Invent. math. 99, 425–454 (1990)

Inventiones mathematicae
© Springer-Verlag 1990

Iteration on Teichmüller space

C. McMullen*
Princeton University, Princeton, NJ 08544, USA

1. Introduction: Uniformization of 3-manifolds

The idea that the search for a geometric structure leads to iteration on Teichmüller space occurs in Thurston's approach to

- the classification of surface mapping classes [Th2] (cf. [FLP, Bers 2]),
- the combinatorics of rational maps [Th6] (cf. [DH]), and
- the uniformization of 3-manifolds [Th5] (cf. [Mor]).

In all three cases, a geometrization problem is formulated as the existence of a fixed point for a self-map f of Teichmüller space.

One might imagine that the behavior of any natural iteration on Teichmüller space can be understood with the theory of Riemann surfaces. This is true in the first two examples because the iteration is defined using surfaces.

In this paper we use Riemann surface techniques to study the third iteration, and provide a new proof of a fundamental step in the geometrization of 3-manifolds. (An expository account appears in [Mc2].)

Here is a simplified rendering of the construction of hyperbolic structures on atoroidal Haken manifolds. A finite union of balls obviously carries a hyperbolic structure. Start gluing them together along incompressible submanifolds of their boundary—any Haken manifold can be built up in this way. By an orbifold trick (using Andreev's theorem), one need only deal with the case of gluing along entire boundary components. At the inductive step, one has a geometrically finite hyperbolic realization N of a 3-manifold M with incompressible boundary, and gluing instructions encoded by an orientation-reversing involution τ: $\partial M \to \partial M$. The construction is completed by the following result:

Theorem 1.1 (Thurston). *M/τ has a hyperbolic structure if and only if the quotient is atoroidal.*

Here is how this theorem is related to iteration on Teichmüller space. By a generalization of Mostow rigidity, there is an identification

$$GF(M) \cong \text{Teich}(\partial M)$$

* Research partially supported by an NSF Postdoctoral Fellowship

between the space $GF(M)$ of geometrically finite realizations of M and the Teichmüller space of ∂M. The *skinning map*

$$\sigma: \mathrm{Teich}(\partial M) \to \mathrm{Teich}(\overline{\partial M})$$

is defined as follows: given a realization $N \in GF(M)$, form quasifuchsian covering spaces for each component of ∂M, and record the conformal structure on the new ends which appear. The gluing instructions determine an isometry

$$\tau: \mathrm{Teich}(\overline{\partial M}) \to \mathrm{Teich}(\partial M) ,$$

and a fixed point for $\tau \circ \sigma$ gives a new hyperbolic structure where the ends to be glued have compatible shape.

We propose to approach Theorem 1.1 via a study of $\|d\sigma\|$, the Teichmüller norm of the derivative of the skinning map. It follows from general principles that $\|d\sigma\| \leqq 1$ since σ is holomorphic and the Teichmüller metric coincides with the Kobayashi metric [Roy].

The case in which M is an interval bundle over a surface proves exceptional: then σ is an isometry.

Otherwise $\|d\sigma\| < 1$ pointwise, so if a fixed point exists it can be found by iteration. Our result can be summarized as follows.

Theorem 1.2 *Assume M is not an interval bundle over a surface.*
1. *If M is acylindrical, $\|d\sigma_N\| < c < 1$ at all points $N \in GF(M)$. Thus any gluing problem for M has a solution.*
2. *In general, $\|d(\tau \circ \sigma)^k_N\| < c([\partial N]) < 1$, where c is a continuous function of the location of ∂N in* moduli *space, and $(\tau \circ \sigma)^k$ denotes some fixed iterate of $\tau \circ \sigma$.*
3. *In this case, either*

   (a) *$\tau \circ \sigma$ has a fixed point, or*
   (b) *as $n \to \infty$ the ends of $(\tau \circ \sigma)^n(N)$ develop short geodesics bounding cylinders joined by τ to form a nontrivial torus in N/τ.*

   Thus the gluing problem has a solution iff M/τ is atoroidal.

Sketch of the proof. Let $Y \to X$ be a large covering space of a Riemann surface X of finite area—e.g., the universal covering, or more generally the covering determined by a proper incompressible subsurface. In [Mc1] we show the natural map $\mathrm{Teich}(X) \to \mathrm{Teich}(Y)$ (obtained by lifting complex structures from X to Y) is a contraction for the Teichmüller metric. Geometrically, this means Teichmüller mappings can be relaxed to maps of lesser dilatation when lifted to a sufficiently large cover.

Such relaxation can be exploited to estimate the contraction of the skinning map. This estimate depends only on the location of ∂N in moduli space, so the contraction is uniform except in the presence of short geodesics.

A short geodesic controls the geometry of a hyperbolic manifold over a large distance, by the Margulis lemma. When combined with the theory of geometric limits of quadratic differentials (cf. Appendix to [Mc1]), one finds a qualitative picture which again forces contraction *unless* the short geodesic lies on one end of a compressing cylinder in the 3-manifold.

In this case, either some iterate of $\tau \circ \sigma$ has significant contraction, or the short geodesics cycle and M/τ is toroidal.

§2 assembles the results on Teichmüller theory and geometric limits of quadratic differentials developed in [Mc1]. §3 discusses hyperbolic 3-manifolds and the skinning map. In §4 we prove a compact convex acylindrical hyperbolic 3-manifold can be deformed so its boundary becomes totally geodesic; the proof is short and it gives a good introduction to the general gluing problem. §5 derives the dependence on moduli, and §6 establishes uniform contraction for acylindrical manifolds. It is here we analyze short geodesics. §7 treats cylindrical manifolds and completes the proof of Theorem 1.2.

Remarks

1. The proof of Theorem 1.2 parallels similar approaches to the mapping class group and to rational maps.

Any iteration f on Teichmüller space, such that the distance from X to $f(X)$ remains bounded as $[X]$ tends to infinity in moduli space, preserves a system of short geodesics on X.

Thus a mapping class whose minimal translation distance is not achieved is reducible.

For critically finite branched covers of the sphere, the amount of contraction of f is controlled by the location of X in moduli space. Hence either $[f^k(X)]$ remains in a compact set—and one constructs a rational map—or one finds a system of annuli providing an obstruction to rationalization.

In the case of 3-manifolds, either one obtain a hyperbolic structure, or one finds a system of simple closed curves bounding cylinders which link up to form a nonperipheral torus. This torus provides a topological obstruction to geometrization.

2. To give some idea of other approaches to the skinning map, we focus on the acylindrical case.

Thurston has shown that the space $AH(M)$ of *all* hyperbolic structures on an acylindrical manifold is compact in the algebraic topology [Th1]. This result has also been obtained using trees and valuations by Morgan and Shalen [MS1, MS2]. To obtain contraction, one first extends the skinning map to $AH(M)$, then uses the theory of geometrically tame ends to show the image lies entirely in $QH_0(M)$, the space of hyperbolic structures quasi-isometric to the initial one N. It follows that σ has a bounded image in Teichmüller space, hence is uniformly contracting.

3. Our techniques do not yield the bounded image theorem, nor do they apply to the case $M = S \times [0,1]$ where S is a surface. This case arises naturally for 3-manifolds which fiber over the circle.

Acknowledgements. The connection between Kra's Θ conjecture and the skinning map was first pointed out to me many years ago by John Hubbard. I'd also like to thank Bill Thurston for explaining much of his original proof, Andrew Casson, Howard Masur and Peter Shalen for suggesting revisions, and Ecole Normale Supérieure, IHES and MSRI for their hospitality while this work was being completed.

2. Riemann surfaces and quadratic differentials

2.1. The infinitesimal form of the Teichmüller metric

Let X be a hyperbolic Riemann surface with finitely generated fundamental group. Each end of X is either a finite volume cusp, or an infinite volume funnel.

The quasiconformal type of X can be encoded by a *pared surface* (S, P), i.e. a compact oriented surface S with specified *parabolic locus* $P \subset \partial S$ designating those components of the boundary corresponding to parabolic cusps. The interior of a pared surface carries a unique quasiconformal structure with infinite modulus ends at P and finite modulus ends at $\partial S - P$.

The *Teichmüller space* $\mathrm{Teich}(S, P)$ parameterizes Riemann surfaces X marked by quasiconformal maps $\phi: \mathrm{int}(S) \to X$ up to isotopy pointwise fixing the ideal boundary. We allow the shorthand $\mathrm{Teich}(S)$ when $P = \partial S$.

The Teichmüller metric is given by

$$d(X_1, X_2) = \inf \log K(\phi_1 \circ \phi_2^{-1})$$

where the infimum is over representative markings ϕ_i and K denotes maximal dilatation.

This metric can be described on an infinitesimal level as follows.

Let $Q(X)$ denote the Banach space of holomorphic quadratic differentials $\phi(z) dz^2$ on X satisfying

$$\|\phi\| = \int_X |\phi(z)| |dz|^2 < \infty .$$

$Q(X)$ is the cotangent space to Teichmüller space at X, and the norm above is the Teichmüller cometric.

There is a natural pairing between $Q(X)$ and the Banach space $M(X)$ of measurable Beltrami differentials $\mu(z) d\bar{z}/dz$ satisfying

$$\|\mu\| = \operatorname{ess\,sup} |\mu(z)| < \infty ,$$

namely

$$\langle \phi, \mu \rangle = \mathrm{Re} \int_X \phi(z) \mu(z) |dz|^2 .$$

Geometrically, $\langle \phi, \mu \rangle$ measures the extent to which the linefields of ϕ and μ are synchronized, averaged with respect to the area form $|\phi\mu|$.

The quotient $M(X)/Q(X)^{\perp}$ is the tangent space to Teichmüller space at X, and the quotient norm is the infinitesimal form of the Teichmüller metric.

This means a Beltrami differential μ gives an infinitesimal quasiconformal deformation of X whose length in the Teichmüller metric is

$$l(\mu) = \sup\{\langle \phi, \mu \rangle : \|\phi\| = 1\} .$$

There are many μ of length zero, giving infinitesimal maps of X to itself. Frequently $l(\mu) < \|\mu\|$, indicating that μ can be replaced by an equivalent map of lesser distortion.

See [Gar] for more details on the analytic foundations of Teichmüller theory.

