ORIGINAL PAPER

Constructing lattice surfaces with prescribed Veech groups: an algorithm

Slade Sanderson[1]

Abstract
The Veech group of a translation surface is the group of Jacobians of orientation-preserving affine automorphisms of the surface. We present an algorithm which constructs all translation surfaces with a given lattice Veech group in any given stratum. In developing this algorithm, we give a new proof of a finiteness result of Smillie and Weiss, namely that there are only finitely many unit-area translation surfaces in any stratum with the same lattice Veech group. Our methods can be applied to obtain obstructions of lattices being realized as Veech groups in certain strata; in particular, we show that the square torus is the only translation surface in any minimal stratum whose Veech group is all of $SL_2\mathbb{Z}$.

1 Introduction

The study of translation surfaces may be approached from numerous different angles— complex analysis, differential geometry, algebraic topology and number theory, for instance, each offer unique insights into these objects. These different perspectives lead to a wealth of questions regarding translation surfaces and their dynamics. One such avenue of inquiry is the well-known action of $SL_2\mathbb{R}$ on the stratum $\mathcal{H}_1(d_1, \ldots, d_\kappa)$ of unit-area translation surfaces with singularities of orders $d_1, \ldots, d_\kappa$. Stabilizers (i.e. Veech groups) and orbits of translation surfaces under this action have proven to be of great interest. Analogues of Ratner's theorems concerning unipotent flows on homogeneous spaces [21] are proven in the celebrated work of Eskin, Mirzakhani and Mohammadi; in particular, $SL_2\mathbb{R}$-orbit closures are affine invariant submanifolds of $\mathcal{H}_1(d_1, \ldots, d_\kappa)$ [8]. Despite having been well-studied since as early as the 1980s, some immediate questions regarding Veech groups have been non-trivial to address. General methods of computing Veech groups were unknown until recently ([3, 4, 6, 7, 20, 29], and the special cases of [9, 23]), and while some universal properties of Veech groups are known—for instance, they are necessarily discrete and non-cocompact [27]—a complete

---

✉ Slade Sanderson
s.b.sanderson@uu.nl

[1] Department of Mathematics, Utrecht University, P.O. Box 80010, 3508TA Utrecht, The Netherlands

classification of which subgroups of $SL_2\mathbb{R}$ are realized as Veech groups remains an open problem [13].

The first non-arithmetic examples were discovered by Veech, who proved that the odd index Hecke triangle groups and an index two subgroup of the even index Hecke triangle groups are realized as Veech groups [27]. More general triangle groups are studied in [2, 11], where the authors show that there are Veech groups commensurable to all triangle groups $\Delta(m, n, \infty)$, $2 \leq m < n < \infty$. On the other hand, triangle groups whose orientation preserving subgroups are never contained in a Veech group are also given in [11]. Infinitely generated Veech groups are shown to exist in [14] and [18], as are non-trivial Veech groups with no parabolic elements in [12]. While this brief survey is far from an exhaustive overview of the literature, it does serve to highlight that the question of realizability of Veech groups has spanned decades and remains a large and interesting problem in the field.

Of particular interest within the space of all translation surfaces are so-called *lattice surfaces*, or *Veech surfaces*, which are those whose Veech groups are lattices (have finite covolume) in $SL_2\mathbb{R}$. These surfaces admit especially nice dynamics [27], but the list of known families of lattice surfaces is relatively short [5]. Lattice surfaces are also interesting as they are precisely those surfaces whose $SL_2\mathbb{R}$-orbits are closed with respect to the analytic topology on the stratum [24, 28], and the projection of a closed orbit to the moduli space of Riemann surfaces is an algebraic curve—called a *Teichmüller curve*—isometrically immersed with respect to the Teichmüller metric (see, say, [31]).

We present an explicit algorithm which constructs all translation surfaces with a given lattice Veech group in any given stratum:

Algorithm 1.1 *Input: Non-negative integers $d_1 \leq \cdots \leq d_\kappa$ whose sum is even and a finite set of generators of a lattice $\Gamma \leq SL_2\mathbb{R}$.*

Output: All translation surfaces $(X, \omega) \in \mathcal{H}_1(d_1, \ldots, d_\kappa)$ with Veech group $SL(X, \omega) = \Gamma$.

An implicit algorithm is given in [25] which enumerates all affine equivalence classes of lattice surfaces in terms of a parameter measuring the infimum of areas of triangles within the surface. Algorithm 1.1 similarly returns lattice surfaces, but it fundamentally differs both in its goal of finding surfaces with specific, prescribed Veech groups and in its method of doing so.

In developing Algorithm 1.1, we obtain a new proof of a finiteness result of Smillie and Weiss (Corollary 1.7, [26]; see also [19]):

Theorem 1.2 (Smillie–Weiss). *There are at most finitely many unit-area lattice surfaces with a given Veech group in any given stratum.*

Our methods essentially reverse the algorithm for computing Veech groups presented in [6] and [7]. There the authors associate to each stratum a canonical (infinite area, disconnected in general) flat surface $\mathcal{O}$, within which any closed, connected translation surface (X, ω) of the stratum may be naturally represented. Data regarding the translation surface and its saddle connections are recorded via pairs of points in $\mathcal{O}$ termed *orientation-paired marked segments*, and the authors show that (X, ω) may be recovered from a certain finite subset of these marked segments—the *marked Voronoi staples* (§3.4 of [7]). Furthermore, a classification of the Veech group of (X, ω) in terms of its orientation-paired marked segments and affine automorphisms of $\mathcal{O}$ is given in Proposition 17 of [7].

