FINITENESS RESULTS FOR FLAT SURFACES: LARGE CUSPS AND SHORT GEODESICS

JOHN SMILLIE AND BARAK WEISS

ABSTRACT. For fixed g and T we show the finiteness of the set of affine equivalence classes of flat surfaces of genus g whose Veech group contains a cusp of hyperbolic co-area less than T. We obtain new restrictions on Veech groups: we show that any non-elementary Fuchsian group can appear only finitely many times in a fixed stratum, that any non-elementary Veech group is of finite index in its normalizer, and that the quotient of $\mathbb{H}$ by a non-lattice Veech group admits arbitrarily large embedded disks. A key ingredient of the proof is the finiteness of the set of affine equivalence classes of flat surfaces of genus g whose Veech group contains a hyperbolic element with eigenvalue less than T.

1. INTRODUCTION

Our objects of study are flat surfaces and their affine automorphism groups. These structures arise in the study of rational polygonal billiards, in Thurston's classification of surface diffeomorphisms in connection with measured foliations, and in complex analysis. The class of flat surfaces is subdivided into translation surfaces and half-translation surfaces, which correspond in the complex analysis literature to abelian and quadratic differentials respectively.

Let $\mathrm{Aff}(M)$ denote the affine automorphism group of a flat surface M. Let $G = \mathrm{SL}(2, \mathbb{R})$ if M is a translation surface, and $\mathrm{PSL}(2, \mathbb{R})$ if M is a half-translation surface. Taking the differential of an automorphism yields a homomorphism $D : \mathrm{Aff}(M) \to G$ with finite kernel, whose image Γ_M is called the *Veech group* of M. Alternatively, Γ_M is the stabilizer, under the G-action, of M in the stratum $\mathcal{H}$ of flat surfaces containing M. For a typical flat surface $\mathrm{Aff}(M)$ (and hence Γ_M) is trivial; however surfaces with non-trivial Veech groups are quite interesting. Much of the interest in the subject was generated by Veech's discovery [Ve2] of Veech groups which are non-arithmetic lattices (recall that Γ is called a *lattice* if $\mathbb{H}/\Gamma$ has finite area). Many additional Veech groups have been described, including infinitely generated ones [HuSc2, Mc], and there are many intriguing questions regarding groups

which may arise as Veech groups. See [HuSc3] for a recent survey of the field, and [SmWe1] for a problem list.

A fundamental question in this regard is to understand the commensurability classes of Fuchsian groups which contain Veech groups (recall that a *Fuchsian group* is a discrete subgroup of G and two of such groups are *commensurable* if their intersection is of finite index in both). We will describe some restrictions on groups commensurable to a Veech group Γ and on the geometry of the corresponding $\mathbb{H}/\Gamma$. The question of which groups within a given commensurability class actually arise as Veech groups is also of considerable interest, see [HuSc1].

It is clear that if Γ is a lattice, there is an upper limit on the radius of a maximal embedded disk. We have the following converse, which provides a new characterization of lattice surfaces.

Theorem 1.1. *If Γ is commensurable to a Veech group and is not a lattice then for every R, $\mathbb{H}/\Gamma$ contains an embedded disk of radius R.*

We derive a purely group-theoretic property of Veech groups:

Theorem 1.2. *A non-elementary Fuchsian group commensurable to a Veech group is of finite index in its normalizer.*

These theorems follow from some finiteness results which involve an upper bound on either the covolume of a cusp for Γ_M, or the eigenvalue of a hyperbolic element in Γ_M. For a Fuchsian group Γ, to a cusp in $\mathbb{H}/\Gamma$ one may associate the *cusp area* — the hyperbolic area of a maximal continuous embedded family of parallel horocycles — and the number of cylinders in the corresponding decomposition of the surface. Moreover these quantities are invariant under affine equivalence. Equivalently, in group-theoretic terms, a cusp in $\mathbb{H}/\Gamma$ is determined by the conjugacy class of a maximal parabolic subgroup $P \subset \Gamma$. Associating to each such P the area $t_0(\Gamma, P)$ and cylinder number $m = m(M, P) \in \mathbb{N}$, we have that for all $g \in G$, $\Gamma_{gM} = g\Gamma_M g^{-1}$ contains the maximal parabolic subgroup gPg^{-1} and

$$t_0(g\Gamma_M g^{-1}, gPg^{-1}) = t_0(\Gamma_M, P), \quad m(gM, gPg^{-1}) = m(M, P).$$

Let

$$\mathrm{SC}(m, T) = \{(M, P) : t_0(\Gamma_M, P) \leq T, \, m(M, P) = m\},$$

where M ranges over all flat surfaces, and P ranges over all maximal parabolic subgroups of Γ_M (SC stands for 'small cusp'). We denote by $\widetilde{\mathrm{SC}}(m, T)$ the corresponding set of affine equivalence classes.

The following holds:

Theorem 1.3. *For any $T > 0$ and any $m \in \mathbb{N}$, $\widetilde{\mathrm{SC}}(m, T)$ is finite.*

Our proof yields an explicit bound on $\#\widetilde{\mathrm{SC}}(m,T)$, see Theorem 4.2.

Let μ be haar measure on G and let Γ be a discrete subgroup of G. We denote the covolume of Γ in G by $\bar{\mu}(\Gamma)$; this number depends only on the conjugacy class of Γ. We deduce:

Corollary 1.4. *Let $\mathcal{H}$ be a stratum of flat surfaces and let $R, T > 0$. Then the following sets are finite:*

  (i) *Affine equivalence classes of $M \in \mathcal{H}$ for which $\mathbb{H}/\Gamma_M$ contains no embedded ball of radius R.*
  (ii) *Affine equivalence classes of $M \in \mathcal{H}$ for which $\bar{\mu}(\Gamma_M) < T$.*

Note that Veech [Ve2] constructs lattice surfaces M_n for all $n \neq 4$ (on different strata) such that $\bar{\mu}(\Gamma_{M_n}) < 2\pi$. Thus one cannot omit the hypothesis that M is contained in a fixed stratum in Corollary 1.4. Assertion (ii) was proved independently by Curt McMullen, using the algebraic geometry of moduli space.

Suppose Γ_M contains a hyperbolic element h. We denote the larger eigenvalue of h by $\lambda(h)$ and call it the *eigenvalue* of (M, h). Also, associated to h are *Markov partitions* of M (see §5). We let $p = p(M, h)$ be the minimal number of parallelograms in a Markov partition. For $T > 0$ and $m \in \mathbb{N}$ we define

$$\mathrm{SMP}(p, T) = \{(M, h) : p(M, h) = p, \lambda(h) < T\}$$

(SMP stands for 'simple Markov partition'). As before these quantities are invariant under affine equivalence: if $h \in \Gamma_M$ is hyperbolic then $ghg^{-1} \in \Gamma_{gM}$ is also hyperbolic with $p(M, h) = p(gM, ghg^{-1})$ and $\lambda(h) = \lambda(ghg^{-1})$. We denote the set of affine equivalence classes in SMP by $\widetilde{\mathrm{SMP}}$. Repeating a folklore argument[1] we obtain:

Proposition 1.5 (Thurston, Veech). *For a fixed $T > 0$ and $p \in \mathbb{N}$, $\widetilde{\mathrm{SMP}}(p, T)$ is finite.*

The existence of Markov partitions is sketched in [FLP] and we explain it in detail in an appendix to our paper. In particular we show that bounding p is equivalent to bounding the genus of M. Note that by [Pe], there are pairs (M_n, h_n) with $\lambda(h_n) \to 1$, i.e. one cannot omit a bound on p (or on the topology of M) from the statement.

Recall that the *geodesic flow* is the restriction of the G-action on $\mathcal{H}$ to the one-parameter subgroup of diagonal matrices. There is a bijective correspondence between G-orbits of pairs (M, h) as above, and surfaces with a periodic trajectory under the geodesic flow, with the length of

---

[1]The argument is probably due to Thurston. We were unable to find a suitable reference, but see [Th] for a hint and [Ry] for more details.

the corresponding trajectory equal to $\log \lambda(h)$. Thus Proposition 1.5 is equivalent to the statement, proved by Veech [Ve1], that the number of periodic geodesic trajectories in $\mathcal{H}$ of length at most T is finite.

For a group Γ and $h \in \Gamma$ we write h^Γ for the conjugacy class of h in Γ. From Theorem 1.3 and Proposition 1.5 we derive a restriction on Veech groups:

Corollary 1.6. *Suppose Γ is commensurable to a Veech group. Then for any $T > 0$, the following sets are finite:*

(I) $\{P \subset \Gamma : P$ *is a maximal parabolic,* $t_0(\Gamma, P) \leq T\}$.
(II) $\{h^\Gamma : h \in \Gamma$ *is hyperbolic with* $\lambda(h) < T\}$.
(III) $\{f^\Gamma : f \in \Gamma$ *is elliptic with cone area at most* $T\}$.

