Rotation Distance, Triangulations, and Hyperbolic Geometry

Daniel D. Sleator
Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213

Robert E. Tarjan
Computer Science Department
Princeton University
Princeton, NJ 08544
and
AT&T Bell Laboratories
Murray Hill, NJ 07974

William P. Thurston
Mathematics Department
Princeton University
Princeton, NJ 08544

Abstract

A *rotation* in a binary tree is a local restructuring that changes the tree into another tree. Rotations are useful in the design of data structures. The *rotation distance* between a pair of trees is the minimum number of rotations needed to convert one tree into the other. In this paper we establish a tight bound of $2n-6$ on the maximum rotation distance between two n-node trees using volumetric arguments in hyperbolic 3-space. Our proof also gives a tight bound on the minimum number of tetrahedra needed to cover a polyhedron in the worst case, and reveals connections among binary trees, triangulations, polyhedra, and hyperbolic geometry.

1. Introduction

A *rotation* in a binary tree is a local restructuring of the tree that changes it into another tree. The *rotation distance* between a pair of trees is the minimum number of rotations needed to convert one tree into the other. The problem addressed in this paper is: what is the maximum rotation distance between any pair of n node binary trees? We show that for all $n \geq 11$ this distance is at most $2n-6$ and that for an infinite set of values of n this bound is tight. To our knowledge the only published work on this problem is by Culik and Wood [2], who defined the concept and showed that the maximum rotation distance is at most $2n-2$. Leighton (private communication) showed that there exist trees such that the rotation distance between them is $7n/4 - O(1)$

Our interest in this problem stems from our attempt to solve "the dynamic optimality conjecture" about the performance of splaying [7]. Splaying is a heuristic for modifying the structure of a binary tree in such a way that accessing and updating the information in the tree is efficient. Although our solution to the problem of maximum rotation distance did not resolve the conjecture about splaying, the results in this paper are interesting for at least two other reasons. First, the combinatorial system of trees and their rotations is a fundamental one that is isomorphic to other natural combinatorial systems. Results concerning this system are of interest from a purely mathematical point of view. Second, the method we use to solve the problem is novel and interesting in its own right, and can potentially be applied to other related problems.

A system that is isomorphic to binary trees related by rotations is that of triangulations of a polygon related by the *diagonal flip* operation. This is the operation that converts one triangulation of a polygon into another by removing a diagonal in the triangulation and adding the diagonal that subdivides the resulting quadrilateral in the opposite way. This type of move was studied by Wagner [11] in the context of arbitrary triangulated planar graphs, and by Dewdney [3] in the case of graphs of genus one. They showed that any such graph can be transformed to any other by diagonal flips, but did not concern themselves with accurately estimating how many flips are necessary.

Our approach to solving the rotation distance problem is based on the observation that any sequence of diagonal flips converting one triangulation of a polygon into another gives a way to dissect (into tetrahedra) a polyhedron formed from the two triangulations. Using hyperbolic geometry we construct polyhedra that require many tetrahedra to triangulate them. (Here and hereafter we use the word "triangulation" in a general sense meaning a dissection of a polytope into simplices of appropriate dimension.) These polyhedra can be used to exhibit pairs of n node trees (for infinitely many values of n) such that the rotation distance between them is $2n-6$.

In section 2 we define the problem on trees, make the connection between trees and triangulations of a polygon, and show that sequences of diagonal flips are related to triangulations of polyhedra. In section 3 we show how to use hyperbolic geometry to obtain a lower bound on the number of tetrahedra required to triangulate any polyhedron. We then construct particular polyhedra that require many tetrahedra to triangulate them. Section 4 contains remarks and some open problems.

2. Definitions and Equivalences

2.1. Binary Trees

A binary tree is a collection of nodes of two types, external and internal, and three relations among these nodes: parent, left child and right child. Every node except a special one called the root has a parent, and every internal node has a left and a right child. External nodes have no children. A tree is said to be of size n if it has n internal nodes. A tree of size n has $n+1$ external nodes. (See [4] for a more complete description of binary trees and tree terminology.) The number of steps required to walk from the root of the tree to a node is the depth of that node. (Each step moves from a node to one of its children.)

A *symmetric order* traversal of the tree visits all of the nodes exactly once. This order can be described by a recursive algorithm as follows: If the node is an internal node, traverse its left subtree in symmetric order, then visit the node itself, then traverse its right subtree in symmetric order. If the node is an external node, then visit it and return. The order in which the nodes are visited is called the *symmetric order permutation* of the nodes (or simply the symmetric order of the nodes).

In a common computer-related application of binary trees the tree is used to store an ordered collection of chunks of information (called items). Each internal node of the tree is labeled with an item, and the order of the items is represented by the symmetric order of the nodes in the tree.

A *rotation* is an operation that changes one binary tree into another. In a tree of size n there are $n-1$ possible rotations, one corresponding to each non-root internal node. Figure 1 shows the general rotation rule and the effect of a particular rotation on a particular tree. The rotation corresponding to a node changes the structure of the tree near that node, but leaves the structure elsewhere intact. A rotation maintains the symmetric order of the nodes of the tree, but changes the depths of some of the nodes. Rotations are the primitives used by most schemes that maintain "balance" in binary trees. [4]

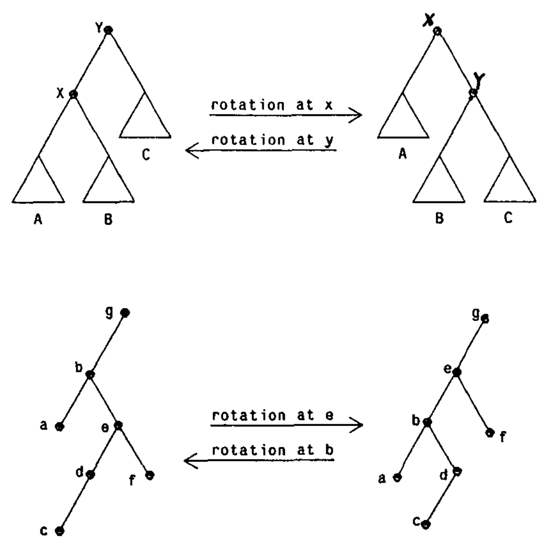

Figure 1. a) The general definition of a rotation. Triangles denote subtrees. The tree shown could be part of a larger tree. b) A rotation in a seven node tree. External nodes are not shown.

A rotation is an invertible operation; that is, if tree T can be changed into T' by a rotation, then T' can be changed back into T by a rotation. The *rotation graph* for trees of size n (denoted $RG(n)$) is an undirected graph with one vertex for each tree of size n, and an edge between nodes T and T' if there is a rotation that changes T into T'.

Any binary tree of size n can be converted to any other by performing an appropriate sequence of rotations. Therefore the rotation graph is connected. We can define the *rotation distance* between two trees as the length of the shortest path in the rotation graph between the two trees. This rotation distance is the minimum number of rotations required to convert one tree into the other.

2.2. Polygon Triangulations

The problems we study can be formulated with respect to a different system of combinatorial objects and their transformations. This alternate formulation is perhaps more natural and also seems to supply more insight.

The Catalan numbers C_n count the number of binary trees of size n as well as the number of planar triangulations of a convex $(n+2)$-gon with no interior vertices. We refer to the $n+2$ sides of the polygon as edges and the chords that divide it into triangles as diagonals. Any triangulation of the $(n+2)$-gon has $n-1$ diagonals and n triangles. We regard the polygon as having a distinguished edge and orientation.

A *diagonal flip* is an operation that transforms one triangulation of a polygon into another. The effect of a diagonal flip is shown in Figure 2, and can be described as follows: A diagonal inside the polygon is removed, creating a face with four sides. The opposite diagonal of this quadrilateral is inserted in place of the one removed, restoring the diagram to a triangulation of the polygon.