Rather than beginning with a lattice surface (X, ω) and using its corresponding orientation-paired marked segments to compute elements of the Veech group $\Gamma = SL(X, \omega)$,

Algorithm 1.1 begins with a lattice $\Gamma \leq \mathrm{SL}_2\mathbb{R}$, simulates normalized subsets of orientation-paired marked segments in $\mathcal{O}$ using Γ, and constructs candidate translation surfaces with Veech group Γ from finite unions of appropriately scaled versions of these 'simulations.' Theorem 1.2 follows from the facts that

  (i) a lattice Γ determines only finitely many such simulations (Lemma 4.8),
 (ii) the orientation-paired marked segments of any unit-area translation surface with Veech group Γ are realized as a finite union of scaled simulations (Theorem 4.9), and
(iii) for any finite collection of simulations there are only finitely many scalars for which the union of scaled simulations could possibly coincide with the orientation-paired marked segments of a unit-area translation surface.

For this latter point, we introduce the technical but important notions of *permissible triples* and *permissible scalars* of particular subsets of $\mathcal{O}$, prove finiteness results regarding them (Proposition 3.5 and Lemma 3.7), and apply these results to orientation-paired marked segments and marked Voronoi staples (Proposition 4.11). Our novel proof differs from that of Smillie and Weiss, who use finiteness results of Markov partitions corresponding to hyperbolic affine automorphisms of surfaces [26].

This work should be viewed in light of the aforementioned open question regarding the realization of subgroups of $\mathrm{SL}_2\mathbb{R}$ as Veech groups: a full implementation[1] of Algorithm 1.1 could lead to the discovery of new lattice surfaces and experimental conjectures as to which groups are realized as Veech groups. Moreover, further work to determine a halting criterion for Algorithm 1.1 would give a general, explicit procedure to determine whether or not a given lattice is a Veech group in any particular stratum (see comments at the end of Sect. 5). At present, these ideas may be used in special cases to obtain obstructions for lattices being realized as Veech groups. In particular, we prove:

Theorem 1.3 *The square torus is the only translation surface in $\cup_{g>0}\mathcal{H}(2g-2)$ with Veech group* $\mathrm{SL}(X,\omega) = \mathrm{SL}_2\mathbb{Z}$.

The paper is organized as follows. Section 2 gives a brief background on translation surfaces (Sect. 2.1) and their Veech groups (Sect. 2.2). Section 3 concerns the canonical surface $\mathcal{O}$ studied in [6] and [7]: definitions and immediate results regarding $\mathcal{O}$ are given in Sect. 3.1, permissible triples and scalars are introduced and studied in Sect. 3.2, and in Sect. 3.3 we recall definitions and results regarding marked segments and Voronoi staples of a translation surface. Section 4 further develops the machinery for Algorithm 1.1 and uses these ideas to give a new proof of Theorem 1.2: a characterization of lattice groups in terms of directions announced by linear transformations is given in Sect. 4.1; Sects. 4.2 and 4.3 explore a group action on marked segments and how a lattice group may be used to simulate orbits of this action; Sect. 4.4 relates marked Voronoi staples to permissible triples; and a new proof of Theorem 1.2 is given in Sect. 4.5. Algorithm 1.1 is presented in Sect. 5 and Theorem 1.3 is proven in Sect. 6.

---

[1] A partial implementation has been written by the author using [22] to confirm some low genus examples with small generating sets, though an efficient version of the algorithm which handles more interesting examples awaits implementation.

2 Background

2.1 Translation surfaces

We begin with basic definitions and notation; see, say, surveys of [31] and [32] for more details.

2.1.1 Translation surfaces: various perspectives

A *translation surface* (X, ω) is a Riemann surface X together with a non-zero, holomorphic one-form ω. Equivalently, a translation surface is defined as a real surface X such that off of a finite set Σ, the surface $X \backslash \Sigma$ is equipped with a translation atlas, and the resulting flat structure extends to all of X to give conical singularities at points of Σ whose angles are positive integral multiples of 2π. Any translation surface may be *polygonally represented* as a collection of polygons in the plane together with an identification of edges in pairs, such that identified edges are parallel, equal length and of opposite orientation. A translation surface is *closed* if it is compact and has no boundary.

In these three perspectives—which we move fluidly between—the zeros of ω, elements of Σ of cone angles $2\pi(d+1)$, $d > 0$, and, after identifying edges, the vertices of angles $2\pi(d+1)$, $d > 0$, of a polygonally represented translation surface each correspond to the same set of *non-removable singularities* of the surface. Throughout, the *singular set* Σ denotes these singularities together with possibly finitely many *removable singularities* (or *marked points*), i.e. points of cone angle 2π. We refer to any element of Σ as a singularity (of cone angle $2\pi(d+1)$, $d \geq 0$), bearing in mind that it may in fact be removable. Points in the complement $X \backslash \Sigma$ are called *regular*.

