Submitted exclusively to the London Mathematical Society
doi:10.1112/0000/000000

A COMBINATION THEOREM FOR CONVEX HYPERBOLIC MANIFOLDS, WITH APPLICATIONS TO SURFACES IN 3-MANIFOLDS

MARK BAKER AND DARYL COOPER

Abstract

We prove the convex combination theorem for hyperbolic n-manifolds. Applications are given both in high dimensions and in 3 dimensions. One consequence is that given two geometrically finite subgroups of a discrete group of isometries of hyperbolic n-space, satisfying a natural condition on their parabolic subgroups, and whose intersection is a separable subgroup, there are finite index subgroups which generate a subgroup that is an amalgamated free product. Constructions of infinite volume hyperbolic n-manifolds are described by gluing lower dimensional manifolds. It is shown that every slope on a cusp of a hyperbolic 3-manifold is a multiple immersed boundary slope. If the fundamental group a hyperbolic 3-manifold contains a maximal surface group not carried by an embedded surface then it contains a freely indecomposable group with second betti number at least 2.

1. *Introduction*

The Klein-Maskit combination theorems [**30**] assert that under certain circumstances two Kleinian groups, Γ_1, Γ_2 corresponding to hyperbolic manifolds M_1, M_2 are subgroups of another Kleinian group Γ corresponding to a hyperbolic manifold M, diffeomorphic to that obtained by gluing M_1 and M_2 along a boundary component. There is also an HNN version, and an orbifold version.

We prove the *convex combination theorem* which allows the gluing of two convex hyperbolic n-manifolds along isometric submanifolds which are not necessarily boundary parallel. There is also an HNN version. One consequence is the *virtual amalgam theorem* which states that that if G and H are two geometrically-finite subgroups of a discrete group in $Isom(\mathbb{H}^n)$, and if their parabolic subgroups are *compatible* (defined in section 5), and if $G \cap H$ is subgroup separable in both G and H, then there are finite index subgroups $G' \subset G, H' \subset H$, each containing $G \cap H$, which generate a geometrically finite group in $Isom_+(\mathbb{H}^n)$, that is the free product of G' and H' amalgamated along $G \cap H$.

One application is proving the existence of certain kinds of surface groups in hyperbolic 3-manifolds. For example we show that given any slope on a torus boundary component of a compact 3-manifold, M, with hyperbolic interior, there is a compact, immersed, oriented, geometrically-finite, π_1-injective, surface not homotopic into the boundary of M whose boundary consists of two components each of which wraps the same number of times, but in opposite directions, around the given slope. In particular ∞ is a multiple immersed boundary slope of every hyperbolic knot. The first author introduced this concept in [**2**] and gave an example of a hyperbolic once-punctured torus bundle with infinitely many immersed boundary slopes. Oertel [**33**] gave the first example of a hyperbolic manifold such that every slope is an immersed boundary slope, and Maher [**28**] proved the latter holds for hyperbolic 2-bridge knots and certain other cases.

Another application is that if the fundamental group of a non-Haken closed hyperbolic 3-manifold contains a surface group then it contains the fundamental group of an irreducible, boundary-irreducible, compact 3-manifold with arbitrarily large second betti number.

2000 *Mathematics Subject Classification* 57M50 (Primary), 30F40 (secondary).

The basic idea is to glue two convex hyperbolic manifolds together to obtain a hyperbolic manifold which is not in general convex. It is well known that every 3-manifold that is not closed admits a non-convex hyperbolic metric, thus such metrics are in general too abundant to provide useful information. What is required is a condition which ensures that the result of the gluing can be thickened to be convex. If this can be done then one obtains a discrete subgroup of isometries of hyperbolic space.

The *convex combination theorem* asserts (roughly speaking) that there is a universal constant κ, independent of dimension, such that if two convex hyperbolic manifolds can be glued together in a way that is compatible with gluing their κ-thickenings, then the resulting manifold can be thickened to be convex.

To deduce the virtual amalgam theorem from the convex combination theorem involves two hyperbolic manifolds, A and B, which are isometrically immersed into a hyperbolic manifold M. One wishes to glue the basepoints of A and B together and the requirement that the identification space is a manifold forces further identifications to be made between A and B. Subgroup separability arguments are used to ensure that certain finite covers of A and B can be glued so that they embed in the resulting identification space. In order to satisfy the thickening hypotheses of the convex combination theorem, one might need to take large finite covers of the manifolds concerned.

The paper is organized as follows. In section 2 we discuss convex hyperbolic manifolds and prove the *convex combination theorem* 2.9. In section 3 we study cusps of geometrically-finite, convex hyperbolic manifolds. In section 4 we introduce the notion of *induced gluing* alluded to in the preceding paragraph and prove the *virtual simple gluing theorem* 4.3 which ensures the manifolds that are glued embed in the resulting space. In section 5 we prove the *virtual amalgam theorem* 5.3 and the *virtual convex combination theorem* 5.1. In section 6 we construct some higher dimensional convex hyperbolic manifolds of infinite volume by gluing lower dimensional ones. In section 7 we show that certain groups are LERF and extend an argument of Scott's that finitely-generated subgroups of surface groups are almost geometric to the case of finitely-generated separable subgroups of three-manifold groups. In section 8 we prove some new results about surface groups in hyperbolic 3-manifolds as well as sketching new proofs of some results of the second author and Long about virtually-Haken Dehn-filling and the existence of surface groups in most Dehn-fillings. In section 9 we apply these tools to the study of immersed boundary slopes.

The convex combination theorem is related to work of Bestvina-Feighn [**4**], Gitik [**18**] and Dahmani [**14**] who proved various combination theorems for (relatively) word hyperbolic groups. The convex combination theorem implies that certain groups are discrete groups of isometries of hyperbolic space, a conclusion which does not follow from the group-theory results mentioned. The work of Bowditch-Mess [**6**] contains many of the ingredients for the non-cusp case. As is often the case, a lot of extra work is needed to handle cusps. This is needed for the application to immersed boundary slopes. It seems possible that there is a common generalization of the Klein-Maskit theorem and the convex combination theorem, but this will have to await a mythical future paper.

The train of ideas in this paper originated with the work of B. Freedman and M.H. Freedman [**16**] who constructed certain closed surfaces by a tubing operation. If the surfaces involved are *far enough apart* in a certain sense (if the tube is *long enough*) then the resulting surface is π_1-injective. There are by now several proofs of this and related facts, and this paper provides new ones.

2. *The Convex Combination Theorem.*

In this section we prove the convex combination theorem. This requires a brief discussion of non-convex hyperbolic manifolds. To do this we need to extend some ideas from the more well-known context of convex to that of non-convex hyperbolic manifolds.

Definition. A *hyperbolic manifold* is a connected manifold with boundary (possibly empty) equipped with a Riemannian metric which is hyperbolic i.e. constant sectional curvature -1.

Warning. We do not assume the holonomy of a hyperbolic manifold is a discrete group of isometries of hyperbolic space.

We will primarily be interested in the case that the boundary is piecewise smooth, but has corners. Let $\tilde{M}$ denote the universal cover of a hyperbolic manifold M. There is a local isometry called the *developing map* and a homomorphism of groups called the *holonomy*

$$dev : \tilde{M} \to \mathbb{H}^n \qquad\qquad hol : \pi_1(M) \longrightarrow Isom(\mathbb{H}^n)$$

such that for all $x \in \tilde{M}$ and all $g \in \pi_1 M$ we have $dev(g \cdot x) = hol(g) \cdot dev(x)$.

Definition. A hyperbolic manifold is *convex* if every two points in the universal cover are connected by a geodesic.

In particular the quotient of hyperbolic space by a discrete torsion-free group of isometries is convex. The following is easy to check:

Proposition 2.1 (characterize convex). *Suppose that M is a hyperbolic manifold. Then the following are equivalent.*
(a) M is convex.
(b) Every path in M is homotopic rel endpoints to a geodesic in M.
(c) The developing map is injective with image a convex subset of hyperbolic space.

Proposition 2.2 (convex has injective holonomy). *If M is a convex hyperbolic manifold then the holonomy is injective. Furthermore $M = dev(\tilde{M})/hol(\pi_1 M)$.*

Proof. It follows from 2.1(c) that the developing map is an isometry onto its image. Since $\pi_1 M$ acts freely by isometries on its universal cover, the holonomy is injective and M is isometric to $dev(\tilde{M})/hol(\pi_1 M)$. □

Proposition 2.3 (local isometry from convex is π_1-injective). *Suppose M and N are hyperbolic manifolds, M is convex, and $f : M \to N$ is a local isometry. Then $f_* : \pi_1 M \to \pi_1 N$ is injective. In particular, if $N = \mathbb{H}^n$ then M is simply connected.*

Proof. It is easy to check that $hol_M = hol_N \circ f_*$. Since M is convex, hol_M is injective by 2.2. Thus f_* is injective. □

Consider two closed geodesics in a hyperbolic surface which intersect in two points. Let A and B be small convex neighborhoods of these geodesics. Then $A \cap B$ is the disjoint union of two discs each of which is convex. More generally we have:

LEMMA 2.4 (intersection of closed convex is convex union). *Suppose* $\{M_i : i \in I\}$ *are metrically-complete, convex hyperbolic manifolds contained in a hyperbolic* n*-manifold* M. *Then every component of* $\cap_{i\in I} M_i$ *is a convex hyperbolic manifold.*

Proof. Let $p : \tilde{M} \to M$ be the universal cover and $\tilde{M}_i$ a component of $p^{-1}(M_i)$. Since M_i is convex it follows from 2.3 applied to $M_i \hookrightarrow M$ that $p| : \tilde{M}_i \to M_i$ is the universal cover.

Let C be a component of $\cap_{i\in I} M_i$. Let $\tilde{C}$ be a component of $p^{-1}(C)$. Consider the components $\tilde{M}_i \subset p^{-1}M_i$ which contain $\tilde{C}$. Since dev_M embeds each $\tilde{M}_i$ it also embeds $K = \cap_i \tilde{M}_i$ into $\mathbb{H}^n$. Thus $dev_M(K) = \cap_i\, dev_M(\tilde{M}_i)$ is a closed (because metrically-complete) convex subset of $\mathbb{H}^n$, and therefore a manifold. Hence K is a convex manifold. Clearly $\tilde{C} \subset K$. Choose $x \in \tilde{C}$ and $y \in K$ then, because K is convex, there is a unique geodesic segment $[x, y]$ in K with endpoints x and y. This segment is in every $\tilde{M}_i$ and therefore $p([x, y])$ is contained in every M_i and thus in C. It follows that $[x, y]$ is contained in $\tilde{C}$ hence $y \in \tilde{C}$. Thus $\tilde{C} = K$. It follows that C is a convex hyperbolic manifold. □

DEFINITION. Suppose M is a metrically-complete convex hyperbolic n-manifold and A is a non-empty, connected subset of M. The *convex hull*, $CH(A)$, of A is defined to be the intersection of all convex manifolds in M which are closed subsets of M and which contain A.

PROPOSITION 2.5 (convex hulls are convex). *If* M *is a metrically-complete convex hyperbolic* n*-manifold and* A *is a non-empty, connected subset of* M, *then* $CH(A)$ *is a convex manifold of some dimension* $k \le n$.

Proof. Since A is connected there is a unique component, C, of $CH(A)$ which contains A. By 2.4 C is a convex manifold which contains A and $CH(A) = C$. □

There are many examples of non-convex hyperbolic manifolds. For example an immersion of a punctured torus into the hyperbolic plane induces a pull-back metric on the punctured torus with trivial holonomy. Every non-compact 3-manifold can be immersed into Euclidean space and hence into $\mathbb{H}^3$. It follows that every such manifold has a hyperbolic metric.

We want to know when a non-convex manifold corresponds to a discrete group of isometries. The preceding examples do not have this property. There are several equivalent ways to describe the desired property, and one involves the notion of thickening.

DEFINITION. A hyperbolic n-manifold N is a *thickening* of a hyperbolic n-manifold, M, if $M \subset N$ and $incl_* : \pi_1 M \to \pi_1 N$ is an isomorphism. If, in addition, N is convex then we say N is a *convex thickening* of M.

The following is easy to check:

PROPOSITION 2.6 (convex thickenings). *Suppose that* M *is a hyperbolic* n*-manifold. Then the following are equivalent.*
(a) The developing map $dev : \tilde{M} \to \mathbb{H}^n$ *is injective.*
(b) The holonomy of M *is a discrete torsion-free group* $\Gamma \subset Isom(\mathbb{H}^n)$ *and there is an isometric embedding* $f : M \to N = \mathbb{H}^n/\Gamma$ *such that* $f_* : \pi_1 M \to \pi_1 N$ *is an isomorphism.*
(c) M *has a convex thickening.*

We will often use the developing map to identify the universal cover, $\tilde{M}$, of a convex manifold with $dev(\tilde{M}) \subset \mathbb{H}^n$. If M is a convex hyperbolic manifold and $K \ge 0$ the K*-thickening* of M is

$$Th_K(M) \;=\; N_K(\tilde{M})/\pi_1 M$$

where $N_K(\tilde{M}) = \{x \in \mathbb{H}^n : d(x, \tilde{M}) \leq K\}$ is the *K-neighborhood* of $\tilde{M}$ in $\mathbb{H}^n$. With this notation $Th_\infty(M) = \mathbb{H}^n/hol(\pi_1 M)$ is the geodesically-complete manifold that is a thickening of M. The following is immediate:

PROPOSITION 2.7 (thickening). *Suppose M is a convex hyperbolic manifold. Then:*
(a) $Th_K(M)$ is a convex thickening of M.
(b) If $x \in M$ and $y \in \partial Th_K(M)$ then $d(x, y) \geq K$.

The convex combination theorem 2.9 gives a sufficient condition to ensure that the union of two convex hyperbolic n-manifolds has a convex thickening, and so has holonomy a discrete subgroup of $Isom(\mathbb{H}^n)$. The following example shows that some additional hypothesis is needed to ensure this.

EXAMPLE. Suppose that $S(\ell, \theta, K) = M_1 \cup M_2$ is homeomorphic to a punctured torus with an incomplete hyperbolic metric and that M_1 and M_2 are hyperbolic annuli isometric to K-neighborhoods of closed geodesics γ_1, γ_2 of length ℓ. Suppose $M_1 \cap M_2$ is a disc and that the angle between the geodesics γ_1, γ_2 is θ. Then given $\theta \in (0, \pi)$, the set of $\ell > 0$ for which there is $K > 0$ such that the developing map $dev : \tilde{S} \to \mathbb{H}^2$ is injective is an interval $[\ell_0(\theta), \infty)$. By Margulis's theorem there is $\mu > 0$ such that $\ell_0 > \mu$ for all θ. Also $\ell_0(\theta) \to \infty$ as $\theta \to 0$. On the other hand it is easy to convince oneself that if $K > 100$ then the developing map is always injective.

As discussed in the introduction, in the setting of geometric group theory there are combination theorems for word-hyperbolic groups where the conclusion is that the combination is word hyperbolic. The next result holds even when the manifold M has no convex thickening. This shows that the property of being a discrete subgroup is distinct from the structure of the abstract group. The may be thought of as the difference is accounted for by the difference between the intrinsic geometry of the word metric and the extrinsic geometry as a subgroup $Isom(\mathbb{H}^n)$.

COROLLARY 2.8 (union of convex gives amalgamated free product). *Suppose $M = M_1 \cup M_2$ is a connected hyperbolic n-manifold which is the union of two convex hyperbolic n-manifolds M_1, M_2 and suppose that $M_1 \cap M_2$ is connected. Given a basepoint $x \in M_1 \cap M_2$ then $\pi_1(M, x) = \pi_1(M_1, x) *_G \pi_1(M_2, x)$ where $G = \pi_1(M_1 \cap M_2, x)$.*

Proof. By 2.4 $M_1 \cap M_2$ is convex. It follows from 2.3 that $M_1 \cap M_2$ is π_1-injective in M. The result follows from Van Kampen's theorem. □

DEFINITION. Suppose that N is a hyperbolic manifold and $M \subset N$ is a submanifold. Given $K > 0$ we say that N contains a *K-neighborhood* of M if for every $p \in M$ and every tangent vector $v \in T_pM$ with $||v|| \leq K$ then $exp_p(v) \in N$.

Thus a K-thickening of a convex manifold is a K-neighborhood. However, for a K-neighborhood we do not assume $\pi_1 M \cong \pi_1 N$. An example is provided by any open subset $M \subset \mathbb{H}^n = N$.

NOTATION. If X is a metric space the notation $\overline{X}$ denotes the metric (not geodesic) completion of X. If a, b are two points in $\mathbb{H}^n$ we denote the geodesic segment with these endpoints by $\overline{ab}$.

The following theorem asserts, very roughly, that there is a universal constant, κ, such that if M is a (probably non-convex) hyperbolic n-manifold which is the union of two convex

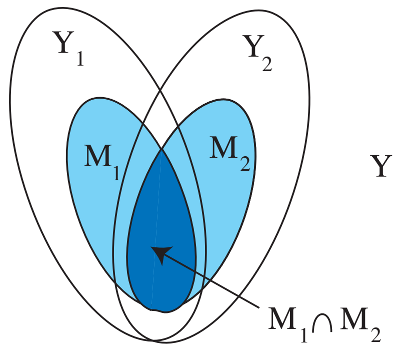

FIGURE 1. *The union of two convex manifolds which have thickenings.*

hyperbolic submanifolds and if M has a κ-thickening with the same topology, then M has a convex thickening. The intuition for this result is lemma 2.12 which says that if two convex sets in hyperbolic space intersect then their convex hull is within a small distance of their union. The fact one can thicken the submanifolds without bumping means that the universal cover is made of convex sets (covers of the submanifolds) which are far apart in some sense. Then the convex hull construction more or less only notices two of the convex sets at any one time and so the convex hull is close to the union.

THEOREM 2.9 (convex combination theorem). *There is a universal constant, κ, called the* **thickening constant** *with the following property. Suppose the following conditions are satisfied:*

(i) *$Y = Y_1 \cup Y_2$ is a connected hyperbolic n-manifold which is the union of two convex n-submanifolds Y_1 and Y_2.*
(ii) *$M = M_1 \cup M_2$ is a connected hyperbolic n-manifold which is the union of two convex n-submanifolds M_1 and M_2.*
(iii) *Y_i is a thickening of M_i for $i = 1, 2$.*
(iv) *Y contains a κ-neighborhood of M.*
No bumping:
(v) *Y_i contains a κ-neighborhood of $\overline{M_i} \setminus int(M_1 \cap M_2)$ for $i \in \{1, 2\}$.*
(vi) *Every component of $Y_1 \cap Y_2$ contains a point of $M_1 \cap M_2$.*

Then M has a convex thickening; in other words there is a hyperbolic n-manifold $N = \mathbb{H}^n/\pi_1 M$ which contains M and $incl_* : \pi_1 M \longrightarrow \pi_1 N$ is an isomorphism. Furthermore, if Y has finite volume then M is geometrically finite. Also $\kappa \leq 6$.

REMARKS.

(i) This theorem may be viewed as extending the generalized Thurston bending construction of Bowditch and Mess [**6**].
(ii) For the statement of this theorem we have not assumed that the manifolds are metrically complete. For example the boundary might be missing. This is why we write $\overline{M}_i$ instead of M_i in condition (v).
(iii) If Y_i is a κ-thickening of M_i then conditions (iii),(iv) and (v)are satisfied.
(iv) Condition (v) implies that at all points, $p \in M_i \setminus (M_1 \cap M_2)$ the restriction of the exponential map, $exp_p|T_p^{\kappa}(M)$, to tangent vectors in $T_p(M)$ of length at most κ is injective.
(v) Conditions (v) and (vi) ensure that the copies of the universal covers of the M_i fit together in hyperbolic space in a treelike way to create $dev(\tilde{M})$.

(vi) If Y contains a rank-2 cusp which contains a rank-1 cusp $C \subset M_1$ then Y does not contain a κ-thickening of C. Attempting to thicken a lower rank cusp of M_1 that is contained in a higher rank cusp of Y causes the thickening to bump into itself within the cusp of Y.

Proof of the Convex Combination Theorem. Here is a sketch of the proof that $\kappa = 6$ satisfies the theorem; the details follow. We need to show that the developing map sends the universal cover, $\tilde{M}$, of M injectively into hyperbolic space. To do this we will show that between any two points in $\tilde{M}$ there is a geodesic in the universal cover, $\tilde{Y}$, of Y connecting them. The image of this geodesic under the developing map is then a geodesic in hyperbolic space and therefore the endpoints are distinct.

We need to understand the universal cover of Y and its image under the developing map. Convexity of Y_1 and Y_2 implies that $\tilde{Y}$ is a union of copies of the universal covers of Y_1 and Y_2. We show $\pi_1 M \cong \pi_1 Y$, and it follows that $\tilde{Y}$ contains the universal cover, $\tilde{M}$, of M and each copy of the universal cover of Y_i contains exactly one copy of the universal cover of M_i.

To show there is such a geodesic we take a shortest path, γ, in the 2-neighborhood (contained in $\tilde{Y}$) of $\tilde{M}$ between the two points. Since M_i is convex, near any point of γ which is in the interior of the 2-neighborhood of either $\tilde{M}_1$ or $\tilde{M}_2$, the path is a geodesic. Thus γ can only fail to be a geodesic at corners that are on the intersection of the boundaries of the 2-neighborhoods. So it suffices to show γ has no corners.

Condition (v) is used to show that the distance between corners is large (bigger than $\kappa - 4$). Thus γ is a union of long geodesic segments each of which starts and ends within a distance of 2 of some convex set, $\tilde{M}_i$. It follows that the midpoint of such a segment is then very close (less than 1) to $\tilde{M}_i$. We deduce that near a corner of γ there is a sub-path, γ', which consists of two long geodesic segments that meet at that corner and one endpoint of γ' is very close to some $\tilde{M}_1$ and the other to some $\tilde{M}_2$.

By condition (vi) the convex sets $\tilde{M}_1, \tilde{M}_2$ intersect near the corner. In hyperbolic space the union of two convex sets which intersect is nearly convex: every point in the convex hull of the union lies within a distance 1 of the union of the convex sets. Thus the geodesic, δ, with the same endpoints as γ', stays less than a distance 2 from the union of $\tilde{M}_1$ and $\tilde{M}_2$. Since γ' is length-minimizing in this set, it follows that $\gamma' = \delta$ and thus γ' does not have any corners. This completes the sketch.

