Markoff triples and quasifuchsian groups.

B. H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
bhb@maths.soton.ac.uk

0. Introduction.

In this paper we aim to study the geometry of certain 2-generator subgroups of $PSL(2, \mathbf{C})$. This area lies close to the interaction of hyperbolic geometry with algebraic number theory and with complex dynamics.

A diophantine equation much studied by number theorists is the "Markoff equation": $x^2 + y^2 + z^2 = xyz$. The rational integral solutions to this equation are closely related to the Markoff and Lagrange spectra, and hence to diophantine approximation of real numbers — see, for example, [17]. (In this context, the Markoff equation is often given as $x^2 + y^2 + z^2 = 3xyz$, which corresponds to solving our version of the equation over $3\mathbf{Z}$. However, since there are there are no non-trivial solutions to $x^2 + y^2 + z^2 = xyz$ over $\mathbf{Z}_3$, this amounts to the same thing.)

The Markoff equation is also closely related to the geometry of the modular group $PSL(2, \mathbf{Z})$: see for example, [23,40]. One way to understand this is to consider the commutator subgroup, G, of $PSL(2, \mathbf{Z})$. Since $PSL(2, \mathbf{Z}) \cong \mathbf{Z}_2 * \mathbf{Z}_3$, one can see that G is a free group on two generators and of index 6 in $PSL(2, \mathbf{Z})$. If we identify $PSL(2, \mathbf{R})$ as the group of orientation preserving isometries of the hyperbolic plane, $\mathbf{H}^2$, then $\mathbf{H}^2/G$ is a punctured torus with a complete finite-area hyperbolic structure, sometimes refered to as the "modular torus". The Markoff equation arises as a trace identity in $PSL(2, \mathbf{R})$, and the integral solutions give the traces of simple closed curves on $\mathbf{H}^2/G$. The traces are related by a simple formula to the lengths of the corresponding closed geodesics. For some applications to the simple length spectrum of the modular torus, see for example [39] and the references therein.

More generally, we can look at other complete finite-area hyperbolic structures on the (topological) once punctured torus. In this case, the traces of simple closed curves arise as solutions to the Markoff equation over the reals.

Such a hyperbolic structure is given by a representation, ρ, of the free group on two generators, $\Gamma = \langle a, b \rangle$ into $PSL(2, \mathbf{R})$ with the property that the commutator $[\rho(a), \rho(b)]$ has trace equal to -2; (note that its sign is well-defined). More generally still, we can consider representations of Γ into $PSL(2, \mathbf{C})$ having this property. This corresponds to solving the Markoff equation over the complex numbers. This will be the main object of study in this paper. Connections between complex Markoff triples and representations are also explored in [26] and [14].

It is often more convenient to work with $SL(2, \mathbf{C})$ rather than $PSL(2, \mathbf{C})$, though the distinction will be of little importance here — it just means being careful with regards sign conventions etc.

Although it will not be a direct concern of us here, it is natural to look for integer solutions to the Markoff equation over other algebraic number fields, for example $\mathbf{Q}(\sqrt{-d})$. See [41] for some discussion of this. They are also considered in [10].

Some of the ideas described in this paper are used in [6] to give a direct proof of McShane's identity [32] for hyperbolic punctured tori (see Theorem 3), and extended in [7] to give similar identities for hyperbolic once-punctured torus bundles fibring over the circle. Similar ideas were also used in [10] to give a classification of arithmetic hyperbolic once-punctured torus bundles over the circle.

Recently, Minsky [33] has given a proof of the ending lamination conjecture for discrete representations of the once-punctured torus group into $PSL(2,\mathbf{C})$. An essential part of the argument makes use of trees of Markoff triples. (In particular he gives an independent proof of a variant of Theorem 1 of this paper). We suspect that the interplay between Markoff triples and representations might shed light on other conjectures concerning the geometry of hyperbolic 3-manifolds — as well as a provide a means of computational exploration of specific cases.

A connection with complex dynamics arises when one asks, for example, which complex solutions to the Markoff equation correspond to representations into $PSL(2,\mathbf{C})$ which are discrete and faithful. This fits into a broader category of questions, concerned with asking when two elements $A, B \in PSL(2,\mathbf{C})$ generate a discrete group. One would like some condition on the traces, for example, of A, B and $[A,B]$. For example, Jørgensen's inequality [25,2] tells us that if A and B generate a discrete non-elementary group, then $|(\operatorname{tr} A)^2 - 4| + |\operatorname{tr}[A,B] - 2| \geq 1$. Possible ways of generalising Jørgensen's inequality were considered by Brooks and Matelski [11], leading them to define the connectedness locus for quadratic maps. These ideas have more recently been taken up by Gehring and Martin [20], and greatly extended. Since the sets of traces corresponding to discrete groups turn out to be "fractal" in nature, one cannot hope for an elementary sharp characterisation. Instead, one finds that, in many cases, they can be described in terms of the filled-in Julia sets or connectedness loci corresponding to certain systems of complex polynomials [20]. Other connections with complex dynamical systems are investigated in [12,13]. See also [27] and [28] for a discussion of discreteness of certain 2-generator groups.

The dynamical set-up described above can be interpreted in terms of the action of the modular group, $PSL(2,\mathbf{Z})$, on triples of complex numbers. (Viewing $PSL(2,\mathbf{Z})$ as $\mathbf{Z}_2 * \mathbf{Z}_3$, this action is generated by the maps $[(x,y,z) \mapsto (y,x,xy-z)]$ and $[(x,y,z) \mapsto (y,z,x)]$.) Note that the quantity $\mu = x^2 + y^2 + z^2 - xyz$ is invariant under this action. (The case of Markoff triples thus corresponds to $\mu = 0$.) ¿From this perspective it is also interesting viewed as a real dynamical system — restricting to triples of real numbers. In the case where $\mu = 0$, there is not much to be said — we essentially get topological conjugates of the standard action of $PSL(2,\mathbf{Z})$ on (the upper-half space model of) the hyperbolic plane. However, if we vary the parameter μ, some interesting behaviour emerges. Some of this can be given geometrical interpretations. It also has applications to physics, for example to the spectrum of the discrete one-dimensional Schrödinger equation with quasiperiodic potential. Such applications are disussed, and some aspects of the dynamics are analysed, for example in [37] and [36].

As we have said, or aim here is to study something of Markoff's equation over the complex numbers, and its relation to 3-dimensional hyperbolic geometry. In the next section, we set out some of the results of the paper.

I am endebted to Colin Maclachlan for introducing me to the wonders of Markoff triples. Conversations with Greg McShane and John Roberts have also been helpful. I would like to thank Makoto Sakuma for his many useful comments on an earlier version of this paper.

1. Outline of results.

In this section, we introduce some terminology and notation, and describe some of our results.

By a *Markoff triple* we mean an ordered triple (x, y, z) of complex numbers satisfying the Markoff equation:

$$x^2 + y^2 + z^2 = xyz. \tag{$*$}$$

We can obviously obtain other Markoff triples by permuting the entries. Slightly less trivially, if (x, y, z) is a Markoff triple, then so are $(x, y, xy - z)$, $(x, xz - y, z)$ and $(yz - x, y, z)$. On repeating such substitutions, we generate an equivalence class of Markoff triples which has a natural "tree" structure. Such a structure will be referred to as a "Markoff map". We may give a formal definition of this term as follows.

Let Σ be a binary tree (a countably infinite simplicial tree, all of whose vertices have degree 3) properly embedded in the plane. By a *complementary* region of Σ, we mean the closure of a connected component of the complement. We write $V(\Sigma)$, $E(\Sigma)$ and Ω for the sets of vertices, edges and complementary regions of Σ respectively. Although we are only interested in the combinatorial structure of this set-up, it is natural to imagine Σ as being dual to the regular tessellation of the hyperbolic plane by ideal triangles. Note that the regions of Ω are in natural bijective correspondence to the ideal vertices of this tessellation. If we choose to put one of these vertices at ∞ in the upper half-space model, then we get an identification of Ω with the rationals $\mathbf{Q} \cup \{\infty\}$. In particular, we see that Ω carries a natural dense cyclic order.

Note that any vertex $v \in V(\Sigma)$ lies in the boundary of precisely three complementary regions, $X, Y, Z \in \Omega$, so that $X \cap Y \cap Z = \{v\}$. An edge of Σ meets four complementary regions X, Y, Z, W in such a way that $e = X \cap Y$ and so that $e \cap Z$ and $e \cap W$ are the endpoints of e. We shall use the notation $e \leftrightarrow (X, Y; Z, W)$ as a shorthand for the above statement; in other words, it is intended to relate the regions X, Y, Z, W to the edge e.

Definition : A map $\phi : \Omega \longrightarrow \mathbf{C}$ is a *Markoff map* if

(1) for all vertices $v \in V(\Sigma)$, the triple $(\phi(X), \phi(Y), \phi(Z))$ is a Markoff triple, where $X, Y, Z \in \Omega$ are the three regions meeting v; and

(2) If $e \in E(\Sigma)$, we have

$$xy = w + z \tag{$**$}$$

where $e \leftrightarrow (X, Y; Z, W)$ and $x = \phi(X)$, $y = \phi(Y)$, $z = \phi(Z)$ and $w = \phi(W)$.

The *trivial Markoff map*, $\underline{0}$, is the map which sends everything to 0.

We shall use Φ to denote the set of all Markoff maps.

Note that, in fact, if the edge relation $(**)$ is satisfied along all edges, then it suffices that the vertex relation $(*)$ be satisfied at a single vertex. In fact we may establish a bijective correspondence between Markoff maps and Markoff triples, by fixing three regions $X_0, Y_0, Z_0 \in \Omega$ which meet at some vertex v_0, and associating to a Markoff map ϕ the triple $(\phi(X_0), \phi(Y_0), \phi(Z_0))$. This process may be inverted by constructing a tree of Markoff triples as outlined in the introduction — given a triple (x, y, z) set $\phi(X_0) = x$, $\phi(Y_0) = y$, $\phi(Z_0) = z$, and extend over Ω as dictated by the edge relations. In this way we get an identification of Φ with the variety in $\mathbf{C}^3$ given by the Markoff equation. In particular, Φ gets a nice topology as a subset of $\mathbf{C}^3$.

An important example of Markoff map begins with the triple $(3, 3, 3)$. In fact, this is "essentially" the only non-trivial Markoff map whose image consists entirely of (rational) integers (Proposition 3.19). The "Markoff numbers", as arise in the study of the Markoff and Lagrange spectra [17], are the numbers which occur in this image. (Strictly speaking, we should divide everything by 3 to get the standard definition of Markoff number.) The well-known Markoff conjecture asserts that the only coincidences of Markoff numbers occuring in this picture arise from the obvious order 6 dihedral group of symmetries of the tree (or equivalently that the greatest element of a triple of Markoff numbers determines uniquely the other two).

Note that there is an action of the Klein-four group, $\mathbf{Z}_2^2$, on Φ obtained by changing two of the signs in a Markoff triple, for example $[(x, y, z) \mapsto (-x, -y, z)]$. (We get the same action, up to automorphism of $\mathbf{Z}_2^2$, no matter at which vertex, $v_0 \in V(\Sigma)$, we choose to perform this operation.) This action is free and properly discontinuous on $\Phi \setminus \{\underline{0}\}$.

Before going on to describe some results about Markoff maps, let's outline the connection with group representations.

Let $\Gamma \cong \mathbf{Z} * \mathbf{Z}$ be a free group on two generators, which we may identify as the fundamental group of the once-punctured torus, $\mathbf{T}$. We define an equivalence relation $\sim$ on Γ by $g \sim h$ if g is conjugate to h or to h^{-1}. We can thus identify $\Gamma/\sim$ with the set of homotopy classes of closed curves on $\mathbf{T}$.

Now, one can show that any (outer) automorphism of Γ is induced by an homeomorphism of $\mathbf{T}$. (This may be seen using the fact that the automorphism group of a free group is generated by Nielsen transformations [29].) In particular, there is a well-defined subset $\hat{\Omega}$ of $\Gamma/\sim$ corresponding to non-trivial non-peripheral simple closed curves on $\mathbf{T}$.

Now one can clearly choose a free generating set, $\{a, b\}$, for Γ so that the $\sim$-classes of a and b belong to $\hat{\Omega}$, and so that $[a, b]$ is a peripheral simple closed curve. ¿From the above observation about automorphisms of Γ we see that this must, in fact, be true of any free generating set. In particular the commutator of any pair of free generators belongs to a particular $\sim$-class in Ω.

It turns out that $\hat{\Omega}$ may be "naturally" identified with Ω. At least, the identification is natural once we have decided to identify a particular element of $\hat{\Omega}$ with some element of Ω. There are several ways to describe this correspondence, as we mention in Section 2.2. For the moment let us just note that it has following property. If $e \in E(\Sigma)$ and $e \leftrightarrow (X, Y; Z, W)$, then there is a free generating set $\{a, b\}$ for Γ so that the regions

$X, Y, Z, W \in \Omega$ correspond respectively to the $\sim$-classes of $a, b, ab, ab^{-1} \in \Gamma$.

Definition : A representation $\rho : \Gamma \longrightarrow SL(2, \mathbf{C})$ is *type-preserving* if $\operatorname{tr} \rho([a,b]) = -2$, for some (and hence any) free generating set $\{a, b\}$.

Given an type-preserving representation, we may define a Markoff map $\phi \in \Phi$ by $\phi(X) = \operatorname{tr} \rho(g)$ where $g \in \Gamma$ represents the simple closed curve corresponding to $X \in \Omega$. The edge and vertex relations follow from the trace identities in $SL(2, \mathbf{C})$:

$$\operatorname{tr} A \operatorname{tr} B = \operatorname{tr} AB + \operatorname{tr} AB^{-1}$$

$$2 + \operatorname{tr}[A, B] = (\operatorname{tr} A)^2 + (\operatorname{tr} B)^2 + (\operatorname{tr} AB)^2 - \operatorname{tr} A \operatorname{tr} B \operatorname{tr} AB.$$

Note that representations conjugate in $SL(2, \mathbf{C})$ give rise to the same Markoff map. Conversely, given any Markoff map, we can recover the type-preserving representation ρ up to conjugacy, for example using Jørgensen's normalisation, see Section 4. Thus we can identify Φ with the set of type-preserving representations into $SL(2, \mathbf{C})$.

Demanding that $\operatorname{tr} \rho([a,b]) = \pm 2$ would be the same as demanding that $\rho([a,b])$ be either parabolic or $\pm I$ (the identity matrix). It is natural also to insist that $\operatorname{tr} \rho([a,b])$ be negative since if it were equal to $+2$, then $\rho(a)$ and $\rho(b)$ would generate an elementary group (see Lemma 4.1), and thus not be of much interest. In the case of the trivial Markoff map, we get that $\rho([a,b]) = -I$. For every non-trivial Markoff map it is parabolic. This explains the term "type-preserving" — it sends peripheral loops to parabolics.

Note that is also makes sense to speak of type-preserving representations of Γ into $PSL(2, \mathbf{C})$, since the sign of the trace of a commutator is well-defined. Such representations are in natural bijective correspondence with elements of $\Phi / \mathbf{Z}_2^2$, where $\mathbf{Z}_2^2$ acts as described earlier.

We say that a Markoff map, ϕ, is *real* if $\phi(\Omega) \subseteq \mathbf{R}$. If ρ is a representation corresponding to a real map, then it can be conjugated so that its image lies in $SL(2, \mathbf{R})$. We say that ρ is *fuchsian*. Writing $\pi : SL(2, \mathbf{C}) \longrightarrow PSL(2, \mathbf{C})$ for the quotient map, we see that $\mathbf{H}^2 / \pi \circ \rho(\Gamma)$ gives a complete finite-area hyperbolic structure on the punctured torus, as described in the introduction. All such structures arise in this way.

More generally, we have the notion of a "quasifuchsian" representation $\rho : \Gamma \longrightarrow SL(2, \mathbf{C})$. This is one which is discrete and faithful, and for which $\pi \circ \rho(\Gamma)$ is geometrically finite without accidental parabolics. We shall elaborate on this definition in Section 4.

It is natural to ask which Markoff maps correspond to quasifuchsian groups. Theorem 4 below gives one way of answering this question, though it is not very satisfactory. One can conjecture a much simpler answer.

Let us state a few results about Markoff maps.

Given $\phi \in \Gamma$ and some $k \geq 0$, define $\Omega(k) = \Omega_\phi(k) \subseteq \Omega$ by

$$\Omega_\phi(k) = \{X \in \Omega \mid |\phi(X)| \leq k\}.$$

Theorem 1 : *If $\phi \in \Phi$ then*
(1) $\Omega_\phi(3)$ is nonempty, and
(2) for any $k \geq 2$ we have that $\bigcup \Omega_\phi(k)$ is connected (as a subset of the plane).

It's not hard to see that if $\Omega_\phi(k) = \Omega$ for some $k \geq 0$ then ϕ must be trivial. We shall see some stronger constraints on this set in Section 5.

A natural hypothesis to put on ϕ is that $\phi^{-1}([-2,2]) = \emptyset$. This means that the corresponding representation has no elliptics or accidental parabolics. Of particular interest are Markoff maps described by the following:

Theorem 2 : *Suppose $\phi \in \Phi$ is such that $\phi^{-1}([-2,2]) = \emptyset$, and $\Omega_\phi(2)$ is finite (possibly empty). Then $\log^+ |\phi|$ has Fibonacci growth.*

Here, $\log^+$ is defined by $\log^+(x) = \max\{0, \log(x)\}$. The term "Fibonacci growth" will be defined in Section 2.1. For the moment, we just note that it implies that $\sum_{X \in \Omega} |\phi(X)|^{-\beta}$ converges for any $\beta > 0$. In particular, $\Omega(k)$ is finite for all k and so $\phi(\Omega) \subseteq \mathbf{C}$ is discrete.

We shall write $\Phi_Q \subseteq \Phi$ for the set of ϕ satisfying the hypotheses of Theorem 2. We shall see (Theorem 3.16) that Φ_Q is an open subset of Φ.

For such ϕ, we have the following version of McShane's identity [32]:

Theorem 3 : *If $\phi \in \Phi_Q$, then $\sum_{X \in \Omega} h(\phi(X)) = \frac{1}{2}$.*

Here $h : \mathbf{C} \setminus [-2,2] \longrightarrow \mathbf{C}$ is defined by $h(x) = \frac{1}{2}(1 - \sqrt{1 - 4/x^2})$, where we adopt the convention that the real part of a square root is always non-negative. Thus, $\Re(h(x)) \leq \frac{1}{2}$. Note that $|h(x)| = O(|x|^{-2})$, and so Theorem 2 tells us that the series converges absolutely. The real case of McShane's identity was expressed in these terms in [6]. We say something of the geometric significance of this result in terms of quasifuchsian groups in Section 4. There is also a version of McShane's identity for once-punctured torus bundles fibring over the circle [7].

We have noted that non-trivial real Markoff maps correspond to fuchsian representations. We shall see that all such representations lie in Φ_Q. We suspect that:

Conjecture A : A Markoff map lies in Φ_Q if and only if corresponds to a quasifuchsian representation.

The "if" part of the conjecture is elementary. The "only if" part may be expressed in a number of different ways. Note that if $\phi \in \Phi_Q$, it would be sufficient to show that the corresponding representation were discrete and faithful. The fact that it is quasifuchsian then follows (see Section 4).