2.1.2 Metric, Lebesgue measure, saddle connections and holonomy vectors

A translation surface (X, ω) comes equipped with a metric defined as follows: given $p, q \in X$, the distance from p to q is given by

$$d(p, q) := \inf_{\gamma} \left| \int_{\gamma} \omega \right|,$$

where the infimum is over all piecewise-smooth curves $\gamma \subset X$ from p to q. The infimum is realized by a straight line segment (with respect to the flat structure) on (X, ω) or by a union of such segments meeting at singularities. Lebesgue measure in the plane pulls back via coordinate maps of the translation atlas of (X, ω) to give a measure on $X \backslash \Sigma$; we extend this measure to all of X by declaring Σ to be a null set.

A *separatrix* is a straight line segment on (X, ω) emanating from a singularity and having no singularities in its interior. A *saddle connection* is a separatrix (of positive length) which also ends at a singularity. The *holonomy vector* of a saddle connection s is defined as $\mathrm{hol}(s) := \int_s \omega$. The collection of all holonomy vectors of (X, ω) forms a subset of the plane with no limit points and whose directions are dense in S^1 (Proposition 3.1, [30]).

2.1.3 Voronoi decomposition

Every translation surface (X, ω) has a (unique) *Voronoi decomposition* subordinate to its singular set Σ, which is described as follows (see also [3, 17]). For any $x \in X$, let $d(x, \Sigma)$

denote the minimum of the set $\{d(x, \sigma)\}_{\sigma \in \Sigma}$. To each $\sigma \in \Sigma$ there is associated an open, connected 2-cell, $\mathcal{C}_\sigma$, comprised of the points $x \in X$ for which $d(x, \Sigma)$ is realized by a unique length-minimizing path ending at σ. The boundary of $\mathcal{C}_\sigma$ is the union of 1-cells of points x for which $d(x, \Sigma)$ is realized by precisely two length-minimizing paths, at least one of which ends at σ, and of 0-cells of points x for which $d(x, \Sigma)$ is realized by three or more length-minimizing paths, at least one of which ends at σ.

2.2 Veech groups

2.2.1 Affine diffeomorphisms and Veech groups

An *affine diffeomorphism* $f : (X_1, \omega_1) \to (X_2, \omega_2)$ between (not necessarily closed nor connected) translation surfaces is a diffeomorphism from X_1 to X_2 sending the singular set Σ_1 of X_1 into that of X_2, and which—on the complement of Σ_1—is locally an affine map of the plane of constant linear part[2] (i.e. a map of the form $v \mapsto Av + b$ for some global $A \in \mathrm{GL}_2\mathbb{R}$ and local $b \in \mathbb{R}^2$). An *affine automorphism* is an affine diffeomorphism from a translation surface (X, ω) to itself. The set of all affine automorphisms of (X, ω), denoted $\mathrm{Aff}(X, \omega)$, forms a group under composition, as does the subset $\mathrm{Aff}^+(X, \omega) \subset \mathrm{Aff}(X, \omega)$ of orientation-preserving elements. The map der : $\mathrm{Aff}(X, \omega) \to \mathrm{GL}_2\mathbb{R}$ sending an orientation-preserving affine automorphism to its linear part in local coordinates gives a group homomorphism. The image of $\mathrm{Aff}^+(X, \omega)$ under this map, denoted $\mathrm{SL}(X, \omega)$, is called the *Veech group* of (X, ω), and its kernel is the group of translations denoted by $\mathrm{Trans}(X, \omega)$. Veech shows in [27] that for closed, connected (X, ω), the Veech group $\mathrm{SL}(X, \omega)$ is a discrete subgroup of $\mathrm{SL}_2\mathbb{R}$—and thus its image in $\mathrm{PSL}_2\mathbb{R}$ is a Fuchsian group—and, moreover, $\mathrm{SL}(X, \omega)$ is always non-cocompact.

2.2.2 Strata, a $\mathrm{GL}_2\mathbb{R}$-action and lattices

We call (X_1, ω_1) and (X_2, ω_2) *translation equivalent* if there is an affine diffeomorphism between them with trivial linear part. For non-negative integers $d_1 \leq \cdots \leq d_\kappa$, let $\mathcal{H}(d_1, \ldots, d_\kappa)$ denote the *stratum* of all closed, connected translation surfaces—up to translation equivalence—with singularities of cone angles $2\pi(d_1 + 1) \leq \cdots \leq 2\pi(d_\kappa + 1)$ and no other singularities. (Recall that we allow for 'removable singularities,' so some d_i may be zero, and with our definitions, $\mathrm{SL}(X, \omega)$ depends on these marked points in Σ. In particular, Algorithm 1.1 applies not only to standard strata where each $d_i > 0$, but it also allows for the construction of lattice surfaces (X, ω) with a specified number of marked points which remain invariant under $\mathrm{Aff}^+(X, \omega)$.) Abusing notation, we denote an element of a stratum by any of its translation-equivalent representatives. By the Riemann-Roch Theorem, every translation surface $(X, \omega) \in \mathcal{H}(d_1, \ldots, d_\kappa)$ has the same genus $g > 0$, and the sum of the d_i equals $2g - 2$.

