COMMUNICATIONS IN
ANALYSIS AND GEOMETRY
Volume 4, Number 2, 181-206, 1996

Hyperbolic 3-Manifolds with Two Generators

COLIN C. ADAMS

We show that if there are two parabolic elements that generate a non-elementary Kleinian group that is not free, then there is a universal upper bound of two on the "length" of each of those parabolics, length being measured in a canonical choice of cusp boundaries. Moreover, there is a universal upper bound of ln(4) on the "distance" between those parabolics, where the distance between them is the distance between a pair of horoballs corresponding to the canonical cusps. We prove a variety of results with these, the most interesting of which is: An orientable finite volume hyperbolic 3-manifold that has fundamental group generated by two parabolic isometries must be a 2-bridge knot or link complement.

1. Introduction.

Since the advent of the study of Kleinian groups and the hyperbolic 3-manifolds and 3-orbifolds that they generate, mathematicians have realized that certain special relations and restrictions apply to these groups. For example, the Shimuzu-Leutbacher Inequality or its generalization Jorgensen's Inequality, puts severe restrictions on the pairs of isometries of hyperbolic 3-space that can generate a discrete non-elementary group.

In this paper, we will determine restrictions on the pairs of isometries that can generate a non-elementary Kleinian group that is not free. In particular, we will begin by proving Theorem 3.2, which states that if there are two parabolic elements that generate a non-elementary non-free Kleinian group, then there is a universal bound of two on the "length" of each of those parabolics, length being measured in a canonical choice of cusp boundaries.

[1]This research partially supported by NSF Grants DMS-9000937 and DMS-9302843 *1991 AMS Subject Classification:* Primary 57M50, Secondary 57M25. *Key words and phrases:* Hyperbolic 3-manifold, hyperbolic 3-orbifold, Kleinian group, hypebolic knot, two generator group.

Moreover, there is a universal bound of $\ln(4)$ on the "distance" between those parabolics, where the distance between them is the distance between a pair of horoballs corresponding to the canonical cusps, each parabolic preserving one horoball in the pair.

It will follow that a Kleinian group can have at most a finite set of pairs of parabolic generators, up to conjugacy of the pair. In the case that the Kleinian group is finite volume, we give a bound on the number of such pairs as a linear function of the total volume in a canonical set of cusps.

We then limit ourselves to the special case of finite volume hyperbolic 3-manifolds that are generated by two parabolic isometries. Perhaps the most interesting result is a proof that if M is a finite volume hyperbolic 3-manifold that is generate by two parabolics, then it must be the complement of a two-bridge link in S^3. This proof does depend on Thurston's Orbifold Theorem, which, at the time of this writing, has not appeared in print.

As an example of the utility of Theorem 3.2, we make a complete determination of all of the pairs of parabolic generators up to conjugacy in the figure-eight knot complement.

Our proof of Theorem 3.2 is geometric. After having discovered this proof, we came across a paper by J.L. Brenner [8] obtaining restrictions on pairs of parabolic generators that generate non-elementary non-free Kleinian groups using non-geometric techniques. At the end of Section 3, we show that Theorem 3.2 and Brenner's result imply one another.

In the last section of the paper, we examine constraints on pairs of generators of non-elementary Kleinian groups that are not free, when the generators are both hyperbolic isometries. For example, we prove that in the case of two hyperbolic generators, it is either the case that one of the two hyperbolic generators has length less than $2\ln(\sqrt{2}+1)$ or it must be that the distance between their axes is less than $2(\ln(\sqrt{2}+1)$. Similar arguments can be applied when there are two generators, one of which is parabolic and one of which is hyperbolic, however, the results are not as strong.

The rest of this section is devoted to background. Every Kleinian group corresponds to a hyperbolic 3-manifold or 3-orbifold that is obtained by taking the quotient of hyperbolic 3-space by the action of the Kleinian group. If there are parabolic isometries in the group, the 3-manifold or 3-orbifold will have cusps. Each cusp is covered by a set of horoballs in hyperbolic space. By shrinking the cusps back if necessary, we can assume that all of the horoballs covering all of the cusps are disjoint. Given a particular choice of a set of cusps for a manifold or orbifold, we define the length of a given parabolic isometry in the fundamental group to be the length of the shortest loop in the cusp boundary that lifts to this parabolic isometry, where the length of the loop is measured with respect to the hyperbolic metric. Since we are only interested in at most one or two parabolic generators, we will be concerned with only one or two of the cusps.

In the case of a manifold or orbifold in which the two parabolic generators occur on the same cusp, we will take that cusp to be a *maximal cusp*, meaning that the cusp has been expanded until it first touches itself on the boundary. Such a cusp lifts to a set of horoballs in H^3, some of which are tangent but all of which have disjoint interiors. We denote the volume of the resulting cusp by $v_c(M)$. The length of the parabolic isometry α, denoted $|\alpha|$, is then defined to be its length with respect to this cusp.

In the case that the two parabolic generators correspond to different cusps, we shrink one of the two cusps back in size until both parabolic generators have the same length with respect to the cusps. We then enlarge the two cusps at rates that keep the two lengths of the parabolic generators equal, until the two cusps either touch one another or one of them touches itself. The length of a parabolic generator then refers to its length with respect to these cusps. In either the case when we are concerned with only a single cusp or when we are concerned with two cusps, we call the particular set of cusps that we have constructed a *canonical set of cusps*.

A finite volume noncompact hyperbolic 3-manifold M has a compact core M' with toroidal boundary components, which is obtained by discarding the interiors of a disjoint set of cusps. A loop in M represents a parabolic isometry in the fundamental group if and only if the loop is freely homotopic into the boundary of M'.

We say that a loop in the complement of a knot is a *meridian* if it is freely homotopic to a curve on the boundary of a regular neighborhood of the knot that bounds a meridional disk in the regular neighborhood. We say that two parabolic elements in a hyperbolic 3-manifold are *parallel* if loops corresponding to them are freely homotopic to curves on the boundary of a single cusp, such that those curves do not intersect one another. Note that two curves that are parallel are either freely homotopic to one another or a power of one is freely homotopic to a power of the other. Given two disjoint horospheres, we define a *connector* to be a geodesic path from one to the other such that it is perpendicular to both. Given two parabolic generators that do not share a fixed point, each one preserves a horoball that covers one of the canonical cusps. Define the distance between the two parabolic generators to be the length of the connector between these two horoballs. When counting distinct pairs of parabolic generators, we will not distinguish between the pair $\{\alpha, \beta\}$ and the pairs $\{\alpha, \beta^{-1}\}$, $\{\alpha^{-1}, \beta\}$ and $\{\alpha^{-1}, \beta^{-1}\}$.

Given a hyperbolic generator α, it corresponds to translation and rotation about a geodesic. We define its length $|\alpha|$ to be the length of the translation along the geodesic.

In all of what follows, we will be working in the upper-half-space model of hyperbolic 3-space. Hence it makes sense to talk about the highest point on a horosphere or geodesic, where highest refers to the point with largest z-coordinate. We call such a highest point the *peak* of the horosphere or geodesic. All size comparisons between pairs of horospheres, or pairs of geodesics are made in reference to which has the higher peak. Given a word $w = a^{n_1}b^{n_2}a^{n_3}...b^{n_m}$ in a free group generated by two elements a and b, where each n_i is a nonzero integer, we define a *syllable* of the word to be a maximal subword of the form a^{n_i} or b^{n_i}.