One of the alternative ways of expressing Conjecture A uses the fact that quasifuchsian space (the space of quasifuchsian representations into $PSL(2,\mathbf{C})$) is connected. Recall that $\mathbf{Z}_2^2$ acts freely properly discontinuously on $\Phi \setminus \{\underline{0}\}$ and note that $\underline{0} \notin \Phi_Q$. It is not hard to see that Φ_Q is invariant under this action. The quotient corresponds to representations to $PSL(2,\mathbf{C})$. Let $\hat{\Phi}_Q^0$ be the connected component of $\Phi_Q/\mathbf{Z}_2^2 \subseteq \Phi/\mathbf{Z}_2^2$ containing the

Markoff map given by the triple $(3,3,3)$ (and hence all real Markoff maps). Let $\Phi_Q^0 \subseteq \Phi$ be the preimage of $\hat{\Phi}_Q^0$ under the quotient map. Clearly $\Phi_Q^0 \subseteq \Phi_Q$.

Theorem 4 : *Those Markoff maps which correspond to quasifuchsian groups are precisely those which lie in* Φ_Q^0.

Since quasifuchsian space is contractible, it follows that Φ_Q^0 consists of four connected components which are permuted under the $\mathbf{Z}_2^2$ action.

We see that Conjecture A is equivalent to asserting that $\Phi_Q^0 = \Phi_Q$, or that $\Phi_Q/\mathbf{Z}_2^2$ is connected.

One can also express Conjecture A in purely geometric terms (Section 4). In this form it generalises to other topological types of surfaces (Conjecture B). It also gives rise to an apparently interesting question about representations to $PSL(2,\mathbf{R})$ (Question C).

The geometry of Φ_Q is also of interest. Some (rather crude) computer experiments show it to have an apparently "fractal" boundary — as one might expect from analogous investigations into 2-generator groups, for example [28,20,12,13].

The behaviour of maps ϕ in the complement of Φ_Q also seems worthy of investigation. Firstly, one would expect (given Conjecture A) those in the boundary of Φ_Q to correspond to discrete representations which are geometrically tame (in the sense of Thurston, see [43,5]) but not geometrically finite, or which are geometrically finite with accidental parabolics. In principle, they therefore provide an experimental testing ground for various conjectures relating to such representations.

We show in Section 5 that $\Phi \setminus \Phi_Q$ has non empty interior, as one might expect, and as computer experiments apparently demonstrate. In fact, the trivial Markoff map is an interior point.

The behaviour of individual Markoff maps lying in this interior seems difficult to analyse. We give some partial results in Section 5, and suggest some further lines of enquiry. One would expect such maps to correspond to non-discrete or non-faithful representations. We develop a little the theory of realisable laminations with respect to a non-discrete representation, and, in the case of a once-punctured torus, conjecturally relate this to the "large scale" behaviour of Markoff maps.

2. Binary trees and simple closed curves.

2.1. Binary trees.

Recall that Σ is a binary tree properly embedded in the plane. We have defined $V(\Sigma)$, $E(\Sigma)$ and Ω, respectively, as the sets of edges, vertices and complementary regions of Σ. Given $e \in E(\Sigma)$, we introduced the notation $e \leftrightarrow (X,Y;Z,W)$ to mean that $e = X \cap Y$ and that $e \cap Z$ and $e \cap W$ are the endpoints of e. We say that the regions X and Y are *adjacent* to the edge e.

A *directed edge*, $\vec{e}$, of Σ can be though of as an ordered pair of adjacent vertices, referred to as the *head* and *tail* of $\vec{e}$. Thus, we shall speak of $\vec{e}$ as being "directed towards" its head. We write $\vec{e} \leftrightarrow (X,Y;Z,W)$ if $e \leftrightarrow (X,Y;Z,W)$ and $e \cap W$ is the head of $\vec{e}$. We write $\vec{E}(\Sigma)$ for the set of all directed edges of Σ. We shall always use e to denote the (unoriented) edge in $E(\Sigma)$ underlying any $\vec{e} \in \vec{E}(\Sigma)$.

Given a subtree, $T \subseteq \Sigma$, we define the set $C(T) \subseteq \vec{E}(\Sigma)$ of directed edges by saying that $\vec{e} \in C(T)$ if and only if $e \cap T$ consists of a single point, that point being the head of $\vec{e}$. If T is finite, then so is $C(T)$. We shall say that a finite set C of directed edges is *circular* if it has the form $C(T)$ for some finite tree T. A trivial example of a circular set consists of the three directed edges whose heads all lie at some given vertex of Σ.

Suppose $\vec{e} \in \vec{E}(\Sigma)$. If we remove the interior of e from Σ, we are left with two disjoint subtrees, which we denote by $\Sigma^{\pm}(\vec{e})$, so that $e \cap \Sigma^{+}$ is the head of $\vec{e}$ and $e \cap \Sigma^{-}$ is its tail. Let $\Omega^{\pm} = \Omega^{\pm}(\vec{e}) \subseteq \Omega$ be the set of regions whose boundaries lie in $\Sigma^{\pm}$, and set $\Omega^{0} = \Omega^{0}(e) = \{X, Y\}$ where $e = X \cap Y$. We see that Ω can be written as a disjoint union: $\Omega = \Omega^{0} \sqcup \Omega^{+} \sqcup \Omega^{-}$. Let $\Omega^{0\pm}(\vec{e}) = \Omega^{0} \cup \Omega^{\pm}$. In other words, $\Omega^{0\pm}$ is the set of regions which meet $\Sigma^{\pm}$.

More generally, suppose $C = C(T) \subseteq \vec{E}(\Sigma)$ is a circular set. Let $\Omega^{0}(C) = \bigcup_{\vec{e} \in C} \Omega^{0}(e)$ be the (finite) set of regions meeting T in at least one edge. We see that Ω can be written as a disjoint union of $\Omega^{0}(C)$ and the sets $\Omega^{-}(\vec{e})$ as $\vec{e}$ varies in C.

Given a directed edge $\vec{e}$, we write $-\vec{e}$ for the same edge pointing in the opposite direction, i.e. we swap head and tail. Note that $\Omega^{+}(\vec{e}) = \Omega^{-}(-\vec{e})$ etc.

Given some $v \in V(\Sigma)$ and $X \in \Omega$ we write $d(v, X)$ for the distance in Σ from v to X; in other words, the number of edges in the shortest path joining v to X. If $\vec{e} \in \vec{E}(\Sigma)$, we define $d : \Omega^{0-}(\vec{e}) \longrightarrow \mathbf{N}$ by $d(X) = d(v, X)$ where v is the head of v. Thus, $d(X) = 0$ if and only if $X \in \Omega^{0}(e)$. Given any $Z \in \Omega^{-}$, there are precisely two regions $X, Y \in \Omega^{0-}$ meeting Z and satisfying $d(X) < d(Z)$ and $d(Y) < d(Z)$. In this case X, Y and Z all meet in a vertex. (Moreover either $d(X)$ or $d(Y)$ must equal $d(Z) - 1$, whereas the other is at most $d(Z) - 2$, provided $d(Z) \geq 2$.)

This observation allows us to use definition by induction on Ω^{0-}, starting with prescribed values on Ω^{0}. More formally, suppose S is a set and $B : S \times S \longrightarrow S$ is a binary operation. If we start with an S-valued function, f, defined on Ω^{0}, we may extend to a map $f : \Omega^{0-} \longrightarrow S$ as follows. Given $Z \in \Omega^{-}$, let $X, Y \in \Omega^{0-}$ be as in the previous paragraph, and such that the triple (X, Y, Z) is ordered consistently with the natural cyclic ordering on Ω. Now we can assume, inductively, that $f(X)$ and $f(Y)$ have already been defined, and we set $f(Z) = B(f(X), f(Y))$. In most of our examples, B is symmetric, i.e. $B(p,q) = B(q,p)$, and so the orientation of X, Y, Z is irrelevant. Note, for example that the function d itself can be defined in this manner: set $S = \mathbf{N}$, $B(p,q) = \max\{p,q\} + 1$ and start with d identically zero on Ω^{0}.

A more interesting example, again with $S = \mathbf{N}$, is obtained by setting $B(p,q) = p + q$, to obtain a map $F_{\vec{e}} : \Omega^{0-} \longrightarrow \mathbf{N}$, with $F_{\vec{e}}(X) = 1$ for each $X \in \Omega^{0}$. We can extend symmetrically to get a map $F_{e} : \Omega \longrightarrow \mathbf{N}$; i.e. we define $F_{e}(X) = F_{\vec{e}}(X)$ for $X \in \Omega^{0-}$ and $F_{e}(X) = F_{-\vec{e}}(X)$ for $X \in \Omega^{+}$.

The main reason for introducing the functions F_{e} is that they provide a means for measuring the growth rates of functions defined on subsets of Ω. The following lemma may be proven by induction on $d(X)$.

Lemma 2.1.1 : *Suppose* $\vec{e} \in \vec{E}(\Sigma)$ *and* $\Omega^0(e) = \{X_1, X_2\}$. *Suppose* $f : \Omega^{0-} \longrightarrow [0, \infty)$.
(1) Suppose f *satisfies* $f(Z) \leq f(X) + f(Y) + c$ *for some fixed constant* $c > 0$, *whenever* $X, Y, Z \in \Omega^{0-}(\vec{e})$ *are three regions meeting at a vertex, and satisfying* $d(X) < d(Z)$ *and* $d(Y) < d(Z)$. *Then* $f(X) \leq (M + c)F_e(X) - c$ *for all* $X \in \Omega^{0-}(\vec{e})$, *where* $M = \max\{f(X_1), f(X_2)\}$.
(2) Suppose f *satisfies* $f(Z) \geq f(X) + f(Y) - c$ *for some fixed constant* $0 < c < m = \min\{f(X_1), f(X_2)\}$, *whenever* X, Y, Z *are as in part (1). Then* $f(X) \geq (m - c)F_e(X) + c$ *for all* $X \in \Omega^{0-}(\vec{e})$. ♢

Corollary 2.1.2 : *Suppose* $f : \Omega \longrightarrow [0, \infty)$ *satisfies an inequality of the form* $f(Z) \leq f(X) + f(Y) + c$ *for some fixed constant* c, *whenever* $X, Y, Z \in \Omega$ *meet at a vertex. Given any edge* $e \in E(\Sigma)$, *there is a constant* $K > 0$, *such that* $f(X) \leq KF_e(X)$ *for all* $X \in \Omega$.

Proof : Apply Lemma 2.1.1(1) to get upper bounds for f on $\Omega^{0-}(\vec{e})$ and $\Omega^{0+}(\vec{e}) = \Omega^{0-}(-\vec{e})$. ♢

Clearly $f = F_{e'}$ satisfies the hypotheses of Corollary 2.1.2, with $c = 0$, for any edge $e' \in E(\Sigma)$. We deduce:

Proposition 2.1.3 : *Given* $e, e' \in E(\Sigma)$, *there is some* $K > 0$ *such that*

$$K^{-1}F_e(X) \leq F_{e'}(X) \leq KF_e(X)$$

for all $X \in \Omega$. ♢

This leads to the following definitions:

Definition : Suppose $f : \Omega \longrightarrow [0, \infty)$, and $\Omega' \subseteq \Omega$.
We say that f has an *upper Fibonacci bound* on Ω' if there is some constant $K > 0$ such that $f(X) \leq KF_e(X)$ for all $X \in \Omega'$.
We say that f has a *lower Fibonacci bound* on Ω' if there is some constant $k > 0$ such that $f(X) \geq kF_e(X)$ for all but finitely many $X \in \Omega'$.
We say that f has *Fibonacci growth* on Ω' if it has both upper and lower Fibonacci bounds on Ω'.
If f has Fibonacci growth on all of Ω, then we shall say simply that it has *Fibonacci growth*.

¿From Proposition 2.1.3, we see that it doesn't matter which edge $e \in E(\Sigma)$ we choose for this purpose.

Note that if Ω' is the union of a finite set of subsets $\Omega_1, \ldots, \Omega_m \subseteq \Omega'$, then f has an upper (lower) Fibonacci bound on Ω' if and only if it has an upper (lower) Fibonacci bound on each Ω_i.

We are principally interested in lower Fibonacci bounds, which will give the convergence of certain series. Let $F = F_e$ for some edge $e \in E(\Sigma)$. If $X, Y \in \Omega$ meet along some edge e', then we see (by induction) that $F(X)$ and $F(Y)$ are coprime. In fact the path in Σ from e' to e can be determined from $F(X)$ and $F(Y)$ by a process of continued subtraction. It is not hard to see from this that F takes any given value $n \in \mathbf{N}$ on precisely $2\phi(n)$ regions of Ω, where ϕ is the Euler function. Thus, for any real $s > 2$, we have

$$\sum_{X \in \Omega} F(X)^{-s} = \frac{2\zeta(s-1)}{\zeta(s)}$$

where ζ is the Riemann zeta function. The right hand side is the result of summing the Dirichlet series $2\sum_{n=1}^{\infty} \phi(n) n^{-s}$.

Although it is amusing that one can actually write down the answer, we are only really interested in the fact that this series converges. We deduce:

Proposition 2.1.4 : *If $f : \Omega \longrightarrow [0, \infty)$ has a lower Fibonacci bound, then $\sum_{X \in \Omega} f(X)^{-s}$ converges for all $s > 2$. (We may have to exclude a finite subset of Ω on which f takes the value 0.)* $\diamondsuit$

Before moving on, let us note that Ω has a natural "3-colouring", i.e. there a partition of Ω into three disjoint subsets, $\{\Omega^1, \Omega^2, \Omega^3\}$ such that any two regions meeting along an edge of Σ lie in different elements of the partition. This partition is unique (up to the permuting the "colours" 1, 2 and 3). All three colours meet at a vertex. We shall return to this subject later.

2.2. Punctured tori.

Recall that $\Gamma = \pi_1(\mathbf{T})$ is the free group on two generators. As in the introduction, we identify the set of homotopy classes of closed curves on $\mathbf{T}$ with $\Gamma/\sim$, where $g \sim h$ if g is conjugate to h or to h^{-1}. Given a particular free generating set, $\{a, b\}$ for Γ, we define the *cyclic word length*, $W(\gamma)$ of $\gamma \in \Gamma/\sim$ as the length of a cyclically reduced word in a, b, a^{-1}, b^{-1} representing γ. Such a word is unique up to cyclic permutations and formal inverses. It has the minimal length among all words representing γ. Let $\hat{\Omega} \subseteq \Gamma/\sim$ be the subset of simple closed curves. An element of $\hat{\Omega}$ may be written as a (cyclically reduced) word either in a and b or in a and b^{-1}. (This follows from the fact that one can represent such a curve as a closed geodesic on a euclidean torus punctured at one point.)

As we have mentioned there is a natural identification of Ω and $\hat{\Omega}$. There are many ways of describing this. For example, we saw, using the upper-half space model, that Ω can be identified with $\mathbf{Q} \cup \{\infty\}$. An identification of $\hat{\Omega}$ with $\mathbf{Q} \cup \{\infty\}$ may be obtained as via its description as the set of rational laminations on the torus. Alternatively, identify any fixed simple closed curve with ∞, and extend this identification using the natural actions of $PSL(2, \mathbf{Z})$ on the sets $\mathbf{Q} \cup \{\infty\}$ and Ω, given that $PSL(2, \mathbf{Z})$ is the mapping class group of $\mathbf{T}$.

A more natural way of carrying out this identifcation is via Harer's complex [24]. (See also [9] for an account via hyberbolic geometry.) This describes a trianguation of the Teichmüller space of a punctured surface, in this case the hyperbolic plane. The (ideal) vertices of this triangulation correspond, in the case of a punctured torus, to simple closed curves.

We can write out explicit representatives in Γ of elements of $\hat{\Omega}$ corresponding to given elements of Ω as follows. Choose some $\vec{e} \in \vec{E}(\Sigma)$. Let $e \leftrightarrow (X_1, X_2; X_3, X_4)$ (recalling the notation introduced in Section 2.1). The elements of $\hat{\Omega}$ corresponding to X_1 and X_2 are represented, respectively, by a pair of free generators, a and b for Γ. Without loss of generality, we can suppose that X_3 and X_4 are represented by ab and ab^{-1}. We can now inductively give representatives for all regions in Ω^{0-} by a process of concatenation: In the notation of Section 2.1, let S be the set of all words in a and b, and given $w_1, w_2 \in S$ define $B(w_1, w_2)$ to be the concatenation of w_1 and w_2. We may do a similar thing for $\Omega^{0+}(\vec{e}) = \Omega^{0-}(-\vec{e})$ replacing b by b^{-1}. Note that all the words arising in this way are cyclically reduced. We deduce:

Lemma 2.2.1 : *Suppose a and b are a pair of free generators for Γ corresponding to regions X_1 and X_2. Let e be the edge $X_1 \cap X_2$. If $\gamma \in \hat{\Omega}$ corresponds to $X \in \Omega$, then $W(\gamma) = F_e(X)$.* $\diamondsuit$

Before we leave this section, we note that the 3-colouring of Ω referred to at the end of Section 2.1 corresponds to partitioning $\hat{\Omega}$ by the three non-trivial $\mathbf{Z}_2$-homology classes of closed curves.

3. Markoff maps.

In this we develop some basic results about Markoff maps, including proofs of Theorems 1–3.

Recall the definition of a Markoff map, $\phi : \Omega \longrightarrow \mathbf{C}$, from Section 1, and the notation $\vec{e} \leftrightarrow (X, Y; Z, W)$ introduced there.

Throughout this section, we shall fix on one Markoff map, ϕ, which we assume to be non-trivial (i.e. not identically zero). We shall adopt the following

Convention : We use upper case latin letters for elements of Ω, and the corresponding lower case letters for the values assigned to them by ϕ; i.e. $x = \phi(X)$, $y = \phi(Y)$ etc.

We begin by showing that ϕ has an upper Fibonacci bound. If three regions $X, Y, Z \in \Omega$ meet around some vertex, then solving the Markoff equation $(*)$ in z, gives us $z = \frac{1}{2}(xy \pm \sqrt{x^2y^2 - 4(x^2 + y^2)})$. ¿From this we easily deduce that $\log^+ |z| \leq \log^+ |x| + \log^+ |y| + \log 2$, where $\log^+ \zeta = \max\{0, \log \zeta\}$. Applying Corollary 2.1.2, we find:

Proposition 3.1 : *If ϕ is a Markoff map, then $\log^+ |\phi|$ has an upper Fibonacci bound on Ω.* $\diamondsuit$

Note that, since ϕ is non-trivial, it is not possible for two adjacent regions to be assigned the value 0. In fact it will simplify our exposition a little if we only deal with the cases where all values are non-zero; i.e. we are effectively assuming that $\phi^{-1}(0) = \emptyset$. However, we will not use this assumption in any essential way, except where other hypotheses imply this anyway. The details of the cases where some of the values are 0 are easily filled in.

This convention will allow us to write the vertex and edge relations, $(*)$ and $(**)$, in a form that is often convenient:

$$\frac{x}{yz}+\frac{y}{zx}+\frac{z}{xy}=1$$
$$\frac{z}{xy}+\frac{w}{xy}=1.$$

We may use ϕ to define a map $\alpha=\alpha_\phi : E(\Sigma)\longrightarrow \vec{E}(\Sigma)$ which assigns to each undirected edge, e, a particular direction or "arrow", $\vec{e}=\alpha(e)$, with underlying edge e. Suppose $e\leftrightarrow (X,Y;Z,W)$. If $|z|>|w|$ then we the arrow on e points from Z to W. In other words, $\alpha(e)\leftrightarrow (X,Y;Z,W)$. Note that the statement $|z|>|w|$ is equivalent to $\Re\left(\frac{z}{xy}\right)>\frac{1}{2}$. In particular, it implies that $2|z|>|xy|$. In fact we have

$$\frac{z}{xy}=\frac{1}{2}\left(1+\sqrt{1-4\left(\frac{1}{x^2}+\frac{1}{y^2}\right)}\right)$$

and

$$\frac{w}{xy}=\frac{1}{2}\left(1-\sqrt{1-4\left(\frac{1}{x^2}+\frac{1}{y^2}\right)}\right),$$

where, as before, we adopt the convention that the real part of the square root of a complex number is non-negative. If $|z|<|w|$ we put an arrow on e pointing from W to Z. If it happens that $|z|=|w|$ then we choose $\alpha(e)$ arbitrarily.