Claim 1. The inclusion $incl : M \hookrightarrow Y$ induces a π_1-isomorphism.

Proof of Claim 1. First we show $incl_*$ is surjective. Suppose that γ is a loop in Y based at a point in M. Then $\gamma = \gamma_1 \cdot \gamma_2 \dots \gamma_n$ with $n \geq 2$ where each γ_i is a path contained either in Y_1 or in Y_2. The endpoints of γ_i are in $Y_1 \cap Y_2$. Using condition (vi) we can homotop γ so that for all i the endpoints of γ_i are contained in $M_1 \cap M_2$. Suppose γ_i is contained in Y_j. Since Y_j is a thickening of M_j we may homotop γ_i into M_j keeping the endpoints fixed. Thus we may homotop γ keeping endpoints fixed into M. Hence $incl_*$ is surjective.

Now we show $incl_*$ is injective. By (iii) Y_i is a thickening of M_i thus $M_i \hookrightarrow Y_i$ induces a π_1-isomorphism. Suppose that γ is an essential loop in M which is contractible in Y. By 2.3 γ is not contained in M_i. Thus $\gamma = \gamma_1 \cdot \gamma_2 \dots \gamma_n$ where each γ_i is a path contained in either M_1 or M_2 and $n \geq 2$. Using convexity of M_1 and M_2 we may assume each γ_i is a geodesic and has both endpoints in $M_1 \cap M_2$. We may suppose that γ is chosen so that n is minimal. This implies that no γ_i is contained in $M_1 \cap M_2$.

Without loss of generality suppose that γ_i is contained in M_1. We claim that γ_i is not contained in Y_2. Otherwise γ_i is a geodesic in the convex manifold Y_2 with both endpoints in the convex submanifold M_2. Consider the universal cover $p : \tilde{Y}_2 \to Y_2$. Then $\tilde{M}_2 = p^{-1} M_2$ is the universal cover of M_2, and in particular is connected. A lift, $\tilde{\gamma}_i$, of γ_i is a geodesic in $\tilde{Y}_2$

with both endpoints in $\tilde{M}_2$. Since $\tilde{M}_2$ is convex it follows that $\tilde{\gamma}_i$ is contained in $\tilde{M}_2$ and thus γ_i is contained in M_2. This contradicts that γ_i is not contained in $M_1 \cap M_2$. It follows, that for each i, that γ_i is contained in exactly one of Y_1 and Y_2.

By 2.3 $Y_i \hookrightarrow Y$ is π_1-injective. Thus the universal cover, $\tilde{Y}$, of Y is a union of copies of the universal covers, $\tilde{Y}_1, \tilde{Y}_2$, of Y_1 and Y_2. By 2.4 the components of $Y_1 \cap Y_2$ are convex, thus π_1-injective into $\pi_1 Y$ by 2.3. It follows that $\pi_1 Y$ is a graph of groups. The copies of the covers of $\tilde{Y}_1$ and $\tilde{Y}_2$ fit together in a tree-like way to give $\tilde{Y}$. This is described in more detail in Serre, [**35**]. It follows that a lift of the path γ to $\tilde{Y}$ does not start and end at the same point. Hence γ is not contractible in Y. This proves claim 1. □

For $i \in \{1, 2\}$ after replacing M_i by its metric completion, we may assume that M_i is a complete metric space. Condition (v) has the following consequence. Suppose that α is an arc in Y which has both endpoints in $M_2 \cap \partial M_1$. If $length(\alpha) \leq 2\kappa$ then α is homotopic rel endpoints into M_2. A similar result holds with M_1 and M_2 interchanged.

Let $\pi : \tilde{Y} \to Y$ denote the universal cover and $dev : \tilde{Y} \to \mathbb{H}^n$ the developing map. Claim 1 implies that $\tilde{M} = \pi^{-1}M$ is the universal cover of M. As observed above, $\tilde{Y}$ is the union of covering translates of the universal covers, $\tilde{Y}_1$ and $\tilde{Y}_2$ of Y_1 and Y_2. Since the inclusion induces an isomorphism between $\pi_1 M_i$ and $\pi_1 Y_i$ it follows that each component of $\pi^{-1}M_i$ is a copy of the universal cover, $\tilde{M}_i$, of M_i. Furthermore every covering translate of $\tilde{Y}_i$ contains exactly one covering translate of $\tilde{M}_i$. The following claim implies that the developing map $dev : \tilde{M} \to \mathbb{H}^n$ is injective. The theorem then follows from proposition 2.6.

Claim 2. Suppose that P_0, P_1 are two points in $\tilde{M}$. Then there is a hyperbolic geodesic, γ, in the interior of $\tilde{Y}$ connecting P_0 and P_1.

Proof of Claim 2. By condition (ii) M_i is convex, so there is a 2-thickening $M_i^+ = Th_2(M_i)$. Thus each component, $\tilde{M}_i$, of $\pi^{-1}(M_i)$ has a 2-thickening $\tilde{M}_i^+$, which, by condition (iv), is contained in $\tilde{Y}$. The covering translates of $\tilde{M}_1$ are pairwise disjoint (and similarly for $\tilde{M}_2$) however the covering translates of $\tilde{M}_1^+$ need not always be disjoint. For example this will happen if M_1 contains a rank-1 cusp which is contained in a rank-2 cusp of M. Thus the natural isometric immersion of M_i^+ into Y is not always injective; the thickening may bump into itself inside $M_1 \cap M_2$. However the no bumping condition implies that every point of intersection of two different translates of $\tilde{M}_1^+$ is contained in some $\tilde{M}_2$ (and similarly with the roles of M_1 and M_2 reversed).

Although M might not be convex we define

$$\tilde{M}^+ = \{ \ x \in \tilde{Y} \ : \ d(x, \tilde{M}) \leq 2 \ \}.$$

This is the union of covering translates of $\tilde{M}_1^+$ and $\tilde{M}_2^+$. We are assuming that M_1 and M_2 are complete metric spaces thus $\tilde{M}_1^+, \tilde{M}_2^+$ and $\tilde{M}^+$ are all metrically complete.

Claim 3. Suppose $\tilde{M}_i$ is any component of $\pi^{-1}(M_i)$. If $\tilde{M}_1^+ \cap \tilde{M}_2^+ \neq \phi$ then $\tilde{M}_1 \cap \tilde{M}_2 \neq \phi$.

Proof of Claim 3. Choose a point x in $\tilde{M}_1^+ \cap \tilde{M}_2^+$. Let $\tilde{Y}_i$ be the component of $\pi^{-1}(Y_i)$ which contains $\tilde{M}_i$. Since $\tilde{M}_i$ is complete there is a point $a_i \in \tilde{M}_i$ which minimizes $d(a_i, x)$. It follows that $\pi(a_i) \in M_i \setminus int(M_i)$. By definition of $\tilde{M}_i^+$ we have $d(a_i, x) \leq 2$, hence $d(a_1, a_2) \leq 4$.

The first case is that $\pi(a_1) \notin int(M_2)$ hence $\pi(a_1) \in M_1 \setminus int(M_1 \cap M_2)$. Condition (v) implies $N_\kappa(a_1) \subset \tilde{Y}_1$. Since $\kappa \geq 4$ it follows that $a_2 \in \tilde{Y}_1$ thus $a_2 \in \tilde{Y}_1 \cap \tilde{Y}_2 \neq \phi$. Let C be the component of $\tilde{Y}_1 \cap \tilde{Y}_2$ which contains a_2. By condition (vi) there is a point of $M_1 \cap M_2$ in $\pi(C)$. Thus $\pi^{-1}(M_1 \cap M_2)$ contains a point, y, in C. However $\pi^{-1}(M_i) \cap \tilde{Y}_i = \tilde{M}_i$ thus $y \in \tilde{M}_1 \cap \tilde{M}_2 \neq \phi$.

The remaining case is that $\pi(a_1) \in int(M_2)$. Since $\pi(a_1) \notin int(M_1)$ it follows that $\pi(a_1) \in M_2 \setminus int(M_1 \cap M_2)$. Let $\tilde{Y}_2'$ be the component of $\pi^{-1}(Y_2)$ which contains a_1. Condition (v) implies $N_\kappa(a_1) \subset \tilde{Y}_2'$. Since $\kappa \geq 4$ it follows that $a_2 \in \tilde{Y}_2'$ thus $\tilde{Y}_2' = \tilde{Y}_2$, and $a_1 \in \tilde{Y}_1 \cap \tilde{Y}_2 \neq \phi$. Let C be the component of $\tilde{Y}_1 \cap \tilde{Y}_2$ which contains a_1. By condition (vi) there is a point of $M_1 \cap M_2$ in $\pi(C)$. The rest of the argument is the same as the first case. This proves claim 3. □

Since $\tilde{M}^+$ is a complete metric space there is a length minimizing path $\gamma : [0,1] \to \tilde{M}^+$ between the points P_0 and P_1.

Claim 4. γ is a geodesic except, possibly, at points in $\partial\tilde{M}_1^+ \cap \partial\tilde{M}_2^+$.

Proof of Claim 4. Consider a point p in the interior of γ. If p is in the interior of $\tilde{M}^+$ then, since γ is length minimizing, γ is a geodesic in a neighborhood of p. Now suppose that p is not in any translate of $\tilde{M}_2^+$. Thus p is in some $\tilde{M}_1^+$. Every point of intersection between distinct copies of $\tilde{M}_1^+$ is contained in some $\tilde{M}_2$. Thus there is an open arc, α, in γ that contains p and α is contained in a unique copy of $\tilde{M}_1^+$. Since $\tilde{M}_1^+$ is convex, and α is length minimizing, it follows that α is a geodesic. A similar conclusion holds if p is not in any translate of $\tilde{M}_1^+$. Thus if γ is not a geodesic near p, then p is in the boundary of $\tilde{M}^+$, and p is also contained in copies of $\tilde{M}_1^+$ and $\tilde{M}_2^+$. Hence it is in the boundaries of these copies. This proves claim 4. □

We will call a point p on γ a *corner* of γ if γ is not a hyperbolic geodesic at p. The following claim proves claim 2 and thus the theorem.

Claim 5. γ has no corners.

Proof of Claim 5. If p is a corner of γ then (by choosing notation for the covering translates) we may assume $p \in \partial\tilde{M}_1^+ \cap \partial\tilde{M}_2^+$. It follows that $\tilde{M}_1 \cap \tilde{M}_2 \neq \phi$. This is because $\pi(p) \notin M_1$ but $\pi(p)$ is a distance of 2 from some point $x \in M_2$ so by condition (v) $\pi(p) \in Y_2$. Similarly $\pi(p) \in Y_1$.

Let $\delta = \overline{pw} \subset \tilde{M}_2^+$ denote the maximal subarc of γ in $\tilde{M}_2{}^+$ which contains p.

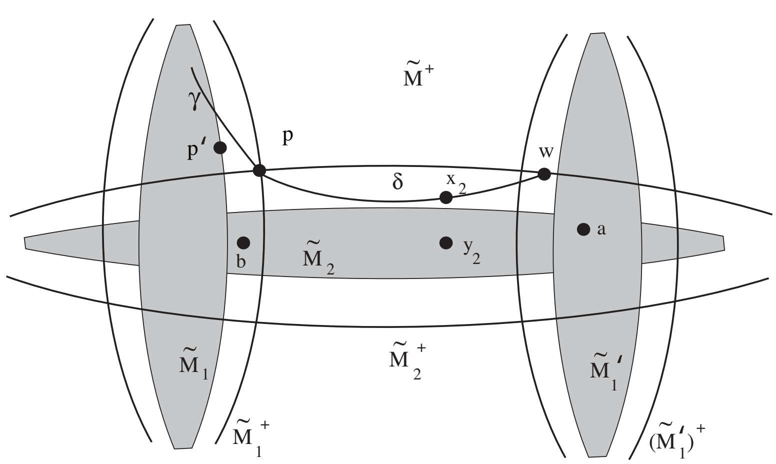

Figure 2. *The developing image*

Claim 6. Either $w \in \tilde{M}_2$ or $length(\delta) \geq \kappa - 2$.

Proof of Claim 6. First we consider the case that w is an endpoint of γ and in addition $\pi(w) \notin M_1$. Then $\pi(w) \in M_2$ thus w is in some translate, $\tilde{M}_2'$, of $\tilde{M}_2$. By definition of δ we have that w is in $\tilde{M}_2^+$. Thus $\tilde{M}_2'$ intersects $\tilde{M}_2^+$ in the point w. Since w is not contained in any translate of $\tilde{M}_1$, condition (v) implies that $\tilde{M}_2' = \tilde{M}_2$. Thus w is in $\tilde{M}_2$ and we are done.

Otherwise w is in some translate, $(\tilde{M}_1')^+$, of $\tilde{M}_1^+$. This is because either w is an endpoint of γ and $\pi(w) \in M_1$ or else, by maximallity of δ, the point w is in the boundary of $\tilde{M}_2^+$ and thus in some $(\tilde{M}_1')^+$ (since distinct translates of $\tilde{M}_2^+$ can only intersect in a translate of $\tilde{M}_1$.) Let $\tilde{Y}_1, \tilde{Y}_1'$ be the components of $\pi^{-1}(Y_1)$ which contain $\tilde{M}_1$ and $\tilde{M}_1'$ respectively.

We first consider the case that $\tilde{M}_1 \neq \tilde{M}_1'$. Let p' be the point in $\tilde{M}_1$ closest to p thus $d(p, p') = 2$. If p' was in the interior of $\tilde{M}_2$ then $d(p, \tilde{M}_2) < d(p, p') = 2$ which contradicts that $d(p, \tilde{M}_2) = 2$. Hence $p' \in closure(\tilde{M}_1) \setminus int(\tilde{M}_1 \cap \tilde{M}_2)$. Since $\tilde{Y}_1$ and $\tilde{Y}_1'$ are not equal they are disjoint. The geodesic δ has one endpoint $p \in \tilde{Y}_1$ and the other endpoint $w \in \tilde{Y}_1'$ thus w is not in $\tilde{Y}_1$. From condition (v) it follows that $d(p', w) \geq \kappa$. Hence $length(\delta) = d(p, w) \geq \kappa - 2$.

The remaining case is that $\tilde{M}_1 = \tilde{M}_1'$ in which case $\tilde{Y}_1 = \tilde{Y}_1'$. If δ is contained in $\tilde{Y}_1$ then δ is a geodesic in the convex set $\tilde{Y}_1$ with both endpoints in the convex subset $\tilde{M}_1^+$. But this implies that δ is contained in $\tilde{M}_1^+$. This in turn means that there is an open interval in γ which contains p, and is contained in $\tilde{M}_1^+$, and this open interval is therefore a geodesic. This contradicts that p is a corner. Hence δ contains a point outside $\tilde{Y}_1$ and then as before we obtain $length(\delta) \geq \kappa - 2$. This proves claim 6. ☐

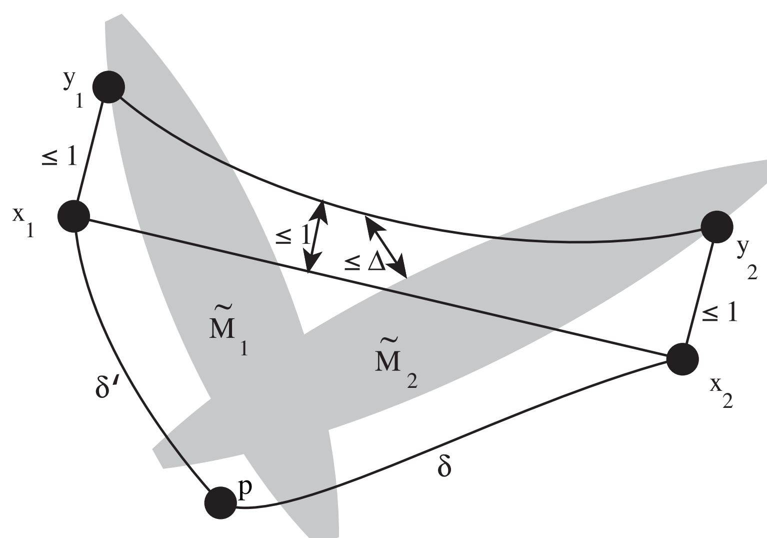

Figure 3. *The path stays close to the corners*

Proof of claim 5, resumed. The geodesic segment δ has both endpoints, w, p within a distance 2 of $\tilde{M}_2$ so we may choose points $a, b \in \tilde{M}_2$ with $d(a, w) \leq 2$ and $d(b, p) \leq 2$. Since $\tilde{M}_2$ is convex there is a geodesic, $\overline{ab}$, in $\tilde{M}_2$. If $length(\delta) \geq \kappa - 2 = 4$ it follows from lemma 2.10 that there is a point $x_2 \in \delta$ and a point $y_2 \in \overline{ab}$ with $d(x_2, y_2) \leq 1$ hence $d(x_2, \tilde{M}_2) \leq d(x_2, y_2) \leq 1$. Otherwise $w \in \tilde{M}_2$ and we choose $x_2 = y_2 = w$ and then $d(x_2, \tilde{M}_2) = d(x_2, y_2) = 0 \leq 1$.

The same argument shows that, if δ' is the maximal segment of γ in $\tilde{M}_1^+$ which contains p, then there are points $x_1 \in \delta'$ and $y_1 \in \tilde{M}_1$ with $d(x_1, \tilde{M}_1) \leq d(x_1, y_1) \leq 1$. Let $\tilde{Y}_i$ be the component of $\pi^{-1}Y_i$ which contains $\tilde{M}_i$. Using the convexity of $\tilde{Y}_1$ and $\tilde{Y}_2$ it is easy to see that the developing map restricted to $\tilde{Y}_1 \cup \tilde{Y}_2$ is an embedding. Thus we may regard $\tilde{Y}_1 \cup \tilde{Y}_2$ as a subset of $\mathbb{H}^n$. By convexity of the distance function (proposition 2.11) it follows that every point on the geodesic $\overline{x_1x_2}$ is less than a distance of 1 from some point on the geodesic $\overline{y_1y_2}$. Since $p \in (\tilde{M}_1)^+ \cap (\tilde{M}_2)^+$ it follows from claim 3 that $\tilde{M}_1 \cap \tilde{M}_2 \neq \phi$. Now $\tilde{M}_1, \tilde{M}_2$ are both convex and have non-empty intersection. Since $y_1 \in \tilde{M}_1$ and $y_2 \in \tilde{M}_2$ it follows from lemma 2.12) that every point on $\overline{y_1y_2}$ is within a distance Δ of $\tilde{M}_1 \cup \tilde{M}_2$. Thus every point on $\overline{x_1x_2}$ is less than a distance $1 + \Delta$ of $\tilde{M}_1 \cup \tilde{M}_2$.

The segment, $\delta' \cup \delta$, of γ between x_1 and x_2 is length minimizing among all paths with the same endpoints in $\tilde{M}_1^+ \cup \tilde{M}_2^+$. Since $1 + \Delta < 2$ the hyperbolic geodesic $\overline{x_1x_2}$ is contained in $\tilde{M}_1^+ \cup \tilde{M}_2^+$ and is the unique length minimizing path in this set with these endpoints. Thus $\delta' \cup \delta = \overline{x_1x_2}$ but this contradicts that p is a corner. This proves claim 5, and thus completes the proofs of claim 2 and the theorem. □

Lemma 2.10. *Suppose that Q is a (not necessarily planar) quadrilateral in hyperbolic space with corners a, b, c, d and geodesic sides $\overline{ab}, \overline{bc}, \overline{cd}, \overline{da}$. Suppose that $|\overline{ad}| \leq 2$, $|\overline{bc}| \leq 2$ and $|\overline{ab}| \geq 4$. Then there are points $w \in \overline{ab}$ and $z \in \overline{cd}$ such that $d(w, z) \leq 1$.*

Proof. The worst case is the symmetric one in the plane, see figure 4. A calculation then gives the result. □

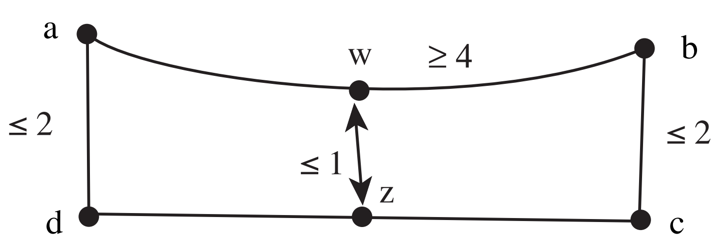

Figure 4. *The worst case.*

Proposition 2.11 (distance function is convex [**39**], p91, 2.5.8). *The distance function $d(x, y)$, considered as a map $d : \mathbb{H}^n \times \mathbb{H}^n \to \mathbb{R}$, is convex. The composition $d \circ \gamma$ is strictly convex for any geodesic γ in $\mathbb{H}^n \times \mathbb{H}^n$ whose projections to the two factors are distinct.*

Lemma 2.12 (convex unions). *Suppose A and B are convex subsets of hyperbolic space $\mathbb{H}^n$ which have non-empty intersection. Then the convex hull $CH(A \cup B)$ is contained in the Δ-neighborhood of $A \cup B$. Here $\Delta = log((3 + \sqrt{5})/2) < 1$ is the thin-triangles constant of $\mathbb{H}^2$.*

Proof. Let X be the union of all the geodesic segments, $[a, b]$ with endpoints $a \in A$ and $b \in B$. We claim that X is the convex hull of $A \cup B$. Clearly X is contained in this convex hull. It suffices to show that X is convex.

Suppose that $p_1 \in [a_1, b_1]$ and $p_2 \in [a_2, b_2]$ are two points in X. Then we need to show that every point, q, on the geodesic segment from p_1 to p_2 is also in X. Since q is in the convex hull of the four points a_1, a_2, b_1, b_2 there is a geodesic $[a, b]$ with endpoints $a \in [a_1, a_2] \subset A$ and $b \in [b_1, b_2] \subset B$ which contains q. Since $[a, b]$ is in X it follows $q \in X$ thus X is convex.