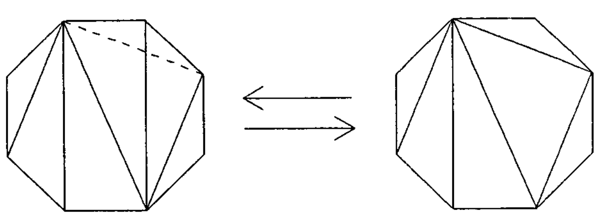

Figure 2. A diagonal flip in a triangulation of an octagon.

Let $TG(n+2)$ be a graph with one node for each triangulation of an $(n+2)$-gon and an edge between two nodes if the two nodes are related by a diagonal flip.

Lemma 1: The graph $TG(n+2)$ is isomorphic to the rotation graph $RG(n)$.

This fact is proven by exhibiting a one to one correspondence between trees of size n and triangulations of an $(n+2)$-gon in which diagonal flips correspond to rotations. This correspondence is shown in Figure 3. A proof of the lemma can be found in [5].

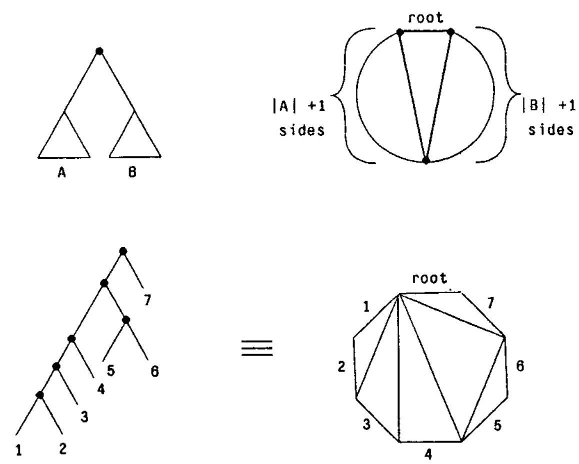

Figure 3. a) The correspondence between a tree and a triangulation. b) An example of a tree and its corresponding triangulation.

2.3. Results on Polygon Triangulations

As we saw in Section 2.2, a study of the rotation distance between trees can be formulated as a study of the distance between triangulations under the diagonal flip operation. Let $d(\tau_1,\tau_2)$ be the minimum number of diagonal flips needed to transform triangulation τ_1 into triangulation τ_2. For convenience, we shall now change our use of the variable "n". We consider triangulations of an n-gon and let $d(n)$ be the maximum distance between any pair of such triangulations. That is, $d(n)$ is the diameter of $TG(n)$ or equivalently of $RG(n-2)$. Figure 4 shows $TG(6)$, whose vertices are the fourteen triangulations of a hexagon. The greatest distance between a pair of triangulations is four; there are several pairs that achieve this distance.

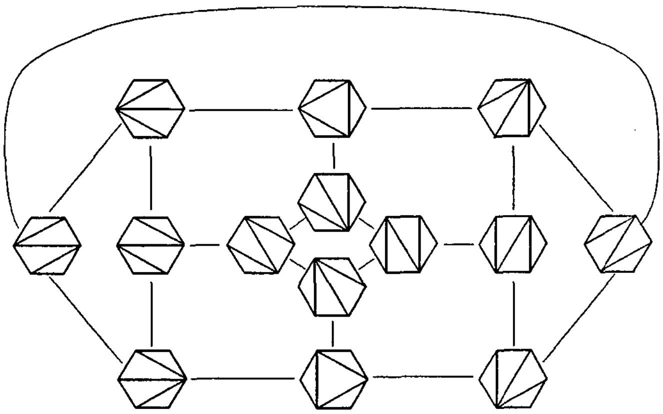

Figure 4. The rotation graph of a hexagon, $RG(6)$.

The added symmetry revealed in the triangulation system that is hidden in the binary tree system enables us to improve Culik and Wood's upper bound on $d(n)$ from $2n-6$ to $2n-10$.

Lemma 2: $d(n) \leq 2n-10$ *for all* $n>12$.

Proof: Any triangulation of an n-gon has $n-3$ diagonals. Given any vertex x of degree $deg(x) < n-3$, we can increase $deg(x)$ by one by a suitable flip. Thus in $n-3-deg(x)$ flips we can produce the unique triangulation all of whose diagonals have one end at x. It follows that given any two triangulations τ_1 and τ_2 we can convert τ_1 into τ_2 in $2n-6-deg_1(x)-deg_2(x)$ flips, where x is any vertex and the degree of x is $deg_1(x)$ in τ_1 and $deg_2(x)$ in τ_2. The average over vertices x of $deg_1(x)$ is $2-6/n$, and of $deg_1(x)+deg_2(x)$ is $4-12/n$. It follows that if $n>12$, there is a vertex x such that $deg_1(x)+deg_2(x) \geq 4$.

□

The following lemma about sequences of diagonal flips will be useful later.

Lemma 3: a) If it is possible to flip one diagonal of τ_1 *creating* τ_1' *so that* τ_1' *has one more diagonal in common with* τ_2 *than does* τ_1 *then there exists a shortest path from* τ_1 *to* τ_2 *in which the first flip creates* τ_1'. *b) If* τ_1 *and* τ_2 *have a diagonal in common, then every shortest path from* τ_1 *to* τ_2 *never flips this diagonal. In fact, any path that flips this diagonal is at least two flips longer than a shortest path.*

Proof: Let S be a sequence of adjacent triangulations connecting τ_1 to τ_2.

$$S = t_0(=\tau_1), t_1, t_2, \cdots, t_k(=\tau_2)$$

Assume that $t_1 \neq \tau_1'$. We shall construct a new sequence of adjacent triangulations S' also connecting τ_1 and τ_2 whose length is no longer than the length of S, and in which the first flip creates τ_1'. This will suffice to prove part 1 of the lemma.

Let l and r be the end points of the diagonal that τ_1' and τ_2 but not τ_1 and τ_2 have in common. Any triangulation τ can be *normalized* with respect to the diagonal (l,r) to create a new triangulation $N(\tau)$. The diagonals of $N(\tau)$ are of three types: 1) $N(\tau)$ contains the diagonal (l,r), 2) $N(\tau)$ contains every diagonal of τ that does not cross the diagonal (l,r) (two diagonals with an endpoint in common are not said to cross), 3) if τ contains a diagonal (a,b) that crosses the diagonal (l,r) then $N(\tau)$ contains the diagonals (a,r) and (b,r). (See Figure 5.)

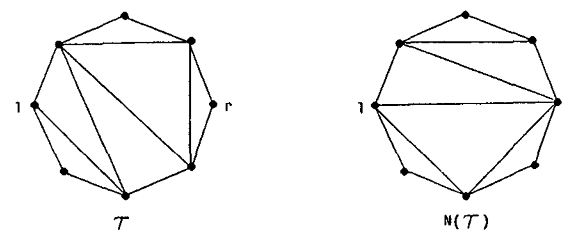

Figure 5. A triangulation τ and its normalized version $N(\tau)$.

Consider the sequence of triangulations

$$N = t_0, N(t_0), N(t_1), \cdots, N(t_k).$$

A straightforward case analysis shows that successive triangulations of this sequence are either identical or adjacent. Eliminating all but one of each group of identical consecutive triangulations in this sequence gives the desired sequence S'. A priori S' might contain $k+2$ triangulations, but this cannot be the case for the following reason. Consider the triangulations t_i and t_{i+1} in S with the property that t_i does not contain diagonal (l,r) and t_{i+1} does. (There must be such a pair since the final triangulation contains the diagonal (l,r) and the initial one does not.) It is easily verified that the triangulations $N(t_i)$ and $N(t_{i+1})$ must be equal, and therefore only occur once in S'. Thus S' contains at most $k+1$ triangulations. Verifying that S' starts and ends with τ_1 and τ_2 and that its second triangulation is τ_1', is straightforward. This completes the proof of part a of the lemma.

