cone* of X to be the orthogonal joins

$$\begin{aligned} \Sigma X &= S^0 * X \\ CX &= pt * X. \end{aligned}$$

G. Moussong [M] has shown that ΣX satisfies $CAT(1)$ if and only if X satisfies $CAT(1)$. Our main theorem (Theorem A10) is a generalization of Moussong's result, but our proof is independent.

Topologically, $X * Y$ may be viewed as the quotient space

$$X \times Y \times [0, \tfrac{\pi}{2}]/\sim$$

where $(a,b,0) \sim (a,b',0)$ and $(a,b,\frac{\pi}{2}) \sim (a',b,\frac{\pi}{2})$. Using this description, we denote points in $X * Y$ as equivalence classes $[a,b,t]$. There are obvious embeddings

$$\begin{array}{ccc} X \xrightarrow{i} X * Y, & \qquad & Y \xrightarrow{j} X * Y \\ a \mapsto [a,b,0] & & b \mapsto [a,b,\frac{\pi}{2}] \end{array}$$

By abuse of notation, we will often denote $[a,b,0]$ by a and $[a,b,\frac{\pi}{2}]$ by b. The reader is cautioned, however, that the embeddings above are not, in general, isometric. (Though, as we will see shortly, they are distance preserving for pairs of points of distance at most π in X or Y.) We will denote the metric on $X * Y$ by d, and the metrics on X and Y by d_X and d_Y respectively.

For points $a = [a,b,0]$ and $b = [a,b,\frac{\pi}{2}]$, the segment

$$\gamma_{a,b} = \{[a,b,t] \mid 0 \le t \le \tfrac{\pi}{2}\}$$

lies entirely in one simplex of $X * Y$ and has length $\frac{\pi}{2}$. We call it the *longitude* from a to b.

LEMMA A1. *$\gamma_{a,b}$ is a geodesic segment from a to b, that is, $d(a,b) = \frac{\pi}{2}$.*

Proof. Let γ be a geodesic segment from a to b. It suffices to show that $\ell(\gamma) \ge \frac{\pi}{2}$. Write $\gamma = (\gamma_1, \gamma_2, \dots \gamma_r)$ where each γ_i lies in a single simplex σ_i of $X * Y$. In σ_i we have

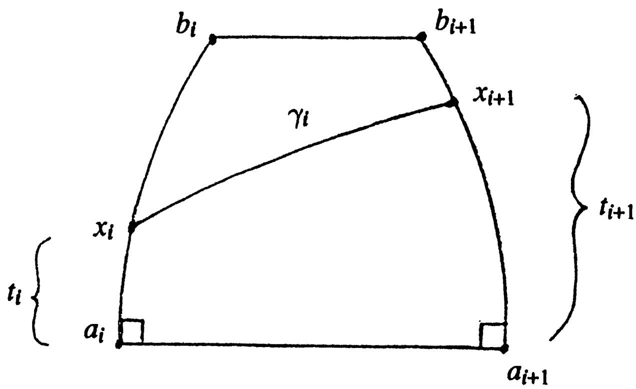

Elementary spherical geometry shows that $\ell(\gamma_i) \geq |t_{i+1} - t_i|$, so $\ell(\gamma) = \Sigma\ell(\gamma_i) \geq \Sigma|t_{i-1} - t_i| \geq \frac{\pi}{2}$. □

As an immediate consequence of Lemma A1, we note that two points in $X * Y$ can have distance at most π. The next lemma gives necessary conditions for two points to have distance exactly π in $X * Y$. Later (Lemma A6), we will show that these conditions are also sufficient.

LEMMA A2. *Let* $x = [a_1, b_1, t_1]$, $y = [a_2, b_2, t_2]$. *Then* $d(x,y) \leq \pi$. *If* $d(x,y) = \pi$ *then one of the following holds:*

(i) $t_1 = t_2 = 0$ *and* $d_X(a_1, a_2) \geq \pi$

(ii) $t_1 = t_2 = \frac{\pi}{2}$ *and* $d_Y(b_1, b_2) \geq \pi$

(iii) $t_1 = t_2 \neq 0, \frac{\pi}{2}$, $d_X(a_1, a_2) \geq \pi$, *and* $d_Y(b_1, b_2) \geq \pi$.