Given $q \in X$ there is $a \in A$ and $b \in B$ such that $q \in [a, b]$. Choose a point $p \in A \cap B$ and consider the geodesic triangle with sides $[a, p], [p, b], [a, b]$. By convexity $[a, p] \subset A$ and $[p, b] \subset B$. The point q on $[a, b]$ is within a distance Δ of some point in $[a, p] \cup [p, b]$ and is thus within a distance Δ of $A \cup B$. □

There is a version of the convex combination theorem which applies to the case of a convex manifold, M, which is locally isometrically immersed by a map which identifies two disjoint submanifolds of M and is otherwise injective:

Corollary 2.13 (HNN convex combinations). *Suppose that :*
(i) M is a convex connected hyperbolic n-manifold and Y is a convex thickening of M.
(ii) X is a hyperbolic n-manifold and $f : Y \to X$ is a locally isometric immersion with image a submanifold $f(Y)$.
(iii) $Y = Z_0 \cup Z_1 \cup Z_2$ is a decomposition into disjoint submanifolds.
(iv) $f|Z_i$ is injective for each i.
(v) $f(Z_1) = f(Z_2)$ is disjoint from $f(Z_0)$.
(6) Define $N_i = M \cap Z_i$ then $f(N_1) = f(N_2)$.
(7) $f(Y)$ contains a κ-neighborhood of $f(M)$.
(8) Y contains a κ-neighborhood of $\overline{N}_0$.
(9) Each component of $f(Z_1)$ contains a component of $f(N_1)$.
Then $f(M)$ has a convex thickening. In addition if Y has finite volume then this thickening is a geometrically finite manifold.

Proof. Set $Q = f(Y)$. There is a unique 2-fold cover $\tilde{Q} \to Q$ with the property that f has two distinct lifts $f_1, f_2 : Y \hookrightarrow \tilde{Q}$ each of which is injective. Then $\tilde{Q} = Y_1 \cup Y_2$ where $Y_i = f_i(Y)$. Define $M_i = f_i(M)$. The hypotheses of the convex combination theorem are satisfied by M_1, M_2, Y_1, Y_2 and $Y_1 \cup Y_2$. It follows that $M_1 \cup M_2$ has a convex thickening. Since $f(M)$ is the quotient of this by an isometric involution it follows that $f(M)$ also has a convex thickening. □

3. *Cusps in Convex Hyperbolic Manifolds.*

In this section we study the geometry of cusps in convex, geometrically-finite, hyperbolic manifolds. The main result we need is 3.5(e) which states that every thin cusp of a convex manifold is contained in a product cusp of some relative thickening. This is used in the proof of the virtual simple gluing theorem 4.3.

If $\Gamma \subset Isom(\mathbb{H}^n)$ is discrete and torsion-free then $M = \mathbb{H}^n/\Gamma$ is a geodesically-complete hyperbolic manifold. If the limit set of Γ contains more than 1 point then, given $\epsilon \geq 0$, we define C^ϵ to be the closed ϵ-neighborhood in $\mathbb{H}^n$ of the convex hull of the limit set of Γ. The *ϵ-thickened convex core of M* is $Core^\epsilon(M) = C^\epsilon/\Gamma$ and when $\epsilon = 0$ this is called the *convex core of M* and we write it as $Core(M)$. The convex core is a convex hyperbolic manifold and a complete metric space.

The manifold M and the group Γ are *geometrically finite* if for some (hence every) $\epsilon > 0$ the volume of $Core^\epsilon(M)$ is finite. In dimension three this is equivalent to M having a finite sided polyhedral fundamental domain. Bowditch gave several equivalent formulations of geometrical finiteness in [**5**].

Definition. A *cusp* is a convex hyperbolic manifold with non-trivial parabolic holonomy. The *rank* of a cusp is the largest rank of a free-abelian subgroup of its fundamental group.

We need to distinguish three kinds of cusp: *complete, thin* and *product*. Suppose that D is a closed horoball in $\mathbb{H}^n$ and Γ is a non-trivial, discrete, torsion-free subgroup of $Isom(\mathbb{H}^n)$ which stabilizes D. Then the quotient $C = D/\Gamma$ is called a *complete cusp*. It is a convex hyperbolic manifold with boundary. The induced metric on the boundary of the horoball D is Euclidean, and Γ acts by Euclidean isometries on ∂D, so the boundary of C is isometric to the Euclidean $(n-1)$-manifold $\mathbb{E}^{n-1}/\Gamma$. By a theorem of Bieberbach such a Euclidean manifold is a flat vector bundle over a closed Euclidean manifold called the *soul*.

Suppose that $N = \mathbb{H}^n/\pi_1 N$ is a geodesically-complete hyperbolic n-manifold. A *cusp in* N is a submanifold C^+ of N which is isometric to a complete cusp. Now suppose that M is a convex hyperbolic n-manifold and $Th_\infty(M) = \mathbb{H}^n/\pi_1 M$ is the corresponding geodesically-complete manifold. Let C^+ be a complete cusp in $Th_\infty(M)$. A *cusp in* M is the intersection $C = C^+ \cap M$. We say that C^+ is the *complete cusp corresponding to* C. Clearly $\pi_1 C \cong \pi_1 C^+$. For simplicity, in this section we will also assume that M is metrically complete. In addition we assume that if M has parabolic holonomy, then a submanifold of M that is designated as a cusp is never all of M. The *cusp boundary* of the cusp C in M is denoted by $\partial_c C$ and equals $(M \setminus int(C)) \cap C = C \cap \partial C^+$ and is a submanifold of ∂C^+.

The horoball D has a codimension-1 foliation by horospheres H_t for $t \geq 0$, such that $\partial D = H_0$ and the distance between H_s and H_t is $|s - t|$. This foliation is preserved by Γ. Thus every cusp has a codimension -1 foliation whose leaves are called *horomanifolds* which are covered by submanifolds of horospheres. The induced metric on a horomanifold is Euclidean.

It follows from [**5**] that if M is a geometrically-finite hyperbolic manifold, $\epsilon > 0$, and if $\mathcal{C}$ is a maximal collection of pairwise disjoint cusps, then $Core^\epsilon(M) \setminus int(\mathcal{C})$ is compact.

Lemma 3.1 (thinning cusps). *Suppose M is a convex hyperbolic manifold and C is a cusp of M. Then $CH(\overline{M \setminus C}) = (\overline{M \setminus C}) \cup CH(\partial_c C)$.*

Proof. Since $\partial_c C \subset \overline{M \setminus C}$ it follows that $(\overline{M \setminus C}) \cup CH(\partial_c C) \subset CH(\overline{M \setminus C})$. Clearly $X \equiv (\overline{M \setminus C}) \cup CH(\partial_c C)$ is closed, thus it only remains to show it is also convex.

Let $\pi : \tilde{M} \to M$ be the universal cover. Let $\tilde{X} = \pi^{-1}(X)$. Given two points in $\tilde{X}$ the convexity of M implies there is a geodesic, γ, in $\tilde{M}$ connecting them. This geodesic is made up of segments; each segment is either contained in $\pi^{-1}(\overline{M \setminus C})$ or in a component of $\pi^{-1}(C)$.

Consider a segment, δ, which is a component of $\gamma \cap \pi^{-1}(C)$. Each endpoint of δ is either in $\pi^{-1}(\partial_c C)$ or is an endpoint of γ and thus in $\pi^{-1}(CH(\partial_c C))$. Each component of $\pi^{-1}(C)$ intersects (and contains) a unique component of $\pi^{-1}(CH(\partial_c C))$. Thus both endpoints of δ are contained in the same component, A, of $\pi^{-1}(CH(\partial_c C))$. Since A is convex δ is contained in A. Thus γ is contained in $\tilde{X}$. It follows from the definition of convex manifold that X is convex.□

Definition. A *product cusp* is a cusp which has a 1-dimensional foliation by geodesic rays orthogonal to the foliation by horomanifolds. Each geodesic ray starts on the cusp boundary and has infinite length. Recall that the definition of cusp includes that the cusp is convex.

Proposition 3.2 (product cusps). *Suppose C is a cusp. The following are equivalent:*

(i) *C is a product cusp.*

(ii) *$\partial_c C$ is a convex Euclidean manifold, possibly with non-empty boundary, and $C = CH(\partial_c C)$.*

(iii) *If we isometrically identify $\mathbb{H}^n$ with the upper half space $x_n > 0$ of $\mathbb{R}^n$ equipped with the metric ds/x_n, then the universal cover of C is isometric to*

$$\Omega \;=\; \{\ (x_1, \ldots, x_n) \in \mathbb{R}^n \;:\; (x_1, x_2, \ldots, x_{n-1}, 1) \in S \textit{ and } x_n \geq 1\ \}$$

where S is a convex $(n-1)$-submanifold of the horosphere $x_n = 1$.

Furthermore, the leaves of the 1*-dimensional foliation of* C *are covered by vertical line segments in* Ω *which are parallel to the* x_n*-axis.*

Proof. The proof follows easily from the observation that $\partial_c C$ is a convex Euclidean manifold iff it's universal cover, S, is a convex $(n-1)$-submanifold of the horosphere $x_n = 1$ iff Ω is a convex subset of $\mathbb{H}^n$. □

Definition. A cusp, C, is *thin* if there is a constant D such that for all $t \geq 0$ the diameter of the horomanifold $H_t/\pi_1 C$ in the cusp is less than $D\exp(-t)$.

Remark. If C is a thin cusp then it has finite volume. The converse is not true in general. For example a 3-dimensional rank-1 cusp for which the diameter of every horomanifold is 1 has finite volume.

Lemma 3.3 (thin cusps). *Suppose* C *is a cusp.*
(i) *Every finite cover of* C *is a thin cusp.*
(ii) *If* C *is a complete cusp then it is thin iff* $rank(C) = dim(C) - 1$.
(iii) *If* $\partial_c C$ *has bounded diameter then* $CH(\partial_c C)$ *is a thin cusp.*
(iv) *If* C *is a product cusp with compact cusp boundary then* C *is thin.*
(v) *If* C *is thin then it has a thickening which is a thin product cusp.*

Proof. (i) is obvious. (ii) follows the observation that the rank of a cusp equals the dimension of the soul of the cusp boundary. Thus $rank(C) = dim(C)-1$ is equivalent to the condition that the cusp boundary is a closed manifold. For (iii), let C^+ be the complete cusp corresponding to C. Then $\partial_c C$ is a bounded subset of the Euclidean manifold $W = \partial C^+$. Now W is a vector bundle over a closed Euclidean manifold. Given $r > 0$ let W_r be the subset of vectors of length at most r. Then W_r is a compact, convex, submanifold of W. Hence $CH(W_r)$ is a product cusp. Now $\partial_c C$ has bounded diameter thus for some $r > 0$ we have $\partial C \subset W_r$. Thus $CH(\partial_c C) \subset CH(W_r)$. Clearly $CH(W_r)$ is thin, therefore $CH(\partial_c C)$ is also thin. This proves (iii), and (iv) follows from (iii).

For (v), let E be the subset of C^+ consisting of all geodesic rays which intersect C and are orthogonal to the horomanifolds. From the description of a cusp in part (iii) of 3.2 it is easy to see that, since C is convex, E is also convex. It is also easy to see that since C is thin (with constant D say) $\partial_c E = E \cap \partial_c C^+$ has diameter at most $2D$. Thus $\partial_c E$ is contained in W_r for r sufficiently large. Now E is contained in the product cusp $CH(W_r)$ and, by part (iii), $CH(W_r)$ is thin. □

Proposition 3.4 (GF cores have thin cusps). *Suppose that* $N = \mathbb{H}^n/\Gamma$ *is a geometrically-finite hyperbolic manifold. Then* $M = Core(N)$ *has thin cusps.*

Proof. If C is a cusp of M then $\partial_c C$ has bounded diameter. This is because otherwise the ϵ-neighborhood of $\partial_c C$ in N has infinite volume. But N is geometrically finite so $C^\epsilon(M)$ has finite volume and contains the ϵ-neighborhood of $\partial_c C$. Since $M = Core(N)$ it follows that $CH(M \setminus C) = M$ and 3.1 implies $C = CH(\partial_c C)$. The result follows from 3.3. □

Although we won't use this fact, in a geometrically-finite hyperbolic 3-manifold every cusp contains a (possibly smaller) cusp which is a product cusp. For rank-2 cusps this is obvious. For rank-1 cusps, see [**40**].

Definition. Suppose that M is a convex hyperbolic manifold and that $C_1, \ldots, C_n$ is a collection of pairwise disjoint cusps in $Th_\infty(M)$. Given $K \geq 0$ the *K-thickening of M relative*

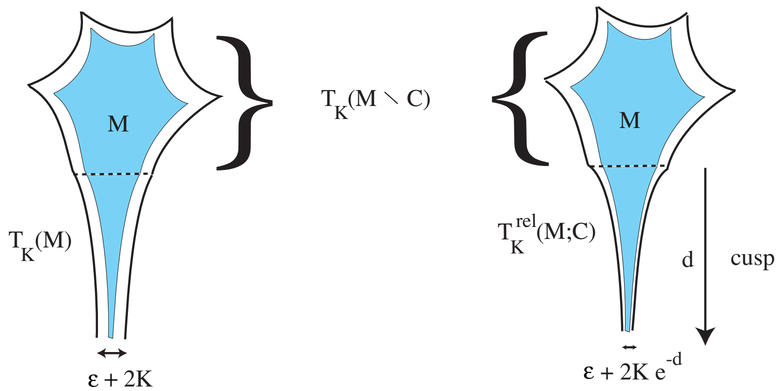

FIGURE 5. *Relative Thickening.*

to $\mathcal{C} = \bigcup_i C_i$ is

$$Th_K^{rel}(M;\mathcal{C}) \;=\; CH(Th_K(M) \setminus int(\mathcal{C})).$$

Since M is convex, $Th_K(M)$ is also convex and, as we shall see below, the difference between it and $Th_K^{rel}(M;\mathcal{C})$ is to replace those cusps in $Th_K(M)$ contained in $\mathcal{C}$ by thin cusps. Usually $\mathcal{C}$ is a maximal collection. In this case, although relative thickening depends on a particular choice of a maximal family of cusps, usually the choice is unimportant. We will therefore use $Th_K^{rel}(M)$ to denote the result of a relative thickening with respect to some maximal set of cusps .

EXAMPLE. Suppose that F is a complete hyperbolic punctured torus. Isometrically embed $\mathbb{H}^2$ into $\mathbb{H}^3$ then the holonomy of F gives a Kleinian group $\Gamma \cong \pi_1 F$. We can regard F as a degenerate hyperbolic 3-manifold of zero thickness. The quotient by Γ of the K-neighborhood of $\mathbb{H}^2$ in $\mathbb{H}^3$ is the K-thickening $M = Th_K(F)$. It is a convex hyperbolic 3-manifold with a rank-1 cusp. The thickness of the cusp everywhere is $2K$. This means that every point in F is contained a geodesic segment in M of length $2K$ which is orthogonal to F. In particular this is an example of a convex 3-manifold of finite volume which has a rank-1 cusp that is not a product cusp. The relative K-thickening of F is the subset of $Th_K(F)$ obtained by replacing the rank-1 cusp of M by a product cusp whose thickness decreases exponentially.

PROPOSITION 3.5 (relative thickenings contain product cusps).
Suppose that M is a metrically-complete convex hyperbolic manifold, and that $\mathcal{C} = \bigcup_i C_i$ is a maximal set of pairwise disjoint cusps in $Th_\infty(M)$ and assume that all these cusps are thin. Then:
(a) For K sufficiently large, $Th_K^{rel}(M;\mathcal{C})$ is a convex thickening of M.
(b) If $M = Core(M)$ is geometrically finite then $Th_K^{rel}(M;\mathcal{C})$ has finite volume.
(c) $Th_K^{rel}(M;\mathcal{C})$contains a K-neighborhood of $M \setminus int(\mathcal{C})$.
(d) $Th_K^{rel}(M;\mathcal{C})$ has thin cusps.
(e) For K sufficiently large each cusp, $M \cap C_i$, of M is contained in a product cusp which is a subset of $Th_K^{rel}(M;\mathcal{C})$.

Proof. Part (a) follows from (e). For (b) observe that $Th_K^{rel}(M;\mathcal{C})$ has finite volume because it is a subset of the convex manifold $Th_K(M)$, and the latter has finite volume because M is geometrically finite. For (d) observe that by 3.1

$$Th_K^{rel}(M;\mathcal{C}) = (Th_K(M) \setminus int(\mathcal{C})) \cup \bigcup_i CH(M \cap \partial_c C_i).$$

Conclusion (d) now follows from 3.3.

For (c) consider the metric ball, B, of radius K in $Th_\infty(M)$ centered on $x \in M \setminus int(\mathcal{C})$. Then $B \subset Th_K(M)$. Hence $B \setminus Th_K^{rel}(M) \subset \mathcal{C}$. Suppose A is a component of $B \cap \mathcal{C}$. We will show that $A \subset Th_K^{rel}(M;\mathcal{C})$. It then follows that $B \subset Th_K^{rel}(M;\mathcal{C})$ which implies (c).

Identify the universal cover of $Th_\infty(M)$ with $\mathbb{H}^n$. Let $\tilde{x}$ be a point in $\mathbb{H}^n$ which covers x and let $\tilde{B}$ be the metric ball in $\mathbb{H}^n$ centered at $\tilde{x}$ and radius K. Thus $\tilde{B}$ projects onto B. Let C be the component of $\mathcal{C}$ which contains A and identify $\mathbb{H}^n$ with the upper half-space model so that the horoball $x_n \geq 1$ projects onto C. Then $\tilde{B} \cap \{ x_n = 1 \}$ is a metric ball in the horosphere $x_n = 1$. It projects onto $A \cap \partial_c C$. The point, p, at infinity ($x_n = \infty$) in the upper half space model is a parabolic fixed point for M and so is in the limit set of $\tilde{M}$. Let Y be the subset of $\mathbb{H}^n$ corresponding to the universal cover of $Th_K^{rel}(M;\mathcal{C})$. It follows that p is in the limit set of Y. Now Y is convex and contains contains $\tilde{B} \cap \{ x_n = 1 \}$ and limits on p. Hence Y contains the solid cylinder, W, of points lying vertically above $\tilde{B} \cap \{ x_n = 1 \}$. Since x is not in the interior of C it follows that $\tilde{x}$ is not in $\{ x_n > 1 \}$. Hence W contains $Z = \tilde{B} \cap \{ x_n \geq 1 \}$, and $Z \subset Y$. The projection of Z to $Th_K^{rel}(M;\mathcal{C})$ is A. This proves (c).

For (e) consider a cusp C in $\mathcal{C}$. Since C is thin 3.3 implies it has a thickening which is a thin product cusp E. For K sufficiently large $\partial_c E \subset Th_K^{rel}(M;\mathcal{C})$. The latter is convex so $E = CH(\partial_c E) \subset Th_K^{rel}(M;\mathcal{C})$ and this proves (e). □

We do not know if every geometrically finite hyperbolic n-manifold has a convex thickening with the property that every cusp disjoint from a sufficiently large compact set is a thin product cusp. Part (e) of the above enables us to bypass this issue in the proof of the virtual simple gluing theorem 4.3.

4. *Induced gluing*

The main result of this section is the virtual simple gluing theorem 4.3. Suppose that we have two locally-isometric immersions of hyperbolic 3-manifolds equipped with basepoints into a hyperbolic 3-manifold M

$$f : (A, a_0) \to (M, m_0) \qquad\qquad g : (B, b_0) \to (M, m_0).$$

We would like to use this information to glue A and B together. For example if both immersions are injective then we might identify A with $f(A)$ and B with $g(B)$. The union $f(A) \cup g(B) \subset M$ may then be regarded as a quotient space of the disjoint union of A and B and it has a hyperbolic metric.

However we will be interested in situations when the immersions are not injective. Furthermore, even when A and B are submanifolds of M, we want to do make the fewest identifications subject to the requirements that the basepoints in A and B are identified and that the identification space is a hyperbolic 3-manifold. Thus if $A \cap B$ is not connected we wish to only identify A and B along the component, C, of $A \cap B$ containing the basepoint m_0. In certain circumstances the fundamental group of the identification space will be a free product of the fundamental groups of A and B amalgamated along a subgroup corresponding to the fundamental group of C.

We give below a very general way of forming an identification space. Even when the resulting identification space is a hyperbolic manifold, it will usually not be convex and might not have a convex thickening.

DEFINITION. Suppose that $f:(X,x_0)\to(Z,z_0)$ and $g:(Y,y_0)\to(Z,z_0)$ are continuous maps of pointed spaces. Define a relation R on the disjoint union, $X\coprod Y$, of X and Y as follows. If $x\in X$ and $y\in Y$ then xRy iff there are paths $\alpha:I\to X$ and $\beta:I\to Y$ such that $\alpha(0)=x_0,\ \alpha(i)=x,\ \beta(0)=y_0,\ \beta(i)=y$ and $f\circ\alpha=g\circ\beta$. We now define a topological space called the *induced gluing*, denoted $S(f,g)$, to be the quotient space $X\coprod Y/\equiv$ obtained by taking the equivalence relation $\equiv$ which is generated by R. It is clear that if two points are identified then they have the same image in Z. Let $\pi_X:X\to S(f,g)$ and $\pi_Y:Y\to S(f,g)$ denote the natural projections. We say that the induced gluing is a *simple gluing* if both π_X and π_Y are injective.

Example (i) If f,g are both embeddings define Z_0 to be the path component of $f(X)\cap g(Y)$ containing the basepoint z_0. Then $S(f,g)$ is obtained from $X\coprod Y$ by identifying $f^{-1}(Z_0)$ with $g^{-1}(Z_0)$ using the homeomorphism $g^{-1}\circ f:f^{-1}(Z_0)\to g^{-1}(Z_0)$. It is clear that this is a simple gluing.

Example (ii) Suppose $f,g:S^1\times[0,2]\to S^1\times S^1$ are given by

$$f(\omega,t)=(\ \omega\ ,\ exp(2\pi it)\ )\qquad\text{and}\qquad \mathrm{g}(\omega,\mathrm{t})=(\ \exp(2\pi\mathrm{it})\ ,\ \omega\ ).$$

Then $S(f,g)$ can be naturally identified with the codomain $S^1\times S^1$. In this case the gluing is not simple, in fact both projections are surjective.

