Circle Quotients and String Art

Richard Evan Schwartz *

April 5, 2001

1 Introduction

Let S^1 be the unit circle. If $\sim$ is an equivalence relation on S^1 we can form the topological quotient space $Q = S^1/\sim$. We call Q a *circle quotient.*

There is a hyperbolic geometry interpretation of circle quotients. We think of the hyperbolic plane, $\boldsymbol{H}^2$, as the open unit disk in the Euclidean plane. The ideal boundary of $\boldsymbol{H}^2$ is S^1. Every pair of distinct points in S^1 determines a unique hyperbolic geodesic having these points as endpoints. Given $\sim$ we produce a collection Γ of geodesics as follows. A geodesic belongs to Γ iff its endpoints are equivalent. Conversely, if a collection Γ of geodesics satisfies a simple condition then there is an equivalence relation $\sim$ which determines Γ as above. In this case we write $Q(\Gamma) = S^1/\sim$.

Sometimes the hyperbolic geometry interpretation is natural because the circle quotient arises in connection with Kleinian groups [**T**]. For instance,

1. If Γ is the set of lifts of a simple closed geodesic on a closed hyperbolic surface, then $Q(\Gamma)$ is homeomorphic to an infinite union of circles, tangent to each other in a tree-like pattern. This is the *Mickey Mouse example.*

2. If Γ is the set of all lifts of a binding on a closed surface then $Q(\Gamma)$ is homeomorphic to S^2, according to a theorem of R.L. Moore. A *binding* is a union of two simple closed geodesics, whose complementary regions are all homeomorphic to disks. This example arises in connection with doubly degenerate limits of quasifuchsian groups.

* Supported by N.S.F. Research Grant

In this paper we will study perhaps the simplest examples of circle quotients which are based on self-intersecting closed curves on a (punctured) hyperbolic surface. Unlike the examples above, many of the quotients we consider are not planar. We encountered the prototypical example while studying *complex hyperbolic* Kleinian groups in [**S1**] and it is the complex projective geometry of S^3, rather than the conformal geometry of S^2, which determines the structure of the corresponding circle quotient. (Essentially no geometry of this sort enters into this paper, however.)

Let Σ be the thrice-punctured sphere, equipped with its usual finite area complete hyperbolic metric. Figure 1.1 shows the *commutator curve* γ on Σ. Here Σ is represented as a twice-punctured plane. We mean for γ to be a closed geodesic. We have the universal covering map $\boldsymbol{H}^2 \to \Sigma$. Let Γ be the set of lifts to $\boldsymbol{H}^2$ of γ. The right hand side of Figure 1.1 shows a sketch of Γ. This is our *prototypical example*.

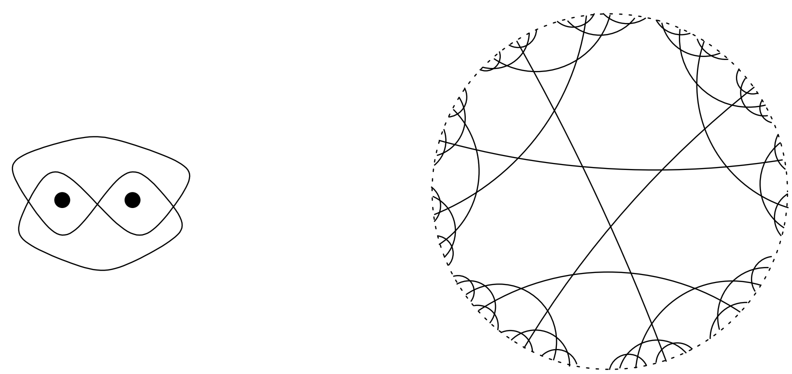

Figure 1.1

Here is a generalization. A *horodisk* in $\boldsymbol{H}^2$ is a disk tangent to S^1 and otherwise contained in $\boldsymbol{H}^2$. A *k-flower* is a union of $k \geq 3$ horodisks, having pairwise disjoint interiors, such that each is tangent to two others, and such that the union has k-fold hyperbolic rotational symmetry. If k is odd (respectively even) we call the flower *odd* (respectively *even*). We say that an *infinite horodisk packing* is an infinite collection $H = \{H_j\}$ of horodisks, having pairwise disjoint interiors, such that all the complementary regions are surrounded by flowers. Figure 1.2 shows part of the horodisk packing in which the complementary regions are all surrounded by 3-flowers.

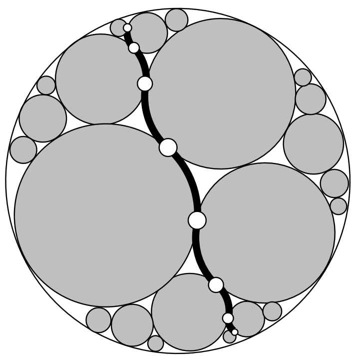

Figure 1.2

Say that a *slalom curve* of H is a regular C^1 bi-infinite path, contained in ∂H, which makes an inflection point at every opportunity. One example is drawn in Figure 1.2. We will see in Corollary 5.2 that slalom curves always have two endpoints in S^1, just like geodesics. If H is the packing in Figure 1.2, and we replace each slalom curve of H by its *geodesic representative* which has the same endpoints, we recover our prototypical collection of geodesics.

In general, we start with a horodisk packing H and consider the collection Γ_H of geodesic representatives of slalom curves of H. It follows from Corollary 5.2 that the relation induced by Γ_H is an equivalence relation. We call $Q(\Gamma_H)$ a *horodisk quotient*. We usually write $Q(H) = Q(\Gamma_H)$.

Our main goal is to describe how to visualize these quotients. We will even explain, in §7, how to build approximations to many of them out of string in a canonical and algorithmic way.

Say that a *tetrahedron space* is a finite collection of tetrahedral subsets of $\boldsymbol{R}^3$ such that every two of the tetrahedra are either disjoint from each other, or intersect in a common vertex, or intersect in a common edge. If Π_0 and Π_1 are tetrahedron spaces we write $\Pi_0 \to \Pi_1$ if each tetrahedron μ_1 of Π_1 is contained in some tetrahedron μ_0 of Π_0. We also insist that $\mu_1 \cap \partial\mu_0$ is either empty, or a common vertex, or a common edge. We say that a *nested sequence* of tetrahedron spaces is a sequence of the form $\Pi_0 \to \Pi_1 \to \Pi_2 \to \Pi_3 \ldots$

Theorem 1.1 (Main Result) *Let H be a horodisk packing, with associated quotient $Q(H)$. There is a nested sequence $\{\Pi_n\}$ of tetrahedron spaces, such that $\Pi = \bigcap_{n=0}^{\infty} \Pi_n$ is homeomorphic to $Q(H)$. Also, the homeomorphism $\Psi : Q(H) \to \Pi$ conjugates Is(H) to a subgroup of $PL(\Pi)$.*

Is(H) is the orientation preserving hyperbolic symmetry group of H. Each element of Is(H) induces a canonical self-homeomorphism of $Q(H)$. In the

Main Result we think of Is(H) as acting on $Q(H)$ in this way. PL(Π) is the group of self-homeomorphisms of Π which extend to piecewise linear maps in a neighborhood of Π.

We wonder about the extent to which the Main Result can be transcribed into a more rigid geometric situation. For instance, our Main Result combines with the work in [**S1**] to show

Corollary 1.2 *The prototypical horodisk quotient is homeomorphic to the limit set of the last discrete complex hyperbolic ideal triangle group. At the same time, there is an embedding of the prototypical horodisk quotient into the* 3*-sphere so that the complement admits a complete metric of constant negative curvature.*

We will discuss this corollary now, but not elsewhere in the paper. Even though we did not know, in [**S1**], that the limit set in question was homeomorphic to the prototypical horodisk quotient, we did have have a concrete description of it. Here we just observe that our description of the prototypical horodisk quotient is identical to our description of the limit set in [**S1**], once the geometric category is changed. The complex projective maps in [**S1**] are changed to piecewise affine maps here. The *hybrid spheres* of [**S1**] are changed to *marked tetrahedra* here. Indeed, the initial motivation for writing this paper was to have a simpler geometric setting in which to elaborate the structure of the neat limit set. In [**S1**] we also proved that the orbifold at infinity for our group was commensurable to the Whitehead link complement, a well known manifold which admits a complete hyperbolic metric of finite volume. This means that the complement of the limit set, in S^3, has a complete hyperbolic metric (of infinite volume).

We now can prove that a vast class of horodisk quotients can be embedded into S^3 so that their complements admit complete metrics of constant negative curvature. See [**S2**].

The Main Result tells us a lot about embedding finite graphs into horodisk quotients. Using the notation from the Main Result, we form the *incidence graph* $G(\Pi_n)$ as follows. $G(\Pi_n)$ has one red vertex placed at the center of each tetrahedron of Π_n and one blue vertex placed at each point in $\boldsymbol{R}^3$ which is a vertex of a tetrahedron in Π_n. We join a red vertex to a blue vertex by an edge iff the corresponding points in space are the center and vertex of the same tetrahedron.

Corollary 1.3 *For any n one can embed into $Q(H)$ a graph $\tilde{G}_n$ which has $G(\Pi_n)$ as a quotient graph.*

A *quotient graph* is obtained by collapsing to points some of the edges in the original graph. A planar graph has only planar quotients.

We will see in §6 that $G(\Pi_n)$ is non-planar for sufficiently large n if H has at least one odd flower. It will follow that $Q(H)$ contains a non-planar embedded graph in this case. It is not hard to see, if H has all even flowers, that one can partition the slalom curves into two sets, so that no two in the first set intersect each other and no two in the second set intersect each other. From here, it is well-known that $Q(H)$ must be planar. In §7.1 we will sketch a self-contained proof, based on our Main Result. In short,

Corollary 1.4 *$Q(H)$ is planar if and only if all the flowers of H are even.*

Here is an overview of the paper. In §2 we introduce *finite horodisk packings* and construct the *horodisk packing graphs*, which are finite graphs associated to the finite packings. The graphs, properly considered as metric spaces, are finite approximations to horodisk quotients.

In §3 we define *marked rectangle spaces*. These are metric spaces, made by gluing together finitely many Euclidean rectangles, which contain isomorphic copies of the horodisk packing graphs constructed in §2.

In §4 we define certain tetrahedron spaces, which we call *marked tetrahedron spaces*. These spaces contain PL embedded copies of the marked rectangle spaces constructed in §3.

In §5 we take the limits of the constructions in §2-4 to prove the Main Result.

In §6 we prove Corollary 1.3. We also prove that $G(\Pi_n)$ is non-planar, for large n, when the associated horodisk packing has at least one odd flower.

In §7 we show how our constructions can be simplified when the horodisk packing either has all even flowers or all odd flowers. The first case leads to a planarity proof and the second case leads to a method for building the corresponding circle quotients out of string. Since the string art topic is rather whimsical, we will only sketch a proof that it works.

I would like to thank Martin Bridgeman, Peter Doyle, Bill Goldman, and David Epstein for conversations on topics relating to this paper. I would also like to thank the anonymous referee, who made a great number of helpful comments and suggestions.

2 Finite Horodisk Packings

2.1 Basic Definitions

Bounded Interstices: Let $\boldsymbol{H}^2$ be the hyperbolic plane. We define *horodisks* and *flowers* in $\boldsymbol{H}^2$ exactly as in the introduction. A *bounded interstice* is the closure of the bounded component of $\boldsymbol{H}^2 - F$, where F is a flower. The *center* of the interstice is the fixed point of the hyperbolic rotation which stabilizes the flower. An *interstitial vertex* is a point of tangency between two of the horodisks in the defining flower. A *bounded interstitial arc* is an arc of one of the horocircles in the defining flower, which connects two interstitial vertices.