Here the *cone area* associated with an elliptic $f \in \Gamma$ is the area of $B(z_f, R)/\langle f \rangle$, where $z_f \in \mathbb{H}$ is the fixed point of f and

$$R = R(f) = \sup \{r : B(z_f, r)/\langle f \rangle \to \mathbb{H}/\Gamma \text{ is injective}\}. \qquad (1)$$

Given a Fuchsian group, it is natural to ask 'how often' it arises as a Veech group. Cyclic parabolic subgroups are associated with cylinder decompositions and are easily described. In any stratum there are infinitely many of them belonging to different G-orbits. Using torus covers one can construct infinitely many flat surfaces M, in different strata, with the same non-elementary Veech group. Also, if $\Gamma = \Gamma_M$ and $g \in G$ normalizes Γ then also $\Gamma = \Gamma_{gM}$. As an application of our results we show that aside from these simple constructions, each group can only appear finitely many times, namely:

Corollary 1.7. *For any stratum $\mathcal{H}$ and any infinite Fuchsian group Γ which is not cyclic parabolic, the set*

$$\{M \in \mathcal{H} : \Gamma_M = \Gamma\} \qquad (2)$$

contains finitely many N-orbits, where N is the normalizer of Γ in G. In particular if Γ is non-elementary, (2) is finite.

2. Basics

In this section we set some notation, and collect standard results.

2.1. **Flat surfaces.** We begin by listing some definitions. For more details we refer the reader to [MaTa, Vo, Zo].

Throughout this paper, S denotes a compact connected orientable surface of genus g. When S is equipped with the structure of a flat surface or a quadratic differential we will usually denote it by M. When confusion may arise we will also use M to denote the underlying surface S. A flat surface admits several equivalent definitions. It may be thought of as an equivalence class of atlases of charts $(U_\alpha, \varphi_\alpha)$ covering all but a finite set $\Sigma = \Sigma_M \subset S$ of *singularities*, such that the transition functions $\varphi_\alpha \circ \varphi_\beta^{-1}$ are of the form $z \mapsto \pm z + c$, and such that for each $\sigma \in \Sigma$ the charts combine to form k-pronged singularity at σ, where $k = k_\sigma \in \mathbb{N}$. If $k_\sigma = 2$ then σ is called a *removable singularity* or *marked point.* We always assume that $\Sigma_M \neq \varnothing$. Two atlases M, M' are *compatible* if $M \cup M'$ is also a quadratic differential. The atlases $M = (U_\alpha, \varphi_\alpha)$, M' are *equivalent* if there is a self-homeomorphism $h : S \to S$ such that $h(\Sigma_M) = \Sigma_{M'}$ and $(h(U_\alpha), \varphi_\alpha \circ h)$ is compatible with M'. A flat surface is called a *translation surface*, or an *abelian differential*, if there is a compatible atlas in which all transition functions are of the form $z \mapsto z + c$. A flat surface which is not a translation surface is called a *half-translation surface*. An *affine automorphism* of M is a self-homeomorphism of S which is affine in each chart. Fix the data giving the number of singularities, the vector $\vec{k} = (k_\sigma)_{\sigma \in \Sigma}$, and the determination whether or not the flat surfaces are translation surfaces; then the set of all quadratic differentials sharing this data is called a *stratum*. Each stratum is equipped with a structure of an affine orbifold, which is locally modelled on a relative cohomology group, see [MaSm, MaTa]. The group G acts on $\mathcal{H}$ by post-composition on each chart in an atlas.

A flat surface inherits from the plane a singular foliation on S called the *horizontal (resp. vertical) foliation*, with each chart foliated into the lines parallel to the x-axis (resp. y-axis). A *saddle connection* is a straight segment joining singularities or punctures, with no singularities or punctures in its interior. The set of all saddle connections in direction θ is denoted by $\mathcal{L}_M(\theta)$.

Throughout this paper we will identify elements of G with matrices of $\mathrm{SL}(2, \mathbb{R})$. We will need the following four one-parameter subgroups of G:

$$g_t = \begin{pmatrix} e^{t/2} & 0 \\ 0 & e^{-t/2} \end{pmatrix}, \quad r_\theta = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix},$$

$$h_s = \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}, \qquad \widetilde{h}_s = \begin{pmatrix} 1 & 0 \\ s & 1 \end{pmatrix}.$$

We say that $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$ are *commensurable* if $\lambda_i/\lambda_j \in \mathbb{Q}$ for all $i, j \in \{1, \ldots, k\}$. If this holds, we denote by $\mathrm{LCM}(\lambda_1, \ldots, \lambda_k)$ the smallest positive number which is an integer multiple of all the λ_i.

2.2. Cylinder decompositions, gluing patterns. A *cylinder* for M is a topological annulus which is isometric to $\mathbb{R}/w\mathbb{Z} \times (0, h)$. It is *maximal* if it is not contained in a larger cylinder, and this implies that both of its boundary components in S contain singularities. The *height, width*, and *inverse modulus* of the cylinder are h, w, and w/h respectively, and a curve which wraps around the cylinder parallel to its boundary is called a *waist curve*. A *cylinder decomposition* is a decomposition of M into maximal cylinders with disjoint interiors. The waist curves of all the cylinders in a cylinder decomposition are parallel. Two cylinder decompositions of M are called *transverse* if the directions of waist curves in each decomposition are different.

2.2.1. *Gluing pattern for a pair of cylinder decompositions.* Suppose for $i = 1, 2$ that we have two transverse cylinder decompositions $M = C_1^{(i)} \cup \cdots \cup C_{m_i}^{(i)}$. For each $i \in \{1, \ldots, m_1\}, j \in \{1, \ldots, m_2\}$, the intersection $C_i^{(1)} \cap C_j^{(2)}$ consists of a finite union of parallelograms. The resulting parallelograms have disjoint interiors and boundary identifications. To describe how the parallelograms are attached to each other we follow ideas of [EsOk].

Suppose first that M is a translation surface. In this case one can consistently label the edges of the parallelograms with labels right, left, top, bottom. Label the parallelograms by $1, \ldots, \ell$, let S_ℓ denote the group of permutations on ℓ elements, and define $\sigma_1, \sigma_2 \in S_\ell$, where $\sigma_1(k_1) = k_2$ (resp. $\sigma_2(k_1) = k_2$) if the right (resp. top) edge of the k_1th parallelogram is attached to the left (resp. bottom) edge of the k_2th parallelogram.

Replacing the labelling of the parallelograms amounts to replacing σ_1, σ_2 with $\tau\sigma_1\tau^{-1}, \tau\sigma_2\tau^{-1}$ for some $\tau \in S_\ell$. We denote the equivalence class of $(\sigma_1, \sigma_2) \in S_\ell \times S_\ell$ under simultaneous conjugations by $[(\sigma_1, \sigma_2)]$, and call $[(\sigma_1, \sigma_2)]$ the *gluing pattern* of M and the given cylinder decompositions. From the gluing pattern it is simple to recover the number of cylinders in each decomposition, the number of connected components in each intersection of cylinders, and the stratum to which M belongs. In particular the following hold:

(i) For $i = 1, 2$, σ_i has m_i cycles, and the length of the jth cycle of σ_1 (resp. σ_2) is $\sum_i a_{ij}$ (resp. $\sum_i a_{ji}$).

(ii) The subgroup of S_ℓ generated by σ_1, σ_2 acts transitively on $\{1, \ldots, \ell\}$.

Moreover M is completely determined by the dimensions and orientation of the parallelogram and the corresponding gluing pattern; indeed such information gives rise to an explicit atlas of charts as in §2.1.

In case M is a half-translation surface one can define a gluing pattern in a more complicated way by subdividing each rectangle into four 'quarter-tiles'. This will not be used in the current paper.

2.2.2. *Gluing pattern for parallelograms.* We will also have occasion to consider flat surfaces made up of finitely many parallelograms, glued along edges. We suppose that M is a translation surface which is a union of closed metric parallelograms $P_1, \ldots, P_m$, with disjoint interiors, such that the two directions of parallel sides are the same for all P_i. For ease of notation we assume that these sides are horizontal and vertical. The interiors of the parallelograms are isometrically embedded, so the singularities of M lie on the boundaries. The connected components of intersections $P_i \cap P_j$ consist of horizontal and vertical segments, which we denote respectively by $\xi_1, \ldots, \xi_{\ell_1}$ and $\eta_1, \ldots, \eta_{\ell_2}$. We do not allow the singularities of M to be in the interior of a segment (if it is we subdivide the segment in two). Each segment is attached to others at its two endpoints. If an endpoint of one of the η_i's is a regular point of M, attached to it are two of the η_i's and at most two of the ξ_j's, and if it is a singularity more segments will be attached. Orienting the ξ-edges from left to right and the η-edges from top to bottom, we obtain a directed graph with two kinds of edges, embedded in M and equipped with two additional structures:

- At each vertex there is a cyclic order for the incident edges, obtained by going around a small neighborhood of the point in M in the counterclockwise direction. Moreover along a cycle, consecutive edges of the same kind have opposite orientations (i.e. if η_i is incoming and η_j is the next η-edge then it is outgoing).
- For each ξ (resp. η) segment there are top and bottom (resp. right and left) labels from $\{1, \ldots, m\}$ indicating the parallelogram glued to the segment on the appropriate side.