The same technique serves to prove part b of the lemma. Let $S = t_0(=\tau_1), t_1, \cdots, t_k(=\tau_2)$ be a sequence that transforms τ_1 into τ_2 in which the first move is one that flips a diagonal (l,r) common to both τ_1 and τ_2. Normalize this sequence with respect to the diagonal (l,r) and eliminate redundancies to create a sequence S'. S' transforms τ_1 into τ_2 in two fewer flips than does S. The reason is that neither the first flip of S, that misaligns (l,r), nor a later flip that aligns (l,r) occurs in S'.

□.

A refinement of the lower bound proof in Lemma 3 for small values of n and a computer search have produced the exact values of $d(n)$ for $n \leq 18$. Here are these values.

n	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
$d(n)$	0	1	2	4	5	7	9	11	12	15	16	18	20	22	24	26

2.4. Triangulations of the Sphere and the Ball

In this section we show that the quantity $d(n)$ is related to the number of tetrahedra that are required to dissect certain polyhedra. First we need a more general definition of a triangulation.

A d-dimensional triangulation is a collection of labeled d-dimensional simplices along with a collection of *gluing rules*. Each gluing rule specifies that one of the $(d-1)$-dimensional facets in this collection of simplices is to be identified with another $(d-1)$-dimensional facet in the collection. Each facet occurs in at most one gluing rule, and a facet cannot be glued to itself. (It may be possible to represent a triangulation in d-dimensional space, but its mathematical structure has nothing to do with how it is embedded in space; e.g. the mathematical structure is in the gluing rules and not in the geometric realization.)

Let T be a d-dimensional triangulation. Some of the $d-1$-dimensional facets of T may not be glued to any other facet. These facets along with certain gluing rules form a $d-1$ dimensional triangulation called the *boundary* of T, denoted ∂T. The gluing rules of ∂T are of two types. Two $d-2$-dimensional facets of ∂T are glued together if there is a gluing rule of T that identifies these two facets of ∂T with each other (or a collection of gluing rules of T that imply by transitivity the equivalence of the two facets of ∂T). Also, if two $d-1$ dimensional simplices of ∂T come from the same d-dimensional simplex of T, then there is a rule to glue their common $d-2$ dimensional simplices together appropriately. Figure 6 shows an example of a two-dimensional triangulation and its boundary.

A *triangulation of the ball* is a 3-dimensional triangulation that is homeomorphic to a ball. That is, there exists a way of rendering each tetrahedron of the triangulation as a curvilinear tetrahedron so that: 1) each tetrahedron is homeomorphic to a ball, 2) the interiors of the tetrahedra are disjoint, 3) the union of all of the tetrahedra form the ball, and 4) if the faces of two tetrahedra are to be glued together then these faces must coincide. (A ball is the set of points $\{(x,y,z) \mid x^2+y^2+z^2 \leq 1\}$.)

A *triangulation of the sphere* is a triangulation that is a boundary of a triangulation of the ball. (A sphere is the set of points $\{(x,y,z) \mid x^2+y^2+z^2 = 1\}$.) If σ is a triangulation of the sphere then there is a way to dissect the sphere into curvilinear triangles whose adjacency is given by the gluing rules of σ. It is sometimes more convenient to think of a triangulation of the sphere as an *embedded triangulated planar graph*. This is a planar graph along with a permutation at each vertex specifying the clockwise order of the edges incident with that vertex, having the additional property that each face is a triangle. Each vertex of the triangulation is a vertex of the graph, and each triangle of the triangulation is a face of the graph.

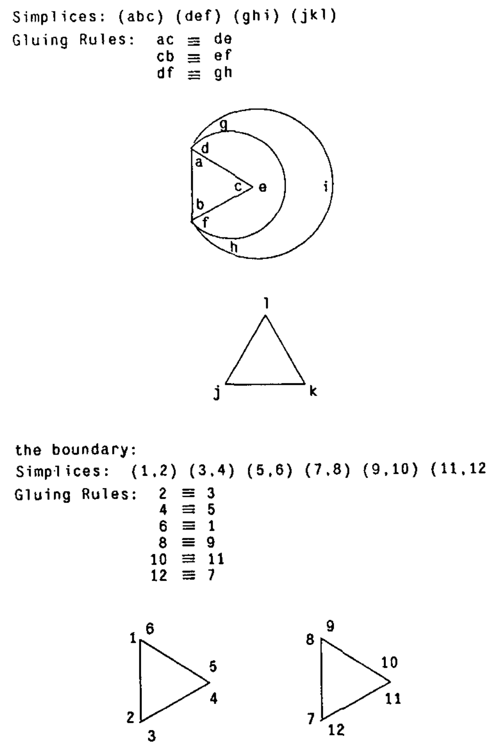

Figure 6. A triangulation and its boundary.

The *boundary of a 3-dimensional polyhedron* is the embedded planar graph obtained by extracting from the polyhedron the edges and vertices and their incidence and embedding relationships.

If σ is a triangulation of the sphere and T is a triangulation of the ball then T is an *exposed triangulation of the ball extending* σ if $\sigma = \partial T$ and all of the vertices of T occur in σ. (None of the vertices of T are inside the ball.)

An exposed triangulation of the ball extending σ is the three dimensional analogy to a triangulation of an n-gon described in previous section. In contrast to the situation in two dimensions there may be several triangulations extending σ containing different numbers of tetrahedra.

The union of two triangulations τ_1 and τ_2, of an n-gon, combined with n extra gluing rules specifying that the boundaries of the two n-gons must coincide gives a triangulation of the sphere, which we denote by $U(\tau_1,\tau_2)$.

For any σ there is an exposed triangulation of the ball extending σ. This allows us to define $t(\sigma)$ to be the minimum number of tetrahedra in any exposed triangulation of the ball extending σ. The following lemma relates triangulations to rotation distance.

Lemma 4: If τ_1 and τ_2 have no diagonal in common, then

$$t(U(\tau_1,\tau_2)) \leq d(\tau_1,\tau_2).$$

Proof: There exists a sequence of $d(\tau_1,\tau_2)$ diagonal flips that changes τ_1 into τ_2. We shall describe how to extract from this sequence an exposed triangulation of the ball extending $U(\tau_1,\tau_2)$ containing $d(\tau_1,\tau_2)$ tetrahedra.

Imagine that there is a planar base with triangulation τ_1 drawn on it. Suppose the first diagonal flip replaces diagonal (a,c) with diagonal (b,d). Create a flat quadrilateral that is the same shape as quadrilateral (a,b,c,d). On the back side of the quadrilateral draw diagonal (a,c). On the front draw diagonal (c,d). Now place the quadrilateral onto the base in the appropriate place with diagonal (a,c) down and (b,d) up. Looking at the base we see a picture of a triangulation which is the result of making the first diagonal flip. For each successive move we create another quadrilateral and place it onto the base. After placing $d(\tau_1,\tau_2)$ such quadrilaterals we will be see τ_2 when we view the base.

The triangulation of the ball that we construct has one tetrahedron for each quadrilateral. The gluing rules are taken from the way the quadrilaterals are stacked. Two triangles are identified with each other if they touch each other in the stack of quadrilaterals.

To finish the proof we need only verify that the resulting triangulation is an exposed triangulation of the ball extending $U(\tau_1,\tau_2)$. The fact that is a triangulation of the ball is made clear by imagining what happens if each quadrilateral is inflated slightly. Each quadrilateral turns into a tetrahedron, and the resulting stack of quadrilaterals is slightly thick in the middle, and is thus homeomorphic to a ball. (This is where we use the assumption that τ_1 and τ_2 have no diagonal in common.) The fact that the triangulation extends $U(\tau_1,\tau_2)$ is obvious because the unglued triangles are exactly those of τ_1 and τ_2.

