A GEOMETRIC APPROACH TO THE COMPLEX OF CURVES ON A SURFACE

YAIR N. MINSKY

1. The complex of curves

Geometry has played an important role in the study of Teichmüller space and its group of symmetries, the Mapping Class Group of a surface. In [13], Harvey introduced a simplicial complex associated to a surface called the *complex of curves*, which was intended to encode the asymptotic geometry of Teichmüller space in analogy with similar constructions for symmetric spaces. This complex was then studied by Harer [11, 12] from a cohomological point of view and by Ivanov [16, 15] with applications to the structure of the mapping class group.

Harvey defined his complex as follows. Let S be a finite genus surface with finitely many punctures, and let "essential simple curve" on S denote a non-trivial homotopy class of simple curves which are also non-peripheral, i.e. cannot be deformed to a puncture. The essential simple curves are the vertices of the complex, and k-simplices are sets of $k+1$ curves that can be realized disjointly. It is immediate from the definition that this complex is finite dimensional, but also that it is usually locally infinite – there are usually infinitely many homotopically distinct curves in the complement of a simple curve.

In joint work with Howard Masur, we have been studying Harvey's complex from the point of view of hyperbolic metric spaces in the sense of Gromov and Cannon (see below). Part of our motivation is an attempt to explain the presence and absence of negative-curvature properties of the Teichmüller space, whereas further long-term motivation includes an improved understanding of the mapping class group, and possible applications to Kleinian groups.

For a few simple cases (such as the torus) we will see in Section 3 that this complex is 0-dimensional, and in those cases we will make an adjustment to the definition that yields a 1-complex. In any case let $K(S)$ denote the 1-skeleton of the complex.

To state our main theorem, we need Gromov's notion of δ-*hyperbolicity*, or "negative curvature" (in the large) in Cannon's terminology, which may be given in terms of the following "thin triangle" property: Let X be a geodesic

Date: May 15, 1996.

Appeared in Proceedings of the 37th Taniguchi Symposium on Topology and Teichmüller Spaces (S. Kojima et. al., ed.), World Scientific, 1996, pp. 149–158.

metric space. For any three points $x, y, z \in X$ let $[xy], [xz], [yz]$ denote geodesics between pairs (not necessarily unique). If $[xz]$ is always contained in a δ-neighborhood of $[xy] \cup [yz]$ for some fixed $\delta \geq 0$, we say that X is δ-hyperbolic. This property encodes many of the most characteristic large-scale properties of spaces such as hyperbolic space $\mathbf{H}^n$, or infinite simplicial trees (see [6, 10, 9]).

Theorem (Masur-Minsky [19]) *The 1-complex $K(S)$ endowed with the metric giving length 1 to edges is δ-hyperbolic for some $\delta(S) \geq 0$*

Remarks. 1. Note that, from the point of view of this theorem, it does not matter if we include the entire complex or just its 1-skeleton, since if we endow each simplex with a standard (say Euclidean) metric then the full complex is quasi-isometric to its 1-skeleton. The property of δ-hyperbolicity is invariant under quasi-isometry, although the specific value of δ is not. 2. The definition of δ-hyperbolicity often includes the additional assumption of *propriety*, i.e. that balls of finite radius are compact. This is clearly false for $K(S)$, and one might emphasize the distinction by saying that $K(S)$ is an improper hyperbolic space.

This article is an exposition of the ideas and motivations behind this theorem. We will discuss a sketch of the proof, and note some remaining questions and future plans.

2. Geometry of Teichmüller space

A persistent, but incomplete, analogy has long been recognized between the geometry of the Teichmüller space $\mathcal{T}(S)$ and that of δ-hyperbolic spaces. For example, $\mathcal{T}(S)$ admits a boundary at infinity similar to that of hyperbolic space (See e.g. [8, 23]), and the Teichmüller geodesic flow on its quotient, the moduli space, is ergodic (see [22]). The analogy was exploited, for example, in Bers [1], Kerckhoff [18], and Wolpert [29]).

On the other hand, Masur [20] first showed that the Teichmüller metric cannot be negatively curved in a local sense, and more recently Masur-Wolf [24] showed that it is not δ-hyperbolic. Although the Weil-Petersson metric on $\mathcal{T}(S)$ has negative sectional curvatures, they are not bounded away from zero [21, 30].