Proof. We first show that $d(x,y) = \pi$ implies $t_1 = t_2$. Say $t_1 < t_2$. Then $t_1 + \left(\frac{\pi}{2} - t_2\right) < \frac{\pi}{2}$. Consider the longitudes connecting b_1, a_1, b_2, and a_2 as indicated:

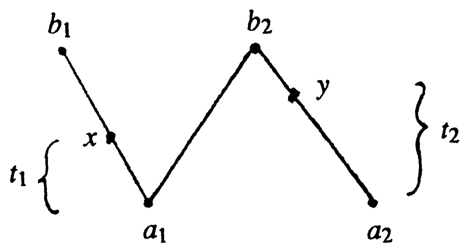

The path from x to y along these longitudes has length $t_1 + \frac{\pi}{2} + \left(\frac{\pi}{2} - t_2\right) < \pi$. Hence $d(x,y) < \pi$. Assume now that $d(x,y) = \pi$.

If $t_1 = t_2 = 0$, then $x, y \in i(X)$ so $d_X(a_1, a_2) \geq d(x, y) = \pi$. Similarly, if $t_1 = t_2 = \frac{\pi}{2}$, then $d_Y(b_1, b_2) \geq d(x, y) = \pi$, so (i) and (ii) follow.

Suppose $t_1 = t_2 \neq 0, \frac{\pi}{2}$, and say $d_X(a_1, a_2) < \pi$. Let γ be a geodesic in X from a_1 to a_2 and let m be its midpoint. The segment γ_1 of γ from a_1 to m has length less than $\frac{\pi}{2}$ and $\gamma_1 * b_1 \subset X * Y$ is isometric to a spherical right triangle.

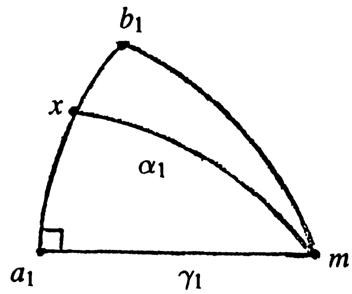

Let α_1 be the geodesic segment in $\gamma_1 * b_1$ from x to m and consider the right spherical triangle with vertices a_1, x, m. Since the legs of this triangle both have length $< \frac{\pi}{2}$, the same is true of the hypotenuse, that is, $\ell(\alpha_1) < \frac{\pi}{2}$. By the same argument, there is a segment α_2 joining m to y with $\ell(\alpha_2) < \frac{\pi}{2}$. Thus (α_1, α_2) is a path from x to y of length $< \pi$, so $d(x, y) < \pi$. □

Next note that there is a retraction

$$r : X * Y \setminus j(Y) \longrightarrow X$$

given by $[a, b, t] \mapsto a$. The following proposition will be fundamental to the analysis of geodesics in $X * Y$.

PROPOSITION A3. *Let γ be a local geodesic in $X * Y$ which contains no points of $j(Y)$ in its interior, int γ. Let $\overline{\gamma}$ denote the closure of r(int γ). Then $\overline{\gamma}$ is either a single point or a local geodesic in X.*

Proof. We begin by analyzing the link of a point $x = [a, b, t]$, $t \neq 0, \frac{\pi}{2}$, in $Z = X * Y$. Recall that $\gamma_{a,b}$ denotes the longitude form a to b. Clearly any simplex $\sigma = \sigma_1 * \sigma_2$ in Z containing a and b must contain $\gamma_{a,b}$. Moreover,

$$\mathrm{Link}(\gamma_{a,b}, \sigma) \cong \mathrm{Link}(a, \sigma_1) * \mathrm{Link}(b, \sigma_2).$$