2.2. *The unfolding map*

A covering of surfaces determines a canonical map between Teichmüller spaces. In this section we describe the contraction of the Teichmüller metric under canonical mappings related to the skinning map; for details and proofs, see [Mc1], Theorems 10.3 and 11.1.

Definitions. A Riemann surface X is of *finite type* (g, n) if it obtained from a compact surface of genus g by introducing n punctures. Similarly we say a pared surface (S, P) is of finite type if its entire boundary is declared parabolic; this characterizes (S, P) such that $\text{Teich}(S, P)$ is finite dimensional.

Let $T \subset S$ be a proper incompressible subsurface of a pared surface (S, P); this means $\pi_1(T)$ injects onto a proper subgroup of $\pi_1(S)$. Adjust T by isotopy so as many boundary components as possible lie in P, and label these Q.

The *unfolding map*

$$\theta: \text{Teich}(S, P) \to \text{Teich}(T, Q)$$

is defined by forming, for each $X \in \text{Teich}(S, P)$, the covering space $Y \to X$ determined by $\pi_1(T)$.

Since T is a *proper* subsurface, the covering is infinite-to-one and Y always has at least one infinite volume end. For example, if T is a disk then $Y \to X$ is the universal covering.

A quasiconformal map ϕ: $X_1 \to X_2$ lifts to a map $\tilde{\phi}$ of the same dilatation between the covering spaces Y_1, Y_2, so θ does not expand the Teichmüller metric. In fact, θ is contracting; that is, $\tilde{\phi}$ can be relaxed to a map of smaller dilatation. The amount of contraction is not uniform over $\text{Teich}(S, P)$; however, one can describe geometrically the location and directions which are contracted the least.

Any contraction obtained by unfolding can be directly exploited to obtain contraction of the skinning map. Our philosophy is to study the simpler operation of unfolding fairly completely before turning to the application to 3-manifolds.

Let $Y = \theta(X)$; then we have a covering map $Y \to X$. Any integrable quadratic differential on Y pushes forward to one on X; this defines an opertor

$$\Theta_{Y/X}: Q(Y) \to Q(X)$$

which is the coderivative $d\theta^*$ at X. By duality the operator norm of the derivative and coderivative agree. We focus our attention on the coderivative because we intend to exploit the geometry of quadratic differentials.

A holomorphic quadratic differential determines a *measure* $|\phi|$ and a *foliation* by the trajectories of vectors v such that $\phi(v, v) > 0$. The operator Θ is like push-forward of measures, except incoherence between foliations on different sheets of Y can cause cancellation. Thus $\|\Theta_{Y/X}\| \leqq 1$.

Theorem 2.1. *Fix the topological data* $(T, Q) \subsetneq (S, P)$. *Then* $\|\Theta_{Y/X}\| < c([X]) < 1$ *where* c *is a continuous function of the location of* X *in* moduli *space.*

Thus the unfolding map is uniformly contracting so long as X has no short geodesics (X is then confined to a compact subset of moduli space [Mum]). Now we describe what happens when short geodesics exist.

The thin, liftable and amenable parts

Fix any small $\varepsilon > 0$. The *thin part* $X_{\text{thin}} \subset X$ is the locus of points at which the injectivity radius of the hyperbolic metric is less than ε. For ε sufficiently small, every component of X_{thin} has cyclic fundamental group, and is either

- a neighborhood of a cusp or
- a long annulus with core curve a short geodesic.

Let $\phi : \text{int}(S) \to X$ be a marking. The *liftable part* X_{lift} of X is the union of those components of X_{thin} and of $X - X_{\text{thin}}$ which are isotopic into $\phi(T)$.

Since the covering $Y \to X$ is determined by $\pi_1(T)$, X_{lift} has a unique isomorphic lift to Y. In fact, the liftable part is largest subsurface of X which lifts to Y and has boundary in the set of injectivity radius ε. The *amenable part* Y_{am} of Y is the union of this lift and the total pre-image of X_{thin} (see Fig. 1).

Theorem 2.2. *If* $\|\Theta_{Y/X}(\phi)\|$ *is close to* 1, *then most of the mass of* $|\phi|$ *is in the amenable part of* Y.

More precisely, if unit norm ϕ_n on $Y_n \to X_n$ satisfy $\|\Theta_{Y_n/X_n}(\phi_n)\| \to 1$, then

$$\int_{(Y_n)_{\text{am}}} |\phi_n| \to 1$$

as $n \to \infty$.

Remark. A covering is *amenable* if there are large sets with small boundary in a graphic caricature of the covering space. For a general covering $\|\Theta\| < 1$ iff the covering is nonamenable [Mc1]. The covering Y determined by the proper subsurface T is nonamenable, but it can degenerate towards amenability as X tends to infinity in moduli space. Intuitively, Y_{am} is the part of Y which is converging to an amenable cover.

As a special case we record:

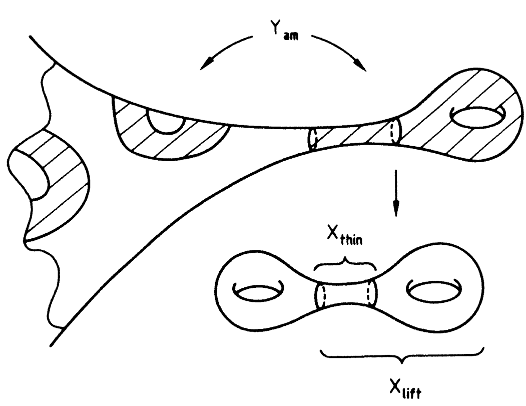

Fig. 1.

Corollary 2.3. *Suppose Y is a disk or a punctured disk. If $\|\Theta_{Y/X}(\phi)\|$ is near 1 then most of the mass of $|\phi|$ lies over the thin part of X.*

Proof. In the case of the universal covering, there is no liftable part. For a punctured disk, the liftable part is a neighborhood of a cusp, which is in the thin part already. □

In the acylindrical case, this corollary reduces the proof of uniform contraction for the skinning map to the study of short geodesics. Using geometric limits of quadratic differentials, we will find a second mechanism of contraction to handle the part of $|\phi|$ lying over X_{thin}. This mechanism applies in the cylindrical case as well, so it is really only the liftable part of X that can cause problems. The liftable part corresponds exactly to the cylindrical part of N, and we will deduce that any convergence difficulty comes from a nonperipheral torus.

2.3. Geometric limits

In this section we recall the existence of geometric limits of quadratic differentials viewed at the scale of the injectivity radius, developed in the Appendix to [Mc1]. Results on the distribution of the mass of a quadratic differential then follow easily by compactness.

Definitions. $\mathbf{P}\mathcal{Q}$ denotes the space of triples $(X, v, [\phi])$ where

- X is a Riemann surface with a complete metric of constant curvature $\kappa \in [-1, 1]$,
- v is a baseframe at a point of X with injectivity radius $\geqq 1$, and
- $[\phi]$ is a nonzero holomorphic quadratic differential determined up to multiplication by a complex unit.

$\mathbf{P}\mathcal{Q}_{g,n}$ denotes the space of integrable ϕ which live on surfaces of finite type (g, n) (genus g with n punctures).

Theorem 2.4. *$\mathbf{P}\mathcal{Q}_{g,n}$ has compact closure in $\mathbf{P}\mathcal{Q}$.*

This means any sequence in $\mathbf{P}\mathcal{Q}_{g,n}$ has a subsequence $(X_n, v_n, [\phi_n])$ converging to $(X, v, [\phi])$ in the *geometric topology*: any compact subset of X admits a baseframe preserving near-isometry into X_n for n large, and there exists constants c_n such that $c_n \phi_n$ converges to ϕ uniformly on compacta (see the Appendix to [Mc1].)

One consequence of compactness is that the measure $|\phi|$ cannot be too wildly distributed on a Riemann surface of fixed type.

Proposition 2.5. *Let $\phi \in Q(X)$ where X is of finite type (g, n).*

Let $B(x, r)$ be an embedded ball in the hyperbolic metric, and let $s > r$. Then the ratio:

$$\frac{\int_{B(x,s)} |\phi|}{\int_{B(x,r)} |\phi|} < C(g, n, s/r) ;$$

i.e. the $|\phi|$-mass of a somewhat larger ball is not too much bigger.

Proof. If not, there is a sequence of balls on varying Riemann surfaces such that the ratio tends to infinity. Scaling the metric so r remains bounded below, we may obtain a geometric limit for which the ratio is infinite (and $\phi \neq 0$), which is absurd. □

Definitions. The *cuspidal thin part* $X_{\mathrm{cusp}}(\varepsilon)$ is the union of the cuspidal components of the ε-thin part of X. Similarly, the *geodesic thin part* $X_{\mathrm{geod}}(\varepsilon)$ is made up of the components associated to short geodesics.

Corollary 2.6. *Suppose* $\|\phi\| = 1$. *Then*

$$\int_{X_{\mathrm{cusp}}(\varepsilon)} |\phi| < \delta(g, n, \varepsilon) ,$$

where $\delta \to 0$ *as* $\varepsilon \to 0$. *In other words, most of the mass of* $|\phi|$ *is outside of the cuspidal thin part.*

Proof. Given any puncture p of X one may easily construct a regular covering space Z of small degree (in fact 4 will do) ramified of order 2 over the puncture. Since ϕ has at worst a simple pole at the puncture, its pull-back to Z is holomorphic across the pre-image of the puncture and defines an integrable differential on the surface $\bar{Z}$ with the pre-image filled in. Now the corollary follows from the result above (and straightforward comparisons of the hyperbolic metric before and after filling in.) □

On the other hand, when there are short geodesics, a definite fraction of the mass of $|\phi|$ lies in the geodesic thin part. This fact will be used in the study of cylindrical manifolds.

Proposition 2.7. *Let* $\phi \in Q(X)$, $\|\phi\| = 1$, *and suppose* X *has a geodesic of length less than* $\varepsilon/2$. *Then*

$$\int_{X_{\mathrm{geod}}(\varepsilon)} |\phi| > \delta(g, n, \varepsilon) > 0 .$$

Proof. If not, we can find a sequence of data (X_n, ϕ_n) for which the integral above tends to zero. Choose basepoints x_n in a component of the $\varepsilon/2$-thick part of X_n whose ϕ_n-mass is bounded below, independent of n; this is possible because the number of components of the thick part is bounded in terms of (g, n), and the mass cannot be concentrated into a cusp, as the preceding proof shows. Now extract a geometric limit (X, x, ϕ), $\phi \neq 0$. It is possible that a geodesic of length less than $\varepsilon/2$ pinches off, but in any case the integral of ϕ over an open component of the ε-thin part of X is zero, a contradiction. □

2.4. Analysis

Finally we record some general principles for estimating the length of a quasiconformal deformation.