The connection between the complex of curves and the geometry of $\mathcal{T}(S)$ is based on the collar lemma (see e.g. [17, 5]) which implies in particular that there is a universal $\epsilon_0 > 0$ such that any collection of simple geodesics of length less than ϵ_0 on a hyperbolic surface are pairwise disjoint. In other words, vertices in $K(S)$ are on a common simplex precisely when the curves they represent can be made simultaneously short for some point in $\mathcal{T}(S)$. If we define the "thin region" of an essential curve α as $T(\alpha) = \{\sigma \in \mathcal{T}(S) : \ell_\sigma(\alpha) < \epsilon_0\}$ (here ℓ_σ denotes hyperbolic length in the metric σ) then $K(\alpha)$ is therefore the *nerve* of the family of regions $\{T(\alpha)\}$. (For a reinterpretation of this in the case of tori and spheres, see section 3).

It has often been noted that these regions where some curve is short are precisely the places where the geometry of $\mathcal{T}(S)$ is non hyperbolic. This is related to the fact that the presence of a short curve α decomposes the surface into pieces that are geometrically nearly independent, thus giving the region $T(\alpha)$ the approximate geometry of a product (see [21, 30, 26]). A product of infinite-diameter spaces cannot be δ-hyperbolic. This is related to the fact that the mapping class group of a higher-genus surface is known not to be δ-hyperbolic because it contains high rank abelian subgroups, generated by elements whose supports lie on disjoint subsurfaces.

On the other hand, outside the thin regions there is evidence of hyperbolic behavior. For example, we show in [27] that a Teichmüller geodesic that stays away from $\cup_\alpha T(\alpha)$ admits a closest-point retraction that satisfies a contraction property (see Section 4.2) which characterizes such projections in a hyperbolic space.

In this context, one interpretation of our main theorem is that $\mathcal{T}(S)$ is *relatively hyperbolic* with respect to the family of regions $T(\alpha)$. This terminology is due to Farb [7], who showed for example that a Kleinian group with finite-volume but non-compact quotient is relatively hyperbolic with respect to the abelian subgroups corresponding to cusps. In other words, all the non-hyperbolicity in $\mathcal{T}(S)$ is "restricted to" the regions $T(\alpha)$, but the way in which these regions are arranged in $\mathcal{T}(S)$ is itself hyperbolic.

3. The low dimensional cases

In a small number of cases, the complex of curves is either trivial or "classically" known, and our Theorem holds vacuously or is fairly easy to prove. If S is a sphere with between 0 and 3 holes, then $K(S)$ is empty, and hyperbolicity for the empty set is a matter of terminology.

If S is a torus, once-punctured torus or 4-times punctured sphere, then any two essential simple curves must intersect, so that $K(S)$ has no edges under the general definition. However, with a small adjustment in the definition K becomes a sensible and familiar 1-complex: Let (α, β) be an edge whenever α and β have the lowest possible intersection number. For the tori this is 1, and for the 4-punctured sphere this is 2.

For the rest of this section assume that S is a torus (with 0 or 1 holes). The sphere case is similar. Choosing an identification of $H_1(S)$ with $\mathbf{Z}^2$, the essential simple curves on S are well known to be in one-to-one correspondence with the slopes of elements in $\mathbf{Z}^2$, namely the rational numbers p/q including $1/0 = \infty$. Thus the 0-skeleton K_0 is identified with $\mathbf{Q} \cup \infty$ in the circle $S^1 = \mathbf{R} \cup \infty$, and the identification is unique up to the action of $\mathrm{SL}_2(\mathbf{Z}) = Isom(\mathbf{Z}^2)$ on the circle.

The full 1-complex K may then be embedded as an ideal triangulation of the unit disk bounded by S^1 ("ideal" means that the vertices are on the boundary) by realizing each edge as a hyperbolic geodesic in the Poincaré model of $\mathbf{H}^2$. To see that this is an embedding, we note first that the unsigned intersection number of curves represented by p/q and r/s (written in lowest terms) is simply the absolute value of the determinant $|ps - qr|$. One may easily check this for $p/q = \infty$ and apply $\mathrm{SL}_2(\mathbf{Z})$ to get the general case. (For the 4-punctured sphere, twice the determinant is the intersection number.) Thus we note that there is an edge from ∞ to 0, and that no other edge crosses this one – for if $|ps-qr| = 1$ the signs of p/q and r/s must agree. Applying $\mathrm{SL}_2(\mathbf{Z})$ we see that no two edges intersect, so the complex embeds. We further check that ∞ is connected by edges to exactly the integers $n/1$, and that $n/1, m/1$ are connected exactly when $n = m \pm 1$. We therefore have a triangulation around ∞, and hence around every vertex – see figure 1. This famous picture, familiar in the study of modular forms, is also known as the Hatcher-Thurston complex for the torus; see [14, 4]. Note also that up to combinatorial equivalence there is only one ideal triangulation of the disk.