It follows that

$$\mathrm{Link}(\gamma_{a,b}, Z) \cong \mathrm{Link}(a, X) * \mathrm{Link}(b, Y)$$

and hence

$$\begin{aligned}\operatorname{Link}(x,Z) &\cong \Sigma(\operatorname{Link}(a,X) * \operatorname{Link}(b,Y))\\ &\cong (\operatorname{Link}(a,X)) * (\Sigma \operatorname{Link}(b,Y))\end{aligned}$$

From this analysis we see that if γ is a ray beginning at x and $v \in \operatorname{Link}(x,Z)$ is its tangent vector at x, then the projection $\overline{\gamma}$ of γ to X is constant (in a neighborhood of x) if and only if $v \in \Sigma \operatorname{Link}(b,Y)$. Whereas if $v \notin \Sigma \operatorname{Link}(b,Y)$, then the tangent to $\overline{\gamma}$ at a is the projection of v on $\operatorname{Link}(a,X)$.

Now let γ be as in the proposition. Write $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_r)$ where each γ_i is contained in a simplex of Z. Clearly the projection of each γ_i is either constant or geodesic in X. Suppose some γ_i projects to a single point $a \in X$. Let $x = [a,b,t]$ be the endpoint of γ_i (= starting point of γ_{i+1}). Then the tangent to γ_i at x is a point $v \in \Sigma \operatorname{Link}(b,Y) \subset \operatorname{Link}(x,Z)$. Since γ is a local geodesic, the tangent to γ_{i+1} at x is a point $w \in \operatorname{Link}(x,Z)$ of distance $\geq \pi$ from v. By Lemma A2, this implies that w also lies in $\Sigma \operatorname{Link}(b,Y)$, and hence γ_{i+1} also projects to a constant in X. Similarly, replacing x by the starting point of γ_i, we can show that γ_{i-1} must project to a constant. We conclude that if some γ_i projects to $a \in X$, then so does all of γ, i.e. $\gamma \subset a * Y$.

Now suppose that no segment of γ projects to a constant. Let x, v, and w be as above; that is, x is the endpoint of γ_i, and v, w are the tangents to γ_i and γ_{i+1} in $\operatorname{Link}(x,Z)$. Identifying

$$\operatorname{Link}(x,Z) = (\operatorname{Link}(a,X)) * (\Sigma \operatorname{Link}(b,Y)),$$

we can write $v = [\alpha, \beta, \tau]$, $w = [\alpha', \beta', \tau']$ with $\alpha, \alpha' \in \operatorname{Link}(a,X)$, $\beta, \beta' \in \Sigma \operatorname{Link}(b,Y)$, $\tau, \tau' \in [0, \frac{\pi}{2}]$. Since γ_i, γ_{i+1} do not project to constants, $\tau, \tau' \neq \frac{\pi}{2}$, and the projections $\overline{\gamma}_i, \overline{\gamma}_{i+1}$ have tangents α, α' at a. Now γ is a local geodesic, so the distance from v to w in $\operatorname{Link}(x,Z)$ is at least π. By Lemma A2, this implies that the distance from α to α' in $\operatorname{Link}(a,X)$ is at least π. This is precisely the condition we need to have $(\overline{\gamma}_i, \overline{\gamma}_{i+1})$ be a local geodesic in X. Repeating this argument for each pair of adjacent segments, we conclude that $\overline{\gamma} = (\overline{\gamma}_1, \overline{\gamma}_2, \dots, \overline{\gamma}_r)$ is a local geodesic in X. □

LEMMA A4. *Let γ be a local geodesic in the cone on X, CX, which begins and ends in $i(X) \subset CX$. If $\ell(\gamma) < \pi$, then γ lies entirely in $i(X)$.*