Unbounded Interstices: An *unbounded interstice* is the closure, in $\boldsymbol{H}^2 \cup S^1$, of a connected component of $(\boldsymbol{H}^2 \cup S^1) - F_2$. Here F_2 is the union of two tangent horodisks. The *center* of this interstice is the point in S^1, contained in the interstice, which is fixed by the hyperbolic reflection which interchanges the two horodisks of F_2. The *interstitial vertex* is the point of tangency of the horodisks of F_2. The *unbounded interstitial arc* is the geodesic ray connecting the interstitial vertex to the center.

Finite Horodisk Packings: A *finite horodisk packing* is a finite union of horodisks, which have pairwise disjoint interiors, such that the complementary regions are all interstices. We normalize so that $(0,0)$ is the center of one of the bounded interstices, and one of the horodisks bounding this interstice has its basepoint at $(1,0)$. (The *basepoint* is the point of tangency with S^1.) We call this interstice the *initial interstice*. Figure 2.1 shows an example on the left, together with 4 interstitial vertices and each type of interstitial arc.

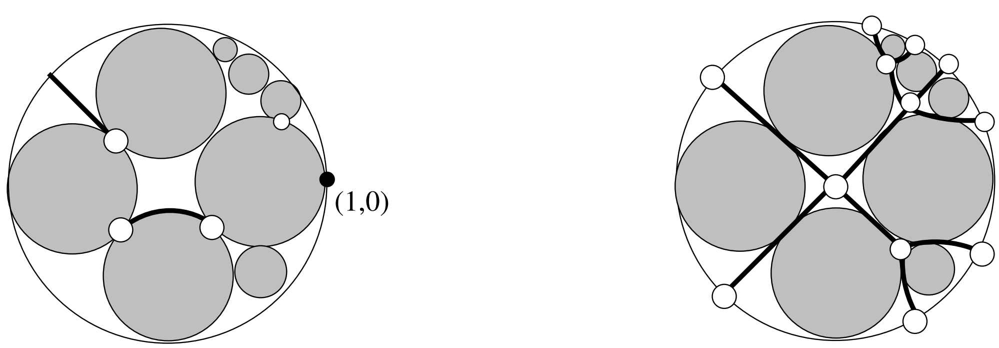

Figure 2.1

The Adjacency Tree: We say that two interstices are *adjacent* if they share a common interstitial vertex. To each horodisk packing H we may assign a finite tree T_H. The nodes of T_H are the centers of the interstices. Two nodes are joined by a geodesic segment or ray iff the corresponding interstices are adjacent. The initial node of T_H is defined to be the initial interstice. In this way, T_H is naturally a finite, embedded, planar, directed tree. Each non-terminal node of T has valence at least 3. We call T_H the *adjacency tree of* H. The right hand side of Figure 2.1 draws this tree.

The Black Order: If v is a node of T_H, other than the initial node, we let $S_v \subset S^1 - (1,0)$ denote the set of terminal nodes $\tilde{v}$ such that the directed path from the initial node to $\tilde{v}$ contains v. Sometimes S_v is called the "set of futures" of v. Suppose v and w are two such nodes, and that there is no directed path, starting from the initial node, which contains both v and w. We write $v \prec w$ iff there are nodes $\tilde{v} \in S_v$ and $\tilde{w} \in S_w$ such that one encounters $\tilde{v}$ before $\tilde{w}$ when travelling counterclockwise around S^1, starting at $(1,0)$. It is easy to see that this definition is independent of the choices of $\tilde{v}$ and $\tilde{w}$, and that either $v \prec w$ or $w \prec v$. Thus, the embedding of T_H into $\boldsymbol{H}^2$ determines a partial order on the nodes of T_H. We call this order the *black order* for reasons which will become clear shortly.

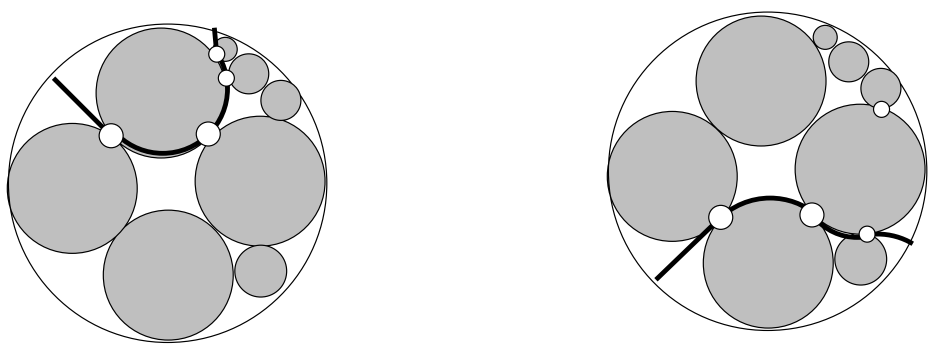

Figure 2.2

Special Curves: Let H be a horodisk packing. Let γ be a C^1 regular bi-infinite path in $\boldsymbol{H}^2$. Suppose also that $\gamma = \gamma_1 \cup \gamma_2 \cup \gamma_3$, where γ_1 and γ_3 are unbounded interstitial arcs, and γ_2 is a union of bounded interstitial arcs. We say that γ is *horo-like* if γ_2 is contained in a single horocircle. We say that γ is *slalom-like* if γ_2 makes an inflection point at each interstitial vertex in its interior. Figure 2.2 shows a horo-like curve and a slalom-like curve.

Horodisk Packing Graphs: We define a graph $G = G(H)$ as follows. The vertices of G are the centers of the unbounded interstices. These points lie in S^1 and are in bijection with the components of $S^1 - H$. Two vertices are joined by a *black edge* if there is a horo-like curve which joins them. Two vertices are joined by a *white edge* if there is a slalom-like curve which joins them. It is easy to see that each vertex is incident to two white edges and two black edges. We call $G(H)$ the *horodisk packing graph* associated to H.

Orienting the Black Edges: We orient the horo-like curves so that they travel clockwise around the horodisks. This gives an orientation to the black edges of $G(H)$. It is easy to see that the union of the black edges is a Hamiltonian circuit for $G(H)$. This circuit visits the vertices of $G(H)$ in the counterclockwise order that they appear on the circle. We define the *first vertex* of $G(H)$ to be the first vertex one encounters when travelling counterclockwise, starting from $(1,0)$. If v and w are vertices of $G(H)$ we write $v \prec w$ if the black Hamiltonian circuit visits v before w, starting from the first vertex. We call this the *black order* on the vertices of $G(H)$. To connect this definition with the one in the previous section: The vertices of $G(H)$ coincide with the terminal nodes of T_H. The black order here coincides with the black order defined above, when restricted to the terminal modes.

Remarks:
(i) Technically, $G(H)$ is a multi-graph, because more than one edge can connect two vertices. In general, many of the graphs we define have this property. We hope that the use of the term *graph* in place of *multi-graph* does not cause confusion.

(ii) If H is contained in a larger horodisk packing H' it is not usually true that $G(H)$ is a subgraph of $G(H')$. We will see in §2.4 that $G(H)$ is always a quotient graph of $G(H')$, however.

(iii) One can ask if there is a natural way to orient the white edges of $G(H)$. This is indeed the case. Furthermore, it turns out that the union of the white edges is also a Hamiltonian circuit. We will explain in the next section how the two kinds of edges play dual roles within the graph.

2.2 Duality

Suppose H is a horodisk packing, with associated horodisk packing graph $G(H)$. Let $G^*(H)$ be the graph obtained by recoloring all the white edges of $G(H)$ black and all the black edges white. Is there a horodisk packing H^* such that $G(H^*)$ is isomorphic to $G^*(H)$, via a color preserving isomorphism?

Let us reformulate the question. Given a horodisk packing H, let (H) denote the union of the interstices of H. In other words, (H) is just the closure of $\boldsymbol{H}^2 - H$. We equip (H) with the path metric induced from the inclusion into $\boldsymbol{H}^2$. Note that the horo-like curves and the slalom-like curves are naturally curves in (H). We say that two finite horodisk packings H_1 and H_2 are *dual* if there is an isometry $\phi : (H_1) \to (H_2)$, which carries horo-like curves to slalom-like curves, and *vice-versa*. We normalize so that the restriction of ϕ to the initial interstice of H_1 coincides with the hyperbolic (and, coincidentally, Euclidean) reflection $r_0(x, y) = (x, -y)$. Except for relatively trivial cases, ϕ would certainly not extend to a self-homeomorphism of $\boldsymbol{H}^2$. If H_1 and H_2 are dual, we write $H_2 = H_1^*$. An isometry from (H_1) to (H_1^*) maps vertices of $G(H_1)$ to vertices of $G(H_1^*)$. Moreover, it induces a bijection between the set of black (respectively white) edges of $G(H_1)$ and the set of white (respectively black) edges of $G(H_1^*)$. In short, such an isometry induces the desired graph isomorphism from $G^*(H_1)$ to $G(H_1^*)$.

Lemma 2.1 (Duality Lemma) *Every finite horodisk packing has a unique dual packing.*

First we prove uniqueness. Suppose that H_2 and H_2' are both dual to H_1. There is an isometry $\psi : (H_2) \to (H_2')$ which carries horocircles to horocircles and interstices to interstices. ψ has a canonical extension to all of $\boldsymbol{H}^2$: To define ψ on the interior of a horodisk of H_2, we choose the unique isometry which extends the action of ψ on the boundary. ψ clearly preserves lengths of paths and hence is an isometry. Since ψ is the identity on an open set, ψ is the identity everywhere.

We now construct a dual packing, H_2. Let A_j denote the union of bounded interstices of H_1 which correspond to nodes of the adjacency tree which are exactly j edges away from the initial node. A_0 is the initial interstice.

Let $\phi_0 : A_0 \to \boldsymbol{H}^2$ be the restriction to A_0 of r_0. Suppose that $\phi_n : A_n \to \boldsymbol{H}^2$ has been defined, and is an isometry on each interstice of A_n. Let b be an interstice of A_{n+1}. Let v be the unique vertex of b which is also a vertex of an interstice a of A_n. There is a unique hyperbolic reflection r which swaps the two horodisks of H_1 which are tangent to each other at v. We define

$$\phi_{n+1}|_b = (r \circ [\phi_n|_a])|_b.$$

Here $[\phi_n|_a]$ is the unique hyperbolic isometry which extends $\phi_n|_a$. We let ϕ be the union of all these maps. By symmetry, ϕ maps slalom-like curves of H_1 into curves of the form $\gamma = \gamma_1 \cup \gamma_2 \cup \gamma_3$, where γ_1 and γ_3 are geodesic rays and γ_2 is contained in a horocircle.

Let a be any bounded interstice of H_1. Let $\sigma_1 \cup ... \cup \sigma_n$ be the union of slalom-like curves which contain the interstitial arcs bounding a. Let h_j be the horodisk containing the horocircular part of $\phi(\sigma_j)$. If σ_j and σ_k intersect in an interstitial vertex of a then h_j and h_k are tangent. From this, and from symmetry, we see that $F(a) = \bigcup h_j$ is an n-flower. If a and b are adjacent bounded interstices then one easily sees that the interstices defined by $F(a)$ and $F(b)$ are likewise adjacent. It follows from this fact that the obvious big union $H_2 = \bigcup F(a)$ is a horodisk packing. By construction, ϕ is the isometry from (H_1) to (H_2) which realizes the duality between H_1 and H_2. ♠

Corollary 2.2 *$G(H)$ is the union of two Hamiltonian circuits, the one made from the white edges and the one made from the black edges.*

We orient the white edges in $G(H)$ using the orientation of the black edges of $G(H^*)$. Thus, $G(H)$ is a directed graph.