Let X be of finite type, and let $\mu \in M(X)$ be an infinitesimal quasiconformal deformation. Since the Teichmüller length

$$l(\mu) = \sup_{\{\phi \in Q(X) : \|\phi\| = 1\}} \operatorname{Re} \int_X \phi\mu$$

is expressed as an integral over X, we can analyze the length by breaking the integral up into pieces.

For $E \subset X$ we introduce the notation:

$$\langle \phi, \mu \rangle_E = \operatorname{Re} \int_E \phi\mu, \quad \|\phi\|_E = \int_E |\phi|, \quad \|\mu\|_E = \operatorname{ess\,sup}_{z \in E} |\mu(z)| .$$

Suppose $\|\phi\| = \|\mu\| = 1$. The *efficiency* of pairing on E is the quotient

$$\langle \phi, \mu \rangle_E / \|\phi\|_E .$$

The efficiency is bounded above by 1. If $\{E_i\}$ is any partition of X into measurable sets, then

$$\langle \phi, \mu \rangle = \sum_i \|\phi\|_E \left(\frac{\langle \phi, \mu \rangle_E}{\|\phi\|_E} \right)$$

expresses the global pairing as a convex combination of efficiencies. If the efficiency is less than $1 - \alpha$ on each E_i in a subcollection whose total $|\phi|$-mass is m, then we obtain a bound of $1 - m\alpha$ for $\langle \phi, \mu \rangle$.

It is sometimes more convenient to use a *covering* by hyperbolic balls, rather than a partition. To justify an estimate in this case, we must pass to a subset of disjoint balls.

Proposition 2.8. *Let $E \subset X$ be a measurable subset of a Riemann surface of type (g, n) and let $\phi \in Q(X)$. Suppose E is covered by a family $\{B_\alpha\}$ of embedded hyperbolic balls. Then there is a subfamily of disjoint balls $\{B_i\}$ such that*

$$\int_{\sqcup B_i} |\phi| \geqq c(g, n) \int_E |\phi|$$

where the constant $c(g, n) > 0$ depends only on the type of X.

Proof. The radius of any B_α is bounded above because the area of X is finite. Let B_0 be a ball whose radius is at least 1/2 as large as any other element of $\{B_\alpha\}$, and choose B_{i+1} inductively so it is at least half as large as any other element of the family which is disjoint from the balls chosen so far.

If we increase the radius of each element of $\{B_i\}$ by a factor of 10, we obtain a cover of E. In doing so, we increase the $|\phi|$-mass of each B_i by at most a factor $C(g, n, 10)$ (Proposition 2.5); so the theorem holds with $c(g, n) = 1/C(10, g, n)$. □

Corollary 2.9. *Let $E \subset X$ be a measurable set with $|\phi|$-mass m, and suppose each point of E is contained in an embedded ball B on which the efficiency*

$$\langle \phi, \mu \rangle_B / \|\phi\|_B < 1 - \alpha .$$

Then the global pairing $\langle \phi, \mu \rangle$ is bounded by $1 - c(g, n) \cdot m\alpha$.

This is immediate.

3. 3-Manifolds

In this section we give a concise description of the skinning map and the gluing problem for hyperbolic 3-manifolds. A more detailed account can be found in [Mor], which also treats the Haken decomposition, pared manifolds, Andreev's theorem, and the overall logic of the proof of the uniformization theorem. Thurston's own papers on the subject [Th1; Th3-Th5], are beginning to appear.

3.1. Kleinian manifolds

Definitions. Let Γ be a *Kleinian group*, that is a discrete subgroup of isometries of hyperbolic 3-space $\mathbf{H}^3$. The action of Γ on hyperbolic space extends to the sphere at infinity which we identify with the Riemann sphere $\hat{\mathbf{C}} = \mathbf{C} \cup \{\infty\}$. $\hat{\mathbf{C}}$ is partitioned into a *domain of discontinuity* Ω and a *limit set* Λ.

Assume Γ is torsion-free and orientation-preserving. The *Kleinian manifold*

$$N = (\mathbf{H}^3 \cup \Omega)/\Gamma$$

is a 3-manifold with a hyperbolic structure on its interior and a complex structure on its boundary. N is oriented by the choice of a fixed orientation for $\mathbf{H}^3$.

Remark. In many ways it is natural to allow finite order and orientation-reversing transformations in Γ. We have chosen not to, for two reasons.

First, a hyperbolic orbifold with finitely generated fundamental group becomes an orientable manifold upon passing to a finite regular cover. When studying the gluing problem, it is enough to find a solution on a finite cover: if it exists, the solution descends by uniqueness.

Secondly, we prefer ∂N to carry a *complex* structure so holomorphic quadratic differentials on ∂N make sense and the Teichmüller metric discussion carries over.

In the sequel we will always assume:

(*) N is *geometrically finite* with *incompressible*, *quasifuchsian boundary.*

Geometrically finite means a unit neighborhood of the convex core of $\mathrm{int}(N)$ has finite volume. It follows that $\Gamma \cong \pi_1(N)$ is finitely generated, and ∂N is a finite union of Riemann surfaces of finite type.

Incompressible boundary means $\pi_1(X)$ injects into $\pi_1(N)$ for each component X of ∂N. Equivalently, the limit set Λ is connected.

Finally, the boundary of N is *quasifuchsian* if the subgroup $\pi_1(X) \subset \pi_1(N)$ is a quasifuchsian group for each component of ∂N.

It is convenient to encode the quasi-isometry type of N by topological data. Let M be a compact oriented 3-manifold with boundary, and designate a *parabolic locus* $P \subset \partial M$ by a compact subsurface satisfying:

- P is a union of incompressible annuli and tori;
- every torus boundary component of ∂M is in P;
- any cylinder in M whose boundary rests on essential curves in P is homotopic rel boundary into P.

The pair (M, P) is a *pared manifold*. For the general definition of a pared manifold, see [Th1]; the conditions above are simplified by the fact that M is orientable and 3-dimensional.

Let $GF(M, P)$ denote the space of realizations of (M, P). $GF(M, P)$ consists of all pairs $(N, [\phi])$ where N is a geometrically finite Kleinian manifold and the marking

$$\phi:(M, \partial M - P) \to (N, \partial N)$$

is an orientation-preserving homotopy equivalence of pairs.

We require that $\phi_*(\pi_1(P))$ is parabolic and represents every parabolic conjugacy class in $\pi_1(N)$. *There are no accidental parabolics*: all cusps of N must be accounted for by P. Homotopic markings are considered equivalent.

The boundary of $N \in GF(M, P)$ is marked by

$$\partial_0 M := \partial M - \text{int}(P) .$$

Recording the conformal structure on ∂N so marked determines a map

$$GF(M, P) \to \text{Teich}(\partial_0 M) .$$

In fact this map is an isomorphism (given our assumption that at least one realization exists); see, e.g., [Kra]. Briefly, the idea is that deformations of a realization N correspond to Γ-invariant measurable complex structures (Beltrami differentials) on $\hat{\mathbb{C}}$, which are in turn parameterized by complex structures on $\partial N = \Omega/\Gamma$ since the limit set has measure zero. (For more general conditions under which this identification can be made, see [Sul].)

One may think of (M, P) as a picture of N truncated at its cusps. A rank 2 cusp gives a finite volume end of N which we cut off by a torus. A rank 1 cusp gives a pair of punctures on ∂N whose peripheral loops are homotopic in N; here we cut along a peripheral cylinder which implements the homotopy. The result is a compact manifold homeomorphic to M; P consists of the cylinders and tori along which we cut.

The condition that ∂N is incompressible translates into the condition that $\partial_0 M$ is incompressible in M.

We assume boundary curves homotopic through M to P are already homotopic to P through ∂M. Then any realization N has quasifuchsian boundary.

We assume $GF(M, P)$ is non-empty. (In the application to the uniformization of 3-manifolds, this is proved inductively).

Any two realizations of (M, P) are quasi-isometric. This may be proved by promoting ϕ to a homeomorphism, quasiconformal map and finally to a real-analytic quasi-isometry (see Hempel [Hem], Marden [Mar], and Douady and Earle [DE].) Thus the quasi-isometry type of (M, P) is specified by its topology.

Remarks.

1. If (S, P) is a compact but disconnected pared surface, we define

$$\text{Teich}(S, P) = \prod \text{Teich}(S_i, P_i) ;$$

the product is over components of S and the metric is the supremum of the Teichmüller metrics on each factor. Adopting a similar convention for $GF(M, P)$, the isomorphism

$$GF(M, P) \cong \mathrm{Teich}(\partial_0 M)$$

holds even if M is disconnected.

2. Since σ can be defined component by component, *until* §7 *we assume* M *is connected.* For the general gluing problem, we will allow M to be disconnected but require M/τ to be connected.

3. There are many interesting manifolds, such a knot complements, which can only be realized as hyperbolic manifolds with cusps. On the other hand, our discussion of the skinning map is insensitive to the parabolic locus, by virtue of Corollary 2.6 which says most of the mass of a quadratic differential is outside of the cuspidal thin part. Thus to follow the argument, little is lost by imagining P is empty and the convex core of N is compact.

3.2. *Fuchsian and quasifuchsian groups*

Let S be a compact oriented surface and let $\bar{S}$ denote S with the orientation reversed. To any Riemann surface X there corresponds a complex conjugate Riemann surface $\bar{X}$, obtained by post-composing every chart $U \to \mathbf{C}$ by complex conjugation. Similarly, there is a natural antiholomorphic map

$$\rho : \mathrm{Teich}(\bar{S}) \cong \mathrm{Teich}(S)$$

which sends $\bar{X}$ to X and remarks it with S; this preserves the requirement that the marking be quasiconformal (rather than antiquasiconformal).

