Rotation Distance, Triangulations of Planar Surfaces and Hyperbolic Geometry

Jin-yi Cai *
Department of Computer Science
State University of New York
Buffalo, New York, 14260
`cai@cs.buffalo.edu`

Michael D. Hirsch†
Department of Math and CS
Emory University
Atlanta, GA 30322
`hirsch@mathcs.emory.edu`

(Extended Abstract)

Abstract

In a beautiful paper, Sleator, Tarjan and Thurston solved the problem of maximum rotation distance of two binary trees. Equivalently they solved the problem of rotation distance of triangulations on the disk. We extend their results to rotation distance of triangulations of other planar surfaces. We give upper and lower bounds for this problem. Equivalently, by duality, one can interpret our results as bounds for rotation distance of the dual graphs of the triangulation graphs. They are the counter parts to binary trees in the case of disk. In the case of the annulus, by cutting along an edge between the inner and outer boundary circles, we obtain rooted binary trees with a distinguished path to a leaf.

The upper bound is obtained by looking at the triangulations in the universal covering space, and the lower bound is obtained by extending and applying the technique of volume estimate in hyperbolic geometry.

*Research supported by NSF grants CCR-9057486 and CCR-9319093, and an Alfred T. Sloan Fellowship.
†Research supported by NSF grant DMS-9007211.

1. Introduction

Binary trees have been used to implement data structures that support every conceivable operations as algorithmic primitives. In particular, they have been used to implement binary search trees to support membership queries, insert and delete etc. Most such data structures are based on the idea that a balanced binary tree affords access time of $O(\log n)$, where n is the number of items processed, and thus they need to be maintained as a balanced tree. Sleator and Tarjan [3, 6] pioneered the idea of self-adjusting trees, and they invented *splay trees*. A splay trees does not need explicit rebalancing after each access to the data items, but rather, it tends to balance itself automatically over time. The key operation in a splay tree is a rotation as in Figure 1.1 (a).

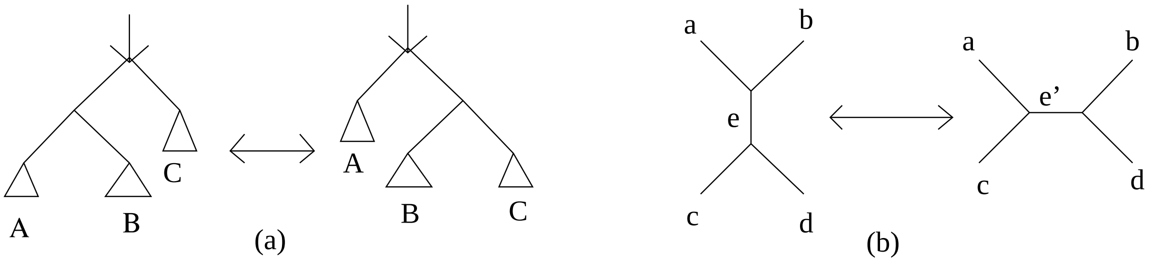

Figure 1.1: (a) Splay tree rotation. (b) Rotation in a 3-regular graph.

Rotation has been realized as a fundamental operation that should be studied carefully. Sleator and Tarjan formulated a dynamic optimality conjecture concerning the performance of splaying [3, 6]. It raises the interesting question: Given two binary trees on n (internal) nodes, what is the maximum distance between the two trees via rotation? In a beautiful paper, Sleator, Tarjan and Thurston [4][5] solved this problem with a tight bound of $2n - 6$. In that paper they first transformed the problem on binary trees to a problem on triangulations on the disk, and then applied techniques of topology and geometry to obtain their bounds. Most remarkable is their use of volume estimate in hyperbolic geometry [1], which yielded their tight bounds. No elementary proof is known.

It is known that the notion of a rotation can be applied to not only binary trees, but any edge with both vertices of degree 3, as in Figure 1.1 (b). Take any triangulation graph of a planar region, and take its dual graph, each triangular face will correspond to a vertex of degree 3. In particular, if we start with a triangulation of the disk, with all its vertices on the boundary circle, and take the dual graph, we obtain precisely the special case of binary trees as in [5][1]. As observed in [5], a rotation in the dual graph (the binary tree) corresponds to the following operation called an edge flip on the triangulation of the disk: Take any quadrilateral formed by two adjacent triangles sharing an edge e, remove e and replace it with the other diagonal of the quadrilateral.