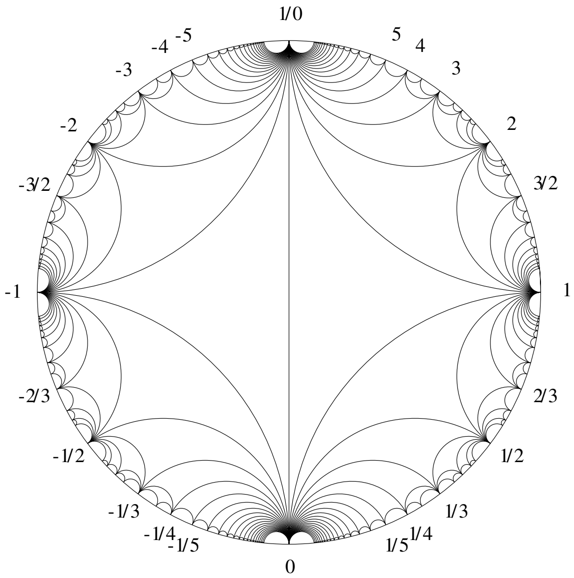

FIGURE 1. The complex of curves on the torus, embedded in the unit disk

Let us now prove K is hyperbolic. We begin by giving a description of the geodesics in K.

Because of its embedding in the disk, K has the property that any edge e separates K into two components. Given two vertices x, y let $E(x, y)$ denote the set of edges that separate x from y. Then not just geodesics but every continuous path in K from x to y must pass through some vertex (or both) of every edge $e \in E(x, y)$.

If $E(x, y)$ is empty then x and y are vertices of the same triangle and the geodesic between them is an edge. From now on assume $E(x, y) \neq \emptyset$. We will check that any geodesic from x to y is in the $\frac{1}{2}$-neighborhood of $E(x, y)$; in fact we will show a bit more.

It is easy to see that $E(x,y)$ admits an order where $e < f$ if e separates x from the interior of f. Successive elements in this order share a vertex, called a *pivot*, and the set of pivots is thus ordered as well. If p is a pivot let $w(p)$ (its "width") be the number of edges in $E(x,y)$ incident to p. (Remark: the sequence of numbers $\{w(p)\}$ is exactly the "left-right" sequence that appears in the classification of homeomorphisms of the torus, and is closely related to continued-fraction expansions for the slopes of curves.)

For $e \in E(x,y)$ consider the ways in which a geodesic $[xy]$ may cross e. If e is incident to no pivots then $E(x,y) = \{e\}$ and x and y are opposite vertices of a quadrilateral consisting of two triangles in K. Thus $d(x,y) = 2$ and there are two geodesics, going around the quadrilateral in the two possible ways.

If e is incident to a pivot p, let $e_1 < \cdots < e_{w(p)}$ be the sequence of edges incident to p (including e), and let v_j be the vertex of e_j different from p.

A shortest path from x to y must enter the configuration $\cup e_i$ at one vertex or the other of e_1 and exit through $e_{w(p)}$. (Note that x may be equidistant from p and v_1 or closer to v_1, but not closer to p, and similarly for y and $p, v_{w(p)}$.) If $w(p) \geq 4$ then the path must pass through p and not around through $v_2, \ldots, v_{w(p)-1}$, but if $w(p) = 2$ or 3 there may be (finitely many) choices. In any case a shortest path cannot pass through any vertices on the other side of an edge (v_j, v_{j+1}) from p, since any such digression would be longer than simply traversing the edge.

Thus we see that although geodesics may not be unique, they must pass through the block of edges around each pivot in one of a finite number of well-defined ways, and in particular they pass entirely through vertices of edges in $E(x,y)$ and hence must remain in a $\frac{1}{2}$-neighborhood of $E(x,y)$.