We call a graph with these structures a *gluing pattern* for a decomposition into parallelograms. The gluing pattern obeys certain obvious restrictions. For example, at a vertex v, if there is no ξ edge between an incoming and an outgoing η edge then the 'right' labels of the two η edges are the same. Note that the gluing pattern determines the stratum containing M. Since each vertex of a gluing pattern is either a corner of a rectangle or a singularity of M, there are only finitely many

gluing patterns when the number of singularities and the number of parallelograms are fixed.

Also associated with a parallelogram decomposition is the *metric data* consisting of the lengths of sides of the P_i's and the lengths of the ξ- and η-edges. Given the gluing pattern and the lengths of the edges, it is possible to calculate the sidelengths of the P_i's; thus the gluing pattern places some restrictions on the metric information. It is clear that the gluing pattern and the metric data taken together uniquely determine the translation surface M.

One could similarly define a gluing pattern for a parallelogram decomposition of a half-translation surface. This will not be used in the present paper and is left to the avid reader.

2.3. Fuchsian groups.

A Fuchsian group is a discrete subgroup of G. Fixing a haar measure μ on G, we define the *covolume* of Γ in G as the measure of a fundamental domain for the action of Γ on G, and denote the covolume by $\bar{\mu}(\Gamma)$. It is easily checked that if M and M' are affinely equivalent then $\bar{\mu}(\Gamma_M) = \bar{\mu}(\Gamma_{M'})$. If $\bar{\mu}(\Gamma_M) < \infty$ then M is called a *lattice surface*.

An element of a Fuchsian group is called *parabolic* if it is conjugate to h_1, and an automorphism $\varphi \in \mathrm{Aff}(M)$ is called *parabolic* if $D\varphi \in \Gamma_M$ is parabolic. The following is well-known (see [Ve2]):

Proposition 2.1. *If $\varphi \in \mathrm{Aff}(M)$ is parabolic and $D\varphi$ fixes the direction θ, then $M \smallsetminus \mathcal{L}_M(\theta)$ is a cylinder decomposition. For any m there is $C = C_m$ such that if m is the number of cylinders in the above decomposition, then there is $k \leq C_m$ such that $\psi = \varphi^k$ preserves the cylinders and the corresponding inverse moduli $\mu_1, \ldots, \mu_m$ for ψ satisfy $\mathrm{LCM}(\mu_1, \ldots, \mu_m) = \mu$, where $D\psi = r_\theta h_\mu r_{-\theta}$ and $n_i = \mu/\mu_i$ is the number of Dehn twists which ψ induces around a waist curve in the ith cylinder.*

For a Fuchsian group Γ we now define cusp areas and cone areas. An infinite cyclic subgroup of G generated by a parabolic element is called *parabolic*. Suppose that Γ contains a parabolic subgroup P and is *non-elementary*, that is, not a finite extension of an abelian group. Suppose also that P is *maximal*, i.e. not properly contained in a parabolic subgroup of Γ. Choose an element $g \in G$ such that

$$gPg^{-1} = \langle h_1 \rangle, \tag{3}$$

and relabelling, replace P and Γ by gPg^{-1} and $g\Gamma g^{-1}$ respectively. Let $\mathbb{H}$ be the complex upper half plane, let $\mathbf{i} = \sqrt{-1}$, let $\mathcal{C}_t$ be the image of the disk $\left\{ z \in \mathbb{H} : \left| z - \frac{t}{2}\mathbf{i} \right| < \frac{t}{2} \right\}$ in $\mathbb{H}/P$, let $\varphi : \mathbb{H}/P \to \mathbb{H}/\Gamma$ be the

natural map, and let

$$t_0 = t_0(\Gamma, P) = \sup\{t > 0 : \varphi|_{\mathcal{C}_t} \text{ is injective}\}. \tag{4}$$

We will show in Propositions 3.1, 3.2 that the set in the right hand side of (4) is nonempty and bounded above, so that t_0 is well-defined, and that t_0 does not depend on the choice of g in (3) and satisfies $t_0(\Gamma, P) = t_0(x\Gamma x^{-1}, xPx^{-1})$ for $x \in G$. We call t_0 the *cusp area* of P in Γ; a simple computation shows that it is equal to the hyperbolic area of $\mathcal{C}_{t_0}$.

2.4. **Orientation double cover.** Suppose M is a half-translation surface. Then there is a topological branched cover $\pi : \widetilde{S} \to S$ of degree 2 such that the pulled-back flat surface $\widetilde{M}$ is a translation surface called the *orientation double cover* of M, see [DoHu]. We have:

Proposition 2.2. *Let $\widetilde{M} \to M$ be the orientation double cover of a half-translation surface M. Then:*

- *Any automorphism of M lifts to an automorphism of $\widetilde{M}$, with the same derivative, so that $\Gamma_M \subset \Gamma_{\widetilde{M}}$. If a parabolic (resp. hyperbolic) affine automorphism φ of M has an associated cylinder decomposition (resp. Markov partition) with k cylinders (resp. parallelograms) then the lifted automorphism $\widetilde{\varphi}$ of $\widetilde{M}$ has an associated cylinder decomposition (resp. Markov partition) with at most $2k$ cylinders (resp. parallelograms).*
- *Taking orientation double covers is G-equivariant, i.e. for $g \in G$, $g\widetilde{M}$ is the orientation double cover of gM.*
- *Given $\widetilde{M}$ there are at most finitely many M such that $\widetilde{M}$ is the orientation double cover of M.*

Proof. A proof of the fact that any affine automorphism lifts can be found in [Ry]. It implies that there is an inclusion $\Gamma_M \subset \Gamma_{\widetilde{M}}$. The statement about the number of cylinders and rectangles follows immediately from the fact that the degree of the cover is 2. The second statement is immediate from the construction of the orientation double cover. For the third statement see e.g. [Vo]. □

2.5. **Perron-Frobenius.** A matrix $A \in \mathrm{Mat}_d(\mathbb{R})$ with all entries non-negative is called *non-negative*, and a non-negative matrix is called *irreducible* if for some k, A^k has all its entries strictly positive. A vector $\vec{v} \in \mathbb{R}^d$ is called *positive* if all its entries are strictly positive.

We will need the following classical result (see e.g. [Ga]).

Proposition 2.3. *Suppose $A \in \mathrm{Mat}_d(\mathbb{R})$ is an irreducible non-negative matrix. Then:*

- *There is a unique (up to scaling) positive eigenvector v^+ of A, and the corresponding eigenspace is one-dimensional.*
- *Let λ be the eigenvalue for which $Av^+ = \lambda v^+$. Then for any other eigenvalue β of A, $|\beta| < \lambda$.*

We denote the eigenvalue λ of Proposition 2.3 by $\lambda(A)$. We will be particularly interested in irreducible non-negative matrices with integer entries.

Proposition 2.4. *Given $T > 0$ and $d \in \mathbb{N}$, the set*

$$\{A \in \mathrm{Mat}_d(\mathbb{Z}) : A \text{ is irreducible non-negative, } \lambda(A) < T\}$$

is finite.

Proof. See [PaPe, Proof of Theorem 6]. $\qquad\qquad\qquad\qquad\qquad\square$

3. SOME HYPERBOLIC GEOMETRY

This section contains some simple propositions in hyperbolic geometry related to cusp areas and cone areas. We will use the notation introduced in §2.3.