In this paper we extend the results of [5] to the case of triangulation graphs on higher genus surfaces. (In this extended abstract, we will only deal with the case where the genus is 1, although most of our proofs are still valid for surfaces with genus greater than 1.) We study triangulation graphs of the annulus, with all its vertices on its boundary inner

[1]Technically, the process of taking the dual graph will identify all leaves of the binary tree. Another technical point is that in [5], one special external edge is chosen, which corresponds to the root, and they obtain rooted binary trees.

and outer circles. We call them exposed triangulations of the annulus. Our problem is the following: Given two exposed triangulations of the annulus with the same number of vertices on the respective boundary circles, what is the smallest number of edge flips to transform one triangulation to the other, maximized over all such pairs of triangulation?

For any exposed triangulation of the annulus, there must be at least two edges that connect the inner and outer circles. If we cut along one such edge, it can be shown that the dual graph corresponds to a rooted binary tree with a distinguished leaf (equivalently, a distinguished path from the root to a leaf). Such trees with a distinguished path appears to have some applications as *fingered trees* in data structures [7]. By duality, applying edge flips to transform the triangulation is equivalent to applying rotations on the dual graph. Our problem can be stated as determining the maximum rotation distance of the dual graphs.

It turns out that, unlike the case for the disk, there are two natural notions of equivalence of triangulations on the annulus, one takes into account the winding number and other doesn't. For each of these cases, we give upper and lower bounds for this problem. For the upper bound we look at the triangulations in the universal covering space; for the lower bound in the case without winding number, we first prove a theorem of topology, and then argue that non-convexity notwithstanding, the technique of volume estimate of [5] is still applicable. We then proceed to apply their technique of volume estimate in hyperbolic geometry, constructing an explicit pair of exposed triangulations of the annulus. The union of these two triangulations is the 1-skeleton of a geodesic triangulation of the surface of a polyhedron homeomorphic to a solid torus. We give a direct lower bound when winding number is taken into account.

2. Definitions and elementary properties

We would like an exposed triangulation of the annulus to be a triangulation of a topological annulus with all the vertices on the boundary. This definition is a bit too restrictive, so we define it instead as a quotient of a periodic triangulation of the strip $\widetilde{A} = \mathbf{R} \times [0, 1]$.

An integer $n \in \mathbf{Z}$ acts on $\widetilde{A}$ by right translation by n (if n is negative then this is of course a left translation). The quotient by this action is topologically the annulus $A^2 = \{\ (x, y) \mid 1 \leq x^2 + y^2 \leq 2\ \}$. An *exposed triangulation* of $\widetilde{A}$ is a triangulation which has all its vertices on the boundary of $\widetilde{A}$. A *periodic triangulation* is an exposed triangulation which is invariant under the $\mathbf{Z}$ action.

We divide the edges of the triangulation into two types: *exterior* edges which are in the boundary of $\widetilde{A}$ and *interior* edges. Interior edges are further classified into *essential* edges which go from the lower boundary to the upper, and *non essential*. Note that there are always at least 2 essential edges. We will use the word "edge" to mean "interior edge" as we are concerned with interior edges almost exclusively.

The two boundary lines of $\widetilde{A}$ (referred to as *upper* and *lower*) are copies of $\mathbf{R}$. The lower (upper) line covers the inner (outer) circle. Each node in the annulus has one representative in each half-open interval $[n, n+1)$ in one of the boundary lines. *We will always assume there are k nodes on the inner boundary and l on the outer*, and that $(0, 0)$ and $(0, 1)$, which we write as 0 and $0'$ respectively, are lifts of nodes in A^2. An edge between two nodes r, s is written rs.

Two triangulations of the strip are the same if they are isotopic (homotopic through embeddings) while keeping the boundary fixed. Depending on whether we view A^2 as an abstract space or as an annulus with an embedding in the plane, there are two different notions of when two such triangulations are equivalent.

Consider the linear transformation τ_n of $\mathbf{R}^2$ which is the fixes the lower boundary of $\widetilde{A}$ and translates the upper boundary by $n \in \mathbf{Z}$. Clearly, τ_n is not the identity on the boundary of the strip, but it projects down to a map on the annulus which is the identity on the boundary. It is the map which twists the outer boundary n full twists leaving the inner boundary fixed. This provides an equivalence relation on periodic triangulations.