□.

What the proof of Lemma 4 tells us is that for every sequence of diagonal flips from τ_1 to τ_2 there is an exposed triangulation of the ball extending $U(\tau_1,\tau_2)$. In fact, the same triangulation of the ball may result from many different sequences of moves from τ_1 to τ_2. It is not the case that every exposed triangulation of the ball extending $U(\tau_1,\tau_2)$ comes from a sequence of diagonal flips. In fact, it is possible to construct exposed triangulations of the ball with the property that no tetrahedron touches the boundary on more than one face, whereas in a triangulation obtained by the construction in the proof of Lemma 4 some tetrahedra touch the boundary on at least two faces.

Let $t(n)$ be the maximum over all n-vertex four-connected triangulations of the sphere σ of the quantity $t(\sigma)$.

Lemma 5: $t(n) \leq d(n)$.

Proof: Let σ be an n-vertex four-connected triangulation of the sphere such that $t(\sigma)$ is maximized. By Tutte's theorem [10] this four-connected triangulated graph must have a Hamiltonian circuit. Draw the triangulation on a sphere. Cut the sphere along the edges of the Hamiltonian circuit. This separates the sphere into two disks, each of which is triangulated. Let these two triangulations be τ_1 and τ_2. Now $\sigma = U(\tau_1,\tau_2)$. By the preceding discussion and Lemma 4,

$$t(n) = t(\sigma) = t(U(\tau_1,\tau_2)) \leq d(\tau_1,\tau_2) \leq d(n).$$

□

To make these concepts more concrete consider the two triangulations τ_1 and τ_2 of a hexagon whose diagonals form a triangle. (See Figure 7.) The triangulation obtained by gluing τ_1 and τ_2 together is the boundary of the octahedron. There are six paths of length four between τ_1 and τ_2 (see Figure 4). Each of these paths gives rise to a triangulation of the octahedron. Three different triangulations of the octahedron are obtained in this way. (Each is produced by two different paths from τ_1 to τ_2.) These triangulations are the ones in which a single edge has been added between a pair of opposite vertices. The octahedron cannot be triangulated with fewer than four tetrahedra because no tetrahedron can contact more than two faces of the boundary.

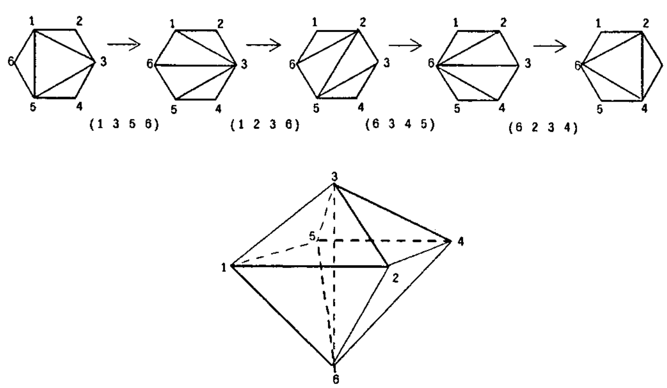

the four tetrahedra: 3 6 1 5, 3 6 2 1, 3 6 4 2, 3 6 5 4

Figure 7. A sequence of diagonal flips and the triangulation that it gives.

By Lemma 5, an upper bound on $d(n)$ is an upper bound on $t(n)$, and a lower bound on $t(n)$ is a lower bound on $d(n)$. In the remainder of this paper we show that $2n-10 \leq t(n)$ (for infinitely many n). Combining this with Lemma 2 gives $t(n) = d(n) = 2n-10$ for infinitely many n.

3. Lower Bounds on $t(n)$

Our approach to deriving accurate lower bounds on $t(n)$ is geometric rather than combinatorial. We convert the combinatorial objects described in the previous section into geometric objects. We then infer properties of the combinatorial objects from the properties of the geometric objects.

Let σ be a triangulation of the sphere that is four-connected Suppose σ is the boundary of a polyhedron P in three dimensional Euclidean space with vertices on the unit sphere. Let T be an exposed triangulation of the ball extending σ. For each tetrahedron Δ of T there is a geometric tetrahedron Δ' whose vertices are the appropriate vertices of P. (By a *geometric tetrahedron* we mean what one normally means by a tetrahedron, a polyhedron with four flat faces. We will use the word *geometric triangle* analogously.) The union of the geometric tetrahedra must contain the polyhedron P. (This will be proven in section 3.2.) Thus the sums of the volumes of the geometric tetrahedra must be at least the volume of P. Let V_Δ be the volume of the largest tetrahedron that can be inscribed in a sphere. Let $vol(P)$ be the volume of P. We know that at least $vol(P)/V_\Delta$ tetrahedra are required to cover P. In other words:

$$\frac{vol(P)}{V_\Delta} \leq t(n)$$

In Euclidean space, this technique does not give interesting bounds because the ratio of the volume of a sphere to the volume of the largest tetrahedron inscribed in the sphere is a small constant. However in hyperbolic space this method does lead to useful results. This is because in hyperbolic space the volume of a tetrahedron is bounded above a constant V_0, while the volume of a polyhedron can grow as a function of the number of vertices. Our problem is thus reduced to finding a polyhedron P with n vertices in hyperbolic space that has large volume.

First we present the necessary fundamentals of hyperbolic geometry. These ideas can be found described in more detail in Coxeter's book [1], Milnor's paper [6], and an expository article by Thurston and Weeks [9].

3.1. Hyperbolic Geometry

In hyperbolic geometry there many lines through a given point parallel to a given line, the sum of the angles of a triangle add to less than 180 degrees, and the circumference of a circle is greater than π times the diameter. There are various ways of mapping hyperbolic space into Euclidean space. These mappings enable us to draw pictures of hyperbolic polyhedra, but these pictures are distorted; two congruent hyperbolic triangles many not look congruent when mapped into Euclidean space.

One mapping of two-dimensional hyperbolic space into the Euclidean plane is called the *upper half-plane model*. In this model all of hyperbolic space is mapped into the upper half of the complex plane (the points with non-negative imaginary parts) and a point at infinity. This mapping is conformal, which means that angles are preserved. The geodesics (straight lines) in hyperbolic space are mapped into the semicircles perpendicular to the real axis, and the vertical lines perpendicular to the real axis. Most of the area of hyperbolic space is mapped into the region near the real axis. See Figure 8.

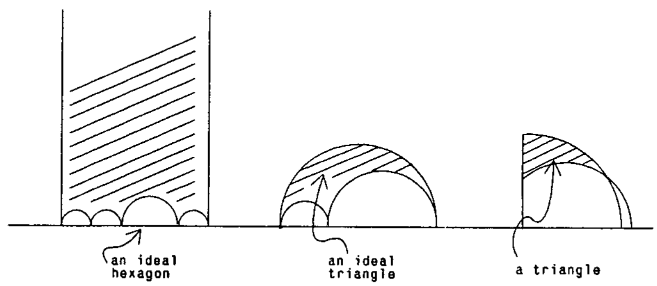

Figure 8. The upper half-plane model of two-dimensional hyperbolic space.

The area of a triangle in hyperbolic space is $\pi - \Sigma$, where Σ is the sum of the angles. An *ideal triangle* is one with three distinct vertices on the real axis or at infinity. All ideal triangles have area π. In fact all ideal triangles are congruent, that is, any ideal triangle can be transformed to any other by a rigid motion. (The rigid motions of the space form a group known as the group of orientation-preserving isometries.)

