in S, otherwise one could enlarge δ. By area considerations, such leaves are finite in number.) From §8.10, γ is realizable if and only if δ is. □

The picture of unrealizable laminations in $\mathcal{PL}_0(S)$ is the following. Let Δ_+ consist of all projective classes of transverse measures (allowing degenerate non-trivial cases) on $\chi_+ = \gamma_+ \cup U_i \epsilon(E_{i+})$. Δ_+ is convex in a coordinate system V_τ coming from any train track τ carrying χ_+.

To see a larger, complete picture, we must find a larger natural coordinate system. This requires a little stretching of our train tracks and imaginations. In fact, it is possible to find coordinate systems which are quite large. For any $\gamma \in \mathcal{PL}_0$, let $\Delta_\gamma \subset \mathcal{PL}_0$ denote the set of projective classes of measures on γ.

PROPOSITION 9.7.2. *Let γ be essentially complete. There is a sequence of train tracks τ_i, where τ_i is carried by τ_{i+1}, such that the union of natural coordinate systems $S_\gamma = U_i V_{\tau_i}$ contains all of $\mathcal{PL}_0 - \Delta_\gamma$.*

The proof will be given presently.

Since τ_i is carried by τ_{i+1}, the inclusion $V_{\tau_i} \subset V_{\tau+1}$ is a projective map (in $\mathcal{ML}_0$, 9.62
the inclusion is linear). Thus S_γ comes naturally equipped with a projective structure. We have not made this analysis, but the typical case is that $\gamma = \Delta_\gamma$. We think of S_γ as a stereographic coordinate system, based on projection from γ. (You may imagine $\mathcal{PL}_0$ as a convex polyhedron in $\mathbb{R}^n$, so that changes of stereographic coordinates are piecewise projective, although this finite-dimensional picture cannot be strictly correct, since there is no fixed subdivision sufficient to make all coordinate changes.)

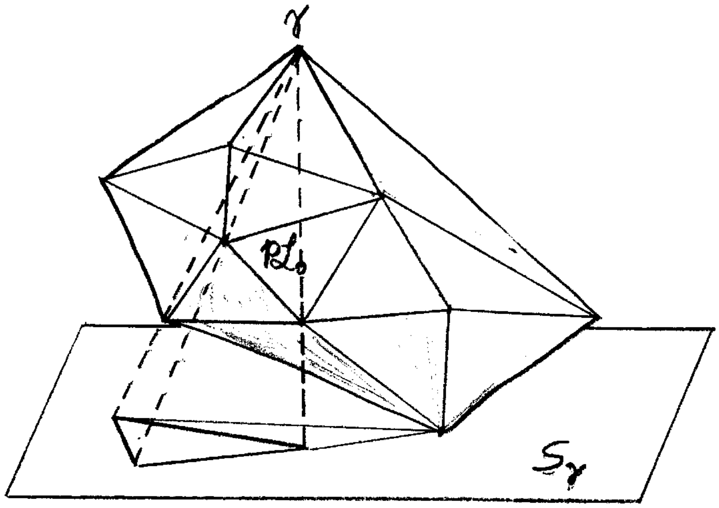

COROLLARY 9.7.3. *$\mathcal{PL}_0(S)$ is homeomorphic to a sphere.*

PROOF THAT 9.7.2 IMPLIES 9.7.3. Let $\gamma \in \mathcal{PL}_0(S)$ be any essentially complete lamination. Let τ be any train track carrying γ. Then $\mathcal{PL}_0(S)$ is the union of two coordinate systems $V_\tau \cup S_\tau$, which are mapped to convex sets in Euclidean space. 9.63
If $\Delta_\gamma \neq \gamma$, nonetheless the complement of Δ_γ in V_τ is homeomorphic to $V_\tau - \gamma$, so $\mathcal{PL}_0(S)$ is homeomorphic to the one-point compactification of S_γ. □

COROLLARY 9.7.4. *When $\mathcal{PL}_0(S)$ has dimension greater than* 1, *it does not have a projective structure.* (In other words, the *pieces* in changes of coordinates have not been eliminated.)

PROOF THAT 9.7.3 IMPLIES 9.7.4. The only projective structure on S^n, when $n > 1$, is the standard one, since S^n is simply connected. The binary relation of antipodality is natural in this structure. What would be the antipodal lamination for a simple closed curve α? It is easy to construct a diffeomorphism fixing α but moving any other given lamination. (If $i(\gamma, \alpha) \neq 0$, the Dehn twist around α will do.) □

REMARK. When $\mathcal{PL}_0(S)$ is one-dimensional (that is, when S is the punctured torus or the quadruply punctured sphere), the PIP structure *does* come from a projective structure, equivalent to $\mathbb{R}P^1$. The natural transformations of $\mathcal{PL}_0(S)$ are necessarily integral—in $\mathrm{PSL}_2(\mathbb{Z})$.

PROOF OF 9.7.2. Don't blink. Let γ be essentially complete. For each region R_i of $S - \gamma$, consider a smaller region r_i of the same shape but with finite points, rotated so its points alternate with cusps of R_i and pierce very slightly through the sides of R_i, ending on a leaf of γ. 9.64

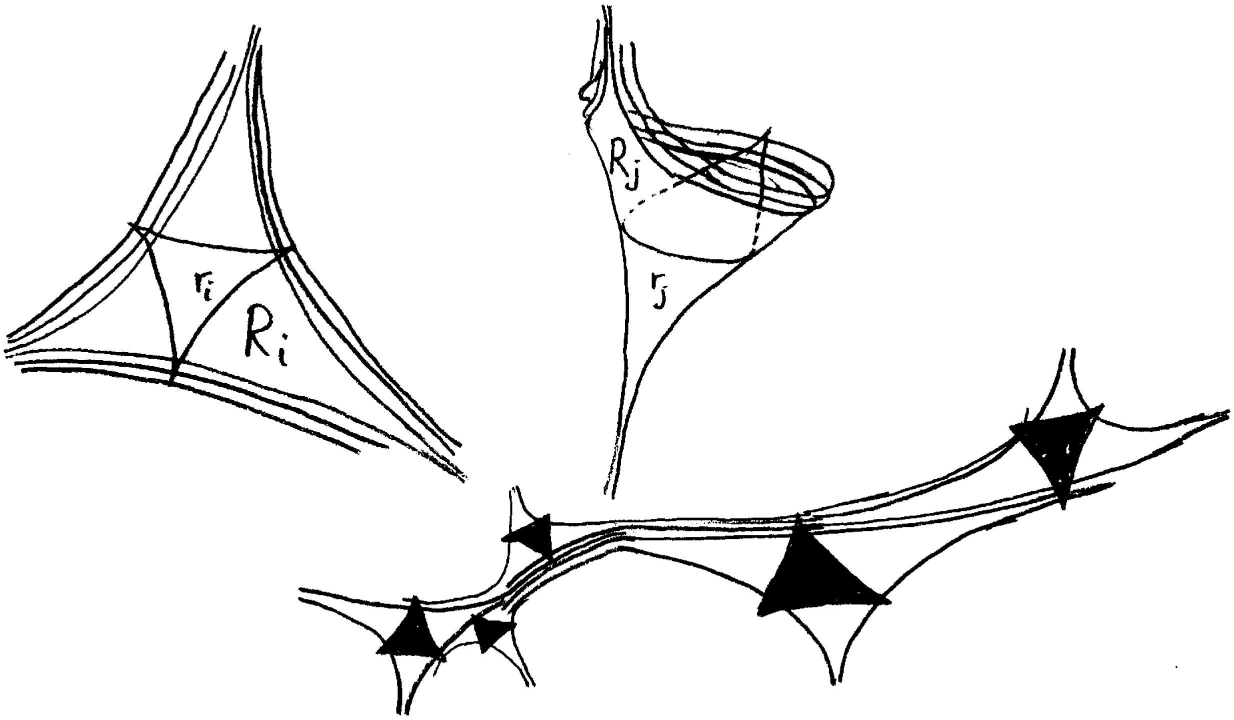

By 9.5.4, 9.5.2 and 9.3.9, both ends of each leaf of γ are dense in γ, so the regions r_i separate leaves of γ into arcs. Each region of $S - \gamma - U_i r_i$ must be a rectangle with two edges on ∂r_i and two on γ, since r_i covers the "interesting" part of R_i. (Or, prove this by area, χ). Collapse all rectangles, identifying the r_i edges with each other, and obtain a surface S' homotopy-equivalent to S, made of $U_i r_i$, where ∂r_i projects to a train track τ. (Equivalently, one may think of $S - U_i r_i$ as made of very wide corridors, with the horizontal direction given approximately by γ).