Now consider a triple of points x, y, z and the triangle $[xyz]$ formed by geodesics $[xy], [yz]$ and $[xz]$. We must show that this triangle satisfies the thin triangle property.

Let e be an edge of $E(x,z)$. If e does not separate x from y then either e separates y from z, or y is a vertex of e. Thus $E(x,z)$ is contained in the union of $E(x,y) \cup E(y,z)$ together with the edges incident to y. It follows that the vertices of any geodesic $[xz]$ are distance 1 from $[xy] \cup [yz]$ and thus that $[xz]$ itself is at most distance 3/2 from $[xy] \cup [yz]$. Thus K is 3/2-hyperbolic.

Let us also remark that the interpretation of $K(S)$ as the nerve of the family of regions $\{T(\alpha)\}$ in Section 2 is still valid, with appropriate definitions. The Poincaré disk in which we embedded $K(S)$ is in fact isometric to the Teichmüller space of the torus (or punctured torus, or 4-punctured sphere). It is not hard to show that there is a constant $\epsilon_1 > 0$ such that for any hyperbolic punctured torus, there are at most two curves of length ϵ_1 or less. If there are two, they must have intersection number 1. Thus $K(S)$ is the nerve of the family of regions $T(\alpha)$ defined using ϵ_1 in place of ϵ_0, and in fact two such regions $T(\alpha), T(\beta)$ meet in a point of tangency in the interior of the edge (α, β) in our embedding of $K(S)$. A similar statement can be made for the 4-punctured sphere and, using extremal length, for the unpunctured torus.

4. Sketch of the proof in the general case

For surfaces S not covered in the previous section, we return to Harvey's original definition for $K(S)$, and the proof becomes more complicated. There are two main components: (a) control of distance in K via nesting of train-tracks, and (b) a projection lemma using Teichmüller geometry.

Short distances in K are relatively easy to understand. That is, $d(\alpha, \beta) = 1$ exactly when α and β are disjoint. $d(\alpha, \beta) = 2$ when α and β intersect, but are both disjoint from an essential curve γ. Finally, $d(\alpha, \beta) > 2$ whenever α and β *bind* S; that is, the components of $S \setminus (\alpha^* \cup \beta^*)$ are disks and once-punctured disks (where α^* and β^* are geodesic representatives in some hyperbolic metric). However, distinguishing between distances of 3 or more requires a further idea.

4.1. Train-tracks and nested polyhedra.

We briefly recall Thurston's notion of a *train-track* (see Penner-Harer [28] for a complete treatment). A train-track is a smoothly embedded 1-complex τ in S (considered up to isotopy) with edges, called *branches*, meeting tangently at the vertices, which are called *switches*. A simple closed curve can be *carried* on τ if it is homotopic to a non-degenerate smoothly immersed curve in τ, and in that case the degree of covering of each branch gives a positive measure on the branches which satisfies the *switch condition*: at each switch, total incoming measure equals total outgoing measure.

Further topological constraints are also imposed on τ to guarantee that curves carried on it are homotopically essential and determined uniquely by their measures, and that each branch is traversed by some closed curve (recurrence). One can further demand (birecurrence) that τ can be realized in a hyperbolic metric on S in such a way that all edges are long and nearly geodesic. A finite number of such train-tracks suffice to carry all the essential simple closed curves in S.

The set of real positive weights on τ satisfying the switch conditions is, modulo scaling, a convex polyhedron $P(\tau)$ in the projective space $P\mathbf{R}^B$ where B is the set of branches of τ. The rational points in $P(\tau)$ correspond to simple curves or unions of simple curves.

We can make the following observation, where a *maximal* train track is one all of whose complementary components are triangles or punctured monogons.

Lemma *Let τ be maximal and let a vertex α of $K(S)$ determine a measure in the interior of $P(\tau)$. If $d(\alpha, \beta) = 1$ then β is carried in τ.*

In other words, a 1-neighborhood of $int(P(\tau))$ is in $P(\tau)$. Here "interior" means the set of measures positive on every branch. A sketch of the proof is this: Assume that τ is realized nearly geodesically and α and β are geodesics, in some hyperbolic metric on S. Since α traverses every branch of τ, β cannot intersect τ essentially since it is disjoint from α. Since τ is maximal any segment of β in the complement of τ can be pushed onto τ and we may conclude that β is carried on τ.