The upper half-space model of three-dimensional hyperbolic space consists of the complex plane plus all the points above the plane in Euclidean three-space plus a point at infinity. The complex plane plus the point at infinity is sometimes called the *sphere at infinity*. A geodesic in hyperbolic three-space is mapped to a semi-circle perpendicular to the complex plane, or a straight line perpendicular to the complex plane going to infinity. The geodesic surfaces are mapped to hemispherical bubbles orthogonal to the complex plane, or planes orthogonal to the complex plane.

An *ideal hyperbolic tetrahedron* is a tetrahedron in which all the vertices are distinct and on the sphere at infinity. Any hyperbolic tetrahedron can be transformed by a rigid motion to one in which three of the vertices are at 0, 1, and ∞ and the other vertex is at a point z in the complex plane. (This motion is possible because all four of the triangles of the tetrahedron are ideal and any ideal triangle can be moved to any other. Note that despite this fact, not all hyperbolic tetrahedra are congruent.) The tetrahedron then looks like three vertical flat walls above the Euclidean triangle $(0,1,z)$ bounded below by part of a hemispherical bubble.

The hyperbolic cross-section of the vertical chimney, in the hyperbolic metric, scales in a way that decreases with increasing height. (Most of the volume of hyperbolic space is near the complex plane.) It can be seen by integrating that the volume of an ideal tetrahedron is finite. Let this volume be denoted by $v(z)$. There are explicit formulas for $v(z)$, from which it can be seen that the maximum is attained at the point $z = \omega$, where ω is defined as

$$\omega = e^{2\pi i/6}$$

(See [6] for a discussion of how to compute hyperbolic volumes.) The tetrahedron of maximum volume is the most symmetrical one. Its base triangle $(0,1,\omega)$ is equilateral, its dihedral angles are all 60 degrees, and its volume is $v(\omega) = V_0 = 1.0149\cdots$

3.2. The Volume of Hyperbolic Polyhedra

In order to follow through with the approach outlined at the beginning of section 3, we must justify our contention that if P is a polyhedron, σ is its boundary, and T is an exposed triangulation of the ball extending σ then the geometric tetrahedra of T must cover P. To do this we need to define the volume of non-simple polyhedra, ones in which the surface may be self-intersecting. (The ideas we develop are the same ones used to prove Brower's fixed point theorem.)

Let σ be a triangulation of the sphere. For concreteness think of σ as though it is embedded in the sphere in some particular way. For any mapping of the vertices of σ to distinct points in three-dimensional hyperbolic space (or Euclidean space for the purposes of this discussion), there is a continuous map f from the sphere into hyperbolic space that 1) maps the vertices of σ to the appropriate places, and 2) maps every triangle of σ one-to-one onto a geometric triangle in hyperbolic space. In other words f maps the surface of the sphere into the surface of some hyperbolic polyhedron. Let the symbol $f(\sigma)$ denote the image of the triangles of σ mapped by f into hyperbolic space.

Imagine that a circulation direction has been chosen on each triangle of σ so that when that triangle is viewed from outside of the sphere the circulation is clockwise. These circulations when mapped by f determine the circulations of the triangles of $f(\sigma)$. This set of triangles and their circulations enable us to define the number of times the polyhedron "wraps around" a point x, and from this we define the volume of the polyhedron.

Consider any continuous path from a point y to a point x. (Assume that the path does not pass through the boundary of any triangle, and that x and y are not in any triangle.) If the path passes through exactly one triangle and the circulation direction is clockwise when looking along the path from y to x, then x is deemed to be inside of the polyhedron, and f is said to wrap around x once.

In general, the *wrapping number* (or *degree*) of $f(\sigma)$ about a point x with respect to a point y (x and y not in any triangle), $w(f(\sigma),x,y)$ is defined as follows:

> Choose some continuous path p from y to x that does not pass through the edges or vertices of any triangle. Each time the path passes through a triangle the circulation direction of the triangle (looking along p) is either clockwise or counterclockwise. The wrapping number at x is the number of times p passes through a triangle in a clockwise fashion, minus the number of times p passes through a triangle in a counterclockwise fashion.

This number only depends on $f(\sigma)$, x and y, and is independent of the path from y to x. This is because the circulations were chosen consistently on the triangles of σ, and σ is a triangulation of the sphere. Furthermore a different choice of the point y will add an integer (independent of x) to the value of the wrapping function. There is a natural way to eliminate the dependence of the wrapping function on y. For a particular choice of y the wrapping function divides space into regions of constant wrapping function value. At least one of these regions must have infinite volume, since the space has infinite volume. The natural wrapping function is one where y is chosen in the region of infinite volume. The wrapping function is thus zero in the infinite region. Denote this wrapping function by $w(f(\sigma),x)$. (If there is more than one infinite region, or if there are infinitely many regions, then the wrapping number is undefined.)

The volume of the polyhedron to which f maps the sphere is defined to be the integral over hyperbolic space of $w(f(\sigma),x)dV$, where dV is the hyperbolic volume element. This generalized definition of volume agrees with the intuitive one for a simple polyhedron in which the wrapping function only takes on the values of zero or one. It also gives us a notion of volume for more complicated collections of triangles, and facilitates the proof of the following lemma.

Lemma 6: Let P be a polyhedron in hyperbolic space whose boundary, σ, is a triangulation of the sphere. Let $vol(P)$ be the volume of P. Then

$$\frac{vol(P)}{V_0} \leq t(\sigma).$$

Proof: Let T be an exposed triangulation of the ball extending σ with $t(\sigma)$ tetrahedra. To be concrete, assume that T and σ have been embedded into the ball in some particular way. Let f be a map from the vertices of σ to the vertices of P as described above with the additional property that it maps all of the triangles of all of the tetrahedra of T to geometric triangles in hyperbolic space. (It is clear that such a function exists because in the embedding of T into the ball each tetrahedron is homeomorphic to a geometric tetrahedron.)

For each tetrahedron Δ of T choose a circulation on its triangles that is clockwise when looking at it from outside of it in its embedding into the ball. For each such tetrahedron there is a wrapping function $w(f(\Delta),x)$ in hyperbolic space which is ± 1 inside the image of Δ and zero outside. The following equation is the key fact we need about wrapping numbers.

$$w(f(\sigma),x) = \sum_{\Delta \varepsilon T} w(f(\Delta),x)$$

To prove this, we choose a point y outside of all the tetrahedra $f(\Delta)$, choose a particular path from y to x, and evaluate both sides of the equation. The left hand side is the sum of the circulations of the triangles of $f(\sigma)$ punctured by the path. The right hand side is the same except it includes more triangles. It includes all the triangles of $f(\sigma)$ plus all the triangles that come from the interior of the ball. However, those that come from the interior of the ball come in pairs, one from each of the two tetrahedra bounding that triangle. The triangles of a pair are the same except that they have opposite circulation directions. Their effects on the wrapping number cancel out. Only the triangles on the boundary matter. This justifies the equation. Figure 9 illustrates the analogous situation in two-dimensions.

Now we can bound the volume of P.

$$\begin{aligned} v(P) &= \int w(f(\sigma),x)dV \\ &= \int \sum_{\Delta \varepsilon T} w(f(\Delta),x)dV \\ &= \sum_{\Delta \varepsilon T} \int w(f(\Delta),x)dV \\ &\leq \sum_{\Delta \varepsilon T} V_0 = t(\sigma)V_0. \end{aligned}$$

□

We have now reduced the problem of finding lower bounds on $t(n)$ to that of finding n vertex hyperbolic polyhedra with large volumes. The remainder of this section is devoted to constructing such polyhedra.