Example (iii) Suppose now that $p:S^1\times[0,2]\to S^1\times[0,2]$ is the 3-fold cyclic cover and f,g are as in example (ii). Then $S(f\circ p,g\circ p)$ is homeomorphic to a torus minus an open disc. It is easy to see that this is a simple gluing. Furthermore this is a modification of example (ii) where the domains are replaced by certain finite covers. This phenomenon of taking a non-simple gluing and making a simple gluing by replacing the spaces by finite covers is generalized below.

Example (iv) Suppose that $X=S^1$ and $Y\cong Z\cong D^2$ and $f(z)=z^2$ and g is a homeomorphism. Then π_Y is a homeomorphism and π_X is a covering map onto ∂D. The gluing is not simple. Furthermore there are no finite covers of the domains which result in a simple gluing as in example (iii).

It is routine to check the following:

LEMMA 4.1 (induced gluing).
Suppose $f:(X,x_0)\to(Z,z_0)$ and $g:(Y,y_0)\to(Z,z_0)$ yield an induced gluing $S(f,g)$.
(a) The induced gluing is simple iff for every $x\in X$ there is at most one $y\in Y$ with xRy and vice versa.
(b) If the induced gluing is simple define subspaces $X_0\subset X,\ Y_0\subset Y$ by $x\in X_0$ and $y\in Y_0$ if xRy. Define $h:X_0\to Y_0$ by $h(x)=y$ if xRy. Then $S(f,g)$ is the quotient space obtained from $X\coprod Y$ by identifying X_0 with Y_0 using h.
(c) If X,Y,Z are smooth n-manifolds with boundary and f,g are immersions, and if the induced gluing is simple, and if $f|\partial X$ is transverse to $g|\partial Y$, then $S(f,g)$ is an n-manifold with boundary.
(d) There is a unique continuous induced map $h:S(f,g)\to Z$ such that $f=h\circ\pi_X$ and $g=h\circ\pi_Y$.
(e) If the induced gluing is simple then $\pi_X(X)\cap\pi_Y(Y)$ is a path connected subspace of $S(f,g)$.

We are concerned with induced gluings of convex hyperbolic manifolds. In this case the gluing of the manifolds is determined by the intersection of the images of their developing maps:

LEMMA 4.2 (geodesic gluing). *Suppose A,B,M are convex hyperbolic manifolds of the same dimension and $f:(A,a_0)\to(M,m_0)$ and $g:(B,b_0)\to(M,m_0)$ are locally isometric immersions. Let $p_A:\tilde{A}\to A$ and $p_B:\tilde{B}\to B$ and $p_M:\tilde{M}\to M$ be the universal covers.*

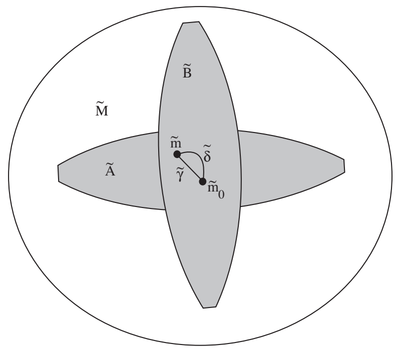

FIGURE 6. *Gluing Convex Manifolds.*

Choose lifts of the base points and maps $\tilde{f} : (\tilde{A}, \tilde{a}_0) \to (\tilde{M}, \tilde{m}_0)$ and $\tilde{g} : (\tilde{B}, \tilde{b}_0) \to (\tilde{M}, \tilde{m}_0)$ covering f and g. Suppose $a \in A$ and $b \in B$, then aRb iff there are $\tilde{a} \in \tilde{A}$ and $\tilde{b} \in \tilde{B}$ covering a and b respectively such that $\tilde{f}(\tilde{a}) = \tilde{g}(\tilde{b})$. Furthermore, if aRb then the paths α, β used in the definition of the R-relation can be chosen to be geodesics.

Proof. First suppose that aRb. Then there are paths α in A from a_0 to a and β in B from b_0 to b with $f \circ \alpha = g \circ \beta$. Let $\tilde{\alpha}, \tilde{\beta}$ be the lifts that start at the respective basepoints $\tilde{a}_0, \tilde{b}_0$ of $\tilde{A}, \tilde{B}$. Then $\tilde{\delta} = \tilde{f} \circ \tilde{\alpha}$ is a lift of $f \circ \alpha$ and $\tilde{g} \circ \tilde{\beta}$ is a lift of $g \circ \beta$ which both start at $\tilde{m}_0$ thus $\tilde{f} \circ \tilde{\alpha} = \tilde{g} \circ \tilde{\beta}$. Setting $\tilde{a} = \tilde{\alpha}(i)$ gives $p_A(\tilde{a}) = a$. Similarly $\tilde{b} = \tilde{\beta}(i)$ implies $p_B(\tilde{b}) = b$. Since $\tilde{f} \circ \tilde{\alpha} = \tilde{g} \circ \tilde{\beta}$ we get $\tilde{f}(\tilde{a}) = \tilde{g}(\tilde{b})$ completing the proof in this direction.

For the converse, given $\tilde{a}, \tilde{b}$ with $\tilde{f}(\tilde{a}) = \tilde{g}(\tilde{b})$, the point $\tilde{m} = \tilde{f}(\tilde{a})$ is in both $\tilde{f}(\tilde{A})$ and $\tilde{g}(\tilde{B})$. These are convex subsets of $\tilde{M}$ each of which may be identified with a convex subset of hyperbolic space; thus their intersection is a convex set, C, that contains both $\tilde{m}_0$ and $\tilde{m}$. By convexity there is a geodesic $\tilde{\gamma}$ in C with endpoints $\tilde{m}$ and $\tilde{m}_0$. Then $\tilde{\alpha} = \tilde{f}^{-1}\tilde{\gamma}$ and $\tilde{\beta} = \tilde{g}^{-1}\tilde{\gamma}$ are geodesics in $\tilde{A}$ and $\tilde{B}$ which project to geodesics $\alpha = p_A\tilde{\alpha} \subset A$ and $\beta = p_B\tilde{\beta} \subset B$. Set $a = \alpha(i)$ and $b = \beta(i)$ then since $\tilde{f} \circ \tilde{\alpha} = \tilde{g} \circ \tilde{\beta}$ it follows that $f \circ \alpha = g \circ \beta$ and hence aRb. □

We now generalize the passage from example (ii) to example (iii). Example (iv) shows that the convexity hypothesis is necessary.

THEOREM 4.3 (virtual simple gluing theorem). *Suppose that $M = \mathbb{H}^n/\pi_1 M$ is a geodesically-complete hyperbolic manifold. Suppose that A and B are geometrically-finite, convex, hyperbolic n-manifolds with finite volume and thin cusps. Suppose $f : (A, a_0) \to (M, m_0)$ and $g : (B, b_0) \to (M, m_0)$ are local isometries. Let $G = f_*\pi_1(A, a_0) \cap g_*\pi_1(B, b_0) < \pi_1(M, m_0)$ and suppose that $G_A = f_*^{-1}(G)$ is separable in $\pi_1(A, a_0)$ and $G_B = g_*^{-1}(G)$ is separable in $\pi_1(B, b_0)$.*

Then there are finite covers $p_A : (\tilde{A}, \tilde{a}_0) \to (A, a_0)$ and $p_B : (\tilde{B}, \tilde{b}_0) \to (B, b_0)$ and maps $\tilde{f} = f \circ p_A : (\tilde{A}, \tilde{a}_0) \to (M, m_0)$ and $\tilde{g} = g \circ p_B : (\tilde{B}, \tilde{b}_0) \to (M, m_0)$ such that $S(\tilde{f}, \tilde{g})$ is a simple gluing of $\tilde{A}$ and $\tilde{B}$ and so that $G_A \subset p_{A*}(\pi_1(\tilde{A}, \tilde{a}_0))$ and $G_B \subset p_{B*}(\pi_1(\tilde{B}, \tilde{b}_0))$.

Proof. First we give a sketch of the proof, the details follow. To understand this outline it may help to refer to the figure 7 below. In claim 1 we show that if A and B are compact there is a constant $L > 0$ such that if $a_1 \in A$ is R-related to $b_1 \in B$ then there are two geodesics, γ_1 in A going from a_0 to a_1, and δ_1 in B going from b_0 to b_1, each of length at most L, and which are identified by f and g.

Using separability there are finite covers $\tilde{A}, \tilde{B}$ of A and B such that the only loops of length at most $2L$ which lift to these covers correspond to elements of G. In claim 2 we show that the induced gluing of these covers is simple as follows.

Suppose two points $\tilde{a}_1, \tilde{a}_2$ in $\tilde{A}$ are R-related to the same point $\tilde{b}$ in $\tilde{B}$. Then there are two (possibly very long) geodesics $\tilde{\alpha}_1, \tilde{\alpha}_2$ in $\tilde{A}$ each of which is identified to a geodesic $\tilde{\beta}_1, \tilde{\beta}_2$ in $\tilde{B}$. The geodesic $\tilde{\alpha}_i$ starts at the basepoint $\tilde{a}_0 \in \tilde{A}$ and ends at $\tilde{a}_i$. The geodesics $\tilde{\beta}_i$ both start at $\tilde{b}_0$ and both end at $\tilde{b}_1 = \tilde{b}_2$. Projecting these geodesics gives geodesics $\alpha_1, \alpha_2 \subset A$ and $\beta_1, \beta_2 \subset B$. Claim 1 implies there are new geodesics, $\gamma_i \subset A$ with the same endpoints as α_i, and $\delta_i \subset B$ with the same endpoints as β_i, such that γ_i is identified to δ_i, and these new geodesics all have length at most L. The loop $\alpha_i \cdot \gamma_i^{-1} \subset A$ is identified to $\beta_i \cdot \delta_i^{-1} \subset B$. Thus they represent elements in G and thus lift to the covers. Since $\tilde{b}_1 = \tilde{b}_2$ we deduce $\delta_1 \cdot \delta_2^{-1}$ lifts to a loop. Since $\delta_1 \cdot \delta_2^{-1}$ has length at most $2L$ it therefore must lie in G. This implies $\gamma_1 \cdot \gamma_2^{-1}$ is also a loop (ie. $a_1 = a_2$) and that it lifts to a loop in $\tilde{A}$ which implies $\tilde{a}_1 = \tilde{a}_2$. Thus there are no self-identifications, proving claim 2.

Finally, in the non-compact case, we first replace A and B by relative thickenings of A and B so that each thin cusp of A or B is contained in a product cusp which in turn is contained in a cusp of the thickenings. Then we truncate along cusp boundaries to obtain compact submanifolds of the relative thickenings. This allows us to find a constant L that applies to points in these submanifolds and the previous argument shows there are no self-identifications in the submanifolds. The fact that thin cusps of A and B are contained in product cusps of the thickenings means any self-identifications within a thin cusp propagate vertically within the larger product cusp in a product like way to give identifications where the cusp meets the compact submanifold at the cusp boundary. But such identifications have been ruled out. Thus there are no self-identifications in the thin cusps of A or B either. This completes the sketch.

The first step is to replace A and B by thickenings so that the thin cusps of A and B are contained in product cusps of the thickenings. We start by renaming A as A^* and B as B^*. Let $\mathcal{C}$ be a maximal collection of pairwise disjoint cusps in M. By choosing $\mathcal{C}$ sufficiently small we may ensure that $C_A = f^{-1}(\mathcal{C})$ and $C_B = g^{-1}(\mathcal{C})$ is a maximal collection of pairwise disjoint cusps in A^* and B^* respectively. We choose the cusps in $\mathcal{C}$ small enough so that $A^* \neq C_A$ and $B^* \neq C_B$ and so the cusps C_A and C_B are thin. By 3.5(e) there is $K > 0$ such that each component of C_A and of C_B is contained in a product cusp contained in $A \equiv Th_K^{rel}(A^*)$ or $B \equiv Th_K^{rel}(B^*)$ respectively. The maps f, g have natural extensions to local isometries $f : A \to M$ and $g : B \to M$. Henceforth we shall use f and g to denote these extensions.

We now remove the cusps to obtain $A^- = A \setminus f^{-1}(int(C))$, and $B^- = B \setminus g^{-1}(int(C))$. Then A^-, B^- are compact. Furthermore each component of C_A is contained in a product cusp contained in $A \setminus A^-$ and similarly for B.

Claim 1. There is $L > 0$ such that if $a_1 \in A^-$ and $b_1 \in B^-$ and $a_1 R b_1$ then there are geodesics $\alpha \subset A$ and $\beta \subset B$ with $length(\alpha) = length(\beta) \leq L$ and $f \circ \alpha = g \circ \beta$ and α has endpoints a_0, a_1 and β has endpoints b_0, b_1.

Proof of Claim 1. Define

$$S = \{(a_1, b_1) \in A \times B : a_1 R b_1\} \qquad \textit{and} \qquad \ell : S \to \mathbb{R}$$

by $\ell(a_1, b_1)$ is the length of some shortest geodesic, α, as above. By 4.1(c) $S(f, g)$ is Hausdorff thus S is a closed subset of $A \times B$. We first show that ℓ is continuous at points of S in the interior of $A \times B$.

If $(a_1, b_1) \in S$ and a_1 is in the interior of A and b_1 is in the interior of B then ℓ is continuous at (a_1, b_1). This is because there are small open balls $U \subset A$ and $V \subset B$ centered on a_1 and b_1 with $fU = gV$. It is clear that each $a \in U$ is R-related to a unique point $b \in V$. Since A and

B are convex there are geodesics, α', β' close to α and β from the basepoints to a and b and $f\alpha' = g\beta'$. This proves continuity of ℓ at interior points.

By enlarging A and B to slightly larger convex manifolds, A^+, B^+ every point in S is in the interior of $A^+ \times B^+$ thus the corresponding function ℓ^+ defined on S^+ is continuous at every point in S. The function ℓ^+ is defined using geodesics in the enlarged manifolds. However if a geodesic in A^+ starts and ends in A then, since A is convex, the geodesic is contained in A. It follows that $\ell^+|S = \ell$ and thus ℓ is continuous on all of S.

Restricting the continuous function ℓ to the compact set $S \cap (A^- \times B^-)$ we obtain the required bound L on ℓ. Observe that the shortest geodesics used above connecting a point in A^- to the basepoint a_0 are not necessarily contained in A^-, they may go into the cusp of A some bounded distance. This proves claim 1. □

Let $S \subset \pi_1(B, b_0)$ be the set of elements represented by based loops of length at most $2L$. Since B is convex S is finite. Since A, B, M are all convex, the maps f_*, g_* are injective by proposition 2.3, and for notational simplicity we will identify G with G_A and G_B. Since G is separable in $\pi_1(B, b_0)$ there is a subgroup, $H < \pi_1(B, b_0)$, of finite index which contains G and contains no element of $S \setminus G$. It follows that if an element $h \in H$ is represented by a based loop of length at most $2L$ then $h \in G$.

Let $p_B : \tilde{B} \to B$ be the cover corresponding to this subgroup. Similarly define $p_A : \tilde{A} \to A$. Let $\tilde{f} = f \circ p_A$ and $\tilde{g} = g \circ p_B$. Define $\tilde{A}^- = p_A^{-1}(A^-)$ and $\tilde{B}^- = p_B^{-1}(B^-)$. Then A^-, B^- are compact. The following claim and lemma 4.1(a) imply the induced gluing, $S(\tilde{f}, \tilde{g})$, of $\tilde{A}^-$ and $\tilde{B}^-$ is simple.

Claim 2. If $\tilde{a}_1, \tilde{a}_2 \in \tilde{A}^-$ are both R-related to $\tilde{b}_1 = \tilde{b}_2 \in \tilde{B}^-$ then $\tilde{a}_1 = \tilde{a}_2$ and similarly with the roles of $\tilde{A}^-$ and $\tilde{B}^-$ reversed.

Proof of Claim 2. Observe that $a_i = p_A(\tilde{a}_i) \in A^-$ and $b_i = p_B(\tilde{b}_i) \in B^-$.

For $i \in \{1, 2\}$ because $\tilde{a}_i R \tilde{b}_i$ there is a geodesic $\tilde{\alpha}_i \subset \tilde{A}$ starting at $\tilde{a}_0$ and ending at $\tilde{a}_i$ and a geodesic $\tilde{\beta}_i \subset \tilde{B}$ starting at $\tilde{b}_0$ and ending at $\tilde{b}_i$ such that $\tilde{f} \circ \tilde{\alpha}_i = \tilde{g} \circ \tilde{\beta}_i$. Project these into A and B to obtain geodesics $\alpha_i = p_A \circ \tilde{\alpha}_i$ in A and $\beta_i = p_B \circ \tilde{\beta}_i$ in B. Observe that $f \circ \alpha_i = g \circ \beta_i$. Thus there are geodesics $\gamma_i \subset A$ and $\delta_i \subset B$ of length at most L with $f\gamma_i = g\delta_i$ such that the endpoints of α_i and γ_i are the same and the endpoints of β_i and δ_i are the same as shown in the diagram.

Consider the loop $\alpha_i \cdot \gamma_i^{-1}$ in A. The image of this loop under f is the same as the image under g of the loop $\beta_i \cdot \delta_i^{-1}$ hence these loops give elements of G. It follows that the loops $\alpha_i \cdot \gamma_i^{-1}$ lift to loops in $\tilde{A}$ based at $\tilde{a}_0$. Similarly the loops $\beta_i \cdot \delta_i^{-1}$ lift to loops in $\tilde{B}$ based at $\tilde{b}_0$.

Claim 3. The loop $\delta = \delta_1 \cdot \delta_2^{-1}$ lifts to a loop $\tilde{\delta}$ in $\tilde{B}$ based at $\tilde{b}_0$.

Proof of Claim 3. This is because

$$\delta = \delta_1 \cdot \delta_2^{-1} = (\delta_1 \cdot \beta_1^{-1}) \cdot (\beta_1 \cdot \beta_2^{-1}) \cdot (\beta_2 \cdot \delta_2^{-1}).$$

The loops $\delta_1 \cdot \beta_1^{-1}$ and $\beta_2 \cdot \delta_2^{-1}$ lift to loops based at $\tilde{b}_0$ by the previous paragraph. The path $\tilde{\beta}_i$ starts at $\tilde{b}_0$ and ends at $\tilde{b}_i$. Using the assumption that $\tilde{b}_1 = \tilde{b}_2$ we see that the $\tilde{\beta}_1, \tilde{\beta}_2$ have the same endpoints thus $\beta_1 \cdot \beta_2^{-1}$ lifts to a loop based at $\tilde{b}_0$. Thus all three loops in the product lift to loops based at $\tilde{b}_0$. This proves claim 3. □

We continue with the proof of claim 2. In what follows $\simeq$ denotes homotopy between maps of an interval keeping endpoints fixed. Since $length(\delta) \leq 2L$, and using the definition of $\tilde{B}$, we see that $[\delta] \in G$. Since $[\delta] \in G$ there is a loop, η, in A based at a_0 such that $f \circ \eta \simeq g \circ \delta$. By sliding one endpoint of η along γ_1 one obtains a path ϵ in A with endpoints a_0, a_1 such that

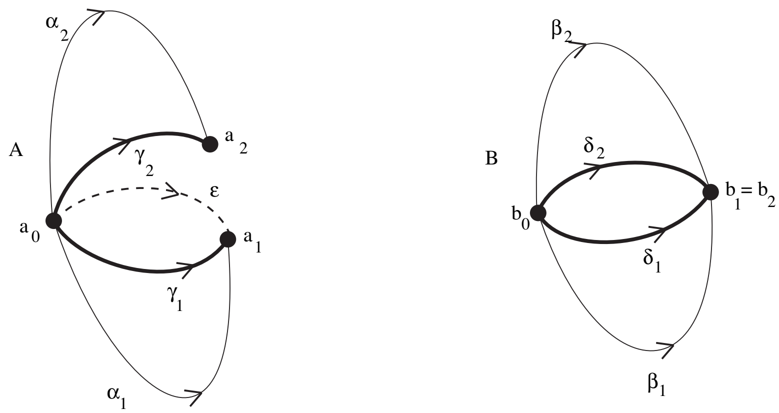

FIGURE 7. *R-Related paths.*

$f \circ (\gamma_1 \cdot \epsilon^{-1}) \simeq g \circ \delta$. By changing basepoints it follows that

$$f \circ (\epsilon^{-1} \cdot \gamma_1) \simeq g \circ (\delta_2^{-1} \cdot \delta_1).$$

Since A is convex we may homotop ϵ keeping its endpoints fixed to be a geodesic.

Combining this with:

$$\begin{array}{rcl} f \circ (\epsilon^{-1} \cdot \gamma_1) & = & (f \circ \epsilon^{-1}) \cdot (f \circ \gamma_1) \\ g \circ (\delta_2^{-1} \cdot \delta_1) & = & (g \circ \delta_2^{-1}) \cdot (g \circ \delta_1) \\ f \circ \gamma_1 & = & g \circ \delta_1 \end{array}$$

it follows that $f \circ \epsilon^{-1} \simeq g \circ \delta_2^{-1}$ and thus $f \circ \epsilon \simeq g \circ \delta_2$. Since $f \circ \epsilon$ and $g \circ \delta_2$ are geodesics in the convex hyperbolic manifold M with the same endpoints, and are homotopic rel endpoints, it follows they are equal. But $f \circ \gamma_2$ is also equal to this geodesic. Thus γ_2 and ϵ are two geodesics in A which start at the same point and map under the local isometry f to the same geodesic. It follows that $\epsilon = \gamma_2$.