By a *sink* we mean a vertex v such that for each of the incident edges e, the arrow $\alpha(e)$ points towards v. A *source* is defined similarly, with the three arrows pointing away from v.

Lemma 3.2 :
(1) There are no sources.
(2) If the three regions $X,Y,Z\in\Omega$ intersect at a sink, then $\min\{|x|,|y|,|z|\}\leq 3$.
(3) Suppose $X,Y,Z\in\Omega$ meet a vertex $v\in V(\Sigma)$, and that the arrows on the edges $X\cap Y$ and $X\cap Z$ both point away from v. Then $|x|\leq 2$.

Proof : In each case we assume that $v\in V(\Sigma)$ is the intersection of the regions $X,Y,Z\in\Omega$.
(1) If v were a source, then we would have $1=\Re\left(\frac{z}{xy}\right)+\Re\left(\frac{y}{zx}\right)+\Re\left(\frac{x}{yz}\right)\geq\frac{1}{2}+\frac{1}{2}+\frac{1}{2}=\frac{3}{2}$.
(2) Suppose v is a sink. Set $p=\Re\left(\frac{z}{xy}\right)$, $q=\Re\left(\frac{y}{zx}\right)$ and $r=\Re\left(\frac{x}{yz}\right)$. Thus $p+q+r=1$, and $p,q,r\leq\frac{1}{2}$. Without loss of generality, $r\leq q\leq p$. It follows that $pq\geq\frac{1}{9}$, for otherwise $1=p+q+r\leq p+2q<p+\frac{2}{9p}\leq 1$. Thus $|z|^2\leq\frac{1}{pq}\leq 9$ and so $|z|\leq 3$ as required.
(3) In this case, $|xy|\leq 2|z|$ and $|xz|\leq 2|y|$ and so $|x|\leq 2$. $\diamondsuit$

Proof of Theorem 1(2) : Suppose $k \geq 2$. Recall the definition $\Omega(k) = \{X \in \Omega \mid |\phi(X)| \leq k\}$. We want to show that $\bigcup \Omega(k)$ is connected.

Suppose it is not connected. Then we can find a path in Σ which joins two distinct components of $\bigcup \Omega(k)$. Let β be such a path of minimal length. Thus β intersects $\bigcup \Omega(k)$ precisely in its two endpoints. Now β consists of a sequence, $e_1, \ldots, e_p$, of edges of Σ. If a region, X, contains one of these edges in its boundary, then we must have $|x| > k$.

Case (1) : $p = 1$.

Let $e_1 \leftrightarrow (X, Y; Z, W)$, so that $Z, W \in \Omega(k)$. Then $k^2 < |xy| \leq |z| + |w| \leq 2k$, implying that $k < 2$.

Case (2) : $p > 1$.

Let $e_1 \leftrightarrow (X, Y; Z, W)$ with $Z \in \Omega(k)$. Since $W \notin \Omega(k)$, we have $|z| < |w|$, and so the arrow on e_1 points towards Z. Similarly the arrow on e_p points towards the other endpoint of β. It follows that there must be a consecutive pair of edges e_i and e_{i+1}, such that the arrows on them point away from their common endpoint. By Lemma 3.2(3), there is a region X' containing both these edges in its boundary, for which $|\phi(X')| \leq 2$, contradicting $X' \notin \Omega(k)$.

This concludes the proof of Theorem 1(2). ♢

Suppose $X \in \Omega$. Then ∂X is a bi-infinite path consisting of a sequence of edges of the form $X \cap Y_n$, where $(Y_n)_{n \in \mathbf{Z}}$ is a bi-infinite sequence of complementary regions. Let $x = \lambda + \lambda^{-1}$ where $|\lambda| \geq 1$. If $x = 2$, then the vertex relation tells us that $y_{n+1} = y_n \pm 2i$, and the edge relation tells us that the $\pm$ sign is constant in n. Similarly, if $x = -2$, then $y_{n+1} = -y_n \pm 2i$, though this time, the $\pm$ sign alternates in n. If $x \notin \{-2, 2\}$ then there are constants $A, B \in \mathbf{C}$ with $AB = x^2/(x^2 - 4)$ such that $y_n = A\lambda^n + B\lambda^{-n}$. Note that $|\lambda| = 1$ if and only if $x \in [-2, 2] \subseteq \mathbf{R}$. We conclude:

Lemma 3.3 :

(1) If $x \notin [-2, 2]$, then $|y_n|$ grows exponentially as $n \to \infty$ and as $n \to -\infty$.

(2) If $x \in (-2, 2)$ then $|y_n|$ remains bounded.

(3) If $x = 2$, then there is some $z \in \mathbf{C}$ such that either $y_n = z + 2ni$ for all n, or $y_n = z - 2ni$ for all n.

(4) If $x = -2$, then there is some $z \in \mathbf{C}$ such that either $y_n = (-1)^n(z + 2ni)$ for all n, or $y_n = (-1)^n(z - 2ni)$ for all n. ♢

Lemma 3.4 : *Suppose β is an infinite ray in Σ consisting of a sequence, $(e_n)_{n \in \mathbf{N}}$ of edges of Σ. Suppose that the arrow on each e_n is directed towards e_{n+1}. Then β meets at least one region X for which $|\phi(X)| < 2$.*

It follows that there must, in fact, be infinitely many such regions. For some applications (such as Theorem 1(1)), we only need a weaker conclusion, (for example, that $|\phi(X)| \leq 3$). For this one can exit from the proof at an earlier stage.

Proof : We shall write $\vec{e}_n = \alpha(e_n)$ for the edge e_n directed towards e_{n+1}. We may partition the set of adjacent regions (those meeting β in at least one edge) into three subsequences, $(X_i^j)_{i\in\mathbf{N}}$, with $j \in \{1,2,3\}$, corresponding to the 3-colouring of Ω described at the end of Section 2.1. If $\vec{e}_n \leftrightarrow (X,Y;Z,W)$, then $|w| \leq |z|$, and Z and W are consecutive elements in the same subsequence. It follows that for each j, the sequence $(|x_i^j|)_{i\in\mathbf{N}}$ is monotonically non-increasing.

Let us assume, for contradiction, that $|x_i^j| \geq 2$ for all i and j.

Now if $\vec{e} = \vec{e}_n$ for some n, and $\vec{e} \leftrightarrow (X,Y;W,Z)$, we can suppose, without loss of generality, that $X \cap Z \subseteq \beta$, so that the arrow on the edge $X \cap Z$ points away from the vertex $X \cap Y \cap Z$.

¿From the monotonicity described above, we can find such an edge $\vec{e}$, so that $|z| \leq |w| \leq |z| + \eta$ for arbitrarily small $\eta > 0$. We see that $|xz| \leq 2|y|$ and $|xy| \leq |z| + |w| \leq 2|z| + \eta$. Thus $|x|^2 \leq 4 + 2\eta/|z| \leq 4 + \eta$ since $|z| \geq 2$. It follows that we can assume that $|x|$ is arbitrarily close to 2.

(This is good enough for many purposes. If we want to show we can find a region so that the norm is less than 2 then we have to plough on.)

Continuing with the same line of argument, using the monotonicity of the three sequences, we can now assume that in the above set-up, $|z|$ and $|w|$ are both arbitrarily close to 2. Since $4 \leq |xy| \leq |z| + |w| \simeq 4$, we have that $|x|$ and $|y|$ are also close to 2. Since X, Y and Z represent all three colours, we can henceforth assume that all numbers have norm close to 2.

Since $\frac{z}{xy} + \frac{w}{xy} = 1$, it follows that $\frac{z}{xy} \simeq \frac{1}{2}$. By a similar argument, applied to the the next edge of β, (i.e. $X \cap Z$) we find that $\frac{y}{zx} \simeq \frac{1}{2}$. Thus $x \simeq \pm 2$. This is true of any region X which contains two consecutive edges e_n and e_{n+1} in its boundary for sufficiently large n.

¿From Lemma 3.2, and the previous discussion, we see that β cannot eventually lie in the boundary of any one region. (In other words each of the three subsequences of regions is indeed infinite.)

It follows that there must be an edge (indeed infinitely many edges), $\vec{e}$ in β with $\vec{e} \leftrightarrow (X,Y;W,Z)$ so that $Y \cap W \subseteq \beta$ and $X \cap Z \subseteq \beta$. Thus, $x \simeq \pm 2$ and $y \simeq \pm 2$. Also $|w| \simeq |z| \simeq 2$. Since $z + w = xy \simeq \pm 4$, we see that $z \simeq w \simeq \pm 2$. Thus the vertex relation, $x^2 + y^2 + z^2 = xyz$, is violated.

This finally contradicts the assumption that all numbers have norm at least 2, thus proving Lemma 3.4. ◊

Proof of Theorem 1(1) : If $\Omega(2) = \emptyset$, then Lemma 3.4 tells us that there must be a sink. We now apply Lemma 3.2(2). This proves Theorem 1(1). ◊

We have seen (Proposition 3.1) that $\log^+ |\phi|$ has an upper Fibonacci bound on Ω. We go on to consider criteria for it to have an lower Fibonacci bound on certain branches of Σ.

Suppose $\vec{e} \in \vec{E}(\Sigma)$ is such that $\Omega^0(e) \cap \Omega(2) = \emptyset$. Since $\bigcup \Omega(2)$ is connected (Theorem 1(2)), we must have either $\Omega(2) \subseteq \Omega^+(\vec{e})$ or $\Omega(2) \subseteq \Omega^-(\vec{e})$. (Possibly $\Omega(2) = \emptyset$.) If $\Omega(2) \subseteq \Omega^-(\vec{e})$ and $\Omega(2) \neq \emptyset$, then $\alpha(e) = -\vec{e}$. (This follows from Lemma 3.2(3), by a similar argument to that for Theorem 1(2) — consider the path β joining e to the nearest point of $\bigcup \Omega(2)$.) Taking the contrapositive, we have:

Lemma 3.5 : *Suppose $\vec{e} \in \vec{E}(\Sigma)$ is such that $\alpha(e) = \vec{e}$ and $\Omega^0(e) \cap \Omega(2) = \emptyset$. Then $\Omega^{0-}(\vec{e}) \cap \Omega(2) = \emptyset$. Moreover the arrow on each edge of Σ^- is directed towards e.*

Proof : ¿From the preceding discussion. The last statement follows by applying Lemma 3.2(3). ♢

Corollary 3.6 : *With the hypotheses of Lemma 3.5, write $f(X) = \log |\phi(X)|$ for any $X \in \Omega^{0-}(\vec{e})$. Let $m = \min\{f(X) \mid X \in \Omega^0(e)\} > \log 2$. Then $f(X) \geq (m - \log 2)F_e(X)$ for all $X \in \Omega^{0-}(\vec{e})$.*

Proof : Suppose $X, Y, Z \in \Omega^{0-}$ meet at a vertex, and $d(X) < d(Z)$ and $d(Y) < d(Z)$ as in Lemma 2.1.1. Now the arrow on $X \cap Y$ points away from Z and so $2|z| \geq |xy|$. Thus $f(Z) \geq f(X) + f(Y) - \log 2$. Apply Lemma 2.1.1(2). ♢

Corollary 3.7 : *If $\Omega(2) = \emptyset$, there is a unique sink, and $\log^+ |\phi|$ has Fibonacci growth.*

Proof : The existence of a sink, v, comes from Lemma 3.4. Its uniqueness comes from Lemma 3.2(3), which in this case shows that there is at most one arrow with its tail at any given vertex. By Proposition 3.1, $\log^+ |\phi|$ has an upper Fibonacci bound on Ω. To get a lower bound, apply Corollary 3.6 to the three edges incident on v.

This proves Corollary 3.7. ♢

To deal with Theorem 2, we shall want to expand somewhat on Corollary 3.6, to consider branches of Σ where at most one region has norm greater than 2. Let us consider the following set-up.

Suppose $X_0 \in \Omega$ and β is an infinite ray lying in the boundary of X_0 consisting of a sequence, $(\vec{e}_i)_{i=0}^{\infty}$, of directed edges so that each $\vec{e}_i$ is directed away from $\vec{e}_{i+1}$. For $i \geq 1$, let v_i be the vertex incident on both e_i and e_{i-1}, and let $\vec{\epsilon}_i$ be the third edge incident on v_i (distinct from e_i and e_{i-1}), and directed towards v_i. Thus $\Omega^{0-}(\vec{e}_0) = \{X_0\} \cup \bigcup_{i=1}^{\infty} \Omega^{0-}(\vec{\epsilon}_i)$. It is easy to see (using Lemma 2.1.1(2)) that a map $f : \Omega^{0-}(\vec{e}_0) \longrightarrow [0, \infty)$ has a lower Fibonacci bound on $\Omega^{0-}(\vec{e}_0)$ if and only if there is some constant $k > 0$ such that for all $n \geq 1$ and for all $X \in \Omega^{0-}(\vec{\epsilon}_n)$, we have $f(X) \geq knF_{\vec{\epsilon}_n}(X)$. As a consequence, we see:

Lemma 3.8 : *Suppose X_0, $\vec{e}_0$ and $(\vec{\epsilon}_n)_{n=1}^{\infty}$ are as described above. Suppose $x_0 \notin [-2, 2]$, and that $\Omega^{0-}(\vec{e}_0) \cap \Omega(2) \subseteq \{X_0\}$. Suppose also that $\vec{\epsilon}_n = \alpha(\epsilon_n)$ for all $n \geq 1$ (which is automatically satisfied if $|x_0| \leq 2$). Then $\log^+ |\phi|$ has a lower Fibonacci bound on $\Omega^{0-}(\vec{e}_0)$.*

Proof : For $n \geq 0$, let e_n be as described in the paragraph before Lemma 3.8, and let $\Omega^0(e_n) = \{X_0, Y_n\}$. Thus $\Omega^0(\epsilon_n) = \{Y_{n-1}, Y_n\}$ for all $n \geq 1$. By Lemma 3.3, $|y_n|$ grows exponentially as $n \to \infty$, and so $\log |y_n|$ is bounded below by some increasing linear function of n, which we can take to have the form $[n \mapsto cn]$ for some $c > 0$. The result now follows by applying Corollary 3.6, and the discussion preceding the statement of the lemma. $\diamondsuit$

Proof of Theorem 2 : We are assuming that $\phi^{-1}[-2,2] = \emptyset$ and that $\Omega(2)$ is finite. By Corollary 3.7, we can suppose that $\Omega(2) \neq \emptyset$. By Proposition 3.1, it suffices to show that $\log^+ |\phi|$ has a lower Fibonacci bound on Ω.

Let $E' \subseteq E(\Sigma)$ be the (finite) set of edges, e, such that $\Omega^0(e) \subseteq \Omega(2)$. If $E' \neq \emptyset$, let $T \subseteq E(\Sigma)$ be the tree spanned by E'. Since $\bigcup \Omega(2)$ is connected, it is easy to see that each vertex of T lies in at least one region of $\Omega(2)$. If $E' = \emptyset$, then, again since $\bigcup \Omega(2)$ is connected, it must be that $\Omega(2)$ consists of a single region, X_0. In this case, take T to be any vertex in the boundary of X_0. Now, let $C = C(T)$ be the circular set of directed edges given by T. Note that $\Omega = \bigcup_{\vec{e} \in C} \Omega^{0-}(\vec{e})$.

Suppose $\vec{e} \in C$. If $\Omega^0(e) \cap \Omega(2) = \emptyset$, then the hypotheses of Corollary 3.6 are satisfied. (Note that the head of $\vec{e}$ lies in T, and is thus in the boundary of some region of $\Omega(2)$.) It follows that $\log |\phi|$ has a lower Fibonacci bound on $\Omega^{0-}(\vec{e})$. If $\Omega^0(e) \cap \Omega(2)$ consists of a single region, X_0, then the hypotheses of Lemma 3.8 are satisfied (with $\vec{e} = \vec{e}_0$), and so again, $\log^+ |\phi|$ has a lower Fibonacci bound on $\Omega^{0-}(\vec{e})$. This concludes the proof of Theorem 2. $\diamondsuit$

We write Φ_Q for the set of Markoff maps satisfying the hypotheses of Theorem 2. If $\phi \in \Phi_Q$, then by Proposition 2.1.4, $\sum_{X \in \Omega} (\log^+ |\phi(X)|)^{-s}$ converges for $s > 2$ (if we exclude those X for which $|\phi(X)| \leq 1$). In particular, we see that $\sum_{X \in \Omega} |\phi(X)|^{-\beta}$ converges for all $\beta > 0$. Thus $\Omega(k)$ is finite for all $k \geq 0$. In other words, ϕ is at most finite-to-one and has discrete image.

Note that the argument of Theorem 2 can be applied to a single branch of the tree Σ.

Proposition 3.9 : *Suppose $\vec{e}$ is a directed edge of Σ such that $\Omega^-(\vec{e}) \cap \Omega(2)$ is finite and $\Omega^-(\vec{e}) \cap \phi^{-1}[-2,2] = \emptyset$. Then $\log^+ |\phi|$ has a lower Fibonacci bound on $\Omega^-(\vec{e})$.*

Proof : Let $T \subseteq \Sigma$ be the (possibly infinite) tree constructed as in the proof of Theorem 2. Let T' the tree spanned by $T' \cap \Sigma^-$ and $e \cap \Sigma^-$. Thus T' is finite, and $-\vec{e} \in C(T')$. As before, we see that $\log^+ |\phi|$ has a lower Fibonacci bound on $\Omega^-(\vec{\epsilon})$ for all $\vec{\epsilon} \in C(T') \setminus \{-\vec{e}\}$.

This proves Proposition 3.9. $\diamondsuit$

Note that (by Lemma 3.5), Proposition 3.9 applies in the case where $\alpha(e) = \vec{e}$ and $\Omega^0(e) \cap \Omega(2) = \emptyset$. Note also that we may conclude that $\sum_{X \in \Omega^-(\vec{e})} |\phi(X)|^{-\beta}$ converges for all $\beta > 0$.

Elaborating on Lemma 3.2, we see that in the case when $x \notin [-2,2]$, so that $|\lambda| > 1$, the sequence $|y_n| = |A\lambda^n + B\lambda^{-n}|$ is monotonic for sufficiently large and sufficiently small n. We conclude:

Lemma 3.10 : *If $X \in \Omega$ and $x = \phi(X) \notin [-2,2]$, then there is a nonempty finite subarc $J(X) \subseteq \partial X$ (i.e. $J(X)$ is the union of finitely many edges or possibly a single vertex of*

∂X) with the property that if e is any edge in ∂X not lying in $J(X)$, then the arrow on e points towards $J(X)$. Moreover, we can assume that $Y \cap X \subseteq J(X)$ for all $Y \in \Omega(2)$. $\diamondsuit$

In the case where $\phi(X) \in [-2, 2]$, we shall set $J(X) = \partial X$. Given $\phi \in \Phi$ with $\Omega(2) \neq \emptyset$, we define $T_0 = \bigcup_{X \in \Omega(2)} J(X)$.

Lemma 3.11 : *T_0 is connected, and the arrow on each edge not in T_0 points towards T_0.*

Proof : This is based on two observations. First, note that if T' is some connected component of T_0, then the arrow on every edge of $C(T')$ points towards T', i.e. $\alpha(e) = \vec{e}$ for all $\vec{e} \in C(T')$. The second observation is that if $v \in V(\Sigma)$ is a vertex not in T_0, then there is at most one arrow leaving v; for if v were the tail of both $\alpha(e_1)$ and $\alpha(e_2)$ where $e_1, e_2 \in E(\Sigma)$, then by Lemma 3.2(3), e_1 and e_2 would both lie in the boundary of some region $X \in \Omega(2)$, and so $v \in J(X)$. The fact that T_0 is connected now follows as in the proof if Theorem 1(2), by considering an arc in Σ which meets T_0 precisely in its two endpoints. The fact that if $e \in E(\Sigma)$ does not lie in T_0 then $\alpha(e)$ points to T_0 follows similarly. $\diamondsuit$

Corollary 3.12 : *If $\phi \in \Phi_Q$, then there is a finite subtree $T_0 \subseteq \Sigma$ with the property that the arrow on each edge not in T_0 points towards T_0.*

Proof : If $\Omega(2) \neq \emptyset$ then the tree T_0 given by Lemma 3.10 is finite. If $\Omega(2) = \emptyset$, then applying Corollary 3.7, we can take to T_0 to be a single vertex.