DEFINITION 2.1. *A* concrete *triangulation of the annulus A^2 is a periodic triangulation of $\widetilde{A}$. An* abstract *triangulation of A^2 is an equivalence class of periodic triangulations of $\widetilde{A}$ under the action of $\mathbf{Z}$ given by τ_n.*

A triangle *of any triangulation of A^2 is an equivalence class of triangles in $\widetilde{A}$. An (interior) edge e of a triangulation of A^2 is on the boundary of two triangles. An* edge flip *on e means to remove e, thus joining the two triangle into a quadrilateral, then putting in a new edge to join the other two corners of the quadrilateral. (In $\widetilde{A}$,of course, this is done for every representative of e.)*

Thus, two concrete triangulations of the A^2 may differ by a twist of the outer boundary and are hence the same triangulations when considered abstractly. The *winding number* serves to distinguish these two concrete triangulations.

2.1. Winding number.

DEFINITION 2.2. *Let T be a triangulation of the annulus with j essential edges e_i where e_i has upper and lower endpoints a_i and b_i, respectively, in the universal cover. The* winding *of an edge e_i is $w(e_i) = a_i - b_i$. The* winding number *of a triangulation is*

$$w = \frac{1}{j} \sum_{i=1}^{j} w(e_i).$$

LEMMA 2.3. *For any two essential edges e_1 and e_2, $|(a_1 - b_1) - (a_2 - b_2)| \leq 1$.*

PROOF. We may assume by a translation of the boundaries that $a_1 = b_1 = 0$. Thus the second edge lies in a fundamental domain which has unit length and $|a_2 - b_2|$ can be at most 1. □

COROLLARY 2.4. *Given a triangulation T with winding number w. If e is an essential edge with endpoints a and b, then $|w - (a - b)| \leq 1$.*

LEMMA 2.5. *Suppose there are $j+1$ essential edges in T. Then an edge flip of one essential edge e for an inessential edge can change the winding number by at most $1/j$.*

PROOF. Let T' be the resulting triangulation. Then $w(T') = \frac{1}{j}[\sum_{i=1}^{j+1} w(e_i) - w(e)]$. Thus Δw, the change in winding number, is

$$\begin{aligned}
\Delta w &= \frac{1}{j+1} \sum_{i=1}^{j+1} w(e_i) - \frac{1}{j} \left[\sum_{i=1}^{j+1} w(e_i) - w(e) \right] \\
&= \frac{1}{j} \left[\sum_{i=1}^{j+1} \left[\frac{j}{j+1} w(e_i) - w(e_i) \right] + w(e) \right] \\
&= \frac{1}{j} \left[\sum_{i=1}^{j+1} \frac{-w(e_i)}{j+1} + w(e) \right] \quad = \quad \frac{1}{j}(-w(T) + w(e))
\end{aligned}$$

Hence, by the previous corollary, $|\Delta w| \leq 1/j$. □

LEMMA 2.6. *Suppose there are j essential edges. Then an edge flip of an essential edge for another essential edge can change the winding number by at most $2/j$. When $j = 2$, such an edge flip always changes the winding number by 1.*

PROOF. Consider the edge that is replaced. The new lower endpoint can change by no more than 1 as can the upper endpoint. Thus the average will change by at most $2/j$.

To see what happens when $j = 2$, note that in this case there are only two nodes in the annulus with essential edges. That means that the upper and lower endpoints will move by 1 each in different directions, thus changing the winding number by 1. □

2.2. The dual graph to a triangulation. Given any exposed triangulation T of the annulus, consider its 1-skeleton as a planar graph. It has a dual graph G with its vertex set the set of faces of T, and two vertices of G are connected by an edge iff they share a common edge in T. There could be more than one edges connecting two vertices in G. It is clear that G consists of a cycle of at least two vertices, and from each vertex of the cycle hangs a rooted binary tree with all its leaves identified. (Such a tree is the dual graph of a disk triangulation.) All these binary trees have at least one vertex (its root), are either growing outward toward the outer face containing ∞ to which all its leaves are identified, or growing inward toward the inner face containing 0 to which all its leaves are identified. The edges on the cycle in G correspond to the essential edges in T.

LEMMA 2.7. *Let $N = k + \ell$ be the number of nodes. There are exactly N interior edges as well as N (triangular) faces.*

PROOF. Let E^o be the number of exterior edges, E^i be the number of interior edges, and F be the number of triangular faces. Then clearly $E^o = N$, $3F = E^o + 2E^i$ and $N - (E^o + E^i) + (F + 2) = 2$ by the Euler formula ($F + 2$ is the total number of faces including those containing 0 and ∞). Thus $F = E^i = E^o = N$, i.e., there are exactly N interior edges as well as N triangular faces. So G has N vertices, not counting the identified leaves. □

One defines a rotation on any edge e of G which has both vertices of degree 3 (any edge corresponding to an interior edge of T) as in Figure 1.1. As both vertices of such an edge have degree 3, the operation of a rotation is well-defined. It corresponds in the primal graph T to the operation of an edge flip, in the sense that we have a commutative diagram:

$$\begin{array}{ccc} T & \xrightarrow{f} & T' \\ \downarrow * & & \downarrow * \\ G & \xrightarrow{r} & G' \end{array}$$

where $*$ denotes duality, f denotes an edge flip, and r denotes the corresponding rotation.