We now have a loop $\gamma_1 \cdot \gamma_2^{-1} \subset A$ such that $f \circ (\gamma_1 \cdot \gamma_2^{-1}) \simeq g \circ \delta$. Since the latter is in G it follows that $\gamma_1 \cdot \gamma_2^{-1}$ lifts to a loop based at $\tilde{a}_0$ and therefore $\tilde{a}_1 = \tilde{a}_2$. This proves claim 2. □

We now finish the proof of the theorem. We have shown that the induced gluing of $\tilde{A}^-$ and $\tilde{B}^-$ is simple. Suppose $a_1, a_2 \in \tilde{A}^*$ are identified by the induced gluing of the corresponding covers $\tilde{A}^*$ and $\tilde{B}^*$. Since $\tilde{A}^* \subset \tilde{A}$ it follows that that a_1, a_2 are in the cusps $\tilde{C}_A$ of $\tilde{A}^*$ that cover C_A. A finite cover of a product cusp is also a product cusp. There are product cusps contained in $\tilde{A}$ which contain $\tilde{C}_A$. A product cusp has a 1-dimensional foliation by rays starting on the cusp boundary. Hence two of these rays are identified by the induced gluing. This implies two point on the cusp boundaries are also identified. The cusp boundaries are in $\tilde{A}^-$ thus there are two points in $\tilde{A}^-$ which are identified and this is a contradiction. Similar remarks apply to B'. Hence the induced gluing of $\tilde{A}'$ and $\tilde{B}'$ is simple. This completes the theorem. □

Here is an example that illustrates what the simple gluing theorem does. Suppose that A and B are each a surface times an interval and are the convex cores of two quasi-Fuchsian 3-manifolds which are immersed into some convex hyperbolic 3-manifold M. By theorem 8.1 the intersection of geometrically finite subgroups is geometrically finite so $G = f_*\pi_1(A, a_0) \cap g_*\pi_1(B, b_0)$ is geometrically finite. By Scott's theorem this group is separable in $\pi_1(A, a_0)$ and $\pi_1(B, b_0)$.

The virtual simple gluing theorem applied with G being this intersection produces finite covers of A and B where the pre-image of this intersection group is embedded. It then identifies the covers along submanifolds corresponding to these subgroups. We call these submanifolds (and the subsurfaces of the quasi-Fuchsian surfaces they correspond to) the *region of parallelism* between A and B.

Assume for simplicity that A and B are embedded and contain spines which are disjoint surfaces S_A, S_B. Then the region of parallelism is the horizontal boundary of the maximal I-bundle, J, between S_A and S_B for which some fiber connects the basepoints of S_A and S_B. (J is a compact, connected, π_1-injective product I-bundle in M that contains the basepoint and has one horizontal boundary component in each of S_A and S_B. The interior of J is disjoint from $S_A \cup S_B$. Also $\pi_1 J = \pi_1 A \cap \pi_1 B$. Possibly J is a disc times an interval, otherwise J is unique up to isotopy fixing S_A and S_B.)

It should be pointed out that the hyperbolic manifold $S(\tilde{f}, \tilde{g})$ produced by a simple gluing is usually not convex, and in general the induced map $S(\tilde{f}, \tilde{g}) \to M$ is not π_1-injective. To achieve π_1-injectivity one needs some extra hypotheses, as in the convex combination theorem.

5. *The virtual amalgam and virtual convex combination theorems.*

In this section we give two results which have the same conclusion as the convex combination theorem but with different hypotheses. The idea is that in some situations when one wishes to glue two convex manifolds the hypotheses of the convex combination theorem are always satisfied by certain finite covers of the manifolds in question.

Definition. Suppose that $M = M_1 \cup M_2$ is a hyperbolic n-manifold which is the union of two convex hyperbolic n-submanifolds, M_1, M_2. Suppose that $\pi : \tilde{M} \to M$ is a finite cover and $\tilde{M}_i$ is a component of $\pi^{-1}(M_i)$ and $\tilde{C}$ is a component of $\pi^{-1}(M_1 \cap M_2)$. The hyperbolic n-manifold obtained from the disjoint union of $\tilde{M}_1$ and $\tilde{M}_2$ by identifying the copy of $\tilde{C}$ in each is called a *virtual gluing* of M_1 and M_2. This equals the simple gluing $S(\pi|\tilde{M}_1, \pi|\tilde{M}_2)$ with basepoints chosen in $\tilde{C}$.

Theorem 5.1 (virtual compact convex combination theorem). *Suppose that $M = M_1 \cup M_2$ is a hyperbolic n-manifold, M_1, M_2 are compact, convex hyperbolic n-manifolds, and C is a component of $M_1 \cap M_2$. Also suppose that $\pi_1 C$ is a separable subgroup of both $\pi_1 M_1$ and $\pi_1 M_2$. Then there is a virtual gluing, N, of M_1 and M_2 along C which has a convex thickening. In particular N is isometric to a submanifold of $\mathbb{H}^n/hol(\pi_1 N)$.*

Proof. Before starting, we remark that a short proof can be given if it is assumed that M is contained in a convex hyperbolic manifold N. In this case one has locally isometric immersions $Y_i = Th^{rel}_{\kappa}(M_i) \hookrightarrow N$. The virtual simple gluing theorem then gives a simple gluing of finite covers $\tilde{Y}_1 \cup \tilde{Y}_2$, and the convex combination theorem now implies this has a convex thickening, giving the result.

We will show that there is a virtual simple gluing of M_1 and M_2 which extends to a virtual gluing of the κ-thickenings. The result then follows from the convex combination theorem (2.9.

By 2.4 C is convex and thus $\pi_1 C$ injects into $\pi_1 M_i$ by 2.3. Thus the universal cover, $\tilde{C}^u$, of C is a submanifold of the universal cover, $\tilde{M}_i^u$, of M_i. Hence we may embed $\tilde{M}_1^u$ and $\tilde{M}_2^u$ in $\mathbb{H}^n$ so that $\tilde{M}_1^u \cap \tilde{M}_2^u \supseteq \tilde{C}^u$. Convexity implies this is an equality. The action of $\pi_1 C$ on $\mathbb{H}^n$ preserves $\tilde{M}_1^u$ and $\tilde{M}_2^u$. Let $N = (\tilde{M}_1 \cup \tilde{M}_2)/\pi_1 C$. Thus N is a submanifold of $\mathbb{H}^n/\pi_1 C$. Define $\tilde{M}_1 \equiv \tilde{M}_1^u/\pi_1 C$ and $\tilde{M}_2 \equiv \tilde{M}_2^u/\pi_1 C$ then $N = \tilde{M}_1 \cup \tilde{M}_2$ and $\tilde{M}_1 \cap \tilde{M}_2 = \tilde{C}$ is a lift of C. Define

$$N^+ = \{ \ x \in \mathbb{H}^n/\pi_1 C \ : \ d(x, N) \leq \kappa \ \}.$$

Then $N^+ = Th_\kappa(\tilde{M}_1) \cup Th_\kappa(\tilde{M}_2)$. Clearly there is a covering map $Th_\kappa(\tilde{M}_i) \to Th_\kappa(M_i)$.

Let $\tilde{C}^+$ be the component of $Th_\kappa(\tilde{M}_1) \cap Th_\kappa(\tilde{M}_2)$ which contains $\tilde{C}$. Since $Th_\kappa(\tilde{M}_1)$ and $Th_\kappa(\tilde{M}_2)$ are convex, 2.4 implies $\tilde{C}^+$ is convex. From 2.3 it follows that $\pi_1\tilde{C}^+ \to \pi_1(\mathbb{H}^n/\pi_1 C)$ is injective. Since $\tilde{C} \subset \tilde{C}^+$ it follows that the inclusion $\tilde{C} \hookrightarrow \tilde{C}^+$ induces an isomorphism of fundamental groups. Thus $\tilde{C}^+$ is a convex thickening of $\tilde{C}$.

We now show that $\tilde{C}^+$ is compact. Otherwise there is geodesic ray λ in $\tilde{C}^+$ which starts at a point $p \in \tilde{C}$ and leaves every compact set. Consider $\lambda_i \equiv \lambda \cap \tilde{M}_i$. Since $p \in \tilde{C} \subset \tilde{M}_i$ it follows that λ_i contains p. By considering the universal cover of the K-thickening of $\tilde{M}_i$, and using convexity of $\tilde{M}_i$, it is easy to see that $\lambda_i = \lambda$. Hence $\lambda \subset \tilde{C}$. But this contradicts that $\tilde{C}$ is compact, and proves $\tilde{C}^+$ is compact.

Claim. There is a finite cover $Y_i \to Th_\kappa(M_i)$ such that the natural map $p_i : \tilde{C}^+ \to Th_\kappa(M_i)$ lifts to an injective map $\tilde{p}_i : \tilde{C}^+ \to Y_i$.

Assuming the claim, consider the hyperbolic manifold $Y = Y_1 \cup Y_2$ obtained by gluing Y_1 and Y_2 along $\tilde{C}^+$. Then Y contains $N = \tilde{M}_1 \cup \tilde{M}_2$ and $Y_i = Th_\kappa(\tilde{M}_1)$. The convex combination theorem now gives the result.

Proof of Claim. Choose a basepoint $x \in C$ and observe that $x \in C \subset M_i \subset Th_\kappa(M_i)$. Under the identification given by the lift, $C \equiv \tilde{C}$, the point x determines a basepoint $\tilde{x} \in \tilde{C} \subset \tilde{C}^+$. In what follows all fundamental groups are based at the relevant basepoint.

Let S_i be the subset of $\pi_1 Th_\kappa(M_i)$ represented by loops of length at most 3 times the diameter of $\tilde{C}^+$. Since $\tilde{C}^+$ is compact S_i is finite. Since $\pi_1 C$ is a separable subgroup of $\pi_1 M_i$ there is finite index subgroup of $\pi_1 M_i$ which contains $\pi_1 C$ and contains no element of $S_i \setminus \pi_1 C$. Let $Y_i \to Th_\kappa(M_i)$ be the corresponding cover.

Clearly the natural map p_i lifts to $\tilde{p}_i : \tilde{C}^+ \to Y_i$. We have a basepoint $\tilde{p}_i(\tilde{x}) \in Y_i$. If $\tilde{p}_i$ is not injective there are distinct points $\tilde{x}_1 \neq \tilde{x}_2 \in \tilde{C}^+$ with $\tilde{p}_i(\tilde{x}_1) = \tilde{p}_i(\tilde{x}_2)$. Then there is a geodesic segment $\tilde{\alpha}$ in $\tilde{C}^+$ with endpoints $\tilde{x}_1$ and $\tilde{x}_2$, of length at most the diameter of $\tilde{C}^+$. Observe that $\tilde{\alpha}$ maps to a loop $\tilde{p}_i \circ \tilde{\alpha}$ in Y_i. Thus $\tilde{p}_i \circ \tilde{\alpha}$ projects to a loop, α, in $Th_\kappa(M_i)$ based at the point $x_1 = p_i(\tilde{x}_1)$. Let $\tilde{\beta}$ be a geodesic segment of minimal length in $\tilde{C}^+$ with endpoints $\tilde{x}$ and $\tilde{x}_1$. Then $\beta = p_i \circ \tilde{\beta}$ is a geodesic segment in $Th_\kappa(M_i)$ starting at x and ending at x_1. Thus $\gamma = \beta.\alpha.\beta^{-1}$ is a loop in $Th_\kappa(M_i)$ based at x of length at most 3 times the diameter of $\tilde{C}^+$. Thus $[\gamma]$ is an element of S_i. It is easy to see that γ lifts to a loop in Y_i based at $\tilde{p}_i(\tilde{x})$ so $[\gamma] \in Im[\pi_1 C \to \pi_1 Th_\kappa(M_i)]$. This implies γ lifts to a loop in $\tilde{C}^+$ based at $\tilde{x}$, which contradicts that $\tilde{\alpha}$ has distinct endpoints. This proves the claim and the theorem. □

Proposition 5.2 (increasing the rank of a cusp). *Suppose that M is a convex hyperbolic n-manifold and C is a thin cusp of M. Suppose that $\Gamma < Isom(\mathbb{H}^n)$ is a discrete group of parabolic isometries which contains $hol(\pi_1 C)$. Then there is a finite-index subgroup $\Gamma' < \Gamma$ which contains $hol(\pi_1 C)$ and a horoball $D \subset \mathbb{H}^n$ stabilized by Γ' with the following property. Set $Q = D/\Gamma'$ then there is a hyperbolic n-manifold $N = M \cup Q$ such that $M \cap Q = C$. Furthermore N has a convex thickening.*

Proof. Let $\pi : \tilde{M} \to M$ be the universal cover. Use the developing map to isometrically identify $\tilde{M}$ with a subset of $\mathbb{H}^n$. Let $\tilde{C}$ be a component of $\pi^{-1}C$. Let D be the horoball which contains $\tilde{C}$ and so that ∂D contains $\tilde{C} \cap \pi^{-1}(\partial_c C)$. Let $D^- \subset D$ be the smaller horoball such that D is a κ-neighborhood of D^-. Then $C^- = (\tilde{C} \cap D^-)/hol(\pi_1 C)$ is a smaller cusp contained in C. The cusp $W = D/\Gamma$ has boundary ∂W which is a Euclidean manifold. There is also a smaller cusp $W^- = D^-/\Gamma$.

By 3.5(e), for K sufficiently large, $Y_1 = Th_K^{rel}(M; C^-)$ contains M. We regard the universal cover $\tilde{Y}_1$ as a subset of $\mathbb{H}^n$ so that it contains $\tilde{M}$ in the natural way. Let C_1 be the image of $D \cap \tilde{Y}_1$ under projection to Y_1. This is a cusp in Y_1 and by choosing K sufficiently large

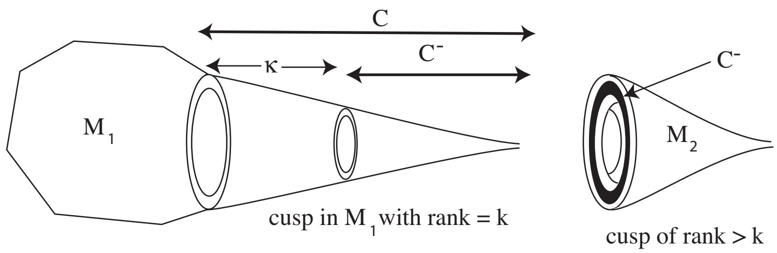

FIGURE 8. *Gluing on a rank-k cusp.*

(and $K \geq \kappa$) we may arrange that C_1 contains C and so $\partial_c C \subset \partial_c C_1$. Thus C_1 is the cusp in Y_1 which naturally corresponds to the cusp C in M_1. Since $hol(\pi_1 C) < \Gamma$ there is a natural local-isometry $f : C_1 \to W$.

CLAIM. There is a finite cover $\tilde{W} \to W$ so that f lifts to an embedding $\tilde{f} : C_1 \to \tilde{W}$.

Assuming this claim, since $C^- \subset C \subset C_1$ it follows that $\tilde{f}| : C^- \to \tilde{W}^-$ is injective. In order to fit with the notation used in the convex combination theorem we now use M_1 to denote M and M_2 to denote $\tilde{W}^-$. We use $\tilde{f}$ to identify $C^- \subset M_1$ with its image in M_2 and set $M = M_1 \cup M_2$ thus $M_1 \cap M_2 = C^-$. Similarly let $Y_2 = Th_\kappa(M_2) = \tilde{W}$ and use $\tilde{f}$ to identify $C_1 \subset Y_1$ with its image in Y_2 then $Y = Y_1 \cup Y_2$ and $Y_1 \cap Y_2 = C_1$.

We now check the 6 hypotheses of the convex combination theorem are satisfied. Using the fact that a relative thickening of a convex manifold is convex, it is easy to check that M_1, M_2, Y_1, Y_2 are all convex which gives conditions (i) and (ii). Condition (iii) is immediate. Since $K \geq \kappa$, by 3.5(c), Y_1 contains a κ-neighborhood of $\overline{M_1 \setminus M_2}$. Also $Y_2 = Th_\kappa(M_2)$ contains a κ-neighborhood of M_2. This implies conditions (iv) and (v) are satisfied. Finally, $C^- = M_1 \cap M_2$ is contained in $C_1 = Y_1 \cap Y_2$, and both are connected, so condition (vi) is satisfied. The convex combination theorem implies that M has a convex thickening $Th_\infty(M)$. This contains Y and thus contains $M \cup Y_2$. Now $M \cap Y_2 = C$ and this proves the theorem with $Q = Y_2$.

Proof of Claim. Since C is a thin cusp C_1 is also a thin cusp. By 3.5(e) we may assume K was chosen large enough that there is a product cusp, P, with $C \subset P \subset C_1$. Clearly the local isometry f has an extension to a local isometry $f : P \to W$. Thus it suffices to prove this extension is injective. A product cusp has a 1-dimensional foliation by rays orthogonal to the cusp boundary. If two distinct points in P have the same image under f then the rays through these points have the same image under f. It follows that if f is not injective there are two points in the cusp boundary $\partial_c P$ which have the same image under f. Thus it suffices to show there is a finite cover of ∂W so that $f| : \partial_c P \to \partial W$ lifts to an embedding.

Since W is a cusp $\pi_1 W \cong \pi_1 \partial W$. The manifold ∂W is Euclidean and so has a fundamental group which is virtually free-abelian. Thus $\pi_1 W$ is subgroup separable. Since P is a product cusp $\partial_c P$ is compact. Since f is a local isometry and $\partial_c P$ is compact and convex it follows from a standard argument there is a finite cover as in the previous paragraph. This proves claim. □

The convex combination theorem sometimes enables one to glue geometrically finite manifolds together to obtain a geometrically finite manifold. This corresponds to forming an amalgamated free product of two geometrically finite groups, amalgamated along their intersection.

Definition. Two subgroups A, B of a group G can be *virtually amalgamated* if there are finite index subgroups $A' < A$ and $B' < B$ such that the subgroup, G', of G generated by A' and B' is the free product of A' and B' amalgamated along $A' \cap B'$. We also say that G' is a *virtual amalgam* of A and B.

Definition. Two non-trivial parabolic subgroups $\Gamma_1, \Gamma_2 < Isom(\mathbb{H}^n)$ are called *compatible* if either:
(i) Γ_1 and Γ_2 stabilize distinct points on the sphere at infinity, or
(ii) $\Gamma_1 \cap \Gamma_2$ has finite index in at least one of the groups Γ_1 or Γ_2.

The second condition is equivalent to saying that, up to taking subgroups of finite index, that one group is a subgroup of the other. Two discrete groups $\Gamma, \Gamma' < Isom(\mathbb{H}^n)$ have *compatible parabolic subgroups* if every maximal parabolic subgroup of Γ is compatible with every maximal parabolic subgroup of Γ'.

Theorem 5.3 (GF subgroups have virtual amalgams). *Suppose that Γ is a discrete subgroup of $Isom(\mathbb{H}^n)$. Suppose that Γ_1 and Γ_2 are two geometrically finite subgroups of Γ which have compatible parabolic subgroups. Suppose that $\Gamma_1 \cap \Gamma_2$ is separable in both Γ_1 and Γ_2. Then Γ_1 and Γ_2 can be virtually amalgamated and the result is a geometrically finite group.*

Proof. We first reduce to the torsion-free case. Since Γ is linear there is a torsion-free subgroup, G, of finite index in Γ. We now replace $\Gamma_1, \Gamma_2, \Gamma$ by the finite index subgroups $G \cap \Gamma_1, G \cap \Gamma_2, G$. It is routine to show that the separability hypothesis is satisfied by these new groups. Thus we may assume that Γ is torsion-free. The next step if to produce convex manifolds, each containing the basepoint, corresponding to these groups

Let $N = \mathbb{H}^n/\Gamma$. Choose a basepoint $\tilde{x} \in \mathbb{H}^n$ and let x be the image in N of $\tilde{x}$. This choice determines an identification $\pi_1(N, x) \equiv \Gamma$. Let $H_i = CH(\Gamma_i \cdot \tilde{x})$ then $M_i = H_i/\Gamma_i$ is a convex hyperbolic manifold which, because Γ_i is geometrically finite, has finite volume. It is easy to see that M_i has thin cusps. Let x_i be the image of $\tilde{x}$ in M_i. The choice of $\tilde{x}$ determines an identification $\pi_1(M_i, x_i) \equiv \Gamma_i$. The inclusion $H_i \subset \mathbb{H}^n$ covers a local isometry $\rho_i : M_i \to N$. Then ρ_{i*} maps $\pi_1(M_i, x_i)$ into a subgroup of $\pi_1(N, x)$ and under the identifications $\pi_1(N, x) \equiv \Gamma$ and $\pi_1(M_i, x_i) \equiv \Gamma_i$ the map ρ_{i*} is inclusion $\Gamma_i \subset \Gamma$. The result now follows from 5.4. □

Theorem 5.4 (Immersed Virtual Convex Combination Theorem). *Suppose that $N = \mathbb{H}^n/\Gamma$ is a hyperbolic manifold and M_1, M_2 are geometrically-finite convex hyperbolic manifolds with thin cusps. Suppose that $\rho_i : M_i \to N$ is a locally isometric immersion. Suppose that $x = \rho_1(x_1) = \rho_2(x_2)$ and choose an identification $\pi_1(N, x) \equiv \Gamma$. This yields identifications $\pi_1(M_i, x_i) \equiv \Gamma_i \subset \Gamma$. Assume that Γ_1 and Γ_2 have compatible cusps. Set $\Gamma_0 = \Gamma_1 \cap \Gamma_2$. Suppose that Γ_0 is a separable subgroup of both Γ_1 and of Γ_2. Then there are finite covers $p_i : \tilde{M}_i \to M_i$ and a connected hyperbolic manifold $\tilde{M}_1 \cup \tilde{M}_2$ which has a convex thickening. Furthermore $\tilde{M}_1 \cap \tilde{M}_2$ is connected and contains a point covering m_1.*

Proof. The strategy is to construct certain convex thickenings of M_1 and M_2 and show that there are finite covers of these thickenings which have a simple gluing. We then apply the convex combination theorem to obtain a convex thickening of this gluing.

Throughout this proof all the convex manifolds we consider are submanifolds of covering spaces of N, and are thus equipped with isometric immersions into N. We wish to choose a consistent system of cusps in all these manifolds. Let $\mathcal{C}$ be a maximal collection of disjoint cusps in N. If $\rho : X \to M$ is an isometric immersion by a *cusp in X* we will mean (until stated otherwise) a non-simply connected component of $\rho^{-1}(\mathcal{C})$. It is possible that some components of $\rho^{-1}(\mathcal{C})$ are balls. When we construct relative thickenings of X it will be relative to these cusps. We shall avoid mentioning these choices of cusps whenever possible.