We will use the same symbol to represent an isometry of hyperbolic space corresponding to an element of the fundamental group and to represent a loop in the equivalence class of loops that corresponds to that element in the fundamental group, when it is not confusing to do so.

Thanks to Alan Reid for several helpful conversations, particularly with regard to Corollary 4.6.

2. Geometric Lemmas.

In this section, we prove a sequence of geometric lemmas and corollaries that will be of use to us in the rest of the paper. All of these lemmas refer to situations occurring in the upper-half-space model of H^3.

Lemma 2.1. *If H_1 and H_2 are two horospheres bounding disjoint horoballs and connected by a connector of length x, and if y_i is the distance on H_i from its peak to the connector for $i = 1, 2$, then $y_1 = 1/(y_2e^x)$.*

Proof. Let H_3 be a horoball centered at ∞ that is tangent to H_2, as in Figure 1a). Then there exists an isometry of hyperbolic space taking H_2 to a horoball H_2' centered at ∞ with height 1 above the $x-y$ plane and taking H_1 to a horoball H_1' centered at the origin. The isometry will take H_3 to a horoball H_3' centered a distance y_2 from the origin, as in Figure 1b). The fact that H_1 and H_2 are a distance x apart forces H_1 to be sent to a horoball of Euclidean diameter e^{-x}.

We will rotate 180° about a geodesic of radius y_2 that is centered above H_1', with one endpoint at the center of H_3'. Then H_1' is sent to a horoball H_1'' centered at ∞ of height y_2^2/e^{-x} as in Figure 1c). The horoball H_2' is sent to a horoball centered at the origin of diameter y_2^2. The horoball H_3' is fixed under the rotation. Then the horizontal line segment in the boundary of H_1'' of hyperbolic length y_1 has a Euclidean length of y_2 and a Euclidean height of y_2^2/e^{-x}, giving it a hyperbolic length of $1/(y_2e^x)$. □

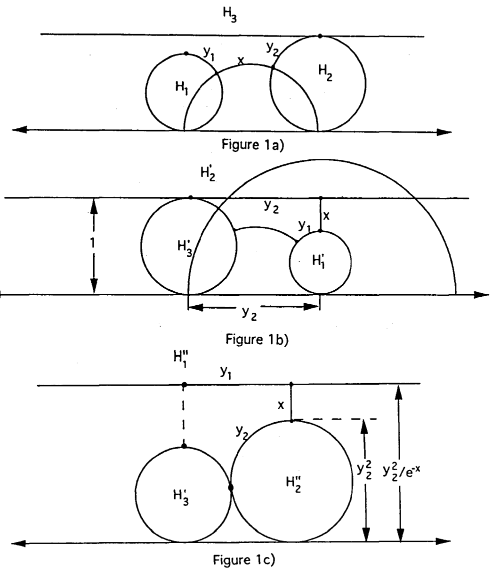

Figure 1.

Corollary 2.2. *Let H be a horosphere connected to a second horosphere at least as large as itself by a connector of length x. Then the distance on H from the connector to its peak is less than or equal to $e^{-x/2}$.*

Proof. When the two horospheres are the same size, the distances from the connector to the peak on each of them are equal. By Lemma 1.1, that distance y then satisfies $y^2 = e^{-x}$, so $y = e^{-x/2}$. As the second horosphere grows in size relative to H, y will shrink. □

Lemma 2.3. *If g_1 and g_2 are two geodesics in the upper-half-space model that intersect perpendicularly, and if y_1 and y_2 are the respective distances on each from the point of intersection to their peaks, then $y_1 \leq \ln(\coth(y_2/2))$.*

Proof. Suppose first that g_1 and g_2 both lie in the same vertical plane. As in Figure 2, y_1 and y_2 each depend on ϕ_1 and ϕ_2, where $\phi_1 + \phi_2 = \pi/2$. Since

$$y_1 = -\ln(\tan(\phi_1/2)),$$

we know

$$\tan((\pi/2 - \phi_2)/2) = e^{-y_1}.$$

Applying trigonometric identities, this becomes

$$e^{-y_1} = \frac{1 - \tan(\phi_2/2)}{1 + \tan(\phi_2/2)}.$$

Since

$$y_2 = \text{-}\ln(\tan(\phi_2/2)),$$

we can substitute e^{-y_2} for $\tan(\phi_2/2)$, obtaining

$$e^{y_1} = \frac{1 + e^{-y_2}}{1 - e^{-y_2}}.$$

Solving for y_1, we obtain $y_1 = \ln(\coth(y_2/2)$.

If we now rotate g_1 through space while fixing g_2 and the point of intersection, so that the two geodesics no longer lie in a common vertical plane, y_1 will decrease, thereby showing that $y_1 \leq \ln(\coth(y_2/2)$. □

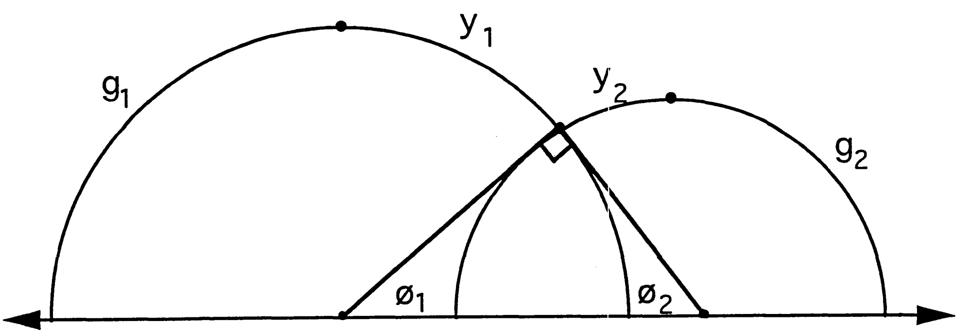

Figure 2.

Lemma 2.4. *Let H and g be a horosphere and a perpendicular geodesic. Then if x and y are the distances on each respectively, from their point of intersection to their peak, then $y = \ln(1/x)$.*

Proof. Let g' be a vertical geodesic that passes through the peak of g and let g'' be a vertical geodesic passing through the center and peak of H. We will take an isometry that sends the center of H to ∞ and ∞ to 0, keeping the designations of the horosphere and geodesics after the isometry acts. We will also normalize so that the Euclidean height of the resultant H above the

$x - y$ plane is 1. The horizontal length along H from its intersection with g'' to its intersection with g' is x. Hence the Euclidean radius of g' is equal to x. Therefore y is the hyperbolic distance obtained as the vertical distance between a point at height x and a point at height 1, which is $\ln(1/x)$. □

Note that if the peak of the geodesic g occurs inside the horosphere H, we will consider the distance along g from its peak to H to be negative.

Lemma 2.5. *Let H_1 and H_2 be two horoballs with disjoint interiors connected to a third horoball H_3 by a pair of connectors of length x. Then one of the connectors must intersect H_1 or H_2 at a point that is a distance of at most 2 from the peak of that horosphere.*

Proof. For convenience, expand H_3 until it becomes tangent with each of H_1 and H_2, which occurs at the same instant. Now we are interested in points of tangency rather than points of intersection with connectors. Suppose first of all that H_1 and H_2 are of the same Euclidean diameter, and that they are tangent to one another. Then the point of tangency with H_3 occurs on the first two balls at points that are within 2 units of their peaks, the points being exactly two units from their peaks when the point where the third ball touches the $x - y$ plane is in line with points where H_1 and H_2 touch the $x - y$ plane. (Note that if three horoballs are pairwise tangent, the distance on any one of them between the points of tangency is exactly one).