2.3 Drawing the Horodisk Packing Graphs

Necklaces: A *closed (respectively open) k-necklace* is a union of two simple closed (respectively open) curves, one black and one white, which string together k vertices in the same order. These graphs are better defined by example. Figure 2.3 shows the cases $k = 3, 4, 5$. A necklace is *horizontal* (respectively *vertical*) if its two curves are oriented in the opposite (respectively the same) directions. The reason for this terminology will emerge in §3.

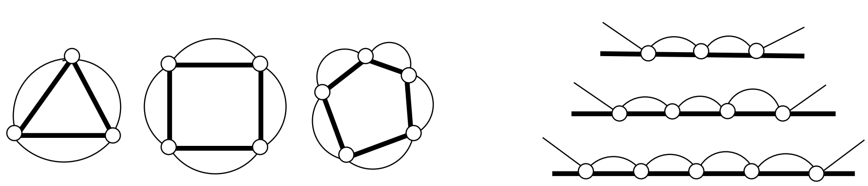

Figure 2.3

Motivating Discussion: If the packing H is a k-flower then so is the dual packing H^*. The horo-like curves and slalom-like curves of H coincide in this case and $G(H)$ must be a closed k-necklace. The isometry $(H) \to (H^*)$ is orientation reversing on the initial interstice of (H), the only one that counts in this case, so that the orientation of the slalom-like curves is opposite from the orientation of the horo-like curves. Hence, $G(H)$ is horizontal.

Suppose that $G(H)$ is given, and we want to construct $G(H')$, where H' is obtained from H by modifying H in the simplest way−that is, by adding a k-flower into one of the unbounded interstices of H. Call this unbounded interstice V. Let v be the vertex of $G(H)$ which corresponds to V. That is, V is the symmetry point of $V \cap S^1$. Outside of V, everything about the two packings agrees. Moreover, exactly two horo-like and two slalom-like curves of either packing enter into V. These 4 curves correspond to the 4 edges of $G(H)$ which are incident to v. All of this tells us that $G(H')$ is obtained from $G(H)$ by cutting out a small neighborhood of v and splicing in a new graph. The spliced-in graph has $k-1$ vertices, since the $k-2$ horodisks inserted into V, to make a k-flower, break up $V \cap S^1$ into $k-1$ smaller arcs.

To find the identity of the spliced-in "mystery graph", we turn the problem inside out. We can obtain H' from a single k-flower by adding all of H (except for two horodisks) to one of the unbounded interstices of this single flower. The same analysis as above shows that we obtain $G(H')$ from a closed k-necklace by cutting out a neighborhood of a single vertex and splicing in another graph. But this cut-open necklace is in exactly the same position as the mystery graph. Hence, the mystery graph is just an open $(k-1)$-necklace.

So, when we pass from $G(H)$ to $G(H')$ we splice in an open necklace at the appropriate vertex. There is only one way to do the splicing so as to make all the orientations match. The only thing we have not determined is: How do we decide if we splice in a horizontal necklace or a vertical necklace? Here is a heuristic idea: If H is a single flower then $G(H)$ is a closed necklace but $G(H')$ is not just a closed necklace. If we splice an open horizontal necklace into a closed horizontal necklace, we just get a longer closed necklace. Hence, we must splice a vertical necklace into the horizontal one. This special case suggests that the general pattern is one of alternation.

General Method: Let T_H be the adjacency tree of the horodisk packing H. Let T_H^k denote the set of nodes of T_H which are exactly k edges away from the initial node. Every two nodes in T_H^k are comparable in the black ordering. Thus, the black partial ordering on the nodes of T_H determines a linear ordering on the nodes of T_H^k.

Now we build $G(H)$. If the initial node of T_H has valence n, we let G_0 be a horizontal closed n-necklace. Once we choose a first vertex of G_0, the black oriented edges of G_0 determine a black order on the vertices of G_0. There is a unique bijection from the vertices of G_0 to the vertices of T_H^1 which respects the two black orders.

Suppose we have constructed a graph G_{k-1}, whose vertices are in a black-order-preserving bijection with the nodes of T_H^k. We create G_k as follows. If v is a vertex of G_{k-1}, we let n_v be the number of edges directed out of the node $\tilde{v}$ of T_H^k which corresponds to v. We cut out a small neighborhood of v and we splice in an open n_v-necklace. We make this necklace horizontal if k is even and vertical if k is odd. Doing these splices at all vertices of G_{k-1}, we create G_k. The first vertex of G_k is defined to be the first vertex of the spliced-in graph which replaces the first vertex of G_{k-1}. This choice determines a black order on the vertices of G_k. There is a unique bijection from the vertices of G_k to the vertices of T_H^{k+1} which respects the two black orders. Thus we construct a sequence of graphs $G_0, ..., G_n$ until we run out of nodes of T_H. The final graph is $G(H)$.

Why Alternation: Why do we alternate between horizontal and vertical necklaces? The reason is: If we have a duality $\phi : (H) \to (H^*)$ then ϕ is orientation preserving or reversing on an interstice, depending on the parity of the distance from the corresponding node to the initial node, in the adjacency tree.

3 Marked Rectangle Spaces

3.1 Basic Definitions

Marked Rectangles: By *rectangle* we always mean the solid 2-dimensional body. We always take the sides of our rectangles parallel to the coordinate axes in the plane. We say that a *marked rectangle* is a rectangle with one pair of opposite sides declared *plain* and one pair declared *dotted*. One of the diagonals of a marked rectangle has negative slope. We call this diagonal *black*. We call the other diagonal *white*. We color a vertex according to the color of the diagonal which contains it. We say that a vertex is *high* if it is contained in the top edge of the rectangle–that is, the horizontal edge whose y-coordinate is larger. The left hand side of Figure 3.1 shows a marked rectangle. We always draw in the black diagonal.

Figure 3.1

Subdivision: We say that a *run* of marked rectangles is a finite sequence $M_1, ..., M_k$ of marked rectangles, such that M_j and M_{j+1} are translation equivalent and intersect in a common dotted edge, for all j. The right hand side of Figure 3.1 shows two examples.

The run $\{M_j\}$ of rectangles canonically determines a single marked rectangle M. As a set, $M = \bigcup M_j$. The plain sides of M are defined to be the vertical (respectively horizontal) ones iff the plain sides of the M_j are the horizontal (respectively vertical) ones. The run shown at the extreme right of Figure 3.1 determines a marked rectangle which is translation equivalent to the one shown in the left side of Figure 3.1.

Inversely, if M is a marked rectangle we define a *subdivision* of M to be a run of $k \geq 2$ marked rectangles which determines M. Here is the secret behind this definition, which relates it to the discussion in §2.3: *Iterated subdivision of marked rectangles alternately produces horizontal and vertical runs.*

Marked Rectangle Patterns: Suppose that P and P' are finite unions of marked rectangles. We write $P \to P'$ if P' is obtained from P by subdividing exactly one marked rectangle of P. We call P' a *strict refinement* of

P. Let M_0 be the marked rectangle whose underlying set is the unit square and whose dotted sides are horizontal. We say that P is a *marked rectangle pattern* if there is a finite sequence: $M_0 \to P_1 \to ... \to P_k = P$.

Black and White Orderings: We inductively define an ordering on the marked rectangles within a pattern. The singleton $\{M_0\}$ obviously has only one order. Suppose that $M_1, ..., M_n$ are the marked rectangles of P listed in their order, and that $P \to P'$. Suppose that M_i is the marked rectangle which, in passing from P to P', is subdivided into the run $\{M_{i1}, ..., M_{ik}\}$. We list the M_{ij} from **left** to **right** (respectively top to bottom) if these marked rectangles run horizontally (respectively vertically). We order the marked rectangles of P' as follows:

$$M_1 < ... < M_{i-1} < M_{i1} < ... < M_{ik} < M_{i+1} < ... < M_n.$$

So, an ordering on P canonically determines an ordering on P'. By induction, then, we define an ordering of the marked rectangles within any marked rectangle pattern. We call this ordering the *black ordering*. We can define the *white ordering* simply by switching the words **left** and **right** in the definition. Figure 3.2 shows an example. The marked rectangles have been pulled apart to show their structure. The big centered numbers show the black ordering. The small corner numbers show the white ordering.

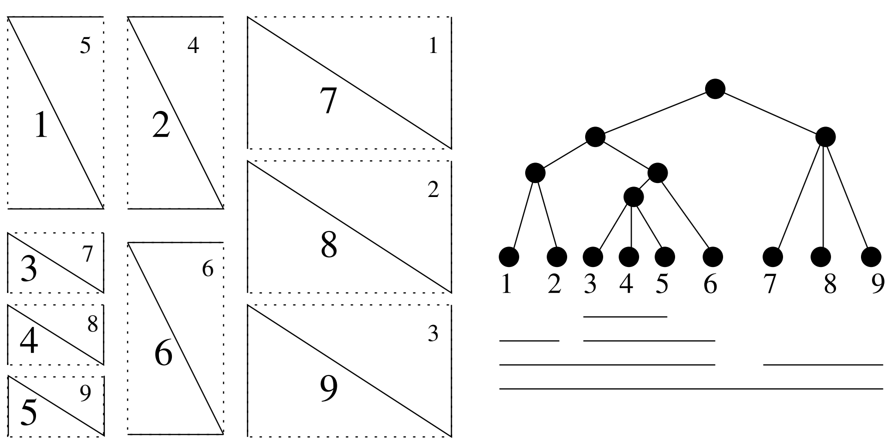

Figure 3.2

Indexing by Trees: Marked rectangle patterns are canonically indexed by finite directed trees. Except for the trivial tree, we insist that non-terminal nodes of our trees have at least two outgoing edges. We assume that these trees are embedded in the plane, with the non-terminal nodes in the upper half plane and the terminal nodes in $\boldsymbol{R}$. The ordering on $\boldsymbol{R}$ gives an ordering on the terminal nodes. Figure 3.2 shows an example. (The lines drawn beneath the tree in Figure 3.2 encode its structure and will be used below.)

Our method prefers the black ordering over the white one. We associate to M_0 the trivial tree. Suppose the rectangle pattern P is associated to the tree T. Suppose also that there is a map from the terminal nodes of T to the marked rectangles of P which respects the black order. Suppose that $P \to P'$ and that P' is obtained from P by subdividing the ith marked rectangle into a run of k. We let T' be the directed tree obtained by lifting the ith terminal node a bit off of $\boldsymbol{R}$ and connecting this node back to $\boldsymbol{R}$ with k new outgoing edges. These edges are then mapped, from left to right, into the M_{ij}, so as to respect the black order. We will let $P(T)$ denote the marked rectangle pattern associated to the tree T.

3.2 Duality Revisited: A Magic Trick

In this section we show a trick which is the secret behind the main result in the chapter, [§3.4, Graph Isomorphism Theorem].

Suppose T is a finite directed tree, drawn so that the outgoing edges come symmetrically downward out of each node, as in Figure 3.2. Given a node $v \in T$, let r_v be the reflection in the vertical line through v. Let T_v be the subtree whose initial node is v. We say that we *reverse* T *at* v if we delete T_v from T and replace it by $r_v(T_v)$. In so doing, we create a new tree, abstractly isomorphic to T, but embedded differently.