We fix $K \geq \kappa$ which will be chosen later and consider the relative thickenings $Y_i = Th_K^{rel}(M_i)$. By 3.5 these are convex and have thin cusps. Since $\pi_1(Y_i, x_i) \cong \pi_1(M_1, x_i)$ the map ρ_i extends to an isometric immersion $\rho_i : (Y_i, x_i) \to (N, x)$. The separability hypothesis enables us to apply the virtual simple gluing theorem 4.3 to these extensions of ρ_1, ρ_2. Thus there are finite covers $\tilde{Y}_i \to Y_i$ and a simple gluing $\tilde{Y} = \tilde{Y}_1 \cup \tilde{Y}_2$. Let $\tilde{M}_i \subset \tilde{Y}_i$ be the corresponding cover of M_i. Clearly $\tilde{Y}_i = Th_K^{rel}(\tilde{M}_i)$. By 4.1(e) $\tilde{Y}_1 \cap \tilde{Y}_2$ and $\tilde{M}_1 \cap \tilde{M}_2$ are both connected. By 2.8 $\pi_1(\tilde{M}_1 \cup \tilde{M}_2)$ is an amalgamated free product. Since Γ_i is geometrically finite, $\tilde{Y}_i$ has finite volume. If we can verify the hypotheses for the convex combination theorem for $\tilde{M}_1 \cup \tilde{M}_2 \subset \tilde{Y}_1 \cup \tilde{Y}_2$, then $\tilde{M}_1 \cup \tilde{M}_2$ has a convex thickening which is geometrically finite, and this proves the theorem.

We now attempt to verify the hypotheses of the convex combination theorem. Conditions (i),(ii),(iii) are obviously satisfied. Since $\tilde{M}_1 \cap \tilde{M}_2$ and $\tilde{Y}_1 \cap \tilde{Y}_2$ are both connected, and both contain the basepoint, condition (vi) holds. However conditions (iv) and (v) are not always satisfied.

Special Case. Every cusp of M_1 and of M_2 has maximal rank $n-1$.

Proof of special case. Observe that if M is a convex manifold with all cusps of maximal rank then thickening and relative thickening coincide: $T_K^{rel}(M) = T_K(M)$. Furthermore this condition is preserved by finite covers. Thus in this case $\tilde{Y}_i$ contains a κ-neighborhood of $\tilde{M}_i$. Hence conditions (iv) and (v) are satisfied, and the theorem follows in this special case. This proves the special case. □

Returning to the general case, the convex combination theorem requires a κ-neighborhood. The obvious way to produce a κ-neighborhood is to use κ-thickenings, but in general this produces cups that are not thin. The virtual simple gluing theorem requires thin cusps, so we can't use it to produce a simple gluing of κ-thickenings. In the general case the proof is necessarily a bit more involved, since it must utilize the hypothesis of compatible cusps.

A *generalized cusp*, C, of $\tilde{Y}$ is a component of the pre-image of a cusp of N under the natural isometric immersion $\tilde{Y} \to N$. Since $\tilde{Y}$ is not convex, C need not be convex. In what follows we will use C_i to denote a cusp in $\tilde{M}_i$ and C_i^* to denote the corresponding (relatively) thickened cusp in $\tilde{Y}_i$. Since $\tilde{M}_i \subset \tilde{Y}_i$ it follows that $C_i \subset C_i^*$.

The strategy is the following. We will construct certain convex thickenings $\tilde{Y}_1^+$ and $\tilde{Y}_2^+$ by selectively thickening certain cusps of $\tilde{Y}_1$ and $\tilde{Y}_2$. This is done in such a way that the simple gluing $\tilde{Y} = \tilde{Y}_1 \cup \tilde{Y}_2$ extends to a simple gluing $\tilde{Y}^+ = \tilde{Y}_1^+ \cup \tilde{Y}_2^+$. The crucial property is that $\tilde{Y}^+$ now contains a κ-neighborhood of $\tilde{M}$. We now sketch how this is done before giving the proof.

We construct $\tilde{Y}_i^+$ by taking a K-thickening of $\tilde{M}_i$ relative to a carefully chosen subset, $\mathcal{C}_i$, of the cusps of $Th_\infty(\tilde{M}_i)$. Those cusps of $\tilde{Y}_i^+$ corresponding to cusps in $\mathcal{C}_i$ are the same as the corresponding cusps of $\tilde{Y}_i$ and are thus thin. The other cusps of $\tilde{Y}_i^+$ are thickenings of the corresponding cusps in $\tilde{Y}_i$ and contain K-neighborhoods of the cusps in $\tilde{M}_i$.

We will show if K is sufficiently large and the cusps $\mathcal{C}$ are sufficiently small, then every generalized cusp in $\tilde{Y}$ is the union of at most one cusp $C_1^* \subset \tilde{Y}_1$ and one cusp $C_2^* \subset \tilde{Y}_2$. A cusp of $\tilde{Y}$ that is disjoint from either $\tilde{Y}_1$ or from $\tilde{Y}_2$ is called a *mono cusp.* A cusp of $\tilde{Y}$ containing a cusp of $\tilde{Y}_1$ and also a cusp of $\tilde{Y}_2$ is called a *double cusp.*

The defining property for the set of complete cusps $\mathcal{C}_i$ is the following. There is a cusp in $\mathcal{C}_i$ corresponding to C_i^* unless the generalized cusp of $\tilde{Y}$ containing C_i^* is a double cusp which is the union of two cusps $C_1^* \subset \tilde{Y}_1, C_2^* \subset \tilde{Y}_2$ (one of which is C_i^*); and $rank(C_i^*) < \max(rank(C_1^*), rank(C_2^*))$. Thus a cusp will not be thickened (is in $\mathcal{C}_i$) only if it is glued in $\tilde{Y}$

to a cusp of strictly larger rank. For example a rank-1 cusp embedded in a rank-2 cusp will not be thickened.

We denote by C_i^+ the cusp in Y_i^+ corresponding to the cusp $C_i^* \subset Y_i$. If K is sufficiently large we show that, for every double cusp, if $rank(C_2^*) \leq rank(C_1^*)$, then $C_2 \subset C_1^+$ (and the corresponding statements with the subscripts 1 and 2 swapped.) Because $C_2 \subset C_1^+$ we can glue C_1^+ onto C_1^* in a way that is metrically compatible with the gluing of C_1^* and C_2^*. We then argue directly that this gives a simple gluing of $\tilde{Y}_1 \cup C_1^+$ to $\tilde{Y}_2$. Doing this to each cusp of $\tilde{Y}$ in turn shows that the simple gluing $\tilde{Y} = \tilde{Y}_1 \cup \tilde{Y}_2$ extends to a simple gluing $\tilde{Y} = \tilde{Y}_1^+ \cup \tilde{Y}_2^+$. Now we give the details.

How to choose K. By 8.1 $\Gamma_0 = \Gamma_1 \cap \Gamma_2$ is geometrically finite. The manifold $\tilde{N} = \mathbb{H}^n/\Gamma_0$ is a cover of N. Since $\Gamma_0 \subset \Gamma_i$ there is a corresponding cover $R_i \to M_i$. Because M_i is convex, the map $\rho_i : M_i \to N$ is covered by an isometric embedding $R_i \hookrightarrow \tilde{N}$ and $\tilde{N} = Th_\infty(R_i)$ is a thickening of R_i.

Since $\tilde{N}$ is geometrically finite it has finitely many cusps. Fix a cusp $P \subset \tilde{N}$. Since R_i is a convex submanifold of $\tilde{N}$ it follows from 2.4 that each component of $R_i \cap P$ is convex. Since $\tilde{N}$ is a thickening of R_i it follows that $\pi_1\tilde{N} \cong \pi_1 R_i$. Thus there is a unique component, $P_i \subset R_i \cap P$, which is a cusp of R_i and, if there are any other components, they are balls.

Now P_i is a cover of some cusp $C_i \subset M_i$. If we relabel so that $rank(C_2) \leq rank(C_1)$ then compatibility of cusps implies that P_2 is a finite cover of C_2. Since M_2 has thin cusps, P_2 is a thin cusp. For $i \in \{1, 2\}$ we choose a geodesic ray $\gamma_i \subset P_i$, starting on the cusp boundary and exiting the cusp P_i, and that is orthogonal to the foliation of P by horomanifolds. Thus the distance between γ_1 and γ_2 decreases exponentially as they go out into the cusp P. Since P_2 is thin the distance of every point in P_2 from γ_2 (and therefore from γ_1) also decreases exponentially as the point goes out into the cusp.

It is now easy to see that if K is sufficiently large then $P_2 \subset Th^{rel}_{(K-\kappa)}(P_1)$. The same argument shows we may choose an even larger thickening constant, K^+, such that $Th^{rel}_K(P_2) \subset Th^{rel}_{K^+}(P_1)$. We choose K and K^+ large enough that this (or the corresponding statement with 1 and 2 switched) holds for each of the finitely many cusps of $\tilde{N}$. In particular, if $rank(C_1) = rank(C_2)$, both inclusions are satisfied.

How to choose $\mathcal{C}$. Observe that as k is increased $Y_i = Th^{rel}_k(M_i)$ gets larger. A large choice of k fattens Y_i a lot, and so its image under ρ_i spill might out into the cusps of N. If k is very large then there are probably many components of $\rho_i^{-1}(\mathcal{C}) \subset Y_i$ which are balls. However, by making the cusps of $\mathcal{C}$ small enough, we arrange that for $k = K^+$ (and hence for $k = K$ also) every component $\rho_i^{-1}(\mathcal{C}) \subset Th^{rel}_k(M_i)$ is a cusp of $Th^{rel}_k(M_i)$.

Making the cusps $\mathcal{C}$ smaller does not affect our choices of K or K^+. This is because K, K^+ just have to be chosen large enough that a certain relative thickening contains certain cusps. But as cusps are made smaller, relative thickenings get bigger. Thus K and K^+ continue to have the defining property as the cusps are made smaller.

With these choices of K, K^+ and $\mathcal{C}$ we now define the virtual simple gluings we will use. Let $Y_i = Th^{rel}_K(M_i)$ and $Z_i = Th^{rel}_{K^+}(M_i)$. There is a virtual simple gluing $\tilde{Z} = \tilde{Z}_1 \cup \tilde{Z}_2$. Let $\tilde{Y}_i$ be the cover corresponding to $\tilde{Z}_i$. Then $\tilde{Z}_i$ is a thickening of $\tilde{Y}_i$. We thus also obtain a virtual simple gluing $\tilde{Y} = \tilde{Y}_1 \cup \tilde{Y}_2 \subset \tilde{Z}$. Notice that the $\tilde{Y}$ gluing involves larger covers of Y_i than would be obtained by just applying the virtual simple gluing theorem to Y_1 and Y_2.

Let C^* be a generalized cusp of $\tilde{Y}$. Suppose that C_i^* is a component of $\tilde{Y}_i \cap C^*$. By our choice of the cusps $\mathcal{C}$, it follows that C_i^* is a cusp of $\tilde{Y}_i$. Since C_i^* is convex we may regard $\pi_1 C_i^*$ as a subgroup of $\pi_1 C^*$. Since $\tilde{Y}_i$ is convex, distinct cusps of $\tilde{Y}_i$ correspond to disjoint $\pi_1\tilde{Y}_i$-orbits of parabolic fixed points. The holonomy of $\tilde{Y}_i \cap C^*$ factors through the holonomy of C^*. The holonomy of C^* fixes a unique point at infinity. Thus $\tilde{Y}_i$ can't have two cusps contained in C^*. Hence C^* is either a mono cusp or a double cusp as asserted in the sketch. The next step is to thicken certain cusps. It is easy to thicken a mono cusp. We must study double cusps.

Claim. Every double cusp in $\tilde{Y}$ can be isometrically embedded into a complete cusp.

Proof of Claim. Suppose that $C^* = C_1^* \cup C_2^* \subset \tilde{Y}$ is a double cusp. Without loss assume that $rank(C_2^*) \leq rank(C_1^*)$. Let $\tilde{C}_1^*, \tilde{C}_2^*$ denote universal covers. Since they are convex we may identify them with their images under the developing map in $\mathbb{H}^n$. These images can be chosen to have the same parabolic fixed point in such a way that they are covers of the isometric immersions $C_i^* \hookrightarrow N$. Let $C_2 \subset C_2^*$ be the corresponding cusp in $\tilde{M}_2$; thus $\tilde{C}_2 \subset \tilde{C}_2^*$. By taking universal covers one sees that the choice of K implies that $\tilde{C}_2 \subset \tilde{C}_1^*$. It follows that $C_2 \subset C_1^*$. Let D denote the cusp in $\tilde{Z}_1$ containing C_1^*. Then by choice of K^+ we have $C_2^* \subset D$. We thus obtain a locally isometric immersion $\rho : C^* \to D$. Restricting this gives embeddings $\rho_i : C_i^* \hookrightarrow D$. We now argue that ρ is injective.

If ρ is not injective then there are $x_i \in C^* \setminus (C_1^* \cap C_2^*)$ with $\rho_1(x_1) = \rho_2(x_2)$. Since D is a thickening of C_1^* we get that $\pi_1 C_2^* \subset \pi_1 D = \pi_1 C_1^*$. Taking universal covers we get a locally isometric immersion $\tilde{\rho} : \tilde{C}^* \to \tilde{D} \subset \mathbb{H}^n$ covering ρ. Since $\pi_1 C_1^* = \pi_1 D$ it follows that $\tilde{C}^*$ contains only one copy of $\tilde{C}_1^*$. Also by convexity, for each copy of $\tilde{C}_2^* \subset \tilde{C}^*$, we get $\tilde{\rho}(\tilde{C}_2^*) \cap \tilde{\rho}(\tilde{C}_1^*)$ is convex. It follows that $\tilde{\rho}(\tilde{C}_1^* \cap \tilde{C}_2^*) = \tilde{\rho}(\tilde{C}_1^*) \cap \tilde{\rho}(\tilde{C}_2^*)$. Thus there are no x_1, x_2 as above. This implies $C_1^* \cup C_2^*$ is embedded in D. This proves claim. □

To complete the proof we now thicken certain generalized cusps of $\tilde{Y}$ as described in the sketch. This is can be done by thickening one generalized cusp, $C^* \subset \tilde{Y}$, at a time. If C^* is a mono cusp, after relabeling, we may assume it is a cusp $C_1^* \subset Y_1$. The thickening consists of replacing C_1^* by another cusp with the same boundary. Since this is disjoint from $\tilde{Y}_2$ the simple gluing clearly extends over this thickening. Otherwise we have a double cusp $C^* = C_1^* \cup C_2^*$. By the claim this is embedded in a complete cusp, D. After relabeling we assume $rank(C_2^*) \leq rank(C_1^*)$. We also saw that D is a thickening of C_1^*. Let $C_1 \subset \tilde{M}_1$ be the cusp corresponding to C_1. Then $Th_K^{rel}(C_1) \subset D$. We replace C_1^* by $C_1^+ = Th_K(C_1) \subset D$. Since $C^* \subset D$ we can glue C_1^+ isometrically onto Y to obtain a hyperbolic manifold. This manifold is the result of the simply gluing of $\tilde{Y}_1 \cup C_1^+$ and $\tilde{Y}_2$.

Finally we must check that $\tilde{Y}^+$ contains a κ-neighborhood of the cusps of $\tilde{M}$. For a cusp of $\tilde{M}$ contained in a mono cusp of $\tilde{Y}$ this is clear. For a double cusp it follows from the fact that $Th_\kappa(C_2) \subset Th_K(C_1)$. □

6. *Some Constructions of Hyperbolic Manifolds.*

In this section we use the convex combination theorem to give constructions of geometrically finite hyperbolic manifolds in dimensions bigger than 3. The basic idea is to take two hyperbolic manifolds of dimensions m, n each of which contains a copy of the same totally geodesic submanifold of dimension p and then glue the manifolds along the submanifold and thicken to get a convex hyperbolic manifold of dimension $m+n-p$. The case that $m = n = p+1$, and the submanifold in question is the boundary of the given manifolds, was described by Bowditch and Mess in [**6**] using their generalization of Thurston's *bending construction.*

Consider a hyperbolic m-manifold, $M = \mathbb{H}^m/G$, where G is a discrete subgroup of $Isom(\mathbb{H}^m)$. We first describe a way to thicken M to obtain a hyperbolic n-manifold with $n > m$. Choose an isometric embedding of $\mathbb{H}^m$ as a totally geodesic subspace of $\mathbb{H}^n$. Let $stab(\mathbb{H}^m) \subset Isom(\mathbb{H}^n)$ be the subgroup which preserves $\mathbb{H}^m$. Let $O(n-m) \subset Isom(\mathbb{H}^n)$ denote the subgroup which acts trivially on $\mathbb{H}^m$. Then there is a splitting:

$$stab(\mathbb{H}^m) \cong Isom(\mathbb{H}^m) \oplus O(n-m).$$

We may thus regard $Isom(\mathbb{H}^m)$ (and thus G) as a subgroup of $Isom(\mathbb{H}^n)$. Given $R > 0$ let

$$N = \{ \ x \in \mathbb{H}^n \ : \ d(x, \mathbb{H}^m) \leq R \ \}.$$

Define $T(M; n, R) = N/G$. This is a convex hyperbolic n manifold with strictly convex boundary. There is a projection $\pi : \mathbb{H}^n \to \mathbb{H}^m$ given by the nearest point retraction. This map is G-equivariant thus we get a map $p : N/G \to M$ which is a Riemannian submersion and is a disc bundle over M.

Example 1. For $i \in \{1, 2\}$ suppose that M_i is a closed hyperbolic 3-manifold which contains a simple closed geodesic C_i of length ℓ such that the holonomy around C_i is a pure translation. We suppose that C_i has a tubular neighborhood V_i in M_i of radius κ, the thickening constant.

Consider the hyperbolic 5-manifolds $M_i^+ = T(M_i; 5, \kappa) = N_i/G_i$ where N_i is a κ-neighborhood of a hyperbolic 3-space $\mathbb{H}^3_i \equiv \tilde{M}_i$ in $\mathbb{H}^5$. We may choose $\mathbb{H}^3_1$ to be orthogonal to $\mathbb{H}^3_2$ and intersect along a geodesic which covers C_1 in M_1 and C_2 in M_2. Then $\tilde{V} = N_1 \cap N_2$ is a neighborhood of $\mathbb{H}^3_1 \cap \mathbb{H}^3_2$. Let $V_i \cong S^1 \times D^4$ denote the image of $\tilde{V}$ in M_i^+. Let M^+ be the non-convex hyperbolic 5-manifold obtained by gluing M_1^+ to M_2^+ by identifying V_1 with V_2 in a way which is covered by the identifications between N_1 and N_2. By the convex combination theorem M^+ has a convex thickening.

Example 2. For $i \in \{1, 2\}$ suppose that M_i is a closed hyperbolic 3-manifold which contains a totally-geodesic 2-sided surface F_i and assume that F_1 is isometric to F_2. Also assume that M_i contains a tubular neighborhood of F_i of radius the thickening constant κ. Now consider the hyperbolic 4-manifolds $M_i^+ = T(M_i; 4, \kappa)$. These may be glued by isometrically identifying neighborhoods of F_1 and F_2 to obtain a hyperbolic 4-manifold, $W = M_1^+ \cup M_2^+$. This is done in such a way that M_1 and M_2 are orthogonal in W. Clearly W has a spine which is obtained by gluing M_1 to M_2 by identifying F_1 with F_2. By the convex combination theorem W has a convex thickening.

Example 3. This time we will construct a hyperbolic 4-manifold by gluing a hyperbolic surface with geodesic boundary component, C_1, along C_1 to a geodesic, C_2, of the same length in a closed hyperbolic 3-manifold and thickening. In order that the resulting 4-manifold have a convex thickening it suffices that C_1 and C_2 both have tubular neighborhoods of radius κ in their respective manifolds. In addition C_2 should be a pure translation.

This example may be modified to allow the holonomy along C_2 to have a small non-zero rotational part, by deforming the holonomy of the surface from a subgroup of $Isom(\mathbb{H}^2)$ into a nearby subgroup of $Isom(\mathbb{H}^4)$ so that C_1 and C_2 have the same holonomy.

Definition. We now generalize the idea that the holonomy along a geodesic is a pure translation. Suppose that P is a geodesically-complete totally-geodesic p-submanifold of a hyperbolic m-manifold M. We say that *the normal bundle of P in M has trivial holonomy* if, whenever a tangent vector in M (based at an arbitrary point in P and orthogonal to P) is parallel translated around an arbitrary loop in P, then the the vector returns to itself. An equivalent formulation is that if we identify the universal cover of M with $\mathbb{H}^m$ and some pre-image of P with a subspace $\mathbb{H}^p \subset \mathbb{H}^m$ then the holonomy of P is contained in the subgroup $Isom(\mathbb{H}^p) \subset stab(\mathbb{H}^m)$.