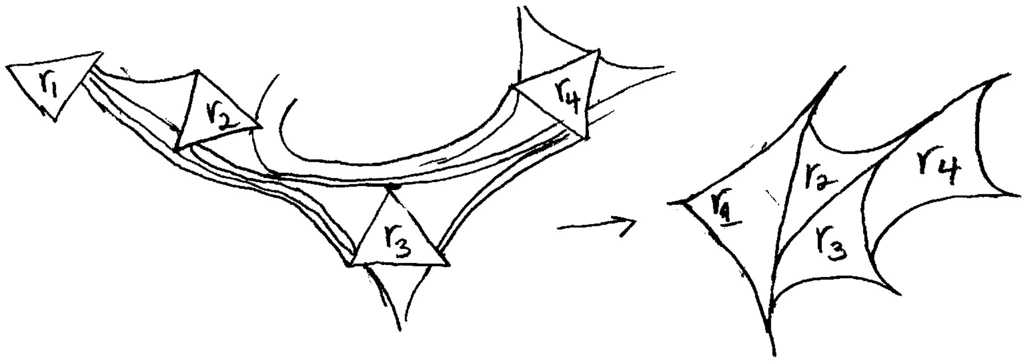

9.65

If we take shrinking sequences of regions $r_{i,j}$ in this manner, we obtain a sequence of train tracks τ_j which obviously have the property that τ_j carries τ_k when $j > k$. Let $\gamma' \in \mathcal{PL}_0(S) - \Delta_\gamma$ be any lamination not topologically equivalent to γ. From the density in γ of ends of leaves of γ, it follows that whenever leaves of γ and γ' cross, they cross at an angle. There is a lower bound to this angle. It also follows that $\gamma \cup \gamma'$ cuts S into pieces which are compact except for cusps of S.

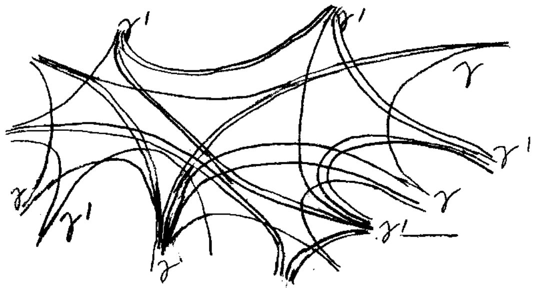

When R_i is an asymptotic triangle, for instance, it contains exactly one region of $S - \gamma - \gamma'$ which is a hexagon, and all other regions of $S - \gamma - \gamma'$ are rectangles. For sufficiently high j, the r_{ij} can be isotoped, without changing the leaves of γ which they touch, into the complement of γ'. It follows that γ' projects nicely to τ_j.

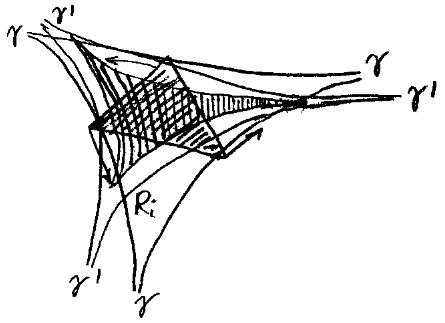

□

Stereographic coordinates give a method of computing and understanding intersection number. The transverse measure for γ projects to a "tangential" measure ν_γ on each of the train tracks τ_i: i.e., $\nu_\gamma(b)$ is the γ-transverse length of the sides of the rectangle projecting to b.

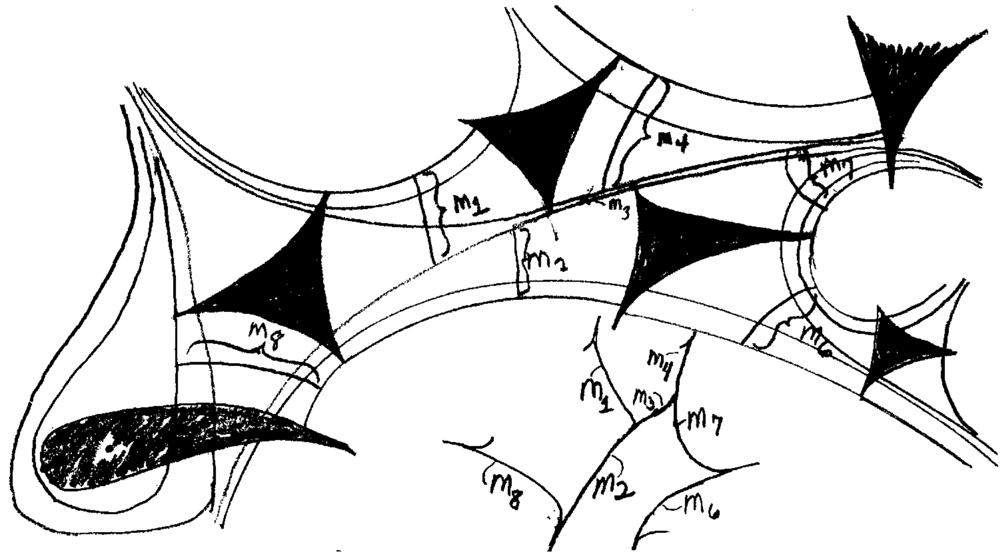

It is clear that for any $\alpha \in \mathcal{ML}_0$ which is determined by a measure μ_α on τ_i

9.7.5.
$$i(\alpha, \gamma) = \sum_b \mu_\alpha(b) \cdot \nu_\gamma(b).$$

Thus, in the coordinate system V_{τ_i} in $\mathcal{ML}_0$, intersection with γ is a linear function.

To make this observation more useful, we can reverse the process of finding a family of "transverse" train tracks τ_i depending on a lamination γ. Suppose we are given 9.67
an essentially complete train track τ, and a non-negative function (or "tangential" measure) ν on the branches of b, subject only to the triangle inequalities

$$a + b - c \geq 0 \quad a + c - b \geq 0 \quad b + c - a \geq 0$$

whenever a, b and c are the total ν-lengths of the sides of any triangle in $S - \tau$. We shall construct a "train track" τ^* dual to τ, where we permit regions of $S - \tau^*$ to be bigons as well as ordinary types of admissible regions—let us call τ^* a *bigon track.*

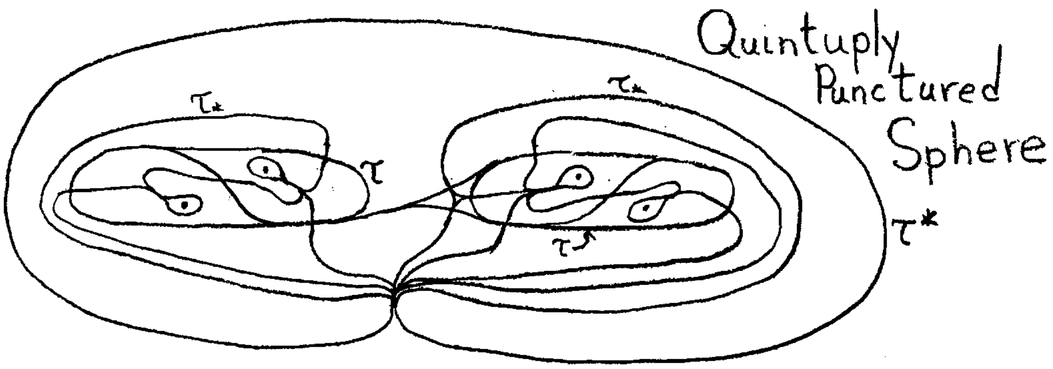

τ^* is constructed by shrinking each region R_i of $S - \tau$ and rotating to obtain a region $R_i^* \subset R_i$ whose points alternate with points of R_i. These points are joined using one more branch b^* crossing each branch b of τ; branches b_1^* and b_2^* are confluent at a vertex of R^* whenever b_1 and b_2 lie on the same side of R. Note that there is a bigon in $S - \tau^*$ for each switch in τ.