There is a natural action of $\mathrm{GL}_2\mathbb{R}$ on each stratum, given by post-composing the coordinate charts of $(X, \omega) \in \mathcal{H}(d_1, \ldots, d_\kappa)$ with the usual action of a matrix in the plane via a linear transformation (one can verify that this action is well-defined with respect to translation-equivalence). Any translation surface may be normalized by the action of a diagonal matrix so as to have unit-area; we denote the collection of unit-area translation surfaces in $\mathcal{H}(d_1, \ldots, d_\kappa)$ by $\mathcal{H}_1(d_1, \ldots, d_\kappa)$. There is a corresponding action of $\mathrm{SL}_2\mathbb{R}$ on

---

[2] The translation structures guarantee that the linear part of a locally affine map is constant on components of (X_1, ω_1); this latter statement in the definition is thus intended for disconnected translation surfaces.
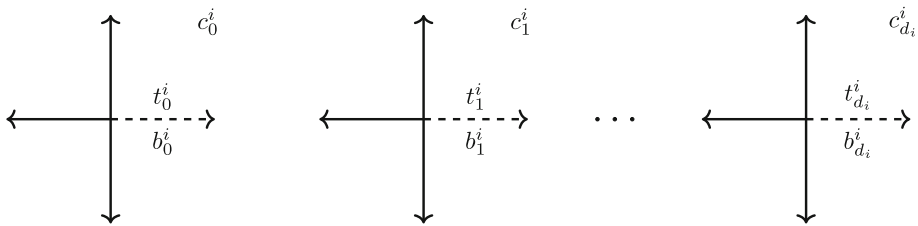
Fig. 1 The surface $\mathcal{O}_i$, where the bottom of the slit, b_j^i, is identified with the top of the slit, t_{j+1}^i, for each $j \in \mathbb{Z}_{d_i+1}$

$\mathcal{H}_1(d_1, \ldots, d_\kappa)$, and the stabilizer of (X, ω) under this $\mathrm{SL}_2\mathbb{R}$-action is isomorphic to the Veech group $\mathrm{SL}(X, \omega)$ as defined in Sect. 2.2.1. We call (X, ω) a *lattice surface* if $\mathrm{SL}(X, \omega)$ is a lattice group, i.e. $\mathrm{SL}(X, \omega)$ has finite covolume in $\mathrm{SL}_2\mathbb{R}$.

Remark Recall from Sect. 2.2.1 that a Veech group is always non-cocompact, and thus a lattice Veech group is non-uniform (i.e. it is a non-cocompact discrete group of finite covolume). Other necessary properties of lattice Veech groups are known; for example, the trace field of $\mathrm{SL}(X, \omega)$ is a totally real number field of degree at most g, where g is the genus of (X, ω) [10, 12, 16]). Throughout, when speaking of a lattice $\Gamma \leq \mathrm{SL}_2\mathbb{R}$ the reader may wish to impose such known restrictions, though they are not strictly necessary for our statements. By a lattice we simply mean a discrete subgroup of $\mathrm{SL}_2\mathbb{R}$ of finite covolume.

3 A canonical surface for each stratum

3.1 Canonical surface

Here we recall the canonical (infinite area and disconnected in general) flat surface associated to each stratum $\mathcal{H}(d_1, \ldots, d_\kappa)$ studied in [6] and [7] and explore some immediate results. For each $1 \leq i \leq \kappa$, let $\mathcal{O}_i := (\mathbb{C}, z^{d_i} dz)$, and set

$$\mathcal{O} = \mathcal{O}(d_1, \ldots, d_\kappa) := \bigsqcup_{i=1}^{\kappa} \mathcal{O}_i.$$

Each component $\mathcal{O}_i \subset \mathcal{O}$ is an infinite translation surface with a sole singularity of cone angle $2\pi(d_i + 1)$ at the origin, which we denote by $0 \in \mathcal{O}_i$ (the context will be clear as to which component an origin 0 belongs). Intuitively, we may think of $\mathcal{O}_i$ as $d_i + 1$ copies of the plane, denoted $c_0^i, \ldots, c_{d_i}^i$, each of which is slit along the non-negative real axis and glued so that the bottom edge, b_j^i, of the slit of c_j^i is identified with the top edge, t_{j+1}^i, of the slit of c_{j+1}^i for each $j \in \mathbb{Z}_{d_i+1}$; see Fig. 1. After making these identifications, we (arbitrarily) consider $t_{j+1}^i \sim b_j^i$ as part of c_{j+1}^i and not of c_j^i. In particular, $c_j^i \cap c_k^i = \{0\}$ for distinct $j, k \in \mathbb{Z}_{d_i+1}$.