Theorem 6.1 (gluing hyperbolic manifolds along isometric submanifolds). *For $i \in \{1, 2\}$ suppose that M_i is a convex hyperbolic m_i-manifold. Suppose that $p < \min(m_1, m_2)$ and P_i is a closed, hyperbolic, totally-geodesic p-submanifold in M_i. Suppose that P_i has a tubular neighborhood in M_i of radius at least the thickening constant κ, and the normal bundle of P_i in M_i has trivial holonomy. Suppose that $f : P_1 \to P_2$ is an isometry. Let M be the path-metric space obtained by gluing M_1 to M_2 by using f to identify P_1 with P_2. Then M has an path-isometric embedding onto a spine of a convex hyperbolic $n = (m_1 + m_2 - p)$-manifold.*

Proof. Let $H_i \equiv \mathbb{H}^{m_i} \subset \mathbb{H}^n$ be a totally geodesic subspace with H_1 orthogonal to H_2. We can identify the universal cover of P with $H_1 \cap H_2$ in a way that is compatible with identifications of the universal cover of M_i with H_i. Then $Y_i = T(M_i; n, \kappa)$ is a thickening whose universal cover, $\tilde{Y}_i$, is identified with the κ-neighborhood of H_i. Since the normal bundle of P has trivial holonomy in M_1 and M_2 it follows that $\pi_1 P$ preserves both H_1 and H_2. Then $Q = (\tilde{Y}_1 \cap \tilde{Y}_2)/\pi_1 P$ embeds isometrically in Y_i as a thickening of P. Let $Y = Y_1 \cup Y_2$ be the simple gluing obtained by identifying the two copies of Q. The convex combination theorem applied to $M_1 \cup M_2 \subset Y_1 \cup Y_2$ then gives the result. □

Theorem 6.2 (virtual gluing hyperbolic manifolds along isometric submanifolds). *For $i \in \{1,2\}$ suppose that M_i is a convex hyperbolic m_i-manifold. Suppose that $p < \min(m_1, m_2)$ and P_i is a closed, hyperbolic, totally-geodesic p-submanifold in M_i, and suppose that the normal bundle of P_i in M_i has trivial holonomy. Finally, suppose that $f : P_1 \to P_2$ is an isometry. Then there are finite covers $\tilde{M}_i$ of M_i and lifts $\tilde{P}_i$ of P_i to $\tilde{M}_i$ with the following property. Let $\tilde{M}$ be the space obtained by gluing $\tilde{M}_1$ to $\tilde{M}_2$ by using f to identify $\tilde{P}_1$ with $\tilde{P}_2$. Then $\tilde{M}$ has a convex thickening which is a hyperbolic $(m_1 + m_2 - p)$-manifold.*

Proof. By 7.5 $\pi_1 P_i$ is a separable subgroup of $\pi_1 M_i$. It follows there is a finite cover, $\tilde{M}_i$ of M_i and a lift of P_i to $\tilde{M}_i$ which has a tubular neighborhood of radius κ. The result follows from the previous theorem. □

7. *Subgroup Separability.*

Two groups G_1 and G_2 are *commensurable* if there is a group H which is isomorphic to finite index subgroups of both G_1 and G_2. Two path-connected topological spaces X, Y are *commensurable* if there are finite sheeted covers $\tilde{X}, \tilde{Y}$ which are homeomorphic. Clearly commensurable spaces have commensurable fundamental groups.

A subgroup H of a group G is *separable in G* if for every $g \in G \setminus H$ there is a subgroup of finite index $K < G$ such that $H \leq K$ and $g \notin K$. The group G is *subgroup separable* if every subgroup is separable and is *LERF* if every finitely generated subgroup is separable.

Definition. Suppose that F is a compact, connected, surface with non-empty boundary and $\chi(F) < 0$. Let $\{\partial_i F\}_{1 \leq i \leq n}$ denote the boundary components of F and let $\{T_i\}_{1 \leq i \leq n}$ denote a collection of distinct tori. The *tubed surface*, X, obtained from F is the 2-complex X obtained obtained by homeomorphically identifying each component $\partial_i F$ with an essential simple closed curve on the torus T_i.

Observe that $\pi_1 X$ is a topological realization of a graph of groups with cyclic edge groups and with vertex groups that are either $\mathbb{Z}^2$ or finitely generated free groups. It follows that X is a $K(\pi, 1)$.

Lemma 7.1 (tubed surface is LERF). *If X is a tubed surface then $\pi_1 X$ is LERF.*

Proof. We first show that all tubed surfaces are commensurable. Let A denote the compact surface obtained by deleting the interior of a disc from a torus. It is easy to check that up to commensurability every compact, connected, surface with negative Euler characteristic and non-empty boundary is commensurable with A. Let Y denote the tubed surface obtained from A. It follows that every tubed surface is commensurable with with Y. Thus all tubed surfaces have commensurable fundamental groups.

R. Gitik proved in theorem (4.4) of [**17**] that an amalgam of a free group, F, and a LERF group, H, is again LERF, provided the amalgamating subgroup is a maximal cyclic subgroup of F. It is easy to see that $\mathbb{Z} \oplus \mathbb{Z}$ is LERF and that $\pi_1 \partial A$ is a maximal cyclic subgroup of $\pi_1 A$. It follows that $\pi_1 Y$ is LERF. Scott shows in lemma (1.1) of [**34**] that the property of being LERF is a commensurability invariant and it follows that the fundamental group of every tubed surface is LERF. □

Definition. If X is a path connected topological space and G is a subgroup of $\pi_1(X)$ we say that *G is embedded in X* if there is a path-connected subspace $Y \subset X$ such that $incl_* : \pi_1 Y \to \pi_1 X$ is injective with image G. We say that *G is virtually embedded in X* if it is embedded in some finite cover of X. This means there is a finite cover $p : \tilde{X} \to X$ and a path-connected subspace $Y \subset \tilde{X}$ such that $incl_* : \pi_1 Y \to \pi_1 X$ is injective with image G.

If $n \geq 5$ and X is an n-manifold then it follows from general position that every subgroup is embedded. In [**34**] Scott introduced the notion of virtually embedded for surface subgroups (but he used the term *almost geometric* which, in the context of 3-manifolds, has certain connotations we prefer to avoid) and used subgroup separability to prove that finitely generated subgroups of surface groups are virtually embedded. Scott's result extends to separable subgroups of 3-manifolds:

Theorem 7.2 (virtually embedded subgroups). *Suppose that M is a connected irreducible 3-manifold and $G < \pi_1 M$ is a finitely generated separable subgroup. Then G is a virtually embedded subgroup.*

Proof. The following argument is standard. Let $p_G : \tilde{M}_G \to M$ be the cover corresponding to G. By Scott's compact core theorem, [**25**], there is a compact submanifold $Y \subset \tilde{M}_G$ such that the inclusion, $\iota : Y \to \tilde{M}_G$, induces an isomorphism of fundamental groups. We may choose a triangulation of M such that Y is a finite sub-complex of the induced triangulation of $\tilde{M}_G$. We will show how to construct a tower of finite covers (as in the proof of the loop theorem)

$$p_n : \tilde{M}_{n+1} \to \tilde{M}_n$$

with $\tilde{M}_0 = M$ and such that the map $f_0 = p_G \circ \iota : Y \to M$ lifts to each cover $f_n : Y \to \tilde{M}_n$. The singular set of f_n is

$$S(f_n) \ = \ \{ \ x \in Y \ : \ \#|f_n^{-1}(f_n(x))| \ > \ 1 \ \}.$$

Observe that $S(f_n)$ is a sub-complex of Y and $S(f_{n+1}) \subset S(f_n)$. We claim if $S(f_n) \neq \phi$ then the cover p_n may be chosen such that $S(f_{n+1})$ is a proper sub-complex of $S(f_n)$. Since Y is a finite complex it then follows that for some $n \geq 0$ that $S(f_n)$ is empty. Then f_n is a π_1-injective embedding of Y into the finite cover $\tilde{M}_n$ which proves the theorem.

To prove the claim, suppose that $a, b \in Y$ are distinct points and $f_n(a) = f_n(b)$. Let γ be a path in Y from a to b. Then $\alpha = [f_n \circ \gamma] \in \pi_1(\tilde{M}_n) \leq \pi_1(M)$, and since Y is a subspace of a covering of M it is clear that α is non-trivial. Since G is separable there is a finite index subgroup $H < \pi_1(M)$ which contains G but does not contain α. The subgroup $H \cap \pi_1 \tilde{M}_n$ has finite index in $\pi_1(\tilde{M}_n)$. Let $p_{n+1} : \tilde{M}_{n+1} \to \tilde{M}_n$ be the cover corresponding to this subgroup. It is clear that $f_n : Y \to \tilde{M}_n$ lifts to $f_{n+1} : Y \to \tilde{M}_{n+1}$ and, since α does not lift, that $f_{n+1}(a) \neq f_{n+1}(b)$. Thus $S(f_{n+1})$ is a proper subset of $S(f_n)$. This proves the claim. □

From this one recovers the following well-known result:

Corollary 7.3 (separable surface subgroups). *Suppose that $f : S \to M$ is a continuous map of a closed surface with $\chi(S) \leq 0$ into an irreducible 3-manifold M and suppose that $f_* : \pi_1 S \to \pi_1 M$ is injective. Suppose that $f_*(\pi_1 S)$ is a separable subgroup of $\pi_1 M$. Then* *there is a finite cover $p : \tilde{M} \to M$ and an embedding $g : S \to \tilde{M}$ such that $p \circ g$ is homotopic to f.*

Proof. By the theorem $f_*(\pi_1 S)$ is virtually embedded so there is a finite cover $p : \tilde{M} \to M$ and an incompressible compact submanifold $Y \subset \tilde{M}$ such that $p_*(\pi_1 Y) = f_*(\pi_1 S)$. Since $\pi_1 Y$ is a surface group Y must have non-empty boundary. By the equivariant sphere theorem, $\tilde{M}$ is also irreducible. Thus we may cap off any sphere components in the boundary of Y by balls. It then follows that Y is irreducible. It is now a theorem of Waldhausen that $Y \cong S \times [0, 1]$. Since M is irreducible it is a $K(\pi_1 M, 1)$, hence the maps $p|S \times 0 \to M$ and $f : S \to M$ are homotopic. □

Corollary 7.4. *Suppose that M is a 3-manifold which is homotopy equivalent to a tubed surface. Then every finitely generated subgroup of $\pi_1 M$ is virtually embedded.*

Theorem 7.5 (totally geodesic is separable). *Suppose that M is a totally geodesic hyperbolic k-manifold immersed in a convex hyperbolic n-manifold N with $k < n$. Also suppose that $\pi_1 N$ is finitely generated. Then $\pi_1 M$ is a separable subgroup of $\pi_1 N$.*

Proof. This can be proved by an extension of the method Long used in the case $k = 2, n = 3$ see [**27**]. □

8. *Surfaces in 3-Manifolds.*

We start with two results that will be needed for some of the 3-manifold applications.

Theorem 8.1 (intersections of GF are GF). *If Γ_1 and Γ_2 are geometrically finite subgroups of a discrete subgroup of $Isom(\mathbb{H}^n)$ then $\Gamma_1 \cap \Gamma_2$ is geometrically finite.*

Proof. This result is theorem 4 in Susskind and Swarup, [**37**]. In dimension at most 3 it is well known, see for example theorem (3.15) of [**32**]. In dimensions 4 and higher there are examples of two geometrically finite groups (which generate a non-discrete group) whose intersection is not finitely generated, see [**36**]. □

A group G has the *finitely generated intersection property* or *FGIP* if the intersection of two finitely generated subgroups is always finitely generated. The following result in this form is due to Susskind [**38**]; see also Hempel [**24**] and [**32**] corollary (3.16).

Theorem 8.2 (GF infinite covolume implies FGIP). *If G is a geometrically finite Kleinian group of infinite co-volume then G has the FGIP.*

The fundamental group of a hyperbolic surface bundle does not have FGIP, so the hypothesis of infinite co-volume is necessary.

The hypothesis of the convex combination theorem that one can thicken without bumping means that there are restrictions on the cusps of the manifolds to be glued. In particular two rank-1 cusps in the same rank-2 cusp of a 3-manifold cannot be thickened without bumping unless they are parallel.

Here is an algebraic viewpoint. Rank-1 cusps give $\mathbb{Z}$ subgroups, and two non-parallel cusps in the same rank-2 cusp will generate a $\mathbb{Z} \oplus \mathbb{Z}$ group in the group, G, generated by the two subgroups. Thus G is not an amalgamated free product of the subgroups.

Suppose M_1 and M_2 are 3-manifolds with rank-1 cusps which intersect in M. We may first glue rank-2 cusps onto each rank-1 cusp of M_1 and M_2.

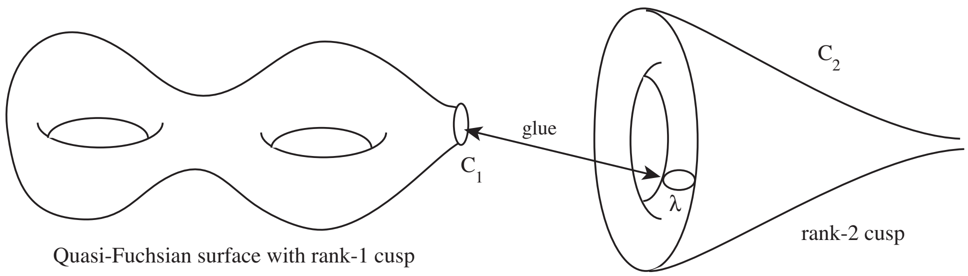

FIGURE 9. *Gluing on a rank-2 cusp.*

This corresponds to an amalgamated free product of a $\mathbb{Z}\oplus\mathbb{Z}$ with $\pi_1 M_i$. This produces two new geometrically finite manifolds M_1^+, M_2^+ which may now be glued. The process of gluing a rank-2 cusp onto a rank-1 cusp can be done with the convex combination theorem:

THEOREM 8.3 (adding a rank-2 cusp). *Suppose that M is a convex hyperbolic 3-manifold and that $f : N \to M$ is a locally-isometric immersion of a geometrically finite hyperbolic 3-manifold N. Suppose that C_2 is a rank-2 cusp of M and C_1 is a component of $f^{-1}(C_2)$ which is a thin rank-1 cusp of N. Then there is a finite cover $\tilde{C}_2$ of C_2 and a geometrically finite 3-manifold $N^+ = N \cup \tilde{C}_2$ with a convex thickening. Furthermore $N \cap \tilde{C}_2 = C_1$, where $C_1 \subset N$ is identified with a subset of $\tilde{C}_2$ using a lift of $f|C_1$.*

Proof. This follows from 5.2. □

REMARK. It is easy to extend this result to the setting where one has immersions of finitely many geometrically finite manifolds, $f_i : N_i \to M$, and finitely many rank-2 cusps $C_{2,j} \subset M$, and it is required to glue some of the rank-1 cusps in $f_i^{-1}(\cup_j C_{2,j}) \subset N_i$ to cyclic covers of $C_{2,j}$ in such a way that if more than one rank-1 cusp is glued onto the same rank-2 cusp then the rank-1 cusps are glued along parallel curves sufficiently far apart in the boundary of the rank-2 cusps.

COROLLARY 8.4 (GF tubed surfaces). *Suppose M is a compact 3-manifold whose interior admits a complete hyperbolic metric. Suppose that S is a compact, connected surface with $\chi(S) < 0$ and $f : S \to M$ is π_1-injective such that each component of ∂S is mapped into some torus boundary component of M. Suppose that $f_*(\pi_1 S)$ is a geometrically-finite subgroup of $\pi_1 M$. Let S^+ denote the tubed surface obtained by gluing one torus onto each boundary component of S. Then f extends to a π_1-injective map $f : S^+ \to M$ such that such that $f_*(\pi_1 S^+)$ is a geometrically-finite subgroup of $\pi_1 M$.*

Proof. Let N be the convex core of $\mathbb{H}^3/f_*(\pi_1 S)$. Then N is geometrically finite and there is an isometric immersion $N \to M$ coming from the fact that $\pi_1 N \subset \pi_1 M$. We apply the previous theorem to add rank-2 cusps to the boundary components of S one at a time. □

DEFINITION. Suppose that N is a convex hyperbolic 3-manifold of finite volume which equals its convex core and that N has spine a tubed surface S^+. In what follows we will the use of the term *tubed surface* to refer to either the 2-complex S^+, or to the geometrically finite hyperbolic 3-manifold N and write this as $Core(S^+)$.

Suppose that $\tilde{N} \to N$ is a finite cyclic cover to which S lifts, then $\tilde{N}$ contains a closed surface, $2S$, homeomorphic to the double of S along its boundary. If the cover has degree bigger than 1 then this surface may be chosen such that it is not homotopic into the boundary of a compact core of $\tilde{N}$. Clearly $2S$ is π_1-injective in $\tilde{N}$. In particular 8.4 implies that $\pi_1 M$ contains a non-peripheral surface group.

Let $N = Core(S^+)$. In [**7**] it was observed that given a rank-2 cusp, C, of N every sufficiently large Dehn-filling of C can be given a Riemannian metric of negative sectional curvature which agrees with the original metric outside C. Suppose $f : N \to M$ is a local isometry and N^+ is a Dehn-filling of N along C. The same filling done on the corresponding cusp of M then gives a local isometry $N^+ \to M^+$ of Dehn-filled manifolds. If this is done to all the cusps of N then, since N^+ is convex, this map is π_1-injective. This was the main step in the proof of the main theorem of [**7**]. It also leads to a quick proof of a virtual-Haken Dehn-filling result of the type in [**8**].

The following is the main tool we need for studying immersed boundary slopes.

Theorem 8.5 (gluing tubed surfaces). *Suppose that W is a convex hyperbolic 3-manifold and for $i \in \{1,2\}$ that M_i hyperbolic 3-manifold which is a tubed surface and $f_i : (M_i, m_i) \to (W, w_0)$ is a local isometry. Set $\Gamma_i = \pi_1(M_i, m_i)$. Then there are finite covers $p_i : \tilde{M}_i \to M_i$ and a simple gluing $\tilde{M} = S(f_1 \circ p_1, f_2 \circ p_2)$ of $\tilde{M}_1, \tilde{M}_2$ such that $\tilde{M}$ has a convex thickening. Furthermore $\Gamma_i^- = p_{i*}(\pi_1(\tilde{M}_i, \tilde{m}_i))$ contains $\Gamma = f_{1*}(\Gamma_1) \cap f_{2*}(\Gamma_2)$. Also $\pi_1 \tilde{M}$ is a free product of Γ_1^- and Γ_2^- amalgamated along Γ.*

Proof. Since tubed surfaces are geometrically finite then, by theorem 8.1, the group Γ is geometrically finite and thus finitely generated. By lemma 7.1 Γ_1 and Γ_2 are LERF. Thus Γ is separable in both Γ_1 and Γ_2. Define $N_i = Th_\kappa(M_i)$ then f_i extends to a local isometry $f_i : N_i \to W$. Applying the virtual simple gluing theorem to N_1, N_2 mapped into W, it follows that there are finite covers $p_i : Y_i \to N_i$ which have a simple gluing Y. Let $p_i| : \tilde{M}_i \to M_i$ be the restriction of the covering p_i. We can now apply the convex combination theorem to $\tilde{M}_1 \subset Y_1$ and $\tilde{M}_2 \subset Y_2$ and deduce that $\tilde{M}$ has a convex thickening. □

The next result is similar to, but has a stronger conclusion than, a special case of corollary 5 of Gitik's paper [**19**], and also (with a little work) to the combination theorem of Bestvina-Feighn [**4**].

Corollary 8.6 (amalgamating QF subgroups). *Let N be a closed hyperbolic 3-manifold. Suppose Γ_1, Γ_2 are two quasi-Fuchsian subgroups of $\pi_1 N$ each isomorphic to the fundamental group of a closed surface with negative Euler characteristic. Suppose $\Gamma_0 = \Gamma_1 \cap \Gamma_2$ is not trivial, and has infinite index in both Γ_1 and in Γ_2. Then there is a compact, convex, hyperbolic 3-manifold M and a locally isometric immersion of M into N. Furthermore M has incompressible boundary, and for every $n > 0$ there is a finite cover, $\tilde{M}$, with $\beta_2(\tilde{M}) \geq n$.*

Proof. By 8.2 Γ_0 is finitely generated, thus separable in both Γ_1 and Γ_2 by Scott's theorem [**34**]. We can now apply the virtual amalgam theorem. Thus there are finite index subgroups $\Gamma_i' \subset \Gamma_i$ such that the subgroup $\Gamma \subset \pi_1 N$ generated by Γ_1' and Γ_2' is their amalgamated free product $\Gamma = \Gamma_1' *_{\Gamma_0'} \Gamma_2'$. The hypotheses imply that Γ_0' has infinite index in both Γ_1' and Γ_2'. Furthermore, $\Gamma_0' = \Gamma_1' \cap \Gamma_2'$ has finite index in Γ_0. Since Γ_0 is torsion-free and non-trivial, it follows that Γ_0' is non-trivial. Therefore this is a non-trivial amalgamated free product decomposition. The groups Γ_1', Γ_2' are surface groups and therefore freely indecomposable. A free product of freely indecomposable groups, amalgamated along a non-trivial subgroup, is freely indecomposable. Define M to be the convex core of $\mathbb{H}^3/\Gamma$. It follows that M has incompressible boundary. The proof of the virtual amalgam theorem shows that M is a thickening of a simple gluing $M_1 \cup M_2$, where M_i is a thickening of the convex core of $\mathbb{H}^3/\Gamma_i'$. Now $M_i \cong S_i \times [-1,1]$ and $M_1 \cap M_2$ does not separated the two boundary components of

M_i, for otherwise $\pi_1 M_i \subset \pi_1(M_1 \cap M_2)$, which contradicts we have a non-trivial amalgamated free product. Thus there is a properly embedded arc $\gamma_i \subset M_i \setminus (M_1 \cap M_2)$ which intersects $S_i \equiv S_i \times 0 \subset S_i \times [-1,1]$ once transversely. It follows that S_1 and S_2 are linearly independent in $H_2(M)$. Hence $\beta_2(M) \geq 2$. Given $n > 0$ it follows from [**10**] that there is a finite cover $\tilde{M}$ with $\beta_2(\tilde{M}) \geq n$. □

Corollary 8.7. *Suppose that M is a closed hyperbolic 3-manifold and S is a closed, connected surface with $\chi(S) < 0$. Suppose that $f : S \to M$ is π_1-injective and not homotopic to an embedding. Also suppose that $f_*(\pi_1 S)$ is a maximal surface subgroup of $\pi_1 M$. Then for all $n > 0$ there is a subgroup $G \subset \pi_1 M$ which is a geometrically finite subgroup of $\pi_1 M$ and such that $\beta_2(G) > n$ and G is freely indecomposable.*

Proof. Thurston showed in [**40**] that a surface subgroup of a closed hyperbolic 3-manifold is either geometrically finite, hence quasi-Fuchsian, or a virtual fiber. First we reduce to the case that S is quasi-Fuchsian. We are indebted to the referee for pointing out the following argument, which is shorter than the original. Suppose that S is a virtual fiber. There is a finite cover $\tilde{M} \to M$ and a pre-image, $\tilde{S}$, of S which is homotopic to a fiber of a fibration of $\tilde{M}$ over the circle. By [**31**] there is a quasi-Fuchsian surface, S', immersed in $\tilde{M}$ transverse to the suspension flow and $k \neq 0$ with $[S'] = k[\tilde{S}] \in H_2(\tilde{M})$. We may choose a maximum quasi-Fuchsian subgroup containing the given one. Suppose this surface is homotopic to an embedded surface. Then, since it's homology class is a multiple of the homology class of a fiber and it is embedded, it is isotopic to the fiber. This contradicts that it is quasi-Fuchsian. Thus we are reduced to the case that S is quasi-Fuchsian.