3. Shortest flip sequences

We are concerned with the number of edge flips it takes to transform one given triangulation into another.

DEFINITION 3.8. *The* flip distance *between two triangulations T, V is the minimum number of flips neeeded to transform T to V, written $\Delta(T, V)$. (This is equivalent to the rotation distance of the dual graph.) A* flip sequence *is a sequence of triangulations $T = T_1, T_2, \ldots, T_n = V$ such that $\Delta(T_i, T_{i-1}) = 1$ for all i.*

As in the case of triangulations on the disk, we have the following lemma.

THEOREM 3.9. *Given T and V triangulations of the annulus, suppose there is an edge flip that can be done to T so as to introduce an edge in V. Then there is a shortest flip sequence from T to V which starts with the given flip.*

PROOF. As the proof is virtually identical as that of [5], we will only sketch it here.

We will consider the triangulations to be periodic in the strip $\widetilde{A}$. We say that the *normalization* of a triangulation T with respect to an edge rs is a triangulation $N(T)$ which has the same edges as T with the following possible exceptions: (1) rs and all its translates by $\mathbf{Z}$ are always in $N(t)$. (2) If uv crosses any translate $r's'$ of rs in T, then $N(t)$ contains the edges us' and $s'v$.

By considering the effect of normalization on a single triangle in T it is easy to prove:

LEMMA 3.10. *Let $N(T)$ be T normalized with respect to rs. Then $N(T)$ is a periodic triangulation of $\widetilde{A}$. If $\Delta(T, T') = 1$ then $\Delta(N(T), N(T')) \leq 1$. $T = N(T)$ iff T contains rs.*

Let $T = T_0, T_1, \ldots, T_n = V$ be a shortest flip sequence and suppose that rs is an edge in V which can be introduced to T by one flip. Normalizing with respect to rs we get the sequence $T = T_0, N(T_0), N(T_1), \ldots, N(T_n)$ in which $\Delta(T, N(T)) = 1$ and $\Delta(N(T_i), N(T_{i-1})) \leq 1$ for all i.

We know that at some point, say at T_j, the edge rs appears for the first time. Then $N(T_{j-1}) = N(T_j)$, so $T, N(T), N(T_1), \ldots N(T_{j-1}), N(T_{j+1}), \ldots, N(T_n)$ is a flip sequence of minimal length, the first flip being the edge rs. □

Informally, this says that when trying to find a short flip sequence, the obvious greedy algorithm, when applicable, is never wrong; it is always efficient to introduce an edge of the final graph. (Of course, finding such an edge is not always possible.)

Similar arguments also show

COROLLARY 3.11. *In a shortest flip sequence from T to V, if an edge of V appears at any intermediate stage then it is never deleted. If any edge is deleted then it is never reintroduced.*

4. Upper bounds

In this section we will demonstrate some upper bounds. As there are two types of problems, there are two upper bounds. The first case is the abstract case in which winding is ignored.

Given triangulations T_1 and T_2, we may assume that $00'$ is an edge in T_1. We will show that there is a canonical triangulation T' which has distance at most $k + l - 1 = N - 1$ from T_1 and $1.5N - 1$ from T_2, proving

THEOREM 4.12. *Any two triangulations are of flip distance at most $\lfloor 2.5N - 2 \rfloor$ in the abstract case.*

PROOF. Let 0 and r be nodes of a triangulation and e an edge not containing 0 or r. If there is a path from 0 to r which crosses only edge e, then e is the opposing side to both 0 and r in triangles of the triangulation. Accordingly, there is an edge flip that can be made flipping e to an edge from 0 to r. We call this *piercing* e with an edge from 0.