Here is a canonical re-embedding of a tree T. Start with $T_0 = T$. Obtain T_1 by reversing T_0 at the initial node. In general, obtain T_{j+1} by reversing T_j at all nodes which are j edges away from the initial node. Let T^* be the final tree obtained. To illustrate this, we perform this re-embedding on the tree shown in Figure 3.2. Rather than draw all the trees, we will just list how the vertices are permuted. In each list, the underline indicates which numbers are to be reversed to get the next list.

$$\underline{123456789} \to \underline{987}\ \underline{654321} \to 789\ \underline{12}\ \underline{3456} \to 789\ 21\ 6\underline{543} \to 789216345$$

We have met this construction before, disguised in hyperbolic clothing. Let T be the adjacency tree to a horodisk packing H, then the main construction of [§2.3, Duality Lemma], which constructs the dual packing H^*, exactly performs the succession of reversals just described. The only difference is that we used hyperbolic geometry rather than Euclidean geometry to effect the reversals. Hence, T^* is isomorphic to the adjacency tree of H^*.

Now for the magic trick: Pair up the labels of T^* with T, using the following suggestive notation: $7^1, 8^2, 9^3, 2^4, 1^5, 6^6, 3^7, 4^8, 5^9$. Now look at the marked rectangles in Figure 3.2 (or Figure 3.3 below). You will see these pairs exactly!

3.3 Marked Rectangle Spaces

Let P be a marked rectangle pattern, consisting of n marked rectangles. We let $\hat{P}$ be the topological space which is the disjoint union of n marked rectangles. We think of the marked rectangles in $\hat{P}$ as being canonically bijective with the marked rectangles of P. We form the space $|P|$ as follows: Identify the low black (respectively white) vertex of the ith rectangle in $\hat{P}$ to the high black (respectively white) vertex of the $(i+1)$st rectangle in $\hat{P}$. Indices are taken cyclically, mod n. The arrows in Figure 3.3 show which points are identified in $|P|$. We are using *both* the black and white orderings to define these identifications.

The union of all the black (respectively white) diagonals is homeomorphic to S^1, and visits every rectangle in $|P|$.

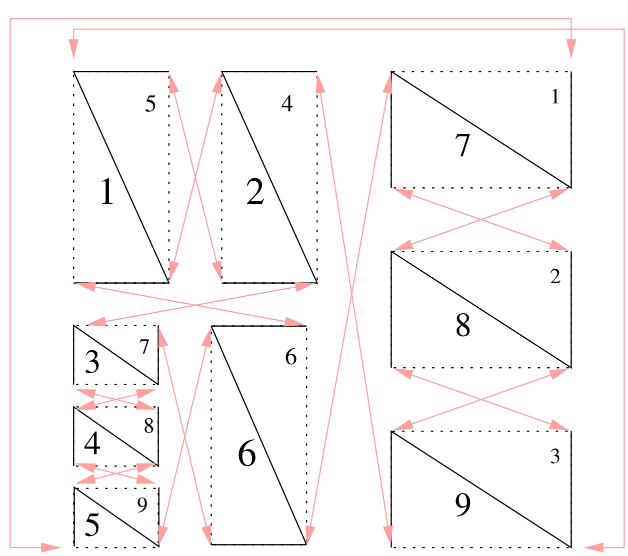

Figure 3.3

To see the black circle in Figure 3.3, start in the first marked rectangle in the black ordering. Trace your finger from high to low along the first black diagonal, then jump to the second high black vertex and repeat. And so forth. When you get to the last black low vertex, you jump back to the first black high vertex. The apparent discontinuities are not discontinuities at all, for the relevant points are identified.

3.4 The Graph Isomorphism Theorem

If H is a finite horodisk packing, H determines a marked rectangle space in the following way. Let τ_H be the adjacency tree of H, as defined in §2.1. Use a stereographic projection μ to identify $\boldsymbol{H}^2$ with the upper half plane, so that the $\mu(1,0)=\infty$. Let $T=\mu(\tau_H)$. The initial node on T is the image under μ of the initial node of τ_H. The interior points of T are all contained in the upper half plane and the terminal nodes are contained in $\boldsymbol{R}$. These are the conditions we required so that we could index a marked rectangle pattern by T. Let $P=P(T)$ be the marked rectangle pattern indexed by T and let $|P|$ be the associated marked rectangle space. We write $|P|=|P|(H)$ to denote that P just depends on the finite horodisk packing H.

There is a graph $G'(H)$ canonically associated to $|P|(H)$. This graph is simply the union of the black and white diagonals, with the centers of rectangles put in as vertices. The corners of the rectangles are not counted as vertices of $G'(H)$. We orient the edges so that they travel from high to low within a rectangle.

Let $G(H)$ be the horodisk packing graph associated to H. By construction, the terminal nodes of the tree T are in canonical bijection with the vertices of $G(H)$ and also with the vertices of $G'(H)$. Hence the vertices of $G(H)$ are in canonical bijection with the vertices of $G'(H)$.

Theorem 3.1 (Graph Isomorphism) *There is a color-preserving graph isomorphism from $G(H)$ to $G'(H)$ which extends the isomorphism of the vertex sets.*

Proof: We use the terminology from §2.3. First, suppose that H is a k-flower. Then the marked rectangle pattern $P(H)$ is a horizontal run of k marked rectangles. It is easy to see that $G'(H)$ must be a horizontal closed k-necklace. For the induction step, suppose that H_1 and H_2 are such that the adjacency tree T_2 is obtained from T_1 by extending a single terminal node by k new outgoing edges. Here $k \geq 2$. To create $P(H_2)$, a unique marked rectangle of $P(H_1)$ is subdivided into k new ones. From this it is easy to see that $G'(H_2)$ is obtained from $G'(H_1)$ by splicing in a horizontal or vertical open k-necklace at the relevant vertex. This necklace is horizontal or vertical depending on whether or not the marked rectangles of the subdivision run horizontally or vertically. Our result now follows from induction, and from the observation that iterated subdivision alternately produces horizontal and vertical runs of marked rectangles. ♠

4 Marked Tetrahron Spaces

4.1 Basic Definitions

Marked Tetrahedra: By *tetrahedron* we mean the convex hull of 4 general position points in $\boldsymbol{R}^3$. A tetrahedron has 6 edges made from three pairs of disjoint edges. We declare one pair *dotted*, one pair *plain* and one pair *colored*. One of the colored edges we call *black* and the other one we call *white*. Each vertex is contained in a unique colored edge. We define a vertex to be the same color as the colored edge that contains it. We call one of the black vertices *high* and one *low*. Likewise for the white vertices. A suitable linear map from $\boldsymbol{R}^3$ to $\boldsymbol{R}^2$ maps a marked tetrahedron onto a marked rectangle, respecting all the markings. Any two marked tetrahedra are eqivalent via a unique affine transformation which respects all the markings.

The Heart: We now define a nice PL embedding of a marked rectangle into a marked tetrahedron μ. We take indices mod 8. Let c be the center (of mass) of μ. Let v_1, v_3, v_5, v_7 be the vertices of μ, labelled so that an uncolored edge e_j connects v_{j-1} and v_{j+1} for all even j. Let w_j be the midpoint of the segment which joins c to the midpoint of e_j. The points w_j are contained in the interior of μ. Let $\underline{\mu}$ be the union of the 8 triangles, defined by the triples (c, x_i, x_{i+1}). Here x stands for either v or w. We call $\underline{\mu}$ the *heart* of μ. Figure 4.1 shows a planar projection. It is easy to see that $\underline{\mu}$ is an embedded topological disk, contained in the interior of μ, except for the vertices. It is also easy to see that there is a canonical piecewise affine map from any marked rectangle into $\underline{\mu}$.

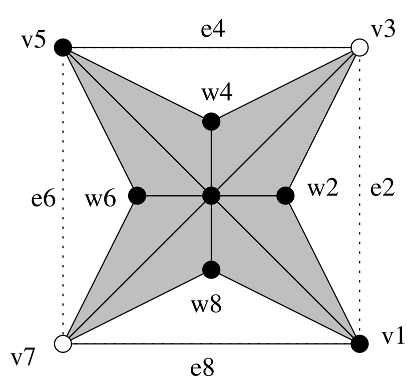

Figure 4.1

Subdivision, Informal Discussion: Soon we will define the *subdivision* of a marked tetrahedron. Our definition incorporates both the subdivision rule in §3.1 and the gluing rules given in §3.3. We will make the construction for a particular marked tetrahedron and then extend the definition to all marked tetrahedra by affine maps.

Before we make our precise construction we paint an informal picture of it. Imagine the usual picture of DNA, with two helical strands, one black and one white, coiling around each other and rising upwards. Picture the horizontal molecular ties connecting the two strands as dotted line segments. The convex hull of the set of two successive ties is a tetrahedron. The black (respectively white) edges of the successive tetrahedra are the edges connecting the black (respectively white) points of two successive ties. The ties form the dotted edges. The plain edges are the other edges. The union of these tetrahedra is roughly our model for a subdivision.

Subdivision of a Marked Tetrahedron: To take advantage of cylindrical coordinates we will temporarily identify $\boldsymbol{R}^3$ with $\boldsymbol{C} \times \boldsymbol{R}$. Let $z_0 = 1 + i$. Let μ be the tetrahedron whose vertices are

$$(z_0, 2)\ (-z_0, 2)\ (\overline{z}_0, -2)\ (-\overline{z}_0, -2).$$

We declare that the plain edges are contained in the horizontal planes $\boldsymbol{C} \times \{\pm 2\}$. (Our subdivision will switch plain and dotted edges, so as to match our informal picture.) The black edge connects $(\overline{z}_0, -2)$ to $(z_0, 2)$. The white edge connects $(-\overline{z}_0, -2)$ to $(-z_0, 2)$. The dotted edges are the remaining edges. Figure 4.2 shows the projection of μ to $\boldsymbol{C}$. The black edge is drawn boldly and the white edge is indicated by a thin white strip. We declare z_0 *high* and $-\overline{z_0}$ *high*. (This definition has nothing to do with their $\boldsymbol{R}$-coordinates in $\boldsymbol{C} \times \boldsymbol{R}$.)

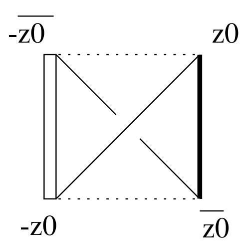

Figure 4.2

Let $n \geq 2$ be fixed. We sometimes omit n from our notation. Let T_n be the set of $n - 1$ points on $[-1, 1]$, which are evenly and maximally spaced.

We have

$$T_2 = \{0\}; \quad T_3 = \{-1, 1\}; \quad T_4 = \{-1, 0, 1\}; \quad T_5 = \{-1, -\frac{1}{3}, \frac{1}{3}, 1\}$$

and so forth. Let t_j be the jth element of T_n.

Let ω be the unit complex number, having smallest possible positive argument, such that $\omega^n \overline{z}_0 = z_0$.

Define the following sequences.

$$B' = \left\{ \left(\frac{1}{2}\, \omega^i \overline{z_0},\ t_i \right) \right\}; \qquad W' = \left\{ \left(\frac{-1}{2}\, \omega^i\, \overline{z_0},\ t_i \right) \right\}.$$

Here i runs from 1 to $n-1$. (The choice of $1/2$ is fairly arbitrary.) Let B be the sequence $\{(\overline{z}_0, -2), B', (z_0, 2)\}$. Note that B_n consists of $n+1$ points in $\boldsymbol{C} \times \boldsymbol{R}$. We think of these as black points. The first point coincides with one of the black vertices of μ. The last point coincides with the other black vertex of μ. We define W in an analogous way, and all the same statements are true, with *white* replacing *black*.