Proposition 3.1. *Let Γ be a Fuchsian group containing the maximal parabolic subgroup $P = \langle h_1 \rangle$. Then $X = \{t > 0 : \varphi|_{\mathcal{C}_t} \text{ is injective}\}$ is nonempty. If Γ is nonelementary then $t_0 = \sup X < \infty$, and there is $\gamma_0 \in \Gamma$ which is a conjugate of h_1, such that for some $s \in \mathbb{R}$, $h_s^{-1} \gamma_0 h_s = \widetilde{h}_{-t_0^2}$. The hyperbolic area of $\mathcal{C}_t$ is t_0.*

Proof. We denote the right-action of Γ on $\mathbb{H}$ by

$$z \cdot \gamma = \frac{az + c}{bz + d}, \quad \text{where } \gamma = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \Gamma.$$

Let $\mathcal{B}_t = \left\{z \in \mathbb{H} : \left|z - \frac{t}{2}\mathbf{i}\right| < \frac{t}{2}\right\}$ and let $H_t = \partial \mathcal{B}_t$. By our conventions $\mathcal{B}_t$ projects to $\mathcal{C}_t$ and H_t is the horocycle based at 0 through $t\mathbf{i}$, which projects to a periodic horocycle on $\mathbb{H}/P$. Let $S \subset H_1$ be a compact segment whose projection to $\mathbb{H}/P$ contains a full period of the horocycle. Since the action of Γ on $\mathbb{H}$ is properly discontinuous, we can write

$$\{\gamma_1, \ldots, \gamma_N\} = \{\gamma \in \Gamma : S \cap S \cdot \gamma \neq \varnothing, \, 0 \cdot \gamma \neq 0\}.$$

To show that $X \neq \varnothing$, take $0 < t \leq 1$ small enough so that $\mathcal{B}_t$ does not intersect $\bigcup_{i=1}^N \mathcal{B}_t \cdot \gamma_i$; such t exist because the $\mathcal{B}_t \cdot \gamma_i$ are not based at 0. Suppose if possible that for some $\gamma \in \Gamma \setminus P$, $\mathcal{B}_t \cap \mathcal{B}_t \cdot \gamma \neq \varnothing$. In particular $\mathcal{B}_1 \cap \mathcal{B}_1 \cdot \gamma \neq \varnothing$. It follows from discreteness of Γ and maximality of P that P is equal to its own normalizer in Γ. This implies that $p = 0 \cdot \gamma \neq 0$; then $H_1 \cap H_1 \cdot \gamma \neq \varnothing$ (see Figure 1) and since

any element of H_1 may be mapped to S by suitable powers of h_1, there are $j, k \in \mathbb{Z}$ such that $S \cap S \cdot \widetilde{\gamma} \neq \varnothing$, where $\widetilde{\gamma} = h_1^j \gamma h_1^k \in \Gamma$. Thus for some $i \in \{1, \dots, N\}$ we have $\widetilde{\gamma} = \gamma_i$ and $\mathcal{B}_t \cap \mathcal{B}_t h_1^{-j} \gamma_i h_1^{-k} \neq \varnothing$. Since $\mathcal{B}_t$ is invariant under h_1, this contradicts the definition of t.


FIGURE 1. Intersection of horocycles

Now if Γ is nonelementary then it contains γ so that $0 \cdot \gamma \neq 0$, so that H_1 and $H_1 \cdot \gamma$ are horocycles based at different points. Then for all sufficiently large t, $H_t \cap H_t \cdot \gamma \neq \varnothing$, which shows $\sup X < \infty$. Further, by definition of t_0, there is $\gamma \in \Gamma$ such that $H_{t_0} \cap H_{t_0} \cdot \gamma \neq \varnothing$ and $\mathcal{B}_{t_0} \cap \mathcal{B}_{t_0} \cdot \gamma = \varnothing$, i.e., H_{t_0} and $H_{t_0} \cdot \gamma$ are horocycles based at different points and tangent to each other. We denote the point of tangency by p_0.

Since $p_0 \gamma^{-1} \in H_{t_0}$, there is s_1 such that
$$p_0 = p_0 \cdot x, \quad H_{t_0} \cdot x = H_{t_0} \cdot \gamma \quad \text{where } x = h_{s_1} \gamma.$$

We now express x in terms of t_0, as follows.

Let $s \in \mathbb{R}$ so that $\widetilde{p}_0 = p_0 \cdot h_s = t_0 \mathbf{i}$. The matrices
$$a = \begin{pmatrix} t_0^{-1/2} & 0 \\ 0 & t_0^{1/2} \end{pmatrix}, \qquad w = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \tag{5}$$

satisfy $\widetilde{p}_0 \cdot a = \mathbf{i} = \mathbf{i} \cdot w$, and $H_1 \cdot w = \widetilde{H}$, where $\widetilde{H} = \{z \in \mathbb{H} : \mathrm{Im} z = 1\} = \{\mathbf{i} \cdot \widetilde{h}_s : s \in \mathbb{R}\}$ (see Figure 2). Thus the matrix $x' = h_s a w a^{-1} h_{-s}$ also satisfies $p_0 = p_0 \cdot x'$, $H_{t_0} \cdot x' = H_{t_0} \cdot \gamma$, and this implies
$$x = \pm h_s a w a^{-1} h_{-s}. \tag{6}$$

FIGURE 2. Tangency of H_{t_0} and $H_{t_0}\gamma$

Let $\gamma_0 = x^{-1}h_1x = \gamma^{-1}h_1\gamma \in \Gamma$. A computation using (6) gives

$$h_s^{-1}\gamma_0 h_s = \widetilde{h}_{-t_0^2}.$$

Finally, note that the hyperbolic area of $\mathcal{C}_{t_0}$ is the same as that of $\mathcal{C}_{t_0} \cdot \gamma h_s$, which by the above is covered bijectively by $\{z = x + \mathbf{i}y \in \mathbb{H} : y \geq t_0, 0 \leq x < t_0^2\}$. Hence its area is

$$\int_0^{t_0^2} \int_{t_0}^\infty \frac{dy}{y^2} \, dx = t_0.$$

$\square$

Proposition 3.2. *The definition of $t_0(\Gamma, P)$ does not depend on the choice of g for which* (3) *holds, and $t_0(x\Gamma x^{-1}, xPx^{-1}) = t_0(\Gamma, P)$.*

Proof. Let the notation be as in the the preceding proof. It was shown in the proof that

$$t_0 = \inf\{t > 0 : H_t \cap \bigcup_{\gamma \in \Gamma \smallsetminus P} H_t\gamma \neq \varnothing\};$$

if we replace g by g' also satisfying (3), then $x = g'g^{-1}$ satisfies $xh_1x^{-1} = h_1$ and hence $x = h_s$ for some s. This implies that $H_tx = H_t$ for all t and hence $H_t \cap H_t\gamma \neq \varnothing$ is equivalent to $H_t \cap H_tx\gamma x^{-1} \neq \varnothing$. This proves the first assertion. The second assertion is an immediate consequence of the first. $\square$

Suppose Γ is a Fuchsian group containing an elliptic element f. The corresponding *cone* in $\mathbb{H}/\Gamma$ is the image of $B(z_f, R(f))$ under the map

$\varphi : \mathbb{H}/\langle f \rangle \to \mathbb{H}/\Gamma$, where $z_f \in \mathbb{H}$ is the fixed point of f and $R(f)$ is defined via (1).

Proposition 3.3. *With the above notation, let $\gamma \in \Gamma$ such that $z_f \gamma$ is a closest point to z_f in the orbit $z_f \Gamma$. Then the commutator $h = [\gamma, f]$ is hyperbolic, its eigenvalue is bounded above by a number depending only on $R(f)$, and the distance from the axis of h to z_f is also bounded by a number depending on $R(f)$.*

Proof. Write $z = z_f, z' = z\gamma$ and $f' = \gamma f \gamma^{-1}$. Clearly $R(f) = d(z, z')/2$. Conjugating Γ with an appropriate element of G and re-labelling we can arrange that $z = t + \mathbf{i}$ and $z' = -t + \mathbf{i}$ are symmetric with respect to the y-axis in $\mathbb{H}$. Using the formula for hyperbolic distance (see e.g. [Ka, Thm. 1.2.6]) we find $\cosh d(z, z') = 1 + 2t^2$.

On the other hand $f = h_t r_\theta h_{-t}$ and $f' = h_{-t} r_\theta h_t$. Computing $h = f' f^{-1}$ directly we find $\operatorname{tr}(h) = 2 + 4t^2 \sin^2 \theta$. Since this number is greater than 2, h is hyperbolic, and since this number is no more than $2 \cosh d(z, z') = 2 \cosh 2R(f)$, and the trace determines the eigenvalue, the second assertion follows.