Define $\boldsymbol{\pi}_i : \mathcal{O}_i \to \mathbb{C}$ by $\boldsymbol{\pi}_i(p) := \int_\gamma z^{d_i} dz$, where γ is any piecewise-smooth curve from 0 to p (note that the integral is independent of path due to Cauchy's integral theorem), and extend to a map $\boldsymbol{\pi} : \mathcal{O} \to \mathbb{C}$ by setting $\boldsymbol{\pi}|_{\mathcal{O}_i} := \boldsymbol{\pi}_i$ for each $i \in \{1, \ldots, \kappa\}$. Observe that $\boldsymbol{\pi}_i(\cdot)$ is a degree $d_i + 1$ branched covering with a single branch point $0 \in \mathbb{C}$, and the extended map $\boldsymbol{\pi}(\cdot)$ is a degree $\sum_{i=1}^{\kappa}(d_i + 1) = 2g - 2 + \kappa$ branched covering with sole branch point $0 \in \mathbb{C}$. Throughout, we denote by z_x and z_y the real and imaginary parts, respectively of $z \in \mathbb{C}$; hence

$\boldsymbol{\pi}(p) = \boldsymbol{\pi}(p)_x + i\boldsymbol{\pi}(p)_y$. Note that on sufficiently small neighborhoods in $\mathcal{O}$ (and away from singularities), the map $\boldsymbol{\pi}(\cdot)$ restricts to give a coordinate chart for the translation atlas associated to $\mathcal{O}$. Moreover, each $\mathcal{O}_i$ is naturally equipped with generalized polar coordinates: every point $p \in c_j^i \subset \mathcal{O}_i$ is determined by a magnitude $|p| := d(0, p) = |\boldsymbol{\pi}(p)|$ and, for $p \neq 0$, a principal argument value $\arg(p) := \arg(\boldsymbol{\pi}(p)) + 2\pi j \in [2\pi j, 2\pi(j + 1))$, where $|z|$ and $\arg(z) \in [0, 2\pi)$ are the Euclidean norm and principal value of the argument function, respectively, of $z \in \mathbb{C}$.

For each $i \in \{1, \ldots, \kappa\}$, define an affine automorphism $\rho_i : \mathcal{O}_i \to \mathcal{O}_i$ which acts as a counterclockwise rotation of angle 2π about the origin; note that ρ_i belongs to $\mathrm{Trans}(\mathcal{O}_i)$ since, in local coordinates given by $\boldsymbol{\pi}(\cdot)$, ρ_i has trivial linear part. We extend ρ_i to $\mathcal{O}$ by acting as the identity on all other components and observe that $\rho_i \in \mathrm{Trans}(\mathcal{O})$. For each integer $d \geq 0$, let $n(d) = \#\{1 \leq i \leq \kappa \mid d_i = d\}$. Having fixed d, let $i_1, \ldots, i_{n(d)}$ be the distinct indices of the $d_1, \ldots, d_\kappa$ which equal d, and let $S_{n(d)}$ denote the symmetric group on the set $\{i_1, \ldots, i_{n(d)}\}$. Extend $S_{n(d)}$ to act as the identity on all other indices, and for each $\alpha \in S_{n(d)}$, let $f_\alpha : \mathcal{O} \to \mathcal{O}$ permute the indices of $\mathcal{O}_1, \ldots, \mathcal{O}_\kappa$ by α, respecting polar coordinates. We again have that $f_\alpha \in \mathrm{Trans}(\mathcal{O})$ for each $\alpha \in S_{n(d)}$ with $n(d) \geq 1$.

Lemma 3.1 *The group Trans($\mathcal{O}$) is in bijective correspondence with*

$$\left(\prod_{\substack{d \geq 0, \\ n(d) \geq 1}} S_{n(d)} \right) \times \left(\prod_{i=1}^{\kappa} C_{d_i+1} \right),$$

where C_{d_i+1} is the cyclic group of order $d_i + 1$.

Sketch Lemma 2 of [7] shows that Trans($\mathcal{O}$) is generated by an action of

$$\prod_{\substack{d \geq 0, \\ n(d) \geq 1}} S_{n(d)}$$

together with an action of $\prod_{i=1}^{\kappa} C_{d_i+1}$ on $\mathcal{O}$, as naturally defined from the definitions of f_α and ρ_i above. That is, any element $\tau \in \mathrm{Trans}(\mathcal{O})$ may be written as a finite composition of various f_α's and ρ_i's. We may rearrange the order of this composition (possibly altering the indices of some ρ_i's) to collect all f_α's with α belonging to the same $S_{n(d)}$ and all ρ_i's corresponding to the same $\mathcal{O}_i$. That τ may be uniquely written in such a way is verified by considering the indices of images of 2π-sectors c_j^i. $\qquad\square$

Lemma 3.2 *Let $f \in \mathrm{Aff}(\mathcal{O})$ with $\mathrm{der}(f) = A$. Then*

$$\boldsymbol{\pi} \circ f = A \cdot \boldsymbol{\pi},$$

where the notation on the right denotes the usual action of a matrix A on the plane as a linear transformation.

Proof Since $\boldsymbol{\pi}(\cdot)$ restricts on sufficiently small neighborhoods to give a (bijective) coordinate map for the translation atlas of $\mathcal{O}$, we have—by definition of f—that for each $v \in \mathbb{R}^2$ in the image of such a neighborhood,

$$\boldsymbol{\pi} \circ f \circ \boldsymbol{\pi}^{-1}(v) = Av + b$$

for some $b \in \mathbb{R}^2$. This composition agrees on the images of the intersections of such neighborhoods; in particular, b depends only on the component $\mathcal{O}_i$ to which the neighborhood

belongs. Choose some neighborhood as above which contains $0 \in \mathcal{O}_i$ on its boundary. The above composition (and each of the maps comprising it) extends continuously to the closure of its domain. Since f sends singularities to singularities, setting $v = 0$ gives $b = 0$. As $\mathcal{O}_i$ was arbitrary, the result follows from the previous observations. $\square$