The tangential measure ν for τ determines a transverse measure defined on the branches of τ^* of the form b^*. This extends uniquely to a transverse for τ^* when S is not a punctured torus. 9.68

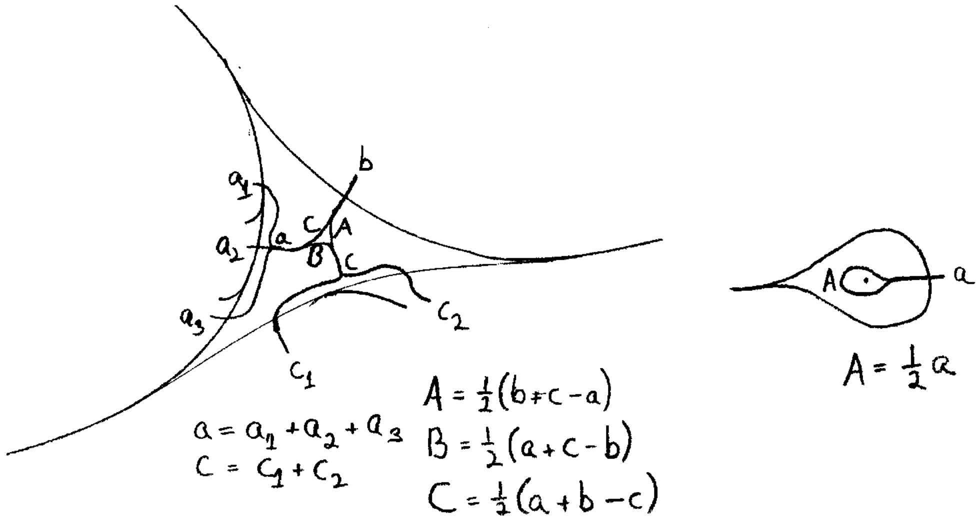

When S is the punctured torus, then τ must look like this, up to the homeomorphism (drawn on the abelian cover of $T - p$):

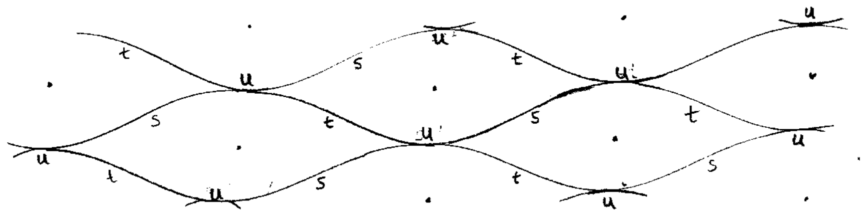

Note that each side of the punctured bigon is incident to each branch of τ. Therefore, the tangential measure ν has an extension to a transverse measure ν^* for τ^*, which is unique if we impose the condition that the two sides of R^* have equal transverse measure.

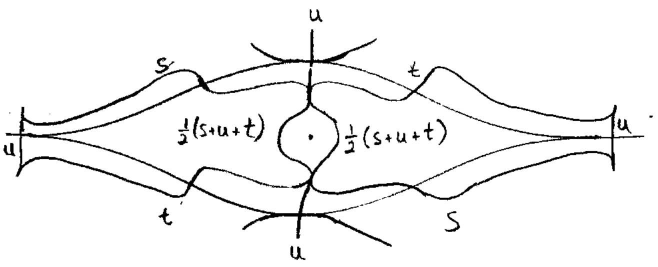

9.69

A transverse measure on a bigon track determines a measured geodesic lamination, by the reasoning of 8.9.4. When τ is an essentially complete train track, an open subset of $\mathcal{ML}_0$ is determined by a function μ on the branches of τ subject to a condition for each switch that

$$\sum_{b\in\mathcal{I}} \mu(b) = \sum_{b\in\mathcal{O}} \mu(b),$$

where $\mathcal{I}$ and $\mathcal{O}$ are the sets of "incoming" and "outgoing" branches. Dually, "tangential" measure ν on the branches of τ determines an element of $\mathcal{ML}_0$ (via ν^*), but two functions ν and ν' determine the same element if ν is obtained from ν' by a process of adding a constant to the incoming branches of a switch, and subtracting the same constant from the outgoing branches—or, in other words, if $\nu - \nu'$ annihilates all transverse measures for τ (using the obvious inner product $\nu \cdot \mu = \sum \nu(b)\mu(b)$). In fact, this operation on ν merely has the effect of switching "trains" from one side of a bigon to the other.

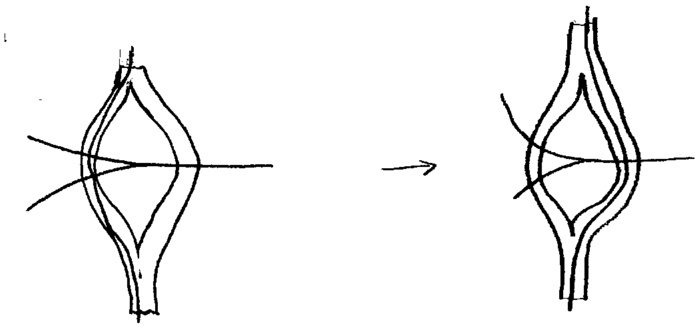

(Some care must be taken to obtain ν' from ν by a sequence of elementary "switching" operations without going through negative numbers. We leave this as an exercise to the reader.) 9.70

Given an essentially complete train track τ, we now have two canonical coordinate systems V_τ and V_τ^* in $\mathcal{ML}_0$ or $\mathcal{PL}_0$. If $\gamma \in V_\tau$ and $\gamma^* \in V_\tau^*$ are defined by measures μ_γ and ν_{γ^*} on τ, then $i(\gamma, \gamma^*)$ is given by the inner product

$$i(\gamma, \gamma^*) = \sum_{b\in\tau} \mu_\gamma(b)\nu_{\gamma^*}(b).$$

To see this, consider the universal cover of S. By an Euler characteristic or area argument, no path on $\tilde{\tau}$ can intersect a path on $\tilde{\tau}^*$ more than once. This implies the formula when γ and γ' are simple geodesics, hence, by continuity, for all measured geodesic laminations.

Proposition 9.7.4. *Formula* 9.7.3 *holds for all* $\gamma \in V_\tau$ *and* $\gamma^* \in V_\tau^*$. *Intersection number is a bilinear function on* $V_\tau \times V_\tau^*$ *(in* $\mathcal{ML}_0$*).* □

This can be interpreted as a more intrinsic justification for the linear structure on the coordinate systems V_τ—the linear structure can be reconstructed from the embedding of V_τ in the dual space of the vector space with basis $\gamma^* \in V_\tau^*$.

Corollary 9.7.5. *If* $\gamma, \gamma' \in \mathcal{ML}_0$ *are not topologically conjugate and if at least one of them is essentially complete, then there are neighborhoods* U *and* U' *of* γ *and* γ' *with linear structures in which intersection number is bilinear.*

9.71

Proof. Apply 9.7.4 to one of the train tracks τ_i constructed in 9.7.2. □

Remark. More generally, the only requirement for obtaining this local bilinearity near γ and γ' is that the complementary regions of $\gamma \cup \gamma'$ are "atomic" and that $S - \gamma$ have no closed non-peripheral curves. To find an appropriate τ, simply burrow out regions of r_i, "transverse" to γ with points going between strands of γ', so the regions r_i cut all leaves of γ into arcs. Then collapse to a train track carrying γ' and "transverse" to γ, as in 9.7.2.