Let S be a compact surface, $(M, P) = (S \times I, \partial S \times I)$. Then

$$GF(M, P) \cong \mathrm{Teich}(S \sqcup \bar{S}) = \mathrm{Teich}(S) \times \mathrm{Teich}(\bar{S}) ,$$

parameterizes the space of *quasifuchsian groups* modeled on S. Here $\bar{S}$ appears because the homotopy of one boundary components to the other through M reverses orientation.

To see there is at least one realization, take a Fuchsian group uniformizing S. The set of all Fuchsian groups in $GF(N)$ is the image of $\mathrm{Teich}(\bar{S})$ under (ρ, id); this is *not* a complex submanifold.

3.3. *The skinning map*

To define the *skinning map*

$$\sigma : GF(M, P) \cong \mathrm{Teich}(\partial_0 M) \to \mathrm{Teich}(\overline{\partial_0 M}) ,$$

let X be a point in Teich $(\partial_0 M)$. From X, construct the corresponding realization N in $GF(M, P)$. Each component of $\partial_0 M$ determines a quasifuchsian covering space of N; since these are parameterized by conformal structures on their boundaries, X determines a point

$$(X, \sigma(X)) \in \mathrm{Teich}(\partial_0 M) \times \mathrm{Teich}(\overline{\partial_0 M}) .$$

Projection onto the second factor is the skinning map.

Geometrically, the surface $\sigma(X)$ forms the conformal boundary attached to the ends of the quasifuchsian covering spaces which were *buried* in N. Unlike X, whose shape is determined by just the asymptotic geometry of N, the conformal structure of $\sigma(X)$ depends on the geometry of N as a whole.

The skinning map is holomorphic, so it does not expand the Teichmüller metric. We will see that σ is actually contracting, unless M is an interval bundle over a surface. Our goal is to use that contraction to produce fixed-points for related maps.

For example, define

$$\rho : \mathrm{Teich}(\overline{\partial_0 M}) \to \mathrm{Teich}(\partial_0 M)$$

by reflection as in §3.2. A fixed point for $\rho \circ \sigma$ is a geometric structure on M for which every boundary component is Fuchsian.

More generally, an orientation-reversing fixed-point free involution

$$\tau : \partial_0 M \to \partial_0 M$$

determines an isometry (which we denote by the same letter):

$$\tau : \mathrm{Teich}(\overline{\partial_0 M}) \to \mathrm{Teich}(\partial_0 M) ;$$

finding a fixed point for $\tau \circ \sigma$ is called the *gluing problem*. When a fixed point exists, it gives a geometric structure for which the ends of N identified by τ can be isometrically sewn together. The result is a geometrically finite realization of M/τ.

4. Acylindrical manifolds with geodesic boundary

Definition. (M, P) is *acylindrical* if every cylinder

$$C : (\mathbf{S}^1 \times I, \mathbf{S}^1 \times \partial I) \to (M, \partial_0 M)$$

resting on essential curves in $\partial_0 M$ is homotopic (as a map of pairs) into ∂M. Otherwise C is a *compressing cylinder.*

In this section the spirit of our approach to the skinning map will be conveyed by a simple test case. We will show that a geometrically finite acylindrical hyperbolic manifold without cusps can be deformed so its convex hull has totally geodesic boundary. Equivalently, there is a quasiconformally conjugate Kleinian group whose domain of discontinuity is a union of circular disks.

This amounts to showing the existence of a fixed point for $\rho \circ \sigma$ when M is acylindrical and $P = \varnothing$. The restriction on P is for simplicity only.

4.1. *Leopard spots*

The shorthand $GF(M)$ denotes $GF(M, \varnothing)$.

Let $N \in GF(M)$. To picture the derivative of the skinning map, consider a Beltrami differential $\mu(z)d\bar{z}/dz$ in $M(\partial N)$. This differential is transported to $\sigma(\partial N)$ as follows.

First, lift μ to the Riemann sphere $\hat{\mathbb{C}}$. For each component X of ∂N select a component $\Omega(X) \subset \Omega$ which covers X. The differential μ is invariant by $\pi_1(N)$, so in particular by the stabilizer $\pi_1(X)$ of $\Omega(X)$. The domain of discontinuity of $\pi_1(X)$ is $\Omega(X) \sqcup \Omega(Y)$, where the latter domain uniformizes a component Y of $\sigma(\partial N)$. Under the action of $\pi_1(X)$, μ descends to a Beltrami differential $d\sigma(\mu)|Y$ which is the image of μ under the derivative of the skinning map.

Here is another description of the derivative from the point of view of Y itself. The domain Ω descends to an open dense subset of Y. By acylindricity, the stabilizers in $\pi_1(N)$ of any two components of Ω intersect in the trivial group. Thus each component maps injectively to Y, and the image is a countable collection of open disks, reminiscent of leopard spots (Fig. 2).

Each component of Ω uniformizes some part of ∂N. Thus, to each spot U on Y there is associated a group of conformal automorphisms G such that $U/G = X'$ for some $X' \subset \partial N$. The differential $d\sigma(\mu)|U$ is obtained by lifting μ from X'. Since N is geometrically finite, its limit set has measure zero and μ is determined by its values on the spots.

The quantity

$$\|d\sigma\| = \sup_{\|\phi\| = 1,\ \|\mu\| = 1} \langle \phi, d\sigma(\mu) \rangle$$

measures the infinitesimal contraction of the skinning map; the supremum is over unit norm $\phi \in Q(\sigma(\partial N))$ and $\mu \in M(\partial N)$.

4.2. *Quasidisks*

The iteration $\rho \circ \sigma$ is particularly tractable because the shape of the spots can be controlled.

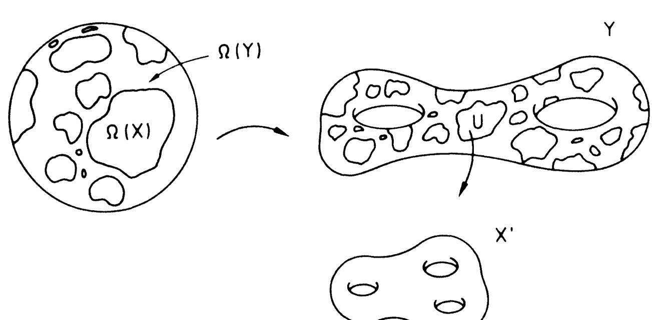

Fig. 2. The skinning map

Definitions. A K-quasidisk $D \subset \mathbf{C}$, $\hat{\mathbf{C}}$ or $\mathbf{H}^2$ is the image of a round disk under a K-quasiconformal map of the Riemann sphere to itself.

If $D \subset X$ is a disk on a Riemann surface, we say D is a K-quasidisk if its lift to the universal cover of X is a K-quasidisk as defined above.

A metrized disk D has *K'-bounded turning* if the ratio

$$\operatorname{diam}(J)/\operatorname{dist}(x, y) \leqq K'$$

for any two points $x, y \in \partial D$, where J is the smaller component of $\partial D - \{x, y\}$.

A disk on the sphere is a quasidisk iff it has bounded turning in the spherical metric. Moreover, K' can be bounded in terms of K and vice versa (see Lehto and Virtanen [LV, II §8] for background).

This property fails already for round balls in the hyperbolic plane: the turning tends to infinity as the diameter grows. However, one may check:

Proposition 4.1. *A K-quasidisk $D \subset \mathbf{H}^2$ whose hyperbolic area is bounded by A has controlled geometry: in the hypebolic metric, both its turning and its diameter are bounded in terms of K and A.*

Remark. If $D \subset E$ are a pair of K-quasidisks on the sphere, then as a disk on the Riemann surface E, D is at worst a K^3 quasidisk. (There is a K-quasiconformal map of the sphere which sends a round ball to D; and the Riemann map from E to $\mathbf{H}^2$ extends to a K^2-quasiconformal map of the sphere.)

4.3. Tame leopards

We will see that under iteration of $\rho \circ \sigma$, the leopard spots remain uniformly quasicircular, so the amount of contraction of σ is controlled by the following result:

Proposition 4.2. *Let X and Y be hyperbolic Riemann surfaces of finite type (g, n) and (g', n').*

Let $\bar{U} \subset Y$ be a K-quasidisk such that $X = U/G$ for a group G of conformal automorphisms of U. Let $\phi \in Q(Y)$ and let $\mu \in M(Y)$ be a unit-norm G-invariant Beltrami differential supported on U.

Then the efficiency of pairing on U satisfies:

$$\frac{\langle \phi, \mu \rangle_U}{\|\phi\|_U} < c(K, g, n, g', n') < 1$$

where c depends only on the amount of quasiconformality and the type of X and Y.

Proof of Proposition 4.2. Consider any sequence of data (U, G, μ, Y, ϕ) satisfying the hypotheses of the theorem: the quasidisk constant K and the genus and number of punctures of Y and U/G are fixed. Suppose the efficiency tends to 1 for this sequence; we will take a geometric limit and deduce a contradiction.

Since the area of Y bounds that of $\bar{U}$, by Proposition 4.1 the turning of $\bar{U}$ is uniformly bounded. Choose a ball $B(x, r) \subset \bar{U}$ of maximal radius in the Poincaré metric on Y. By bounded turning, the diameter of $\bar{U}$ is comparable to τ. Replacing the hyperbolic metric on Y by a flatter scalar multiple if necessary, we may assume that the radius r of B is constant throughout the sequence. In the new metric the injectivity radius at x is uniformly bounded below, so we may pass to a subsequence and obtain a geometric limit of the pointed surfaces (Y, x).

By compactness of $\mathbf{P}\mathcal{Q}_{g,n}$, we may rescale the differentials ϕ so they converge (to a nonzero limit) as well. Note that scaling ϕ by a positive constant does not alter the efficiency of ϕ on U.

Next we extract a geometric limit of the K-quasidisks $\bar{U}$. Any Hausdorff limit contains a ball of radius r, so U cannot degenerate to a point. Nor can U expand to an unbounded disk, since its diameter is comparable to r. And since the turning of $\bar{U}$ is bounded, the limit is still a disk.

The image of x on $X = U/G$ lies on a non-trivial closed loop of uniformly bounded Poincaré length, since the injectivity radius of X is bounded above. It follows that any geometric limit of G contains at least an infinite cyclic group acting conformally and freely on U.