In Lemma 3 of [7], it is shown that $\mathrm{der}(\mathrm{Aff}(\mathcal{O})) = \mathrm{GL}_2\mathbb{R}$. To each $A \in \mathrm{GL}_2^+\mathbb{R}$ we associate a canonical $f_A \in \mathrm{Aff}^+(\mathcal{O})$ with $\mathrm{der}(f_A) = A$ as follows. Fix some $g_A \in \mathrm{Aff}^+(\mathcal{O})$ with $\mathrm{der}(g_A) = A$, and choose $\tau \in \mathrm{Trans}(\mathcal{O})$ so that the composition $f_A := \tau \circ g_A$ satisfies for each $1 \le i \le \kappa$ both (i) $f_A(\mathcal{O}_i) = \mathcal{O}_i$ and (ii) for each point $p \in \mathcal{O}_i$, the angle measured counterclockwise from p to $f_A(p)$ is non-negative and less than 2π. Existence of τ follows from Lemma 3.1. Notice that $\mathrm{der}(f_A) = A$ since $\mathrm{der}(\cdot)$ is a group homomorphism, $\tau \in \mathrm{Trans}(\mathcal{O}) = \ker(\mathrm{der})$ and $\mathrm{der}(g_A) = A$. Furthermore, conditions (i) and (ii) guarantee that f_A is unique, regardless of the initial choice of g_A (and subsequent choice of τ).

We also note that any $g_A \in \mathrm{Aff}^+(\mathcal{O})$ with $\mathrm{der}(g_A) = A \in \mathrm{GL}_2^+\mathbb{R}$ may be written uniquely as $g_A = \tau' \circ f_A$ for some $\tau' \in \mathrm{Trans}(O)$ (namely $\tau' = \tau^{-1}$ for τ as above), with f_A the canonical affine automorphism associated to A.

Definition 3.1 With notation as above, let

$$\mathrm{Aff}_C^+(\mathcal{O}) := \{f_A \in \mathrm{Aff}^+(\mathcal{O}) \mid A \in \mathrm{GL}_2^+\mathbb{R}\}$$

be the collection of canonical affine automorphisms of $\mathcal{O}$. For any $r \in \mathbb{R}_+$ and $p \in \mathcal{O}$, let

$$rp := f_{D(r)}(p),$$

where $f_{D(r)} \in \mathrm{Aff}_C^+(\mathcal{O})$ with $D(r) = \begin{pmatrix} r & 0 \\ 0 & r \end{pmatrix}$.

In particular, if $p \in \mathcal{O}_i$ with polar coordinates $(|p|, \theta)$, then $rp \in \mathcal{O}_i$ with polar coordinates $(r|p|, \theta)$. Note then that $\boldsymbol{\pi}(rp) = r\boldsymbol{\pi}(p)$.

While $\mathrm{Trans}(\mathcal{O})$ is generally non-abelian, we have the following commutativity result:

Lemma 3.3 *The subgroup $\mathrm{Trans}(\mathcal{O}) \le \mathrm{Aff}(\mathcal{O})$ belongs to the centralizer of $\mathrm{Aff}_C^+(\mathcal{O})$, i.e.*

$$\tau \circ f_A = f_A \circ \tau$$

for each $\tau \in \mathrm{Trans}(\mathcal{O})$ and $f_A \in \mathrm{Aff}_C^+(\mathcal{O})$.

Proof Let $\tau \in \mathrm{Trans}(\mathcal{O})$, $f_A \in \mathrm{Aff}_C^+(\mathcal{O})$ and $p \in \mathcal{O}$. We must show