Since f^{-1} fixes z and f' rotates around z', the distance which z is moved by $h = f' f^{-1}$ is bounded in terms of $R(f)$. On the other hand the order of f bounds $\sin \theta$ from below and and hence implies a lower bound on the displacement of h. Now let p_1, p_2 be the points on the axis of h closest to z and zh respectively, so that $p_2 = p_1 h$; consider the quadrilateral joining z, p_1, p_2, zh. We have bounded $d(p_1, p_2)$ from below and $d(z, zh)$ from above, and the angles at p_1, p_2 are right angles. From this it is easy to deduce an upper bound on the length $d(p_1, z)$. $\square$

4. FINITENESS OF SMALL CUSPS

In this section we prove a more precise version of Theorem 1.3. The idea is that when Γ_M is non-elementary, a cusp in $\mathbb{H}/\Gamma_M$ gives rise to two transverse cylinder decompositions on M, one an image of the other under an affine automorphism, and the cusp area bounds the combinatorics of the corresponding intersection pattern. We will first consider translation surfaces, so denote

$$\mathrm{SC}_1(m, T) = \{(M, P) \in \mathrm{SC}(m, T) : M \text{ is a translation surface}\},$$

and by $\widetilde{\mathrm{SC}}_1(m, T)$ the corresponding sets of G-orbits.

For any $T > 0$ and $m \in \mathbb{N}$, let $\mathcal{N}(m, T)$ denote the set of pairs (A, D), where $A, D \in \operatorname{Mat}_m(\mathbb{Z}_+)$, A is symmetric, D is diagonal, DA is non-negative irreducible and $\lambda(DA) < T$. The following is an immediate corollary of Proposition 2.4.

Proposition 4.1. *For any $m \in \mathbb{N}$ and $T > 0$, $\#\mathcal{N}(m, T) < \infty$.*

Given a symmetric $A = (a_{ij}) \in \text{Mat}_m(\mathbb{Z}_+)$, let $\ell = \sum_{i,j} a_{ij}$. Denote by $\mathcal{P}(A)$ the set of simultaneous conjugacy classes $[(\sigma_1, \sigma_2)]$ for which (i) and (ii) of §2.2 hold, and such that σ_1 and σ_2 are conjugate in S_ℓ.

Theorem 4.2. *For fixed $T > 0$ and $m \in \mathbb{N}$,*

$$\#\widetilde{\text{SC}}_1(m, T) \leq \sum_{(A,D) \in \mathcal{N}(m, C_m T)} \#\mathcal{P}(A).$$

Proof. We will construct a map

$$\Psi : \text{SC}_1(m, T) \to X,$$

where

$$X = \{(A, D, [(\sigma_1, \sigma_2)]) : (A, D) \in \mathcal{N}(m, C_m T), \, [(\sigma_1, \sigma_2)] \in \mathcal{P}(A)\},$$

such that $\Psi(M, P) = \Psi(M', P')$ if and only if there is $g \in G$ such that $M' = gM$ and $P' = gPg^{-1}$; in other words, Ψ induces an injective map defined on $\widetilde{\text{SC}}_1(m, T)$.

Let $(M', P') \in \text{SC}_1(m, T)$, that is M' is a flat surface, $\Gamma_{M'}$ is non-elementary, and $P' \subset \Gamma_{M'}$ is maximal parabolic. By a conjugation we may assume that P' is the cyclic group generated by h_1. Let s, γ_0, h_1, t_0 be as in Proposition 3.1, let a be as (5), let $g = a^{-1}h_{-s}$, and let $M = gM'$. Then $\Gamma_M = g\Gamma_{M'}g^{-1}$ contains the two elements

$$\gamma_1 = h_{t_0}, \qquad \gamma_2 = \widetilde{h}_{-t_0}, \tag{7}$$

and there is $\gamma \in \Gamma_M$ such that

$$\gamma_2 = \gamma^{-1}\gamma_1\gamma. \tag{8}$$

Moreover $P = gP'g^{-1}$ is generated by γ_1.

By Proposition 2.1, there is a parabolic affine automorphism $\varphi_1 \in \text{Aff}(M)$ and $k \leq C_m$ such that $D\varphi_1 = \gamma_1^k = h_{kt_0}$, and φ_1 preserves the cylinders in a horizontal cylinder decomposition $M = C_1 \cup \cdots \cup C_m$. Let $\psi \in \text{Aff}(M)$ such that $\gamma = D\psi$, and let $\varphi_2 = \psi^{-1}\varphi_1\psi \in \text{Aff}(M)$, so that $D\varphi_2 = \gamma_2^k$. Let $C_i' = \psi^{-1}(C_i)$. Then $M = C_1' \cup \cdots \cup C_m'$ is a vertical cylinder decomposition invariant under φ_2.

Let $\vec{a}, \vec{w}$ (resp. $\vec{a}', \vec{w}'$) be positive vectors in $\mathbb{R}^m$ recording the heights and circumferences of the cylinders $C_1, \ldots, C_m$ (resp. $C_1', \ldots, C_m'$). It follows from (7) and (8) that $\gamma = h_{s_1}w\widetilde{h}_{s_2}$ for some $s_1, s_2 \in \mathbb{R}$, where w is as in (5). Since the action of h_s (resp. $\widetilde{h}_s$) does not affect the heights and circumferences of horizontal (resp. vertical) cylinders, and since w affects a rotation by $\pi/2$, we get

$$\vec{a} = \vec{a}', \quad \vec{w} = \vec{w}'.$$

Now let $A = (a_{ij}) \in \mathrm{Mat}_m(\mathbb{Z})$ where a_{ij} is the number of connected components of $C_i \cap C'_j$. By connectedness of M, A is irreducible. Since ψ exchanges the roles of C_i and C'_j, A is symmetric. By construction, the gluing pattern $[(\sigma_1, \sigma_2)]$ corresponding to the two cylinder decompositions above is in $\mathcal{P}(A)$. We have

$$\vec{w} = \vec{w}' = A\vec{a}. \tag{9}$$

Now let n_i be the number of Dehn twists induced by φ_1 around a waist curve in C_i. By Proposition 2.1 we have $kt_0 = n_i w_i / a_i$, that is, if $D = \mathrm{diag}(n_1, \ldots, n_m) \in \mathrm{Mat}_m(\mathbb{N})$, then

$$kt_0 \vec{a} = D\vec{w} = DA\vec{a}.$$

That is, $kt_0 = \lambda(DA)$, $(D, A) \in \mathcal{N}(kT)$ and $k \leq C_m$. Altogether we have $(D, A, [(\sigma_1, \sigma_2)]) = \Psi(M', P) \in X$, so to prove the theorem it remains to show that the G-orbit of (M', P') may be reconstructed from $(D, A, [(\sigma_1, \sigma_2)])$. Note that $\vec{a}$ is uniquely determined up to scaling as the positive eigenvector of DA, and $\vec{w}$ is determined from $\vec{a}$ and A via (9). The scaling parameter is also uniquely determined by the requirement that M has area one. So (A, D) determine $\vec{a}$ and $\vec{w}$, that is, the heights and circumferences of the rectangles, and $[(\sigma_1, \sigma_2)]$ determines how they are to be glued to each other. This determines M, and P is the maximal parabolic subgroup of Γ_M leaving the horizontal direction fixed. $\square$

Proof of Theorem 1.3. By Proposition 2.2, for each $(M, P) \in \mathrm{SC}(m, T)$, the orientation double cover $\widetilde{M}$ also has P as a parabolic subgroup of $\Gamma_{\widetilde{M}}$, and the corresponding cylinder decomposition has at most $2m$ cylinders. Thus, if $\widetilde{P} \subset \Gamma_{\widetilde{M}}$ is a maximal parabolic subgroup containing P, then (M, P) gives rise to $(\widetilde{M}, \widetilde{P}) \in \mathrm{SC}_1(2m, T)$. Moreover for a fixed $(\widetilde{M}, \widetilde{P})$ there are only finitely many (M, P) from which it arises in this way. The Theorem follows. $\square$

Proof of Corollary 1.4. In a given stratum there is an upper bound on the number of cylinders in a cylinder decomposition. See [KeMaSm, Lemma, pg. 302] or Smillie's improved bound as presented in [Na]. Suppose by contradiction that $M_1, M_2, \ldots$ are infinitely many affinely inequivalent surfaces, such that for each i, $\mathbb{H}/\Gamma_i$ does not contain an embedded ball of radius R. Here $\Gamma_i = \Gamma_{M_i}$. By Theorem 1.1 the M_i are lattice surfaces, hence for each i, Γ_i contains at least one maximal parabolic P_i. By Theorem 1.3, $t_0(\Gamma_i, P_i) \to \infty$. A cusp with cusp area t_0 contains an embedded ball of radius $R(t_0)$, where $R(t_0) \to \infty$ as

$t_0 \to \infty$ – a contradiction proving (i). For a lattice surface, the cusp area is bounded above by the covolume, implying (ii). $\qquad\square$

Remark 4.3. *It is also possible to deduce Corollary 1.4 and the finiteness of* $\widetilde{\mathrm{SC}}(m, T) \cap \mathcal{H}$ *for any stratum* $\mathcal{H}$*, from Proposition 1.5, as follows. To any cusp one associates the element* $h = \gamma_1 \gamma_2$ *as in (7) and proves it is hyperbolic, with* $\lambda(h)$ *bounded by a number depending on the corresponding cusp area. Our Theorem 1.3 is stronger in that it does not assume a bound on the topology of the surface, but only on the number of cylinders in the corresponding cylinder decomposition.*

5. Hyperbolic affine automorphisms and Markov partitions

In this section we prove Proposition 1.5. As noted in the introduction, the result is not new and presumably our argument is also well-known. We include it since it does not appear in the literature.