Finally, we extract a weak limit (in the L^∞-norm) of the differentials μ. The limiting differential is G-invariant and satisfies $\|\mu\| \leqq 1$. Since the quadratic differentials ϕ converge uniformly on U, the efficiency on U is continuous under geometric limit.

Suppose the limiting efficiency is 1; then $\mu = \bar{\phi}/|\phi|$. Since μ is G-invariant, so is $\phi|U$. Now ϕ is holomorphic on a neighborhood of $\bar{U}$, so $\int_U |\phi| < \infty$. But $\phi|U$ cannot be simultaneously nonzero, integrable and invariant under the infinite group G. □

Corollary 4.3. *If M is acylindrical (and realizable), there is a unique N in $GF(M)$ whose convex hull has totally geodesic boundary.*

Proof of the corollary. Let N_0 be a realization of M.

Since σ is nonexpanding and ρ is an isometry, the Teichmüller distance $d(\partial N_i, \partial N_{i+1}) \leqq d(\partial N_0, \partial N_1)$, where $N_i = (\rho \circ \sigma)^i(N)$. This quantity measures the Teichmüller distance between the inner and outer surfaces of each quasifuchsian boundary group; thus each component of the domain of discontinuity of N_i is the K-quasiconformal image of a round disk, where K is independent of i. The same holds true for all N along the path

$$P = \bigcup_0^\infty (\rho \circ \sigma)^i(p) ,$$

where p is a smooth path joining N_0 to N_1. To prove the corollary, it suffices to show $\|d\sigma_N\| < c < 1$ for all $N \in P$, since then P has finite length and N_i converges at a geometric rate to a fixed point for $\rho \circ \sigma$.

Let $N \in P$. By the Remark of §4.2, each spot U on $\sigma(\partial N)$ is at worst a K^3-quasidisk (on the sphere, U lifts to a K-quasidisk inside another one which uniformizes a component of $\sigma(\partial N)$). Proposition 4.2 provides a bound uniformly less than one on the efficiency of pairing between ϕ and μ on every spot U. But $\|d\sigma_N\|$ is a convex combination of these quantities (cf. §2.4), so $\|d\sigma_N\| < c < 1$ for a fixed c (determined by K and the topology of ∂M). □

5. Dependence on moduli

In this section we show the amount of contraction of the skinning map depends only on the location of ∂N in moduli space. As a preliminary, we discuss the analogue of leopard spots in the cylindrical case.

5.1. Stripes

Here is a more precise description of the mottling which appears on a general skinned surface; in the cylindrical case, not only spots but stripes may appear.

Notation:
$N \in GF(M, P)$,
$\Omega = \Omega(\pi_1(N))$ is its domain of discontinuity,
$X \subset \partial N$ is uniformized by a component
$\Omega(X)$ of Ω with stabilizer $\pi_1(X)$, and
$\Omega(\pi_1(X))$ is the domain of discontinuity of $\pi_1(X)$;
from this data form the two manifolds:
$CX = (\mathbf{H}^3 \cup \Omega(\pi_1(N))/\pi_1(X)$, and
$QX = (\mathbf{H}^3 \cup \Omega(\pi_1(X))/\pi_1(X)$.

$CX \subsetneq QX$ and the interiors of CX and QX are isomorphic. If $CX = QX$ then $\pi_1(N)$ is at most a degree 2 extension of $\pi_1(X)$ and M is an interval bundle over a surface.

Construct CN and QN by taking the disjoint union of CX and QX over all components of ∂N.

$$\partial QN = \partial N \cup \sigma(\partial N),$$

indeed QN is the quasifuchsian manifold used to define the skinning map. On the other hand

$$\partial CN = \partial N \cup BN,$$

where

$$BN = \cup(\Omega - \Omega(X))/\pi_1(X).$$

is a union of subsurfaces of $\sigma(\partial N)$.

Suppose M is acylindrical. If $P = \emptyset$, we have seen that BN is just a collection of disks (the leopard spots). Otherwise BN may also contain punctured disks coming from peripheral cylinders in M. (The limit set Λ is a Sierpiński curve in first case—no components of Ω share a boundary point. In the second case components may touch at one point.)

Now we describe the general case. Since ∂N is incompressible, BN consists of incompressible subsurfaces of $\sigma(\partial N)$. Every simple curve on BN is homotopic through CN to a unique curve on ∂N. Unless the curve encloses a puncture, the trace of the homotopy descends to give an compressing cylinder in N. Conversely, any compressing cylinder gives an essential, non-peripheral curve on BN.

We claim there is a compact incompressible possibly disconnected pared surface $(T, Q) \subset \partial_0 M$ which accounts for all components of BN which are not disks. More precisely, suppose U is a non-simply connected component of BN covering a component X of ∂N. Then there is:

1. a component $(S, \partial S)$ of $\partial_0 M$ marking X, and
2. a unique component (T', Q') of (T, Q) such that
3. T' is a subsurface of S (with parabolic locus Q' = the maximal boundary isotopic into ∂S) and
4. $U \to X$ is the covering determined by $T \subsetneq S$.

This sort of assertion can be approached via the theory of the characteristic submanifold (cf. Jaco-Shalen [JS], Thurston [Th4]). Here is a geometric approach, adapted from Thurston's course. Whenever the boundaries of two components of Ω meet, form the convex hull of their intersection with respect to the Poincaré metric on each component. The result is a collection of disjoint convex sets, invariant under $\pi_1(N)$, which descend to an embedded convex subsurface $Y \subset \partial N$. Y is almost a geometric realization for T—one need only include, in addition, peripheral annuli whenever two components touch in just one point.

Proposition 5.1. *All but a finite number of components of BN are disks.*

Proof. BN is an incompressible subsurface of $\sigma(\partial N)$, so the number of components with negative Euler characteristic is bounded. The remaining components are disks, punctured disks and annuli. Only a finite number of annuli appear in T; thus the modulus of any annular component is bounded below and their total number is finite. Finally the number of punctured disks is bounded by the number of punctures of ∂N. □

Proposition 5.2. *Either M is an interval bundle over a surface, or every component*

$$(T', P') \subset (S, \partial S) \subset \partial_0 M$$

is a proper subsurface of a component S of $\partial_0 M$.

Proof. If $T' = S$, then, $CX = QX$ for the component X of ∂N marked by S and N has a quasifuchsian covering space of degree at most 2. □

5.2. Contraction from unfolding

Assume M is not an interval bundle over a surface; here is one way to see skinning map is actually contracting.

Let

$$\phi : \partial N_1 \to \partial N_2$$

be an extremal quasiconformal map between two points in $\mathrm{Teich}(\partial N)$. Then there is a quasiconformal map

$$\psi : \sigma(\partial N_1) \to \sigma(\partial N_2)$$

obtained by lifting ϕ to a map between corresponding components of BN_1 and BN_2 and completing by continuity. The dilatation of ψ is the same as that of ϕ, but—because BN is an infinite-sheeted covering space of N—ψ fails to be a Teichmüller mapping. (The dilatation of any Teichmüller mapping, such as ϕ, is dual to a finite norm integrable quadratic differential; and the lift of this differential to BN has infinite norm.)

Thus the optimal map from $\sigma(N_2)$ to $\sigma(N_2)$ has less distortion than ψ, and therefore the skinning map shrinks the Teichmüller metric.

We would like to know how much the dilatation of ψ improves when it is pulled taut. Since it seems difficult to take advantage of the way various spots fit together, a simplified approach is to pull ψ as taut as possible without changing its values on the boundary of each component of BN.

This reduces to the problem of relaxing the lift of a Teichmüller map to a large covering surface. As we have seen in §2.2, for coverings determined by a proper subsurface such relaxation is always possible.

Let us formulate this contraction on the level of the cotangent space to Teichmüller space. The coderivative

$$d\sigma^* : Q(\sigma(\partial N)) \rightarrow Q(\partial N)$$

is given by

$$d\sigma^*(\phi) = \sum_U \Theta_{U/X}(\phi \,|\, U)$$

where the sum is over all components $U \subset BN \subset \sigma(\partial N)$, each of which covers a component $X \subset \partial N$. Therefore

$$\|d\sigma^*\| \leqq \sup_{U/X} \|\Theta_{U/X}\| \ .$$

Theorem 5.3. *Assuming (M, P) is not an interval bundle over a surface,*

$$\|d\sigma_N\| < c([\partial N]) < 1 \ ,$$

where c is a continuous function of the location of ∂N in moduli space.

Proof. Each $U \rightarrow X \subset \partial N$ is either a universal covering or a covering determined by a *proper* subsurface $(T', P') \subset (S, \partial S) \subset \partial_0(M)$. By Theorem 2.1, for fixed (T', P') there is a bound for $\|\Theta_{U/X}\|$ depending only on $[X]$. (T, P) has only a finite number of components, so $\|\Theta_{U/X}\| < c([\partial N]) < 1$ and this gives a bound for $\|d\sigma\| = \|d\sigma^*\|$. □

Remark. When M is disconnected, the contraction from one component eventually propagates to the rest; then the same result holds with σ replaced by an appropriate iterate of $\tau \circ \sigma$, as stated in Theorem 1.2.

Refinements. The *amenable* part BN_{am} of BN is the union of the amenable parts of $U \rightarrow X$ for each component U of BN (see §2.2).

A similar application of Theorem 2.2 yields:

Proposition 5.4. *Let $\phi \in Q(\sigma(\partial N))$, $\|\phi\| = 1$. If $\|d\sigma^*(\phi)\|$ is close to 1, then most of the mass of $|\phi|$ lies in the amenable part of BN.*

Proposition 5.5. *If M is acylindrical, $\|\phi\| = 1$ and $\|d\sigma^*(\phi)\|$ is close to 1, then most of the mass of $|\phi|$ lies over the thin part of N.*

Proof. In this case BN consists of disks and punctured disks; apply Corollary 2.3. □

This reduces the analysis of the acylindrical case to properties of short geodesics, which we turn to next.

6. Uniform contraction for acylindrical manifolds

This section establishes the second mechanism of contraction, and applies it to the acylindrical case.

6.1. Inefficiency over the thin part

Definitions. Let $N \in GF(M, P)$ be a realization of an acylindrical pared manifold.

Recall that the geodesic thin part $\partial N_{\mathbf{geod}}(\varepsilon)$ consists of the components of the ε-thin part associated to short geodesics, rather than cusps.