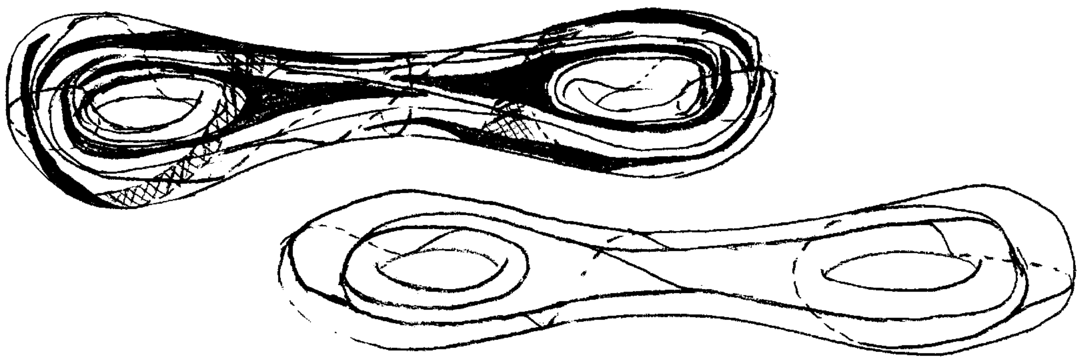

What is the image of $\mathbb{R}^n$ of stereographic coordinates S_γ for $\mathcal{ML}_0(S)$? To understand this, consider a system of train tracks

$$\tau_1 \to \tau_2 \to \cdots \to \tau_k \to \cdots$$

defining S_γ. A "transverse" measure for τ_i pushes forward to a "transverse" measure for τ_j, for $j > i$. If we drop the restriction that the measure on τ_i is non-negative, still it often pushes forward to a positive measure on τ_j. The image of S_γ is the set of 9.72

such arbitrary "transverse" measures on τ_1 which eventually become positive when pushed far enough forward.

For $\gamma' \in \Delta_\gamma$, let $\nu_{\gamma'}$ be a "tangential" measure on τ_1 defining γ'.

PROPOSITION 9.7.6. *The image of S_γ is the set of all "transverse," not necessarily positive, measures μ on τ_1 such that for all $\gamma' \in \Delta_\gamma$, $\nu_{\gamma'} \cdot \mu > 0$.*

(Note that the functions $\nu_{\gamma'} \cdot \mu$ and $\nu_{\gamma''} \cdot \mu$ are distinct for $\gamma' \neq \gamma''$.)

In particular, note that if $\Delta_\gamma = \gamma$, the image of stereographic coordinates for $\mathcal{ML}_0$ is a half-space, or for $\mathcal{PL}_0$ the image is $\mathbb{R}^n$. If Δ_γ is a k-simplex, then the image of S_γ for $\mathcal{PL}_0$ is of the form int $(\Delta^k) \times \mathbb{R}^{n-k}$. (This image is defined only up to projective equivalence, until a normalization is made.)

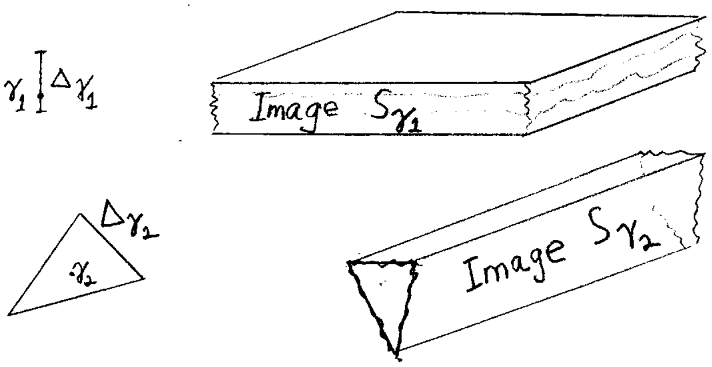

PROOF. The condition that $\nu_{\gamma'} \cdot \mu > 0$ is clearly necessary: intersection number $i(\gamma', \gamma'')$ for $\gamma' \in \Delta_\gamma$, $\gamma'' \in S_\gamma$ is bilinear and given by the formula 9.73

$$i(\gamma', \gamma'') = \nu_{\gamma'} \cdot \mu_{\gamma''}.$$

Consider any transverse measure μ on τ_1 such that μ is always non-positive when pushed forward to τ_i. Let b_i be a branch of τ_i such that the push-forward of μ is non-positive on b_i. This branch b_i, for high i, comes from a very long and thin rectangle ρ_i. There is a standard construction for a transverse measure coming from a limit of the average transverse counting measures of one of the sides of ρ_i. To make this more concrete, one can map ρ_i in a natural way to τ_j^* for $j \leq i$.

(In general, whenever an essentially complete train track τ carries a train track σ, then σ^* carries τ^*

$$\sigma \to \tau$$
$$\sigma^* \leftarrow \tau^*.$$

To see this, embed σ in a narrow corridor around τ, so that branches of τ^* do not pass through switches of σ. Now σ^* is obtained by squeezing all intersections of branches of τ^* with a single branch of σ to a single point, and then eliminating any bigons contained in a single region of $S - \sigma$.)

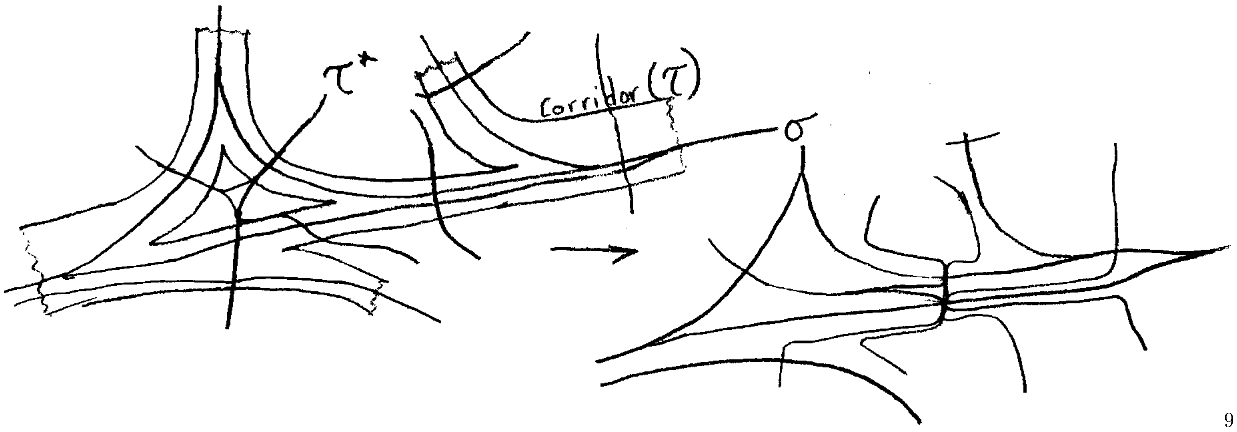

9.74

On τ_1^*, ρ_i is a finite but very long path. The average number of times ρ_i tranverses a branch of τ_1^* gives a function ν_i which almost satisfies the switch condition, but not quite. Passing to a limit point of $\{\nu_i\}$ one obtains a "transverse" measure ν for τ_1^*, whose lamination topologically equals γ, since it comes from a transverse measure on τ_i^*, for all i. Clearly $\nu \cdot \mu \leq 0$, since ν_i comes frm a function supported on a single branch b_i^* of τ_i^*, and $\mu(b_i) < 0$. □

For $\gamma \in \mathcal{ML}_0$ let $Z_\gamma \subset \mathcal{ML}_0$ consist of γ' such that $i(\gamma, \gamma') = 0$. Let C_γ consist of laminations γ' not intersecting γ, i.e., such that support of γ' is disjoint from the support of γ. An arbitrary element of Z_γ is an element of C_γ, together with some measure on γ. The same symbols will be used to denote the images of these sets in $\mathcal{PL}_0(S)$.

Proposition 9.7.6. *The intersection of Z_γ with any of the canonical coordinate systems X containing γ is convex.* (In $\mathcal{ML}_0$ or $\mathcal{PL}_0$.)