This proves Corollary 3.12. $\diamondsuit$

Note that the conclusion of Corollary 3.12 may be rephrased as follows. Suppose $T \subseteq \Sigma$ is any finite tree containing T_0, then $\vec{e} = \alpha(e)$ for all $\vec{e}$ in the circular set $C(T)$.

We are now in a position to give a proof of McShane's identity for elements of Φ_Q.

Suppose $\vec{e} \in \vec{E}(\Sigma)$ with $e = X \cap Y$, and the head of $\vec{e}$ at $e \cap Z$. Set $\psi(\vec{e}) = z/xy$. Exactly as in [6], we see that if C is a circular set of directed edges, then $\sum_{\vec{e} \in C} \psi(\vec{e}) = 1$. Also, if $\vec{e} = \alpha(e)$, then $\psi(\vec{e}) = h(x, y)$ where $h(x, y) = \frac{1}{2}\left(1 - \sqrt{1 - 4\left(\frac{1}{x^2} + \frac{1}{y^2}\right)}\right)$. (Recall that the real part of a square root is non-negative.) Note that $h(x, y) \to h(x)$ as $y \to \infty$ for any fixed x.

Proof of Theorem 3 : The proof is essentially the same as that given in [6]. To make the analysis work, we need to observe that $h(x, y) = O(|x|^{-2} + |y|^{-2})$, $h(x) = O(|x|^{-2})$ and that $\sum_{X \in \Omega} |\phi(X)|^{-2}$ converges (Theorem 2). (This replaces the fact that $\sum_{X \in \Omega} h(\phi(X)) \leq \frac{1}{2}$ in [6].)

We now choose an exhaustion of Σ by suitable finite subtrees T_n as in [6]. Note that, provided $T_n \supseteq T$, we have $\sum\{h(x, y) \mid e = X \cap Y, \vec{e} \in C(T_n)\} = 1$. We complete the proof as in the real case. $\diamondsuit$

Again, there is a version of Theorem 3 applicable to a single branch of the tree Σ. This is quoted in [7].

Proposition 3.13 : *Suppose $\vec{e}$ is a directed edge of Σ such that $\Omega^-(\vec{e}) \cap \Omega(2)$ is finite and $\Omega^{0-}(\vec{e}) \cap \phi^{-1}[-2,2] = \emptyset$. Then,*

$$\psi(\vec{e}) = \sum_{X \in \Omega^0(e)} h(\phi(X)) + 2 \sum_{X \in \Omega^-(\vec{e})} h(\phi(X)).$$

Proof : By Proposition 3.9, we know that $\sum_{X \in \Omega^-(\vec{e})} |h(\phi(X)|^{-2}$ converges.

We apply the arguments of [6], this time taking an exhaustion of Σ^- by a sequence of finite subtrees $(T_n)_{n \in \mathbf{N}}$. We can assume that each T_n meets e in its tail, so $-\vec{e} \in C(T_n)$. Let $C'_n = C(T_n) \setminus \{-\vec{e}\}$, so that $\psi(\vec{e}) = \sum_{\vec{\epsilon} \in C'_n} \psi(\vec{\epsilon})$. Letting $n \to \infty$ we replace the right hand side by a a sum of quantities $h(\phi(X))$ where X ranges over an appropriate subset of regions of Ω. In this sum, each region of $\Omega^-(\vec{e})$ gets counted twice, while the two regions of $\Omega^0(e)$ each get counted once.

This proves Proposition 3.13. ◊

The next objective is to show that Φ_Q is an open subset of Φ. This works essentially because one can recognise that an Markoff map belongs to Φ_Q from a finite amount of information. One way to do this is to give an explicit descriptions of trees satisfying the conclusion of Corollary 3.12.

To begin with, let us elaborate further on Lemma 3.3. Suppose $x = \lambda + \lambda^{-1}$ with $|\lambda| > 1$, and $y_n = A\lambda^n + B\lambda^{-n}$ with $AB = x^2/(x^2-4)$. It is a simple exercise to show that there is a continuous function $H : \mathbf{C} \setminus [-2,2] \longrightarrow (0,\infty)$ such that there are numbers $n_0 \leq n_1 \in \mathbf{Z}$ so that $|y_n| \leq H(x)$ if and only if $n_0 \leq n \leq n_1$ and so that $|y_n|$ is monotonically decreasing on $(-\infty, n_0 - 1]$ and monotonically increasing on $[n_1 + 1, \infty)$. We can assume that $H(x) \geq 2$ for all x. Thus,

Lemma 3.14 : *Suppose $X \in \Omega$ and $(Y_n)_{n \in \mathbf{Z}}$ is the bi-infinite sequence of regions meeting X. If $x = \phi(X) \notin [-2,2]$ and $r \geq H(x)$, set $J_r(X) = \bigcup \{X \cap Y_n \mid |y_n| \leq r\}$. Then $J(X) = J_r(X) \subseteq \partial X$ has the property described by Lemma 3.10.* ◊

Of course, it would not be hard to write down an explicit formula which would serve to define $H(x)$. We shall not bother to do this here since it seems to hard to find a simple expression which is anywhere close to optimal. Note that $J_r(X)$ always has at least one edge. If $x \in [-2,2]$ we shall set $H(x) = \infty$. Thus, if $\phi(X) \in [-2,2]$ we take $J_r(X) = \partial X$.

With a view to showing that Φ_Q is open, we introduce the following, somewhat arbitrary construction. Given $t \geq 0$, let $T(t) = T_\phi(t)$ be the union of all the arcs $J_{H(x)+t}(X)$ as X varies in $\Omega(2+t)$. The arguments of Lemma 3.11 show that $T(t)$ is connected. Also, if $T(t) \neq \emptyset$, then the arrow on any edge not in $T(t)$ points towards $T(t)$. By Theorem 1(1), we will always have $T(t) \neq \emptyset$ for $t \geq 1$.

Note that we can describe $T(t)$ directly in terms of its edges. Suppose $e = X \cap Y \in E(\Sigma)$. Then e is an edge of $T(t)$ if and only if either $|x| \leq 2+t$ and $|y| \leq H(x) + t$ or $|y| \leq 2+t$ and $|x| \leq H(y) + t$.

Lemma 3.15 : *For any fixed $t \geq 0$, we have $\phi \in \Phi_Q$ if and only if $T(t)$ is finite.*

Proof : If $\phi \in \Phi_Q$, then each arc $J_{H(x)+t}(X)$ has finitely many edges. Also $\Omega(2+t)$ is finite. We see that $T(t)$ is finite. Conversely, if $T(t)$ is finite, then $\phi^{-1}[-2,2] = \emptyset$. Since each $J_{H(x)+t}(X)$ has at least one edge, we see that $\Omega(2+t)$ is finite. Thus $\phi \in \Phi_Q$. This proves Lemma 3.15. ♢

This gives a finite criterion for recognising that a given Markoff map ϕ lies in Φ_Q. If we find a finite non-empty component of $T(t)$, then we know that it must be the whole of $T(t)$ and so $\phi \in \Phi_Q$. Such a component must exist if $t \geq 1$. This is essentially the criterion used in [10] to rule out certain Markoff triples as corresponding to double limit groups. We need to check that this criterion is indeed an open property.

Recall how the topology on Φ is defined. If $X, Y, Z \in \Omega$ meet at some vertex of Σ, then the correspondence $\phi \leftrightarrow (\phi(X), \phi(Y), \phi(Z))$ identifies Φ as the complex variety $\{(x,y,z) \in \mathbf{C}^3 \mid x^2 + y^2 + z^2 = xyz\}$. We take the subspace topology on Φ. If X', Y', Z' are another three regions meeting around a vertex, then we get a different identification. However the maps $(x,y,z) \leftrightarrow (x',y',z')$ are polynomial. In particular, we get the same topology on Φ. Note that the variety is smooth away from $\underline{0}$.

Theorem 3.16 : *Φ_Q is an open subset of Φ.*

Proof : Fix any $t_1 > t_0 > 1$. Suppose $\phi \in \Phi_Q$. Now $T(t_1)$ is a finite subtree of Σ. Clearly, if $t_2 > t_1$ is sufficiently large (depending on ϕ), then the finite tree $T(t_2)$ must contain $T(t_1)$ in its interior, i.e. it contains $T(t_1)$, together with all the edges of the circular set $C(T(t_1))$.

Given $\phi' \in \Phi$, write $T'(t)$ for $T_{\phi'}(t)$. Now, if ϕ' is sufficiently close to ϕ, then $T'(t_0) \cap T(t_2) \subseteq T(t_1)$. Also, since $T(1)$ is a non-empty subtree of $T(t_2)$, it follows that $T'(t_0) \cap T(t_2) \neq \emptyset$, provided again that ϕ' is sufficiently close to ϕ. Since $T'(t_0)$ is connected, we must have $T'(t_0) \subseteq T(t_2)$. Thus $T'(t_0)$ is finite, and so $\phi' \in \Phi_Q$.

This proves Theorem 3.16. ♢

In particular, we see that Φ_Q is locally connected, and so all its connected components are also open.

For some refinements of Theorem 3.16, see Proposition 5.8, and the related discussion.

We mentioned in Chapter 1 that Φ admits a $\mathbf{Z}_2^2$-action. We choose an identification of the non-trivial elements, g_1, g_2, g_3 of $\mathbf{Z}_2^2$ with the colours $1, 2, 3$ of the 3-colouring, $\{\Omega^1, \Omega^2, \Omega^3\}$ of Ω. Given $\phi \in \Phi$, we define $g_i\phi$ by $(g_i\phi)(X) = \phi(X)$ if $X \in \Omega^i$ and $(g_i\phi)(X) = -\phi(X)$ otherwise. It is easily seen that $\mathbf{Z}_2^2$ acts freely and properly discontinuously on $\Phi \setminus \{\underline{0}\}$. Clearly, Φ_Q is invariant under this action, and $\Phi_Q/\mathbf{Z}_2^2$ is an open subset of $\Phi/\mathbf{Z}_2^2$. Conjecture A stated in Chapter 1 is equivalent to the assertion that $\Phi_Q/\mathbf{Z}_2^2$ is connected. We say more about this in Chapter 4.

We finish this chapter with a brief discussion of real Markoff maps. Write $\Phi^{\mathbf{R}} = \{\phi \in \Phi \mid \phi(\Omega) \subseteq \mathbf{R}\}$ for the set of such maps. It is easily seen from the edge relations that $\phi \in \Phi^{\mathbf{R}}$ if and only if $x, y, z \in \mathbf{R}$, where X, Y, Z are three regions meeting at a particular vertex. Note that if ϕ is non-trivial, then all three of x, y and z must be non-zero. In fact, since $xyz = x^2+y^2+z^2 > 0$, either all three are positive, or else one is positive and the other two negative. Let us suppose that all three are positive. Then $xy - z = (x^2 + y^2)/z > 0$. Similarly $yz - x$ and $zx - y$ are positive, so it follows inductively that $\phi(X) > 0$ for all $X \in \Omega$. We conclude:

Lemma 3.17 : *Any non-trivial real Markoff map is equivalent under the $\mathbf{Z}_2^2$-action to one which takes only positive real values.* $\diamondsuit$

We write $\Phi^{\mathbf{R}+} = \{\phi \in \Phi \mid \phi(\Omega) \subseteq (0, \infty)\} \subseteq \Phi^{\mathbf{R}}$.

Suppose (x, y, z) is a positive real Markoff triple corresponding to $\phi \in \Phi^{\mathbf{R}+}$. The fact that $z^2 - xyz - (x^2 + y^2) = 0$ has a real root in z tells us that $x^{-2} + y^{-2} < 1/4$. Similarly, $y^{-2} + z^{-2} < 1/4$ and $z^{-2} + x^{-2} < 1/4$. In particular, we have $x, y, z > 2$. Thus $\Omega(2) = \emptyset$ and so $\phi \in \Phi_Q$. We see that all non-trivial real Markoff maps lie in Φ_Q.

It is a simple excercise to verify that the space $\{(x, y, z) \in \mathbf{R}^3 \mid x, y, z > 0, x^2+y^2+z^2 = xyz\}$ is diffeomorphic to $\mathbf{R}^2$, and properly embedded. In summary we see:

Proposition 3.18 : *The space $\Phi^{\mathbf{R}}$ has five connected components, namely the singleton $\{\underline{0}\}$ together with the four images of $\Phi^{\mathbf{R}+}$ under the $\mathbf{Z}_2^2$-action. Each of these four images is diffeomorphic to $\mathbf{R}^2$. Moreover $\Phi^{\mathbf{R}} \subseteq \Phi_Q \cup \{\underline{0}\}$.* $\diamondsuit$

Finally, let us consider integer valued Markoff maps. Suppose $\phi(\Omega) \subseteq \mathbf{Z}$. If ϕ is non-trivial, we may as well suppose that $\phi(\Omega) \subseteq \mathbf{N}$. We know that $\phi(X) > 2$ and so $\phi(X) \geq 3$ for all $X \in \Omega$. Also, by Theorem 1(1), there is some $Z \in \Omega$ with $z = \phi(Z) = 3$. Choose $X \in \Omega$ to minimise $\phi(X)$ among all regions adjacent to Z. Let Y be one of the two regions adjacent to both X and Z. We have $x \leq y$ and $x \leq zx - y = 3x - y$, so $y \leq 2x$. Writing $y = \mu x$ with $1 \leq \mu \leq 2$, we have $(3\mu - 1 - \mu^2)x^2 = 9$. But $3\mu - 1 - \mu^2 \geq 1$, and so $x \leq 3$. Thus $x = 3$, and y is either 3 or 6. In the latter case, $zx - y = 9 - 6 = 3$, so we may as well take $y = 3$. We have shown:

Proposition 3.19 : *If ϕ is a positive integer valued Markoff map, then we can find three regions $X, Y, Z \in \Omega$ meeting at a vertex, so that $\phi(X) = \phi(Y) = \phi(Z) = 3$.* $\diamondsuit$

This is the well-known result, due to Markoff [31], that any non-trivial integer valued Markoff triple can be reduced to $(3, 3, 3)$ by a sequence of operations of the form $(x, y, z) \mapsto (x, y, xy - z)$ and $(x, y, z) \mapsto (-x, -y, z)$ together with permutations of x, y and z. See [17] for more discussion of this.

4. Quasifuchsian groups.

In this chapter we relate some of the results on Markoff maps to spaces of representations. First we make a few observations about matrices in $SL(2, \mathbf{C})$ all of which seem to be well-known. We omit proofs of the first two.

Lemma 4.1 : *Suppose $A, B \in SL(2, \mathbf{C})$. Then A and B have a common eigenvector if and only if* $\mathrm{tr}[A, B] = 2$. $\diamondsuit$

Lemma 4.2 : *Suppose $A, B, A', B' \in SL(2, \mathbf{C})$ satisfy $\operatorname{tr} A = \operatorname{tr} A'$, $\operatorname{tr} B = \operatorname{tr} B'$ and $\operatorname{tr} AB = \operatorname{tr} A'B'$, then either $\operatorname{tr}[A, B] = \operatorname{tr}[A', B'] = 2$, or there is some $P \in SL(2, \mathbf{C})$ such that $A' = PAP^{-1}$ and $B' = PBP^{-1}$.* ♢

Lemma 4.3 : *Suppose $x, y, z \in \mathbf{C}$, then there exist $A, B \in SL(2, \mathbf{C})$ such that $\operatorname{tr} A = x$, $\operatorname{tr} B = y$ and $\operatorname{tr} AB = z$.*

Proof : One way to see this would be to write $A = \begin{pmatrix} a & b \\ b & a \end{pmatrix}$ and $B = \begin{pmatrix} \alpha & \beta \\ \gamma & \alpha \end{pmatrix}$, so that $a = x/2$, $b = a^2 - 1$ and $\alpha = y/2$. We can solve for β and γ from $\beta + \gamma = (z - 2a\alpha)/b$ and $\beta\gamma = \alpha^2 - 1$. This works unless $x = \pm 2$ (so that $b = 0$). By interchanging the roles of x, y and z, this deals with all cases except when $x, y, z \in \{-2, 2\}$. In these cases we can just write out explicit matrices (though they do not concern us here since they cannot arise from Markoff triples). This proves Lemma 4.3. ♢

Note that if $x = y = z = 0$ we get the matrices $\begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$, $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ and $\begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$. In this case, we have $AB = -BA$, or equivalently $[A, B] = -I$. We see that A and B generate the quaternionic group of order 8. Note also that the matrices A and B can be assumed to vary continuously in (x, y, z) in a neighbourhood of $(0, 0, 0)$.

If (x, y, z) is a Markoff triple, then $\operatorname{tr}[A, B] = -2$ (from the trace identity quoted in Chapter 1). In this case we get nice alternative normalisation due to Jørgensen [26], namely:

$$A = \frac{1}{z} \begin{pmatrix} xz - y & x/z \\ xz & y \end{pmatrix} \qquad B = \frac{1}{z} \begin{pmatrix} yz - x & -y/z \\ -yz & x \end{pmatrix}.$$

These matrices are unique up to simultaneous conjugacy, by Lemma 4.2. Note that this particular representation blows up when $z = 0$. We can put this right by interchanging x, y and z, unless $x = y = z = 0$. But we have already observed that the matrices can be assumed to vary continuously up to conjugacy in a neighbourhood of $(0, 0, 0)$.

Recall that a representation ρ from the free group on two generators, $\Gamma = \langle a, b \rangle$ to $SL(2, \mathbf{C})$ is called "type preserving" if $\operatorname{tr} \rho[a, b] = -2$. The space of representations carries a natural (algebraic) topology, which can be described, for example, as arising from the embedding in $SL(2, \mathbf{C}) \times SL(2, \mathbf{C})$ given by $[\rho \mapsto (\rho(a), \rho(b))]$. We take the quotient topology on the space of representations up to conjugacy.

Given a Markoff triple, (x, y, z), we construct a representation ρ by $\rho(a) = A$ and $\rho(b) = B$ using Lemma 4.3. By Lemma 4.2, ρ is unique up to conjugacy in $SL(2, \mathbf{C})$. In this way, we get an identification of the space Φ with the space of type-preserving representations defined up to conjugacy. ¿From our earlier discussion, we see that the topologies on these spaces agree.

In the case of the trivial Markoff map, we get a representation ρ_0, whose image in $SL(2, \mathbf{C})$ is isomorphic to the quaternionic group (and so projects in $PSL(2, \mathbf{C})$ to a Klein-four group). In particular, $\rho_0([a, b]) = -I$. In all other cases the image of the commutator will be a parabolic element of $PSL(2, \mathbf{C})$.

The $\mathbf{Z}_2^2$ action on Φ corresponds to multiplying one or other or both of the matrices $\rho(a)$ and $\rho(b)$ by -1. Thus, the quotient $\Phi/\mathbf{Z}_2^2$ may be identified with the space of conjugacy classes of type preserving representations into $PSL(2,\mathbf{C})$. (Note that the sign of the trace of commutator in $PSL(2,\mathbf{C})$ is well-defined.) We identify $PSL(2,\mathbf{C})$ as the group of orientation preserving isometries of hyperbolic 3-space, $\mathbf{H}^3$.