Since S is not homotopic to an embedding there are two conjugates, $A \neq B$ of $f_*(\pi_1 S)$ such that $C = A \cap B \neq 1$. Given a Kleinian group H we denote its limit set by $\Lambda(H)$. Since A is quasi-Fuchsian $\Lambda(A) \cong S^1$. The subgroup of $\pi_1 M$ which stabilizes $\Lambda(A)$ is a surface group. By maximallity this group is A. Since $A \neq B$ it follows that $\Lambda(A) \neq \Lambda(B)$. By [**37**] $\Lambda(A \cap B) = \Lambda(A) \cap \Lambda(B)$. This is therefore a proper non-empty subset of $\Lambda(A)$ and therefore not homeomorphic to a circle. Hence $C = A \cap B$ is not the fundamental group of a closed surface. Thus C has infinite index in A and in B. The result now follows by applying 8.6 to the groups A and B. □

Theorem 8.8 (gluing two rank-1 cusps: algebraic version). *Suppose that G is a torsion-free Kleinian group. Suppose that H is a geometrically finite subgroup of G and that P, P' are two maximal parabolic subgroups in H each of which is infinite cyclic. Suppose that T is a maximal rank-2 parabolic subgroup of G which contains P and there is $\gamma \in G$ such that $P' = \gamma P \gamma^{-1}$. Then there is $n > 0$ with the following property. Suppose that $\alpha \in T$ and the subgroup of T generated by P and α has finite index at least n. Set $\beta = \gamma\alpha$ then the subgroup of G generated by H and β is the HNN extension $H *_P$ amalgamated along P and P' given by $< H, \beta : \ \beta P \beta^{-1} = P' >$. Furthermore this group is geometrically finite.*

Proof. This follows from 8.9 by taking convex hyperbolic manifolds corresponding to the given groups. □

A geometric formulation of this result involves a generalization of the notion of *spinning* an annulus boundary component of a surface around a torus boundary component of a 3-manifold that contains the surface. This idea was introduced in [**16**] and used in [**7**], [**8**]. We recall the construction. Suppose S is a compact surface immersed in a 3-manifold M and that ∂S has two components, α and β, which both lie on a torus $T \subset \partial M$. We suppose the immersion maps α and β to the same loop on T. Attach an annulus A to the boundary of S and choose an immersion of the annulus into M so that it wraps some number $n \geq 0$ times around T. We

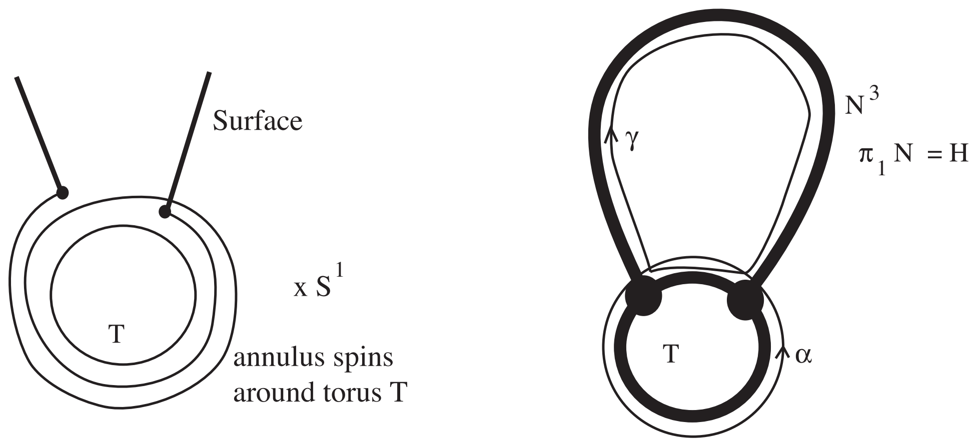

FIGURE 10. *Gluing two rank-1 cusps after spinning.*

describe this process by saying the two boundary components α and β have been glued using an annulus that spins n times around the torus T.

This operation can be extended to the case where we are given an immersion $f : N \to M$ of a 3-manifold N into M and there are two annuli $A_0, A_1 \subset \partial N$ which both map m times around the same annulus in some torus $T \subset \partial M$. To be more precise, if $f : N \to M$ is the immersion we require there is a homeomorphism $h : A_1 \to A_0$ such that $f|A_1 = (f|A_0) \circ h$. Then we glue $A \times [0,1]$ to N to obtain a 3-manifold N^+ by identifying $A \times i$ with A_i for $i = 0, 1$ and choose an extension of the immersion f over $A \times [0,1]$ which spins A round T a total of n times. The meaning of this last statement is the following. Identify A_0 with $S^1 \times [0, s]$ and T with $S^1 \times S^1$ in such a way that $f|A_0$ is given by the map $f(e^{i\theta}, x) = (e^{im\theta}, e^{ix})$. Then define $f : A \times [0,1] \to S^1 \times S^1$ by $f((e^{i\theta}, x), y) = (e^{i\theta}, e^{i(x+2n\pi y)})$. We describe N^+ and the resulting immersion $f : N^+ \to M$ as obtained from the immersion of N into M by *gluing two annuli in ∂N together after spinning n times around the torus T*. We can now state a geometric reformulation of the preceding result:

THEOREM 8.9 (gluing two rank-1 cusps: geometric version). *Suppose that M is a geodesically complete hyperbolic 3-manifold. Suppose that N is a convex hyperbolic 3-manifold with thin cusps and $f : N \to M$ is a locally isometric immersion. Suppose that C is a rank-2 cusp of M and C_1, C_2, are two components of $f^{-1}(C)$ which are thin rank-1 cusps. Suppose that $f_*\pi_1 C_1$ and $f_*\pi_1 C_2$ are the same subgroup of $\pi_1 C$. Then for $n > 0$ sufficiently large, the result of gluing C_1 to C_2 after spinning n times around C gives an isometric immersion of a geometrically finite 3-manifold P into M.*

Proof. Let $C^- \subset C$ be a smaller cusp such that C contains a κ-neighborhood of C^-. Let C_i^- be the component of $f^{-1}(C^-)$ contained in C_i. Set $Q = Th_\kappa^{rel}(N; \{C_1^-, C_2^-\})$. There is an extension of f to an isometric immersion $f : Q \to M$. Let C_i^+ be the cusp in Q which contains C_i^- and is a component of $f^{-1}(C)$. Then C_i^+ is a thin cusp.

Choose $\alpha \in \pi_1 C$ such that the subgroup of $\pi_1 C$ generated by α and $f_*\pi_1 C_i$ has finite index and this index is minimal. Given $n > 0$ let G_n be the finite index subgroup of $\pi_1 C$ generated by $f_*\pi_1 C_i$ and α^n. Let $D \to C$ be the cover corresponding to G_n. Observe that there is a lift $g_i : C_i^+ \to D$ of the map $f| : C_i^+ \to C$. Since C_i^+ is thin, for n sufficiently large, we may choose the lifts g_1 and g_2 to have disjoint images.

Let $R = Q \cup D$ be the hyperbolic 3-manifold obtained by using g_1 and g_2 to identify C_1^+ and C_2^+ with their images. This is a simple gluing of convex manifolds. It contains a κ-neighborhood of the simple gluing $N \cup D^-$ where C_1^-, C_2^- are glued to D^- by the restrictions of g_1, g_2. The convex combination theorem applied to these gluings implies $N \cup D^-$ has a convex thickening R.

Topologically we have glued two rank-1 cusps of N to $T^2 \times [0, \infty)$. Let H be the subgroup of $\pi_1 R$ generated by $\pi_1 N$, together with a loop consisting of an arc in N running from C_1 to C_2 and an arc in D. Let P be the cover of Q corresponding to this subgroup. Then P is the required manifold. □

For example, if N is a quasi-Fuchsian 3-manifold then P is just a thickened version (regular neighborhood) of the surface obtained by spinning two boundary components of the surface round T. The manifold P contains an accidental parabolic.

9. *Multiple Immersed Boundary Slopes.*

Definition. An *immersed slope* α on a torus T is an element of $(H_1(T;\mathbb{Z})-0)/(\pm 1)$. Thus α is uniquely represented by the homotopy class of an unoriented loop, also denoted α, in T which is not contractible. If α is homotopic in T to a simple closed curve then α is called a *slope*.

Suppose M is a compact 3-manifold and T is a torus boundary component of M. An *immersed boundary slope* is an immersed slope α on T such that there is a compact, connected, orientable surface S with non-empty boundary and a π_1-injective map $f : S \to M$ such that:

(i) for every boundary component $\beta \subset \partial S$ we have $f_*([\beta]) = [\alpha]$.
(ii) f is not homotopic rel ∂S to a map into ∂M.

If S is embedded we say that α is a *boundary slope*. If S does not lift to a fiber of a fibration of a finite cover $\tilde{M}$ then α is a *strict boundary slope*. If α is a boundary slope which is not a strict boundary slope and S is embedded then either S is a fiber of a fibration of M or else S separates M into two components each of which is a twisted I-bundle over a surface. In this case we say that S is a *semi-fiber*. If α is a *slope* and for some $n > 0$ the immersed slope $n \cdot \alpha$ is an immersed boundary slope we say that α is a *multiple immersed boundary slope* or MIBS.

Hatcher showed [**23**] that if M is compact and has boundary a torus then there are only finitely many boundary slopes. In [**3**] it was shown that if M is Seifert fibered then every immersed boundary slope is also a boundary slope and there are only two boundary slopes; the longitude (rationally null-homologous slope) and the slope of a fiber. The same result holds for MIBS. In this section we will show that if the interior of M admits a complete hyperbolic metric then every slope is a MIBS.

Here is an outline of what follows. Suppose S, S' are two quasi-Fuchsian surfaces each with two boundary components which correspond to different boundary slopes α, β. One would like to cut and cross join these surfaces along an arc connecting the two boundary components in each surface to produce a surface with boundary components $\alpha + \beta$ and $\alpha - \beta$. Finally one would like to tube two copies of this surface together using the $\alpha - \beta$ slopes to obtain an immersed surface with two boundary components both with slope $\alpha + \beta$.

There are many problems with trying to do this directly, the most obvious one being that in general one can't expect to produce a π_1-injective surface this way. Instead we first add rank-2 cusps to the surfaces. Then finite covers of these tubed surfaces can be glued to give a geometrically finite manifold with a torus boundary component having the property that every slope on this torus is homologous to a cycle on the union of the other torus boundary components. This 2-chain is represented by an incompressible surface. Finally two copies of this surface are glued by spinning and gluing all the boundary components except those on the chosen torus. This gives the desired immersion.

Culler and Shalen [**11**],[**13**] used character varieties to show:

Theorem 9.1. *Suppose that M is a compact, orientable 3-manifold with boundary a torus T and that the interior of M admits a complete hyperbolic structure of finite volume. Then there are two strict boundary slopes on T.*

Combining this with lemma (2.3) of [**7**] yields the following:

Addendum 9.2. *There are two incompressible, ∂-incompressible, quasifuchsian surfaces S_1, S_2 in M with boundary slopes $\alpha_1 \neq \alpha_2$.*

The existence of the following cover is perhaps independently interesting.

Theorem 9.3. *Suppose that M is a compact, orientable 3-manifold with boundary a torus T and that the interior of M admits a complete hyperbolic structure of finite volume. Suppose that $\partial M = T$ is a torus. Then there is an infinite cover $p : \tilde{M} \to M$ such that $\pi_1 \tilde{M}$ is finitely generated and there are distinct tori $T_1, \ldots, T_n \subset \partial \tilde{M}$ with the property that*

$$0 = incl_* : H_1(T_1; \mathbb{Q}) \to H_1(\tilde{M}, \cup_{i=2}^n T_i; \mathbb{Q}).$$

Proof. We will use the same notation for a manifold and for its interior. By addendum 9.2 there are two quasi-Fuchsian surfaces S_1, S_2 in M with distinct boundary slopes α_1, α_2. By corollary 8.4 we may construct two geometrically finite manifolds M_1, M_2 with spines that are the tubed surfaces obtained from them. Choose basepoints $m_i \in M_i$ which map to the same point in the cusp of M. Then apply theorem 8.5 to $\pi_1(M_1, m_1)$ and $\pi_1(M_2, m_2)$. This gives a geometrically finite manifold, N, obtained by gluing finite covers of M_1, M_2 along a submanifold. The manifold N has a thickening, $N \subset \tilde{M}$, which is a covering $p : \tilde{M} \to M$.

There is a torus $T_1 \subset \tilde{M}$ which corresponds to the rank-2 cusp obtained by gluing the rank-2 cusps of M_1 and M_2 that contain the basepoints. Hence for $i \in \{1, 2\}$ there is a component $\tilde{S}_i \subset p^{-1} S_i$ such that $T_1 \cap \tilde{S}_i \neq \phi$. Each component of $T_1 \cap \tilde{S}_i$ is a loop $\tilde{\alpha}_i$ that covers α_i. Clearly $\tilde{\alpha}_1, \tilde{\alpha}_2$ generate $H_1(T_1; \mathbb{Q})$. The result follows from consideration of the algebraic sum of m copies of $\tilde{S}_1$ and n copies of $\tilde{S}_2$. □

Theorem 9.4 (All slopes are MIBS). *Suppose that M is a compact, orientable 3-manifold with boundary a torus T and that the interior of M admits a complete hyperbolic structure of finite volume. Then there is a subgroup of finite index in $H_1(T; \mathbb{Z})$ such that every non-trivial element in this subgroup is an immersed boundary slope for a geometrically finite surface with exactly two boundary components. Thus every slope on T is a MIBS.*

Proof. We apply 9.3 to obtain a cover $p : \tilde{M} \to M$ and a torus $T_1 \subset \partial \tilde{M}$ with the property stated. Set $K = \ker[incl_* : H_1(T_1; \mathbb{Z}) \to H_1(\tilde{M}, \cup_{i=2}^n T_i; \mathbb{Z})]$, then K has finite index in $H_1(T_1)$. Let n be the index of $p_*(K)$ in $H_1(T)$. Given an essential loop α on T representing some slope, then $n \cdot \alpha$ lifts to a loop β on T_1 with $[\beta] \in K$. Thus there is a compact, connected, 2-sided, incompressible surface S properly embedded in $\tilde{M}$ such that $S \cap T_1 = \beta$. For $i = 0, 1$ let S_i be a copy of S and $\beta_i \subset \partial S_i$ the boundary component corresponding to β. For each boundary component $\gamma_0 \subset \partial S_0$ with $\gamma_0 \neq \beta_0$ attach the boundary components of an annulus to γ_0 and γ_1 to obtain a surface R with two boundary components β_0, β_1. Immerse R into $\tilde{M}$ identifying the two copies S_i with S and by spinning each annulus around the appropriate torus in $\partial \tilde{M}$ enough times to ensure the resulting immersed surface is π_1-injective and geometrically finite. That this can be done follows from 8.9. The composition $R \to \tilde{M} \to M$ is the desired surface. □

Acknowledgements. Both authors thank the Université de Rennes 1 and UCSB for hospitality and partial support during the time this paper was written. This work was also partially supported by NSF grants DMS0104039 and DMS0405963. The authors thank the referee for doing a splendid job.

References

1. I. Agol, D.D. Long and A.W. Reid, 'The Bianchi groups are separable on geometrically finite subgroups', *Ann. of Math.* (ii) 153 (2001), no. 3, 599–621.

2. M.D. Baker, 'On boundary slopes of immersed incompressible surfaces', *Ann. Inst. Fourier (Grenoble)* 46 (1996), no. 5, 1443–1449.

3. M.D. Baker and D. Cooper, 'Immersed, virtually-embedded, boundary slopes', *Topology Appl.* 102 (2000), no. 3, 239–252.

4. M. Bestvina and M. Feighn, 'A combination theorem for negatively curved groups', *J. Differential Geom.* 35 (1992), no. 1, 85–101.
Addendum and correction to: "A combination theorem for negatively curved groups" J. Differential Geom. 43 (1996), no. 4, 783–788.

5. B.H. Bowditch, 'Geometrical finiteness for hyperbolic groups.' *J. Funct. Anal.* 113 (1993), no. 2, 245–317

6. B.H. Bowditch and G. Mess, 'A 4-Dimensional Kleinian Group.',*Trans. AMS*, Vol. 344, No. 1. (Jul., 1994), pp. 391–405.

7. D. Cooper and D.D. Long, 'Some Surface Subgroups Survive Surgery', *Geom. Topol.* 5 (2001), 347–367.

8. D. Cooper and D.D. Long, 'Virtually Haken Dehn-filling', *J. Differential Geom.* 52 (1999), no. 1, 173–187.

9. D. Cooper, D.D. Long and A.W. Reid, 'Bundles and finite foliations', *Invent. Math.* 118 (1994), no. 2, 255–283.

10. D. Cooper, D.D. Long and A.W. Reid, 'Essential closed surfaces in bounded 3-manifolds', *J. Amer. Math.* Soc. 10 (1997), no. 3, 553–563.

11. M. Culler and P.B. Shalen, ' Boundary slopes of knots', *Comment. Math. Helv.* 74 (1999), no. 4, 530–547.

12. M. Culler and P.B. Shalen, ' Singular Surfaces, Mod 2 Homology, and Hyperbolic Volume, I',*preprint*, 2007, arXiv GT/0506396v3

13. M. Culler, C. McA. Gordon, J. Luecke and P.B. Shalen, ' Dehn surgery on knots', *Ann. of Math.* (ii) 125 (1987), no. 2, 237–300.

14. F. Dahmani, 'Combination of convergence groups.' Geom. Topol. 7 (2003) 933-963

15. C. E. Fan , 'Injectivity Radius Bounds in Hyperbolic 3-Manifold Cores I', *preprint*, arXiv:math.GT/9907058

16. B. Freedman and M.H. Freedman, 'Haken finiteness for bounded 3-manifolds, locally free groups and cyclic covers', *Topology* 37 (1998), no. 1, 133–147.

17. R. Gitik, 'Graphs and separability properties of groups', J. Algebra 188 (1997), no. 1, 125–143.

18. R. Gitik, 'On the combination theorem for negatively curved groups', *Internat. J. Algebra Comput.* 6 (1996), no. 6, 751–760.
Corrected reprint of "On the combination theorem for negatively curved groups.',',Internat. J. Algebra Comput. 7 (1997), no. 2, 267–276.
Errata: "On the combination theorem for negatively curved groups" Internat. J. Algebra Comput. 7 (1997), no. 2, 265.

19. R. Gitik, 'Ping-Pong on Negatively Curved Groups', *J. Algebra* 217 (1999), 65–72

20. J. Hempel, '3-Manifolds', *Ann. of Math. Studies*, No. 86. Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1976.

21. J. Hass, J.H. Rubinstein and S. Wang, 'Boundary slopes of immersed surfaces in 3-manifolds', *J. Differential Geom.* 52 (1999), no. 2, 303–325.

22. J. Hass, S. Wang and Q. Zhou, 'On finiteness of the number of boundary slopes of immersed surfaces in 3-manifolds', *Proc. Amer. Math. Soc.* 130 (2002), no. 6, 1851–1857

23. A. Hatcher, 'On the boundary curves of incompressible surfaces',*Pacific J. Math.* 99 (1982) 373-377.

24. J. Hempel, 'The finitely generated intersection property for Kleinian groups', Knot theory and manifolds (Vancouver, B.C., 1983), 18–24, *Lecture Notes in Math.*, 1144, Springer, Berlin, 1985.

25. J. Hempel, '3-Manifolds', *Ann. of Math. Studies*, No. 86. Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, (1976).

26. E. Kang and J.H. Rubinstein, 'Ideal triangulations of 3-manifolds I: spun normal surface theory', *G & T monographs*, (2004), vol 7, 235–265

27. D.D. Long, 'Immersions and embeddings of totally geodesic surfaces.' *Bull. London Math. Soc.* 19 (1987), no. 5, 481–484.

28. J. Maher, 'Virtually embedded boundary slopes', *Topology Appl. 95* (1999), no. 1, 63–74.

29. A. I. Malcev, 'On isomorphic matrix representations of infinite groups', *Mat. Sb.* 8, 405-422 (1940).

30. B. Maskit, ' On Klein's combination theorem. IV, ' *Trans. Amer. Math. Soc.* 336 (1993), no. 1, 265–294.

31. J.D. Masters, 'Thick surfaces in hyperbolic 3-manifolds' *Geom. Dedicata* 119 (2006), 17–33.

32. K. Matsuzaki and M. Taniguchi, 'Hyperbolic manifolds and Kleinian Groups' *Oxford Science pub.* (1998)

33. U. Oertel, 'Boundaries of π_1-injective surfaces' *Topology Appl.* 78 (1997), no. 3, 215–234.
34. P. Scott, 'Subgroups of surface groups are almost geometric' *J. London Math. Soc.* (ii) 17 (1978), no. 3, 555–565.
Correction to: "Subgroups of surface groups are almost geometric" J. London Math. Soc. (ii) 32 (1985), no. 2, 217–220.
35. J.P. Serre, 'Trees.', *Springer-Verlag*, Berlin, 2003
36. P. Susskind, 'An infinitely generated intersection of geometrically finite hyperbolic groups.', *Proc. Amer. Math. Soc. 129 (2001)*, no. 9, 2643–2646
37. P. Susskind and G.A. Swarup,' Limit sets of geometrically finite hyperbolic groups.' *Amer. J. Math.* 114 (1992), no. 2, 233–250.
38. P. Susskind, 'Kleinian groups with intersecting limit sets.' *J. Analyse Math.* 52 (1989), 26–38.
39. W.P. Thurston, 'Three-Dimensional Geometry and Topology' *Princeton Univ. Press*, (1997), Ed Silvio Levy
40. W.P. Thurston, 'The Geometry and Topology of Three-Manifolds,' *Princeton notes, 1977. at http://www.msri.org/publications/books/gt3m/*

Mark Baker
Math Dept.
Université de Rennes 1
France

mark.baker@univ-rennes1.fr

Daryl Cooper
Math Dept.
UCSB
Santa Barbara
CA 93106
USA

cooper@math.ucsb.edu