Bounds on Volume Increase Under Dehn Drilling Operations

Martin Bridgeman *

June 14, 1995

Abstract

In this paper we investigate how the volume of hyperbolic manifolds increases under the process of removing a curve, that is, Dehn drilling. If the curve we remove is a geodesic we are able to show that for a certain family of manifolds the volume increase is bounded above by $\pi \cdot l$ where l is the length of the geodesic drilled. Also we construct examples to show that there is no lower bound to the volume increase in terms of a linear function of a positive power of length and in particular volume increase is not bounded linearly in length.

1 Introduction

Let M and M' be three-dimensional manifolds such that M' is homeomorphic to $M - \alpha$, where α is a simple closed curve in M. Then we say that M' is obtained by Dehn drilling on M and α is the curve drilled.

If M and M' are complete hyperbolic manifolds of finite volume then Gromov ([Th]) showed that the volume of M' is greater than the volume of M. Therefore hyperbolic Dehn drilling always increases volume, but what this increase depends on is unknown. In the following we study the case of drilling out a geodesic and show how the volume increase is related to the length of the geodesic drilled.

Conjecture: 1 *Dehn drilling a geodesic increases volume by at most $\pi \cdot L$ where L is the length of the geodesic to be drilled.*

In this paper we will prove this conjecture for a class of Dehn drillings as well as give motivation for why the conjecture should be true in general. We also show that in general there is no linear lower bound on the volume increase by showing that there is no linear lower bound for the set of Dehn drillings we consider.

*Research at MSRI is supported in part by NSF grant DMS-9022140.

2 Motivation

The motivation for the conjecture and the idea behind the proof in the special case to be considered is the following. If M' is obtained by hyperbolic Dehn drilling a geodesic α in M then we first construct a continuous one-parameter family of hyperbolic cone-manifolds $M_\theta, 0 < \theta \leq 2\pi$ interpolating between M and M'. By interpolating we mean that $M_{2\pi} = M$, $\lim_{\theta \to 0} M_\theta = M'$ and M_θ is homeomorphic to M with cone axis α and cone angle θ.

Having constructed a one-parameter family of cone-manifolds interpolating between M and M', we apply results of Craig Hodgson ([Ho]) which describe the derivative of volume of a smooth one-parameter family of cone manifolds.

Before we give Hodgson's formula, we state the Sclafli formula for the variation of volume of a one-parameter family of polyhedra in a space of constant curvature, on which Hodgson's formula is based.

Theorem 1 *(Sclafli, Hodgson [Ho]) Let X_t be a smooth one-parameter family of polyhedra in a simply connected n-dimensional space of constant curvature K. Then the derivative of the volume of X_t satisfies the equation:*

$$(n-1)KdVol(X_t) \quad = \quad \sum_F V_{n-2}(F)d\phi_F$$

where the sum is over all co-dimension two faces of X_t, V_{n-2} is the $(n-2)$-dimensional volume and ϕ_F is the dihedral angle at F.

In the cone manifold case we have the following.

Theorem 2 *(Hodgson[Ho]) Let C_t be a smooth family of (curvature K) cone-manifold structures on a manifold with fixed topological type of singular locus. Then the derivative of volume of C_t satisfies*

$$(n-1)KdVol(C_t) \quad = \quad \sum_\Sigma V_{n-2}(\Sigma)d\phi_\Sigma$$

where the sum is over all components Σ of the singular locus of C and θ_Σ is the cone angle along Σ.

Therefore in our case where the singular locus is a single cone axis α we have

$$\begin{aligned} dV &= -\frac{1}{2}l(\theta)d\theta \\ \Delta V &= \frac{1}{2}\int_0^{2\pi} l(\theta)d\theta \end{aligned}$$

where θ is the cone angle along α, $l(\theta)$ is the length of α in M_θ and where ΔV is the volume increase under Dehn drilling.

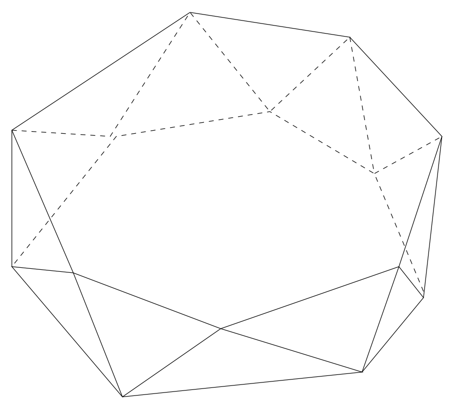

Figure 1: Ideal hyperbolic drum

Therefore if the function l is monotonically increasing in θ then it follows that

$$\Delta V \leq \pi \cdot L$$

where L is the length of the geodesic α in M.

The conjecture and the proof of the conjecture for the class of Dehn drillings we consider, is based on this two-fold approach of interpolation and monotonicity.

Basic Polyhedra

We will now prove that for a Dehn drilling operation on *basic polyhedra*, the conjecture is true. This Dehn drilling operation has some interesting properties including giving a new way to enumerate all basic polyhedra and hence all link projections ([B2]). A basic polyhedron is an ideal hyperbolic polyhedron P with all dihedral angles being right angles. Therefore all vertices have valence four. Let $\mathcal{B}$ be the set of all basic polyhedra. Examples of basic polyhedra are the ideal hyperbolic drums $\{T_n\}_{n\geq 3}$ which have two faces being regular ideal n-gons, corresponding to the top and bottom of the drum and every other face being an ideal triangle (figure 1).

As all the dihedral angles are $\pi/2$, basic polyhedra can also be thought of as orbifolds. In [B1] it was shown that if P is a basic polyhedron then it has a four-fold cover (in the orbifold sense) L_P which is a hyperbolic link complement in S^3 having a boundary component for each vertex of P.

Surgery and Dehn drilling

If P is a basic polyhedron and F is a non-triangular face of P then we define *surgery* on P as follows. Choose two non-adjacent edges e_1, e_2 of F and pinch them together to form a new combinatorial polyhedron $\overline{P}$ (figure 2). By Andreev's theorem ([Th]), $\overline{P}$ has a realization P'

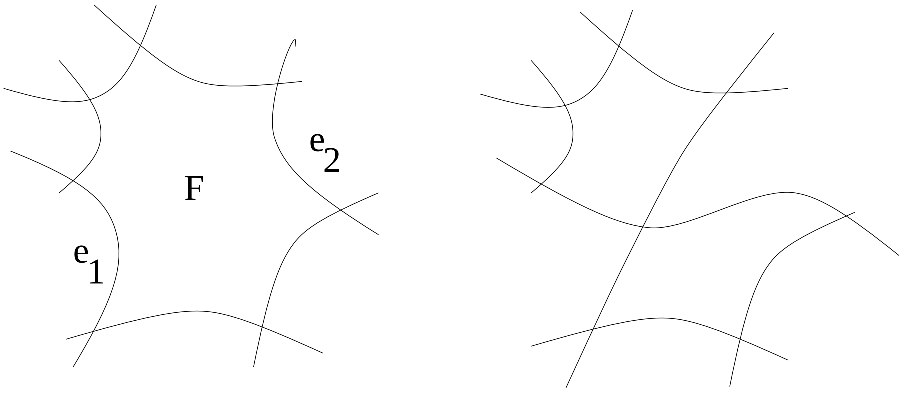

Figure 2: Surgery

which is an ideal hyperbolic polyhedron with all dihedral angles right angles and therefore is also a basic polyhedron. We say P' is obtained by surgery on P and therefore surgery is an operation defined on the set $\mathcal{B}$.

In the sense of orbifolds, P' is obtained by Dehn drilling the unique geodesic perpendicular to both e_1 and e_2 in P and if we lift to the four-fold covers, we see that $L_{P'}$ is obtained by Dehn drilling a simple closed geodesic in L_P ([B2]). Thus surgery corresponds to a Dehn drilling operation.

We define a partial ordering on $\mathcal{B}$ by saying $P \prec P'$ if P' is the result of repeated surgeries starting with P. This partial ordering has as initial objects the set $\{T_n\}_{n\geq 3}$ mentioned above ([B2]) and therefore we can enumerate all basic polyhedra by repeated surgery on these initial objects.

Theorem 3 *Surgery increases volume by at most $\pi \cdot L$ where L is the length of the geodesic drilled.*

To prove this theorem we must first interpolate between P and P' where P' is obtained by surgery on P.