Suppose now that we allow one of H_1 or H_2 to expand. However, in order to prevent H_1 and H_2 from overlapping, we have to shrink the second ball at the same time. This second ball will still have its point of tangency with H_3 a distance of at most 2 from its peak. Finally, the case where H_1 and H_2 are not tangent is realized by taking the case where they are tangent and then shrinking either H_1 or H_2. The smaller ball will still have its point of tangency with H_3 a distance no more than 2 from its peak. □

3. A Pair of Parabolic Generators.

Lemma 3.1. *If α and β are two parabolic isometries that generate a non-elementary Kleinian group, then relative to the canonical set of cusps, their lengths are both equal to each other and their lengths are both at least 1.*

Proof. In the case that both parabolic generators correspond to the same cusp, we know that the cusp has been expanded until it touches itself. Therefore, a horoball H_1 centered at ∞ that covers the cusp must be tangent to another horoball H_2. The horoball H_2 must be moved off itself by any parabolic isometry fixing ∞, and therefore that isometry must correspond to a path of length at least one in the cusp boundary. (See [2] for more details.)

Any two such parabolic elements are conjugate in the group of all hyperbolic isometries. In particular, given two parabolic isometries that do not share the same fixed point on S^2_∞ (as our generators cannot, since the group is non-elementary), there exists a geodesic g' perpendicular to the geodesic g between their fixed points, such that the isometry r obtained by rotating 180° about the geodesic g' conjugates the one parabolic isometry to the other. Since the isometry r switches the two generators under conjugation, it conjugates the entire group back to itself. Let H_1 and H_2 be the two horoballs covering the canonical cusp at the endpoints of g. If we can show that r sends H_1 to H_2, then we will have shown that the two parabolic isometries have the same length.

However, since H_1 and H_2 both cover the same cusp, there exists an element j in the Kleinian group G that identifies H_1 to H_2. Then rjr^{-1} sends the center of H_2 to the center of H_1. Since rjr^{-1} is an element of G, it must send H_2 to H_1. If r does not send H_2 to H_1, then either $r(H_2)$ is contained in H_1 or $r(H_2)$ contains H_1. In either case $rjr^{-1}(H_2) \neq H_1$, a contradiction.

Suppose now that the two parabolic generators correspond to distinct cusps. By our choice of canonical cusps, we are assured that the two generators have the same length. If the two canonical cusps touch one another, then when either one is lifted to a horoball centered at ∞, it has a tangent horoball which must be moved off itself by any parabolic isometry fixing ∞. Hence, each parabolic isometry has length at least 1.

However, it could be the case that one of the canonical cusps touched itself before it touched the other cusp. By the same argument that we used in the case of a single cusp, any parabolic isometry corresponding to this cusp will have a length of at least 1. As in the case of a single cusp, there exists an isometry r that conjugates the one parabolic generator to the other, and therefore conjugates the entire group to itself. Since the two parabolics have the same lengths with respect to the two cusps, the isometry r must send the horoballs covering the first canonical cusp to the horoballs covering the second canonical cusp. In particular, if the first cusp touches itself, so must the second cusp. Hence the parabolic isometries corresponding to the second cusp must also be of length at least 1. □

Note that a non-elementary hyperbolic 3-manifold or 3-orbifold that is generated by two parabolic isometries always has a $Z_2 + Z_2$ action on it, one Z_2 action coming from the isometry r described in the preceding proof, and a second Z_2 action coming from 180° rotation about the geodesic g connecting the two fixed points of the parabolic isometries, sending α to α^{-1} and β to β^{-1} via conjugation. (See [16], p. 5.17 for more details.)

The geodesic g must project to a non-singular curve in the manifold, in order that it be an involution axis. (If it were singular, the lift of the involution to H^3 would have to simultaneously rotate about two or more axes that crossed one another, which is not possible.)

Theorem 3.2. *Suppose that α and β are a pair of parabolic isometries that generate a non-elementary Kleinian group G that is not free. Then, with respect to the canonical choice of cusps, $1 \leq |\alpha| = |\beta| < 2$. Moreover, the distance x between α and β must be less than* $\ln(4/|\alpha|^2)$. *In particular,* $x < \ln(4)$.

Proof. Given the two parabolic elements α and β, each one must correspond to a particular cusp, possibly the same cusp. By Lemma 3.1, we know that the two parabolic elements have the same length with respect to the canonical cusps, and that length is at least 1.

For simplicity, we choose a basepoint for the fundamental group of the manifold or orbifold on the boundary of one of these cusps. One of the two parabolic generators α can be realized as a loop α' entirely contained in the cusp boundary, and therefore lifting to a path that is entirely contained in a single horosphere. The second parabolic generator β can be realized as a loop β' that begins at the basepoint, travels along a path ω passing though the complement of the interior of the cusps from the basepoint back to a point on the cusp boundary, follows a closed loop in the cusp boundary, call it δ, and then returns to the basepoint along ω^{-1}. Lifting the path ω to a path ω' in H^3, we obtain a path that begins and ends in the horospheres covering the cusp boundary. If ω' begins and ends on the same horosphere, α' and β' lift to two parabolics that share a fixed point on the sphere at ∞. However, this would imply that the group G was elementary.

Hence, we can assume that ω' connects two distinct horopheres. We may homotope ω' to a connector of the two horoballs. Call that connector γ. The homotopy of ω' projects to a homotopy in the manifold, and induces a homotopy of β' onto a loop β'' that consists in following a geodesic path from the basepoint on the boundary of the canonical cusps, through the manifold or orbifold, back to the boundary of the canonical cusps, around a cusp boundary and then back to the basepoint along the geodesic path. Let x be the length of the connector γ.

It could be the case that α' and β'' intersect one another, however that intersection can occur only in the cusp boundary. In particular, this means that α' and δ' are not parallel curves in the cusp boundary. In this case, we will homotope a small segment of δ slightly into the complement of the cusps at each of their points of intersection, so that they are now disjoint.

Now we take the preimage of α' and $\beta'' = \gamma\delta\gamma^{-1}$ in H^3 and look at the graph K that results. First note that K is connected. If this were not the case, we could form a path in H^3 from one component of K to another, such that the path began and ended at lifts of the basepoint in M. Such a path projects to a loop in the manifold which must be homotopic to a loop given as a word in α and β. The homotopy must lift to a homotopy of the path into K, a contradiction to the fact that the path connects two distinct components of K.

Since the Kleinian group G is not a free group, we know that there exists a word w in α, α^{-1}, β and β^{-1} that is trivial. This implies that there must be a nontrivial loop in K corresponding to w. Let

$$w = \alpha^{n_1}\beta^{n_2}\alpha^{n_3}...\beta^{n_m}.$$

Given a choice of a lift of the basepoint, each syllable of w has a horosphere covering a cusp that is associated to it in H^3, as follows. A lift of α^{n_1} which begins at the lift of the basepoint will be a path in a horosphere containing that basepoint. Lifting β^{n_2} to a path that begins at the end of the lift of α^{n_1}, we will obtain a path that leaves the original horosphere along a connector, travels to a new horosphere, travels along the surface of that horosphere a distance equal to n times the length δ, and then leaves along a connector, depositing us at a third horosphere. The lift of α^{n_3} would then be a path on the surface of this new horosphere. We continue in this way. Hence we associate to each syllable a horosphere, which in the case of α^{n_i} is the horosphere that the lift of α^{n_i} lies in and in the case of β^{n_j} is the intermediate horosphere that the lift of β^{n_j} travels along. Choosing one of the horospheres in the sequence to be centered at ∞ in the upper-half-space model, we know that there exists a smallest horosphere in the sequence of horospheres. That is to say, there must be a horosphere H in the sequence that is connected to two horospheres that are at least as large as it is, and such that one of them is strictly larger than it is. If there is a set of consecutive smallest horospheres, take one on the end of the sequence. We count the horosphere centered at ∞ as the horosphere of largest (in fact, infinite) radius.