Let $\mu \in M(\partial N)$ and $\phi \in Q(\sigma(\partial N))$ satisfy $\|\mu\| = \|\phi\| = 1$.

We say $x \in \sigma(\partial N)$ *lies over* the geodesic thin part if

$$x \in U \subset BN \subset \sigma(\partial N) ,$$

U covers a component $X \subset \partial N$, and under this covering x maps to the geodesic thin part.

Theorem 6.1. *For ε sufficiently small*:

If $x \in \sigma(\partial N)$ lies over the geodesic thin part $\partial N_{\mathbf{geod}}(\varepsilon)$, then there exists an embedded hyperbolic ball $B \subset \sigma(\partial N)$ centered at x such that the efficiency

$$\frac{\langle \phi, d\sigma(\mu) \rangle_B}{\|\phi\|_B} \leqq c(\partial_0 M) < 1 .$$

Corollary 6.2. *The skinning map for an acylindrical pared manifold is uniformly contracting.*

Proof. By definition of the Teichmüller metric,

$$\|d\sigma_N\| = \sup \langle \phi, d\sigma(\mu) \rangle = \sup \langle d\sigma^*(\phi), \mu \rangle$$

where $\|\phi\| = \|\mu\| = 1$. Let $E \subset BN \subset \sigma(\partial N)$ denote the pre-image of the geodesic thin part of ∂N.

If ϕ and μ pair very efficiently, then $\|d\sigma^*(\phi)\|$ is near 1 and therefore most of the mass of $|\phi|$ lies over the thin part of ∂N (Proposition 5.5). By Corollary 2.6 most of the mass of $|d\sigma^*(\phi)|$ lies outside the cusps, so $|\phi|$ is concentrated over the *geodesic* thin part. Thus it suffices to bound the pairing when the $|\phi|$-mass of E is near 1.

By Proposition 6.1 above, we may cover E with balls on which the pairing is inefficient by a definite amount, independent of N. Then a uniform bound on the global pairing follows by Corollary 2.9: we can extract disjoint inefficient balls whose $|\phi|$-mass is still substantial. □

Corollary 6.3. *If (M, P) is acylindrical, any gluing problem has a solution.*

Proof. The composition $\tau \circ \sigma$ is a uniformly contracting map on a complete metric space. □

Remark. The same result for disconnected (but compact) acylindrical pared manifolds follows immediately.

6.2. Problematic example

It is easy to imagine something going wrong over the thin part.

Here we will sketch an arrangement of disks on a Riemann surface X, each covering a fixed surface Y with a short geodesic, such that the lift of a Teichmüller mapping on Y is nearly a Teichmüller mapping on X. Then we indicate how this scenario will is ruled out for Kleinian groups. The proof of Theorem 6.1 simply involves making this argument precise.

As a hyperbolic surface, the quotient of the strip $\Sigma = \{z : 0 < \mathrm{Im}(z) < 1\}$ by a small translation $T(z) = z + \varepsilon$ is an annulus Y_ε with a short geodesic. Think of Y_ε as a caricature of a compact surface with a short geodesic.

The extremal quasiconformal map between two annuli Y_ε and Y_δ is an affine stretch. When lifted to Σ, the stretch is simply a map which preserves vertical distances and stretches horizontal lengths by a factor of δ/ε.

Fix a Teichmüller mapping $f: X \to X'$, with associated quadratic differential ϕ. In the $|\phi|$ metric, we can think of f as preserving the lengths of the vertical trajectories of ϕ and stretching the horizontal trajectories by a fixed factor.

Now cover X by a full-measure set of disks U, each of which is a ϕ-rectangle which is long in the horizontal direction (Fig. 3). We can map Σ to U so that the horizontal trajectories of ϕ approximately correspond to horizontal lines on Σ. Then an affine stretch on Y lifts to a very good approximation of a Teichmüller mapping on X.

The example can easily be modified so Y is compact. Clearly such examples must be avoided if we are to establish uniform contraction for the skinning map.

Here is why they do not occur. Suppose $\gamma \in \pi_1(N)$ corresponds to a short geodesic on a component Y of ∂N. The fixed points of γ, when transported to $\sigma(\partial N)$, turn out to be close together in the Poincaré metric on $\sigma(\partial N)$. The thin part of Y lifts to a narrow strip, perhaps lying close to a geodesic joining the fixed points.

So far, the picture is still problematic. But suppose we stand back a little farther. Recall that $d\sigma(\mu)$ is obtained by pushing down a differential invariant under the full

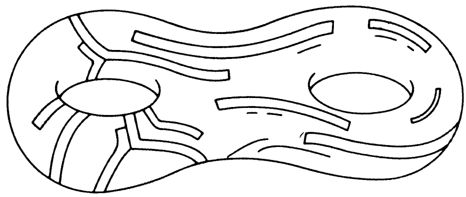

Fig. 3.

group $\pi_1(N)$, in particular under γ. This invariance forces the linefield to bend quite a bit near the pair of fixed points of γ.

More precisely, by choosing the viewing scale correctly we may take a geometric limit in which $\langle\gamma\rangle$ converges to a parabolic subgroup. Then the linefield of μ is like a family of tangent circles passing through the fixed point. On the other hand, the scale is small enough that the leaves of the foliation of $\phi \in Q(\sigma(\partial N))$ are nearly parallel. The limiting pairing $\langle\phi, \mu\rangle$ is clearly inefficient (Fig. 4).

The final picture reminds one of the punchline in Sullivan's proof that the limit set of a finitely generated group carries no invariant linefield [Su1].

6.3. Shape of short geodesics

The argument sketched above will now be justified.

Let N be a quasifuchsian manifold. To begin, we check that a geodesic short in the Poincaré metric on ∂N corresponds to a nearly parabolic element in $\pi_1(N)$.

Any hyperbolic $\gamma \in \pi_1(N)$ has a natural *complex* length $\mathscr{L}$, defined as follows. Put the repelling and attracting fixed points of γ at 0 and ∞ respectively, so $\gamma(z) = \lambda z$, $|\lambda| > 1$, and arrange that 1 is in the limit set Λ. There is a unique arc in $\Lambda - \{0, \infty\}$ connecting 1 to λ. Determine the complex length $\mathscr{L} = \log(\lambda)$ by analytic continuation along the arc, starting with $\log(1) = 0$.

Note that

- $\operatorname{Re}\mathscr{L}$ is the length of the closed geodesic in N homotopic to γ; while
- $\operatorname{Im}\mathscr{L}$ is the rate of spiraling of the limit set about the repelling fixed point.

γ also determines closed geodesics on the two Riemann surfaces at infinity for N; let L_+ and L_- denote their lengths in the respective Poincaré metrics.

Proposition 6.4. *The complex length and the Poincaré lengths at infinity satisfy*:

$$\frac{2\operatorname{Re}\mathscr{L}}{|\mathscr{L}|^2} \geqq \frac{1}{L_+} + \frac{1}{L_-}.$$

In particular, $|\mathscr{L}| \leqq 2L_+$, which implies:

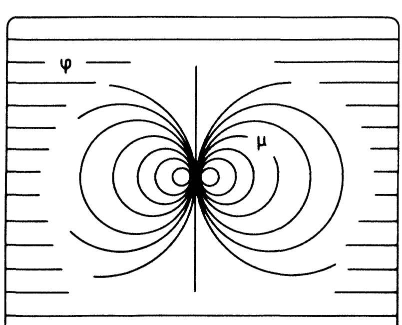

Fig. 4. Inefficient pairing

Corollary 6.5. *A geodesic which is short in the Poincaré metric at infinity is nearly parabolic in the quasifuchsian manifold: its complex length is near* 0, *and its multiplier* λ *is near* 1.

Proof of the Proposition. This is a well-known extremal length argument, compare [Bers1, Theorem 3], [Th3, Proposition 1.3]. Form the torus

$$T = [\hat{\mathbf{C}} - \{\text{the fixed points of } \gamma\}]/\langle\gamma\rangle\ ;$$

T may be identified with the quotient of $\mathbf{C}$ by the lattice $\mathbf{Z}2\pi i \oplus \mathbf{Z}\mathscr{L}$. The domain of discontinuity descends to a pair of disjoint annuli A_+ and A_- on T, conformally isomorphic to right cylinders whose moduli (the ratios height/circumference) are π/L_+ and π/L_-.

These annuli lie in the homotopy class on T corresponding to $\mathscr{L}$. A cylinder of maximum modulus in this homotopy class is the complement of a geodesic (with respect to the flat metric); its circumference is $|\mathscr{L}|$, its height is $2\pi \mathrm{Re}\,\mathscr{L}/|\mathscr{L}|$, and its modulus bounds the sum of the moduli of A_+ and A_-. . □

Proof of Theorem 6.1. Let $Y = \sigma(X)$ be a component of $\sigma(\partial N)$. Choose a quasifuchsian subgroup $\pi_1(X) \subset \pi_1(N)$ uniformizing X and Y, with limit set $\Lambda(X)$ and domain of discontinuity $\Omega(X) \sqcup \Omega(Y)$.

Let $x \in \Omega \cap \Omega(Y)$ lie over the geodesic thin part of $\partial N = \Omega/\pi_1(N)$. Then there is a hyperbolic element $\gamma \in \pi_1(N)$ which translates x by a small distance ε in the Poincaré metric on Ω.

As in the proof of Proposition 4.2, our program will be to consider a sequence of pictures as $\varepsilon \to 0$, and show all geometric limits exhibit inefficiency.

Let α and ω denote the fixed points of γ. Normalize coordinates on the Riemann sphere so α is the origin, infinity belongs to $\Lambda(X)$, and the center of the hyperbolic ball is translated a fixed distance ε_0 by γ. (This is possible for all ε sufficiently small).

Note that the set of Möbius transformations translating the center of the ball distance ε_0 form a compact set. By the preceding corollary, as $\varepsilon \to 0$ any limiting γ is parabolic.

We claim there exists an $\varepsilon_0(g, n)$, depending only on the type (g, n) of X, such that with this normalization the unit disk

$$\Delta = \{z : |z| < 1\}$$

(the lower hemisphere) is contained in $\Omega(Y)$ and injects into $Y = \Omega(Y)/\pi_1(X)$.