3 Interpolating Surgery

Let P be a basic polyhedron with F a non-triangular face of P and e_1, e_2 be two non-adjacent edges of F. Also let g be the unique geodesic in F perpendicular to e_1, e_2. To interpolate between P and P' we need a continuous one-parameter family of polyhedra $P_\theta, 0 < \theta \leq \pi$ which have combinatorial type of $P \cup \{g\}$, with dihedral angle θ along g and all other edges having the same dihedral angles as the corresponding edge of P and with $P_\pi = P, \lim_{\theta\to 0} P_\theta = P'$. Intuitively P_θ is obtained by bending P along g an amount $\pi - \theta$. To find such a one-parameter family of polyhedra we use results of Hodgson and Rivin ([HR]) which characterizes compact hyperbolic polyhedra in terms of their *spherical duals*.

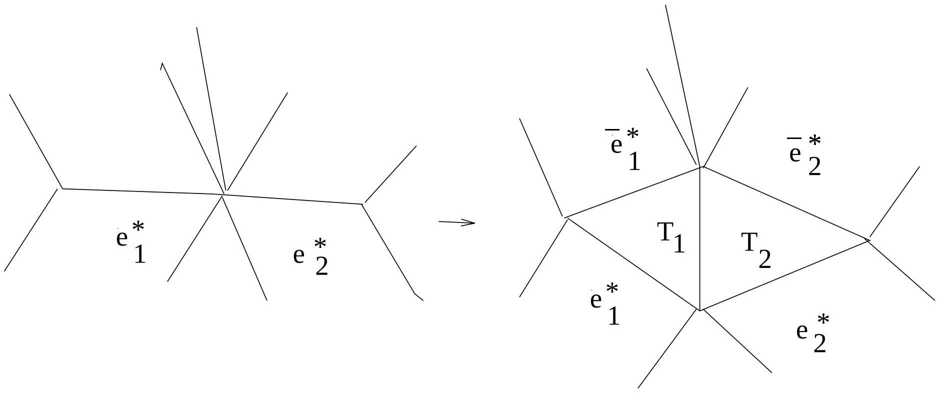

Figure 3: Cutting open S to make S_θ

Spherical Dual of a Polyhedron

If P is a hyperbolic polyhedron then we form its spherical dual P^* which is a spherical complex combinatorially dual to the polyhedron P as follows. Each vertex corresponds to a spherical polygon with edges having lengths being the exterior dihedral angles of the edges incident to the vertex and with face angles being the complementary angles of the face angles at the vertex (complement of $\alpha = \pi - \alpha$). Similarly an ideal vertex corresponds to a spherical hemisphere with vertices dividing the equator into arcs of length less than π. Let $S = P^*$ and define spherical complexes S_θ by cutting S open at e_1^*, e_2^* (dual edges to e_1, e_2), and attaching a spherical bigon B with angle $\pi - \theta$. B is further divided into two triangles T_1, T_2 by joining the two points on the equator by an edge g^*. If P_θ is realisable as a hyperbolic polyhedron then its spherical dual is S_θ with g^* dual to g. Note that each of e_1^*, e_2^* split into two edges $e_i^*, \overline{e}_i^*, i = 1, 2$ (figure 3). Thus we have a continuous one-parameter family of spherical complexes $S_\theta, 0 \leq \theta \leq \pi$ which interpolates between the spherical duals of P and P'. Before we can interpolate between P and P' we need to know more about the map between hyperbolic polyhedra and spherical complexes.

The following theorem due to Craig Hodgson and Igor Rivin gives a characterization of compact hyperbolic polyhedra in terms of their spherical duals.

Theorem 4 *([HR]) A metric space (M, g) homeomorphic to S^2 can arise as the spherical dual P^* of a compact hyperbolic polyhedron P in H^3 if an only if the following hold.*

(a) The metric has constant curvature 1 away from a finite collection of cone points c_i.

(b) The cone angles at c_i are greater than 2π.

(c) The lengths of closed geodesics of (M, g) are all strictly greater than 2π.

Moreover P is uniquely determined.

Approximating by compact Polyhedra

Since the polyhedron P_θ is not compact, we will find a family of compact hyperbolic polyhedra $P_\theta^t, t > 0$ approximating P_θ by applying the above result. We define spherical complexes S_θ^t by modifying S_θ so that it is the dual of a compact hyperbolic polyhedron as follows. We increase the length of every edge of S_θ which does not belong to B by a factor of $(1+t)$ and add a vertex v_i at the pole of each hemisphere H_i, and join v_i by an edge to each vertex on the equator of H_i. Note that if we let $t = 0$ then $S_\theta^0 = S_\theta$. It will be shown, using the above characterization (theorem 4), that S_θ^t is the dual of a compact hyperbolic polyhedron P_θ^t. Furthermore P_θ^t has combinatorial type of P_θ with its ideal vertices truncated and the truncating faces are perpendicular to the faces of P_θ^t intersecting it. Thus in the Klein model, all the planes intersecting a truncating face meet in a single point outside S_∞^2 corresponding to the unique point perpendicular to the truncating face. A *generalized hyperbolic polyhedron* is when we allow vertices to be outside S_∞^2 in the Klein model and we call the vertices outside of S_∞^2 *hyper-ideal.* In this paper we will not distinguish between a generalized polyhedron and the unique polyhedron obtained by truncating the hyper-ideal vertices. Thus P_θ^t corresponds to a generalized hyperbolic polyhedron with combinatorial type P_θ and with the vertices corresponding to ideal vertices of P_θ being hyper-ideal. Intuitively as t tends to 0 these hyper-ideal vertices tend to S_∞^2 and become ideal.

Lemma: 1 *Given $\theta, 0 < \theta \leq \pi$ then there exists $\delta > 0$ such that the spherical complex S_θ^t is the dual of a unique compact hyperbolic polyhedron P_θ^t for $0 < t < \delta$.*

Proof : Obviously S_θ^t is a spherical complex satisfying (a) and (b) of the characterization (theorem 4), and therefore we need only show that there are no geodesics of length $\leq 2\pi$. Let γ^t be a geodesic in S_θ^t. If γ^t intersects the interior of a hemisphere H in a single arc then $l(\gamma^t|_{H^o}) = \pi$. Therefore if γ^t intersects more than two hemispheres in their interiors then $l(\gamma^t) \geq 3\pi$.

Case 1: γ^t doesn't intersect the interior of bigon B. Then we get a curve in S_π^t by removing the bigon and stitching back the edges. This curve we also call γ^t and it is a geodesic in S_π^t.

If γ^t intersects the interior of two hemispheres H_1, H_2 and if γ^t doesn't contain an edge then H_1, H_2 intersect in points a distance $\geq \pi$ apart giving a contradiction. Thus γ^t must contain an edge and $l(\gamma^t) \geq 2\pi + \pi/2 > 2\pi$.

If γ^t intersects the interior of only one hemisphere H, then it must enter and leave H along an edge. We replace the arc in H^o by an edge path of the same length by going along the boundary of H. This gives us a new geodesic which we still call γ^t satisfying

$$l(\gamma^t) \quad \geq \quad \pi/2 + \pi/2 + \pi/2(1+t) + \pi/2(1+t) = 2\pi + t\pi > 2\pi$$

Therefore all geodesics in S_θ^t not intersecting the interior of B have length bounded away from 2π.

Case 2: γ^t intersects interior of B. There are a number of possibilities.

If γ^t enters B at either the apex of T_1 or T_2 then it has to leave at the other apex and $l(\gamma^t|_{B^o}) = \pi$. We replace this arc by edge path $e_1^* \cup e_2^*$ or $\overline{e}_1^* \cup \overline{e}_2^*$ depending on which gives a geodesic as a result (at least one must give a geodesic). Thus we have a geodesic which has the same length as γ^t and doesn't intersect the interior of B. Therefore γ^t has length $> 2\pi$ by the first case.

If γ^t intersects the interior of two hemispheres H_1, H_2 and both are not incident to B then γ^t must contain an edge of S_θ^t and therefore $l(\gamma^t) \geq 2\pi + \pi/2 > 2\pi$. If either H_1, H_2 are incident to B only at an apex then γ^t intersects B in an arc from one apex to the other and has been already shown to be $> 2\pi$. Therefore either γ^t contains an edge and is longer than 2π or H_1, H_2 intersect at a point p along their boundaries where γ^t traverses from H_1 to H_2. But H_1, H_2 correspond to hemispheres in S_π^t whose intersection is contained in $e_1^* \cup e_2^*$ and thus H_1, H_2 cannot intersect at point p. Therefore any geodesic in S_θ^t intersecting the interior of two or more hemispheres has length bounded away from 2π.