Since any two horospheres that are connected by a lift of γ are connected by a geodesic of length x, we know by Corollary 2.2 that each of the two endpoints of the connectors corresponding to the two ends of the syllable must be within a distance $e^{-x/2}$ of the peak of H. In particular, the length of the parabolic isometry corresponding to that syllable as measured in H must be less than $2e^{-x/2}$. This implies that in particular, the length of that parabolic isometry is less that 2. Moreover, since any parabolic isometry has length at least 1, we know that the syllable must be a primitive element given by either $\alpha, \alpha^{-1}, \beta$ or β^{-1}. Since Lemma 3.1 gives us that $|\alpha| = |\beta|$, we have

$$1 \leq |\alpha| = |\beta| < 2.$$

In fact, we actually know more, namely that

$$1 \leq |\alpha| < 2e^{-x/2}.$$

Solving

$$|\alpha| < 2e^{-x/2},$$

for x, we have that $x < \ln(4/|\alpha|^2)$. In particular, since $|\alpha| \geq 1$,

$$x < \ln(4).$$

□

We are particularly interested in the special case that G is the fundamental group of a hyperbolic 3-manifold. Note that a 2-generator subgroup of the fundamental group of an irreducible atoroidal 3-manifold must either be free, free abelian or finite index (c.f. [11], Theorem VI.4.1). Since a hyperbolic 3-manifold is both irreducible and atoroidal, every non-elementary subgroup of the fundamental group of a hyperbolic 3-manifold that is generated by two parabolic elements must either be free or finite index. Thus, Theorem 3.2 is giving us restrictions on the pairs of parabolic elements that can generate finite index subgroups.

After having proved Theorem 3.2, we discovered a paper by J.L. Brenner [8] where he proves that if $(1\mu|01)$ and $(10|\mu 1)$ are matrix representations for two parabolic isometries that generate a Kleinian group, and if $|\mu| \geq 2$, then the group is free. Here, we show how Theorem 3.2 can be derived from that result. Let α and β be two parabolic isometries that generate a non-free non-elementary Kleinian group G. Since G is non- elementary, α and β have distinct fixed points and hence we can normalize so that α and β fix ∞ and 0 respectively and have matrix representations of the forms $(1\mu|01)$ and $(10|\mu 1)$. By Brenner's result $|\mu| < 2$. Expand the cusps until the two horospheres centered at 0 and ∞ become tangent. Note that in order to do this, we may have had to allow a cusp to overlap itself or, if two cusps are present, to have the two cusps overlap each other. When there are two cusps present, we can choose the rate of expansion in each cusp so that the two parabolic generators move points in the two horospheres an equivalent distance.

Note that conjugation by $(01|10)$ sends the second generator to the first. Since the isometry corresponding to $(01|10)$ fixes the point $(0,0,1)$ in the upper-half-space model, the horosphere centered at ∞ must be plane $z = 1$ and the horosphere centered at 0 must be a sphere of Euclidean diameter 1. Since $|\mu| < 2$, each of the two parabolic generators moves points in these two horospheres a distance less than 2. As we shrink back the cusp or cusps to obtain a maximal disjoint set of cusps by removing any overlap, the length that each generator moves points in the horospheres shrinks by a factor of $e^{-x/2}$ where x is the distance between the horospheres. Hence, if the group is not free, each parabolic generator moves points in the horospheres corresponding to a canonical set of cusps a distance less than $2e^{-x/2}$, which is certainly less than 2. We can then derive the fact that

$$x < \ln(4/|\alpha|^2) \leq \ln(4)$$

exactly as in the proof of Theorem 3.2. Similarly, Theorem 3.2 can be used to prove Brenner's result.

4. Corollaries and Related Results.

Note that any two-generator Kleinian group always has infinitely many distinct pairs of generators, even after conjugacy. If α and β are a particular pair of generators for the fundamental group, so are α and $\alpha^n\beta$, for instance. Infinitely many of these pairs cannot be conjugate, as $\alpha^n\beta$ must have increasing "length" as n increases. Hence, the next corollary is a bit surprising.

Corollary 4.1. *A two generator finite volume Kleinian group has at most finitely many distinct pairs of parabolic generators, up to conjugacy of the* *pair. In particular, if M is a one-cusped or two-cusped hyperbolic 3-manifold and if n represents the number of non-conjugate pairs of parabolic generators, and $v_c(M)$ is the volume in one of the canonical cusps, then*

i). For a 1-cusped hyperbolic 3-manifold,

$$\begin{cases} n \leq 3(\frac{32v_c(M)}{\sqrt{3}} - 6) & \text{if } v_c(M) < 2 \\ n \leq \frac{32v_c(M)}{\sqrt{3}} - 6 & \text{if } v_c(M) \geq 2 \end{cases}$$

ii). For a 2-cusped hyperbolic 3-manifold,

$$\begin{cases} n \leq 15(\frac{32v_c(M)}{\sqrt{3}} - \frac{5}{2}) & \text{if } v_c(M) < 2 \\ n \leq \frac{32v_c(M)}{\sqrt{3}} - \frac{5}{2} & \text{if } v_c(M) \geq 2. \end{cases}$$

Proof. A finite volume Kleinian group has at most a finite number of cusps, each of which is generated by a $Z + Z$ subgroup. For each cusp individually, and then for each pair of cusps, we take a canonical set of cusps, and see that there are finitely many pairs of parabolic generators in each case.

In particular, given a specific choice of a single cusp or a pair of cusps and the corresponding canonical cusps, the shortest nontrivial curve in the canonical cusps' boundaries has length at least 1. The set of nontrivial curves of length less than 2 in any one cusp boundary is finite. Hence the total number of pairs of nontrivial curves such that both have length less than 2 is finite. We now look at connectors between the horospheres covering cusps. Each pair of parabolic generators corresponds to a pair of curves in the cusp boundary as above and a connecting geodesic between cusps that lifts to a connector of length less than $\ln(4)$. Lifting one of the cusps to a horoball centered at ∞ of height 1 above the boundary plane in the upper-half-space model of H^3, each of these connecting geodesics lifts to a connector that connects the horoball centered at ∞ with a horoball of diameter at least $1/4$. Since each distinct horoball of diameter at least $1/4$ requires a certain amount of area in the fundamental domain for the cusp subgroup fixing ∞ and since the area of the fundamental domain of the cusp is finite, there can be at most finitely many distinct connecting geodesics. Hence there are only finitely many ways to put together pairs of nontrivial curves in the cusp boundaries with connecting geodesics between them, yielding only finitely many pairs of parabolic generators (α, β), up to conjugation of the pair.