Let b_i be the ith point in B. Likewise define w_i. Let μ_i be the convex hull of the points $b_i, w_i, b_{i+1}, w_{i+1}$. We define the dotted edges to be the horizontal ones. We define the black edge to be the one bounded by two b-vertices. Likewise we define the white edge. The other two edges are plain. The vertices already have colors. For $k = 1, n$, we define a vertex of μ_k to be high if it coincides with a high vertex of μ. For $i = 1, ..., n-1$, we define a vertex of μ_{i+1} to be high if the corresponding vertex of μ_i is low.

Our construction is now done. We call $\{\mu_1, ..., \mu_n\}$ the n-th subdivision of μ and we extend this definition to all marked tetrahedra by affine maps. It is easy to see that $\mu_i \subset \mu$ for all i. Also $\mu_i \cap \mu_j = \emptyset$ if $|i - j| \geq 2$ and μ_i and μ_{i+1} share a common dotted edge. Finally a dotted edge of μ_k coincides with a plain edge of μ for $k = 1, n$.

Remark: One might worry that our definition of high and low runs into a problem with parity. We write $b_j \uparrow \mu_k$ to denote the sentence "b_j is high, considered as a vertex of μ_k." The symbol $b_j \downarrow \mu_k$ has the opposite meaning. Here is an analysis of two consecutive cases, $n = 2, 3$.

$$b_1 \downarrow \mu \quad b_1 \downarrow \mu_1 \quad b_2 \uparrow \mu_1 \quad b_2 \downarrow \mu_2 \quad b_3 \uparrow \mu_2 \quad b_3 \uparrow \mu.$$

$$b_1 \downarrow \mu \quad b_1 \downarrow \mu_1 \quad b_2 \uparrow \mu_1 \quad b_2 \downarrow \mu_2 \quad b_3 \uparrow \mu_2 \quad b_3 \downarrow \mu_3 \quad b_4 \uparrow \mu_3 \quad b_4 \uparrow \mu$$

No problem.

4.2 Simple Marked Tetrahedron Spaces

Suppose T and T' are finite unions of marked tetrahedra. We write $T \to T'$ if T' is obtained from T by subdividing a single marked tetrahedron. We call T' a *strict refinement* of T. Let μ_0 be the regular tetrahedron, with some marking chosen. We say that a *simple marked tetrahedron space* is a finite collection P of marked tetrahedra such that $\mu_0 \to P_1 \to ... \to P_n = P$. We call these objects SMTSs for short.

We would like to define the *black order* on the marked tetrahedra within an SMTS, but we are temporarily thwarted because we do not have good notions of left and right, or top and bottom, in space. However, looking carefully at the definition given in §3.2 we only had to have the notion of *left to right* and *top to bottom*. Our markings allow us to define these concepts. Suppose that e_1 and e_2 are two edges, either both plain or both dotted. We say that e_1 *precedes* e_2 if e_1 contains the high black vertex and e_2 contains the low black vertex. This notion is compatible with the planar definition, for the left and top edges of a marked rectangle, however they be marked, always contain the high black vertex.

Now that we have our notion of precedence, we can define the *black ordering* on the marked tetrahedra within a space inductively. For the inductive step, suppose μ is a marked tetrahedron which is subdivided into $\mu_1, ..., \mu_k$. We order the μ_j so that the first one touches the first plain edge of μ and the last one touches the last (that is, second) plain edge of μ. The rest of the construction is exactly the same as in §3.2.

Just as in §3.2, each SMTS is indexed by a directed tree, such that there is a canonical bijection between the terminal nodes of the tree and the marked tetrahedra within the SMTS. The construction is the same, and so we omit the details. If T is such a tree, we let $\widetilde{\Pi}(T)$ be the corresponding SMTS.

We now discuss the relationship between a marked rectangle space (MRS), defined in §3.3, and an SMTS. Every SMTS has a *heart*. Namely, one takes the union of all the hearts of the individial marked tetrahedra. One can see, by induction on the complexity of T, that there is a canonical bijection between the rectangles of the MRS, $|P|(T)$, and their PL copies within the heart of $\widetilde{\Pi}(T)$. One would like to say, simply, that this map is induced by a PL homeomorphism from $|P|(T)$ to the heart of $\widetilde{\Pi}(T)$. This is almost the case.

Note that an MRS has 2 special points. One of these points is the equivalence class consisting of the first high black vertex and last low black vertex.

The other special point is the equivalence class consisting of the first high white vertex and the last low white vertex. Let $|\widetilde{P}|(T)$ be the same space as $|P|(T)$, except that the points comprising these equivalence class are not identified. We call $|\widetilde{P}|(T)$ a *simple marked rectangle space* (SMRS).

Lemma 4.1 *There is a canonical PL homeomorphism between the SMRS and the heart of the corresponding SMTS. This homeomorphism respects the two black orders.*

Proof: Our definitions have been set up precisely for this result. The result follows from induction and simply from unravelling the definitions. ♠

4.3 Marked Tetrahedron Spaces

We are interested in MRSs and not SMRSs, so we need to improve the definition of an SMTS. Really, we just need a way to close up the outside points of the biggest marked tetrahedra in our configuration. Our solution is not that canonical, but it is perhaps the best that can be done in $\boldsymbol{R}^3$.

Say that a *marked tetrahedral k-ring* is a collection $\mu_1, ..., \mu_k$ of marked tetrahedra, having pairwise disjoint interiors, such that μ_j and μ_{j+1} share a common dotted edge. Here indices are taken mod k. We also insist that the relevant high points of μ_j are matched with the relevant low points of μ_{j+1}. Intuitively, one thinks of these tetrahedra as coming from a subdivision, except that the ends are wrapped around and brought together.

Lemma 4.2 *For every $k \geq 3$ there exists a marked tetrahedral k-ring.*

Proof: Let $v_1, ..., v_k$ be the vertices of a regular planar k-gon. Consider the line segments $s_j = v_j \times [0, 1] \subset \boldsymbol{R}^3$. Perturb these line segments slightly so that no two are parallel. Let T_j be the convex hull of s_j and s_{j+1}, suitably marked. Indices are taken mod k. If the perturbation is small the union of the T_j has all the required properties. ♠

We always distinguish, within a marked tetrahedral ring, a marked tetrahedron which we call *leftmost*. Say that a *marked tetrahedron space* (MTS) is a finite collection P such that, in the sense of the previous section, $R \to P_1 \to ... \to P_n = P$. Here R is a marked tetrahedral ring. Every MTS has a *black ordering*. The choice of the leftmost marked tetrahedron within R allows us to define the black order for tetrahedral rings. After this, the induction step is the same as for an SMTS. Each MTS is indexed by a finite directed tree. The only change is that the initial node of this tree must have valence at least 3. We let $\Pi(T)$ be the MTS associated to T.

The spaces $\widetilde{\Pi}(T)$ and $\Pi(T)$ differ only in that two extra dotted edges of $\Pi(T)$ is glued together. There is the same difference between $|\widetilde{P}|(T)$ and $|P|(T)$. Hence

Lemma 4.3 *There is a canonical PL homeomorphism between the MRS and the heart of the corresponding MTS. This homeomorphism respects the two black orders.*

Let H be a finite horodisk packing with adjacency tree T_H We write $\Pi(H) = \Pi(T_H)$. Let $G(H)$ be the horodisk packing graph. Recall from the introduction that $G(\Pi)$ is the graph whose red vertices are centers of tetrahedra in Π and whose blue vertices are vertices of tetrahedra in Π. A red vertex is connected to a blue vertex iff the corresponding center and the corresponding vertex belong to the same tetrahedron. Thus, $G(\Pi)$ is a graph embedded into $\boldsymbol{R}^3$, whose edges are straight line segments. The following is immediate from the Graph Isomorphism Theorem and from Lemma 4.3.

Lemma 4.4 (Realization Lemma) *If one places a new vertex at the center of each edge of $G(H)$ then there is a canonical graph isomorphism from $G(H)$ to $G(\Pi)$. The original vertices of $G(H)$ are mapped to the red vertices of $G(\Pi)$ and the added vertices in $G(H)$ are mapped to the blue vertices of $G(\Pi)$.*

5 The Main Result

5.1 Basic Definitions

All Definitions Extended: Recall from the introduction that an *infinite horodisk packing* H is an infinite collection of horodisks, having pairwise disjoint interiors, such that every complementary region is a bounded interstice. We normalize as in the finite case. We define *interstitial arcs*, *interstitial vertices* and *slalom curves* as in the finite case. The only difference is that there are no unbounded interstices. We define the *adjacency tree* just as in the finite case. Here this is an infinite directed tree with no terminal nodes.

Notation: Now we set up some notation which we will use throughout the chapter.

• Let $T = T_H$ be the adjacency tree of H. Let $T^n \subset T$ denote the set of nodes which are at most n away from the initial node.

• Let H_n denote the union of flowers corresponding to nodes of T^n. Note that H_n is a finite horodisk packing for all n, and $H_n \subset H_{n+1}$. Obviously, $H = \bigcup H_n$.

• Let U_n denote the set of unbounded interstices of H_n. Each interstice $u \in U_n$ defines the closed arc $u \cap S^1$. Let U_n^∞ be the union of these arcs. Note that U_n^∞ is a partition of S^1 and U_{n+1}^∞ refines U_n^∞ in the ordinary sense that partitions refine each other.

• Let $G_n = G(H_n)$ be the horodisk packing graph associated to H_n. The arcs of U_n^∞ are canonically bijective with the vertices of G_n.

• Let $\Pi_n = \Pi(H_n)$ be the marked tetrahedron space associated to H_n. Let $\phi_n : G_n \to \Pi_n$ be the embedding from the Realization Lemma. Using ϕ_n we see that the arcs of U_n^∞ are canonically bijective with the tetrahedra of Π_n. This bijection is *coherent*: Nested intervals correspond to nested marked tetrahedra.

5.2 Endpoints of Slalom Curves

Let $|U_n|$ be the maximum Euclidean diameter of a region in U_n.

Lemma 5.1 $\lim_{n\to\infty} |U_n| = 0$.

Proof: Every region $u \in U_n$ is determined by two tangent horodisks A_u and B_u. There are only finite many horodisks in H which have diameter greater than ϵ. If we choose n sufficiently large then at least one of A_u or B_u will have diameter less than ϵ. In either case, it is not hard to see that u has diameter at most $C\sqrt{\epsilon}$ for some universal constant C. ♠

Corollary 5.2 *Every slalom curve of H has two distinct accumulation points in S^1. If two slalom curves of H share an endpoint then they coincide. Finally, the endpoint of a slalom curve is never the basepoint of a horodisk.*

Proof: Let γ be a slalom curve of H. First, since we can always renormalize H by an isometry to pick a new initial interstice, we can assume that γ contains an interstitial arc of the initial interstice of H. Let γ_n denote the portion of γ which is contained in the union of bounded interstices of H_n. The set $\gamma - \gamma_n$ consists of two infinite rays. One of these rays is contained in some unbounded interstice a_n of U_n and the other is contained in some b_n of U_n. Since $|U_n| \to 0$, we see that $\cap a_n$ is a single point $a_\infty \in S^1$. Likewise $\cap b_n = b_\infty \subset S^1$. It is easy to see that $a_1 \neq b_1$ and that $a_2 \cap b_2 = \emptyset$. Hence $a_\infty \neq b_\infty$. These are obviously the two accumulation points of γ.