If γ^t only intersects the interior of one hemisphere H and it enters B^o along an edge and leaves along an edge then it must intersect B^o along the equatorial edge. Therefore there is a geodesic in S_π^t of length $l(\gamma^t) - (\pi - \theta)$ obtained by removing the bigon as before and taking what is left of γ^t as the curve. This must be of length $> 2\pi$, therefore γ^t has length $\geq 2\pi + (\pi - \theta) > 2\pi$. If γ^t enters through H and leaves along an edge at vertex v of B (v is not an apex), then we first replace the interior arc in H by a boundary edge path b of H of the same length, we call this new curve γ^t also. Now γ^t contains two sides of a spherical triangle $\overline{T}$ in B. H intersects B in an edge e and let v_1 be the vertex of B contained in $b \cap e$. Now replace the part of γ^t consisting of two sides of the spherical triangle $\overline{T}$ by the edge of B joining v, v_1 (figure 4). This gives a geodesic of smaller length that is an edge path and therefore $l(\gamma^t) > 2\pi$.

If γ^t intersects no hemisphere in its interior then γ^t is an edge geodesic containing the equatorial edge of B. Removing the bigon we get a geodesic in S_π^t of smaller length which implies $l(\gamma^t) > 2\pi$.

Since no geodesic of S_θ^t has length $\leq 2\pi$, then S_θ^t is the spherical dual of a unique compact hyperbolic polyhedron P_θ^t (by theorem 4). ■

To find what is the combinatorial structure of P_θ^t we need to find a cell division of S_θ^t into convex spherical polygons by geodesics of length less than π and whose endpoints are cone points. This is how we described S_θ^t originally, but this may not be the only cell division possible. To show that it is, consider the apex v of one of the triangles T_1, T_2. The only other vertices a distance less than π from v are those vertices connected to v by an edge in the original cell division and therefore the edges from v must be a subset of the original edges from v. Also all edges must be included as if any are left out then there is an angle $\geq \pi$ subtended at v. Now considering a vertex v_1 on the equator we have an angle of π inside the bigon B at v_1. Therefore the equatorial edge of B must also be included in the division. Using the fact that any geodesic intersecting the interior of a hemisphere must have length $\geq \pi$ and that no vertex can have an angle of $\geq \pi$ subtended at it, it is easy to show that all other edges of the original division are necessary. There cannot be any more edges as any

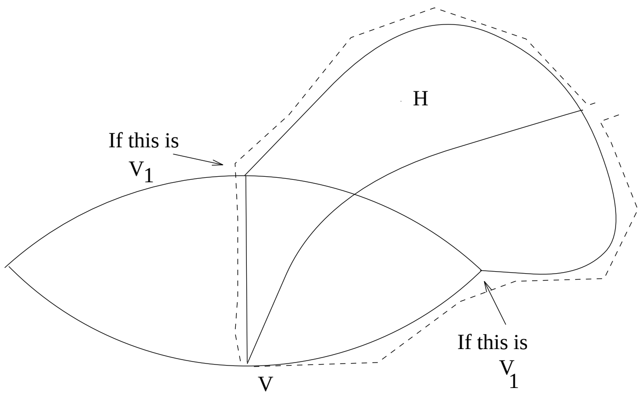

Figure 4: Replacing geodesic by edge geodesic

other edge from vertex to vertex would cross another edge in an interior point which can't happen. Therefore P_θ^t has dual given by the original cell division of S_θ^t.

P_θ^t is a general hyperbolic polyhedron with combinatorial type $P \cup g$ and hyper-ideal vertices corresponding to the ideal vertices of P. The dihedral angles between a truncating plane and the planes it intersects are all $\pi/2$. Also the dihedral angle along the edge corresponding to g is θ and except for the four edges incident to g which have dihedral angles $\pi/2$, all others have dihedral angle $\pi/2(1-t)$. Similarly the face angles on all truncating faces are $\pi/2(1-t)$ except where the four edges incident to g intersect. The face angles on the two faces perpendicular to g, at the vertices incident to g, are θ, and all other face angles are $\pi/2$. As t tends to zero all dihedral angles approach $\pi/2$ except along g which has fixed dihedral angle θ.

4 Monotonicity and Volume Increase

Lemma 1 shows that P_θ^t exists for $0 < t \leq k(\theta)$, $0 < \theta \leq \pi$ where $k(\theta) > 0$. We can choose $k(\theta) = c \cdot \theta$ where c is a positive constant. Hodgson and Rivin showed ([HR]) that the map from the space of compact hyperbolic polyhedra to their spherical duals is a homeomorphism and hence P_θ^t is a continuous two-parameter family of polyhedra in the above domain. When all the dihedral angles are acute we have continuity by Alexandrof ([A]) and therefore $lim_{t\to 0} P_\pi^t = P$. Thus we extend P_θ^t by letting $P_\pi^0 = P$. Also if $0 < \theta \leq \pi/2$ then $lim_{t\to 0} P_\theta^t = P_\theta^0$ where P_θ^0 is the unique hyperbolic polyhedron with combinatorial type of $P \cup \{g\}$ and dihedral angles $\pi/2$ except along g which has dihedral angle θ. We can now take the limit as (t, θ) tends to $(0, 0)$ and we get that $lim_{(t,\theta)\to(0,0)} P_\theta^t = P'$. Therefore P_θ^t is well defined and continuous on the triangle $0 \leq t \leq c \cdot \theta$, $0 \leq \theta \leq \pi$ except for along the open interval of length $\pi/2$ from $(0, \pi/2)$ to $(0, \pi)$.

We now choose a continuous one-parameter family of polyhedra interpolating between P and P' by taking the image of a piecewise linear path joining $(0, \pi)$ and $(0, 0)$ in the domain of $P_{(\cdot)}^{(\cdot)}$ (figure 5). Thus the one-parameter family X interpolating between P and P' is the product of four continuous one-parameter families $X_i, i = 1, \ldots, 4$ obtained by taking the image of the path joining the four points $(0, \pi), (0, t), (c/t, t), (c/t, 0), (0, 0)$ by alternately horizontal and vertical lines.

The volume change from P to P' ΔV satisfies

$$\Delta V = \Delta V_1 + \Delta V_2 + \Delta V_3 + \Delta V_4$$

where ΔV_i is the volume change over X_i.

By continuity given $\epsilon > 0$ we can choose a $t > 0$ such that

$$\begin{aligned} |\Delta V_i| &\leq \epsilon \quad i \neq 2 \\ |l^t(\pi) - l^0(\pi)| &\leq \epsilon \end{aligned}$$

where $l^0(\pi) = l$, the length of the original geodesic in P.

If l^t is monotonic then

$$\begin{aligned} 0 \leq \Delta V_2 &\leq 1/2(\pi - t/c) l^t(\pi) \leq \pi/2 \cdot l^t(\pi) \\ \Rightarrow 0 \leq \Delta V &\leq \pi/2 \cdot l^t(\pi) + 3\epsilon \leq \pi/2 \cdot l + \tilde{\epsilon} \end{aligned}$$

Therefore by choosing smaller values of t we have

$$\Delta V \leq \pi/2 \cdot l$$

Taking the four-fold cover of P given by the hyperbolic link complement L_P we have

$$\Delta V' = 4 \cdot \Delta V \leq 4 \cdot (\pi/2 \cdot l) = 2\pi l = \pi L$$

where $\Delta V'$ is the volume increase from L_P to $L_{P'}$ and L is the length of the geodesic drilled from L_P to obtain $L_{P'}$.

Monotonicity

The following theorem is based on Cauchy Rigidity and is the step needed to prove monotonicity, but first a definition.

Definition: 1 *A* signed polyhedron P *is a polyhedron whose edges are labeled with either* $+$, $-$ *or* 0.

A vertex v of a signed polyhedron has *type* given by listing the signs on the edges incident to v in a clockwise or anti-clockwise manner (with the usual equivalence) and a face F is called a zero-face if all its edges are marked with zero.

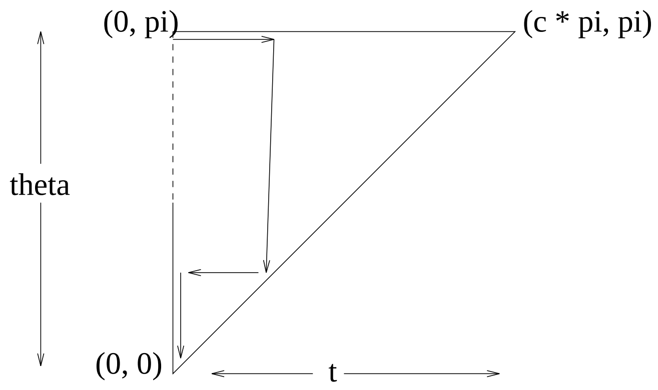

Figure 5: 2-parameter domain and path from P to P'

Theorem 5 *Let P be a trivalent signed polyhedron such that, except for two disjoint faces F_1, F_2, every face is either a zero-face or has at least four sign changes. Then the space obtained by collapsing all zero-edges $\tilde{P}$ is topologically a bouquet of spheres where either*

(1) Each sphere has the trivial cell division (a point and a disc).

or

(2) There is exactly one sphere $\tilde{S}$ in the bouquet with a non-trivial cell division. $\tilde{S}$ has two cells with no sign changes and every other cell has exactly four sign changes. Also all vertices of $\tilde{S}$ have type $(++-)$, $(--+)$ or $(+-+-)$ and therefore valence three or four.