We say that a non-trivial element $A \in SL(2,\mathbf{C})$ is *elliptic* if $\operatorname{tr} A \in (-2,2)$, *parabolic* if $\operatorname{tr} A \in \{-2,2\}$ and *loxodromic* if $\operatorname{tr} A \notin [-2,2]$. We define the *complex translation distance* $l(A) \in \mathbf{C}/4\pi i\mathbf{Z}$ of loxodromic element by setting $\operatorname{tr} A = 2\cosh(l(A)/2)$ and demanding that $\Re(l(A)) > 0$. It may be given the following geometrical interpretation. When A acts on $\mathbf{H}^3$, it preserves a bi-infinite geodesic, which it translates through a distance $\Re(l(A))$ and rotates the normal bundle though an angle of $\Im(l(A))$. We call $L(A) = \Re(l(A))$ the *real translation distance* A. It is the length of the simple closed geodesic in $\mathbf{H}^3/\langle A\rangle$.

Note that for any $\zeta \in \mathbf{C}$, we have $|\cosh\zeta|^2 = \cosh^2(\Re\zeta) - \sin^2(\Im\zeta) = \sinh^2(\Re\zeta) + \cos^2(\Im\zeta)$. Thus $|\sinh(\Re\zeta)| \leq |\cosh\zeta| \leq \cosh(\Re\zeta)$. In particular, $|\operatorname{tr} A| \leq 2\cosh(L(A)/2)$, and so $\log^+|\operatorname{tr} A| \leq L(A)$.

Let $S_{g,p}$ be the (topological) surface of genus g with p punctures. We assume that the Euler characteristic, $2-2g-p$, is negative. Let $\Gamma_{g,p} = \pi_1(S_{g,p})$. (Thus $\Gamma = \Gamma_{1,1}$ is the free group on two generators.) Thus the set of homotopy classes of closed curves on $S_{g,p}$ may be identified with $\Gamma_{g,p}/\sim$, where $\sim$ identifies conjugates and conjugates of inverses. We can define a *type preserving* representation $\rho : \Gamma_{g,p} \longrightarrow SL(2,\mathbf{C})$ as one which sends peripheral elements to parabolics, and such that $\rho(\Gamma_{g,p})$ is non-elementary, i.e. its elements do not have a common eigenvector. The latter clause ensures that we have agreement with our earlier definition for $\Gamma = \Gamma_{1,1}$ (using Lemma 4.1).

Given any fixed finite generating set for $\Gamma_{g,p}$, and $\gamma \in \Gamma_{g,p}/\sim$, we define $W(\gamma)$ to be the length of the shortest cyclic word, in the generators and their inverses, representing γ. If we were to take a different finite generating set to get a function, W', there would be a constant $c \geq 1$ such that $c^{-1}W(\gamma) \leq W'(\gamma) \leq cW(\gamma)$ for all $\gamma \in \Gamma_{g,p}/\sim$. For this reason, it will not matter to us which generating set is chosen, so we shall stay with some particular function W. Suppose $\rho : \Gamma \longrightarrow SL(2,\mathbf{C})$ is a type preserving representation. Note that $\operatorname{tr}(\rho(\gamma))$ and $L(\rho(\gamma)$ are well defined for $\gamma \in \Gamma_{g,p}/\sim$. It is easily seen that there is some constant K, depending on ρ such that $L(\rho(\gamma)) \leq KW(\gamma)$ for all $\gamma \in \Gamma_{g,p}/\sim$.

Of particular interest are quasifuchsian representations. A type preserving representation, ρ, is *quasifuchsian* if it is discrete and faithful with so that $\rho(\Gamma_{g,p})$ is geometrically finite with no accidental parabolics. The last assertion means that if $\rho(g)$ is parabolic, then g is peripheral. There are many ways of characterising quasifuchsian groups, see for example [30].

For closed surfaces ($p=0$), a type preserving representation is quasifuchsian if and only if there is some $k > 0$ such that $L(\rho(\gamma)) \geq kW(\gamma)$ for all $\gamma \in \Gamma_{g,0}/\sim$. However, this fails for $p > 0$, since a closed curve may wrap itself many times around a puncture. We content ourselves with observing:

Lemma 4.4 : *If $\rho : \Gamma_{g,p} \longrightarrow SL(2,\mathbf{C})$ is quasifuchsian, then there is some $k > 0$ such that for all non-peripheral simple closed curves, γ, on $S_{g,p}$, we have $L(\rho(\gamma)) \geq kW(\gamma)$.* $\diamondsuit$

In other words, we shall restrict our attention to those closed curves which are simple.

It's not hard to see that the converse to Lemma 4.4 holds provided we assume that ρ is discrete and faithful. Indeed in this case, we can weaken the hypotheses, using the work of Bonahon (Proposition 4.5). It seems natural to ask whether the assumption of discreteness is really essential:

Question B : Suppose $\rho : \Gamma_{g,p} \longrightarrow SL(2, \mathbf{C})$ is a type preserving representation such that there is a constant $k > 0$ so that $L(\rho(\gamma)) \geq kW(\gamma)$ for all non-peripheral simple closed curves γ on $S_{g,p}$. Then, must ρ be discrete and faithful (and hence quasifuchsian)?

We shall see that Question B is equivalent to Conjecture A in the case of punctured tori (Proposition 4.9).

Proposition 4.5 : *Given any $g, p \in \mathbf{N}$, with $2 - 2g - p < 0$, there is some constant $M(g,p) > 0$ such that the following holds. Suppose $\rho : \Gamma_{g,p} \longrightarrow SL(2, \mathbf{C})$ is a discrete faithful type preserving representation with no accidental parabolics, such that the set of homotopy classes of simple closed curves, γ, on $S_{g,p}$ satisfying $L(\rho(\gamma)) \leq M(g,p)$ is finite. Then ρ is quasifuchsian.*

Proof : If ρ is not quasifuchsian, then it follows from [5] that the quotient of $\mathbf{H}^3$ by $\rho(\Gamma_{g,p})$ must have at least one simply degenerate end. (Alternatively, this is known directly from work of Thurston in the case of limits of geometrically finite groups.) Every neighbourhood of such an end contains infinitely many pleated surfaces. Such a pleated surface carries an intrinsic hyperbolic metric. It is easily seen that the shortest closed curves on these surfaces cannot all lie in the same homotopy class.

This proves Proposition 4.5. $\diamondsuit$

This argument gives $M(g,p)$ as the maximum length of a shortest closed curve on $S_{g,p}$ among all possible complete finite-area hyperbolic structures on $S_{g,p}$. (Note that a shortest closed curve is necessarily simple.)

(Note that in general, a simply degenerate end, as defined geometrically, corresponds to some subsurface, S', of $S_{g,p}$ onto which the pleated surfaces can be retracted. The boundary of S' in $S_{g,p}$ consists of simple closed curves which get sent to accidental parabolics. Since we are assuming that there are no accidental parabolics in our case, we see that S' must in fact be the whole of $S_{g,p}$. In particular, all the surfaces have the same topological type.)

In the case of once-punctured tori, this maximum is attained for the torus $\mathbf{H}^2/G$, where G is the commutator subgroup of $PSL(2, \mathbf{Z})$ as described in the introduction. This gives $M(1,1) = 2\cosh^{-1}(3/2) = 2\log((3+\sqrt{5})/2)$. It is possible to improve somewhat on this — see Proposition 4.10.

The space of representations up to conjugacy carries a natural topology. ¿From Ahlfors-Bers deformation theory, it is well-known that the subspace of quasifuchsian representations into $PSL(2, \mathbf{C})$ is diffeomorphic to $\mathbf{R}^{12g-12+4p}$ — the product of two copies of the Teichmüller space for $S_{g,p}$ (see [3] or [19]). See [34] for further discussion.

Let's return to the particular case of punctured tori. We write Φ_D for the subset of Φ

corresponding to discrete faithful representations, and Φ_{QF} for the subset corresponding to quasifuchsian representations.

¿From Lemma 4.4, we see that $\Phi_{QF} \subseteq \Phi_Q$. Also, $\Phi_{QF}/\mathbf{Z}_2^2$ is diffeomorphic to $\mathbf{R}^4$. In particular, Φ_{QF} has four components permuted under the $\mathbf{Z}_2^2$-action.

Lemma 4.6 : Φ_{QF} *is open in* Φ.

Proof : It was shown by Marden [30] that geometrically finite representations without accidental parabolics are structurally stable among all type preserving representations of a group into $PSL(2,\mathbf{C})$. In particular (in this dimension), geometrically finite representations form an open subset of the space of type-preserving representations given the algebraic topology. This proves Lemma 4.6. $\diamondsuit$

(Alternatively, one could prove Lemma 4.6 using the Invariance of Domain theorem.) Another proof of Marden's theorem and a discussion of the higher dimensional situation is given in [8].

A result, due, with varying degrees of generality, Chuckrow, Marden, Jørgensen and Wielenberg (see for example [34] or [8]), tells us that the property of being discrete and faithful is closed in the algebraic topology (provided our group is not virtually abelian). Thus:

Lemma 4.7 : Φ_D *is closed in* Φ. $\diamondsuit$

Lemma 4.8 : $\Phi_Q \cap \partial\Phi_{QF} = \emptyset$.

Proof : Suppose a representation ρ corresponds to some Markoff map $\phi \in \Phi_Q \cap \partial\Phi_{QF}$. By Lemma 4.7, ρ must be discrete. Since $\phi \in \Phi_Q$, we have that $\Omega(k)$ is finite for all k and so there are only finitely many homotopy classes of non-peripheral simple closed curves, γ, on the punctured torus, $S_{1,1}$, for which $L(\rho(\gamma)) \leq M(1,1)$. By Proposition 4.5, ρ must be quasifuchsian, and so $\phi \in \Phi_{QF}$, contradicting the fact that Φ_{QF} is open (Lemma 4.6). $\diamondsuit$

(In fact, one can avoid the use of Bonahon's theorem, by using Proposition 4.9 below.)

Proof of Theorem 4 : Since Φ_{QF} is connected, we see that it must be precisely the connected component, $\Phi_Q^0/\mathbf{Z}_2^2$, of $\Phi_Q/\mathbf{Z}_2^2$ which contains $(3,3,3)$ and hence all real Markoff maps. Thus, $\Phi_{QF} = \Phi_Q^0$.

This proves Theorem 4. $\diamondsuit$

Recall that Conjecture A asserted that $\Phi_{QF} = \Phi_Q$. It is thus equivalent to the statement that $\Phi_Q = \Phi_Q^0$, i.e. that $\Phi_Q/\mathbf{Z}_2^2$ is connected.

In general, it seems difficult to verify that a particular 2-generator group has to be discrete. Some discussion of the matter for quasifuchsian groups of punctured tori is given in [27]. A finite sufficient condition for discreteness is described in [38]. For some further discussion of Markoff triples and their relationship with representations, see [14] and [35].

The following relates Conjecture A to Question B.

Proposition 4.9 : *Suppose $\rho : \Gamma \longrightarrow SL(2, \mathbf{C})$ corresponds to a Markoff map in Φ_Q. Then there is some $k > 0$ such that $L(\rho(\gamma)) \geq kW(\gamma)$ for all simple closed curves, γ, on the punctured torus.*

Proof : Let $\phi \in \Phi_Q$ be the corresponding Markoff map. Recall that $W(\gamma)$ is the minimal cyclically reduced word length with respect to some generating set. Since it doesn't matter which generating set we choose, we may as well take it to be a free basis. Thus Lemma 2.2.1 tells us that for the appropriate edge $e \in E(\Sigma)$, we have $W(\gamma) = F_e(X)$ where γ corresponds to $X \in \Omega$. By Theorem 2, there is a constant $k > 0$ such that $\log^+ |\phi(X)| \geq kF_e(X)$ for all $X \in \Omega$. Recall that $\phi(X) = \operatorname{tr} \rho(\gamma)$, so from the inequality $\log^+ |\operatorname{tr} A| \leq L(A)$, we see that $L(\rho(\gamma)) \geq \log^+ |\phi(X)| \geq kF_e(X) = kW(\gamma)$ as required. ♢

This shows that Question B reduces to Conjecture A in the case of punctured tori. We can also weaken the hypetheses of Proposition 4.5 in this case:

Proposition 4.10 : *Suppose $\rho : \Gamma \longrightarrow SL(2, \mathbf{C})$ is discrete faithful and type-preserving, and that the set of homotopy classes of simple closed curves, γ, for which $|\operatorname{tr} \rho(\gamma)| \leq 2$ is finite, then ρ is quasifuchsian.*

Proof : By Theorem 2, the set of γ such that $|\operatorname{tr} \rho(\gamma)| \leq k$ is finite for all k. Apply Proposition 4.5 using the fact that $L(\rho(\gamma)) \leq \log^+(\operatorname{tr} \rho(\gamma))$. This proves Proposition 4.10. ♢

¿From the inequality $|\cosh \zeta| \geq \sinh \Re \zeta$, we see that if $|\operatorname{tr} \rho(\gamma)| \leq 2$, then $L(\gamma) \leq 2 \sinh^{-1} 1 = 2 \log(1 + \sqrt{2})$. This means that we can replace the constant $M(1, 1) = 2 \log((3+\sqrt{5})/2)$ in Proposition 4.5 by $M(1, 1) = 2 \log(1+\sqrt{2})$. Since $1+\sqrt{2} < (3+\sqrt{5})/2$, Proposition 4.10 is indeed a strengthening of Proposition 4.5. As already mentioned, I suspect the assumption of discreteness is superfluous.

We should say a few words about real Markoff maps. We have already seen that the special case of the trivial Markoff map corresponds to a representation whose image in $PSL(2, \mathbf{C})$ is a Klein-four group acting by half-turns with a fixed point in $\mathbf{H}^3$. In all other cases we have:

Proposition 4.11 : *A representation corresponds to a Markoff map in $\Phi^{\mathbf{R}} \setminus \{\underline{0}\}$ if and only if it is (conjugate to) a fuchsian representation.*

Recall that a "fuchsian representation" is a representation into $SL(2, \mathbf{R})$ which is discrete and faithful. Note that all fuchsian representations are quasifuchsian.

Proof : Clearly the traces of a fuchsian representation are all real. Conversely if ρ corresponds to an element of $\Phi^{\mathbf{R}} \setminus \{\underline{0}\}$, then the Jørgensen normalisation shows that if can be conjugated into $SL(2, \mathbf{R})$. There are several ways to see discreteness. For example, we know that $(\Phi^{\mathbf{R}} \setminus \{\underline{0}\})/\mathbf{Z}_2^2$ is connected and a subset of $\Phi_Q/\mathbf{Z}_2^2$ (Proposition 3.18). It follows that $\Phi^{\mathbf{R}} \setminus \{\underline{0}\} \subseteq \Phi_Q^0$. By Theorem 4, ρ is quasifuchsian and hence discrete, as required. ♢

We have effectively shown that a (non-elementary) type preserving representation of

Γ into $PSL(2,\mathbf{R})$ is necessarily discrete. For other surfaces, Question B suggests the following:

Question C : Suppose $\rho : \Gamma_{g,p} \longrightarrow PSL(2,\mathbf{R})$ is a non-elementary type preserving representation such that the image of every non-peripheral simple closed curve is hyperbolic (i.e. loxodromic). Then must ρ be fuchsian (i.e. discrete and faithful)?

Note that, by a theorem of Goldman [22], a type-preserving representation is fuchsian precisely if has the correct Euler class. We have already observed that Question C is true for punctured tori.

We next move on to give an interpretation of McShane's identity for quasifuchsian groups. It will be more convenient in what follows to consider representations into $PSL(2,\mathbf{C})$ (as opposed to $SL(2,\mathbf{C})$).

Note that if $\Re\zeta > 0$, then $\Re(1/(1+e^\zeta)) < 1/2$, and so $\frac{1}{1+e^\zeta} = h(2\cosh(\zeta/2))$, where $h(\xi) = \frac{1}{2}\left(1 - \sqrt{1-\frac{4}{\xi^2}}\right)$ (with the usual convention that the real part of a square root is non-negative). It follows that if $A \in PSL(2,\mathbf{C})$, then $\frac{1}{1+e^{l(A)}} = h(\operatorname{tr} A)$, where $l(A)$ is the complex translation distance of A. (Note that both $e^{l(A)}$ and $h(\operatorname{tr} A)$ are well-defined.) We may rephrase Theorem 3 as follows. Suppose $\rho : \Gamma \longrightarrow PSL(2,\mathbf{C})$ is a representation corresponding to an element of $\Phi_Q/\mathbf{Z}_2^2$, then

$$\sum \frac{1}{1+e^{l(\rho(\gamma))}} = \frac{1}{2}$$

where the sum is taken over all homotopy classes of non-peripheral simple closed curves on the punctured torus. Moreover, the convergence is absolute. In the case of quasifuchsian groups the identity follows by analytic continuation from the fuchsian case. The above result is possibly more general (if Conjecture A is false).

This identity has a very nice geometrical interpretation, which provided the inspiration for the result in the fuchsian case, and is well set out in the original paper [32]. Since the quasifuchsian case was not discussed there, we give a brief account below. To do this we need to begin with a few observations about the symmetries of type preserving representations.

Suppose $A, B \in PSL(2,\mathbf{C})$ are two matrices with $\operatorname{tr}[A,B] \neq 2$, then there is an involution $Q \in PSL(2,\mathbf{C})$ such that $QAQ^{-1} = A^{-1}$ and $QBQ^{-1} = B^{-1}$. Let $P = AQ$ and $R = QB$. We have $P^2 = Q^2 = R^2 = I$, and $A = PQ$ and $B = QR$. Also $[A,B] = ABA^{-1}B^{-1} = (PRQ)^2$.

We may regard $\Gamma \cong \mathbf{Z} * \mathbf{Z}$ as a normal subgroup of index 2 in $\mathbf{Z}_2 * \mathbf{Z}_2 * \mathbf{Z}_2$, by identifying $a = pq$ and $b = qr$ where p, q, r respectively generate the three factors of $\mathbf{Z}_2$ in $\mathbf{Z}_2 * \mathbf{Z}_2 * \mathbf{Z}_2$. It follows from the above discussion that any admissible representation $\rho : \Gamma \longrightarrow PSL(2,\mathbf{C})$ extends to a representation $\hat\rho : \mathbf{Z}_2 * \mathbf{Z}_2 * \mathbf{Z}_2 \longrightarrow PSL(2,\mathbf{C})$. Thus $\hat\rho(\mathbf{Z}_2 * \mathbf{Z}_2 * \mathbf{Z}_2)$ normalises $\rho(\Gamma)$. In particular the normaliser of $\rho(\Gamma)$ contains a square root of the commutator $[\rho(a), \rho(b)]$.

We consider $PSL(2,\mathbf{C})$ as acting on the upper half-space model of $\mathbf{H}^3$. We thus identify the ideal boundary of $\mathbf{H}^3$ with $\mathbf{C} \cup \{\infty\}$. If $\rho : \Gamma \longrightarrow PSL(2,\mathbf{C})$ is a type-preserving representation, we can normalise so that $\rho([a,b])$ is given by the translation $[\zeta \mapsto \zeta + 2]$ on $\mathbf{C}$. Thus ∞ is a parabolic fixed point of $\rho(\Gamma)$. ¿From our earlier discussion, we see that $[\zeta \mapsto \zeta + 1]$ lies in the normaliser of $\rho(\Gamma)$.

First, let us recall from [32] the picture in the fuchsian case. Our description will be fairly minimal — we refer to the original paper for elaboration. In this case we have a representation $\rho_0 : \Gamma \longrightarrow PSL(2,\mathbf{R})$. We use the upper half space-model for $\mathbf{H}^2$ with ideal boundary $\mathbf{R} \cup \infty$ and normalise ρ_0 accordingly. Thus $\mathbf{H}^2/\rho_0(\Gamma)$ is a hyperbolic once-punctured torus. We shall work mostly in the quotient $\mathbf{H}^2/\mathbf{Z}$), where the action of $\mathbf{Z}$ is generated by the translation $[\zeta \mapsto \zeta + 1]$. This surface has one cuspidal end, and another which is compactified by the "ideal circle", $C_0 = \mathbf{R}/\mathbf{Z}$.