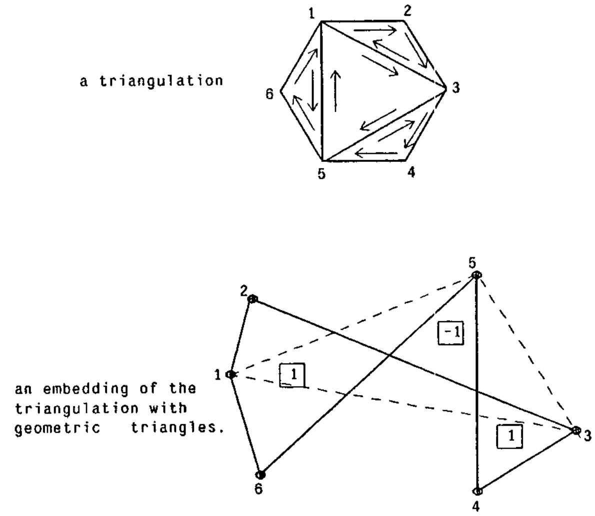

Figure 9. A two-dimensional version of the wrapping function used in the proof of Lemma 6.

3.3. First Examples: $2n - O(n^{1/2})$

There is a tessellation of hyperbolic space consisting of copies of the simplex of maximal volume. This tessellation can be constructed by choosing a maximal simplex to start with, reflecting it through its faces, reflecting these through their faces, etc. For any finite union of these tetrahedra whose boundary is a sphere, we obtain a polyhedron. The triangulation we have is automatically minimal since all the simplices are disjoint and have maximal volume.

Consider the special case when all the simplices have a common vertex. We may assume that this point is the point at infinity in the upper half space model. Each of the tetrahedra lies above an equilateral triangle in the tessellation of the plane by equilateral triangles. Any set of triangles whose union is homeomorphic to a disk will determine such a polyhedron. Consider the case when the polygon is hexagonal, with k edges on a side. The hexagon has $6k^2$ triangles, hence the polyhedron has $6k^2$ tetrahedra. (See Figure 10.) The hexagon has $3k^2+3k+1$ vertices, so the polyhedron has $3k^2+3k+2$ vertices (including the one at infinity.) In particular we obtain

$$t(n) \geq 2n - O(n^{1/2}).$$

Note that we actually get explicit lower bounds for each n, not just those n of the form $3k^2+3k+2$ by using other triangulations.

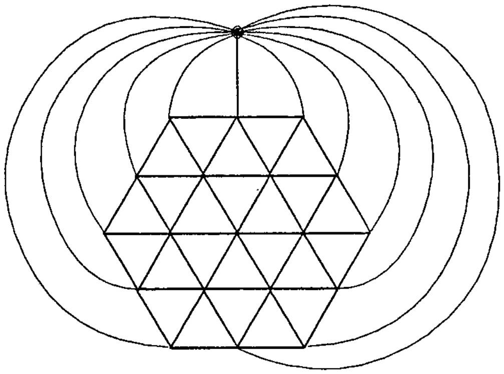

Figure 10. The boundary of the polyhedron used to show that $t(n) \geq 2n - O(n^{1/2})$.

3.4. Better Lower bound: $2n - O(\log n)$

To construct polyhedra that require more simplices for a given number of vertices we must eliminate the vertex of high degree. (Roughly speaking, polyhedra in which the vertices are spread over the sphere at infinity as uniformly as possible have the largest volumes.) A natural sequence of triangulations for this purpose can be derived from a regular icosahedron. Divide each face of the icosahedron into k^2 equilateral triangles, giving $20k^2$ triangles in all, and $n = 10k^2 + 2$ vertices.

This gives a combinatorial definition of a likely triangulation. How can it be mapped into hyperbolic space so as to enclose a large volume? The Riemann mapping theorem gives us a way to do this. Corresponding to the icosahedron, there is a subdivision of the sphere into triangles bounded by segments of great circles, obtained by projecting the edges of the icosahedron out to the sphere. The Riemann mapping theorem implies that there is a unique *conformal* map of the faces of the icosahedron to the spherical triangles, sending vertices of the icosahedron to the corresponding vertices of the spherical triangles. By symmetry, these maps determined on individual triangles piece together to give a map h of the entire surface of the icosahedron to the sphere. This map is conformal everywhere except at the vertices of the icosahedron. Note that it is conformal even on the edges of the icosahedron because they can be flattened out (locally) in the plane. Define the ideal hyperbolic polyhedron $P(k)$ to have its vertices at those places on the sphere at infinity in hyperbolic space to which h maps the vertices of the subdivided icosahedron.

Lemma 7: The volume of $P(k)$ *is* $2nV_0 - O(\log n)$, *where* $n = 10k^2 + 2$.

Proof: To make the estimate of volume, pick a vertex of $P(k)$ of order 6 which is as far as possible from vertices of order 5, and arrange this vertex to be at infinity in the upper half-space model. (This is a rigid rotation of $P(k)$ in hyperbolic space.) Triangulate $P(k)$ as the union of cones from the vertex at infinity to the triangles with finite vertices. (This is called a *cone-type triangulation.*) Call a vertex of $P(k)$ "bad" if it is a vertex of the icosahedron or if it is the vertex mapped to infinity. Now h can be thought of as a map from the icosahedron to **C** (the complex plane) that is conformal everywhere except at the bad vertices. For large k the triangles far away from bad vertices get mapped by h to triangles that are nearly equilateral (because h is conformal). Figure 11 shows how the vertices of $P(k)$ near a vertex of the icosahedron get mapped to the complex plane by h. If the vertices were vertices of true equilateral triangles then the tetrahedra formed by coning them to infinity would all be congruent to the tetrahedron of maximal volume. We must show that the deficit caused by the fact that the triangles are not quite equilateral is small.

The shape of a triangle Δ with vertices p,q,r is described by the complex number $s(\Delta) = \dfrac{(r-p)}{(q-p)}$. (The triangle $(0,1,s(\Delta))$ is congruent to Δ.) The volume of the hyperbolic tetrahedron $C(\Delta)$ formed by coning Δ to the point at infinity depends only on the shape parameter of Δ. The volume deficit of $C(\Delta)$, defined as $V_0 - vol(C(\Delta))$ (where V_0 is the maximum volume of a tetrahedron), satisfies

$$V_0 - vol(C(\Delta)) = K_1(|s(\Delta) - \omega|^2) + O(|s(\Delta) - \omega|^3)$$

for some constant K_1. Our goal is to show that the cumu-

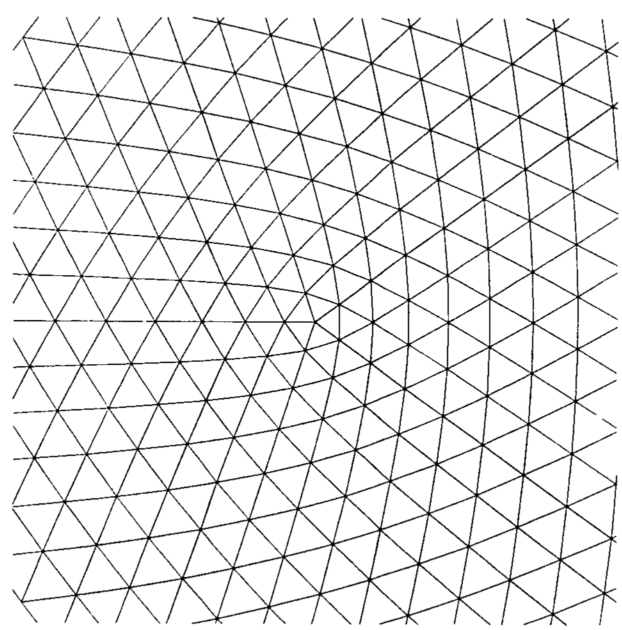

Figure 11. The vertices of $P(k)$ near a vertex of the icosahedron.

lative volume deficit is $O(\log(n)) = O(\log(1/\varepsilon))$ where ε is the side length of a triangle in $P(k)$. We will now digress to evaluate the deficit of a triangle obtained by applying an arbitrary complex analytic function f to the vertices of an equilateral triangle. Subsequently we will apply this result using the map h.