Proof : Let Σ be the set of edges of P and $\Sigma = \Sigma_+ \cup \Sigma_- \cup \Sigma_0$ be the partition into $+$, $-$ and 0 edges. If $|\Sigma_0| = 0$ then construct a vector field on P as follows. In a neighborhood of Σ the vector field is transverse to Σ_+ and parallel to Σ_-. The index of a vertex v can easily be calculated from its type. One way is to read the type of a vertex from left to right, count $-1/4$ for sign changes, $-1/2$ for same sign and then adding 1. This is never positive if the valence is at least three and if the valence of v is greater than four its index is always negative. In our case the vertices all have valence three but in general only vertices of type $(++-), (--+)$ and $(+-+-)$ (which have index zero), have non-negative indices (figure 6).

The vector field is extended over a face F as follows. As we go round the edges of F the sign changes $2n$ times , which means that there are n disjoint $+$ segments joined together by n disjoint $-$ segments. Choose a center point c_F in the interior of F and for each $+$ segment we join an interior point of it to the center c_F by a prong and extend the vector field in a neighborhood of these n prongs in the obvious way. Cutting along the prongs divides F into n pieces over which the vector field can easily be extended. The index of the critical point c_F is $1 - n/2$ (figure 7). Note that if the face has no sign changes then the vector field can still be extended with c_F having index 1, in keeping with the formula. The interior of each face has at most one critical point and all other critical points are vertices of P. Each vertex contributes a non-positive index and each face $F \neq F_1, F_2$ has at least four sign changes,

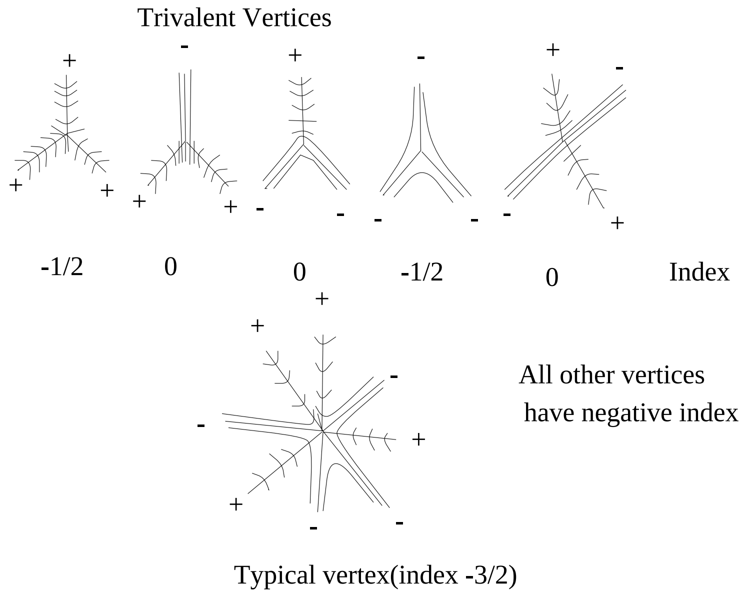

Figure 6: Indices of vertices

which implies c_F has non-positive index. In order to have the sum of the indices equal to two (the Euler number of P), the indexes of both c_{F_1}, c_{F_2} must be one and therefore they are of constant sign. Also all vertices and centers c_F , $F \neq F_1, F_2$ must have index zero. Therefore all vertices have type $(++-), (--+)$ and each F has exactly four sign changes.

If $|\Sigma_0| \neq 0$ then we cannot find a vector field extending over Σ_0. We avoid this problem by collapsing these zero-edges. First consider $\overline{\Sigma}_0 \subseteq \Sigma_0$ to be the collection of edges belonging to loops in Σ_0. This decomposes P into regions $\{R_i\}$ which have disjoint interiors and are n_i-connected regions whose boundaries consist of zero-edges. Some regions correspond to a zero-face which we call zero-regions. To find a cell division with no zero-edges we must collapse Σ_0. Each region R_i becomes a sphere $\tilde{S}_i$ in a bouquet of spheres $\tilde{P}$. The zero-regions become spheres with the trivial cell division of a point and a disc and every other face F collapses to a cell $\tilde{F}$ of some $\tilde{S}_i$.

We collapse Σ_0 by first collapsing $\overline{\Sigma}_0$. Each R_i becomes a sphere S_i. If R_i is not a zero-face then we have each component of ∂R_i, β_{ij} collapses to a vertex v_{ij} on a sphere S_i. As R_i is not a face there is an edge path γ_{ij} splitting R_i into two regions R^1_{ij}, R^2_{ij} with endpoints on β_{ij}. If there are no other paths from β_{ij} to the interior then either R_i is the union of two faces or there are two faces which intersect in two separate edges (figure 8). Both give contradictions. Therefore collapsing β_{ij} gives vertex v_{ij} with valence at least three. Thus the vertices on S_i after collapsing $\overline{\Sigma}_0$, have valence at least three. If we now collapse $\Sigma_0 - \overline{\Sigma}_0$ to get sphere $\tilde{S}_i$ in $\tilde{P}$ the valences cannot go down as the only zero-edges left on S_i form a collection of disjoint trees.

Therefore if $\tilde{S}_i$ does not correspond to a zero-face it has a decomposition into faces all but two of which have at least four sign changes and all vertices have valence at least three. Thus

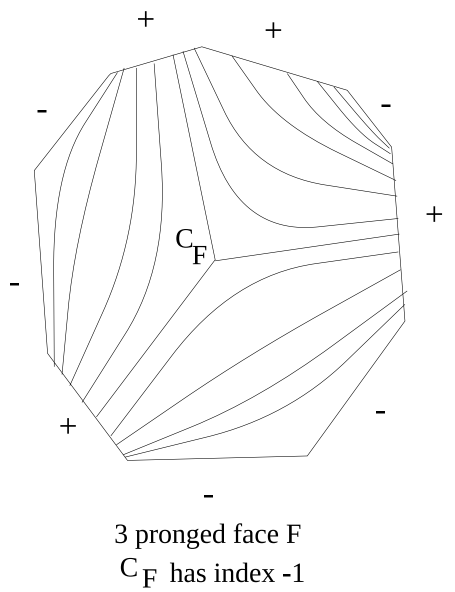

Figure 7: Extend vector field over face

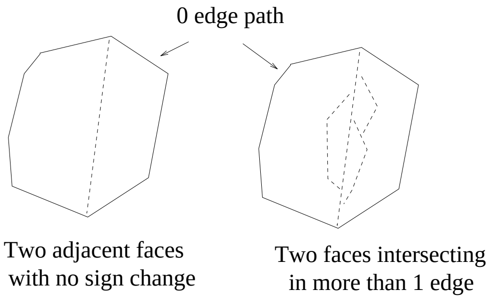

Figure 8: Valence after collapse is at least three

as before we can give $\tilde{S}_i$ a vector field which implies that there can be at most one sphere $\tilde{S}$ which has a non-trivial cell division. On $\tilde{S}$ the vertices all have type $(++-), (--+)$ or $(+-+-)$, the faces corresponding to F_1, F_2 give rise to faces $\tilde{F}_1, \tilde{F}_2$ of $\tilde{S}$ which have constant sign on their boundaries and all other faces have exactly four sign changes. ■

Now we apply the above theorem. If we compare the corresponding faces of $P^t_{\theta_1}$ and $P^t_{\theta_2}$ then apart from the two faces F_1, F_2 which are incident and perpendicular to g , all others have identical face angles. Let Q be the marked polyhedron combinatorially equivalent to both $P^t_{\theta_1}$, $P^t_{\theta_2}$ and marked with $+$, $-$ or 0 depending on whether the length of the edge in $P^t_{\theta_1}$ is bigger, smaller or equal in length to the corresponding edge of $P^t_{\theta_2}$. By the below result in planar hyperbolic geometry (lemma 2), any face $F \neq F_1, F_2$ is either a zero-face or has at least four sign changes. Therefore by the above (theorem 5), both faces F_1, F_2 have no sign changes. If $\theta_2 > \theta_1$ then again using the below lemma, both faces must have constant sign $+$. If the edge g is $+$ then we get two $(+++)$ vertices and if the edge g is 0 then after collapsing we get a $(++++)$ vertex, both giving contradictions.

Therefore g must be marked with $-$ and thus l^t is monotonic in θ.