Corresponding to each $X \in \Omega$, there is a unique simple closed geodesic $\gamma(X)$ on $\mathbf{H}^2/\rho(\Gamma)$. There is also a unique properly embedded bi-infinite geodesic arc, $\beta(X)$, running from cusp to cusp, with the property that $\beta(X) \cap \gamma(X) = \emptyset$. Among the lifts of $\beta(X)$ to $\mathbf{H}^2$, we select those which have ∞ as an ideal endpoint. These geodesics project to a single geodesic arc, $\beta_0(X)$ on $\mathbf{H}^2/\mathbf{Z}$, which has one end going up the cusp, and the other endpoint, $p_0(X)$ on the ideal circle C_0. Similarly, we lift $\gamma(X)$ to a set of bi-infinite geodesics in $\mathbf{H}^2$, and select those which are "nearest to ∞", i.e. which can be connected to ∞ by an arc which does not meet any other lift of $\gamma(X)$. This subset projects of a single bi-infinite geodesic $\gamma_0(X)$ in $\mathbf{H}^2/\mathbf{Z}$. Now $\gamma_0(X)$ separates an open hyperbolic half-space, $H(X)$, from the cusp of $\mathbf{H}^2/\mathbf{Z}$. Let $\alpha_0(X) = C_0 \setminus H(X)$. Thus, $\alpha_0(X)$ is an arc of C_0, centred on the point $p_0(X)$, and whose endpoints coincide with those of $\gamma_0(X)$.

Now, as X varies in Ω, the arcs $\alpha_0(X)$ are all disjoint. Their cyclic order on C_0 agrees with the natural cyclic order on Ω. The complement, $R_0 = C_0 \setminus \bigcup_{X \in \Omega} \alpha_0(X)$ is a Cantor set. Taking the total length of C_0 to be 1, a computation shows that the length of the arc $\alpha_0(X)$ is $2/(1+e^{l(\rho_0(\gamma))})$ where $l(\rho_0(\gamma))$ is the the length of the closed geodesic $\gamma(X)$ on $\mathbf{H}^2/\rho_0(\Gamma)$. McShane's identity is thus equivalent to asserting that the Lebesgue measure of the Cantor set R_0 is 0.

In fact, R_0 has Hausdorff dimension 0. This follows from the result of [4] which was the starting point in the original proof in [32]. If we prefer, we can go backwards. Given McShane's identity, it follows that R_0 has Lebesgue measure 0. To see that it has Hausdorff dimension 0, we need, in addition, that the lengths of the complementary segments decay exponentially. This follows from the fact that $\log^+ |\phi|$ has Fibonacci growth (Theorem 2), given that the length of $\alpha_0(X)$ is $h(\phi(X)) = O(|\phi(X)|^{-2})$.

Now suppose that $\rho : \Gamma \longrightarrow PSL(2,\mathbf{C})$ is quasifuchsian, normalised as described earlier. The limit set Λ of $\rho(\Gamma)$ is a quasicircle passing through ∞. Now, $\Lambda \setminus \{\infty\}$ is invariant under the translation $[\zeta \mapsto \zeta + 1]$, and so it projects to a quasicircle, $C = (\Lambda \setminus \{\infty\})/\mathbf{Z}$ in the cylinder $\mathbf{C}/\mathbf{Z}$. If we fix a particular fuchsian representation ρ_0 as above, we get a natural dynamically defined identification of C_0 with C. This identification extends to a quasiconformal homeomorphism of the cylinder $\mathbf{C}/\mathbf{Z}$.

Given any $X \in \Omega$, we may define, by an analogous procedure, a point $p(X) \in C$, and an arc $\alpha(X) \subseteq C$ "centred" on $p(X)$. These may be alternatively described as the images of $p_0(X)$ and $\alpha_0(X)$ under the above identification. Again, the arcs $\alpha(X)$ are all disjoint as X varies in Ω, and the complement $R = C \setminus \bigcup_{X \in \Omega} \alpha(X)$ is a Cantor set — the image of R_0 under the identification. This Cantor set also has Hausdorff dimension 0. This can be seen, for example, from the fact that the property of having Hausdorff dimension 0 is preserved under quasiconformal homeomorphisms [21]. (See also [1] for some refinements of this result.)

Of course, C will not be rectifiable (unless ρ happens to be fuchsian). To put this right, we replace each arc $\alpha(X)$ by the straight line segment, $\alpha'(X)$, with the same endpoints as $\alpha(X)$, and in the same homotopy class relative to its endpoints. (By "straight" we mean geodesic with respect to the standard euclidean structure on $\mathbf{C}$.) The resulting curve is rectifiable (though not necessarily embedded). In fact a computation shows that the complex number represented by the segment $\alpha'(X)$ (i.e. the difference of its endpoints in $\mathbf{C}$) is equal to $2h(\phi(X)) = 2/(1 + e^{l(\rho(\gamma))})$ where l is complex translation distance, and γ is the curve corresponding to X. Thus, summing over all $X \in \Omega$, McShane's identity tells us that the total displacement is equal to 1, as we expect from this geometric picture.

The proof offered in [6] can also be interpreted in this picture. If the regions $X, Y, Z \in \Omega$ meet at some vertex of Σ, then the complex number given by difference $p(Y) - p(X)$, taking proper account of the homotopy class of the arc joining $p(X)$ to $p(Y)$ in C, is equal to z/xy. This explains the vertex identity, since $\frac{z}{xy} + \frac{y}{zx} + \frac{x}{yz} = (p(Y) - p(X)) + (p(X) - p(Z)) + (p(Z) - p(Y))$ which is equal to 1, since we have wandered once around the curve C.

The points $p(X)$ also have interpretations as ideal vertices of certain ideal triangulations of (the convex core of) the quotient manifold $\mathbf{H}^3/\rho(\Gamma)$. These triangulations are perhaps more interesting in the case of double limit groups, and are discussed in [18] and [7].

5. Non-discrete representations.

In this section, we shall consider Markoff maps lying in the complement, $\Phi \setminus \Phi_Q$. As one might expect, these prove more difficult to analyse. We get started here with a few results about such maps, though most are somewhat weaker than what one might hope for, and we suggest a few directions for further inquiry. One of the basic consequences of our analysis is that the interior of $\Phi \setminus \Phi_Q$ is non-empty. In fact, $\underline{0} \in \mathrm{int}(\Phi \setminus \Phi_Q)$ (See Corollary 5.6).

Recall that $PSL(2, \mathbf{Z})$ acts faithfully on Σ and Ω and hence on Φ in the obvious way: $\gamma(\phi)(X) = \phi(\gamma X)$, where $X \in \Omega$, $\phi \in \Phi$ and $\gamma \in PSL(2, \mathbf{Z})$. Clearly, the sets Φ_Q, Φ_{QF}, Φ_D and $\Phi^{\mathbf{R}}$ are all invariant under this action. It's easy to see (via Theorem 2, for example) that $PSL(2, \mathbf{Z})$ acts properly discontinuously on Φ_Q. The action on $\Phi \setminus \Phi_Q$ is much more complicated, and seems more interesting from a dynamical point of view. We suspect that, in some sense, "most" orbits in this set are dense. What we are able to show is quite a lot weaker. For example, the orbit of any $\phi \in \Phi$ sufficiently close to $\underline{0}$ must accumulate at $\underline{0}$. The fact that $\underline{0} \in \mathrm{int}(\Phi \setminus \Phi_Q)$ is a simple consequence. Another question of interest seems to be what subgroups of $PSL(2, \mathbf{Z})$ can stabilise an element of $\Phi \setminus \{\underline{0}\}$. The only examples I know are either finite or virtually cyclic.

We saw in Section 4, that $\partial\Phi_{QF}$ is a subset of the closed set $\Phi_D \setminus \Phi_{QF}$, and it has recently been shown by Minsky [33] that these sets are in fact equal; i.e. that all discrete faithful representations are limits of quasifuchsian representations. (Minsky's result is specific to the case of a once-punctured torus — the issue for more complicated surfaces remains open.) ¿From the earlier work of Thurston, Bonahon and others [43,5], quite a lot is known about the geometry of a representation, ρ, corresponding to some $\phi \in \partial\Phi_{QF}$. For a start, $\phi^{-1}(-2,2) = \emptyset$, and $\phi^{-1}\{-2,2\}$ consists of at most two elements. These correspond to "accidental parabolics" of ρ. The quotient, $M = \mathbf{H}^3/\rho(\Gamma)$ is homeomorphic to a punctured torus times the real line. Each of the two "ends" of M (corresponding to the topological ends of the real line) is either simply degenerate or geometrically finite. In the latter case the end may or may not contain an accidental parabolic. (In this respect, the case of the once-punctured torus is special, since there is not enough room for a single topological end to split into different geometrically finite or simply degenerate pieces separated by accidental parabolics — as might happen for a more complicated surface.) The group is termed "geometrically finite", "a Bers boundary group" or "a double limit group" according to whether neither, one or both of its ends are simply degenerate. A quasifuchsian group is thus a geometrically finite group without any accidental parabolics. We shall return to a discussion of these cases later.

An important class of double limit groups are those for which the corresponding Markoff map is invariant under a cyclic subgroup of $PSL(2, \mathbf{Z})$. This subgroup translates a certain bi-infinite path in Σ. This "axis" meets infinitely many regions of $\Omega(2)$. A concrete example of such a map is that corresponding to the Markoff triple $\left(\frac{3+\sqrt{-3}}{2}, \frac{3-\sqrt{-3}}{2}, \frac{3-\sqrt{-3}}{2}\right)$. This is invariant under (a conjugate of) the subgroup of $PSL(2, \mathbf{Z})$ generated by $\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$. Such Markoff maps are discussed in [10] and [7].

We can also explicitly describe various Markoff maps corresponding to non-discrete representations which turn out to be of interest.

Suppose we have some $X \in \Omega$ with $x = \phi(X) \in (-2, 2)$, so that, by definition, $\phi \in \Phi \setminus \Phi_Q$. Let $(Y_n)_{n\in\mathbf{Z}}$ be the bi-infinite sequence of regions meeting X, as in Lemma 3.3. We thus have $x = \lambda+\lambda^{-1}$ and $y_n = A\lambda^n+B\lambda^{-n}$, where $|\lambda| = 1$ and $AB = x^2/(x^2-4)$. In particular, $|y_n|$ is bounded. Given any $K \geq 0$ and any $x \in (-2, 2)$, it is possible to find such a Markoff map with $|y_n| \geq K$ for all n. If K is large enough, then we will have that if e is any edge of Σ which meets X in a single vertex, then $\alpha(e)$ points towards that vertex. Provided $K \geq 2$, Lemma 3.5 then tells us that $\Omega^{0-}(\alpha(e)) \cap \Omega(2) = \emptyset$. (In fact, by Proposition 3.9, $\log^+ |\phi|$ has a lower Fibonacci bound on $\Omega^{0-}(\alpha(e))$.) We see that $\Omega(2) = \{X\}$. In particular, $|\phi|$ has a unique minimum at X. Note that this minimum may be arbitrarily close to, or equal to, 0.

Another interesting class is that of what we shall call *imaginary* Markoff maps, $\Phi^I \subseteq \Phi$. We shall need to look at these in some detail, as they arise as special cases in the proofs of later results. We may define Φ^I as follows. Recall the 3-colouring of Ω described at the end of Section 2.1. We define Φ^I to be the set of those $\phi \in \Phi \setminus \{\underline{0}\}$ such that for some element, $\Omega' \subseteq \Omega$, of the partition of Ω by the 3 colours, we have $\phi(\Omega') \subseteq \mathbf{R}$ and $\phi(\Omega \setminus \Omega') \subseteq i\mathbf{R}$. Note that it is sufficient to verify this at a single vertex of Σ; i.e. we want that the corresponding Markoff triple consists of one real and two purely imaginary numbers.

We claim that $\Phi^I \cap \Phi_Q = \emptyset$. In fact,

Proposition 5.1 :
(1) Suppose that $\phi \in \Phi^I$ and $\phi^{-1}(0) = \emptyset$. Then the orbit of ϕ under $PSL(2, \mathbf{Z})$ accumulates at $\underline{0}$.
(2) There is some neighbourhood, N^I, of $\underline{0}$ in Φ^I such that if $\phi \in N^I$ then, again, the orbit of ϕ accumulates at $\underline{0}$.

Note that some restriction on ϕ is necessary in Proposition 5.1, in addition to it lying in Φ^I. As the previous example showed, it is possible to find $\phi \in \Phi^I$ for which the $|\phi|$ is equal to 0 at a single $X \in \Omega$, but is otherwise bounded away from 0. For example, start with the Markoff triple $(0, \mu, i\mu)$ for sufficiently large $\mu \in [0, \infty)$. The orbit of such a map cannot accumulate at $\underline{0}$. It is not clear whether all such examples arise in this way.

Before we set about proving Proposition 5.1, let us make a few remarks about general Markoff maps, ϕ. Note that the statement that the orbit of ϕ accumulates at $\underline{0}$ is really an intrinsic property of ϕ. It says that we can find vertices of Σ so that the corresponding Markoff triple (x, y, z) is arbitrarily close to $(0, 0, 0)$. Note that if x and y are both small, then so also is z. Thus, in fact, it's sufficient to find edges e_n such that $\max\{|\phi(X)| \mid X \in \Omega^0(e_n)\}$ tends to 0.

There is a convenient way of representing Markoff triples, which give rise to alternative coordinates for Ω. Given a Markoff triple (x, y, z) with $x, y, z \neq 0$, set $a = x/yz$, $b = y/zx$ and $c = z/xy$, so that $a+b+c = 1$. Conversely, given $a, b, c \in \mathbf{C} \setminus \{0\}$ with $a+b+c = 1$, we can recover (x, y, z) up to changes of sign (i.e. up to the action of $\mathbf{Z}_2^2$) using the formulae $x^2 = 1/bc$, $y^2 = 1/ca$, $z^2 = 1/ab$.

Suppose $\phi \in \Phi$, and that the triple (x, y, z) arises at the vertex $v = X \cap Y \cap Z$ of Σ. We write $k(v) = |a| + |b| + |c|$ and $m(v) = \min\{|a|, |b|, |c|\}$. If we move to the adjacent vertex $v' = Y \cap Z \cap W$ along the edge $Y \cap Z$, then we should replace the triple (a, b, c) by $\left(1 - a, \frac{ab}{1-a}, \frac{ac}{1-a}\right)$. The only problem arises if $a = 1$, in which case we can deduce that $w = \phi(W) = 0$. So, for example, if we know that $\phi^{-1}(0) = \emptyset$, then we can move around all of Σ in this way.

Suppose that $\Re(a) > 1/2$. Then $|1 - a| < |a|$. Since $|1 - a| \leq |b| + |c|$, we have

$$\begin{aligned} |b| + |c| - |1 - a| &\leq \left|\frac{a}{1 - a}\right| (|b| + |c| - |1 - a|) \\ &= \left|\frac{ab}{1 - a}\right| + \left|\frac{ac}{1 - a}\right| + |a|. \end{aligned}$$

Thus,

$$k(v) = |a| + |b| + |c| \leq \left|\frac{ab}{1 - a}\right| + \left|\frac{ac}{1 - a}\right| + |1 - a| = k(v').$$

We have equality only if $|b| + |c| = |1 - a|$. Since $b + c = 1 - a$, this can happen only if b is a positive real multiple of c.

Note that to show that the orbit of ϕ accumulates at $\underline{0}$, it's enough to show that we can find $v \in V(\Sigma)$ with $k(v)$ arbitrarily large while $k(v)$ is bounded away from 0. $\{k(v) \mid v \in V(\Sigma)\}$ has no upper bound.

Let's return to the case of an imaginary Markoff map $\phi \in \Phi^I$. Now we can order the triple (a, b, c) so that a is the largest positive number, and c is the least negative number; in other words, $|b| \leq a$ and $c \leq b$ (which implies $c < 0$). In particular, $a \geq (1+k(v))/4 > 1/2$. It follow that there is a neighbouring vertex v' for which $k(v') \geq k(v)$.

The idea for the proof of Proposition 5.1 is thus to start at some vertex $v_0 \in V(\Sigma)$, and to define a path in Σ consisting of an infinite sequence $(v_n)_{n\in\mathbf{N}}$ of vertices, so that v_{n+1} is obtained from v_n in the manner described above. More formally, we write $v_n = X_n \cap Y_n \cap Z_n$ in such a way that $|b_n| \leq a_n$ and $c_n \leq b_n$ where $a_n = x_n/y_n z_n$, $b_n = y_n/z_n x_n$ and $c_n = z_n/x_n y_n$. We let v_{n+1} be the other endpoint of the edge $Y_n \cap Z_n$. This gives us an infinite sequence, provided we never run into a region in $\phi^{-1}(0)$. Recall that $k(v_n) = |a_n| + |b_n| + |c_n|$, and $m(v_n) = \min\{|a_n|, |b_n|, |c_n|\}$.

Lemma 5.2 : *If $(v_n)_{n\in\mathbf{N}}$ is a sequence of vertices as described above, then $k(v_n) \to \infty$ and $m(v_n)$ is bounded away from 0.*

Proof : We know that $k(v_n)$ is monotonically non-decreasing. Suppose, for contradiction, that $k(v_n)$ is bounded. It thus tends to some number $k = 1 + 2h$ where $h > 0$.

Suppose first, that $b_n \to 0$. Then $a_n \to 1 + h$ and $c_n \to -h$. Now $\left|\frac{a_n}{1-a_n}\right| \to 1 + \frac{1}{h} > 1$. The numbers $a_{n+1}, b_{n+1}, c_{n+1}$ are defined so that $|b_{n+1}| \leq a_{n+1}$ and $c_{n+1} \leq b_{n+1}$. It thus follows that, for all sufficiently large n, we have $b_{n+1} = \frac{a_n b_n}{1-a_n}$, and so $|b_{n+1}| \geq |b_n|$, contradicting the fact that $b_n \to 0$.

It now follows that we can find some subsequence of vertices, $(v_{n_i})_{i\in\mathbf{N}}$ so that $(a_{n_i}, b_{n_i}, c_{n_i})$ converges on some triple (a, b, c) with $|b| \leq a$, $c \leq b$ and with $b \neq 0$. Note that $a \neq 1$; otherwise (since $|c_n| + |b_n|$ is bounded below) we would find that when a_{n_i} is sufficiently close to 1, the quantity $k(v_{n_i+1})$ could be make arbitrarily large, contradicting the assumption that $k(v_n)$ is bounded. Now let (a', b', c') be the triple obtained by ordering the numbers $\left\{1 - a, \frac{ab}{1-a}, \frac{ac}{1-a}\right\}$ so that $|b'| \leq a'$ and $b' \leq c'$. Thus, $(a_{n_1+1}, b_{n_1+1}, c_{n_1+1}) \to (a', b', c')$. Now $k = |a| + |b| + |c| = |a'| + |b'| + |c'|$, and so (from an earlier discussion) b is a positive real multiple of c, i.e. $b < 0$. Now since $a + b + c = 1$, we have $a > 1$ and so $b' > 0$. However, repeating the above argument with the subsequence (v_{n_i+1}) in place of (v_{n_i}), we deduce that $b' < 0$, and so get a contradiction.

The only way out of this mess is to admit that $k(v_n) \to \infty$.