Proof. It suffices to give the proof in $\mathcal{ML}_0$. First consider the case that γ is a simple closed curve and $X = V_\tau$, for some train track τ carrying γ. Pass to the cylindrical covering space C of S with fundamental group generated by γ. The path of γ on C is embedded in the train track $\tilde{\tau}$ covering τ. From a "transverse" measure m on $\tilde{\tau}$, construct corridors on C with a metric giving them the proper widths. 9.75

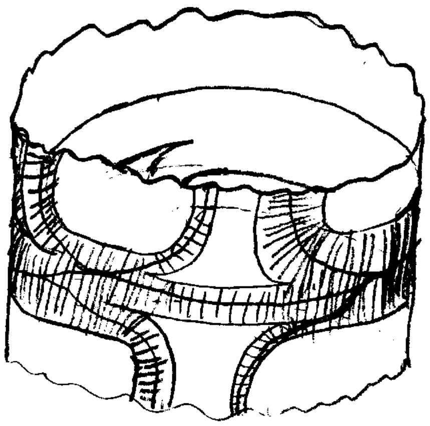

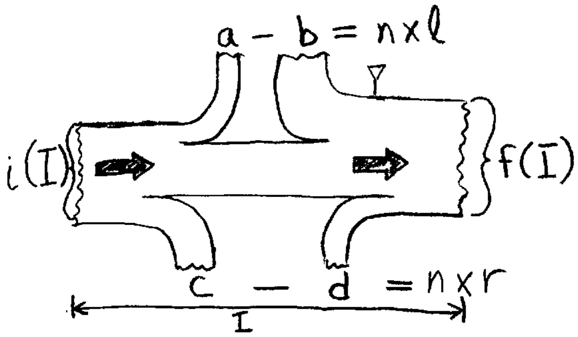

For any subinterval I of γ, let $\mathrm{nxr}(I)$ and $\mathrm{nxl}(I)$ be (respectively) the net right hand exiting and the net left hand exiting in the corresponding to I; in computing this, we weight entrances negatively. (We have chosen some orientation for γ). Let $i(I)$ be the initial width of I, and $f(I)$ be the final width.

If the measure m comes from an element γ', then $\gamma' \in Z_\gamma$ if and only if there is no "traffic" entering the corridor of γ on one side and exiting on the other. This implies the inequalities

$$i(I) \geq \mathrm{nxl}(I)$$

and

$$i(I) \geq \mathrm{nxr}(I)$$

for all subintervals I.

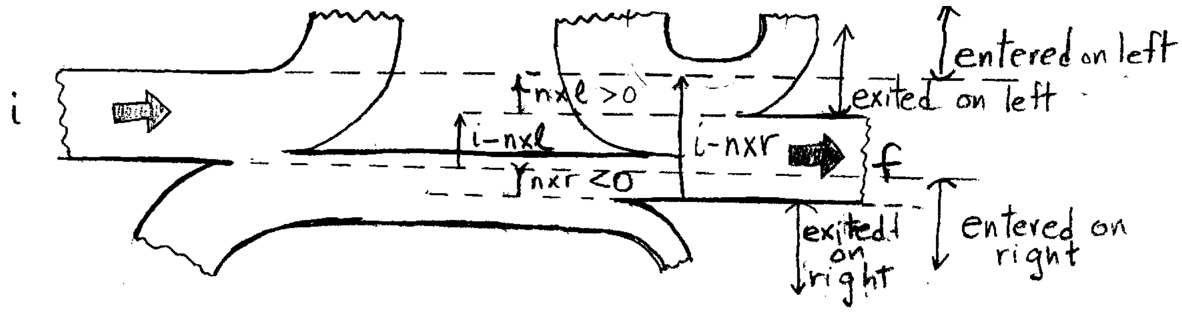

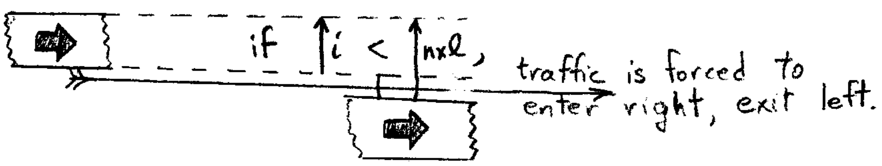

9.76

It also implies the equation

$$\mathrm{nxl}(\gamma) = 0,$$

so that any traffic travelling once around the corridor returns to its inital position. (Otherwise, this traffic would spiral around to the left or right, and be inexorably forced off on the side opposite to its entrance.)

Conversely, if these inequalities hold, then there is some trajectory going clear around the corridor and closing up. To see this, begin with any cross-section of the corridor. Let x be the supremum of points whose trajectories exit on the right. Follow the trajectory of x as far as possible around the corridor, always staying in the corridor whenever there is a choice.

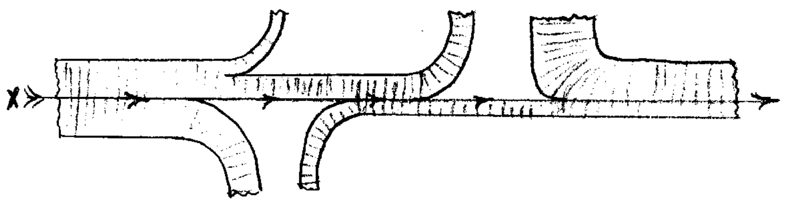

The trajectory can never exit on the left—otherwise some trajectory slightly lower would be forced to enter on the right and exit on the left, or vice versa. Similarly, it can't exit on the right. Therefore it continues around until it closes up.

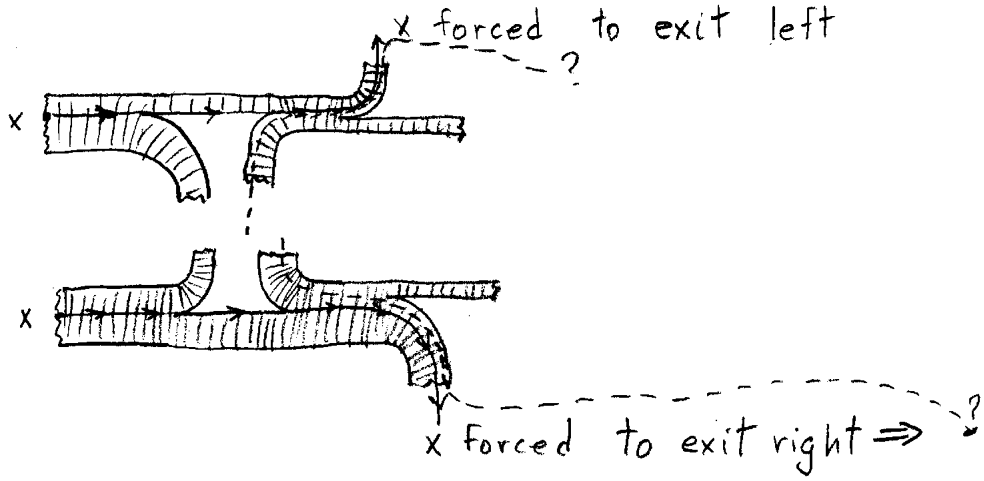

9.77

Thus when γ is a simple closed curve, $Z_\gamma \cap V_\tau$ is defined by linear inequalities, so it is convex.

Consider now the case $X = V_\tau$ and γ is connected but not a simple geodesic. Then γ is associated with some subsurface $M_\gamma \subset S$ with geodesic boundary defined to be the minimal convex surface containing γ. The set C_γ is the set of laminations not intersecting int (M_γ). It is convex in V_τ, since

$$C_\gamma = \bigcap \{Z_\alpha | \alpha \text{ is a simple closed curve } \subset \text{ int } (M_\gamma)\}.$$

A general element γ' of Z_γ is a measure on $\gamma \cup \gamma''$, so Z_γ consists of convex combinations of Δ_γ and C_γ: hence, it is convex.

If γ is not connected, then Z_γ is convex since it is the intersection of $\{Z_{\gamma_i}\}$, where the γ_i are the components of γ.

The case X is a stereographic coordinate system follows immediately. When $X = V_\tau^*$, consider any essentially complete $\gamma \in V_\tau$. From 9.7.5 it follows that V_τ^* is linearly embedded in S_γ. (Or more directly, construct a train track (without bigons) carrying τ^*; or, apply the preceding proof to bigon track τ^*.) $\square$

REMARK. Note that when γ is a union of simple closed curves, C_γ in $\mathcal{PL}_0(S)$ is homeomorphic to $\mathcal{PL}_0(S - \gamma)$, regarded as a complete surface with finite area—i.e., 9.78
C_γ is a sphere. When γ has no component which is a simple closed curve, C_γ is convex. Topologically, it is the join of $\mathcal{PL}_0(S - \bigcup S_\gamma)$ with the simplex of measures on the boundary components of the S_{γ_i}, where the S_{γ_i} are subsurfaces associated with the components γ_i of γ.