Theorem 5 gives more information about the pattern on a marked polyhedron P. The region R corresponding to $\tilde{S}$ is obtained by removing the zero-faces of P and has no zero-loops with interior edges. When we collapse the boundary curves of R to get sphere S then any vertex of S has valence three or four. Any zero-edge path of more than one edge will give a vertex in $\tilde{S}$ with valence greater than four and therefore every zero-edge in S is isolated. These isolated zero-edges in S must join two trivalent vertices which have alternating signs as you go around a neighborhood of the zero-edge.

In the case of the markings on Q we obtain figure 9.

The following is the technical lemma needed above. A proof can be found in the referenced paper.

Lemma: 2 *([HR]) Let $E_1, E_2, \ldots, E_n$ be lines in H^2 defining a convex polygonal curve C. Let $A_i = E_i \cap E_{i+1} (i = 1, \ldots, n-1)$ be the vertices of C , let $\nu_i = E_i^{\perp} (i = 1, \ldots, n)$ be the outward unit normals to the Minkowski inner product $\langle \cdot, \cdot \rangle$, and let l_i be the length of $E_i \cap C = A_{i-1}A_i (i = 1, \ldots, n-2)$. Let $E'_1, E'_2, \ldots, E'_n$ be lines in H^2 defining another such curve C' with A'_i, ν'_i, l'_i defined as above. Assume that*

(i) $< A_i =< A'_i$ for $i = 2, \ldots, n-1$ and

(ii) $l'_i \geq l_i$ for $i = 2, \ldots, n-1$.

Then

$$\langle \nu'_1, \nu'_n \rangle \quad \leq \quad \langle \nu_1, \nu_n \rangle,$$

with equality if and only if $A_1 A_2 \ldots A_{n-1}$ is congruent to $A'_1 A'_2 \ldots A'_{n-1}$.

An easy consequence of this lemma is the following. Let F_a, F_b be two hyperbolic n-gons with face angles equal. If we label an n-gon F with $+$, 0 or $-$ by comparing lengths of corresponding sides of F_a, F_b, then either F has all sides marked with 0, in which case F_a, F_b

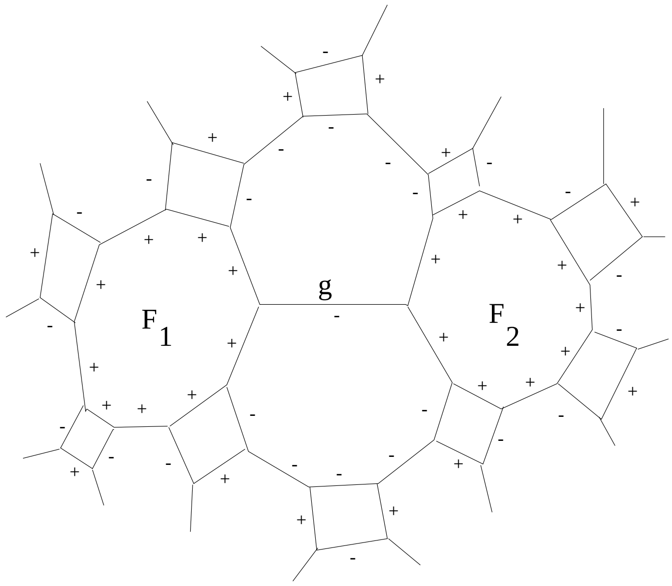

Figure 9: How edge lengths in P^t_θ vary with θ

are congruent, or, F has at least four sign changes as we go around the perimeter. This shows that Q satisfies the hypothesis of the theorem.

To show that F_1, F_2 are marked $+$ we need to relate the Minkowski inner product to face angles. If two edges of a compact hyperbolic polygon intersect in an angle $\theta < \pi$ then the inner product of their outward normals is $\cos(\theta_{ext})$ where θ_{ext} is the exterior angle of θ. Now the convex polygons p_1, p_2 corresponding to F_1 in $P^t_{\theta_1}, P^t_{\theta_2}$ respectively, have the same face angles except at the vertex incident to g, where p_i has the angle θ_i. Thus if the sides of p_1 where all smaller or equal in length to the sides of p_2 then we would have by the above lemma that the inner product of the outward normals of the two sides intersecting at g in p_1 is greater than or equal to the same inner product in p_2, that is $\cos(\pi - \theta_1) \geq \cos(\pi - \theta_2)$. But $\theta_1 < \theta_2$, and therefore the sides of p_1 must be greater than or equal in length to the corresponding sides in p_2. This gives the needed step to show that both F_1, F_2 are of constant sign $+$.

5 Combinatorial Bound on Volume Increase

Here we will show that bending along a geodesic g on a face F which is perpendicular to two edges of F increases volume at most linearly in the number of edges of F.

The idea is that as θ, the exterior dihedral angle at g varies, F is still intrinsically an ideal n-gon. If $l(\theta)$ is long then this forces F to be narrow somewhere across g. This implies that the two faces F_1, F_2 obtained by bending F along g (which meet with exterior dihedral angle θ), have faces F_3, F_4 adjacent to F_1, F_2 respectively which are connected by a short path on the bent face F. As F_3, F_4 cannot intersect, this shows that $l(\theta)$ must decrease as we bend otherwise F_3, F_4 would be forced to intersect.

We will describe the proof for when we have a continuous one-parameter family of polyhedra P_θ interpolating between P and P'. We obtain the result in general by approximating as before by the one-parameter family X^t.

Long Geodesics give Thin Polygons

An ideal hyperbolic quadrilateral has two perpendiculars and it can be shown using elementary hyperbolic geometry that the hyperbolic sine's of the half lengths of these are reciprocal. The following lemma is a generalization of this fact.

Lemma: 3 *If F is an ideal hyperbolic n-gon and g the unique perpendicular to edges e_1, e_2 of F then $\exists\ e_3, e_4$ separated by g with distance d between them satisfying*

$$\sinh(d/2) \cdot \sinh(\frac{(\pi-2)l}{2\pi(n-3)}) \leq 1$$

where l is the length of g.

Proof : Choose $\epsilon > 0$ such that the one-sided ϵ-neighborhood of g has the same area as F.

$$\Rightarrow l\sinh(\epsilon) = (n-2)\pi$$

If we take N_ϵ the two-sided $\epsilon-$neighborhood of g then for each $x \in g$ there is an arc γ_x perpendicular to g centered at x of radius ϵ. We define three subsets of g as follows:

$$g_i = \{x \in g : |\gamma_x \cap \partial F| = i\}$$

Therefore $l = l_0 + l_1 + l_2$ where $l_i =$ length of $g_i, i = 0, 1, 2$.

Each of the sets g_i is a collection of intervals. If $A_1 =$ Area of F outside N_ϵ (cusp region) and $A_2 =$ area of F consisting of points whose perpendicular geodesic to g has length less than ϵ (narrow region) then

$$\begin{aligned}
\text{Area}(F) &= \text{Area}(F \cap N_\epsilon) + A_1 \\
\text{Area}(F \cap N_\epsilon) &= 2l_0\sinh(\epsilon) + l_1\sinh(\epsilon) + A_2 \\
\Rightarrow l\sinh(\epsilon) &= 2l_0\sinh(\epsilon) + l_1\sinh(\epsilon) + A_1 + A_2 \\
\Rightarrow l_2\sinh(\epsilon) &= l_0\sinh(\epsilon) + A
\end{aligned}$$

where $A = A_1 + A_2$ (cusp plus narrow region).

If e is one of the $n-2$ edges of F not intersecting g, drop perpendiculars to g from the endpoints of e on S^1_∞ to the two points x, y on g. This forms a quadrilateral Q_e with area π and base $[x, y]$. Let A_e be the area of the region in Q_e which is either a distance greater than ϵ from g or whose perpendicular to g has length less than ϵ. A_e is the contribution of

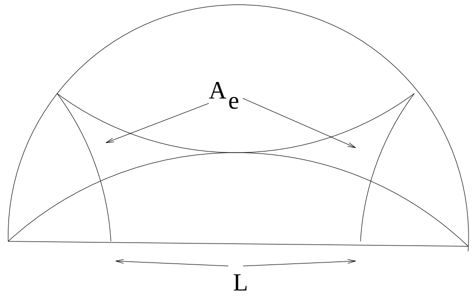

Figure 10: A_e is minimum

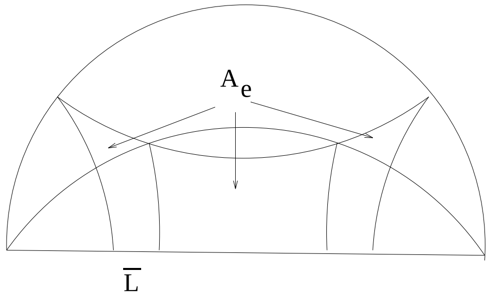

Figure 11: General picture of A_e

Q_e to A. If e doesn't intersect N_ϵ then the A_e is minimum when e is tangent to ∂N_ϵ (figure 10). If L is the length of the base of the quadrilateral tangent to ∂N_ϵ then

$$\begin{aligned} \sinh(L/2) &= \frac{1}{\sinh(\epsilon)} \\ \Rightarrow \quad A_e &\geq \pi - L\sinh(\epsilon) = \pi - \frac{L}{\sinh(L/2)} \\ \Rightarrow \quad A_e &\geq \pi - 2 \end{aligned}$$