It remains to see that $m(v_n)$ is bounded below. Now, for all sufficiently large n, we have $k(v_n) \geq 7$. Thus, $a_n \geq (1 + k(v_n))/4 \geq 2$, and $|a_n - 1| \geq 1$. Now, $m(v_{n+1}) \geq \min\{|a_n - 1|, m(v_n)\} \geq \min\{1, m(v_n)\}$, and the result follows easily. $\diamondsuit$

Proof of Proposition 5.1 : It is sufficient to find vertices $v \in V(\Sigma)$ for which $k(v)$ is arbitrarily large, while $m(v)$ remains bounded.

(1) Choose any $v_0 \in V(\Sigma)$. Since $\phi^{-1}(0) = \emptyset$, we construct the sequence (v_n) of Lemma 5.2.

(2) Let $v = X \cap Y \cap Z$ be some fixed vertex of Σ. Let (x, y, z) be the Markoff triple corresponding to $\phi \in \Phi^I$.

Suppose first that $x, y, z \neq 0$. Then $abc = \pm 1/xyz$, so we can suppose that $k(v) \geq 7$ (say). Thus, we may construct the sequence (v_n) starting with $v_0 = v$. Since $k(v_n)$ is monotonically non-decreasing, there is no risk that a_n is ever equal to 1 (since $a_n \geq (1 + k(v_n))/4 \geq 2$.

Suppose now that $x = 0$. Without loss of generality, we have $y = \mu$ and $z = i\mu$ where $\mu \in \mathbf{R} \setminus \{0\}$. Let $W \in \Omega$ be the region meeting $Y \cap Z$ in the opposite endpoint, v' from v. Thus $w = \phi(W) = i\mu^2$. If ϕ is sufficiently close to $\underline{0}$, then we can suppose that $|\mu| < 1$, and so the triple (a, b, c) at v' is given by $(\mu^{-2}, 1, -\mu^{-2})$. We can also suppose that $k(v') > 7$ (say), so we may construct our sequence (v_n) starting with $v_0 = v'$.

This proves Proposition 5.1. ♢

Corollary 5.3 : $\Phi_Q \cap \Phi^I = \emptyset$.

Proof : Suppose $\phi \in \Phi_Q \cap \Phi^I$. By definition, $\phi^{-1}(0) = \emptyset$ so the orbit of ϕ accumulates at $\underline{0}$. This contradicts Theorem 2. ♢

So far the examples of Markoff maps in $\Phi \setminus \Phi_Q$ we have considered have been non-generic. To show that $\mathrm{int}(\Phi \setminus \Phi_Q)$ is non-empty, we shall need a more robust criterion for recognising that a particular map is not in Φ_Q. The main result we are aiming at is Theorem 5.5. Let us begin with a few general observations.

A convenient way to represent a general Markoff triple (x, y, z) with $x \neq \pm 2$ is given by:

$$\begin{aligned} x &= 2\cosh\beta \\ y &= 2\coth\beta\cosh\gamma \\ z &= 2\coth\beta\cosh(\beta+\gamma), \end{aligned}$$

where $\beta, \gamma \in \mathbf{C}$. Note moreover, that if $x \notin (-2, 2)$, then $\beta \notin i\mathbf{R}$. Suppose that $X \in \Omega$, and that $(Y_n)_{n\in\mathbf{Z}}$ is the bi-infinite sequence of regions adjacent to X. Assuming $x \notin [-2, 2]$, we write $x = 2\cosh\beta$ and $y_n = 2\coth\beta\cosh(\beta n + \gamma)$. Since $\beta \notin i\mathbf{R}$, we can write $\beta n + \gamma = \beta t + i\psi$, where $t = n + t_0$, and $t_0, \psi \in \mathbf{R}$ are fixed. Without loss of generality, we can assume that $|t_0| < 1/2$ (if not, relabel the sequence (Y_n) by shifting the subscripts). In the case where $|x|$ is small, it will be convenient to write $\beta = \delta + (\theta + \pi/2)i$, where $\delta, \theta \in \mathbf{R}$ are both small. We obtain $x = 2i\sinh(\delta + i\theta)$ and $y_n = 2\tanh(\delta + i\theta)\cosh(\delta t + ((\theta + \pi/2)t + \psi)i)$.

Lemma 5.4 : *There exist fixed constants $\epsilon_0 > 0$ and $\mu < 1$ such that the following holds. Suppose $\phi \in \Phi \setminus \Phi^I$, and that $X \in \Omega$ with $|\phi(X)| < \epsilon_0$ and $\phi(X) \notin \mathbf{R}$. Then there is some region $Y \in \Omega$, adjacent to X, such that $|\phi(Y)| \leq \mu|\phi(X)|$ and $\phi(Y) \notin \mathbf{R}$.*

Proof : Set $x = 2i\sinh(\delta + i\theta)$, with $\delta, \theta \in \mathbf{R}$ small. Let (Y_n) be the bi-infinite sequence of regions meeting X. Then, in the above notation, we have $y_n/x = -i\,\mathrm{sech}(\delta + i\theta)\cosh(\delta t + ((\theta + \pi/2)t + \psi)i)$, and so

$$|y_n/x|^2 = |\,\mathrm{sech}(\delta + i\theta)|^2(\sinh^2\delta t + \cos^2((\theta + \pi/2)t + \psi)).$$

Now, we can suppose that $|\operatorname{sech}(\delta + i\theta)|$ is arbitrarily close to 1. Also, since $|t - n| = |t_0| \leq 1/2$, we can suppose that for all n in a certain range, $-n_0 \leq n \leq n_0$, we have that $\sinh \delta t$ is arbitrarily small. We are free to chose n_0 at the outset; $n_0 = 3$ will serve for our purposes here.

Suppose $n, n' \in [-n_0, n_0] \cap \mathbf{Z}$, and $t = n + t_0$ and $t' = n' + t_0$. If $n - n'$ is even, then the quantities $\cos^2((\theta + \pi/2)t + \psi)$ and $\cos^2((\theta + \pi/2)t' + \psi)$ are almost equal; whereas if $n - n'$ is odd, then they sum almost to 1. Thus, again without loss of generality, we can suppose that for all even $n \in [-n_0, n_0]$ we have that $\cos^2((\theta + \pi/2) + \psi)$ is not much bigger than 1/2. It follows that for such n, $|y_n/x|$ is not much bigger than $\frac{1}{\sqrt{2}}$ and so is certainly less than μ for some fixed $\mu < 1$.

It remains to worry about the possibility that for all such n, y_n will be real. Suppose then, that $y_{-2}, y_0, y_2 \in \mathbf{R}$. ¿From the edge relations, we find that xy_{-1} and xy_1 are real. Thus $x^2 y_0 = x(y_{-1} + y_1) = xy_{-1} + xy_1 \in \mathbf{R}$, so $x^2 \in \mathbf{R}$. We can suppose that $|x| < 2$, so if $x \in \mathbf{R}$, we see easily that $\phi = \underline{0}$. Thus we can assume $x \in i\mathbf{R}$, so $\phi \in \Phi^I$. ♢

Theorem 5.5 :

(1) There exists $\epsilon_0 > 0$ so that if $\phi \in \Phi$ with $\inf\{|\phi(X)| \mid X \in \Omega\} < \epsilon_0$, then $\phi \notin \Phi_Q$.

(2) There is a neighbourhood, N, of $\underline{0}$ in Φ such that if $\phi \in N$, then the orbit of ϕ under the action of $PSL(2, \mathbf{Z})$ accumulates at $\underline{0}$.

Proof :

(1) By Corollary 5.3, we can suppose that $\phi \notin \Phi^I$. Starting with any region $X_0 \in \Omega$ with $|x_0| \leq 0$, Lemma 5.4 gives us a sequence of regions $(X_n)_{n \in \mathbf{N}}$ with $|x_n| \to 0$.

(2) We can suppose that $N \cap \Phi^I \subseteq N_I$, so that by Proposition 5.1(2), we can assume that $\phi \notin \Phi^I$. For a fixed vertex $v = X \cap Y \cap Z$ of Σ we have $|\phi(X)|, |\phi(Y)|, |\phi(X)| \leq \epsilon_0 < 2$. If $x, y, z \in \mathbf{R}$ then ϕ would be trivial. Thus we can assume that $x \notin \mathbf{R}$. We now construct the sequence (X_n) as in part (1), starting with $X_0 = X$. Let e_n be the edge $X_n \cap X_{n+1}$. Thus $\max\{|\phi(X)| \mid X \in \Omega^0(e_n)\} \to 0$. So, as observed earlier, the orbit of ϕ accumulates at $\underline{0}$. ♢

¿From either part(1) or part(2) of Theorem 5.5, we get

Corollary 5.6 : $\underline{0} \in \operatorname{int}(\Phi \setminus \Phi_Q)$. ♢

Of course, these results are much weaker than what one might expect to be true. There are many questions remaining, for example: Is $\operatorname{int}(\Phi \setminus \Phi_Q)$ dense in $\Phi \setminus \Phi_Q$? Does $\Phi \setminus \Phi_Q$ or $\operatorname{int}(\Phi \setminus \Phi_Q)$ contain a dense orbit? Are all orbits meeting some neighbourhood of $\underline{0}$ dense? Are "most" orbits dense, in some sense?

In addition to the dynamics on the space $\Phi \setminus \Phi_Q$, one might also consider the behaviour of individual Markoff maps lying in $\Phi \setminus \Phi_Q$. In general, these seem difficult to analyse. One possible direction is to study how they behave on the "circle at infinity", as we go on to explain.

Let Σ^∞ be the cyclically ordered Cantor set of ends of Σ. Let P be the circle obtained by collapsing each pair of adjacent "boundary points" of Σ^∞ to a point. Another way to describe P is to imagine Σ as being dual to the regular tessellation of the hyperbolic plane, $\mathbf{H}^2$, by ideal triangles, as mentioned in Chapter 1. Now, P may be identified with the ideal boundary, $\partial\mathbf{H}^2$. We may thus think of Ω as being the "rational points" of P. Note that in this context, $\mathbf{H}^2$ is naturally thought of as the Teichmüller space of the punctured torus, and so P is the projective lamination space (see for example [16]). Since a "rational lamination" is essentially a simple closed curve, this gives us the identification of Ω and $\mathcal{C}$ referred to in Chapter 1.

Given $\phi \in \Phi$, we may associate a closed subset $L(\phi) \subseteq P$ as follows. If $\phi \in \Phi_Q$, we set $L(\phi) = \emptyset$. If $\phi \notin \Phi_Q$, we begin by putting an orientation on certain edges of Σ as follows. If $e \in E(\Sigma)$, we give e the direction $\vec{e}$ if $\log^+|\phi|$ has a lower Fibonacci bound on $\Omega^-(\vec{e})$. Thus, each edge is oriented in at most one direction. Also, for each vertex $v \in V(\Sigma)$, there are at most three possibilities. Perhaps none of the incident edges are oriented; perhaps one incident edge points towards v while the other two are unoriented; or perhaps two edges point towards v while the other points away. Let $T \subseteq \Sigma$ be the union of all unoriented edges. If $T \neq \emptyset$, then it is a tree, all of whose vertices have degree 2 or 3. Thus, T can be thought of as the "convex hull" of some closed subset $\tilde{L}(\phi) \subseteq \Sigma^\infty$. Moreover all the edges of $\Sigma \setminus T$ point towards T. If $T = \emptyset$, then all edges of Σ point towards some $p \in \Sigma^\infty$. In this case set $\tilde{L}(\phi) = \{p\}$. (It's not hard to see from Lemma 3.3, that p cannot be a boundary point of Σ^∞.) In each case we set $L(\phi)$ to be the projection of $\tilde{L}(\phi)$ to P.

There are other ways one might attempt to associate to ϕ a closed subset of P, but $L(\phi)$ seems fairly natural. It has a number of other descriptions. For example, note that $p \in P \setminus L(\phi)$ if and only if either (1) $p \in P \setminus \Omega$, and given any $K \geq 0$, there is a neighbourhood N of p in P such that $|\phi(X)| \geq K$ for all $X \in N \cap \Omega$; or (2) $p \in \Omega$, $\phi(p) \neq \pm 2$, and given any $K \geq 0$, there is a neighbourhood N of p in P such that $|\phi(X)| \geq K$ for all $X \in N \cap \Omega \setminus \{p\}$.

By definition, $L(\phi) = \emptyset$ if and only if $\phi \in \Phi_Q$. Note that $L(\underline{0}) = P$. We shall see (Proposition 5.7) that if $\phi \neq \underline{0}$, then $L(\phi)$ has empty interior. If $\phi \in \partial\Phi_{QF}$, so that ϕ corresponds to a representation ρ in the boundary of quasifuchsian space, then $L(\phi)$ consists of either one or two elements. In this case, a rational point of $L(\phi)$ corresponds to a curve which has degenerated to an accidental parabolic. An irrational point of $L(\phi)$ corresponds to the ending lamination of a simply degenerate end. We have also seen examples of $\phi \in \mathrm{int}(\Phi \setminus \Phi_Q)$ for which $L(\phi)$ consists of a single rational point. (For example start with the Markoff triple $(0, \mu, i\mu)$ for large enough $\mu \in \mathbf{R}$. The fact that ϕ lies in the interior of $\Phi \setminus \Phi_Q$ follows from Theorem 5.5(1).) One might expect that for generic $\phi \in \Phi \setminus \Phi_Q$, $L(\phi)$ should be a Cantor set, though I have no explict example of a Markoff map for which this is the case. It also seems reasonable to ask to what extent $L(\phi)$ determines ϕ. For example:

Question D : If $L(\phi) = L(\phi')$ and $|L(\phi)| \geq 2$, then must ϕ and ϕ' be equal up to the action of $\mathbf{Z}_2^2$?

Proposition 5.7 : *If $\phi \neq \underline{0}$, then $L(\phi)$ has empty interior.*

Proof : We first observe that if $X \in \Omega$, and $x = \phi(X) \notin [-2, 2]$, then $X \notin L(\phi)$. This follows easily from Lemma 3.8.

Now if $J \subseteq L(\phi)$ were some interval, then there must be three regions $X, Y, Z \in J \cap \Omega$ which meet at some vertex of Σ. ¿From the above observation, we see that $x, y, z \in [-2, 2]$. Since (x, y, z) is a Markoff triple, we must have $x = y = z = 0$, so $\phi = \underline{0}$. This proves Proposition 5.7. $\diamondsuit$

It would be interesting to understand something of the way in which $L(\phi)$ varies with $\phi \in \Phi$.

Let $\mathcal{F} = \mathcal{F}(S)$ be the set of all closed subsets of P. We give $\mathcal{F}$ the Hausdorff topology. (By this we mean the topology induced by the Hausdorff distance on $\mathcal{F}$, starting with some metric on P. Since P is compact the topology on $\mathcal{F}$ thus defined is independent of the choice of metric on P.) The set $\mathcal{F}$ is also carries a partial order given by set inclusion. We can thus speak of a map to $\mathcal{F}$ as being "upper (or lower) semicontinuous".

Proposition 5.8 : *The map $[\phi \mapsto L(\phi)] : \Phi \longrightarrow \mathcal{F}$ is upper semicontinuous.*

In other words, given any $\phi \in \Phi$, and an open set, $M \subseteq P$ with $L(\phi) \subseteq M$, then there exists a neighbourhood N of ϕ in Φ such that for all $\phi' \in N$, $L(\phi') \subseteq M$.

Note that this is a generalisation of Theorem 3.16 (that Φ_Q is open), since if $\phi \in \Phi_Q$, we can take $M = \emptyset$. In this case, Proposition 5.8 gives us a neighbourhood, N, of ϕ with $N \subseteq \Phi_Q$.

To prove Proposition 5.8, we shall need the following:

Lemma 5.9 : *Suppose $\vec{e} \in \vec{E}(\Sigma)$. The set of ϕ in Φ for which ϕ has Fibonacci growth on $\Omega^-(\vec{e})$ is open in Φ.*

Proof : This is just a refinement of the argument used to prove Theorem 3.16 (cf. Proposition 3.9). Given $\phi \in \Phi$, recall the definition of $T(t)$ given before Lemma 3.15. We saw that for all $t > 0$, $T(t)$ is a subtree of Σ. Exactly as with Lemma 3.15, we see that ϕ has Fibonacci growth on $\Omega^-(\vec{e})$ if and only if $T(t) \cap \Sigma^-$ is finite. We now go through the proof of Theorem 3.16, restricting the discussion to Σ^- and Ω^-.

This proves Lemma 5.9. $\diamondsuit$

We shall also need the idea of the "impression" of a branch $\Omega^-(\vec{e})$. Given $\vec{e} \in \vec{E}(\Sigma)$, let $I(\vec{e})$ be the closure, in P of $\Omega^-(\vec{e})$ thought of as a subset of P. Thus $I(\vec{e})$ is a closed interval in P. Note that given $\phi \in \Phi$, the open set $P \setminus L(\phi)$ is covered by the set of impressions, $I(\vec{e})$, as $\vec{e}$ varies over those directed edges for which ϕ has Fibonacci growth on $\Omega^-(\vec{e})$. These impressions are all disjoint from $L(\phi)$. (In fact, we could restrict to those $\vec{e} \in \vec{E}(\Sigma)$ for which $e \cap T$ is the head of $\vec{e}$, where $T \subseteq \Sigma$ is the "convex hull" of $\tilde{L}(\phi)$ as described in the definition of $L(\phi)$.)

Proof of Proposition 5.8 : Suppose $\phi \in \Phi$. Let M be an open subset of P containing $L(\phi)$. We may cover $S \setminus M$ by a finite number of impressions, $I(\vec{e}_1), \ldots I(\vec{e}_n)$, where ϕ has Fibonacci growth on $\Omega^-(\vec{e}_i)$ for each $i \in \{1, \ldots, n\}$. Now by Lemma 5.9, there is a neighbourhood, N, of ϕ in Φ, such that if $\phi' \in N$, then ϕ' has Fibonacci growth on $\Omega^-(\vec{e}_i)$

for each $i \in \{1, \ldots, n\}$. It follows that $L(\phi')$ does not meet any of the impressions $I(\vec{e}_i)$. Thus $L(\phi) \subseteq M$.

This proves Proposition 5.8. $\diamondsuit$

In general, we would not expect L to be lower semicontinuous. However, we do have:

Proposition 5.10 : *The map* $[\phi \mapsto L(\phi)] : \Phi \longrightarrow \mathcal{F}$ *is continuous at* $\underline{0}$.

In other words, if $\{I_1, \ldots, I_n\}$ is a covering of P by (non-trivial) intervals, then there is a neighbourhood, N, of $\underline{0}$ in Φ such that for all $\phi \in N$, $L(\phi)$ meets each I_i. Note that we can suppose that each I_i has the form $I(\vec{e}_i)$ for some $\vec{e}_i \in \vec{E}(\Sigma)$. The proof of Proposition 5.10 is based on the observation that if ϕ is sufficiently close to $\underline{0}$, then then we can assume that for each i, $|\phi(X_i)|$ and $|\phi(Y_i)|$ are arbitrarily small, where $X_i \cap Y_i = e_i$. For the case $\phi \notin \Phi^I$, the idea is now to apply Lemma 5.4 (or rather a slight refinement which gives us, instead of just one region, Y, adjacent to X, some finite number, p, of such regions, provided that $\phi(X) \notin \mathbf{R}$, and $|\phi(X)| \leq \epsilon_0(p)$ for some $\epsilon_0(p)$ depending on p — the method of proof is the same). This ensures that if $|\phi(X_i)|$ and $|\phi(Y_i)|$ are small enough, then ϕ cannot have Fibonacci growth on $\Omega^-(\vec{e}_i)$ so $L(\phi)$ must meet $I_i = I(\vec{e}_i)$. The case where $\phi \in \Phi^I$ calls for Lemma 5.2 in place of Lemma 5.4, but the argument is similar. We leave the reader to fill in the details of the proof.