Consider the equilateral triangle Δ_t with vertices 0, t, and ωt in the complex plane, and suppose that f is a holomorphic (complex analytic, or conformal) embedding of the disk of radius $R > |t|$ about 0 into the complex plane. How can we estimate the volume deficit of the tetrahedron spanned by ∞ together with f of the vertices of Δ?

In the general, we may change f by postcomposing with a translation so that f fixes the origin. Let us expand f as a power series

$$f(z) = az + bz^2 + cz^3 + \cdots$$

and solve for the coefficients of the power series expansion of the shape parameter $s(t) = f(\omega t)/f(t)$ of the image triangle. We can write

$$\begin{aligned} s(t) &= \omega + At + Bt^2 + \cdots \\ &= \frac{a\omega t + b\omega^2 t^2 + c\omega^3 t^3 + \cdots}{at + bt^2 + ct^3 + \cdots} \end{aligned}$$

and solve for the coefficients. We obtain

$$\omega b + Aa = b\omega^2$$

and

$$\omega c + Ab + Ba = c\omega^3.$$

The term of real interest is

$$A = \frac{-b}{a} = -\frac{f''(0)}{2f'(0)}$$

since $\omega^2 - \omega = -1$. The next term is

$$B = \frac{1}{a}(c(\omega^3 - \omega) - \frac{-b}{a}b) = -(\omega+1)\frac{c}{a} + \frac{b^2}{a^2}.$$

Thus, the shape of the image triangle is

$$s(t) = \omega - t\frac{f''(0)}{2f'(0)} + O(t^2).$$

How does the error term depend on f? First, we claim that the error is uniformly bounded by $O(t^2)$ independent of f defined on a fixed disk of radius R. In fact, the set of all holomorphic embeddings of the disk of radius R into $\mathbf{C}$ is compact in the appropriate topology, that is, any sequence of embeddings has a subsequence which converges to an embedding. The errors could not get worse and worse, or else the limit function would not have an estimate of the form $O(t^2)$.

The dependence of the error term on R can now be easily deduced. A disk of radius R can be mapped to a disk of radius S by a complex affine map --- that is, a complex linear map followed by a translation. The parameter t is multiplied by the ratio of the radii of the disks under such an affine map. Consequently, the error term above is $O((t/R)^2)$.

Now we return to the estimate of the volume deficit. Suppose we have a triangle in some subdivision of the icosahedron that has no vertices in common with the icosahedron, and no vertex mapped to infinity. Let R be the distance of the triangle from the nearest bad vertex in the metric of the icosahedron, and let ε be the size of a side of the triangle (in the metric of the icosahedron). Then its volume deficit is $K_2(\varepsilon^2 |\frac{h''}{2h'}|^2) + O((\frac{\varepsilon}{R})^3)$.

Since ε^2 is proportional to the area of the triangle, the total volume deficit can be approximated by an integral.

$$\int_{\text{good triangles}} K_3 |\frac{h''}{2h'}|^2 dA + O(1).$$

By good triangles, we mean those that have no bad vertex. The contribution of the bad triangles (those with a bad vertex) is only a constant, and has been included in the $O(1)$ error term. The contribution of the part of the icosahedron that is farther than a fixed distance ε_0 away from a bad vertex is also bounded by a constant. This is because the integrand is continuous and bounded except near the bad vertices. The only contribution left to evaluate is that of the annular regions of inner radius ε and outer radius ε_0 centered on the bad vertices. (Where ε is the triangle mesh size.)

Near the bad vertices h behaves like z^β where $\beta = 6/5$ if the bad vertex is an icosahedron vertex and $\beta = -1$ if the bad vertex is the one mapped to infinity. (The local coordinates are chosen so that the bad point is at the origin.) The entire deficit is estimated to within an additive constant by the sum (over bad vertices) of the integrals of $K_3 \frac{1}{4}(\beta - 1)^2 \frac{1}{|z|^2}$, over annular regions centered at the bad vertices with fixed outer radius ε_0 and inner radius approximately equal to the mesh size ε. The value of each of these integrals is

$$K_3 \frac{1}{4}(\beta - 1)^2 2\pi(\ln\varepsilon_0 - \ln\varepsilon).$$

Since $-\ln\varepsilon = \ln(1/\varepsilon) = O(\log n)$, we have bounded the deficit by $O(\log n)$. This completes the proof that $vol(P(k)) = nV_0 - O(\log(n))$.

□

3.5. The structure of minimal triangulations

We will now apply the fact that $P(k)$ has a volume deficit which is small compared to the number of tetrahedra to completely determine the minimal extensions of $P(k)$ to the ball, provided k is sufficiently large.

Theorem 1: For sufficiently large k the exposed triangulation of the ball extending the boundary of $P(k)$ having the minimum number of tetrahedra is a cone-type triangulation.

Proof: Suppose that T is any minimal exposed triangulation of the ball extending the boundary of $P(k)$. For any fixed ε, at most $O(\log n)$ of the simplices of T can have volume less than $V_0-\varepsilon$. In particular, only $O(\log n)$ of its simplices can touch the boundary on two faces, since such simplices have volume roughly $(2/3)V_0$. For every simplex that does not touch the boundary on any face, there must be one that touches on two faces, so the number of such simplices is also bounded by $O(\log n)$. Furthermore only $O(\log n)$ of the simplices can have four or more edges common with the boundary for their deficit is also large.

A "good" simplex is one with exactly one of its faces and three of its edges on the boundary. A simplex is "bad" if it is not good. From the preceding paragraph we conclude that there are $O(\log n)$ bad simplices, thus there are at least $2n-O(\log n)$ good simplices. Let m be a map from the good faces of the polyhedron to the vertices of the polyhedron. (A face of $P(k)$ is good if it is covered by a good simplex.) The vertex $m(f)$ is the fourth vertex of the simplex with face f on the boundary of $P(k)$.

If f and g are adjacent good faces of the boundary, and if $m(f)\neq m(g)$, then there is a bad simplex which contains the common edge of f and g. (It is bad because if one of its triangles is on the boundary then four of its edges must be on the boundary.) In this situation call the common edge of f and g a "bad" edge; there are at most $O(\log n)$ bad edges, because for each one there is a bad simplex. Each component of the union of bad faces and bad edges is contained in a disk of radius $\dfrac{O(\log n)}{k}=\dfrac{O(\log k)}{k}$ with respect to the metric of the icosahedron. Therefore the union of bad faces and edges is contained in a disjoint union B of such disks.

Let G be the union of those good tetrahedra whose face on $P(k)$ does not intersect B. Since the total number of bad faces and edges is considerably smaller than the diameter of the sphere, there is one big component of G and possibly a bunch of smaller components. We throw out the smaller components from G. G agrees with a cone-type triangulation $C(v^*)$, coned to a single distinguished vertex v^*, and G contains most of the tetrahedra of T.

What remains are a number of blemishes -- components where T does not agree with $C(v^*)$, each having a known boundary triangulation but a mysterious interior triangulation with at most $O(\log(k)^2)$ simplices. To complete the analysis, we will show that the triangulations of these blemishes in fact must also agree with $C(v^*)$, given that they are minimal triangulations.

The blemishes can be sorted into three types, depending on their intersection with $P(k)$. Any blemish which does not exactly fit these types can be enlarged until it does, using portions of the known triangulation G, provided k is large enough. Here are the three kinds of boundary triangulations for possible blemishes, up to isomorphism:

(a) A portion of the triangulation of the plane by equilateral triangles, with one extra vertex v^* coned to its boundary.