If e intersects ∂N_ϵ at points x', y' then project down onto g to get points x_1, y_1 (figure 11). Now if $\overline{L}$ is the length of either $[x, x_1]$ or $[y_1, y]$ (both have the same length) then

$$\begin{aligned} \overline{L} &\leq L/2 \\ \Rightarrow A_e &= \pi - 2\overline{L}\sinh(\epsilon) \\ \Rightarrow A_e &\geq \pi - 2 \\ \Rightarrow A &\geq (n-2)(\pi - 2) \end{aligned}$$

Since we have that $l_2 \sinh(\epsilon) = l_0 \sinh(\epsilon) + A$ then

$$\begin{aligned} l_2 &\geq \frac{A}{\sinh(\epsilon)} \geq \frac{(n-2)(\pi-2)}{\sinh(\epsilon)} \\ \Rightarrow l_2 &\geq \frac{\pi-2}{\pi} \cdot l \end{aligned}$$

The subset g_2 consists of at most $(n-3)$ intervals. Therefore by the above, one of these intervals $[a, b]$ must have length D satisfying

$$D \geq \frac{\pi-2}{\pi(n-3)} \cdot l$$

Consider the quadrilateral Q formed by taking the two perpendicular arcs γ_a, γ_b to g, centered at a, b respectively, of length 2ϵ. Choose the other sides α, β to join the endpoints of γ_a, γ_b. There are two edges e_3, e_4 of F separated by g, which enter and leave Q through the sides γ_a, γ_b. To show that e_3, e_4 are close note that they are closer than α and β are. The two geodesics α, β describe an ideal quadrilateral by extending α, β to the sphere at infinity. This quadrilateral has a diagonal of length D' which is longer than D whose other diagonal has length d'. If d is the distance between e_3 and e_4 then

$$\sinh(D/2) \leq \sinh(D'/2) = \frac{1}{\sinh(d'/2)} \leq \frac{1}{\sinh(d/2)}$$

Therefore we have

$$\sinh(d/2) \cdot \sinh\left(\frac{(\pi-2)l}{2\pi(n-3)}\right) \leq 1 \tag{1}$$

■

This lemma shows that long geodesics must give thin polygons.

Thin Bent Plane Forces Intersection

In the hyperbolic plane consider a geodesic γ represented by a circle C of radius R intersecting the unit circle at right angles. The two geodesics passing through the origin o meeting γ on S^1_∞ form an angle of ϕ at o (visual angle). If x is the Euclidean distance from o to γ and d is the hyperbolic distance then we have by simple planar geometry

$$\begin{aligned} R &= \frac{1-x^2}{2x} \\ \tan(\phi/2) &= R \\ 2\tanh^{-1}(x) &= d \\ \Rightarrow \quad x &= \tanh(d/2) \\ \Rightarrow \quad R &= \frac{1}{\sinh(d)} \\ \Rightarrow \quad \tan(\phi/2) &= \frac{1}{\sinh(d)} \end{aligned}$$

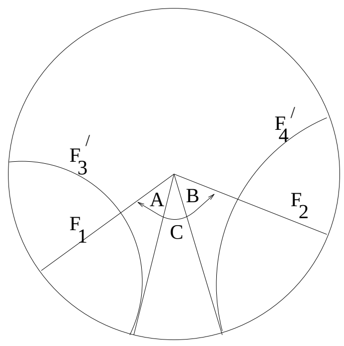

Figure 12: Sideview of the planes

In P_θ we have planes F_1, F_2 intersecting with dihedral angle ϕ and let F_3, F_4 be the two planes intersecting in edges e_3, e_4 respectively of lemma 3. We also have planes F_3', F_4' defined by geodesics α, β which intersect F_1, F_2 at right angles.

If F_3', F_4' intersect then F_3, F_4 must also intersect. Therefore since planes F_3, F_4 cannot intersect then neither can F_3', F_4'. Looking along the geodesic g, F_3', F_4' are given by two geodesics γ_3, γ_4 and F_3', F_4' by two geodesics passing through the origin and perpendicular to γ_3, γ_4 respectively (figure 12).

Both γ_3, γ_4 have the same visual angle ϕ'. If they are not to intersect then

$$\begin{aligned} \phi'/2 &\leq \phi/2 \\ \Rightarrow \tan(\phi'/2) &\leq \tan(\phi/2) \\ \Rightarrow \frac{1}{\sinh(d'/2)} &\leq \tan(\phi/2) \end{aligned}$$

where d' is the length of the short curve on F joining F_3', F_4'. By lemma 3 we have

$$\sinh\left(\frac{(\pi-2)l(\phi)}{2\pi(n-3)}\right) \leq \tan(\phi/2)$$

Therefore

$$l(\phi) \leq \frac{2\pi(n-3)}{(\pi-2)} \cdot \sinh^{-1}(tan(\phi/2)) \tag{2}$$

Therefore we obtain a combinatorial bound on the volume increase by integrating the above equation. This integration is possible as although $\sinh^{-1}(\tan(\phi/2))$ tends to infinity at π it is still integrable. Therefore we have

$$\Delta V \leq \frac{\pi(n-3)}{(\pi-2)} \cdot \int_0^\pi \sinh^{-1}(\tan(\phi/2)d\phi = K \cdot (n-3) \tag{3}$$

where K is some fixed constant. We can derive this bound in general using the one-parameter family of approximations X^t in the same manner as before.

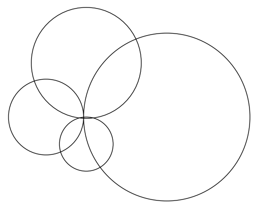

Figure 13:

6 Non-existence of lower bound for volume increase

If the volume increase under Dehn drilling is bounded below by a constant times some positive power of the length of the geodesic drilled then we have that

$$\Delta V \geq c \cdot l^a$$

where c, a are fixed positive constants and l is the length of the geodesic drilled. We have shown in the previous section, that the volume increase by drilling out a geodesic belonging to an ideal n-gon of a basic polyhedron is bounded linearly in n. Therefore to prove that no lower bound exists we need to find a sequence of basic polyhedra each containing a face of bounded valence and the set of lengths of geodesics perpendicular to two sides unbounded. Let $\{P_k\}_{k=0}^{\infty}$ be a sequence of basic polyhedra and N a number such that each P_k has a marked face F_k which is an n_k-gon, $n_k < N$. Also let l_k be the length of the longest geodesic α_k in F_k perpendicular to two edges of F_k. Then the volume increase ΔV_k, by drilling out α_k is bounded above by some constant K depending only on N. If the sequence l_k is not bounded then by the lower bound for volume increase $\Delta V_k \geq c \cdot l_k^a$. Therefore ΔV_k is unbounded giving a contradiction. We now construct the required family using circle packings.

Circle Patterns and Packings

If P is a hyperbolic polyhedron then in the upper half-space model each face F of P lies on a hemisphere H_F which intersects S^2_∞ in a circle C_F. Furthermore the angle at which two circles C_{F_1}, C_{F_2} intersect equals the dihedral angle between the faces F_1, F_2. Thus looking for a polyhedron of prescribed dihedral angles and combinatorics is equivalent to finding a circle pattern with prescribed angles of intersection and combinatorial pattern. We will use the terminology circle packing to refer to a collection of circles with intersections being tangential.

In the case of constructing a basic polyhedron we require that the circles meet in groups of four with opposite circles tangent and neighboring circles meeting at right-angles (figure 13).

We start our circle pattern with a circle C_0 of radius one centered at the origin O. We further take circles C_1, C_1' of radius r such that both are perpendicular to C_0 and have centers on the positive, negative real axis respectively. Thus there are unique circles C_2, C_2' which are perpendicular to C_0, with centers on the positive, negative imaginary axis respectively and are tangent to both C_1, C_1'. This configuration of five circles depends only on our choice of r and therefore we have a one-parameter family $\mathcal{C}_r$ of circle patterns. By simple planar hyperbolic geometry the radius of both C_2, C_2' is $1/r$. This collection of five circles describes an ideal quadrilateral on the hyperbolic plane H_0 associated with circle C_0. Furthermore r is a coordinate system for the shape of the quadrilateral with it becoming narrow in the horizontal direction as r tends to infinity and becoming narrow in the vertical direction r tends to 0.