Let M be a flat surface and let φ be a hyperbolic affine automorphism (in a different terminology, φ represents a *pseudo-Anosov homeomorphism*). Let $h = D\varphi$ be the corresponding element of Γ_M, and let $\lambda = \lambda(h)$. By applying $g \in G$, let us assume that h expands the x-axis and contracts the y-axis by a factor of λ. A *Markov partition* for φ is a covering $M = P_1 \cup \cdots \cup P_p$ by closed rectangles, with disjoint interiors, horizontal and vertical sides, such that the following *Markov property* is satisfied: for each i, the image of any vertical (resp. horizontal) side of any P_i under φ (resp. φ^{-1}) is contained in a vertical (resp. horizontal) side of one of the P_j. This means that if $\mathrm{int}\,\varphi(P_i)$ intersects the interior of some P_j then it extends all the way through to both sides. It is known that a Markov partition exists. A sketch of proof is given in [FLP], and we have provided more details in an appendix.

The *intersection matrix of* φ is the $p \times p$ integer matrix $A = (a_{ij})$, where a_{ij} is the number of connected components of $\mathrm{int}\,\varphi(P_i) \cap \mathrm{int}\,P_j$. One has:

Proposition 5.1.
- *A is irreducible.*
- *$\lambda = \lambda(A)$.*
- *The vector recording the widths (resp. heights) of the rectangle P_i is a positive eigenvector for A (resp. A^{tr}).*

Proof. It follows from the Markov property that the i, jth entry of A^n counts the number of components in the intersection $\varphi^n(P_i) \cap P_j$. Suppose by contradiction that A is not irreducible, so that there is a sequence $n_k \to \infty$ and $1 \le i, j \le p$ so that

$$\varphi^{n_k}(P_i) \cap P_j = \varnothing. \tag{10}$$

Let σ be a horizontal segment in P_i and consider its image σ_k under φ^{n_k}. Then σ_k is a horizontal segment contained in $\varphi^{n_k}(P_i)$ whose length tends to infinity. Passing to a subsequence we can assume the left endpoints of σ_k converge to $x \in M$, so that any point along an infinite horizontal ray ℓ issuing from x is a limit of points of σ_k. Since φ^{-1} is an affine automorphism on M which contracts horizontal saddle segments, there are no horizontal saddle connections on M, so the horizontal line flow on M is minimal. This implies that ℓ is dense in M. In particular for large enough k, $\sigma_k \cap \operatorname{int} P_j \neq \varnothing$, contradicting (10).

Now let w_i be the width of P_i, so that λw_i is the width of $\varphi(P_i)$. By the Markov property $\varphi(P_i)$ passes a_{ij} times through P_j, so that $\lambda w_i = \sum_j a_{ij} w_j$. That is,
$$\lambda \vec{w} = A \vec{w}.$$
Since $\vec{w}$ is a positive vector, $\lambda = \lambda(A)$. The proof for heights is almost identical. $\qquad\square$

For a non-negative irreducible $A \in \operatorname{Mat}_p(\mathbb{Z})$, let $\mathcal{G}(A)$ denote the gluing patterns (as in §2.2.2) of Markov partitions arising from hyperbolic affine automorphisms whose intersection matrix is A. We have:

Proposition 5.2. $\mathcal{G}(A)$ *is finite.*

Proof. Since φ (resp. φ^{-1}) preserves the boundaries of rectangles, but contracts vertical (horizontal) saddle connections, there are there are no vertical or horizontal saddle connections for M. Thus each rectangle contributes at most 4π to the total angle around the singularities of M. In other words the number of rectangles bounds the number of singularities. Since each edge in a gluing pattern is either bounded by a singularity or by a full size of a rectangle, this also bounds the number of edges in a gluing pattern. Thus the number of gluing patterns is finite. $\qquad\square$

We now want to show that intersection matrix of φ determines the metric data for the gluing pattern. Let ξ_i be the horizontal segments of the gluing pattern corresponding to $P_1, \ldots, P_p$. We will refine our Markov partition to obtain another partition with the Markov property. Each ξ_i is a connected component of the intersection of two of the P_j, say P_1, P_2. Consider the vertical segments (if any) which issue from the endpoints of ξ_i into the interior of P_1, P_2. Any such segment divides a rectangle vertically into two rectangles, with the same height as before. At least one of the new rectangles will have width ξ_i. We denote the resulting rectangle decomposition by $M = Q_1 \cup \cdots \cup Q_r$, and claim that this decomposition also has the Markov property. To see this, note first that subdividing rectangles by vertical lines does not affect the

requirement for φ^{-1}, which involves only horizontal sides. Now suppose $Q \subset P_1, P_2$ are as in Figure 3, $Q' \subset P'$ such that $\operatorname{int} \varphi(Q) \cap \operatorname{int} Q' \neq \varnothing$. There are two cases to consider:

(1) If $\varphi(P_1)$ does not reach the top of P' then also $\operatorname{int} \varphi(P_2) \cap \operatorname{int} P' \neq \varnothing$ and by the Markov property for the P_i, both $\varphi(P_1)$ and $\varphi(P_2)$ go across P' so $\varphi(Q)$ goes across P' and hence across Q'.
(2) If $\varphi(P_1)$ reaches the top of P' and P'' lies above P' then $\varphi(P_1)$ goes across P' and $\varphi(P_2)$ goes across P'' so that $\varphi(Q)$ goes across the smaller of the two, i.e., $\varphi(Q)$ goes across Q'.
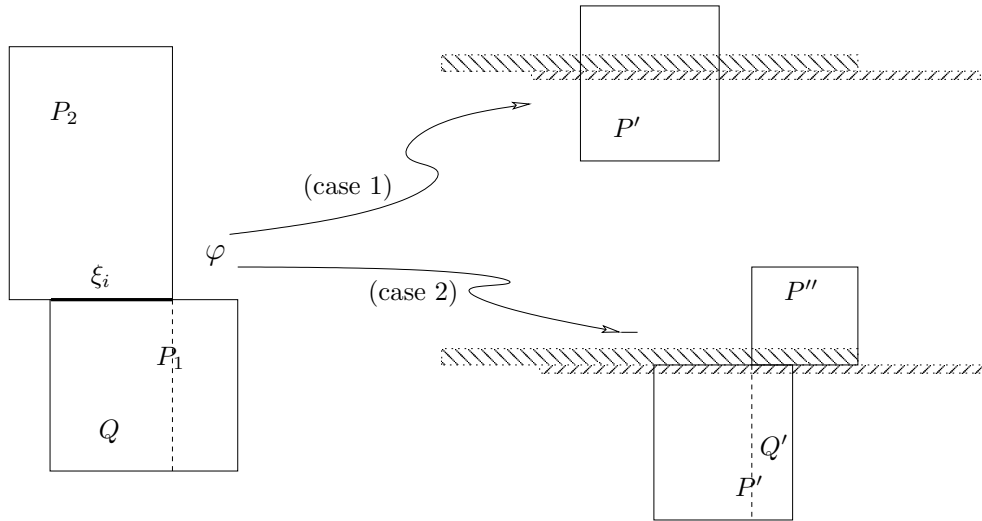

FIGURE 3. Markov property for the Q_j

Considering all cases in this manner proves the claim.