Proof. Since $\ell(\gamma) < \pi$, γ cannot pass through the cone point. Let $\overline{\gamma}$ be the projection of γ onto X, so $\gamma \subset C(\overline{\gamma})$. We can "develop" γ and $\overline{\gamma}$ onto the upper hemisphere S^2_+ of a standard 2-sphere as follows. View S^2_+ as the cone on a circle S^1 of length 2π. Divide γ into segments $\gamma = (\gamma_1, \dots, \gamma_n)$ such that the projection $\overline{\gamma}_i$ of each segment lies in a single simplex of X. Map these onto consecutive geodesic segments $\overline{\beta} = (\overline{\beta}_1, \dots, \overline{\beta}_n)$ of S^1 of length $\ell(\overline{\beta}_i) = \ell(\overline{\gamma}_i)$. Then $C(\overline{\gamma}_i)$

maps isometrically onto $C(\overline{\beta}_i) \subset S^2_+$ in the obvious manner. Under this isometry, γ_i corresponds to a geodesic segment β_i in $C(\overline{\beta}_i)$. Clearly $\beta = (\beta_1, \ldots, \beta_n)$ is a local geodesic in S^2_+ of length $\ell(\beta) = \ell(\gamma) < \pi$. Since β begins and ends on the equator of S^2_+, it must lie entirely in the equator. That is, $\beta = \overline{\beta}$ and hence, $\gamma = \overline{\gamma}$. □

LEMMA A5. *Let $\gamma, \overline{\gamma}$ be as in Proposition A3. If $\ell(\overline{\gamma}) \geq \pi$, then $\ell(\gamma) \geq \pi$.*

Proof. It suffices to show that if $\ell(\overline{\gamma}) = \pi$ then $\ell(\gamma) \geq \pi$. As in the previous lemma, $\overline{\gamma}$ can be developed onto a geodesic segment $\overline{\beta}$ of length π in S^1, so that γ corresponds to a local geodesic β in $\overline{\beta} * Y$ of length $\ell(\beta) = \ell(\gamma)$. But since $\ell(\overline{\beta}) = \pi$, $\overline{\beta}$ may be viewed as the cone on its endpoints. Thus,

$$\overline{\beta} * Y = C(\partial\overline{\beta}) * Y = C(\partial\overline{\beta} * Y).$$

The endpoints of β are the same as the endpoints of $\overline{\beta}$, hence, they lie in $\partial\overline{\beta} * Y$, whereas the interior of β clearly does not lie in $\partial\overline{\beta} * Y$. It follows from Lemma A5 that $\ell(\beta) \geq \pi$, so $\ell(\gamma) \geq \pi$. □

LEMMA A6. *The converse of Lemma A2 holds. That is, $d(x,y) = \pi$ if and only if one of the conditions (i), (ii) or (iii) holds.*

Proof. (i) Let $a_1, a_2 \in i(X)$ be such that $d_X(a_1, a_2) \geq \pi$ and let γ be a geodesic in $X * Y$ connecting them. Then either γ contains a point in $j(Y)$ in which case $\ell(\gamma) \geq \pi$ (Lemma A1) or its projection $\overline{\gamma}$ is a geodesic from a_1 to a_2 in X so $\ell(\overline{\gamma}) \geq \pi$, and by Lemma A5, $\ell(\gamma) \geq \pi$.

(ii) This is similar to (i).

(iii) Let $x = [a_1, b_1, t]$, $y = [a_2, b_2, t]$ with $d_X(a_1, a_2) \geq \pi$ and $d_Y(b_1, b_2) \geq \pi$, and let γ be a geodesic in $X * Y$ from x to y. If γ contains no points of $j(Y)$, then it projects onto a geodesic $\overline{\gamma}$ in X connecting a_1 and a_2. Hence $\ell(\overline{\gamma}) \geq \pi$, so $\ell(\gamma) \geq \pi$. Similarly, if γ contains no points of $i(X)$ then $\ell(\gamma) \geq \pi$. On the other hand, if γ intersects both $i(X)$ and $j(Y)$, say at points $a \in i(X)$ and $b \in j(Y)$, then

$$\ell(\gamma) = d(x,a) + d(a,b) + d(b,y) \geq t + \frac{\pi}{2} + \left(\frac{\pi}{2} - t\right) = \pi.$$