$$\tau \circ f_A(p) = f_A \circ \tau(p).$$

Since $\mathrm{der}(\tau \circ f_A) = \mathrm{der}(f_A \circ \tau) = A$, both sides of the previous line are sent under $\boldsymbol{\pi}(\cdot)$ to the same point $z \in \mathbb{C}$ (Lemma 3.2). Suppose $p \in c_k^i$ and $\tau(\mathcal{O}_i) = \mathcal{O}_j$ with $\tau(c_k^i) = c_\ell^j$. Let r_k^i denote the set of points in $\mathcal{O}_i$ with argument $2\pi k$ (i.e. r_k^i is the ray along the positive real axis in c_k^i). Note that $\boldsymbol{\pi}(\cdot)$ sends each of $\tau \circ f_A(c_k^i)$ and $f_A \circ \tau(c_k^i)$ bijectively onto $\mathbb{C}$, and these sets are completely determined by the images $\tau \circ f_A(r_k^i)$ and $f_A \circ \tau(r_k^i)$ in $\mathcal{O}_j$: the former are the sets of points within angle 2π counterclockwise of the respective images of the ray r_k^i. Since there is only one point in each 2π-sector which maps to $z \in \mathbb{C}$ under $\boldsymbol{\pi}(\cdot)$, it suffices to show that $\tau \circ f_A(r_k^i) = f_A \circ \tau(r_k^i)$. By definition of f_A and r_k^i, the image $f_A(r_k^i)$ belongs to c_k^i, and thus $\tau \circ f_A(r_k^i)$ belongs to c_ℓ^j. But also $\tau(r_k^i)$ belongs to c_ℓ^j by assumption, and again by definition of f_A, the image $f_A \circ \tau(r_k^i)$ belongs to c_ℓ^j as well. Since the images $\tau \circ f_A(r_k^i)$ and $f_A \circ \tau(r_k^i)$ are rays emanating from 0 in c_ℓ^j and pointing in the same direction, we have $\tau \circ f_A(r_k^i) = f_A \circ \tau(r_k^i)$ and the result follows. $\square$
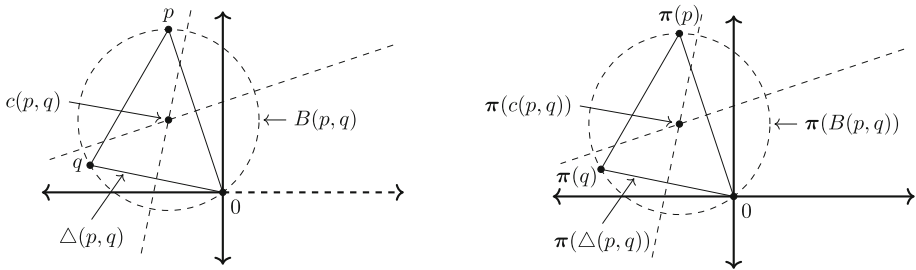
Fig. 2 Left: The triangle $\triangle(p, q)$, circumcenter $c(p, q)$, and ball $B(p, q)$ determined by p and q in $\mathcal{O}_i$. The circumcircle $C(p, q)$ is the boundary of $B(p, q)$, and the straight dashed lines through $c(p, q)$ are the boundaries of the half-spaces $H(p)$ and $H(q)$. Right: The image in $\mathbb{C}$ under $\pi(\cdot)$

3.2 Permissible triples

Here we introduce terminology and present results regarding particular subsets of the canonical surface $\mathcal{O} = \mathcal{O}(d_1, \ldots, d_\kappa)$. While the material of this subsection becomes technical, the underlying notions are rooted in elementary Euclidean geometry. As we shall see in Sects. 4.4 and 4.5, the results proven here will be crucial for our proof of Theorem 1.2. We begin with definitions and notation.

Definition 3.2 The *half-space* determined by $p \in \mathcal{O}_i$ is

$$H(p) := \{q \in \mathcal{O}_i \mid d(0, q) \leq d(p, q)\}.$$

Note that $H(p)$ is convex in the sense that the length-minimizing path between any two points in $H(p)$ is also contained in $H(p)$. Unless otherwise noted, a (closed or open) θ-*sector* of $\mathcal{O}_i$ is a (closed or open) sector of infinite radius, centered at $0 \in \mathcal{O}_i$, and of angle θ. For the following definition and subsequent discussion, see Fig. 2.

Definition 3.3 Let $p, q \in \mathcal{O}_i \backslash \{0\}$ be two points with distinct arguments in the same open π-sector of $\mathcal{O}_i$. The *triangle* determined by p and q, denoted $\triangle(p, q)$, is the union of straight line segments from $0 \in \mathcal{O}_i$ to p, from p to q, and from q to 0. Let

$$c(p, q) := \partial H(p) \cap \partial H(q)$$
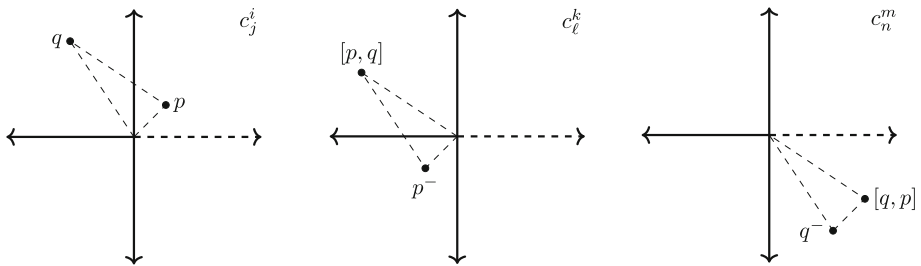
denote the *circumcenter* determined by p and q,

$$C(p, q) := \{z \in \mathcal{O}_i \mid d(z, c(p, q)) = |c(p, q)|\}$$

denote the *circumcircle* determined by p and q, and

$$B(p, q) := \{z \in \mathcal{O}_i \mid d(z, c(p, q)) < |c(p, q)|\}$$

denote the *ball* determined by p and q.

Locally, and away from a non-removable singularity, the geometry on $\mathcal{O}_i$ is Euclidean, so the circumcenter, circumcircle and ball of Definition 3.3 may be viewed as isometric copies of their Euclidean namesakes in the plane. To be more precise, Lemma 4 of [7] implies that for any two points $p_1, p_2 \in \mathcal{O}_i$ belonging to the same closed π-sector, the distance $d(p_1, p_2)$ between p_1 and p_2 in $\mathcal{O}_i$ equals the distance between their respective images $\pi(p_1)$ and $\pi(p_2)$ in the plane. Hence on any such sector, $\pi(\cdot)$ is an isometry. Using Euclidean geometry, we find that $\pi(c(p, q))$ is the center of the Euclidean circle $\pi(C(p, q))$ of radius $|c(p, q)|$,

Fig. 3 The points $[p, q]$ and $[q, p]$, as determined by $\{p, p^-\}$ and $\{q, q^-\}$

and $\pi(B(p, q))$ is the open ball whose boundary is $\pi(C(p, q))$. Moreover, $\pi(C(p, q))$ is the circumcircle of the Euclidean triangle $\pi(\triangle(p, q))$ with vertices $0, \pi(p), \pi(q)$, and $0, \pi(p), \pi(q) \in \pi(C(p, q))$ implies $0, p, q \in C(p, q) = \partial B(p, q)$.