We can also refine using the η_i instead of the ξ_i, cutting rectangles horizontally, to obtain another Markov partition $T_1, \ldots, T_t$. We define matrices B and C by letting b_{ij} (resp. c_{ij}) be the number of connected components of $\operatorname{int} \varphi(Q_i) \cap \operatorname{int} Q_j$ (resp. $\operatorname{int} \varphi(T_i) \cap \operatorname{int} T_j$). Repeating the proof of Proposition 5.1 we obtain:

Proposition 5.3.   • *B and C are irreducible.*
   • $\lambda = \lambda(B) = \lambda(C)$.
   • *The vector recording the lengths of the segments ξ_i (resp. η_i) is a positive eigenvector for B (resp. C).*

We are now in a position to formulate and prove a more precise version of Proposition 1.5. We will consider separately translation and

half-translation surfaces, so we write

$$\mathrm{SMP}_1(p, T) = \{(M, h) \in \mathrm{SMP}(p, T) : M \text{ is a translation surface}\},$$

and denote the corresponding set of G-orbits by $\widetilde{\mathrm{SMP}}_1(p, T)$. Let $\mathcal{M}(p, T)$ be the set of positive integer matrices A with $\lambda(A) < T$. This is a finite set by Proposition 2.4. Then we have:

Theorem 5.4. *For a fixed $T > 0$ and $p \in \mathbb{N}$,*

$$\# \widetilde{\mathrm{SMP}}_1(p, T) \leq \sum_{A \in \mathcal{M}(p,T)} \# \mathcal{G}(A).$$

Proof. Given $(M, h) \in \mathrm{SMP}_1(p, T)$, we have constructed $A \in \mathcal{M}(p, T)$ and a gluing pattern in $\mathcal{G}(A)$, and these data only depend on the affine equivalence class of M. Thus we have defined a map

$$\Phi : \widetilde{\mathrm{SMP}}_1(p, T) \to \cup_{A \in \mathcal{M}(p,T)} \mathcal{G}(A),$$

and it remains to show that Φ is injective, i.e. that M and h are uniquely determined by the matrix A and the gluing pattern, up to an element of G. Applying an element $g \in G$ we can assume that the parallelograms in the Markov partition corresponding to (M, h) are rectangles with the expanding direction horizontal, and moreover the width of the widest rectangle can be normalized to be 1. With this choice of g we need to show that the matrix and gluing pattern uniquely determine gM. Since the widths of the rectangles are in the unique positive eigendirection of A, and by our scaling convention, the matrix A determines the widths of the rectangles. The heights also span the unique eigendirection for A^{tr} so are determined by A up to scaling. By the requirement that gM has area 1, the heights are uniquely determined by A. From A and the gluing pattern one determines the matrices B and C, and the lengths of the ξ_i and η_j give a positive eigenvector of B and C which is thus determined up to scaling. Since the sidelengths of the rectangles P_i are determined, and can also be calculated using the lengths of the ξ_i and η_i, the ξ_i and η_j are also uniquely determined. Since we have specified its gluing pattern and geometry, the flat surface gM is uniquely determined. $\square$

Proof of Proposition 1.5. By Proposition 2.2, for each $(M, h) \in \mathrm{SMP}(p, T)$ the orientation double cover $\widetilde{M}$ has a hyperbolic affine automorphism φ such that the corresponding number of parallelograms is at most $2p$. I.e., $(M, h) \in \mathrm{SMP}(p, T)$ gives rise to $(\widetilde{M}, h) \in \mathrm{SMP}_1(2p, T)$. Also for a fixed $(\widetilde{M}, h)$ the number of (M, h) covered in this way is finite. The Proposition follows. $\square$

6. Restrictions on Veech groups

Proof of 1.6. It is easily checked that (I) and (II) depend only on the commensurability class of Γ, so to prove them we can assume that $\Gamma = \Gamma_M$ for a flat surface M. Since the maximal number of cylinders for a cylinder decomposition on M is bounded, assertion (I) follows immediately from Theorem 1.3. Similarly assertion (II) is immediate from Proposition 1.5.

To prove assertion (III) we note that in a Veech group there is a uniform upper bound on the order of an elliptic element by Hurwitz's theorem, hence such a bound also exists in any group commensurable with a Veech group. So it remains to show that a Fuchsian group Γ in which (II) holds and orders of elliptic elements are uniformly bounded, (III) holds. To see this, given T suppose we have a list $z_1, z_2, \ldots$ of fixed points for elliptics $f_1, f_2, \ldots$ in Γ, such that the corresponding cone areas are no more than T. Since the order of the f_i is bounded above, we have a uniform upper bound for $R(f_i)$ defined via (1). Let h_i be the hyperbolic element corresponding to f_i via Proposition 3.3. Since the eigenvalue of the h_i is bounded, by (II) the list $h_1, \ldots$, contains only finitely many conjugacy classes. Conjugating the f_i's and correspondingly the h_i's we may assume that there are only finitely many distinct elements in the list $h_1, h_2, \ldots$. Given $h = h_i$ in this list, let A be its axis, and let $A_0 \subset A$ be a compact fundamental domain for the action of h on A. By a further conjugation of f_i assume that the closest point to z_i on A is in A_0. In view of the upper bound on the order of the f_i's, by Proposition 3.3 z_i is within bounded distance of A_0. By discreteness of Γ there are only finitely many z_i's and hence finitely many f_i's. □

Remark 6.1. *The argument proving (III) can be adapted to prove the following analogue of Theorem 1.3 and Proposition 1.5: for any stratum $\mathcal{H}$ and any T, the set of affine equivalence classes of pairs (M, f), where $M \in \mathcal{H}$ and $f \in \Gamma_M$ is elliptic, with corresponding cone area at most T, is finite.*

For the proof of Theorem 1.2 we will need the following Lemma.

Lemma 6.2. *Suppose Λ_0 is a Fuchsian group, Γ a normal subgroup of Λ_0, $\gamma \in \Gamma$ a non-central element such that for any $\lambda \in \Lambda_0$ there is $\tau \in \Gamma$ with $\gamma^\lambda = \gamma^\tau$. Then Γ is of finite index in Λ_0.*

Proof of Theorem 1.2 (assuming Lemma 6.2). Let Γ be a non-elementary Fuchsian group commensurable to a Veech group, let Λ be the normalizer of Γ in G, and suppose Λ/Γ is infinite. Since Γ is non-elementary

it contains a hyperbolic element γ [Ka, Thm. 2.4.4], and Λ is Fuchsian [Ka, Thm. 2.3.8]. Let

$$\Lambda_0 = \{\lambda \in \Lambda : \exists \tau \in \Gamma \text{ s.t. } \gamma^\lambda = \gamma^\tau\}.$$

In the action of Λ on conjugacy classes of Γ, Λ_0 is the stabilizer of the conjugacy class of γ. In particular Λ_0 is a subgroup of Λ. According to Lemma 6.2, Λ_0/Γ is finite so Λ/Λ_0 is infinite. This implies that the conjugates $\{\gamma^\lambda : \lambda \in \Lambda\}$ comprise infinitely many conjugacy classes, in other words Γ contains infinitely many hyperbolic elements which are conjugate in Λ, hence have the same eigenvalue, but are not conjugate in Γ. This contradicts Corollary 1.6. $\qquad\square$

Proof of Lemma 6.2. Let C be the centralizer of γ in Λ_0 and let $C_0 = C \cap \Gamma$. Since Λ_0 is Fuchsian, C is cyclic (see [Ka, §2.3]) and therefore C_0 is of finite index in C. By assumption for every $\lambda \in \Lambda_0$ there is $\tau \in \Gamma$ such that $\lambda^{-1}\tau \in C$, which implies that $\Lambda_0 = C\Gamma$. Thus C maps onto Λ_0/Γ. Since this surjection factors through C/C_0, Λ_0/Γ must be finite. $\qquad\square$

Proof of Corollary 1.7. Given $\mathcal{H}$, for any $M \in \mathcal{H}$ and any hyperbolic $h \in \Gamma_M$, the minimal number $p(M, h)$ of parallelograms in a corresponding Markov partition can assume only finitely many values by Proposition A.1. Let Γ contain a hyperbolic element h. By Proposition 1.5 the number of G-orbits of M such that $h \in \Gamma_M$ is finite. In particular the set in (2) intersects only finitely many G-orbits. If M_1 and M_2 are flat surfaces in $\mathcal{H}$ such that $gM_1 = M_2$ and $\Gamma_{M_1} = \Gamma_{M_2}$ then $g \in N$; and if Γ is non-elementary then N/Γ is finite by Corollary 1.2. $\qquad\square$

Proof of Theorem 1.1. Given R let $T = 2R$. Enlarging T we ensure that a cusp of area at least T contains an embedded ball of radius R. Enlarging T further, in light of the upper bound on the order of elliptic elements in Γ, we ensure that any cone of area at least T contains an embedded ball of radius R. Thus we may assume $\mathbb{H}/\Gamma$ contain neither a cusp of volume at least T, nor a cone of area at least T; in view of Corollary 1.6, $\mathbb{H}/\Gamma$ has only finitely many cusps and cones, and only finitely many closed geodesics of length less than T.

Now let $\mathcal{N}_1 \subset \mathbb{H}/\Gamma$ be the union of closed geodesics of length less than T, cusps of area less than T, and cones of area less than T. Then $\mathcal{N}_1$ has finite area. Let $\mathcal{N}$ be a neighborhood of $\mathcal{N}_1$ which is large enough so that a closed loop intersecting the complement of $\mathcal{N}$ and homotopic to either a geodesic of length less than T, a loop around a cusp of area less than T, or a loop around a cone point of area less

than T, must have length at least $2R$. Since $\mathbb{H}/\Gamma$ has only finitely many cusps and cones, the area of $\mathcal{N}$ is also finite, so $\mathbb{H}/\Gamma \neq \mathcal{N}$.