PROPOSITION A7. *If $a_1, a_2 \in X$ with $d_X(a_1, a_2) < \pi$, then any geodesic γ in $X * Y$ from $i(a_1)$ to $i(a_2)$ lies entirely in $i(X)$. In particular, the embedding $i : X \to X * Y$ takes local geodesics to local geodesics and preserves distances for pairs of points of distance $\leq \pi$.*

Proof. Let $a_1, a_2 \in X$. We have already seen that $d_X(a_1, a_2) = \pi$ implies $d(a_1, a_2) = \pi$ (where, as usual, we abuse notation by letting a_1, a_2 denote both the points in X and their images in $X * Y$). Suppose $d_X(a_1, a_2) < \pi$ and let γ be a geodesic in $X * Y$ connecting a_1 and a_2. Then $\ell(\gamma) = d(a_1, a_2) \leq d_X(a_1, a_2) < \pi$, hence γ contains no points of $j(Y)$. If $\gamma \not\subset i(X)$, we may assume without loss of generality that γ intersects $i(X)$ only at its endpoints a_1 and a_2. (If not, replace a_2 by the first intersection point of $(\operatorname{int}\gamma)$ and $i(X)$.) Then γ can be projected onto both X and Y.

Denote these projections by $\overline{\gamma}_X, \overline{\gamma}_Y$, respectively. Now subdivide γ into segments $\gamma = (\gamma_1, \dots, \gamma_n)$ such that each γ_i is contained in a single simplex of $X * Y$. Let $\overline{\gamma}_{X,i}$, $\overline{\gamma}_{Y,i}$ denote the projections of γ_i onto X and Y, respectively. Then, arguing as in the proof of Lemma A4, $\overline{\gamma}_X = (\overline{\gamma}_{X,1}, \dots, \overline{\gamma}_{X,n})$ can be developed onto a local geodesic segment $\overline{\beta}_X = (\overline{\beta}_{X,1}, \dots, \overline{\beta}_{X,n})$ of S^1 with $\ell(\overline{\beta}_{X,i}) = \ell(\overline{\gamma}_{X,i})$. Likewise, $\overline{\gamma}_Y$ can be developed onto a segment $\overline{\beta}_Y = (\overline{\beta}_{Y,1}, \dots, \overline{\beta}_{Y,n})$ of S^1 with $\ell(\beta_{Y,i}) = \ell(\overline{\gamma}_{Y,i})$. For each $i, j, \overline{\gamma}_{X,i} * \overline{\gamma}_{Y,j}$ is contained in a single simplex of $X * Y$ hence, it is isometric to $\overline{\beta}_{X,i} * \overline{\beta}_{Y,j} \subset S^1 * S^1 = S^3$. Now by Lemma A5, $\ell(\gamma) < \pi$ implies $\ell(\overline{\beta}_X) = \ell(\overline{\gamma}_X) < \pi$ and $\ell(\overline{\beta}_Y) = \ell(\overline{\gamma}_Y) < \pi$. Hence, the union $\cup_{i,j} \overline{\beta}_{X,i} * \overline{\beta}_{Y,j} = \overline{\beta}_X * \overline{\beta}_Y$ is a spherical tetrahedron Δ^3 in S^3 containing no pair of antipodal points. The isometries $\overline{\beta}_{X,i} * \overline{\beta}_{Y,j} \cong \overline{\gamma}_{X,i} * \overline{\gamma}_{Y,j}$ give rise to a projection $D : \Delta^3 \to \overline{\gamma}_X * \overline{\gamma}_Y \subset X * Y$. The path γ is the image under D of a unique path $\beta = (\beta_1, \dots, \beta_n)$ with $\beta_i \subset \overline{\beta}_{X,i} * \overline{\beta}_{Y,i}$. Since γ is geodesic in $X * Y$, its lift β must be geodesic in Δ^3 (a shorter path in Δ^3 would project to a shorter path in $X * Y$). Now by construction, $\overline{\beta}_X$ is a geodesic in Δ^3 (in fact it is one edge of the tetrahedron) with the same endpoints as β. But geodesics in Δ^3 are unique, hence, $\beta = \overline{\beta}_X$. It follows that $\gamma = \overline{\gamma}_X \subset i(X)$. □

LEMMA A8. *(i) Let γ be a local geodesic in $X * Y$. If γ contains a longitude, then γ is made up entirely of longitudes.*
(ii) If $a \in i(X)$ and $b \in j(Y)$, then the only local geodesic of length $< \pi$ from a to b is the longitude $\gamma_{a,b}$.