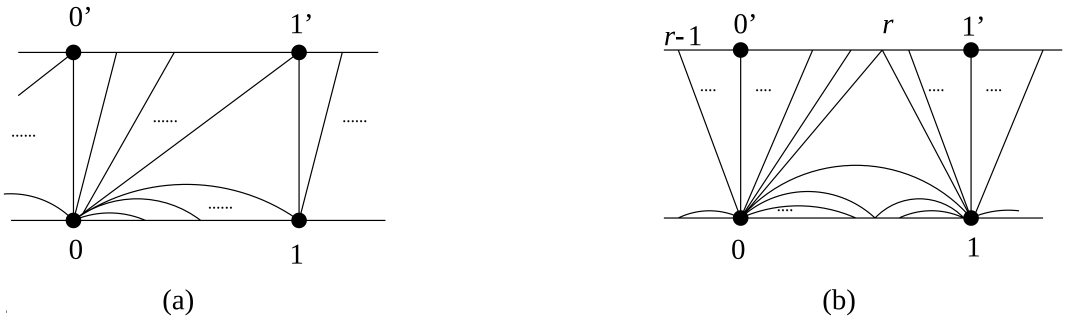

Figure 4.2: (a) The final stage for T_1 and T_2. (b) An intermediate stage for T_2.

$00'$ is an edge of T_1 by hypothesis. We can pierce all edges not meeting 0 with an edge from 0, one at a time. After this, all edges have 0 as an endpoint, and we can assume that the edges leave 0 all to the right or all to the left of $00'$. Each possibilities takes at most $k+l-1 = N-1$ flips. After this we have reached the canonical position shown in Figure 4.2 (a), or its mirror image.

$00'$ may not be an edge in T_2, so we must make it one. If 0 is on no essential edge, choose an arc from 0 to the upper boundary which crosses a minimal number of edges. Pierce these edges with an edge from 0 until 0 is connected to the upper boundary. Call the upper endpoint r. If 0 is on an essential edge, let r be the other endpoint of this edge.

Now, moving leftwards from $0r$, pierce all other edges attached to the upper boundary to make a new edge to 0. Continue until the node $r-1$ is reached. Finish connecting all the lower nodes to 0 (including -1). In these three stages, no new edge has been pierced and all edges (N of them, by lemma 2.7) go to 0. Therefore we have done at most N flips.

At this point (see Figure 4.2 (b)) the upper edge of T_2 is attached to 0 in a way that differs from the final configuration of T_1 in that $0'$ is in the middle of the group of edges connected to 0 instead of being the outermost node in each group. Similarly, 0 in T_2 does not have all its edges to the lower boundary leaving to one side.

We can now pierce all edges either to the right or to the left of $00'$ to get to the same final position as T_1, choosing the direction which minimizes flips. One of the directions has at most $\lfloor (N-2)/2 \rfloor$ flips required ($00'$ and $01'$ don't get flipped), which proves that T_1 and T_2 can be changed to the same picture in $2.5N-2$ flips. □

The concrete case is similar. Again, we may assume that $00'$ is an edge in T_1, $w(T_1) = 0$ and $w(T_2) = w > 0$. We are able to show that there is a canonical triangulation which is distance at most $N+l+w$ from T_1 and $N+k/2$ from T_2, proving:

THEOREM 4.13. *Any two triangulations are of flip distance at most* $\lfloor 3.25N + |w(T_1) - w(T_2)| \rfloor$.

PROOF. The proof is quite similar, but complicated by the winding, and we only provide a sketch. T_1 can still be made to look like Figures 4.2 (b), but it is no longer easy to get T_2 to Figure 4.2 (b). Because of the winding 0 is not close to $0'$, but we can get fairly easily to a situation like Figures 4.2 (b) except with some other integer i' on the upper boundary. This takes less than $2N$ flips.

From Lemma 2.6 we know that it is best to have few essential edges so as to change the winding number quickly. This can be done with no more than $2l+k/2$ flips. At this point T_1

and T_2 are the same abstract triangulation, but they differ by appproximately w in winding number (use the results of §2.1 to control the growth of w). It takes w more flips to make them agree. This gives us a total of $N + N + l + k/2 + l + w = 2.5N + 3l/2 + w$ flips.. Since k and l are interchangeable, we can assume $l \leq N/2$ and it takes at most $\lfloor 3.25N - 2 + w \rfloor$. □

5. Lower bounds

We apply the hyperbolic volume estimate technique of Sleator, Tarjan and Thurston to construct two triangulations of the annulus that are provably "far apart".

Suppose there is a flip sequence of length m from one triangulation to another. Place the first triangulation on an annulus A^2. For each edge flip, the original edge borders two adjacent triangles. After the edge flip, the original edge disappears together with the two triangles, and the opposite two vertices are linked together by a new edge, creating two new adjacent triangles.

Imagine the quadrilateral in which this occurs is "fattened" to a tetrahedron, where the two original triangles are on the lower side of the tetrahedron and the two new triangles are on the upper side of the tetrahedron. Paste a (topological) tetrahedron to the quadrilateral with two original triangles identified with the two faces of the tetrahedron's lower side, exposing the tetrahedron's upper side. At this point, the view from above is the triangulation after one flip. We do this for each flips successively, and in the end of this sequence of m edge flips we have obtained the second triangulation.