For the second part of the lemma, let γ and γ' be two such slalom curves which share an endpoint a. Let $\{a_n\}$ and $\{a'_n\}$ be the corresponding sequences of unbounded interstices, such that $a = \cap a_n = \cap a'_n$ is the common endpoint. It is easy to see that $a_{n+2} \cap S^1$ is contained in the interior of $a_n \cap S^1$. In particular, a_∞ is contained in the interior of a_n. Since U_n^∞ is a partition of S^1, we must have $a_n = a'_n$ for large n. Since this is true, in particular, for two consecutive choices of n, we see that γ and γ' share an interstitial arc. But this implies that $\gamma = \gamma'$.

For the third part, suppose that a is the basepoint of a horodisk h of H. Since a is contained in the interior of $a_n \cap S^1$ for large n, we see that $h \subset a_n$. This contradicts Lemma 5.1. ♠

5.3 The Interstitial Space

Let (H) is the union of interstices of H. We equip H with the path metric, induced with the hyperbolic metric on $\boldsymbol{H}^2$. We define an *end* of (H) to be a nested sequence of regions

$$u_0 \supset u_1 \supset u_2 \supset u_3 ...; \qquad u_n \in U_n.$$

We put a metric (and hence a topology) on the set of ends by saying that the distance between two ends is 2^{-n} if they agree exactly up to the first n terms. Let H_∞ denote the space of ends, equipped with this metric. It is easy to see that H_∞ is a Cantor set.

Every horocircle determines two distinct ends. One simply lists out the unbounded interstices entered by the horocircle. For cosmetic purposes, we pad out the beginning of the sequences so that they start with an element of U_0. Every slalom curve also determines two distinct ends, in the same way. There are two natural equivalence relations on H_∞. Given $x, y \in H_\infty$ we write $x \sim_1 y$ iff x and y are the two ends of a horocircle. We write $x \sim_2 y$ iff x and y are the two ends of a slalom curve. Let $\sim$ be the union of the two relations $\sim_1$ and $\sim_2$. That is, $x \sim y$ if and only if $x \sim_j y$ for some $j = 1, 2$.

Lemma 5.3 *$Q(H)$ is canonically homeomorphic to $H_\infty / \sim$.*

Proof: First, we claim that $H_\infty / \sim_1$ is canonically homeomorphic to S^1. This is almost a tautology. The inclusion map $(H) \to \boldsymbol{H}^2$ induces a continuous surjection $\pi : H_\infty \to S^1$. It is easy to see that π identifies two ends of H_∞ iff they are equivalent under $\sim_1$. Hence π is a continuous bijection. Note that H_∞ is compact, and hence so is $H_\infty / \sim_1$. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism. Hence, π is a homeomorphism. In sum: $Q(H) = S^1 / \sim_2 = (H_\infty / \sim_1) / \sim_2 = H_\infty / \sim$. This completes the proof. ♠

Lemma 5.4 *Let $\rho = \{r_n\}$ and $\sigma = \{s_n\}$ be two ends of H_∞.*

1. *If ρ and σ are the two ends of a horocircle (respectively slalom curve) then for sufficiently large n the centers of r_n and s_n are joined by horo-like (respectively slalom-like) curves.*

2. *If ρ and σ are inequivalent in H_∞ then there is some n such that the centers of r_n and s_n are joined neither by a horo-like curve of H_n nor a slalom-like curve of H_n.*

Proof: For large n the sequences $\{r_n\}$ and $\{s_n\}$ simply list out which regions the horocircle (or slalom curve) enters. The first half of our lemma is obvious from this. For the second half, suppose there is a a horo-like curve which joins the relevant centers of r_n and s_n, for infinitely many n. The limit of these horo-like curves converges, on a subsequence, to a horocircle of H. The convergence may be taken in the Hausdorff topology on closed subsets of $\boldsymbol{H}^2$. The key point to this convergence is that the points on S^1, corresponding to our ends, are distinct. Hence there is a lower bound to the set of diameters of our horo-like curves. A very similar argument works for slalom-like curves. ♠

5.4 Duality Again

This section is not needed elsewhere in the paper. We say that two packings H and H' are *dual* if there is an isometry from (H) to (H') which interchanges horocircles of H with slalom curves of H' and *vice versa*. We normalize this isometry as in §2.3.

Lemma 5.5 (Infinite Duality Lemma) *H has a unique dual packing H^*. The horodisk quotients $Q(H)$ and $Q(H^*)$ are homeomorphic.*

Proof: The uniqueness proof is exactly the same as for the Duality Lemma of §2.3. For existence, let $\{H_n\}$ be the sequence of finite horodisk packings approximating H and let $\{H_n^*\}$ be the sequence of dual packings. Let $\phi_n : (H_n) \to (H_n^*)$ be the sequence of isometries. It is easy to see, from the proof of the Duality Lemma, that these maps are all compatible, and that the limit map exists. Moreover, $H_n^* \subset H_{n+1}^*$ for all n. It follows that $H^* = \bigcup H_n^*$ is dual to H and that limit map $\lim \phi_n$ implements the duality.

Now consider $Q(H)$ and $Q(H^*)$. Let H_∞ be the set of ends of (H) and let H_∞^* be the set of ends of (H^*). The isometry $i : (H) \to (H^*)$ induces a homeomorphism from H_∞ to H_∞^*. By construction, $x \sim_j y$ in H_∞ if and only if $i(x) \sim_{(3-j)} i(y)$ in H_∞^*. Combining this information with Lemma 5.3 we see that i descends to a continuous map from $Q(H)$ to $Q(H^*)$. The whole process is invertible, so that i induces a homeomorphism. ♠

5.5 Geometry of Iterated Subdivision

In this section we prove several technical results about subdivision of tetrahedra. Let $|S|$ denote the Euclidean diameter of a set $S \subset \boldsymbol{R}^3$. We begin with two preliminary results.

Let μ be the marked tetrahedron used in §4.2 to define the subdivision. Say that a tetrahedron $\widetilde{\mu}$ is *internal* to μ if $\widetilde{\mu} \subset \mu$ and $\widetilde{\mu} \cap \partial\mu$ is either empty, or a vertex common to both tetrahedra.

Lemma 5.6 *Suppose $\widetilde{\mu}$ is a tetrahedron which is internal to μ. There is a universal constant $\alpha \in (0,1)$ which has the following property: Let L be any line which contains two distinct points of $\widetilde{\mu}$. Then $|L \cap \widetilde{\mu}| < \alpha |L \cap \mu|$.*

Proof: The result follows from compactness if $\widetilde{\mu}$ is contained in the interior of μ. So, consider the case when they share a vertex v. We translate so that $v = 0$. Suppose $\{L_n\}$ is a sequence of lines such that $|L_n \cap \widetilde{\mu}|/|L_n \cap \mu|$ converges to 1. By compactness, this can only happen if $L_n \cap \mu$, as a set, shrinks to 0. Let D_n be the dilation such that $|D_n(L_n \cap \mu)| = 1$. The scale factors for the D_n increase unboundedly. From this it follows that $D_n(\mu)$ converges, in the Hausdorff topology on closed subsets, to a strictly convex cone C. Indeed, μ and C coincide in a neighborhood of 0. Likewise $D_n(\widetilde{\mu})$ converges to a strictly convex cone $\widetilde{C}$. Note that $\widetilde{C} \subset C$, and $\partial\widetilde{C} \cap \partial C = \{0\}$. Let $M_n = D_n(L_n \cap \mu)$. The segments M_n all have unit length, and have both endpoints on ∂C. It is easy to see that a subsequence converges to a limit segment M. Since dilations preserve ratios of lengths, we have $|M \cap \widetilde{C}| = |M \cap C| = 1$. But then $\partial C \cap \partial\widetilde{C}$ contains the endpoints of M. This is a contradiction. ♠

Lemma 5.7 *Let μ be the marked tetrahedron used in the definition of subdivision. There exists a finite list of tetrahedra $\widetilde{\mu}_0, ..., \widetilde{\mu}_k$, internal to μ, having the following property: If μ' is any tetrahedron in any subdivision of μ and μ'' is any tetrahedron in any subdivision of μ' then $\mu'' \subset \widetilde{\mu}_j$ for some j.*

Proof: By construction, there is a single compact subset K, contained in the interior of μ, such that $\mu' \subset K$ unless μ' is either the first or last tetrahedron in the relevant subdivision. This K does not depend on the number of tetrahedra in the subdivision. We can fit K inside a tetrahedron $\widetilde{\mu}_0$ which is contained in the interior of μ. Now suppose μ' is either the first or the last tetrahedron in some subdivision of μ. If μ is subdivided into $k \geq 3$ tetrahedra then the first and last tetrahedra in the subdivision are the same independent of k. For this reason, there are only 4 different choices for μ'. Running the same argument as above, we see that there are tetrahedra $\widetilde{\mu}_1, ..., \widetilde{\mu}_4$, internal to μ, such that $\mu'' \subset \widetilde{\mu}_j$, for some j, unless μ'' is the first or last tetrahedron in the subdivision of μ'. In this final case, it is easy to see that μ'' itself is internal to μ. Again there are only 4 choices for μ'' given μ'. Hence there are only finitely many choices of μ''. By adding these choices to our list of internal tetrahedra we complete the proof. ♠

Corollary 5.8 (Shrinking Lemma) *There is some $\alpha < 1$ which has the following property: Let μ_0 be a marked tetrahedron. Let μ_1 be a marked tetrahedron in a subdivision of μ_0. Let μ_2 be a marked tetrahedron in a subdivision of μ_1. Then $|\mu_2| < \alpha|\mu_0|$. Here α does not depend on these tetrahedra.*

Proof: Let T be the affine map which maps μ_0 to μ, our model marked tetrahedron, in the way which preserves the markings. Let $\mu' = T(\mu_1)$ and $\mu'' = T(\mu_2)$. Let $\widetilde{\mu}_j$ be the tetrahedron, internal to μ, which contains μ''. Here we are using the notation from the previous lemma. Let $a, b \in \mu_2$ be points which realize $|\mu_2|$. Let λ be the line which contains a and b. Let $L = T(\lambda)$. By convexity, $\mu_0 \cap L$ is a line segment bounded by two points a' and b'. We have

$$\frac{|\mu_2|}{|\mu_0|} \leq \frac{|a-b|}{|a'-b'|} = \frac{|L \cap \widetilde{\mu}|}{|L \cap \mu|} < \alpha_j < 1.$$

The equality comes from the fact that an affine map is a similarity when restricted to a line. The last inequality comes from Lemma 5.6. Taking $\alpha = \max(\alpha_0, ..., \alpha_k)$ finishes the proof. ♠

5.6 Putting it Together

Referring to the notation of §5.1, note that Π_{n+1} is obtained from Π_n by subdividing each marked tetrahedron in Π_n. Let $\Pi = \bigcap \Pi_n$. We will construct a homeomorphism $\Psi : Q(H) \to \Pi$. In outline:

1. We construct a canonical map $\widetilde{\Psi}: H_\infty \to \Pi$.
2. We show that $\widetilde{\Psi}$ is surjective and continuous.
3. We show that $\widetilde{\Psi}$ respects the equivalence relation $\sim$ and never identifies points which are inequivalent under $\sim$.

These results show that $\widetilde{\Psi}$ induces a quotient map $\Psi: \Pi \to Q(H)$ which is a continuous bijection. Since the domain is compact and the range is Hausdorff, Ψ is a homeomorphism.

Definition of the Map: Suppose $\rho = \{r_i\}$ is an end of H_∞. Let v_n be the vertex of G_n which is the center of the unbounded interstice r_n. Let μ_n be the marked tetrahedron of Π_n which contains $\phi_n(v_n)$. By construction $\{\mu_n\}$ is a nested sequence of tetrahedra. Compare the discussion at the end of §5.1. From the Shrinking Lemma, $\bigcap \mu_n$ is a single point. We define $\widetilde{\Psi}(\rho) = \bigcap \mu_n$.