We will now obtain an explicit bound on the number of such pairs in the case that M is a hyperbolic 3-manifold with one cusp. Since the length of any curve in the cusp is at least 1, there are at most three possibilities for each of α and β such that $1 \leq |\alpha| = |\beta| < 2$. Either α corresponds to the shortest nontrivial curve in the cusp, or to the shortest nontrivial curve that is linearly independent from the shortest curve. This follows from the fact that if we have normalized the lift of the cusp boundary so that it occurs as a plane of height 1, a fundamental domain to the cusp subgroup must be a parallelogram that is large enough to contain two disks of diameter 1, those disks being the projections of two full-sized horoballs. (See [2]). This gives at most three possibilities for α, and three possibilities only occurs for a hexagonal cusp. Once α has been determined, there are at most two choices for δ (δ defined as in the proof of Theorem 3.2), the option of two choices occurring only when there exists a linearly independent translation in the cusp subgroup of the same length as the shortest translation. Since we have not specified which element in the pair is which, we have at most a total of three distinct pairs of possibilities for α and δ. Note that it is always the case that $v_c(M)/ \geq \sqrt{3}/2$ (c.f.[2]). If $v_c(M) \geq 2$, then there is at most one parabolic translation (ignoring inverses) with length less than 2 and therefore there is only one possibility for each of α and δ.

Now we would like to determine a bound on the number of distinct connecting geodesics for these pairs of elements. Given a cusp of volume $v_c(M)$, then after normalizing so that a horosphere covering the cusp boundary is a horizontal plane at height 1, the cusp subgroup fixing ∞ has a fundamental domain that is a parallelogram of area $2v_c(M)$. Within that parallelogram each ball of diameter 1/4 projects to a disk that requires an area of $\sqrt{3}/32$, when attempting to pack them in so that they respect the cusp subgroup. Hence, we can fit at most $2v_c(M)32/\sqrt{3}$ disks of diameter at least 1/4 into the fundamental domain. By choosing an orientation on each connecting geodesic, we see that each such lifts to two vertical geodesics modulo the action of the subgroup fixing ∞, one oriented up and one oriented down. Hence the horoballs corresponding to a given connecting geodesic come in pairs, and there are only $32v_c(M)/\sqrt{3}$ possible distinct vertical connectors of length less than $\ln(4)$. But because there must always be a pair of full-sized balls in the cusp diagram, and since a ball of diameter 1/4 must have center a distance 1/2 from the center of a full-sized ball, the disk of diameter 1 that is the projection of each full-sized ball acts as if it has room for 7 disks of diameter 1/4. Hence, if there are n connecting geodesics, all of length less than $\ln(4)$, there is enough room for the equivalent of $14 + 2(n-1)$ balls, all of diameter at least 1/4. Each such ball contributes an area of $\sqrt{3}/32$ to the area of the cusp. Since $2v_c(M)$ is the area of the cusp, we have

$$2v_c(M) \geq \sqrt{3}/32(14 + 2(n-1)).$$

Thus, the number of connectors of length less than $\ln(4)$ is bounded above by $32v_c(M)/\sqrt{3} - 6$. Therefore, when the cusp volume is less than 2, the total number of pairs of parabolic generators up to conjugacy is bounded above by $3(32v_c(M)/\sqrt{3}-6)$. When the cusp volume is at least 2, the total number of pairs of parabolic generators is bounded above by $32v_c(M)/\sqrt{3} - 6$.

If M is a hyperbolic 3-manifold with two cusps, let C_1 and C_2 be the canonical set of cusps. Let α correspond to C_1 and β correspond to C_2. If C_1 touches itself, then by symmetry, C_2 touches itself. As in the case of a single cusp, we then have two possibilities for α and two for β.

On the other hand, if the cusp C_1 touches the cusp C_2, but neither cusp touches itself, there is only one full-sized ball per cusp diagram and we are only assured of a volume of $\sqrt{3}/4$ in each cusp. In particular, this means that there are five possibilities for α. Checking the number of possibilities for δ, given a particular choice of α, we have a total of fifteen possibilities for the pair of α and β.

Note that the symmetry of the manifold ensures that the volume in each of the two canonical cusps is equal. If that volume is greater than 2, we again have at most one choice for α and one choice for δ. We now obtain an upper bound on the number n of connectors of length less than $\ln(4)$ exactly as we did for the one-cusped case. However, there is now potentially only one full-sized horoball in the cusp diagram for either cusp. Hence, if $v_c(M)$ is the cusp volume in either one of the two canonical cusps we have that

$$2v_c(M) \geq \sqrt{3}/32(7 + 2(n-1)).$$

so $n \leq 32v_c(M)/\sqrt{3} - 5/2$. □

Note that Corollary 4.1 implies that if α and β are a pair of parabolic generators, then $\alpha^n\beta$ and $\alpha\beta^n$ must be hyperbolic isometries for all but a finite number of values of n. In fact, a stronger result is true. The isometries corresponding to $\alpha^n\beta$ and $\alpha\beta^n$ must be hyperbolic isometries for all non-zero values of n. This follows from the proof of Theorem 3.1 of [1], where it is shown that if the product of two parabolic isometries that do not share a fixed point is another parabolic isometry, then the two parabolic isometries preserve a common circle on the sphere at ∞. In our case, the two parabolic isometries are α and β^n or α^n and β. However, then α and β generate a Fuchsian group, contradicting our assumption that the fundamental domain for the group is finite volume.

Lemma 4.2. *Let M be a non-elementary finite volume hyperbolic 3-manifold with fundamental group generated by two parabolic elements, both corresponding to the same cusp. Then the following hold:*

i) M has a single cusp.

ii) Any pair of parabolic generators are freely homotopic to one another or their inverses.

iii) There exists an involution of the manifold that lifts to a parabolic isometry.

Proof. Since the two parabloic generators have distinct fixed points, we know that the manifold has a $Z_2 + Z_2$ action on it, as explained in Section 3, before Theorem 3.2. Each of these isometries of the manifold is realized in the universal cover by an elliptic isometry that permutes the horoballs covering the canonical cusp. In particular, we will normalize so that a lift L_1 of one of the three involutions fixes the horoball at ∞ and the horoball centered at 0 while lifts L_2 and L_3 of the second and third involution switch those two horoballs. However, there must exist an isometry T in the fundamental group that takes the horoball centered at 0 to the horoball centered at ∞. Therefore L_1, TL_2 and TL_3 are all lifts of the three involutions that fix a single common point on the boundary of hyperbolic space and that preserve each horosphere centered at that point. Hence, each of these lifts must either be elliptic or parabolic. Suppose that two of them are elliptic. Then the third nontrivial isometry lifts to the product of these two, which is a parabolic isometry. Therefore, at least one of these three involutions lifts to a parabolic isometry, proving iii). Suppose that TL_2 is the parabolic lift of an involution. But L_2 must conjugate α to β or β^{-1}, and by switching β for β^{-1} if necessary, we will assume L_2 conjugates α to β. Then, since TL_2 and α commute,

$$\alpha = TL_2\alpha(TL_2)^{-1} = T\beta T^{-1}.$$

Hence α and β are conjugate in the fundamental group. This implies that they are freely homotopic as loops in the manifold.