Now we recall that a train-track τ' may be *carried* on τ if it can be homotoped so that all its edges are immersed in τ. Then every curve carried on τ' is carried on τ, and $P(\tau') \subset P(\tau)$. In view of the Lemma, if we now have a sequence $\tau_1, \ldots, \tau_n$ such that $P(\tau_{j+1}) \subset int(P(\tau_j))$, then if α is carried on τ_n and β is not carried on τ_1, we may conclude $d(\alpha, \beta) \geq n$.

In particular one may use this idea to show that $\mathrm{diam}(K) = \infty$, a fact which may not have been obvious before, in view of the local infiniteness of the complex.

A generalization of this observation to non-maximal train tracks allows us to go the other way – given two curves far apart we can find a sequence of train-tracks whose polyhedra are nested (in a slightly weaker way) and separate the two curves.

4.2. **A projection lemma.** In a δ-hyperbolic space X, geodesics have the following *contraction property*. Any geodesic L in X admits a projection $\pi_L : X \to L$ taking each point in X to a closest point on L, and π is contracting in this sense: $\mathrm{diam}(\pi(B_x(d(x, L))))$ is uniformly bounded for all $x \in X$, where $B_x(R)$ is an R-neighborhood of x. This is easy to see for example when X is a tree, or $\mathbf{H}^n$.

Conversely, one may prove that if X contains a family of paths that connect any two points in X and each admitting projections with this contraction property (with uniform bounds), then the paths are quasi-geodesics and X is δ-hyperbolic.

Our proof of hyperbolicity for K uses this approach, showing in particular that a family of paths in K induced by the family of Teichmüller geodesics in Teichmüller space has the contraction property.

For any point σ in $\mathcal{T}(S)$, let $\Phi(\sigma)$ denote the set of curves of minimal extremal lengths in σ. It is not hard to see that this is a finite set, of uniformly bounded diameter in $K(S)$. With this, we can convert any Teichmüller geodesic L to a path in K, simply by choosing some element of $\Phi(t)$ for each $t \in L$ (we do not require paths to be continuous, since hyperbolicity only requires precision within some bounded constant). By slight abuse of notation, we call this path $\Phi \circ L$.

To describe the projection from $K(S)$ to $\Phi \circ L$, we use the notion of *balance*. The Teichmüller geodesic L is determined by a holomorphic quadratic differential q on S, whose horizontal and vertical foliations are respectively stretched and shrunk as we move along L in one direction. Thus for every curve α and point on L there is a horizontal and vertical length with respect to these foliations. As long as α is not completely vertical or horizontal, there is a unique point $t_\alpha \in L$ where its horizontal and vertical lengths are equal, and we say α is balanced at t_α. Our "projection" map is then defined by taking α to a point in $\Phi(t_\alpha)$.

To prove that this projection has the contraction property, we need to prove something like this. Let α be a point in $\Phi \circ L$ and let β, γ be two other points in $K(S)$ which are far from α but relatively close to each other (that is, $d(\beta,\gamma)/d(\beta,\alpha)$ is bounded by a suitable constant), and suppose that β projects to α. We would like to know that the projection of γ is a bounded distance from α.

The basic idea is that, as we move along L, the transition of any curve μ from mostly vertical to mostly horizontal occurs in a bounded distance as measured in $K(S)$ (although not as measured in Teichmüller distance), once we reach the point where μ does not consist of nearly vertical segments filling up almost the entire surface. As soon as the nearly vertical part of μ misses some definite horizontal arc in S (as measured in the quadratic differential metric), we show that a subsurface filled by the vertical parts begins to steadily shrink, making a bounded number of topological reductions until it is a disk, and the boundary components of the resulting sequence of nested surfaces yield a bounded path in $K(S)$.

Using this argument, we show that if β is balanced and γ is far from balanced (say more vertical than horizontal), we can move a bounded distance to a point where β is almost completely horizontal and γ is still almost completely vertical, and both fill up the whole surface.

At this point we use the train-track machinery to reach a contradiction. The distance of β from α implie the existence of a long sequence of nested train-track polyhedra which contain β, and the bound on $d(\beta,\gamma)$ then traps γ deep inside a long subsequence of these polyhedra (the actual proof is complicated by the fact that the train-tracks may not be maximal). It follows that there is a train-track which simultaneously carries both β and γ, but which intersects any short curve on the surface many times. In particular the intersection number of β and γ with each other is considerably smaller than their intersection numbers with any short curve. We then show that this contradicts the fact that one of them is nearly horizontal and the other nearly vertical. We conclude that in fact β and γ are balanced a bounded distance apart in $\Phi \circ L$.