(b) A triangulation obtained from a regular pentagon by first subdividing into five triangles, then subdividing these into congruent subtriangles, then coning its boundary to one extra vertex v^*.

(c) A blemish coming from the vicinity of v^* on $P(k)$, which is obtained from a portion of the triangulation of the plane with equilateral triangles by coning its boundary to a vertex v^* of the triangulation. Note that in this case, the boundary of the blemish is not a sphere; it may be chosen to be homeomorphic with a sphere with the north pole and the south pole identified.

Case (a) is easiest to take care of. The triangulation $C(v^*)$ of the interior of the blemish has all its interior edges of order 6. This triangulation is isomorphic to a portion of the triangulation of hyperbolic space by simplices of maximal volume, as described in section 3.3. Any triangulation other than the one where all the simplices are coned to v^* would involve a simplex of less than maximum volume, therefore at least one extra simplex would be required.

Case (c) is also easy to take care of, with a simple observation. We may embed the boundary of the blemish first so all its vertices except v^* agree with the vertices of the equilateral triangulation of the plane, but v^* is mapped to the point at infinity in upper half space. Now each simplex of the cone triangulation $C(v^*)$ has maximal volume. Again, this is must be a minimal triangulation, and furthermore it is the only possible minimal triangulation extending the boundary of the blemish.

What remains to examine are the possible 12 blemishes of type (b). We can embed the boundary of any such blemish in hyperbolic space as follows. The vertex v^* is mapped to the point at infinity. We slit open the pentagon along a line segment from its center to a vertex, lay out the result to a portion of the triangulation of the plane by equilateral triangles with the vertex of the slit at the origin, and raise to the 6/5 power to close up the slit again.

In this embedding, the simplices of the cone triangulation $C(v^*)$ are not regular, but they are nearly regular, and numerics comes to the rescue. Let j be the number of edges in the subdivision of each edge of the pentagon, so that the triangulation of the pentagon has $5j^2$ triangles. Then when the constants are evaluated in the method above, the volume deficit is estimated by

$$\mathrm{Def}(j) = \frac{\pi}{50}\log(j) + O(1).$$

Here is a table giving values of the volume deficit, the estimated deficit, and the error.

j	Triangles	Deficit	Estimated	Error
1	5	.087935	.000000	.087935
2	20	.122353	.036293	.086060
4	80	.158141	.072586	.085554
8	320	.194304	.108879	.085424
16	1280	.230564	.145172	.085391
32	5120	.266849	.181466	.085383
64	20480	.303140	.217759	.085381

The deficit increases by about .0363 each time j is doubled. The value of j when the deficit finally reaches 1 is therefore greater than 2^{25}. At that time, when the direct volume estimate finally would admit one tetrahedron less, the number of tetrahedra would be more than $5{\cdot}2^{50}$.

A simple recursive argument finishes the proof. We will show that the only minimal triangulation of the type (b) blemish is $C(v^*)$. In a minimal triangulation there can only be a tiny proportion of tetrahedra which touch the boundary at at two adjacent triangles of the pentagon, since the deficit is quite small. In fact if the number of triangles is less than $5{\cdot}2^{50}$ then the numerical calculation above shows that the minimal triangulation must be $C(v^*)$. This is because any simplex that does cone from a triangle to vertex v^* has a deficit of about 1/3, and that is more than the total deficit allowed.

If the number of triangles is much more than $5{\cdot}2^{50}$ then there may be some simplices that do not touch any triangle of the pentagon, but only a tiny fraction of the simplices can do this. Almost all the triangles of the pentagon around its boundary are faces of tetrahedra with the fourth vertex at v^*, since otherwise the tetrahedra whose faces are the "vertical" faces from an edge along the pentagon to v^* do not have a face within the pentagon. As before, it follows that all but a tiny fraction of the triangles in the blemish are faces of tetrahedra which have a fourth vertex at v^*.

Thus, the blemish is triangulated by a cone-type triangulation to v^* except for small type (a) and type (b) subblemishes within it. The previous argument shows that the only way to triangulate the type (a) subblemished minimally is to cone to v^*. The type (b) subblemish is smaller than the one we started with, so by induction the only way to triangulate it minimally is to cone to v^*.

□

If the sizes of the deficits were not quite so small, this argument would not work for small j, and we could only deduce that the minimal number of simplices was within an additive constant of the number for $C(v^*)$. For example, consider the sequence of subdivisions of the tetrahedron or of the octahedron instead of the icosahedron. For a tetrahedron, the numbers do not work out: in fact, the cone-type triangulations of subdivisions of the tetrahedron can be improved by first knocking out the corners. This fact manifests itself in a table in which the deficit is more than 1 even for small values of j. For an octahedron the numbers do work, and we could have used this instead of the icosahedron to prove a similar theorem.

The following theorem is an immediate consequence of Theorem 1.

Theorem 2: $t(n) = 2n-10$ *for sufficiently large values of* n *of the form* $n=10k^2+2$.

4. Remarks and Questions

It remains open what size n must be before $2n-10$ is the maximum distance between triangulations. We conjecture that this is the correct answer for all $n>12$.

We can prove a slightly stronger result than Theorem 2 by chopping the surface of the icosahedron up into triangles in different ways. In this way we can extend the result to some values of n not of the form $10k^2+2$. For positive integers i and j we can chop each face of the icosahedron into three j by j triangles and an i by i triangle. (This technique is like Escher's way of generating one tessellation of the plane from another by changing the shape of the pieces.) The number of faces in the resulting triangulation of the ball is $20(i^2+3j^2)$. The number of vertices is $10i^2+30j^2+62-60i$ if $i\geq j$ and is $10i^2+30j^2+62-30j$ if $i\leq j$. Our proof is valid for all sufficiently large values of n of this form.

Problem: find a purely combinatorial proof.

Acknowledgements

We thank Mark Manasse for his contribution to the program that efficiently found shortest paths in the rotation graph. Dick Jerrard, Greg Nelson and Jim Reeds made suggestions that improved the exposition of these results.

References

[1] H. M. S. Coxeter, *Non-Euclidean geometry* Univ. Toronto Press, Toronto, 1942.

[2] K. Culik and D. Wood, "A note on some tree similarity measures," Info. Process. Lett. 15 (1982), 39-42.

[3] A. K. Dewdney, "Wagner's theorem for torus graphs," Discrete Math. 4 (1973), 139-149.

[4] D. E. Knuth, *The Art of Computer Programming, Vol. 3: Sorting and Searching,* Addison-Wesley, Reading, MA, 1973.

[5] C. W. Lee, "Catalan polytopes and triangulations of the n-gon," Technical Report Number 36, Abteilung Fur Mathematik der Ruhr-Universitat Bochum, December 1984.

[6] J. Milnor, "Hyperbolic geometry: the first 150 years," Bulletin of the American Mathematical Society 6, 1 (January 1982), 9-24.

[7] D. D. Sleator and R. E. Tarjan, "Self-adjusting binary search trees," J. Assoc. Comput. Mach. 32, 3 (July 1985), 652-686.

[8] R. E. Tarjan, *Data Structures and Network Algorithms,* Society of Industrial and Applied Mathematics, Philadelphia, PA 1983.

[9] W. P. Thurston, J. R. Weeks, "The mathematics of three-dimensional manifolds," Scientific American 251, 1 (July 1984), 108-120.

[10] W. T. Tutte, "A theorem on planar graphs," Trans. Am. Math. Soc. 82, (1956), 99-116.

[11] K. Wagner, "Bemerkungen zum vierfarbenproblem," Jber. Deutsch. Math.-Verein. 46 (1936), 26-32.