To extend $\mathcal{C}_r$ to the circle pattern of a basic polyhedron we perform a moebius transformation M on $\mathcal{C}_r$. M consists of a rotation of $\pi/2$ about the origin followed by an expansion by a real number k_r. The factor k_r is chosen so that $M(C_1)$ is tangent to the original C_1. For r large these four tangent points are the only intersections between $\mathcal{C}_r$ and $M(\mathcal{C}_r)$ and so we extend circle pattern $\mathcal{C}_r$ by adding $M(\mathcal{C}_r)$ to form circle pattern $\overline{\mathcal{C}_r} = \mathcal{C}_r \cup M(\mathcal{C}_r)$ (figure 14).

Now take the collection of hemispheres $\mathcal{H}_r$ associated with $\overline{\mathcal{C}_r}$ and let N_r be the region between the two hemispheres $H_0, M(H_0)$ cut out by $\mathcal{H}_r$. N_r is an infinite volume hyperbolic orbifold with four infinite volume funnels. Each of these funnels meet S^2_∞ in a quadrilateral bounded by circles which are tangent at points of intersection. It can be easily shown that all four of these quadrilaterals are equivalent up to moebius transformation to a quadrilateral Q_r.

To form a basic polyhedron from N_r we must truncate the funnels of N_r. We do this by first finding a packing of Q_r by a finite number of circles such that the interstices are all triangular. Then having packed each of the four copies of Q_r, we are left with a finite number of triangular interstices which we truncate by adding the unique circle passing through the three vertices of each interstice. These circles are also perpendicular to the three sides of the interstices and therefore we have truncated the four infinite volume funnels of N_r to get a basic polyhedron P_r.

Unfortunately it is not always possible to find a finite packing of an arbitrary circular quadrilateral but using work of Robert Brooks([Br]) we will show that there is a sequence $\{r_n\}$ such that Q_{r_n} is packable and r_n monotonically tends to infinity.

We define the continued fraction $c(Q)$ of a quadrilateral Q by using the greedy algorithm. Assume that the four sides have been labelled left, right, top, bottom. We first pack Q from left to right by starting at the left side of Q and packing in the largest circle tangent to the left, top and bottom of Q. Now Q contains a smaller quadrilateral Q_1 which we again pack with the largest circle tangent to the left, top and bottom sides. After n_1 circles have been added the remaining quadrilateral Q_{n_1} is too narrow to fit any more circles which are tangent to the left, top and bottom so we now pack Q_{n_1} from top to bottom with circles tangent to the top, left and right sides. Again after n_2 circles have been packed into Q_{n_1} we are left with a quadrilateral Q_{n_1,n_2} too wide to pack with a circle tangent to the top, left and right. The packing algorithm is defined recursively by now applying the algorithm to the quadrilateral Q_{n_1,n_2} (figure 15). To each quadrilateral we get a sequence of integers $\{n_i\}$ and we define the function c on the space of quadrilaterals to be the real number with these