Proof. (i) Let $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ where each γ_i is contained in a simplex of $X * Y$. Suppose some γ_i is a longitude from $a \in i(X)$ to $b \in j(Y)$. Now

$$\operatorname{Link}(b, X * Y) = X * \operatorname{Link}(b, Y)$$

where the image of X in this link corresponds to the longitudes from b to $i(X)$. In particular, $a \in X$ corresponds to γ_i. Since γ is a local geodesic, γ_i and γ_{i+1} must represent points of distance $\geq \pi$ in $\operatorname{Link}(b, X * Y)$; hence they must both lie in X (Lemma A2). Thus γ_{i+1} is also a longitude. A similar argument using $\operatorname{Link}(a, X * Y)$ shows that γ_{i-1} is also a longitude. We conclude inductively that all the segments $\gamma_1, \dots, \gamma_n$ are longitudes.

(ii) Let $a \in i(X)$, $b \in j(Y)$ and let γ be a geodesic of length $< \pi$ connecting a to b. Assume for the moment that γ contains no points of $i(X)$ or $j(Y)$ in its interior. Then γ projects to geodesic segments $\overline{\gamma}_X$ and $\overline{\gamma}_Y$ in X and Y. As in the proof of Lemma A7, we can develop $\overline{\gamma}_X$ and $\overline{\gamma}_Y$ onto segments $\overline{\beta}_X, \overline{\beta}_Y$ of S^1 so that $\Delta^3 = \overline{\beta}_X * \overline{\beta}_Y$ is a spherical tetrahedron containing no pair of antipodal points, and such that the projection $D : \Delta^3 \to \overline{\gamma}_X * \overline{\gamma}_Y$ restricts to an isometry on each simplex. Then γ lifts to a geodesic β in Δ^3 which begins at the initial point $\overline{a}$ of $\overline{\beta}_X$ and ends at the endpoint $\overline{b}$ of $\overline{\beta}_Y$. But there is a unique such geodesic in Δ^3, namely the longitude from $\overline{a}$ to $\overline{b}$. It follows that $\gamma = D(\beta)$ is the longitude from a to b.

More generally, if γ is any geodesic from a to b of length $< \pi$, then γ contains a segment γ' connecting $a' \in i(X)$ and $b' \in j(Y)$ with no points of $i(X)$ or $j(Y)$ in its interior. By the discussion above, γ' must be a longitude, hence, by part (i), γ is made up entirely of longitudes. Since $\ell(\gamma) < \pi$, and each longitude has length $\frac{\pi}{2}$, we conclude that γ is a single longitude. □

We are now ready to prove our main theorems. Recall that the systole of X, sys(X), is the greatest lower bound of the lengths of closed geodesics in X.

THEOREM A9. *Suppose X and Y are finite piecewise spherical polyhedra. Then sys$(X * Y) \geq 2\pi$ if and only if sys$(X) \geq 2\pi$ and sys$(Y) \geq 2\pi$.*

Proof. It follows immediately from Proposition A7 that if X or Y have systole $< 2\pi$, then $X * Y$ has systole $< 2\pi$. Assume now that sys$(X) \geq 2\pi$ and sys$(Y) \geq 2\pi$, and let γ be a closed local geodesic in $X * Y$.