Let $\pi : \mathbb{H} \to \mathbb{H}/\Gamma$ be the natural map, and let $\bar{z} \in \mathbb{H}$ so that $z = \pi(\bar{z}) \notin \mathcal{N}$. We claim that $\pi|_{B(\bar{z},R)}$ is injective. Otherwise there is a segment $\bar{\sigma}$ connecting two distinct points in $B(\bar{z}, R)$, mapping to a closed loop σ in $\mathbb{H}/\Gamma$ of length less than $2R$. Either σ has a shortest representative which is a geodesic of length less than $2R$, or σ is freely homotopic to a curve around a cusp or cone, that is a curve in $\mathcal{N}_1$. But then the definition of $\mathcal{N}$ ensures that the length of σ is greater than T, and this is a contradiction. $\qquad\square$

Appendix A. The construction of Markov partitions

Proposition A.1. *For any stratum $\mathcal{H}$ of flat surfaces, there are R_1 and R_2 such that for any surface $M \in \mathcal{H}$ and any hyperbolic $\varphi \in \mathrm{Aff}(M)$:*

- *any Markov partition for φ has at least R_1 parallelograms;*
- *there is a Markov partition for φ with at most R_2 parallelograms.*

Proof. Let $k = \sum_{\sigma \in \Sigma} k_\sigma$ so that $k\pi$ is the total angle around all the singularties of a surface in $\mathcal{H}$. Suppose $M \in \mathcal{H}$ and $\varphi \in \mathrm{Aff}(M)$ is hyperbolic, with a Markov partition into parallelograms $P_1, \ldots, P_p$. By conjugating we can assume the P_i are rectangles with horizontal and vertical sides. The horizontal and vertical directions are contracted by either φ or φ^{-1}, but these maps preserve the set of saddle connections, so there are no horizontal or vertical saddle connections on M. In particular each edge of each P_i can only contain one singularity, so that the total angle around singularities coming from each P_i is at most 4π. This implies that $p \geq R_1 = k/4$, proving the first assertion.

To prove the second assertion, by Proposition 2.2 it suffices to consider strata of translation surfaces. Also, if we have a Markov partition for φ^r, the common refinement of its images under φ^i, $i = 0, \ldots, r$, provides a Markov partition for φ. Thus by passing to a suitable finite power of φ, which can be taken to depend only on $\mathcal{H}$, we can assume that φ fixes all the singularities and all the critical leaves issuing from singularities. Assume without loss of generality that the contracting direction for φ (resp. φ^{-1}) is vertical (resp. horizontal).

Since there are no horizontal or vertical saddle connections, the foliation in both the horizontal and vertical direction is minimal, so the leaves intersect any transverse segment. We will construct the edges of the partition in 3 steps.

(1) Let γ be a short closed segment starting at a singularity and going along a horizontal leaf, say from left to right, and let λ be a union of closed segments starting at each singularity along the

vertical leaves (going both up and down) until the first point of intersection with γ. Denote the segments comprising λ by $\lambda_1, \ldots, \lambda_k$.

(2) The intersection $\gamma \cap \lambda$ is finite, so there is a *terminal point p* which is the right-most intersection point of $\lambda \cap \gamma$. Remove from γ the subsegment to the right of p. We retain the name γ for the shorter segment.

(3) Suppose λ_j ends at p. Continue it further until its next intersection with γ. We retain the names λ, λ_j for this new collection. Since M has no vertical saddle connections, the new intersection point of λ_j with γ is not one of the previous ones, and no two of these new intersection points coincide.

We first claim that each connected component of $M \smallsetminus (\gamma \cup \lambda)$ is a rectangle. Take a segment $I \subset \gamma$ which is bounded by either a singularity and the first point where a λ_ℓ comes down to γ, or by two consecutive points at which λ_ℓ's come down to γ. Consider the union of vertical leaves which begin at points of I and move upwards. These form a strip, which near γ is bounded on both sides by segments in λ. By choice of the λ_j all these leaves all hit γ again 'at the same time', i.e. none of these leaves hit a singularity or a terminal point before returning to γ. The construction also ensures that the segments from λ on both sides of the strips also extend until the next intersection with γ. Thus the strip is a rectangle bounded by $\gamma \cup \lambda$.

We now claim that this partition into rectangles has the Markov property. Since φ^{-1} fixes each singularity and maps each critical leaf to itself, and since it contracts the horizontal direction, γ is mapped into itself. This implies that lower boundaries of rectangles map into lower boundaries under φ^{-1}. For identical reasons each λ_j is mapped into itself by φ.

Now note that the endpoints of lower edges of rectangles on γ are the terminal point p, the singularity, and one additional point for each downward pointing edge in λ. In total we have $2 + k/2$ endpoints which gives $1 + k/2$ rectangles. In particular the number of rectangles depends only on $\mathcal{H}$. $\qquad \square$

References

[DoHu] A. Douady and J. Hubbard, *On the density of Strebel differentials*, Inv. Math. **30** (1975) 175–179.

[EsOk] A. Eskin and A. Okounkov, *Asymptotics of numbers of branched coverings of a torus and volumes of moduli spaces of holomorphic differentials*, Inv. Math. **145** (2001), no. 1, 59–103.

[FLP] A. Fathi, F. Laudenbach, and V. Poenaru, **Travaux de Thurston sur les surfaces**, vol. 66-67, Asterisque, 1979.

[Ga] F. R. Gantmacher, **Theory of matrices, Vol. 2** Chelsea publishing, 1959.

[HuSc1] P. Hubert and T. Schmidt, *Invariants of translation surfaces*, Ann. Inst. Fourier **51**, 2 (2001) 461–495.

[HuSc2] P. Hubert and T. Schmidt, *Infinitely generated Veech groups*, Duke Math. J. **123** (2004), no. 1, 49–69.

[HuSc3] P. Hubert and T. Schmidt, *An introduction to Veech surfaces,* in **Handbook of Dynamical Systems, Vol. 1B**, B. Hasselblatt and A. Katok (eds.) 501–526 Elsevier (2006).

[Ka] S. Katok, **Fuchsian groups**, Univ. of Chicago Press (1992).

[KeMaSm] S. Kerckhoff, H. Masur and J. Smillie, *Ergodicity of billiard flows and quadratic differentials*, Ann. Math. **124** (1986) 293–311.

[MaSm] H. Masur and J. Smillie, *Hausdorff dimension of sets of nonergodic measured foliations*, Ann. Math. **134** (1991) 455-543.

[MaTa] H. Masur and S. Tabachnikov, *Rational billiards and flat structures*, in **Handbook of dynamical systems, Enc. Math. Sci. Ser.** (2001).

[Mc] C. McMullen, *Teichmüller geodesics of infinite complexity*, Acta Math. **191** (2003) 191–223.

[Na] Y. Naveh, *Tight upper bounds on the number of invariant components on translation surfaces*, preprint (2006).

[PaPe] A. Papadopoulos and R. C. Penner, *Enumerating pseudo-Anosov foliations*, Pac. J. Math **142** (1990) 159–173.

[Pe] R. C. Penner, *Bounds on least dilatations*, Proc. Amer. Math. Soc. **113** (1991) 443–450.

[Ry] E. Rykken, *Expanding factors for pseudo-Anosov homeomorphisms*, Michigan Math. J. **46** (1999), no. 2, 281–296.

[SmWe1] J. Smillie and B. Weiss, *Finiteness results for flat surfaces: a survey and problem list*, (2006) to appear in **Partially hyperbolic dynamics, laminations, and Teichmüller flow** (Proceedings of a conference, Fields Institute, Toronto Jan 2006), G. Forni, M. Lyubich, and M. Shub (eds.)

[SmWe2] J. Smillie and B. Weiss, *Characterizations of lattice surfaces*, preprint (2008).

[Th] W. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. AMS (new series) **19** no. 2 (1988) 417–431.

[Ve1] W. A. Veech, *The Teichmüller geodesic flow*, Ann. of Math. **124** (1986) 441–530.

[Ve2] W. A. Veech, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Invent. Math. **97** (1989), no. 3, 553–583.

[Vo] Ya. B. Vorobets, *Planar structures and billiards in rational polygons: the Veech alternative*, (Russian) Uspekhi Mat. Nauk **51** (1996), no. 5(311), 3–42; translation in Russian Math. Surveys **51** (1996), no. 5, 779–817.

[Zo] A. Zorich, *Flat surfaces*, in **Frontiers in number theory, physics and geometry,** P. Cartier, B. Julia, P. Moussa and P. Vanhove (eds), Springer (2006).

Cornell University, Ithaca, NY smillie@math.cornell.edu

Ben Gurion University, Be'er Sheva, Israel 84105 barakw@math.bgu.ac.il