To see this, consider the convex core of $\mathbf{H}^3/\pi_1(X)$. Its boundary has the intrinsic geometry of a hyperbolic surface, so through any point there exists an essential loop of length less $L(g, n)$.

Since M is acylindrical, γ *does not lie in* $\pi_1(X)$.

Thus the Margulis lemma guarantees a large hyperbolic distance between the convex hull of $\Lambda(X)$ and any point of $\mathbf{H}^3$ where the translation length of γ is small. In particular a geodesic joining a point of $\Lambda(X)$ to ∞ must stay far from the center of the ball. It follows that $\Lambda(X)$ is confined to a small polar cap (see Fig. 5).

Equivalently,

$$\Omega(Y) \supset \Delta_R = \{z : |z| < R\}$$

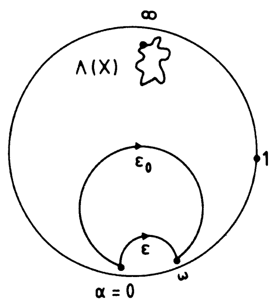

Fig. 5. The quasifuchsian limit set is confined to a polar cap

for some large R. Choosing ε_0 sufficiently small, we may conclude that the radius of Δ in the Poincaré metric on $\Omega(Y)$ is small, so we have injectivity so long as 0 maps to the thick part of Y. On the other hand, if 0 maps to the thin part of Y corresponding to an element $\delta \in \pi_1(X)$, then δ is geometrically close to a parabolic transformation fixing ∞ and translating the center of the ball a large hyperbolic distance (again by Margulis); it follows that Δ injects into Y in this case as well.

We claim that as $\varepsilon \to 0$, x tends to the origin. Indeed, the full limit set Λ is connected and contains both 0 (a fixed point of γ) and ∞, so the Poincaré metric on Ω is bounded below in terms of the spherical metric. By assumption γ translates x distance ε in the Poincaré metric, so its spherical translation is also small. But γ is close to a parabolic transformation fixing the origin and translating ∞ a definite distance, so only points near the origin have small spherical translation.

Let $B(x) \subset \Omega(Y)$ be the largest closed ball, in the Poincaré metric on $\Omega(Y)$, which is centered at x and which is contained in the lower hemisphere Δ. (Since $x \to 0$, there is such a ball for all ε sufficiently small).

To complete the proof, suppose the efficiency of pairing between ϕ and μ on $B(x)$ tends to 1 as $\varepsilon \to 0$; we will extract a geometric limit and obtain a contradiction.

We claim that as $\varepsilon \to 0$, any Hausdorff limit B of $B(x)$ contains the origin in its interior. For clearly $B(x)$ contains a definite neighborhood of 0 when x is sufficiently near 0.

Think of ϕ and μ as differentials on $\hat{\mathbf{C}}$ invariant under $\pi_1(X)$ and $\pi_1(N)$ respectively. Since $B(x)$ injects into Y, we may apply compactness of $\mathbf{P}\mathcal{Q}_{g,n}$ to obtain, after scaling and passing to a subsequence, a limiting quadratic differential $\phi \neq 0$ holomorphic on a limiting region B with 0 in its interior. We may extract a weak limiting μ, invariant under a limiting γ, which will be a parabolic transformation fixing the origin.

Now γ stabilizes a compact round ball $B' \subset B$, tangent to its direction of translation at the origin. If the efficiency of pairing on B is 1, then $\phi|B'$ is γ-invariant. But this contradicts finiteness of $\int_{B'} |\phi|$.

In passing to a geometric limit, the only uniformly used was the type (g, n) of X, which ranges through only a finite number of possibilities once $\partial_0 M$ is fixed. □

7. Cylindrical manifolds

Let (M, P) be a possibly disconnected pared manifold with gluing data τ, such that M/τ is connected.

The quotient M/τ is *atoroidal* if every torus

$$T: \mathbf{S}^1 \times \mathbf{S}^1 \to M/\tau ,$$

injective on fundamental group, is homotopic to a map into the parabolic locus P/τ.

Assume M is not an interval bundle over a surface. In this section we establish:

Theorem 7.1. *The gluing problem has a solution if and only if M/τ is atoroidal.*

7.1. *Creating a torus*

Suppose the gluing problem has no solution. How will we find a torus in M/τ? The idea is to look at short geodesics on ∂N, and see they bound cylinders in N joined by τ to give an essential torus in the quotient.

Proposition 7.2. *Let $N \in GF(M, P)$, and let $\langle \alpha_i, \beta_i, \gamma_i \rangle : i = 1, \ldots, n$ denote homotopy classes of simple geodesics on $\langle \partial N, BN, \partial N \rangle$ respectively. Suppose*

- *α_i lifts to β_i,*
- *β_i is homotopic through QN to γ_i and*
- *$\tau(\gamma_i) = \alpha_{i+1}$ (or α_1 if $i = n$).*

Then M/τ is toroidal.

Proof. The homotopy from β_i to γ_i, when mapped down to N, gives a compressing cylinder C_i connecting α_i to γ_i. These are joined end-to-end by τ to form a torus in N/τ, injective on fundamental group by incompressibility of ∂N. This torus is non-peripheral since it is centered on geodesics, so M/τ is toroidal. □

7.2. *Liftable short geodesics*

In this section we show that *liftable* short geodesics exist when σ fails to contract uniformly.

Indeed, the proof of Theorem 6.1 (relating short geodesics to inefficient pairing) uses the acylindrical assumption at only one point—which we have italicized. The ingredient actually used (in the notation of the proof) is that γ is not an element of $\pi_1(X)$. Geometrically, this means the Theorem generalizes as follows:

Theorem 7.3. *Let $x \in BN \subset \sigma(\partial N)$ lie over a geodesic component of the ε-thin part of ∂N (for ε sufficient small); then:*

Either the component of the thin part lifts to BN at x, or the pairing

$$\frac{\langle \phi, d\sigma(\mu) \rangle_B}{\|\phi\|_B} < c(\partial_0 M) < 1$$

for some ball B centered at x.

Definition. The *geodesic lifted part* $\mathscr{L}(BN, \varepsilon)$ is the union of all possible lifts to BN of components of $(\partial N)_{\text{geod}}(\varepsilon)$ and $(\partial N) - (\partial N_{\text{geod}}(\varepsilon))$. (See Fig. 6).

Corollary 7.4 *Let* $\phi \in Q(\sigma(\partial N))$, $\|\phi\| = 1$. *If* $\|d\sigma^*(\phi)\|$ *is close to* 1, *then most of the mass of* $|\phi|$ *is in the geodesic lifted part.*

More precisely: fix $\eta > 0$. *Then for all* $\varepsilon > 0$ *sufficiently small, there exists a* $\delta > 0$ *such that if* $\|\phi\| = 1$ *and* $\|d\sigma^*(\phi)\| > 1 - \delta$, *then*

$$\int_{\mathscr{L}(BN, \varepsilon)} |\phi| > 1 - \eta \, .$$

Proof. We have already seen that $\|d\sigma^*(\phi)\|$ close to 1 implies most of the mass of $|\phi|$ is in the amenable part BN_{am} (Proposition 5.4). Now the amenable part is the union of the liftable part and the total pre-image of the thin part. By Corollary 2.6, the $|\phi|$-mass lying over the cuspidal thin part is small if both ε and δ are small. By Theorem 7.3, over the geodesic thin part ϕ pushes down inefficiently unless the thin part actually lifts, and the corollary follows. □

7.3. Completion of the proof

Putting the pieces together, we now solve the general gluing problem.

The idea of the proof is straightforward: if $d\sigma^*(\phi)$ has nearly the same norm as ϕ, then most of the mass of ϕ has been communicated along cylinders corresponding to short geodesics. The number of short geodesics is bounded above. To avoid contraction over several iterates, these cylinders must eventually be matched up by τ to form a torus.

Proof of Theorem 7.1. The gluing problem has no solution if M/τ is toroidal, since an incompressible torus in a hyperbolic manifold is homotopic to a cusp.

We will prove the converse under the assumption that M is not an interval bundle over a surface.

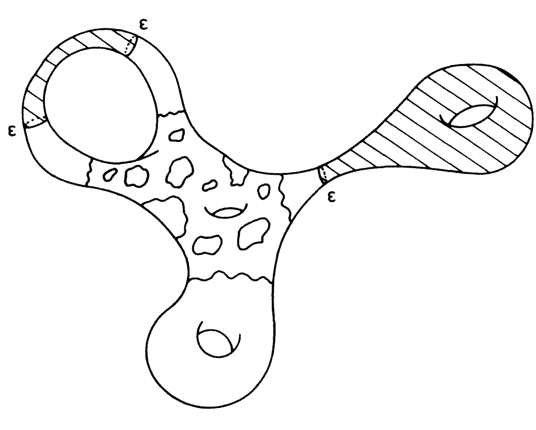

Fig. 6. The geodesic lifted part (*shaded*)

It is useful to restrict attention to the subset $GF(M, P, L)$ where the Teichmüller distance from N to $\tau \circ \sigma(N)$ is bounded by L. Fix L sufficiently large that $GF(M, P, L)$ is non-empty.

Since the $\tau \circ \sigma$ is non-expanding, it maps $GF(M, P, L)$ into itself. Moreover, if N lies in $GF(M, P, L)$, then so does the Teichmüller geodesic segment joining N to $\tau \circ \sigma(N)$. (This follows from the triangle *equality* for 3 points on a geodesic.)

If at all points of $GF(M, P, L)$, $\|d(\tau \circ \sigma)^k\| < c < 1$ for some fixed iterate k, the forward images of this geodesic segment have finite total length.Then $\tau \circ \sigma$ has a unique fixed point and the gluing problem is solved.

Otherwise a high iterate still fails to contract uniformly. Dually formulated, this means that for any K and $\delta > 0$, there are N_k in $GF(M, P, L)$ and $\phi_k \in Q(N_k)$ such that

$$1 + \delta \geqq \|\phi_k\| \geqq 1$$

for $k = 0, 1, 2, \ldots, K$; where

$$N_{k+1} = \tau \circ \sigma(N_k) \quad \text{and} \quad \phi_k = d(\tau \circ \sigma)^*(\phi_{k+1}) .$$

We will show this implies M/τ is toroidal.