Now we are in a position to form an image of the set of unrealizable laminations for $\rho\pi_1 S$. Let $U_+ \subset \mathcal{PL}_0$ be the union of laminations containing a component of χ_+ and define U_- similarly, so that γ is unrealizable if and only if $\gamma \in U_+ \cup U_-$. U_+ is a union of finitely many convex pieces, and it is contained in a subcomplex of $\mathcal{PL}_0$ of codimension at least one. It may be disjoint from U_-, or it may intersect U_- in an interesting way.

EXAMPLE. Let S be the twice punctured torus. From a random essentially complete train track,

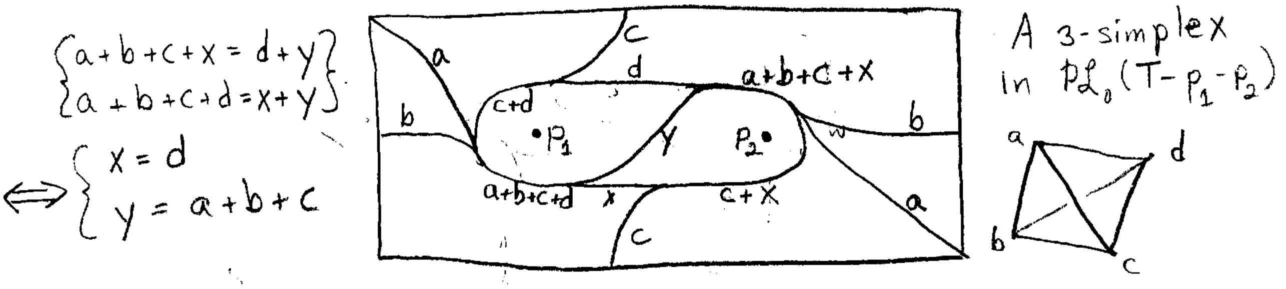

we compute that $\mathcal{ML}_0$ has dimension 4, so $\mathcal{PL}_0$ is homeomorphic to S^3. For any simple closed curve α on S, C_α is $\mathcal{PL}_0(S - \alpha)$,

9.79

where $S - \alpha$ is either a punctured torus union a (trivial) thrice punctured sphere, or a 4-times punctured sphere. In either case, C_α is a circle, so Z_α is a disk.

Here are some sketches of what U_+ and $U-$ can look like.

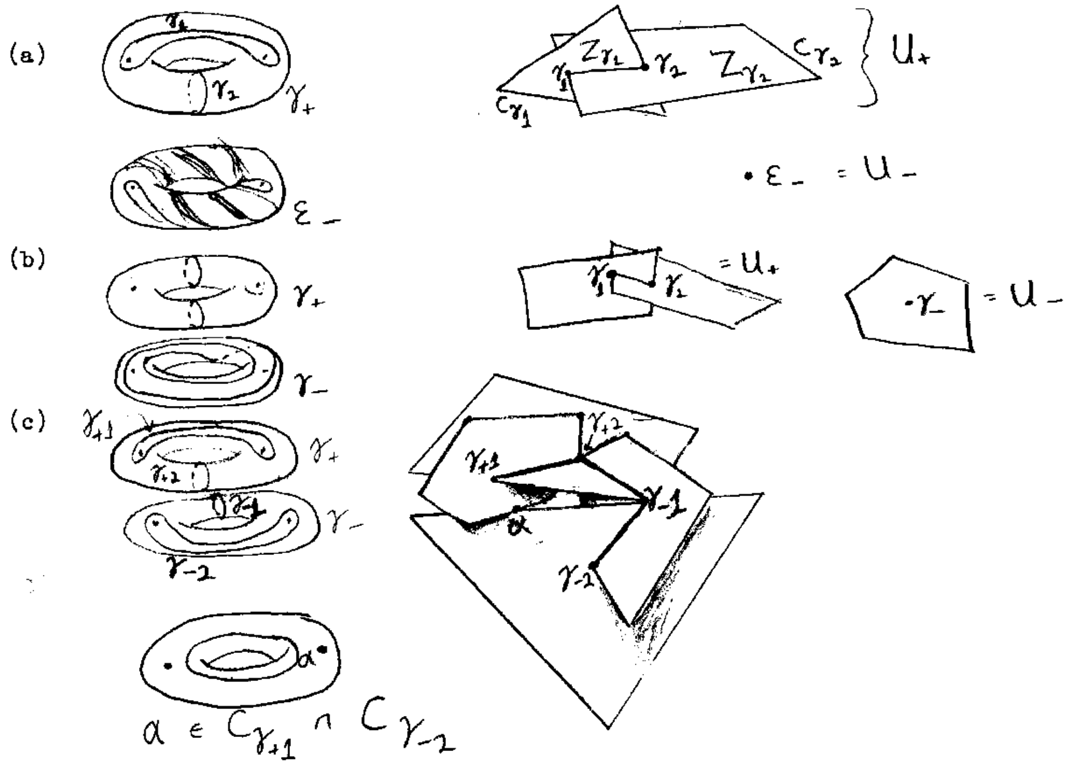

Here is another example, where S is a surface of genus 2, and $U_+(S) \cup U_-(S)$ has the homotopy type of a circle (although its closure is contractible):

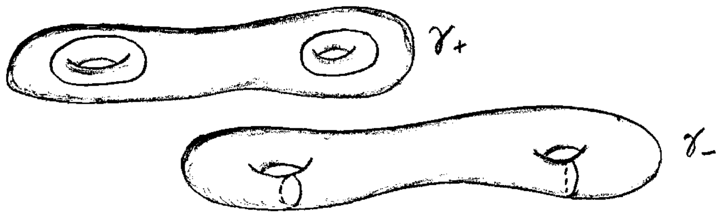

9.80

In fact, $U_+ \cup U_-$ is made up of convex sets $Z_\gamma - C_\gamma$, with relations of inclusion as diagrammed:

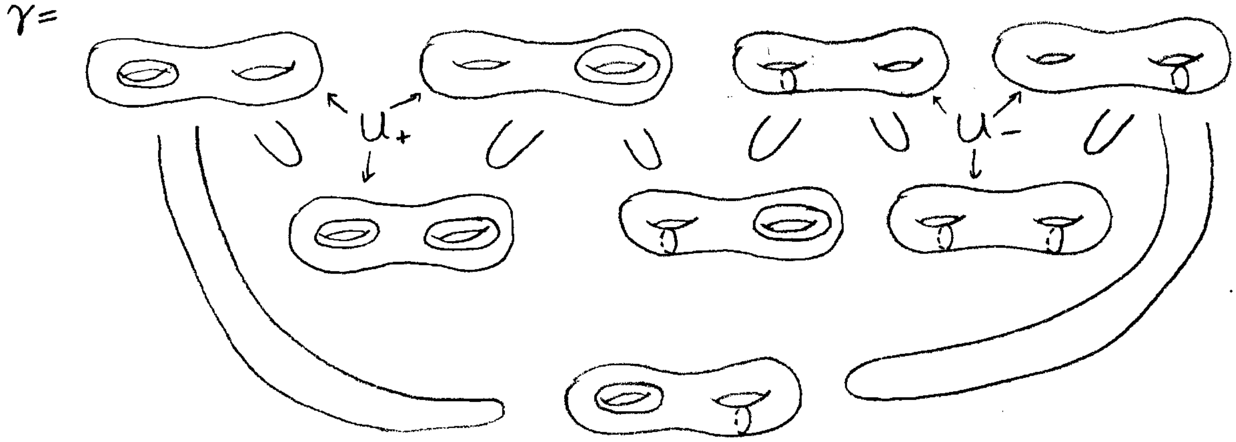

The closures all contain the element α; hence the closure of the union is starlike:

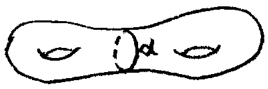

9.9-1

9.9. Ergodicity of the geodesic flow

We will prove a theorem of Sullivan (1979): There is no §9.8

THEOREM 9.9.1. *Let M^n be a complete hyperbolic manifold (of not necessarily finite volume). Then these four conditions are equivalent:*

(a) *The series*

$$\sum_{\gamma \in \pi_1 M^n} \exp\bigl(-(n-1)\, d(x_0, \gamma x_0)\bigr)$$

diverges. (Here, $x_0 \in H^n$ is an arbitrary point, γx_0 is the image of x_0 under a covering transformation, and $d(,)$ is hyperbolic distance).