With the contraction property thus established, we conclude simultaneously that $K(S)$ is hyperbolic, and that the class of paths determined by Teichmüller geodesics are quasi-geodesics in $K(S)$. We note that this fact is also easy to verify in the simpler cases discussed in Section 3.

5. Further comments and questions

δ-hyperbolicity is useful often because of the control one obtains on local geometry using *asymptotic* data. For example, two geodesic segments whose endpoints are some fixed distance apart remain a bounded distance apart along their whole lengths (we call them "fellow travelers"). Using this idea, there is a canonical *boundary at infinity* ∂X for a hyperbolic space X, consisting of equivalence classes of infinite geodesic rays, so that two geodesics whose endpoints at infinity are the same are fellow travelers.

One is thus led to ask, for example, what is the boundary at infinity for $K(S)$. In the case of the torus it is not hard to see that the boundary may be identified with the irrational numbers. In general, one can show that a sequence of curves α_n whose distance from a fixed base point is going to infinity must (up to subsequence) converge to a maximal geodesic lamination, and that if two sequences converge to distinct laminations then they determine different points at infinity. However when two sequences converge to the same geodesic lamination one has to consider the possibility that they converge to it in two very different ways – in particular, it may be that the lamination supports two distinct projective classes of transverse measures, and that as measured laminations the sequences have different limits. Thus the question remains whether the boundary can be identified with the space of maximal geodesic laminations, or maximal geodesic laminations with projective measure classes, or something else.

For a proper δ-hyperbolic space X, the union $X \cup \partial X$ in a natural topology is compact. The same is not true in our case, because of the lack of local finiteness. This in fact indicates a basic difficulty in applying hyperbolicity to study the complex of curves: namely, what good is it to know that two geodesics are a bounded distance apart, if a ball of bounded distance still contains an infinite number of vertices? Note that this question is nicely resolved in the case of the complex for the torus, where we had an explicit description for geodesics in spite of the fact that the link of every vertex was infinite. In particular we note that for every pivot on the path between two vertices, if the width w of the pivot was sufficiently large (at least 4) then the geodesic was forced to travel through, and not around, the pivot. This idea generalizes to higher genus. That is, given one geodesic passing through some vertex x, if its entry and exit points are sufficiently far apart in the link of x (which we note is itself a complex of curves for a subsurface) then we can show that any fellow travelling geodesic must pass through x as well. This opens the way for an inductive study of $K(S)$, since as soon as two geodesics are known to pass through the same vertex, they both induce paths in the complex corresponding to the link of that vertex. In an upcoming paper we hope to develop this idea, and apply it to a study of the mapping class group.

Harvey's complex arises naturally in the study of Kleinian groups, where one of the intriguing open problems is Thurston's Ending Lamination Conjecture, which proposes that a hyperbolic 3-manifold is uniquely determined by its topological type and a list of *end invariants* encoding the asymptotic geometry of its ends. For example, if N is a geometrically infinite hyperbolic 3-manifold homotopy-equivalent to a surface S then the ending laminations of N describe the limits of sequences of closed curves on S whose geodesic representatives in N leave every compact set in N (see Bonahon [2]).

A solution to Thurston's conjecture would require some way to convert the asymptotic information in the ending laminations to more precise geometric information. For example, one may ask to understand the structure of the set K_L of vertices of $K(S)$ whose geodesic lengths in N are at most L, for some fixed $L > 0$. In view of the hyperbolicity theorem one might seek to show that K_L is quasi-geodesic. This result was shown in [25] for the case that S is a punctured torus, and played a role in our proof of the Ending Lamination Conjecture for that case. (Bowditch [3] has also considered the relation between $K(S)$ and representations in $\mathrm{PSL}_2(\mathbf{C})$ when S is a punctured torus, showing for example that a set defined similarly to K_L is connected in $K(S)$, even if the representation is not discrete).

In the general case there is some hope that a similar analysis of K_L can go through, and become a first step towards a solution of the conjecture.