Surjectivity and Continuity: The maps ϕ_n are surjective onto the centers of tetrahedra of Π_n, and these tetrahedra shrink to points as $n \to \infty$, by the Shrinking Lemma. This proves that $\widetilde{\Psi}$ is surjective.

If ρ_1 and ρ_2 are two ends which are less than 2^{-n} apart then $\widetilde{\Psi}(\rho_1)$ and $\widetilde{\Psi}(\rho_2)$ are contained in the same marked tetrahedron of Π_n. The diameter of this tetrahedron is exponentially small, by the Shrinking Lemma. Hence, $\widetilde{\Psi}$ is a continuous map.

Interaction with the Equivalence Relation: Now let us show that $\widetilde{\Psi}$ respects the equivalence relation $\sim_1$. Suppose $\rho = \{r_n\}$ and $\sigma = \{s_n\}$ are the two ends of a horocircle. By Lemma 5.4, a horo-like curve of H_n connects the center of r_n to s_n for n sufficiently large. That is, the corresponding vertices of G_n are joined by a black edge. Hence $\widetilde{\Psi}(\rho)$ and $\widetilde{\Psi}(\sigma)$ are contained in tetrahedra of Π_n which share a vertex. Since this is true for all large n, and since these tetrahedra shrink to points, we must have $\widetilde{\Psi}(\rho) = \widetilde{\Psi}(\sigma)$. This limit must be a black vertex of infinitely many tetrahedra. The same argument works for $\sim_2$, with *white* replacing *black*.

Finally, let us show that $\widetilde{\Psi}$ doesn't identify any points which are inequivalent under $\sim$. If the ends ρ and σ are not equivalent under $\sim$ then by Lemma 5.4 there is some n such that r_n and s_n are unbounded interstices whose centers are joined neither by a horo-like curve nor a slalom-like curve. By the Realization Lemma, two tetrahedra in Π_n share a vertex if and only if they correspond to adjacent vertices in G_n. Therefore, the map ϕ_n maps the vertices corresponding to r_n and s_n into marked tetrahedra which are disjoint from each other. It now follows from our definition of $\widetilde{\Psi}$ that $\widetilde{\Psi}(\rho) \neq \widetilde{\Psi}(\sigma)$.

Remark: We have also shown that Π contains the union of all the vertices of all the Π_n. Moreover, we have shown that the composition map $\Psi \circ \pi$ maps basepoints of horodisks onto the black vertices, and endpoints of slalom curves onto the white vertices. If we had made all the constructions in this chapter with respect to the dual packing, we could construct all the same objects, except that all the black vertices would be colored white and all the white vertices would be colored black.

Symmetries: Having completed our outline, we know that $\Psi : Q(H) \to \Pi$ is a homeomorphism. To finish the proof of the Main Result, we consider symmetries of the infinite packing H.

Suppose $\gamma : \boldsymbol{H}^2 \to \boldsymbol{H}^2$ is a hyperbolic symmetry of H. We give the same name to the induced self-homeomorphism of $Q(H)$. The map $\Psi \circ \gamma \circ \Psi^{-1}$ is a self-homeomorphism of Π. Call this map ψ.

We say that a *partition interval* is an interval in $\bigcup U_n^\infty$. These are the closures of the connected components of $S^1 - H_n$, for all the n. Each partition interval a corresponds canonically with a marked tetrahedron μ_a, as discussed in the last point of §5.1. Let ψ_a be the restriction of ψ to $\Pi \cap \mu_a$. Consider the action of γ on U_0^∞, the first partition of S^1 into intervals. For each interval $a \in U_0^\infty$ there is a finite collection of partition intervals $a_1, ..., a_k$ such that $a = a_1 \cup ... \cup a_k$ and $b_k = \gamma(a_k)$ is a partition interval. This follows from the fact that γ preserves H.

From the affinely natural way we make our subdivision construction, we see that ψ_{a_j} is the restriction of the unique marking preserving affine map from μ_{a_k} to μ_{b_k}. (The orientation preserving nature of γ makes sure that the notions of "high vertex" and "low vertex" are not reversed.) We perform the same analysis for each of the finitely many intervals of U_0^∞.

The analysis above shows that there is a finite union X of tetrahedra such that $\Pi \subset X$ and such that ψ extends to an affine map on each tetrahedron within X. Now X is not quite an open neighborhood of Π. However, given that every two tetrahedra in X are either disjoint or share at most one edge, it is easy to thicken the extension up a bit, in the complement of X, to get a PL extension in a neighborhood of Π.

6 Graphs

6.1 Embedding Paths

Let $\Pi = \bigcap \Pi_n$ be as in §5. A Π-*tetrahedron* is a marked tetrahedron which belongs to Π_n for some n. In the proof of the Main Result, we saw that Π contains every vertex of every Π-tetrahedron. If μ is a Π-tetrahedron, let $\Pi_\mu = \Pi \cap \mu$. Let $V(\mu)$ be the vertex set of μ. It follows from Lemma 5.7 that $\Pi_\mu \cap \partial\mu = V(\mu)$.

Let $\phi : S^1 \to \Pi$ be the composition of the projection $S^1 \to Q(H)$ with the homeomorphism $\Psi : Q(H) \to \Pi$, given in the Main Result.

Lemma 6.1 *Suppose x and y are the endpoints of an arc $I \subset S^1$, and $\phi(x) \neq \phi(y)$. Then, there is an embedded path $\gamma \subset \phi(I)$ which joins x to y.*

Proof: This result is a special case of the well-known result that a compact Hausdorff space X is path connected if there is a continuous surjection from an interval to X. To keep this paper self-contained, we prove the special case at hand. For ease of exposition we assume that neither x nor y is the endpoint of a slalom curve. We also assume that I is contained in a semicircle.

We order the slalom curves, which have both endpoints in I, according to the distance between their endpoints. If several coincide by this measure, we order these arbitrarily. Let $I_0 = I$. In general, if I_n is a finite union of closed arcs, we create I_{n+1} by deleting the open subarc bounded by the first slalom curve which has both endpoints in one of the arcs of I_n. Since, by Lemma 5.2, two slalom curves never share an endpoint, the two endpoints in question are contained in the interior of the relevant arc of I_n. Hence, I_{n+1} is again a finite union of closed arcs. Lemma 5.4 implies that this process goes on forever. $C = \bigcap I_n$ is totally disconnected, and hence a Cantor set.

Our ordering has been chosen so that no two points in C are identified, except for each of the two points bounding an arc of $S^1 - C$. For instance, if two points in different intervals of I_1 are identified, then the slalom curve which makes the identification would have been listed and used before the one we used to define I_1 in the first place. The same argument works, in an inductive way, when passing from I_n to I_{n+1}. The identifications that are made on C merely "close up the gaps". Thus, $\phi(C)$ is an embedded path joining $\phi(x)$ to $\phi(y)$. ♠

Lemma 6.2 *Any two vertices of μ can be joined by a path in Π_μ which avoids the other vertices of μ.*

Proof: Suppose that v and w are both white vertices of μ. Let I' be the interval of U_n^∞ which corresponds to μ via our constructions in §5. By the construction of Ψ, the map from the Main Result, $\phi(I') = \Pi_\mu$ and ϕ maps the endpoints of I' to the black vertices of μ. Now, ϕ is injective on the basepoints of horodisks by Lemma 5.2. Therefore, there is an arc I, contained in the *interior* of I', whose endpoints are mapped to v and w. Applying the previous result to I gives us a path connecting v to w. Since ϕ is injective on basepoints of horodisks, and since both black vertices have pre-images not in I, we see that our path avoids these black vertices.

Now suppose v and w are joined by a plain edge. Let μ' be the marked tetrahedron in the subdivision of μ which shares this plain edge. Note that μ' is disjoint from $V(\mu) - \{v, w\}$. There is some arc J' such that $\phi(J') = \Pi_{\mu'}$. There is some sub-arc $J \subset J'$ such that $\phi(J) \subset \mu'$ and ϕ maps the endpoints of J to v and w. Applying the previous result to J gives us a path, embedded in μ', which joins v to w. Though this path may contain other vertices of μ', it does not contain any other vertices of μ.

The other cases are not used below. In brief, the "black edge case" follows from the "white edge case" and from the Duality Principle; the "dotted edge case", like the "plain edge case", is treated by passing to a subdivision. We omit the details. ♠

6.2 Embedding Graphs

In this section we prove Corollary 1.3.

Say that the *I-graph* is just the letter I. It has 6 vertices. Likewise we define the *X-graph*. Note that the X-graph is a quotient graph of the I-graph. We first prove that we can embed either an X-graph or an I-graph in Π_μ so that the valence-1 vertices map bijectively to $V(\mu)$.

Figure 6.3 shows our construction. For ease of exposition we will assume that, in passing from Π_n to Π_{n+1}, the tetrahedron μ is subdivided into 3 tetrahedra, as shown schematically in Figure 6.3. Some of the white vertices have been marked. The arrows indicate that points are actually identified. Let μ_1', μ_2', μ_3' be the three marked tetrahedra in the subdivison of μ, represented from left to right in Figure 6.3. If η is some given marked tetrahedron and a and b are vertices of η we use the notation $(a \to b|\ \eta)$ to stand for the sentence "connect a to b in Π_η with a path which avoids the $V(\eta) - \{a, b\}$."

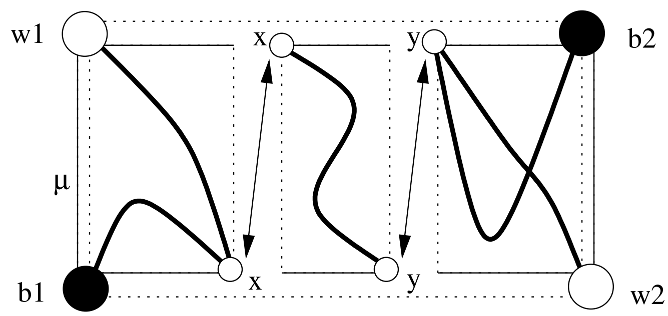

Figure 6.3

The following 5 constructions are made possible by Lemma 6.2:

$$(x \to w_1|\ \mu_1') \quad (x \to b_1|\ \mu_1') \quad (x \to y|\ \mu_2') \quad (y \to w_2|\ \mu_3') \quad (y \to b_2|\ \mu_3')$$

After trimming away some of our set, we have an embedded I-graph. If μ is divided into $k > 3$ tetrahedra, then the middle step above is expanded into $k - 2$ similar steps. If μ is divided into 2 tetrahedra then the middle stap is eliminated. In this case, we could produce an X-graph rather than an I-graph.

We start with the incidence graph $G(\Pi_n)$. Within each tetrahedron μ of Π_n, the graph $G(\Pi_n)$ is just an X-graph. We replace this X-graph with the graph we have embedded into Π_μ. Doing this replacement, we arrive at the $\widetilde{G}_n$ mentioned in Corollary 1.3.

6.3 Non-Planar Incidence Graphs

We begin with an example. Figure 6.4 shows the marked rectangle space $|P|(H_1)$ when H is the prototypical horodisk packing. The arrows indicate identifications. The graph made from thick lines is a subgraph of $G'(H_1)$. We have colored the vertices so as to reveal that this subgraph is exactly $K_{3,3}$ (the complete bipartite graph) and hence non-planar. So, $G'(H_1)$ is non-planar. Now, $G(\Pi_1)$ is obtained from $G'(H_1)$ by inserting a vertex at the center of each edge of $G'(H_1)$. Hence, $G(\Pi_1)$ is non-planar.