We begin by choosing K and δ. Let

$$C = (\text{the number of components of } \partial_0 M)$$

$$S = (\text{a bound on the number of disjoint simple closed geodesics on } \partial N_k)$$
$$S' = (\text{a bound on the number of disjoint simple closed geodesics on } BN_k) .$$

Of course S and S' depend only on the topology of (M, P).

Choose

$$K = C + S .$$

This number of iterations will allow ϕ_k to first propagate to a component of M which is not an interval bundle over a surface; thereafter, we will define a collection of simple closed curves, and we have allowed enough time to force them to cycle.

Our choice of δ depends on two other quantities: ε, a length to specify 'very short' geodesics; and m, an amount of mass which we will consider 'negligible'.

There is a universal constant ε_0 such that the geodesics of length less than ε_0 on a hyperbolic surface are disjoint simple curves. Pick $\varepsilon > 0$ such that

$$\log(\varepsilon) + KL < \log(\varepsilon_0) .$$

We claim any closed geodesic of length less than ε on ∂N_k lies in one of at most S homotopy classes on $\partial_0 M$. Indeed, the Teichmüller distance from ∂N_k to ∂N_0 is at most KL, so a geodesic of length less than ε on ∂N_k has length at most ε_0 on ∂N_0 (cf. [Gar. §8.5]) and therefore belongs to one of at most S homotopy classes.

Notation. If α is a geodesic of length less than $\varepsilon/2$ on a Riemann surface X, $X(\alpha)$ will denote the corresponding component of the ε-thin part of X.

To choose m, first choose $\zeta > 0$ such that whenever a component X of ∂N_k carries geodesics shorter than $\varepsilon/2$, one of them, α, satisfies

$$\frac{|\phi_k|\text{-mass of } X(\alpha)}{|\phi_k|\text{-mass of } X} > \zeta .$$

This is possible by Proposition 2.7, using the fact that $X_{\text{geod}}(\varepsilon)$ has at most S components. Now set

$$m = \frac{\zeta}{C(2S')^K}\,.$$

Finally, choose $\delta > 0$ sufficiently small that for all k, all but m of the mass of $|\tau^*\phi_k|$ resides in the geodesic lifted part $\mathscr{L}(BN_k, \varepsilon/2)$. This is possible by Corollary 7.4.

Having made these choices, we now construct curves α_k, β_k and γ_k as in Proposition 7.2 to prove that M/τ is toroidal. The construction proceeds from two observations.

I. For some $k \leqq C$, there is a geodesic α_k on ∂N_k, shorter than $\varepsilon/2$, such that the $|\phi_k|$-mass of $\partial N_k(\alpha_k)$ is greater than ζ/C.

Indeed, for $k = C$, some component X of ∂N_k has $|\phi_k|$-mass at least $1/C$. This component may belong to an interval bundle over a surface; however since M/τ is connected, and not itself an interval bundle, after at most C iterates of $d(\tau \circ \sigma)^*$, we arrive at a k such that X does not belong to an interval bundle. Since $1/C > m$, X must contain a component of $\mathscr{L}(BN_k, \varepsilon/2)$ as a *proper* incompressible subsurface. In particular, this component must carry a geodesic of length less than $\varepsilon/2$ in the Poincaré metric on BN_k. By the Schwarz lemma, inclusions contract the Poincaré metric, so X itself has such a short geodesic. Thus (I) follows from the definition of ζ.

II. Suppose α_k is geodesic on ∂N_k, shorter than $\varepsilon/2$, such that the $|\phi_k|$-mass of $\partial N_k(\alpha)$ is μ. Assume

$$\mu' = (\mu - m)/S' > 0\,.$$

Then α_k lifts to a geodesic β_k on BN_k, of equal length, such that the $|\tau^*(\phi_{k+1})|$-mass of $BN_k(\beta_k)$ is at least μ'.

To prove II, we simply inquire, where did the mass of $\partial N_k(\alpha)$ come from under $d\sigma^*$? All but m of the mass of $|\tau^* \phi_{k+1}|$ belongs to the geodesic lifted part, which has at most S' components. Since pushforward does not increase mass, one of the lifts of $\partial N_k(\alpha)$ to BN_k has mass at least μ'.

Using I and II, we give the construction.

By I there is a $k = k_0 \leqq C$ and a short geodesic α_k with the $|\phi_k|$-mass of $\partial N_k(\alpha_k)$ at least ζ/C. For simplicity of notation, assume $k_0 = 0$. (We will only be considering k in the range $[k_0, k_0 + |S|]$, and since $k_0 + S \leqq C + S = K$, this amounts to a relabeling.)

For $k = 0, 1, 2, \ldots |S| - 1$, we construct β_k, γ_k and α_{k+1} from α_k by induction, such that:

1. α_k is a geodesic on ∂N_k, no longer than $\varepsilon/2$;
2. The $|\phi_k|$ mass of $\partial N_k(\alpha_k)$ is at least $\zeta/C(2S')^k$;
3. α_k lifts to β_k on BN_k;
4. β_k is homotopic through BN_k to γ_k; and
5. $\tau(\gamma_k) = \alpha_{k+1}$.

1 and 2 are true for α_0. Given α_k satisfying 1 and 2, apply II to construct β_k satisfying 3, and define γ_k and α_{k+1} by 4 and 5. We must verify that 1 and 2 hold for α_{k+1}.

By II the $|\tau^*\phi_{k+1}|$ mass of $BN_k(\beta_k)$ is at least $\zeta/C(2S')^{k+1}$, since $m = \zeta/C(2S')^K$ is less than half the $|\phi_k|$-mass of $\partial N_k(\alpha_k)$. Since the inclusion $BN_k \subset \partial\sigma(N_k)$ is a contraction, homotopic to β_k there is a possibly shorter geodesic β'_k on $\sigma(N_k)$ whose corresponding component of the ε-thin part has at least as much mass.

Recall

$$\sigma(\partial N_k) \in \mathrm{Teich}(\overline{\partial_0 M})\ ;$$

the curve β'_k is marked by the homotopy class γ_k on ∂N_k via an orientation-reversing homotopy through QN_k. Thus β'_k is in the homotopy class of $\tau(\gamma_k) = \alpha_{k+1}$ on the boundary of $\tau \circ \sigma(N_k) = N_{k+1}$. Assertions 1 and 2 for α_{k+1} then follow from their validity for β'_k.

We now have $S + 1$ short geodesics α_k, ranging among at most S distinct homotopy classes on $\partial_0 M$, so two are equal. By Proposition 7.2, M/τ is toroidal. □

References

[Bers1] Bers, L.: On boundaries of Teichmüller spaces and on kleinian groups: I. Ann. Math. **91**, 570–600 (1970)

[Bers2] Bers, L.: An extremal problem for quasiconformal maps and a theorem by Thurston. Acta Math. **141**, 73–98 (1978)

[DE] Douady, A., Earle, C.J.: Conformally natural extension of homeomorphisms of the circle. Acta Math. **157**, 23–48 (1986)

[DH] Douady, A., Hubbard, J.: A proof of Thurston's topological characterization of rational maps. Institute Mittag-Leffer preprint, 1985

[FLP] Fathi, A., Laudenbach, F., Poénaru, V.: Travaux de Thurston sur les surfaces. Asterisque **66–67**, 1979

[Gar] Gardiner, F.: Teichmüller theory and quadratic differentials. New York: Wiley Interscience 1987

[Hem] Hempel, J.: 3-manifolds. Ann. Math. Stud. **86**, Princeton University Press: 1976

[JS] Jaco, W., Shalen, P.: Seifert fibered spaces in 3-manifolds. Mem. Am. Math. Soc. vol. 220, 1979.

[Kra] Kra, I.: Deformation spaces. In: A crash course on Kleinian groups, Lecture Notes Math. Vol. 400. Berlin-Heidelberg-New York: Springer 1974

[LV] Lehto, O., Virtanen, K.J.: Quasiconformal mappings in the plane. Berlin-Heidelberg-New York: Springer 1973

[Mar] Marden, A.: The geometry of finitely generated kleinian groups. Ann. Math. **99**, 383–462 (1974)

[Mc1] McMullen, C.: Amenability, Poincaré series and quasiconformal maps. Invent. Math. **97**, 95–127 (1989)

[Mc2] McMullen, C.: Riemann surfaces and the geometrization of 3-manifolds (Submitted) Ann. New York Acad. Sci.

[Mor] Morgan, J.: On Thurston's uniformization theorem for three dimensional manifolds, pp. 37–125. In: The Smith conjecture. New York: Academic Press 1984

[MS1] Morgan, J., Shalen, P.: Valuations, trees, and degenerations of hyperbolic structures, I. Ann. Math. 120, 401–476 (1984).

[MS2] Morgan, J., Shalen, P.: An introduction to compactifying spaces of hyperbolic structures by actions on trees. In: Geometry and topology, Lecture Notes Math., Vol. 1167, pp. 228–240. Berlin-Heidelberg-New York: Springer 1985

[Mum] Mumford, D.: A remark on Mahler's compactness theorem. Proc. Am. Math. Soc. **28**, 289–294 (1971)

[Roy] Royden, H.L.: Automorphisms and isometries of Teichmüller space, pp. 369–384. In: Advances in the theory of Riemann surfaces. Princeton University Press, 1971

[Sul] Sullivan, D.: On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions. In: Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference, Ann. Math. Stud. Princeton University Press: 1981, pp. 465–496

[Th1] Thurston, W.P.: Hyperbolic structures on 3-manifolds I: Deformations of acylindrical manifolds. Ann. Math. **124**, 203–246 (1986)

[Th2] Thurston, W.P.: On the geometry and dynamics of diffeomorphisms of surfaces. Bull. Am. Math. Soc. **19**, 417–432 (1988)

[Th3] Thorston, W.P.: Hyperbolic structures on 3-manifolds II: Surface groups and 3-manifolds which fiber over the circle. (To appear) Ann. Math.

[Th4] Thurston, W.P.: Hyperbolic structures on 3-manifolds III: Deformation of 3-manifolds with incompressible boundary. (To appear) Ann. Math.

[Th5] Thurston, W.P.: Hyperbolic structures on 3-manifolds IV: Construction of hyperbolic manifolds. In preparation

[Th6] Thurston, W.P.: On the combinatorics and dynamics of iterated rational maps. Preprint

Oblatum 25-IV-1988 & 17-V-1989