We remark that there is also a sort of converse to Proposition 5.10; namely that if $L(\phi)$ is sufficiently close to P in $\mathcal{F}$, then ϕ must be close to $\underline{0}$. More precisely, given any neighbourhood, N, of $\underline{0}$ in Φ, there is a neighbourhood, $\mathcal{M}$, of P in $\mathcal{F}$, such that if $\phi \in \Phi$ and $L(\phi) \in \mathcal{M}$ then $\phi \in N$. The argument is roughly as follows. Fix some vertex $v = X \cap Y \cap Z$ of Σ, and suppose that $L(\phi)$ is close to P. Using Proposition 3.9, we can show that every edge within a given distance of v in Σ must be adjacent to some region W with $|\phi(W)| \leq 2$ (since both $I(\vec{e})$ and $I(-\vec{e})$ can be assumed to meet $L(\phi)$). This implies that each of $\phi(X)$, $\phi(Y)$ and $\phi(Z)$ lie close to the interval $[-2, 2]$. Since they form a Markoff triple, they must all in fact be close to 0. Again the details are left to the reader.

It would be nice to give a geometric interpretation to some of the results of this section. One possible direction is to investigate the realisability of geodesic laminations. These notions were introduced by Thurston [43]. His account has been elaborated on in [15]. The latter will serve as a reference to the results about laminations quoted in this section. Also the techniques developed there can be used to supply details to the arguments outlined below. These and other published accounts I know of only consider the case of discrete representations. We describe here how some of the theory might be developed in general, allowing for non-discrete representations. We suggest some natural questions arising from these considerations.

We may as well start in a general situation with an orientable surface $S = S_{g,p}$ of genus g and p punctures. To define the various concepts, it will be useful to fix some complete finite area hyperbolic structure on S. We thus identify $S = \mathbf{H}^2/\Gamma_{g,p}$ where we have chosen an action of $\Gamma_{g,p} = \pi_1(S)$ on $\mathbf{H}^2$. It turns out that all the notions we define are actually independent of this choice.

A *(geometric) lamination* on S consists of a set, Λ, of disjoint simple geodesics on S such that $\bigcup \Lambda \subseteq S$ is closed. By "simple geodesic" we understand either a simple closed geodesic, or a bi-infinite geodesic with no self intersections. An element of Λ is referred to as a *leaf*. The set $\bigcup \Lambda$ is the *support* of Λ. A *minimal lamination* is one which contains no proper non-empty sublamination, and is not a bi-infinite geodesic with both ends going up a cusp of S. A minimal lamination has compact support. It is either a single simple closed curve or else is uncountable. In the latter case the support is locally homeomorphic to a Cantor set times an interval. Every lamination on S consists of a finite disjoint union of minimal sublaminations, together with a finite number of additional bi-infinite geodesics whose ends either go up a cusp or "spiral" onto a minimal sublamination.

As suggested earlier, a the notion of a lamination turns out to be a well defined topological concept, defined independently of any particular structure on S. (In fact, we can think of structure we have just defined as being really the "geometric realisation" of a particular lamination for a given hyperbolic structure. This realisation varies continuously as we move about in Teichmüller space.) Let $\mathcal{G} = \mathcal{G}_{g,p}$ be the set of all laminations on S. Now $\mathcal{G}$ carries a natural (Chabauty) topology — see [15]. It will be convenient here to restrict attention to the set $\mathcal{G}^0 = \mathcal{G}^0_{g,p}$ of laminations with compact support. It will not be hard to see how one might proceed in general.

According to Thurston [43], a "realisation" of a lamination, Λ, in a hyperbolic 3-manifold, M, is essentially a π_1-injective map from S to M such that each leaf of Λ gets mapped locally homeomorpically onto a geodesic in M. We can thus speak of a lamination being "realisable" with respect to a given homotopy class of π_1-injective maps from S to M. (These maps are usually taken to by "type-preserving".) In fact, given such a class of maps, we may as well work in the cover of M given by the subgroup of $\pi_1(M)$ which is the image of $\Gamma = \pi_1(S)$. Thus, realisability is really a property of the lamination and a representation ρ from Γ to $PSL(2, \mathbf{C})$, or equivalently to $SL(2, \mathbf{C})$.

To generalise this to non-discrete lamination, we work in the universal cover. We write $S = \mathbf{H}^2/\Gamma$. A lamination Λ can thus be thought of as the quotient of a Γ-invariant set, $\tilde{\Lambda}$, of disjoint geodesics in $\mathbf{H}^2$. Now fix some representation $\rho : \Gamma \longrightarrow SL(2, \mathbf{C})$, and consider a (continuous) Γ-equivariant map $R : \mathbf{H}^2 \longrightarrow \mathbf{H}^3$. Such a map will be coarsely lipschitz when restricted to any Γ-invariant subset of $\mathbf{H}^2$ which projects to a compact subset of S (for example the support of a lamination in $\mathcal{G}^0$). By "coarsely lipschitz" we mean that distances in the range are bounded by a linear function of distances in the domain.

Definition : The Γ-equivariant map $R : \mathbf{H}^2 \longrightarrow \mathbf{H}^3$ is a *realisation* of the compactly supported lamination Λ with respect to the representation ρ if there are constants k_1 and k_2 such that for any $\lambda \in \tilde{\Lambda}$ and any $x, y \in \lambda$, we have $d(R(x), R(y)) \geq k_1 d(x, y) - k_2$.

Note that since $R|\bigcup \tilde{\Lambda}$ is coarsely lipschitz, for each leaf $\lambda \in \tilde{\Lambda}$, the map $R|\lambda$ is a quasi-geodesic in $\mathbf{H}^3$. In fact the maps $R|\lambda$ are uniformly quasigeodesic as λ varies over each $\lambda \in \Lambda$. (It is not hard to see that it is enough for R to be quasigeodesic for a finite set of leaves consisting of some choice of leaf from each minimal sublamination together with the set of spiralling leaves.)

Suppose $F : \mathbf{H}^2 \longrightarrow \mathbf{H}^3$ is another Γ-equivariant map. Now $R|\bigcup \tilde{\Lambda}$ and $F|\bigcup \tilde{\Lambda}$ remain a bounded distance apart. We see that R is a realisation of Λ if and only if F is.

Definition : A (compactly supported) lamination is realisable if admits a realisation.

In such a case any Γ-equivariant map from $\mathbf{H}^2$ to $\mathbf{H}^3$ will serve as a realisation. Also, it's not hard to see that whether or not a given lamination is realisable does not depend on the hyperbolic metric on S — if we change to a different point in Teichmüller space, then the new metric on $\bigcup \tilde{\Lambda}$ will be quasiisometric to the old, so the property of a map being (uniformly) quasigeodesic remains unchanged.

Defining realisations in terms of quasigeodesics, rather than geodesics, makes certain arguments simpler and more natural (e.g. Propostion 5.11), though we shall see that the result is the same.

Let $G(\mathbf{H}^n)$ be the set of (unoriented) geodesics in $\mathbf{H}^n$, thought of as subsets of $\mathbf{H}^n$. Now $G(\mathbf{H}^n)$ caries a natural topology (for example, it may be identified as a quotient of $\partial\mathbf{H}^n \times \partial\mathbf{H}^n$ minus the diagonal by the involution which swaps the two coordinates). Thus if Λ is a lamination, then $\tilde{\Lambda}$ is a a closed Γ-invariant subset of $G(\mathbf{H}^2)$, and thus inherits a topology.

Now, if R is a realisation of Λ, and $\lambda \in \tilde{\Lambda}$, then since $R|\lambda$ is a quasigeodesic, it must remain a bounded distance from a unique geodesic $G(\lambda) \in G(\mathbf{H}^3)$. It's not hard to see that the map $G : \tilde{\Lambda} \longrightarrow G(\mathbf{H}^3)$ is continuous, and independent of the choice of realisation R. Given $x \in \lambda \in \tilde{\Lambda}$, let $R'(x)$ be the nearest point on $G(\lambda)$ to x. Since G is continuous, it follows that $R' : \bigcup \tilde{\Lambda} \longrightarrow \mathbf{H}^3$ is also. It is clearly Γ-equivariant. Moreover it extends to a Γ-equivariant map from $\mathbf{H}^2$ to $\mathbf{H}^3$, also denoted by R'. If $\lambda \in \tilde{\Lambda}$, then $R'(\lambda) = G(\lambda)$. In other words we see that if Λ is realisable, then we can choose the realisation so that the image of each leaf in $\tilde{\Lambda}$ is a geodesic in $\mathbf{H}^3$.

With some amount of additional fussing around, one can find a realisation which sends each leaf of $\tilde{\Lambda}$ homeomorphically onto a geodesic in $\mathbf{H}^3$. This shows that our definition of realisability agrees with the usual one in the case of discrete representations (see Definition 5.3.4 of [15]).

Recall that the space $\mathcal{G}^0$ of compactly supported laminations carries a natural topology.

Proposition 5.11 : *For any representation ρ, the set of realisable (compactly supported) laminations is open in $\mathcal{G}^0$.*

Proof : (Sketch) Suppose $\Lambda \in \mathcal{G}^0$ is realisable. Given any constant $h > 0$, there is a neighbourhood $\mathcal{N}$ of Λ in $\mathcal{G}^0$, such that if $\Upsilon \in \mathcal{N}$, then every segment of a leaf of length at most h in $\tilde{\Upsilon}$ lies close to a leaf of $\tilde{\Lambda}$. This follows more or less directly from the definition of the Chabauty topology [15]. By "close" we mean close in the Hausdorff distance — so each segment lies in a neighbourhood of radius 1, say, about the other segment. Now, if R is a realisation of Λ, then for each $\lambda \in \Upsilon$, the path $R|\lambda$ must be quasigeodesic over all segments of length at most h. We now use the well-known fact that being quasigeodesic in a hyperbolic space is really a "local" property, in the sense that if we choose h large enough in relation to the quasigeodesic constants, then these paths must be globally (and uniformly) quasigeodesic. The result now follows — taking care to choose our various constants in the right order. ♢

Let's now consider the case of a once punctured torus. We restrict attention to minimal laminations. Now every minimal lamination (on any surface) carries a transverse measure of full support. On a punctured torus such a measure will be unique up to a multiplicative constant. We can thus identify the space of minimal laminations with the projective lamination space P. Moreover the topologies agree. We can thus ask which points of P correspond to realisable laminations. It is not hard to see that a realisable lamination must lie in $P \setminus L(\phi)$, where ϕ is the Markoff map corresponding to our representation ρ. It seems likely that the converse holds:

Conjecture E : $L(\phi)$ consists precisely of those laminations which are not realisable with respect to the representation corresponding to ϕ.

Returning to a general surface, $S = S_{g,p}$, the following variant of Question B (Section 4) suggests itself:

Conjecture F : An admissible representation is quasifuchsian if and only if every minimal lamination is realisable.

It's well known that the "only if" direction is true. The converse is also true if we restrict attention to discrete representations. This follows since, by the work of Bonahon [5], if a discrete representation is not quasifuchsian, then it must have an accidental parabolic or a simply degenerate end. The simple closed curve corresponding to an accidental parabolic is not realisable, nor is the "ending lamination" of a simply degenerate end. (Here the term "end" is defined geometrically. In general an end will retract topologically onto some subsurface of S whose bounding curves give rise to accidental parabolics.)

We finish this section with a hint at some tentative connections with physics. (I'm indebted to John Roberts for introducing me to these ideas.)

Suppose we fix a point, $p \in P$. An immediate corollary of Proposition 5.8 is that $\{\phi \in \Phi \mid p \in L(\phi)\}$ is a closed subset of Φ. We suggested that for a "generic" $\phi \in \Phi \setminus \Phi_Q$, one might expect $L(\phi)$ to be a Cantor set. Put together, this suggests that, if $[t \mapsto \phi_t]$ is a "generic" one-real-parameter family of Markoff maps, then $\{t \mid p \in L(\phi_t)\}$ should be a Cantor set (if it is non-empty).

This appears to be related to a conjecture concerning the spectrum of the discrete one-dimensional quasiperiodic Schrödinger operator. A recent survey of work in this area can be found in [42]. Basically, this is a discrete version of the usual Schrödinger operator, which acts on the space of discrete wave functions from the integers to the reals. It depends on a "potential" which is a map, $V : \mathbf{Z} \longrightarrow \mathbf{R}$. In the case of interest, V takes only two values, $V_0, V_1 \in \mathbf{R}$, and is such that $\{V^{-1}(V_0), V^{-1}(V_1)\}$ is a quasiperiodic partition of the integers, $\mathbf{Z}$. One can generate such a partition by taking a line of gradient θ in the euclidean plane, $\mathbf{R}^2$, and reading off the successive intercepts with lines of the form $\{n\} \times \mathbf{R}$ or $\mathbf{R} \times \{n\}$ for $n \in \mathbf{Z}$ (cf. [40]). (Phyisicists seem to be mainly interested in those sequences arising from lines passing through the origin, $(0,0)$.) Given such a potential, V, the spectrum is the set of eigenvalues or "energies", E, which admit a bounded eigenfunction.

There is a conjecture that this spectrum is always equal to the "dynamical spectrum". The latter can be interpreted in terms of Markoff maps (or generalised real versions thereof, as discussed in the introduction, with $\mu = 4 + (V_0 - V_1)^2$.) Note that a quasiperiodic potential, V, determines a gradient, θ (which can be thought of as the ratio of V_1's to V_0's), which in turn determines a projective lamination $p \in P$. The potential, together with an energy, $E \in \mathbf{R}$, gives rise, in a natural way to a generalised real Markoff map, ϕ_E (corresponding to the triple $(E - V_0, E - V_1, (E - V_0)(E - V_1) - 2)$, as explained in [36]). If we fix V, and vary E, we get a one parameter family $[E \mapsto \phi_E]$ of such maps. The dynamical spectrum is then essentially the set $\{E \mid p \in L(\phi_E)\}$. A major conjecture in this field is that this spectrum is always a Cantor set. Some analysis of this question is given in [36]. There seem to be parallels between this work and some of the results of this section.

References.

[1] K.Astala, *Area distorsion of quasiconformal mappings* : Acta Math. **173** (1994) 37–60.

[2] A.F.Beardon, *The geometry of discrete groups* : Graduate Texts in Mathematics, No. 91, Springer-Verlag (1982).

[3] L.Bers, *Uniformization, moduli and Kleinian groups* : Bull. London Math. Soc. **4** (1972) 257–300.

[4] J.Birman, C.M.Series, *Geodesics with bounded intersection are sparse* : Topology **24** (1985) 217–225.

[5] F.Bonahon, *Bouts des variétés hyperboliques de dimension 3,* : Ann. Math. **124** (1986), 71–158.

[6] B.H.Bowditch, *A proof of McShane's identity via Markoff triples* : Bull. London Math. Soc. **28** (1996) 73–78.

[7] B.H.Bowditch, *A variation of McShane's identity for once-punctured torus bundles* : Topology **36** (1997) 325-334.

[8] B.H.Bowditch, *Spaces of geometrically finite representations* : to appear in Ann. Acad. Sci. Fenn. Math.

[9] B.H.Bowditch, D.B.A.Epstein, *Natural triangulations associated to a surface* : Topology **27** (1988) 91–117.

[10] B.H.Bowditch, C.Maclachlan, A.W.Reid, *Arithmetic hyperbolic surface bundles* : Math. Annalen **302** (1995) 31–60.

[11] R.Brooks, J.P.Matelski, *The dynamics of 2-generator subgroups of* $PSL(2,\mathbf{C})$: in "Riemann surfaces and related topics: proceedings of the 1978 Stony Brook conference", ed. I. Kra and B. Maskit, Annals of Mathematics Studies No. 97, Princeton University

Press and University of Tokyo Press (1981) 65–71.

[12] S.Bullett, C.Penrose, *Mating quadratic maps with the modular group* : Invent. Math. **115** (1994) 483–511.

[13] S.Bullett, C.Penrose, *Perturbing circle packing Kleinian groups as correspondences* : preprint.

[14] J.O.Button, *Quasifuchsian space of the once-punctured torus* : thesis, University of Warwick (1994).

[15] R.D.Canary, D.B.A.Epstein, P.Green, *Notes on notes of Thurston* : in “Analytic and geometric aspects of hyperbolic space” London Math. Soc. Lecture Notes Series No. 111, ed. D. B. A. Epstein, Cambridge University Press (1987).

[16] A.J.Casson, S.A.Bleiler, *Automorphisms of surfaces after Nielsen and Thurston* : L.M.S. Student Texts No. 9, Cambridge University Press (1988).

[17] T.W.Cusick, M.E.Flahive, *The Markoff and Lagrange spectra* : Amer. Math. Soc. Mathematical Surveys and Monographs No. 30, Providence, Rhode Island (1989).

[18] W.Floyd, A.Hatcher, *Incompressible surfaces in punctured torus bundles* : Topology and its appl. **13** (1982) 263–282.

[19] F.P.Gardiner, *Teichmüller theory and quadratic differentials* : Wiley (1987).

[20] F.W.Gehring, G.J.Martin, *Commutators, collars and the geometry of Möbius groups* : J. d’Analyse Math. **63** (1994) 175–219.

[21] F.W.Gehring, J.Väisälä, *Hausdorff dimension and quasiconformal mappings* : J. London Math. Soc. **6** (1973) 504-512.

[22] W.M.Goldman, *Topological components of spaces of representations,* : Invent. Math. **93** (1988) 557–607.

[23] A.Haas, *Diophantine approximation on hyperbolic Riemann surfaces* : Acta Math. **156** (1986) 33–82.

[24] J.L.Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface* : Invent. Math. **84** (1) (1986) 157–176.

[25] T.Jørgensen, *On discrete groups of Möbius transformations* : Amer. J. Math. **98** (1976) 739–749.

[26] T.Jørgensen, *On pairs of punctured tori* : Unpublished manuscript.

[27] L.Keen, *Teichmuller spaces of punctured tori 1 & 2* : Complex variable theory appl. **2** (1983) 199–221.

[28] L.Keen, C.Series, *The Riley slice of Schottky space* : Proc. London Math. Soc. **69** (1994) 72–90.

[29] R.C.Lyndon, P.E.Schupp, *Combinatorial group theory* : Springer Verlag (1977).

[30] A.Marden, *The geometry of finitely generated Kleinian groups* : Ann. Math. **99** (1974)

383–462.

[31] A.A.Markoff, *Sur les formes binaires indéfinies* : Math. Annalen **17** (1880) 379–399.

[32] G.McShane, *A remarkable identity for lengths of curves* : thesis, University of Warwick (1991).

[33] Y.N.Minsky, *The classification of punctured torus groups* : preprint, Stony Brook (1995).

[34] J.-P. Otal, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension trois* : Astérisque No. 235, Sociét'e Mathématique de France (1996).

[35] J.R.Parker, C.M.Series, *Bending formulae for convex hull boundaries* : to appear in J. d'Analyse.

[36] J.A.G.Roberts, *Escaping orbits in trace maps* : to appear in Physica A.

[37] J.A.G.Roberts, M.Baake, *Trace maps as 3D reversible dynamical systems with an invariant* : J. Stat. Phys. 74 (1994) 829–888.

[38] T.Sasaki, *A fundamental domain for some quasifuchsian groups* : Osaka J. Math. **27** (1990) 67–80.

[39] P.S.Schaller (=P.Schmutz), *The Markoff torus has maximal length spectrum* : preprint, Lausanne (1995).

[40] C.M.Series, *The geometry of Markoff numbers* : Math. Intelligencer, Vol. 7 No. 3 (1985) 20–29.

[41] J.Silverman, *The Markoff equation* $X^2 + Y^2 + Z^2 = aXYZ$ *over quadratic imaginary fields* : J. Number Theory **35** (1990) 72–104.

[42] A.Sütő, *Schrödinger difference equations with determinisitic ergodic potientials* : Chapter 17 of "Quasicrystals and beyond" (ed. F.Axel, D.Gratias), Springer (1995) 481–549.

[43] W.P.Thurston, *The geometry and topology of three-manifolds* : Notes, Princeton University (1979).