In particular, they are freely homotopic to the same curve in the boundary of the manifold, proving ii). If we do a Dehn filling along a primitive element in the boundary of the manifold that generates this curve, we obtain a 3-manifold with trivial fundamental group. If there was more than one cusp, we can discard the interiors of the remaining cusps, to obtain a compact 3-manifold with trivial fundamental group and toroidal boundary components, a contradiction. Hence, this manifold has only a single cusp. □

Theorem 4.3. *Let M be a non-elementary finite volume hyperbolic 3-manifold. Then the fundamental group of M is generated by two parabolic elements if an only if M is the complement of a two-bridge link in S^3 that is not a 2-braid. The two generators correspond to meridians.*

Proof. Corollary 4.2 shows that if both of the parabolic generators correspond to one cusp, they must be freely homotopic. Since each parabolic element must have length at least 1 and Theorem 3.2 says that the generators have length less than 2, each of the generators must correspond to a loop in the boundary of the manifold that yields a primitive element in the fundamental group of the boundary. In particular, if we glue one or two solid tori to the one or two cusp boundaries so that a meridian of each solid torus goes to the one or two boundary curves corresponding to the parabolics, we obtain a compact 3-manifold with trivial fundamental group. This manifold has a nontrivial involution coming from extending any one of the three nontrivial involutions on M to the surgered manifold. By Thurston's Orbifold Theorem, (which has not yet appeared in print but see [17] and [10]), a compact 3-manifold with an involution that has fixed point set of dimension at least one satisfies the Geometrization Conjecture. In particular, the surgered manifold must be S^3. Hence, M is a knot or link complement in S^3 and the Dehn filling was along a meridian. However, Corollary 6.3 of [7] shows that a link complement in S^3 generated by two meridians must be a 2-bridge link complement.

Note that every 2-bridge link complement has fundamental group generated by two meridians. If the 2-bridge link complement is hyperbolic, those meridians will lift to parabolic isometries. But every 2-bridge link is a prime non-splittable alternating link, and in [12], it is proved that a prime non-splittable alternating link that is not a 2-braid is hyperbolic. Hence, we know that a compact orientable hyperbolic 3-manifold has fundamental group generated by two parabolics if and only if it is the exterior of a 2-bridge link in S^3 that is not a 2-braid. □

A compact orientable 3-manifold with boundary consisting of tori is said to be *tunnel number one* if it is obtained by gluing a 2-handle to a genus two handlebody, and the two-handle is known as an *unknotting tunnel.* Any pair of generators for the fundamental group of the handlebody can then serve as a pair of generators for the fundamental group of the manifold. In [15], it was conjectured that a compact orientable 2-generator 3-manifold is tunnel number one. Since every two-bridge link complement is tunnel number one, Theorem 4.3 confirms the conjecture in the case that the manifold is hyperbolic and the two generators are parabolic. In fact, it was this conjecture that motivated the original work on this paper. Since there are many manifolds that are tunnel number one but that are not two-bridge knot or link complements, we have the following corollary to Theorem 4.3.

Corollary 4.4. *There exist hyperbolic 3-manifolds of tunnel number one and unknotting tunnels within them, such that among all of the pairs of generators corresponding to the given unknotting tunnel, there is no pair, both of the elements of which are parabolic.*

Corollary 4.5. *A meridian in a hyperbolic 2-bridge knot is represented by a nontrivial loop in the maximal cusp boundary of length less than 2.*

Proof. This follows immediately from Theorem 3.2 and the fact that the fundamental group of the complement of a 2-bridge knot is generated by two meridians. □

As mentioned previously, when counting distinct pairs of parabolic generators, we will not distinguish between the pair $\{\alpha, \beta\}$ and the pairs $\{\alpha, \beta^{-1}\}$, $\{\alpha^{-1}, \beta\}$ and $\{\alpha^{-1}, \beta^{-1}\}$.

Corollary 4.6. *There are exactly two distinct pairs of parabolic isometries, up to conjugation of the pair, that generate the fundamental group of the figure-eight knot complement.*

Proof. Taking a maximal cusp for the figure-eight knot complement, the shortest translation in the cusp boundary has a length of 1 and corresponds to a meridian, while the next shortest translation has a length of $2\sqrt{3}$. Hence, Theorem 3.2 implies that both α and δ are meridians of length 1. There are two connectors between horoballs of length 0 and four connectors between horoballs of length $\ln(3)$. See Figure 3. ([5] contains many additional cusp diagrams, and an explanation of what they represent.) Each connector, once it is given an orientation, appears twice in the cusp diagram, once coming into the cusp, and once leaving the cusp. (A connector of length less than $\ln(4)$ corresponds to a horoball of Euclidean diameter greater than $1/4$ where the largest horoballs in the cusp diagram have been normalized to have diameter 1.) There are many more connectors of length exactly $\ln(4)$, however, we needn't consider them.

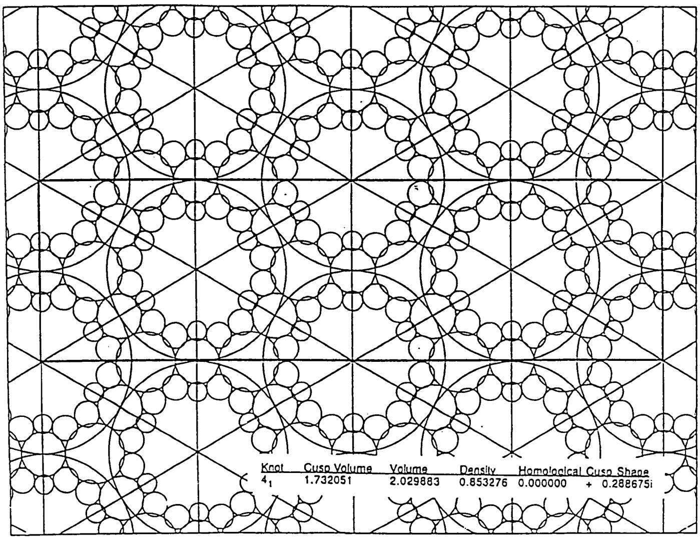

Figure 3: The cusp diagram for the figure-eight knot.

The two pairs of parabolic generators corresponding to the two connectors of length 0 each generate the entire group. In the decomposition of the figure-eight knot complement into two ideal regular tetrahedra, with two equivalence classes of edges, the two connectors of length 0 appear as the two equivalence classes of edges, each class corresponding to an edge in the manifold that is perpendicular to a point where the maximal cusp boundary is tangent to itself. (See [16] for this decomposition of the figure-eight knot complement.)

The four connectors of length $\ln(3)$ correspond to the four edges obtained by gluing one face of one tetrahedron to the corresponding face of the opposite tetrahedron, and then running an edge perpendicular to this conjoined face from the opposite vertex on one of the tetrahedra to the opposite vertex on the other. Each of the four faces on one of the tetrahedra generates an edge in this manner. Symmetries of the figure-eight knot complement allow us to send any one of the four connectors of length $\ln(3)$ to any other. Therefore it is enough to show that one of these connectors produces a pair of parabolic elements that together do not succeed in generating the entire fundamental group of the figure-eight knot complement in order to see that there are only two pairs of parabolic generators.

The fundamental group for the figure-eight knot complement has presentation $\langle a, b : b^{-1}a^{-1}bab^{-1}aba^{-1}b^{-1}a \rangle$ where a and b are meridians. The subgroup corresponding to one of these four connectors of length $\ln(3)$ is generated by b and $c = a^{-1}b^{-1}aba$.

As in [14], we will take the homomorphism sending a to $(11|01)$ and b to $(1\ 0| - \omega\ 1)$ where ω is a complex number such that $\omega^2 + \omega + 1 = 0$. Then

c is sent to

$$(-2\omega^{\vee} - 3\omega| - \omega^{2\vee} - 2\omega^2).$$

Alan Reid pointed out that the quotient by the ideal $\langle\sqrt{-3}\rangle$ turns this into a map onto the non-abelian group $PSL_2(F_3)$, where F_3 is the field with three elements. However, the element c goes to $(1\ 0| - \omega^2\ 1)$ which commutes with the image of b. Hence they generate an abelian subgroup, which must be proper. Therefore, b and c generate a proper subgroup of the fundamental group of the figure-eight knot complement. □

Note that we can state the above result without reference to the hyperbolic geometry. Namely, there are exactly two pairs of generators, up to conjugation of pairs in the fundamental group of the figure-eight knot complement, such that each element in a pair is freely homotopic into the boundary of a regular neighborhood of the missing knot.