We call a triangulation of a solid torus an exposed triangulation if all its vertices are on its boundary. Thus, if the two triangulations of A^2 share no common edge, they bound a topological solid torus with an exposed triangulation by m tetrahedra. The solid torus' surface triangulation is the union of the two given triangulations.

Suppose we have a polyhedral solid torus embedded in the hyperbolic space $\mathbf{H}^3$. Suppose the surface of the polyhedron has a geodesic triangulation whose 1-skeleton, as a graph, has the property that its vertex set can be decomposed into two disjoint cycles, and these two cycles are disjointly homotopic to longitudes—dividing the boundary torus into two triangulated annuli. Let τ be a sequence of m edge flips which transforms one triangulation to the other. Construct an abstract (topological) solid torus $ST^2 \cong B^1 \times S^1$ by the process described in the last paragraph to get an exposed triangulation of ST^2 by m tetrahedra, whose surface triangulation is the union of the two given triangulations.

We now define a map from ST^2 to $\mathbf{H}^3$. The map is defined locally one tetrahedron at a time. The image of a tetrahedron in ST^2 is the geodesic tetrahedron spanned by the 4 vertices in $\mathbf{H}^3$. Since the geodesic path connecting two vertices and the geodesic surface connecting three vertices in the hyperbolic space $\mathbf{H}^3$ is unique, the map is well defined. Also the map on the surface of ST^2, namely the torus T^2, is a map from T^2 to T^2, and has degree 1 [2].

The following theorem is critical:

THEOREM 5.14. *Let $f : ST^2 \to \mathbf{H}^3$ be a continuous, piece-wise smooth map from the solid torus ST^2 to 3-dimensional hyperbolic space $\mathbf{H}^3$, where a copy of the solid torus $\widetilde{ST^2}$ is embedded. Let $T^2 = \partial(ST^2) \cong S^1 \times S^1$ and $\widetilde{T^2} = \partial(\widetilde{ST^2})$ be the boundary tori of the respective solid tori, and suppose the restriction $f|_{T^2}$ of f to its boundary maps T^2 into $\widetilde{T^2}$*

with $\deg(f|_{T^2}) \neq 0$. *Then every point in the solid torus* $\widetilde{ST^2}$ *is covered by* f, *i.e.,* $\forall y \in \widetilde{ST^2}$, $\exists x \in ST^2$, *such that* $f(x) = y$.

PROOF. The proof is by contradiction. Let $y_0 \in \widetilde{ST^2}$ be a point of the solid torus not covered by the map f. Since the map is continuous and the domain set ST^2 is compact, the image set is also compact, and hence, there exists a ball $B_\varepsilon = \{\ y \mid \mathrm{dist}(y, y_0) \leq \varepsilon\ \}$ around y_0 which is disjoint from the image set $f(ST^2)$.

Let ρ be the radial projection from $\mathbf{H}^3 - B_\varepsilon$ to the boundary sphere ∂B_ε of B_ε. We define a new map $g : ST^2 \to \partial B_\varepsilon$ by composition $g = \rho \circ f$, then, since $\deg(\rho|_{T^{2\prime}}) = 1$, the restriction $g|_{T^2} = \rho \circ f|_{T^2}$ has degree $\deg(g|_{T^2}) \neq 0$.

However, since g is continuous and piece-wise smooth, when g is restricted to the "core circle" $\{0\} \times S^1$, the image $g(\{0\} \times S^1)$ does not cover the entire sphere ∂B_ε. By continuity, there exists some small $\delta > 0$, such that the image $g(T^2_\delta)$ of g on the "thin" torus $T^2_\delta = \{\ x \mid |x| = \delta\} \times S^1$ does not cover ∂B_ε, and hence has degree 0. But since g itself is a homotopy between $g|_{T^2}$ and $g|_{T^2_\delta}$, $\deg(g|_{T^2}) = \deg(g|_{T^2_\delta}) = 0$. This contradiction proves that the map must cover the solid torus $\widetilde{ST^2}$. □

The same proof would work in Euclidean space; the reason we chose to consider it in hyperbolic space instead is due to the following crucial fact: Every geodesic tetrahedron in 3-dimensional hyperbolic space has a bounded volume, which is no more than $V_0 = 1.014941606\ldots$, the volume of a certain ideal tetrahedron with all four vertices at the infinity plane.