(b) *The geodesic flow is not dissipative. (A flow ϕ_t on a measure space (X, μ) is dissipative if there exists a measurable set $A \subset X$ and a $T > 0$ such that $\mu(A \cap \phi_t(A)) = 0$ for $t > T$, and $X = \cup_{t\phi_t(A)}$.)*

(c) *The geodesic flow on $T_1(M)$ is recurrent. (A flow ϕ_t on a measure space (X, μ) is recurrent when for every measure set $A \subset X$ of positive measure and every $T > 0$ there is a $t \geq T$ such that $\mu(A \cap \phi_t(A)) > 0$.)*

(d) *The geodesic flow on $T_1(M)$ is ergodic.*

Note that in the case M has finite volume, recurrence of the geodesic flow is immediate (from the Poincaré recurrence lemma). The ergodicity of the geodesic flow in this case was proved by Eberhard Hopf, in ??. The idea of (c) $\to$ (d) goes back to Hopf, and has been developed more generally in the theory of Anosov flows ??. 9.9-2

Corollary 9.9.2. *If the geodesic flow is not ergodic, there is a non-constant bounded superharmonic function on M.*

Proof of 9.9.2. Consider the Green's function $g(x) = \int_{d(x,x_0)}^{\infty} \sin h^{1-n} t \, dt$ for hyperbolic space. (This is a harmonic function which blows up at x_0.) By (a), the series $\sum_{\gamma \in \pi_1 M} g \circ \gamma$ converges to a function, invariant by γ, which projects to a Green's function G for M. The function $f = \arctan G$ (where $\arctan \infty = \pi/2$) is a bounded superharmonic function, since arctan is convex. □

Remark. The convergence of the series (a) is actually equivalent to the existence of a Green's function on M, and also equivalent to the existence of a bounded superharmonic function. See (Ahlfors, Sario) for the case $n = 2$, and [ ] for the general case.

Corollary 9.9.3. *If Γ is a geometrically tame Kleinian group, the geodesic flow on $T_1(H^n/\Gamma)$ is ergodic if and only if $L_\Gamma = S^2$.*

Proof of 9.9.3. From 9.9.2 and 8.12.3. □

Proof of 9.9.1. Sullivan's proof of 9.9.1 makes use of the theory of Brownian motion on M^n. This approach is conceptually simple, but takes a certain amount of technical background (or faith). Our proof will be phrased directly in terms of geodesics, but a basic underlying idea is that a geodesic behaves like a random path: its future is "nearly" independent of its past. 9.9-2a

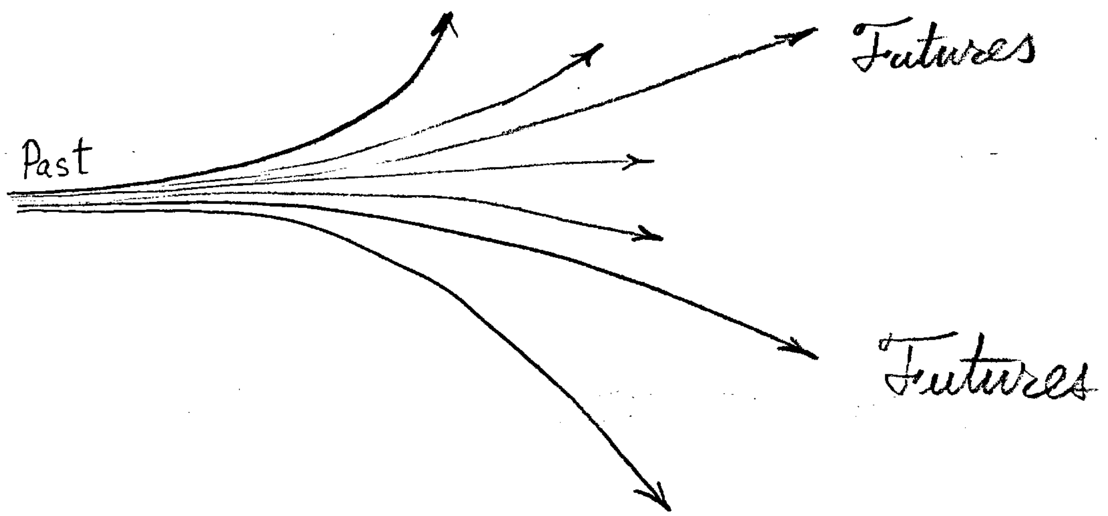

9.9-3

(d) → (c). This is a general fact. If a flow ϕ_t is not recurrent, there is some set A of positive measure such that only for t in some bounded interval is $\mu(A \cap \phi_t(A)) > 0$.

Then for any subset $B \subset A$ of small enough measure, $\cup_t \phi_t(B)$ is an invariant subset which is proper, since its intersection with A is proper.

(c) $\to$ (b). Immediate.

(b) $\to$ (a). Let B be any ball in H^n, and consider its orbit ΓB where $\Gamma = \pi_1 M$. For the series of (a) to diverge means precisely that the total apparent area of ΓG as seen from a point $x_0 \in H^n$, (measured with multiplicity) is infinite.

In general, the underlying space of a flow is decomposed into two measurable parts, $X = D \cup R$, where ϕ_t is dissipative on D (the union of all subsets of X which eventually do not return) and recurrent on R. The reader may check this elementary fact. If the recurrent part of the geodesic flow is non-empty, there is some ball B in M^n such that a set of positive measure of tangent vectors to points of B give rise to geodesics that intersect B infinitely often. This clearly implies that the series of (a) diverges.

The idea of the reverse implication (a) $\to$ (b) is this: if the geodesic flow is dissipative there are points x_0 such that a positive proportion of the visual sphere is not covered infinitely often by images of some ball. Then for *each* "group" of geodesics that return to B, a definite proportion must eventually escape ΓB, because future and past are nearly independent. The series of (a) can be regrouped as a geometric progression, so it converges. We now make this more precise.

Recall that the term "visual sphere" at x_0 is a synonym to the "set of rays" emanating from x_0. It has a metric and a measure obtained from its identification with the unit sphere in the tangent space at x_0. 9.9-4

Let $x_0 \in M^n$ be any point and $B \subset M^n$ any ball. If a positive proportion of the rays emanating from x_0 pass infinitely often through B, then for a slightly larger ball B', a definite proportion of the rays emanating from *any* point $x \in M^n$ spend an infinite amount of time in B', since the rays through x are parallel to rays through x_0. Consequently, a subset of $T_1(B')$ of positive measure consists of vectors whose geodesics spend an infinite total time in $T_1(B')$; by the Poincaré recurrence lemma, the set of such vectors is a recurrent set for the geodesic flow. (b) holds so (a) $\to$ (b) is valid in this case. To prove (a) $\to$ (b), it remains to consider the case that almost every ray from x_0 eventually escapes B; we will prove that (a) fails, i.e., the series of (a) converges.

Replace B by a slightly smaller ball. Now almost every ray from almost every point $x \in M$ eventually escapes the ball. Equivalently, we have a ball $B \subset H^n$ such that for every point $x \in H^n$, almost no geodesic through x intersects ΓB, or even $\Gamma(N_\epsilon(B))$, more than a finite number of times.

Let x_0 be the center of B and let α be the infimum, for $y \in H^n$, of the diameter of the set of rays from x_0 which are parallel to rays from y which intersect B. This infimum is positive, and very rapidly approached as y moves away from x_0.