Corollary 4.7. *If M is a hyperbolic 3-manifold with a choice of cusps such that all parabolics have lengths with respect to the cusps of at least 2, then all two-parabolic generator subgroups are either free or free abelian.*

Proof. If M is a one-cusped manifold, this is immediate. If M has multiple cusps, then for any pair of cusps in the manifold, we can obtain a canonical set of cusps, without shrinking one of the two cusps. In particular, by our involutions, we obtain a canonical set of cusps such that all parabolics have length at least two with respect to them. Hence, all two-parabolic generator subgroups are either free or free abelian. □

This next lemma demonstrates that the relators for a Kleinian group generated by two parabolic generators are restricted to certain types of words.

Lemma 4.8. *Let α and β be a pair of parabolic generators generating the fundamental group of a hyperbolic 3-manifold or 3-orbifold. Let w be a cyclically reduced word in α and β that is trivial. Then:*

i) w contains at least two one letter syllables.

ii) w contains a subword consisting of a pair of adjacent syllables in w of the form $\alpha^i\beta^j, \alpha^j\beta^i, \beta^i\alpha^j$ and $\beta^j\alpha^i$ where $i = \pm 1$ and $j = \pm 1, \pm 2$ or ± 3.

Proof. Given the word, we have a corresponding set of horoballs, one for each syllable. Divide the sphere at ∞ into subsets of points, one for each horoball in the set, consisting of all of the points that are farther from that horoball than from any other horoball in the set. When we say that a point x is farther from a given horoball H, we mean that if a horoball H' is centered at x and then expanded, the last horoball from the set that it hits is H. This construction will result in at least two non-empty regions on the sphere at ∞. Hence, there are at least two distinct choices for the point at ∞ in the upper-half-space model resulting in two different horoballs in the set appearing as the smallest horoball relative to this choice. The syllables corresponding to each of these horoballs must be of the form $\alpha, \beta, \alpha^{-1}$ or β^{-1} by the proof of Theorem 3.2, proving i).

To see ii), choose the center of one of the horoballs in the set to be the point at ∞. Let H_1 be the smallest horoball in the set. As in the proof of Theorem 3.2, the syllable corresponding to H_1 must be either $\alpha^{\pm 1}$ or $\beta^{\pm 1}$. For convenience, say it is α. Each of the two corresponding connectors leaving H_1 must lead to horoballs at least as large as H_1.

Let H_2 and H_3 be these two connected horoballs. By Lemma 2.5, one of these two neighboring horoballs, call it H_2, is intersected by its connector at a point that is a distance at most 2 from its peak. In order that the horoball following H_2 in the chain be no smaller than H_1, the parabolic translation corresponding to H_2 must have length at most 4. Since $|\beta| \geq 1$, the syllable corresponding to this horosphere is β^n, where n must equal $\pm 1, \pm 2$ or $\pm 3, \pm 4$. We will show that we can find a pair of syllables where we have eliminated the possibility that $n = \pm 4$. If either of the adjacent horoballs has its connector a distance less than 2 from its peak, $n \neq \pm 4$. If neither of the adjacent horoballs have the connector a distance less than 2 from their peaks, then they are tangent to one another. In this case, $|\alpha| = 1$ and x has length 0 to prevent $|\alpha|$ from being shorter than 1. If $n = \pm 4$, then $|\beta| = 1$. If each of H_2 and H_3 corresponds to a syllable of the form $\beta^{\pm 4}$, then the next horoball in either direction must correspond to $\alpha^{\pm 1}$, or else H_1 is not the smallest horoball in the set. Going one ball further out in either direction, either the distance on that ball from the point of tangency to the peak is strictly less than 2, in which case we can show that this ball corresponds to the syllable β^i with $i = \pm 1, \pm 2$ or ± 3, or we repeat the construction two balls out further still in each direction. Unless we eventually come across a ball corresponding to the syllable β^i with $\beta = \pm 1, \pm 2$ or ± 3, the construction continues ad infinitum with all centers of balls in a line and the sequence of balls never closing up into a cycle, a contradiction. □

Corollary 4.9. *Let M be a one-cusped hyperbolic 3-manifold that has at least n distinct pairs of parabolic generators up to conjugacy of the pairs. Then $v_c(M) \geq \sqrt{3}(n+6)/32$ and $\mathrm{vol}(M) \geq (n+6)v_0/16$, where v_0 is the volume of an ideal regular tetrahedron, approximately equal to* 1.01494.

Proof. Since all of the elements in the pairs of generators are all freely homotopic, there must be at least n distinct connectors. If we normalize so that a horosphere covering the cusp is given as a plane at height 1, then since a connector lifts to a vertical geodesic that connects the horosphere centered at ∞ with a horosphere of diameter at least $1/4$ and since each connector must lift to two distinct vertical geodesics up to the action of the cusp subgroup, it must be the case that there are $2n$ balls of diameter at least $1/4$ up to the action of the cusp subgroup. Since there is always a pair of full-sized balls, and since the center of a ball of diameter $1/4$ can only come within a distance of $1/2$ of the center of a ball of diameter 1, we can pack a disk of diameter 1 with seven disks of diameter $1/4$, replace the balls of diameter $1/4$ with disks of diameter $1/4$, and pack these disks in the plane. Each disk of diameter $1/4$ contributes a total of $\sqrt{3}/32$ to the area of a fundamental parallelogram, yielding a total area of at least $(14+2(n-1))\sqrt{3}/32$. Hence the total volume is at least half of this, namely $\sqrt{3}(n+6)/32$. By Meyerhoff's application of Boroczky's horoball packing arguments in hyperbolic 3-space (cf. [13]), we have that the volume of the manifold is at least

$$v_c(M)2v_0/\sqrt{3} = (n+6)v_0/16.$$

□

Note that Theorem 4.3 implies that all of the manifolds considered in Corollary 4.9 are 2-bridge link complements.

5. A Pair of Hyperbolic Generators.

We are interested in determining what can be said in the case that we have a pair of generators, both of which are hyperbolic.

Lemma 5.1. *Let α and β be a pair of hyperbolic generators for a non-elementary Kleinian group that is not free, such that $|\alpha| \leq |\beta|$. If x is the distance between the images of the axes of the two hyperbolic generators in* *the quotient of hyperbolic space by the group, then* $x \leq 2\ln(\coth(|\alpha|/4))$. *(Equivalently,* $|\alpha| \leq 2\ln(\coth(x/4))$.

Proof. Let g_α and g_β be the geodesics in the quotient orbifold corresponding to α and β and let g_x be the geodesic segment of length x that is the shortest path from one to the other. Lifting g_α, g_β and g_x to H^3 we have a graph in H^3 that must be connected, since α and β generate the entire group, and that contains nontrivial cycles, since the group is not free, as in the proof of Theorem 3.2. In particular, there must be a cyclic sequence of geodesics, alternately covering g_α and g_β that are connected consecutively by lifts of g_x. By moving the point at ∞ if necessary, we can be sure that the lifts of g_α and g_β in the sequence do not all have peaks at the same height. Let g be the geodesic with the lowest peak from among them. If there is more than one geodesic with a peak at the same height as the peak of g, we can choose g so that one of the geodesics in the sequence that is connected to it by a lift of g_x has a strictly higher peak than it does. The geodesic g either covers g_α or g_β. For convenience, let c represent the length of the geodesic g_α or g_β that it covers.