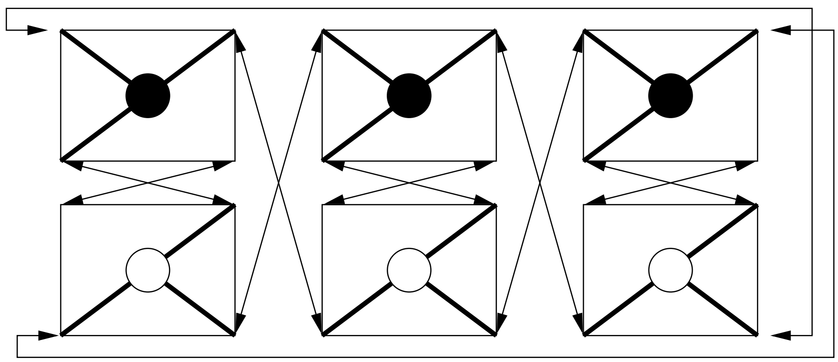

Figure 6.4

In general, suppose that H has some odd flower. Let H' be the horodisk packing, isometric to H, which has this odd flower as its initial interstice. It is not hard to see that $G(\Pi_n)$ has $G(\Pi_1')$ as a quotient graph for n sufficiently large. This works as long as n is at least two more than the distance from the odd flower of H to the initial node of H, in terms of the adjacency tree. Since planar graphs have only planar quotients, $G(\Pi_n)$ is non-planar provided that $G(\Pi_1')$ is non-planar. Therefore, we can assume without loss of generality that the initial interstice of H is an odd flower. In this case we will prove that $G(\Pi_1)$ is non-planar.

When we build the marked rectangle pattern for H_1, we start with an odd horizontal run $M_1, ..., M_k$ of marked rectangles and then subdivide each M_i into a vertical run $\{M_{ij}\}$. We color the center of M_{i1} black for all i. We color all other centers white. In the graph $G'(H_1)$, all the white centers within a vertical run are connected together. Also, each black (respectively white) center in the ith vertical run is connected to some white (respectively black) center the $(i-1)$st vertical run and some white (respectively black) center in the $(i+1)$st vertical run.

From this description, we see that $G'(H_1)$ has a subgraph $G''(H_1)$, and this subgraph has the following quotient graph, X_k: The vertices of X_k are $b_1, ..., b_k, w_1, ..., w_k$. Taking indices mod k, an edge connects each b_j to w_{j-1}, w_j, and w_{j+1}. (This forces the same connections, with the b's and w's reversed.) When k is odd X_k is easily seen to be non-planar. Since X_k is non-planar, $G''(H_1)$ is non-planar. Hence $G'(H_1)$ is non-planar. $G(\Pi_1)$ is obtained from $G'(H_1)$ simply by inserting vertices at the midpoints of all the edges of $G'(H_1)$. Hence $G(\Pi_1)$ is non-planar.

7 Special Cases

7.1 Even Packings and Planarity

Here we sketch a proof that $Q(H)$ is planar when H has only even flowers. Say that a *marked quadrilateral* is a quadrilateral with the same markings as a marked rectangle. The left half of Figure 7.1 shows the subdivision rule for breaking a marked quadrilateral into 3 smaller ones. (We do not want to specify the exact geometry of this subdivision.) The main point of this construction is that the union of the three black (respectively white) diagonals still connects the black (respectively white) points of the original marked quadrilateral.

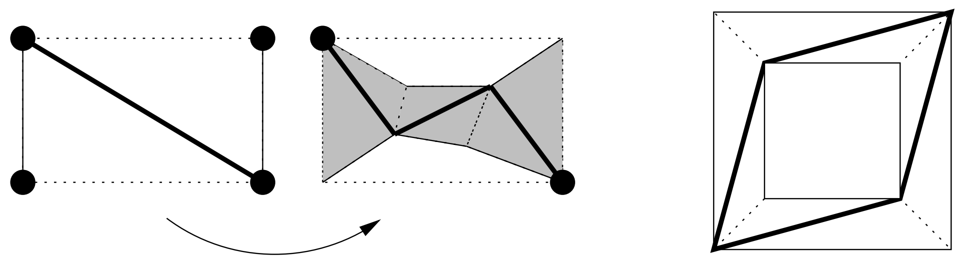

Figure 7.1

A similar subdivision rule can be made for any odd $k \geq 3$. In analogy to what we did in §4.3 we define a *quadrilateral r-ring* to be a succession of $r \geq 4$ marked quadrilaterals joined cyclically, along common dotted edges. Here $r \geq 4$ is always even. The right hand side of Figure 7.1 shows the case $r = 4$. The black and white diagonals can be chosen so as to make closed polygons. Using the alternative definitions, everything we did in §4-5 goes through, except perhaps the Shrinking Lemma. If all nested sequences of quadrilaterals shrink to points, then $Q(H)$ is homeomorphic to $\bigcap \Pi_n$, which is obviously planar using the new definitions.

The game is to choose the geometry of the subdivision, subject to the fixed combinatorics, so that nested sequences all shrink to points. Here we explain without proof one way to do this. Define the enormous sequence $\{S_n\}$ by the formula $S_0 = 1$ and $S_{n+1} = 10^{S_n}$. Suppose q is a quadrilateral of Π_n, and we need to subdivide q into $q_1, ..., q_k$. We do the subdivision so that the diameter of $q_2 \cup ... \cup q_{k-1}$ is less than $1/S_n$ times the length of the shortest side of q. We center this tiny union within $1/S_n$ of the midpoint of the longest dotted side of q. This method is fun to contemplate.

7.2 Odd Packings and String Art

We saw in §7.1 that we could define a simpler subdivision rule when we were dealing with horodisk packings in which all the flowers were even. The same thing is true when all the flowers are odd, though the main application of the construction is a whimsical one: It gives a method for building the corresponding circle quotient out of string.

Say that a *skew pair* is a pair of line segments in $\boldsymbol{R}^3$ which are contained in skew lines. We call one of the line segments *black* and one of them *white*. A marked tetrahedron is really an elaboration of this structure. If we start with a marked tetrahedron we form a skew pair by taking the colored edges.

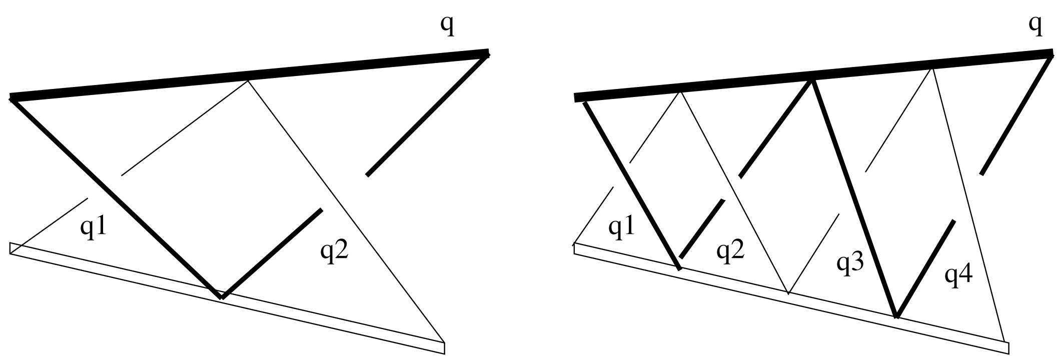

Figure 7.2

If q is a skew pair and k is an even integer, we form k new skew pairs, $q_1, ..., q_k$, as shown in Figure 7.2 for $k = 2$ and $k = 4$. The thickest pair of segments is q. The skew pairs q_j move sequentially along these segments, from left to right, as indicated by the labellings. The endpoints of the q_j are meant to be evenly spaced on the relevant segment of q.

The key point of our construction is that the union of the black (respectively white) segments of the q_j connects up the two endpoints of the black (respectively white) segment of q when k is even. Another virtue of our construction is that it is affinely natural, since the restriction of an affine map to a line preserves the notion of "even spacing". We call the set $\{q_j\}$ the *subdivision* of q.

We define a *k-skew ring* to be a collection of k skew pairs, joined end to end, so that the colors match. This definition works for any $k \geq 3$. However, since the subdivision rule for skew pairs works only when k is even, it seems natural only to consider k-skew rings when k is odd.

Given a finite horodisk packing H, we make some of the same constructions as in §4. We start with a k-skew ring, where k is the number of horodisks bounding the initial interstice. Then we subdivide the individual skew pairs according to the combinatorics of the packing. For instance, for the prototypical horodisk packing, we start with a skew 3-ring and then iteratively replace each skew pair by two smaller ones, as shown on the left side of Figure 7.2.

Performing this procedure for the finite horodisk packings $\{H_n\}$, which approximate the infinite horodisk packing H, we produce a sequence $\{B_n\}$ of black curves and a sequence $\{W_n\}$ of white curves. Each curve in this sequence is an PL embedding of a circle into $\boldsymbol{R}^3$. We now give a heuristic argument that the limit of B_n, in the Hausdorff topology on closed subsets of $\boldsymbol{R}^3$, is homeomorphic to $Q(H)$. The same statement is true for the limit of the W_n.

Recall that Π_n is the finite collection of marked tetrahedra associated to H_n. Let $\widehat{\Pi}_n$ be the union of the convex hulls of the skew pairs produced at the nth stage of our construction. From a combinatorial point of view, Π_n and $\widehat{\Pi}_n$ are identical. Geometrically, however, there is a difference: Every tetrahedron in Π_{n+1} is contained in a tetrahedron of Π_n, except possibly for one edge. This is no longer true for $\widehat{\Pi}_{n+1}$ and $\widehat{\Pi}_n$.

On the positive side, every tetrahedron of $\widehat{\Pi}_{n+1}$ is at contained in a tetrahedron of $\widehat{\Pi}_n$. Furthermore, it is not hard to show that every tetrahedron in $\widehat{\Pi}_{n+3}$ is, in the sense of Lemma 5.7, internal to a tetrahedron of $\widehat{\Pi}_n$. Consider the infinite intersection $\widehat{\Pi} = \bigcap \widehat{\Pi}_n$. If we knew that every nested sequence of tetrahedra shrinks to a point then essentially the same proof given in §5 would show that $Q(H)$ and $\widehat{\Pi}$ are homeomorphic. Furthermore, this shrinking implies that $\widehat{\Pi}$ is the Hausdorff limit of either of the two sequences of curves we defined.

We will not prove in general that nested sequences shrink to points. We will, however, sketch a proof when there is a bound to the size of the flowers of H. In this case, a variant of the Shrinking Lemma is true. As we mentioned above, it is not hard to check that every tetrahedron of $\widehat{\Pi}_{n+3}$ is internal to some tetrahedron in $\widehat{\Pi}_n$, in the sense of Lemma 5.7. Knowing this, and using the bound on the sizes of the flowers, one can produce a version of Lemma 5.7, for the *third* rather than *second* subdivision. From here, the proof of the Shrinking Lemma is the same.

8 References

[**S1**] R.E. Schwartz, *Degenerating the Complex Hyperbolic Ideal Triangle Groups*, Acta Mathematica **186** (2001)

[**S2**] R.E. Schwartz, *Hyperbolic Structures on the Complements of Selected Fractals*, preprint, 2001.

[**T**] W. Thurston, *Three Dimensional Manifolds, Kleinian Groups, and Hyperbolic Geometry*, Bulletin of the A.M.S., 1983.