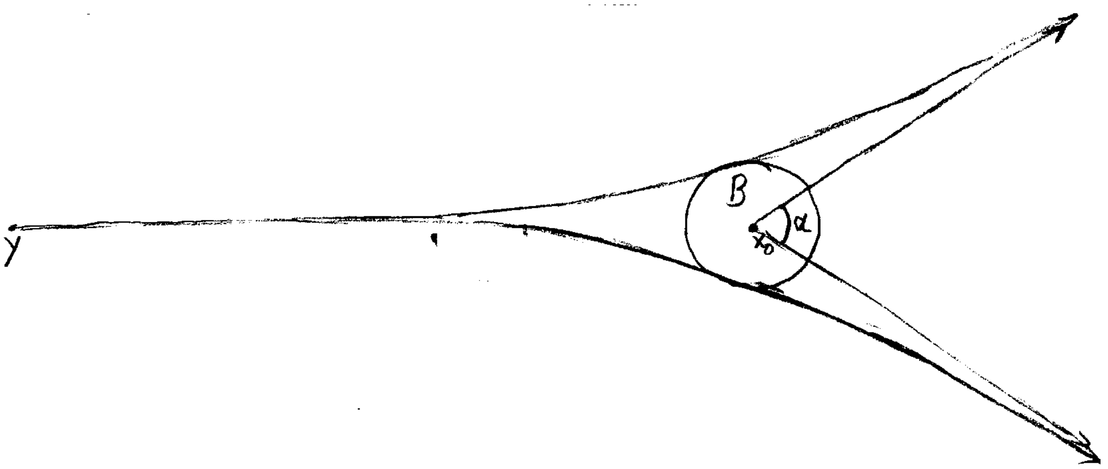

9.9-5

Let R be large enough so that for every ball of diameter greater than α in the visual sphere at x_0, at most (say) half of the rays in this ball intersect $\Gamma N_{\in}(B)$ at a distance greater than R from x_0. R should also be reasonably large in absolute terms and in comparison to the diameter of B.

Let x_0 be the center of B. Choose a subset $\Gamma' \subset \Gamma$ of elements such that: (i) for every $\gamma \in \Gamma$ there is a $\gamma' \in \Gamma'$ with $d(\gamma' x_0, \gamma x_0) < R$. (ii) For any γ_1 and γ_2 in Γ', $d(\gamma_1 x_0, \gamma_2 x_0) \geq R$.

Any subset of Γ maximal with respect to (ii) satisfies (i).

We will show that $\sum_{\gamma' \in \Gamma'} \exp(-(n-1)\, d(x_0, \gamma' x_0))$ converges. Since for any γ' there are a bounded number of elements $\gamma \in \Gamma$ so that $d(\gamma x_0, \gamma' x_0) < R$, this will imply that the series of (a) converges.

Let $<$ be the partial ordering on the elements of Γ' generated by the relation $\gamma_1 < \gamma_2$ when $\gamma_2 B$ eclipses $\gamma_1 B$ (partially or totally) as viewed from x_0; extend $<$ to be transitive.

Let us denote the image of γB in the visual sphere of x_0 by B_γ. Note that when $\gamma' < \gamma$, the ratio $\operatorname{diam}(B_{\gamma'}) / \operatorname{diam}(B_\gamma)$ is fairly small, less than $1/10$, say. Therefore $\cup_{\gamma' < \gamma} B_{\gamma'}$ is contained in a ball concentric with B_γ of radius $10/9$ that of B_γ.

Choose a maximal independent subset $\Delta_1 \subset \Gamma'$ (this means there is no relation $\delta_1 < \delta_2$ for any $\delta_1, \delta_2 \in \Delta_1$). Do this by successively adjoining any γ whose B_γ has largest size among elements not less than any previously chosen member. Note that area $(\cup_{\delta \in \Delta} B_\delta) / \operatorname{area}(\cup_{\gamma \in \Gamma'} B_\gamma)$ is greater than some definite (a priori) constant: $(9/10)^{n-1}$ in our example. Inductively define $\Gamma'_0 = \Gamma'$, $\gamma'_{i+1} = \Gamma'_i - \Delta_{i+1}$ and define $\Delta_{i+1} \subset \Gamma_i$ similarly to Δ_1. Then $\Gamma' = \cup_{i=1}^{\infty} \Delta_i$. 9.9-6

For any $\gamma \in \Gamma'$, we can compare the set B_γ of rays through x_0 which intersect $\gamma(B)$ to the set C_γ of parallel rays through γX_0.

Any ray of B_γ which re-enters $\Gamma'(B)$ after passing through $\gamma'(B)$, is within ϵ of the parallel ray of C_γ by that time. At most half of the rays of C_γ ever enter $N_\epsilon(\Gamma' B)$.

The distortion between the visual measure of B_γ and that of C_γ is modest, so we can conclude that the set of reentering rays, $B_\gamma \cap \bigcup_{\gamma' < \gamma} B_{\gamma'}$, has measure less than 2/3 the measure of B_γ.

We conclude that, for each i,

$$\begin{aligned} &\text{area}\,\big(\bigcup_{\gamma\in\Gamma'_{i+1}} B_\gamma\big) - \text{area}\,\big(\bigcup_{\gamma\in\Gamma'_i} B_\gamma\big) \\ &\geq 1/3\ \text{area}\,\big(\bigcup_{\delta\in\Delta_{i+1}} B_\delta\big) \\ &\geq 1/3\cdot(9/10)^{n-1}\,\text{area}\,\big(\bigcup_{\gamma\in\Gamma'_i} B_\gamma\big). \end{aligned}$$

The sequence $\{\text{area}(\bigcup_{\gamma\in\Gamma'_i} B_\gamma)\}$ decreases geometrically. This sequence dominates the terms of the series $\sum_i \text{area} \cup_{\delta\in\Delta_i} B_\delta = \sum_{\gamma\in\Gamma'} \text{area}(B_\gamma)$, so the latter converges, which completes the proof of (a) $\to$ (b). 9.9-7

(b) $\to$ (c). Suppose $R \subset T_1(M^n)$ is any recurrent set of positive measure for the geodesic flow ϕ_t. Let B be a ball such that $R \cap T_1(B)$ has positive measure. Almost every forward geodesic of a vector in R spends an infinite amount of time in B. Let $A \subset T_1(B)$ consist of all vectors whose forward geodesics spend an infinite time in B and let ψ_t, $t \geq 0$, be the measurable flow on A induced from ϕ_t which takes a point leaving A immediately back to its next return to A.

Since ψ_t is measure preserving, almost every point of A is in the image of ψ_t for all t and an inverse flow ψ_{-t} is defined on almost all of A, so the definition of A is unchanged under reversal of time. Every geodesic parallel in either direction to a geodesic in A is also in A; it follows that $A = T_1(B)$. By the Poincaré recurrence lemma, ψ_t is recurrent, hence ϕ_t is also recurrent.

(c) $\to$ (d). It is convenient to prove this in the equivalent form, that if the action of Γ on $S^{n-1}_\infty \times S^{n-1}_\infty$ is recurrent, it is ergodic. “Recurrent” in this context means that for any set $A \subset S^{n-1} \times S^{n-1}$ of positive measure, there are an infinite number of elements $\gamma \in \Gamma$ such that $\mu(\gamma A \cap A) > 0$. Let $I \subset S^{n-1} \times S^{n-1}$ be any measurable set invariant by Γ. Let $-B_1$ and $B_2 \subset S^{n-1}$ be small balls. Let us consider what I must look like near a general point $x = (x_1, x_2) \in B_1 \times B_2$. If γ is a “large” element of Γ such that γx is near x, then the preimage of γ of a product of small ϵ-ball around γx_1 and γx_2 is one of two types: it is a thin neighborhood of one of the factors, $(x_1 \times B_2)$ or $(B_1 \times x_2)$. (γ must be a translation in one direction or the other along an axis from approximately x_1 to approximately x_2.) Since Γ is recurrent, almost every point $x \in B_1 \times B_2$ is the preimage of elements γ of both types, of an infinite number of 9.9-8 points where I has density 0 or 1. Define

$$f(x_1) = \int_{B_2} \chi_I(x_1, x_2)\, dx_2,$$

where χ_I is the characteristic function of I, for $x_1 \in B_1$ (using a probability measure on B_2). By the above, for almost every x_1 there are arbitrarily small intervals around x_1 such that the average of f in that interval is either 0 or 1. Therefore f is a characteristic function, so $I \cap B_1 \times B_2$ is of the form $S \times B_2$ (up to a set of measure zero) for some set $S \subset B_1$.

Similarly, I is of the form $B_1 \times R$, so I is either $\emptyset \times \emptyset$ or $B_1 \times B_2$ (up to a set of measure zero). $\square$