Then there are two lifts of g_x that connect g to two geodesics in the sequence, both of which either cover g_α or g_β. Suppose first of all that the two lifts γ_1 and γ_2 of g_x intersect g on opposite sides of the peak of g. Then one of them, say γ_1, is a distance greater than or equal to $c/2$ on g from the peak. Similarly, if γ_1 and γ_2 intersect g on the same side of the peak of g, then one of them, say γ_1, is a distance of at least $c/2$ on g from the peak of g.

By Lemma 2.3, the distance from the intersection of γ_1 and g to the peak of γ_1 is no more than $\ln(\coth(c/4))$. If the three geodesics g, γ_1 and γ_2 are all in a vertical plane, then it is clear that the distance down the other side of γ_1 to the next geodesic in the sequence must be no more than $\ln(\coth(c/4))$ in order for the next geodesic in the sequence to have its peak at least as high as the peak of g. So $x \leq 2\ln(\coth(c/4))$. On the other hand, suppose that the three geodesics are not in a vertical plane. If the third geodesic is not in a plane with the first two, its peak is even lower than when it does lie in a plane with the first two. If the second geodesic is not in a plane with the first, then the distance to its peak goes down, forcing x to be even smaller than $2\ln(\coth(c/4))$. Since $2\ln(\coth(x/4))$ is a monotonically decreasing function, and since $|\beta| \geq |\alpha|$, we have that

$$x \leq 2\ln(\coth(|\alpha|/4)).$$

Note that $2\ln(\coth(x/4))$ is a monotonically decreasing function that is its own inverse, so we obtain the alternative inequality that $|\alpha| \leq 2\ln(\coth(x/4))$. □

Unlike the case for parabolic elements, where we have a lower bound of one on the shortest length of the parabolic with respect to a certain set of cusps, there is no universal lower bound on the length of a hyperbolic element. However, there is a lower bound on the lengths of the hyperbolic elements in a given finite volume manifold.

Corollary 5.2. *If I_0 is the length of the shortest geodesic in a hyperbolic 3-manifold or 3-orbifold M, then for any pair of distinct geodesics such that the corresponding pair of hyperbolic elements do not generate a free group, the distance between those geodesics x satisfies $x \leq 2\ln(\coth(I_0/2))$.*

Proof. The fact that the geodesics are distinct ensures that the group is non-elementary. The result then follows immediately from Lemma 5.1. □

Lemma 5.3. *Let α and β be hyperbolic elements that generate a non-elementary Kleinian group that is not free. Let x be the distance in the quotient orbifold between the geodesics that correspond to α and β. Then at least one of $|\alpha|, |\beta|$ and x is strictly less than $2\ln(\sqrt{2}+1)$.*

Proof. We have a cyclic sequence of geodesics that intersect at right angles, the odd ones alternately covering g_α and g_β while the even ones cover x. We can choose a point at ∞ so that they do not all have peaks at the same height. Let g be the geodesic in the sequence with the lowest height. We can choose g so that at least one of its neighboring geodesics has a strictly higher peak.

Let g_1 and g_2 be the two neighbors to g in the sequence. Let c be the distance between g_1 and g_2 on g. If g_1 and g are both in the same vertical plane, then, in order for the peak of g_1 to be at least as high as the peak of g, the angle ϕ in Figure 4 must be at least $\pi/4$, making the distance on g from the point of intersection to the peak at most $\ln(\sqrt{2}+1)$.

As we rotate g_1 out of the same vertical plane with g, the peak of g_1 will occur at a lower level. In order to keep the peak of g_1 above the peak of g, it is necessary to shrink the distance from the point of intersection to the peak on g. Hence, the distance from the point of intersetion of g_1 and g to the peak of g is always at most $\ln(\sqrt{2}+1)$. The same is true for g_2, implying

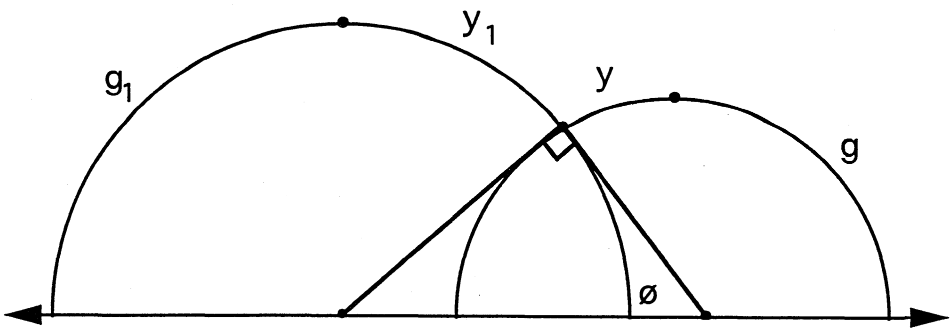

Figure 4.

that c is at most $2\ln(\sqrt{2}+1)$. Since one of g_1 or g_2 must have a peak that is strictly higher than the peak of g, we have that $c < 2\ln(\sqrt{2}+1)$. The length c could either be equal to the length x, equal to $n(|\alpha|)$ or equal to $n(|\beta|)$ where n is a positive integer. Hence, one of $x, |\alpha|$ or $|\beta|$ is less than $2\ln(\sqrt{2}+1)$. □

References.

[1] C. Adams, *Thrice-punctured spheres in hyperbolic 3- manifolds*, Trans. Amer. Math. Soc., **287** (1985), 645-656.

[2] C. Adams, *The noncompact hyperbolic 3-manifold of minimal volume*, Proc. Amer. Math. Soc., **100** (1987), 601-606.

[3] C. Adams, *Dehn filling hyperbolic 3-manifolds*, Pac. J. Math., **165** (1994), 217-238.

[4] C. Adams, *Unknotting tunnels in hyperbolic 3- manifolds*, Math. Ann., **302** (1995), 177-195.

[5] C. Adams, M. Hildebrand and J. Weeks, *Hyperbolic invariants of Knots and Links*, Trans. Amer. Math. Soc., **326** (1991), 1-56.

[6] S. Bleiler and C. Hodgson, *Spherical space forms and Dehn filling*, Knots 90, de Gruyter, (1992), 425-433.

[7] M. Boileau and B. Zimmerman, *The π-orbifold group of a link*, Math. Z., **200** (1989), 187-208.

[8] J.L. Brenner, *Quelques groups libres de matrices*, C.R. Acac. Sci. Paris, **241** (1955), 1689-1691.

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

[10] C. Hodgson, *Geometric structures on 3-dimensional orbifolds: Notes on Thurston's Proof*, preprint.

[11] W. Jaco and P.Shalen, *Seifert fibered spaces in 3-manifolds*, Memoirs of the Amer. Math. Soc., **21** (1979).

[12] W. Menasco, *Closed incompressible surfaces in alternating knot and link complements*, Topology, **23** (1984), 37- 44.

[13] R. Meyerhoff, *Sphere-packing and volume in hyperbolic 3-space*, Comment. Math. Helv., **61** (1986), 271-278.

[14] R. Riley, *A quadratic parabolic group*, Math. Proc. Camb. Phil. Soc., **77** (1975), 281-288.

[15] M. Scharlemann, *Tunnel number one knots satisfy the Poenaru Conjecture*, Topology and its Applications, **18** (1984), 235-258.

[16] W. Thurston, *The Geometry and Topology of 3-Manifolds*, lecture notes, (1978) (revised version to be published by Princeton University Press.)

[17] W. Thurston, *Three manifolds with symmetry*, research announcement.