Finally we remark that we do not know a purely combinatorial proof of the hyperbolicity theorem which does not use Teichmüller theory. In addition, except in the low-dimensional cases the constant δ is completely non-constructive, as it depends on bounds obtained from compactness arguments on spaces of quadratic differentials. It would be very interesting to find a proof which places some bounds on δ, or at least describes its asymptotic dependence on the topological type of S.

References

1. L. Bers, *An extremal problem for quasiconformal mappings and a theorem by Thurston*, Acta Math. **141** (1978), 73–98.
2. F. Bonahon, *Bouts des variétés hyperboliques de dimension 3*, Ann. of Math. **124** (1986), 71–158.
3. B. Bowditch, *Markoff triples and quasifuchsian groups*, University of Southampton Preprint #249, 1995.
4. B. Bowditch and D. B. A. Epstein, *Natural triangulations associated to a surface*, Topology **27** (1988), 91–117.
5. P. Buser, *Geometry and Spectra of Compact Riemann Surfaces*, Birkhäuser, 1992.
6. J. Cannon, *The theory of negatively curved spaces and groups*, Ergodic theory, symbolic dynamics, and hyperbolic spaces (Trieste, 1989), Oxford Univ. Press, 1991, pp. 315–369.
7. B. Farb, *Relatively Hyperbolic and Automatic Groups with Applications to Negatively Curved Manifolds*, Ph.D. thesis, Princeton University, 1994.
8. A. Fathi, F. Laudenbach, and V. Poenaru, *Travaux de Thurston sur les surfaces*, vol. 66-67, Asterisque, 1979.
9. E. Ghys and P. de la Harpe, *Sur les groupes hyperboliques d'aprés Mikhael Gromov*, Birkhäuser, 1990.
10. M. Gromov, *Hyperbolic groups*, Essays in Group Theory (S. M. Gersten, editor), MSRI Publications no. 8, Springer-Verlag, 1987.
11. J. Harer, *Stability of the homology of the mapping class group of an orientable surface*, Ann. of Math. **121** (1985), 215–249.
12. ______, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), 157–176.

13. W. J. Harvey, *Boundary structure of the modular group*, Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (I. Kra and B. Maskit, eds.), Ann. of Math. Stud. 97, Princeton, 1981.
14. A. E. Hatcher and W. P. Thurston, *A presentation for the mapping class group*, Topology **19** (1980), 221–237.
15. N. V. Ivanov, *Automorphisms of complexes of curves and of Teichmüller spaces*, Preprint.
16. ———, *Complexes of curves and the Teichmüller modular group*, Uspekhi Mat. Nauk **42** (1987), 55–107.
17. L. Keen, *Collars on Riemann surfaces*, Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland 1973), Ann. of Math. Studies 79, Princeton, 1974, pp. 263–268.
18. S. Kerckhoff, *The Nielsen realization problem*, Ann. of Math. **117** (1983), 235–265.
19. H. Masur and Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, In preparation.
20. H. A. Masur, *On a class of geodesics in Teichmüller space*, Ann. of Math. **102** (1975), 205–221.
21. ———, *The extension of the Weil-Petersson metric to the boundary of Teichmüller space*, Duke Math. J. **43** (1976), 623–635.
22. ———, *Transitivity properties of the horocyclic and geodesic flows on moduli space*, J. D'analyse Math. **39** (1981), 1–10.
23. ———, *Two boundaries of Teichmüller space*, Duke Math. J. **49** (1982), no. 1, 183–190.
24. H. A. Masur and M. Wolf, *Teichmüller space is not Gromov hyperbolic*, MSRI preprint No. 011-94, 1994.
25. Y. Minsky, *The classification of punctured-torus groups*, Preprint.
26. ———, *Extremal length estimates and product regions in Teichmüller space*, Stony Brook IMS Preprint #1994/11, to appear in Duke Math. J.
27. ———, *Quasi-projections in Teichmüller space*, Stony Brook IMS Preprint #1994/16, to appear in J. für die Reine und Angew. Math.
28. R. Penner with J. Harer, *Combinatorics of train tracks*, Annals of Math. Studies no. 125, Princeton University Press, 1992.
29. S. A. Wolpert, *Geodesic length functions and the Nielsen problem*, J. Differential Geom. **25** (1987), 275–296.
30. ———, *The hyperbolic metric and the geometry of the universal curve*, J. Differential Geom. **31** (1990), 417–472.

State University of New York, Stony Brook, NY 11794
E-mail address: `yair@math.sunysb.edu`