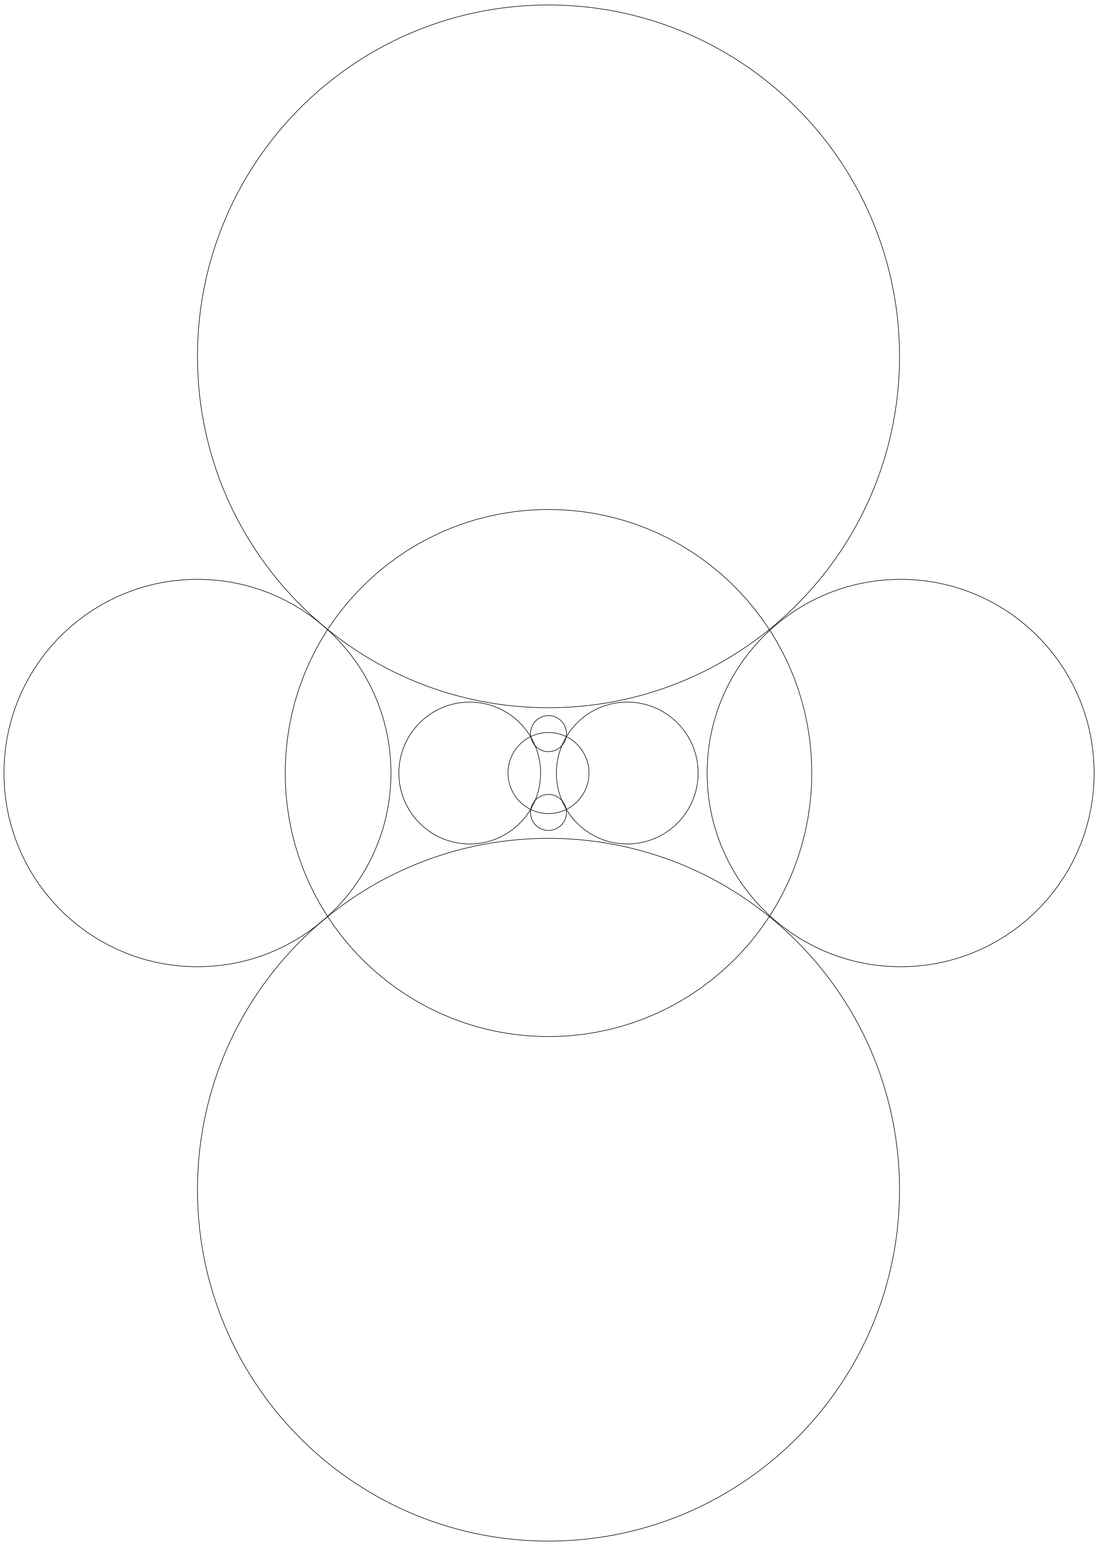

Figure 14: Circle pattern

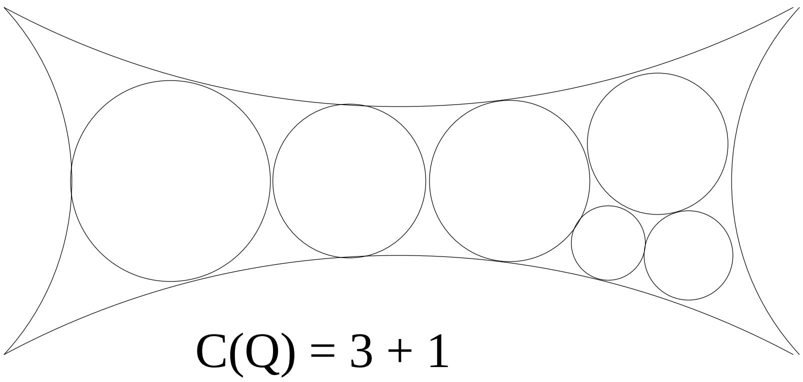

$$C(Q) = 3 + \cfrac{1}{1 + \cfrac{1}{2 + \cfrac{1}{\cdots\cdot}}}$$

Figure 15: Greedy algorithm

as its continued fraction expansion, that is

$$c(Q) = n_1 + \cfrac{1}{n_2 + \cfrac{1}{n_3 + \cfrac{1}{n_4 + \cdots}}}$$

Note that the greedy algorithm gives a finite packing of Q if and only if $c(Q)$ is rational. Now let $\mathcal{S}$ be the space of all configurations of four circles up to moebius transformation which have disjoint interiors and such that the complement of the interiors is two circular quadrilaterals. Then Robert Brooks showed ([Br]) that the map $\overline{c} : \mathcal{S} \to \Re^2$ defined by mapping a configuration of circles to the pair of continued fractions of its two circular quadrilateral regions, is not only continuous but a homeomorphism. This implies that c is continuous and therefore $c(Q_r)$ is a continuous function of r. It can be shown that as r tends to infinity the modulus of Q_r tends to infinity. Therefore $c(Q_r)$ tends to infinity and by continuity there exists a sequence $\{r_n\}$ such that $c(Q_{r_n}) = n$. Therefore Q_{r_n} can be packed with n circles from left to right.

Now basic polyhedron P_n is defined by truncating the funnels of N_{r_n} as described above. Figure 16 is a combinatorial picture of P_n. The polyhedra P_n have the property of each containing an ideal quadrilateral face F_n whose modulus tends to infinity as a function in n. Thus F_n contains a geodesic α_n, perpendicular to two sides of F_n whose length monotonically increases to infinity as a function of n. Therefore this sequence of polyhedra is a counter-example to the existence of a lower bound to volume increase under Dehn drilling. We get a similar counter-example for the manifold case by looking at the four-fold covers of these polyhedra when viewed as orbifolds.

Conclusion

The two bounds we derive are very different. One is a bound on the volume increase by the length of the geodesic drilled and the other is a purely combinatorial bound. I believe that these techniques can be applied to a wider class than I have done and further hope to prove the conjecture at the outset by decomposing the manifold into polyhedra and extending the results in this paper.

References

[A] Alexandroff *Convex Polyhedra*

[B1] M. Bridgeman *Volume increase under Dehn drilling operations* Phd. thesis Princeton June 1994

[B2] M. Bridgeman *The Structure and Enumeration of Link Projections* To Appear Trans. A.M.S. (1995)

[Br] R. Brooks *On the Deformation Theory of Classical Schottky Groups* Duke Mathematical Journal December 1985

[Th] W. P. Thurston *The Geometry and Topology of 3-manifolds*

[HR] Craig Hodgson , Igor Rivin *A characterization of Compact Convex Polyhedra in Hyperbolic 3-Space*

[Ho] Craig Hodgson *Degeneration & Regeneration of Geometric structures on Three-Manifolds* Phd. thesis Princeton June 1986

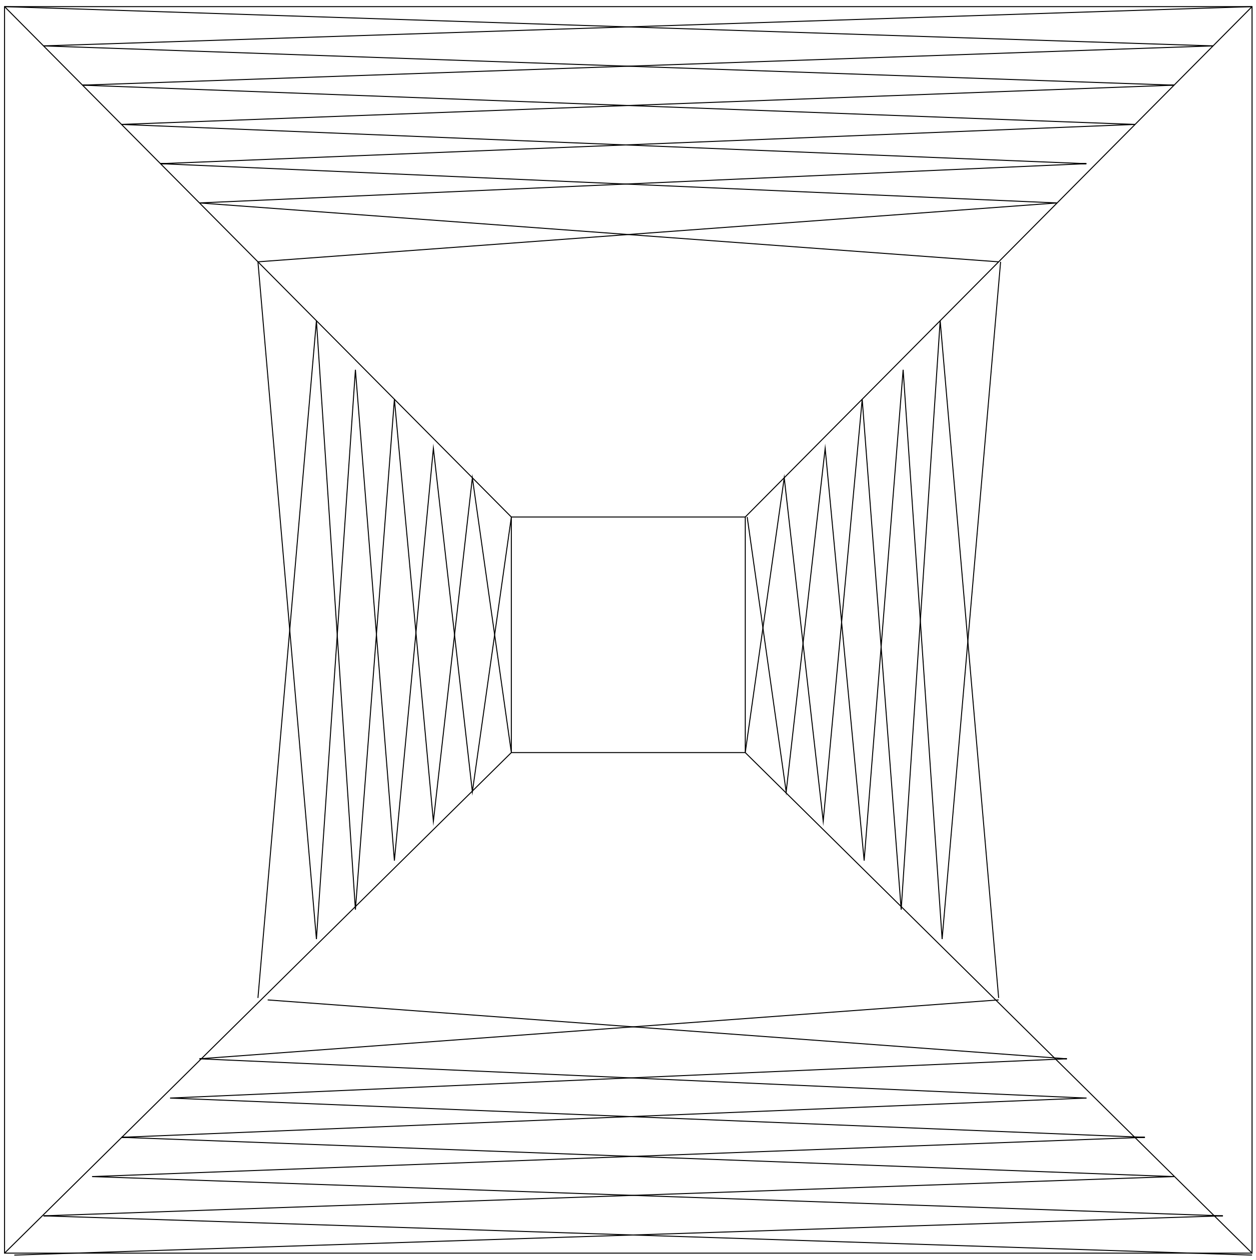

Figure 16: Polyhedron with thin ideal quadrilateral face