By applying this fact to our polyhedron in $\mathbf{H}^3$, we know that the union of m geodesic tetrahedra covers the given polyhedron, and thus mV_0 is an upper bound on the total volume of the polyhedron. Thus, if we can construct such a polyhedron with n vertices and volume at least $(2n - o(n))V_0$, we will have proved a lower bound of $m \geq 2n - o(n)$.

We now describe our construction, an adaptation of Sleator, Tarjan and Thurston's construction to get a polyhedron with n vertices and volume at least $(2n - O(\sqrt{n}))V_0$.

We work in the upper half-space Poincaré model of $\mathbf{H}^3$. Here geodesic paths are semicircles or Euclidean lines perpendicular to the (infinity) plane $z = 0$, and geodesic surfaces are hemispheres or Euclidean planes perpendicular to plane $z = 0$. An ideal tetrahedron with maximum volume V_0 has its vertices at the infinity plane, $0, 1, e^{2\pi i/6}$ and ∞, the infinity point thought of as lying above all points ($z = \infty$) and belonging to all Euclidean lines perpendicular to $z = 0$.

The general shape of the polyhedron is an infinite prism with a hole. The infinite prism has a equilateral triangular base (Figure 5.3) on the infinity plane $z = 0$ and rises to the infinity point $z = \infty$. (Strictly speaking it is not a prism, since its upper "face" is really a single point $z = \infty$, but this should not cause any confusion.) Then we dig a "tunnel" in this prism: let a large semicircular arch with a small square base go through the prism at high altitude; delete the intersection from the prism. (If you imagine the apex of the St. Louis Arch puncturing a triangular Empire State Building, you've got the idea.)

The base of the infinite prism is triangulated by equilateral triangles each of which is congruent to the ideal triangle on vertices $0, 1, e^{2\pi i/6}$ (Figure 5.3). They are triangulated via hemispheric bubbles each of which passes through three vertices of each equilateral triangle. On the three sides of the infinite prism, all vertices on the infinity plane $z = 0$ rise to connect to the infinity point $z = \infty$ via straight lines. On two faces of these three sides of the prism, there is a square deleted, marking the entrance and the exit of the tunnel. The surface triangulation is properly repaired so that we still get a geodesic surface triangulation.

More specifically, let a, b, c, d be the labels of the four vertices of the entrance of the tunnel, clockwise facing the tunnel starting with a as the label for the upper left corner of the square of the entrance, and let a', b', c', d' be the labels of the four vertices of the exit square of the tunnel in the corresponding positions, so that aa', bb', cc' and dd' are connected via the semicircular geodesic paths of the arch. Choose the position and the size of the base squares of the arch so that a and d have the same (x, y) coordinates and project to the same vertex which is also a vertex α of the triangulation of the prism's base. Similarly b and c, a' and d', and b' and c' also project to some three vertices, β, α' and β' respectively. We also let α and α' be both adjacent to the same corner vertex δ of the triangular base as in Figure 5.3.

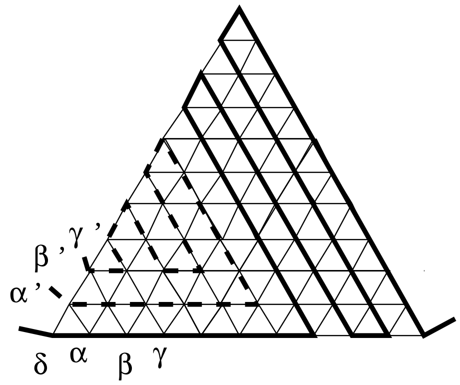

Figure 5.3: The triangulated "base" of the polyhedron.

Let γ and γ' be on the circumference of the triangular base connecting to β and β' respectively, as in Figure 5.3. The surface triangulation before the tunnel was dug has geodesic paths as edges from α, β, α' and β' to ∞. They are replaced by α to d, β to c, α' to d', β' to c', the two squares with edges (by geodesics) ab, bc, cd and da, and $a'b'$, $b'c'$, $c'd'$ and $d'a'$, and four geodesics a, b, a' and b' all to ∞. We also connect aa', bb' cc' and dd' via semicircular edges of the arch. To make it a triangulation again, we need to add some more. Let's connect (by geodesics) d to β and d' to β'; ab', bc', cd' and da'. Finally, we will add edges (by geodesics) a, d, a' and d' to δ, b and c to γ, and b' and c' to γ'. It can easily be checked that this is a triangulation.