Case 1. $\gamma \cap j(Y) = \emptyset$. In this case, we can project γ onto X. The projection $\overline{\gamma}_X$ is either a closed local geodesic in X or a single point $\overline{\gamma}_X = a \in X$. In the former case, $\ell(\overline{\gamma}_X) \geq 2\pi$ (since sys$(X) \geq 2\pi$), so by Lemma A5, $\ell(\gamma) \geq 2\pi$. In the latter case, $\gamma \subset a * Y \cong C(Y)$. Any ray out of the cone point a is a longitude and hence intersects $j(Y)$. This contradicts our assumption that $\gamma \cap j(Y) = \emptyset$, so we may assume that γ does not contain a. But then we also have a projection $\overline{\gamma}_Y$ of γ onto $j(Y)$. This projection cannot be a single point (since that would mean that γ was a longitude), hence, $\overline{\gamma}_Y$ is a closed local geodesic in Y. It follows that $\ell(\overline{\gamma}_Y) \geq 2\pi$, so by Lemma A5, $\ell(\gamma) \geq 2\pi$.

Case 2. $\gamma \cap i(X) = \emptyset$. Same as case 1.

Case 3. γ contains points of both $i(X)$ and $j(Y)$. Write $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_r)$ where each γ_i connects a point in $i(X)$ to a point in $j(Y)$. Note that r is necessarily even since γ is closed. If some γ_i is a longitude, then by Lemma A8 (i), all the γ_i's are longitudes and hence have length $\frac{\pi}{2}$. If $r = 2$, then $\gamma = (\gamma_{a,b}, \gamma_{a,b}^{-1})$. But clearly this is not a local geodesic at the endpoints a, b of $\gamma_{a,b}$. Hence, we must have $r \geq 4$, so $\ell(\gamma) \geq 2\pi$. If no γ_i is a longitude, then by Lemma A8 (ii), $\ell(\gamma_i) \geq \pi$ for all i, so $\ell(\gamma) \geq r\pi \geq 2\pi$. □

Finally, to prove our last theorem, we review some facts from Chapter 1. Recall that a geodesic metric space Z has curvature ≤ 1 if Z satisfies $CAT(1)$ *locally*. Suppose Z is a finite piecewise spherical polyhedron. Then by Lemma 1.3, X satisfies $CAT(1)$ globally if and only if Z has curvature ≤ 1 and systole $\geq 2\pi$. On the other hand, by Theorem 3.1, parts (1) and (2), Z has curvature ≤ 1 if and only if the link of every point in Z satisfies $CAT(1)$. (While the proof of some parts of Theorem 3.1 use properties of suspensions, parts (1) and (2) use only an independent result of Ballman [GH, Chap 10].) Thus Z satisfies $CAT(1)$ if and only if $\mathrm{sys}(Z) \geq 2\pi$ and $\mathrm{Link}(x, Z)$ satisfies $CAT(1)$ for all $x \in Z$.

THEOREM A10. *Assume X and Y are nonempty, finite, piecewise spherical polyhedra. Then the following are equivalent.*

(i) $X * Y$ *has curvature* ≤ 1.

(ii) $X * Y$ *satisfies* $CAT(1)$.

(iii) X *and* Y *both satisfy* $CAT(1)$.

Proof. ii) $\Rightarrow$ i). Since curvature ≤ 1 means, by definition, that $CAT(1)$ holds locally, this implication is obvious.

i) $\Rightarrow$ iii). First consider the special case $X = S^0$, $X * Y = \Sigma Y$. Let x_0 be a cone point of ΣY. Then $\mathrm{Link}(x_0, \Sigma Y) = Y$. By the discussion above, therefore, curvature ≤ 1 for ΣY implies $CAT(1)$ for Y. For general X, let x be a point in a maximal simplex σ of X and let $n = \dim \sigma$. Then $\mathrm{Link}(x, X) = \mathrm{Link}(x, \sigma) = S^{n-1}$, so $\mathrm{Link}(x, X * Y) = \mathrm{Link}(x, X) * Y = \Sigma^n Y$. If $X * Y$ has curvature ≤ 1, then $\Sigma^n Y$ satisfies $CAT(1)$, hence Y satisfies $CAT(1)$.

iii) $\Rightarrow$ ii). Assume X and Y both satisfy $CAT(1)$. Then X and Y have systole $\geq 2\pi$, hence, by Theorem A9, the same holds for $X * Y$. Thus, it suffices to show that the link of every point in $X * Y$ satisfies $CAT(1)$. For this, we proceed by induction on the dimension of $X * Y$. If $\dim(X * Y) = 1$, then links of points in $X * Y$ are dimension 0 and satisfy $CAT(1)$ vacuously. Suppose $\dim(X * Y) = n$. Let $z = [a, b, t]$ be a point in $X * Y$. Then $\mathrm{Link}(z, X * Y)$ is one of the following:

(1) $\Sigma(\mathrm{Link}(a, X) * \mathrm{Link}(b, Y))$, if $t \neq 0, \frac{\pi}{2}$

(2) $\mathrm{Link}(a, X) * Y$, if $t = 0$

(3) $X * \mathrm{Link}(b, Y)$, if $t = \frac{\pi}{2}$.

Now X and Y satisfy $CAT(1)$ by assumption and hence, so do $\mathrm{Link}(a, X)$ and $\mathrm{Link}(b, Y)$ for all $a \in X$, $b \in Y$. By induction, we conclude that the same holds for $\mathrm{Link}(z, X * Y)$. $\square$

DEPARTMENT OF MATHEMATICS, THE OHIO STATE UNIVERSITY, COLUMBUS, OH 43210-1174

REFERENCES

[AB] J. Alonzo and M. Bridson, Semihyperbolic groups, preprint, 1990.

[Bo] N. Bourbaki, *Groupes et Algèbres de Lie,* Hermann, Paris, 1968, Chapters IV–VI

[B] G. Bredon, *Introduction to Compact Transformation Groups,* Academic Press, New York, 1972.

[Br] M. Bridson, On the existence of flat planes in spaces of nonpositive curvature, preprint, 1992.

[Bri] E. Brieskorn, Beispiele differential topologie von singularitäten, *Invent. Math.*, **2** (1966), 1–74.

[Bro] K. S. Brown, *Buildings,* Springer-Verlag, New York, 1989.

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

[CE] J. Cheeger and D. Ebin, *Comparison Theorems in Riemannian Geometry,* North-Holland, Amsterdam-Oxford, 1975.

[Ch] C. Chevalley, Invariants of finite groups generated by reflections, *Amer. J. Math.*, **77** (1955), 778–782.

[C] H. S. M. Coxeter, *Regular Complex Polytopes,* Cambridge Univ. Press, Cambridge, 1974.

[GH] E. Ghys and P. de la Harpe (eds.), *Sur les Groupes Hyperboliques d'après Mikhael Gromov*, *Prog. Math.*, vol. 83, Birkhäuser, Boston, 1990.

[G] M. Gromov, Hyperbolic groups, *Essays in Group Theory*, (S. M. Gersten, ed.), Springer-Verlag, New York, 1987, pp. 75–264.

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

[Gv] K. Grove, Metric differential geometry, *Differential Geometry, Proceedings of the 1985 Nordic Summer School in Lyngby (Denmark), Lecture Notes in Math.*, vol. 1263, Springer-Verlag, New York, 1986, pp. 171–227.

[H] A. Haefliger, Groupoides d'holonomie et classifiants, *Asterisque*, **116** (1984), 170–97.

[Hm] H. Hamm, Exotische Sphären als Umgebungsränder in Speziellen komplexen Räumen, *Math. Ann.*, **197** (1972), 44–56.

[Mi] J. Milnor, Singular Points of Complex Hypersurfaces, *Ann. of Math. Stud.*, vol. 61, Princeton Univ. Press, Princeton, 1968.

[M] G. Moussong, Hyperbolic Coxeter groups, Ph.D. thesis, Ohio State Univ., 1988.

[S] I. Satake, On a generalization of the notion of manifold, *Proc. Nat. Acad. Sci.*, **42** (1956), 359–363.

[Sp] B. Spieler, Non-positively curved orbihedra, Ph.D. thesis, Ohio State Univ., 1992.

[T] W. Thurston, The geometry and topology of 3-manifolds, lecture notes, Princeton, 1977.