Let there be $n - 9$ points on the base of the infinite prism; together with 8 points of the tunnel entrance and exit and the infinity point ∞ there are a total of n points. There are $2n - O(\sqrt{n})$ ideal tetrahedra with volume V_0 each. This can be easily seen as each vertex in the interior of the base borders 6 equilateral triangles and each such equilateral triangle has three vertices. The boundary has only $O(\sqrt{n})$ vertices and the tunnel takes away a negligible amount of volume since it intersects at high altitude (most of the volume of hyperbolic space is concentrated near the infinity plane $z = 0$.)

Thus our polyhedron has the claimed volume $(2n - O(\sqrt{n}))V_0$. It is clear that it is topologically a solid torus.

We now verify that the 1-skeleton graph of the surface triangulation can be decomposed into a disjoint union of two cycles with k and l vertices, where $k + l = n$. And these two cycles are both isotopic to the the "longitude" $abcd$. (Another way to say this is that the cycles have unit winding number with respect to the arch.) Thus, between them are two annuli, which can be homotoped to each other within the solid torus. Furthermore, we can take k and l to be any value between $O(1)$ and $n - O(1)$, subject to the above condition. Since we are only proving our bound up to an additive term of $O(\sqrt{n})$, we will only verify that k and l can take any value with an increment of $O(\sqrt{n})$.

One cycle descends from $z = \infty$ to both δ and another corner point and weaves its way to cover a portion of the triangular base of the infinite prism. The area it covers is the complementary area of a smaller triangular region, which is covered by the other cycle. (See Figure 5.3.) This second cycle weaves out the smaller triangular region leaving at α' and β', where they go straight up to meet d' and c' respectively. There it includes the 8 vertices in the following order, $d', d, a, a', b', b, c, c'$. Clearly, they together cover all the vertices and they are both isotopic to $abcd$. We have proved the following Theorem.

THEOREM 5.15. *There exist two (abstract) triangulations T_1 and T_2, each with (k, ℓ) vertices, and $N = k+\ell$, such that the rotation distance between the two is at least $2N - O(\sqrt{N})$.*

Similar to [5], we could improve our construction, say, using the "subdivided" hyperbolic icosahedron to improve on the term $O(\sqrt{n})$. But since our bound is not tight for the leading term, there is little point to pursue it at this point.

It is not clear how to apply the volume technique in general to account for the winding of concrete triangulations. However for this concrete case we have a simpler analysis.

THEOREM 5.16. *Given two triangulations, T and V, each with N essential edges, let $w = |w(T) - w(V)|$. Then when $w > 2\log N + 4$ the rotation distance $d(T, V)$ between T and V is lower bounded by $2N + w - 2\log N - 4$, where the $\log$ is base e.*

PROOF. Suppose we have a sequence of edge flips from T to V. Let j be the minimum over the sequence of the number of essential edges in each triangulation. Note that $j \geq 2$.

The sequence must contain at least $2(N - j)$ flips in which an essential edge is traded for a non-essential edge ($N - j$ to get to j and $N - j$ to get back to N). By Lemma 2.5 this can change the winding by a total of

$$\Delta w = 2 \sum_{i=j}^{N-1} 1/i < 2 \log N.$$

By Lemma 2.6 it takes at least another $(w - 2\log N) \cdot \frac{1}{2/j}$ flips to change w enough.

Hence, the total number of flips $d(T, V)$ is at least $2(N - j) + (w - 2\log N)j/2 = 2N + j(w/2 - (\log N + 2))$. When $w > 2\log N + 4$ this is minimized by $j = 2$. So the number of flips is at least $2N + w - 2\log N - 4$. □

Acknowledgements

We wish to thank Claire Kenyon and Laci Lovász for several long discussions. We also thank Danny Sleator, Bob Tarjan and Andy Yao for enlightening discussions on the subject. Thanks also to Bill Thurston for asking about the winding number case.

References

[1] H. S, M. Coxeter, *Introduction to Geometry*, 1969, Wiley, New York.

[2] W. S. Massey, *Singular Homology Theory*, 1980, Springer-Verlag, New York.

[3] D. D. Sleator and R. E. Tarjan, Self-adjusting binary trees, J. ACM, 32, 1985, 652–686.

[4] D. D. Sleator, R. E. Tarjan and W. P. Thurston, Rotation distance, triangulations, and hyperbolic geometry, *Proc. ACM STOC*, 1986, 122–135.

[5] D. D. Sleator, R. E. Tarjan and W. P. Thurston, Rotation distance, triangulations, and hyperbolic geometry, J. AMS, vol 1, Num 3, 1988, 647–681.

[6] R. E. Tarjan, *Data structures and network algorithms*, SIAM, PA, 1983.

[7] R. E. Tarjan, Private communications.