Remark We use the same notation and terminology from Definition 3.3 for the analogous objects determined by two points in the Euclidean plane $\mathbb{C}$.

Definition 3.4 Define the set of *oppositely projected pairs* of points in $\mathcal{O}$ by

$$\mathbb{P}(\mathcal{O}) := \{\{p, p^-\} \subset \mathcal{O} \mid \pi(p^-) = -\pi(p) \neq 0\},$$

and for any subset $P \subset \mathbb{P}(\mathcal{O})$ of oppositely projected pairs, define the *forgotten* version of P by

$$P_F := \bigcup_{\{p, p^-\} \in P} \{p, p^-\};$$

that is, the pairing inherent to P is 'forgotten' in the set $P_F \subset \mathcal{O}$. For $r \in \mathbb{R}_+$ and $P \subset \mathbb{P}(\mathcal{O})$, let $rP := \{\{rp, rp^-\} \mid \{p, p^-\} \in P\}$ and $rP_F := (rP)_F$. Call a subset $P \subset \mathbb{P}(\mathcal{O})$ *limit-point free* if its forgotten version P_F has no limit points in $\mathcal{O}$.

Recall that $\pi(\cdot)$ is a degree $2g - 2 + \kappa$ branched covering with a single branch point $0 \in \mathbb{C}$. Thus for any $p \in \mathcal{O}$ with $\pi(p) \neq 0$, there are precisely $2g - 2 + \kappa$ points $p^- \in \mathcal{O}$ in the fiber above $-\pi(p)$. That is,

$$\left|\{p^- \mid \{p, p^-\} \in \mathbb{P}(\mathcal{O})\}\right| = 2g - 2 + \kappa.$$

The subsets $P \subset \mathbb{P}(\mathcal{O})$ of interest to us do not have this multiplicity of elements:

Definition 3.5 Call a nonempty subset $P \subset \mathbb{P}(\mathcal{O})$ of oppositely projected pairs *distinctive* if for any $p \in P_F$,

$$\left|\{p^- \mid \{p, p^-\} \in P\}\right| = 1,$$

or, equivalently, for any $p \in P_F$, there is a unique point $p^- \in P_F$ satisfying $\pi(p^-) = -\pi(p)$.

Now let $(p, q) \in \mathcal{O}^2$ be an ordered pair of two regular points with different arguments which belong to the same open π-sector of the same component of $\mathcal{O}$ (see the left-hand side of Fig. 3), and consider the difference $\pi(q) - \pi(p)$ of their images under $\pi(\cdot)$ in $\mathbb{C}$. As above, there are exactly $2g - 2 + \kappa$ points $u \in \mathcal{O}$ satisfying $\pi(u) = \pi(q) - \pi(p)$. Each of the $2g - 2 + \kappa$ oppositely projected pairs $\{p, p^-\} \in \mathbb{P}(\mathcal{O})$ naturally announces a unique such point u, namely the u lying nearest to p^-. A distinctive set P determines a unique point p^- and hence a unique point u. We fix special notation for this point u in the following:

Definition 3.6 Let $P, Q \subset \mathbb{P}(\mathcal{O})$ be distinctive subsets of oppositely projected pairs and $(\{p, p^-\}, \{q, q^-\}) \in P \times Q$. If p and q belong to the same open π-sector of the same component of $\mathcal{O}$ and satisfy $\arg(p) \neq \arg(q)$, then we denote by $[p, q]$ the unique element of $\mathcal{O}$ for which

(i) $\boldsymbol{\pi}([p, q]) = \boldsymbol{\pi}(q) - \boldsymbol{\pi}(p)$, and
(ii) $[p, q]$ belongs to the same open π-sector of the same component of $\mathcal{O}$ as the point p^-.

See Fig. 3.

Remark The reader should note that the point $[p, q]$ crucially depends on the distinctive set P, as distinctiveness uniquely determines the point p^- in Definition 3.6. However, as P shall be clear by context it is absent from the notation $[p, q]$.

Notice from Definition 3.6 that $[p, q]$ is defined if and only if $[q, p]$ is defined. The following technical result is needed for Proposition 3.6 below and is due to the fact that for distinctive sets, the points p and q naturally determine three isometric triangles in $\mathcal{O}$. See again Fig. 3.

Lemma 3.4 *Let $P, Q, U \subset \mathbb{P}(\mathcal{O})$ be distinctive subsets of oppositely projected pairs. For any $(\{p, p^-\}, \{q, q^-\}) \in P \times Q$ for which $[p, q] \in \mathcal{O}$ is defined, we have that $\{[p, q], [q, p]\} \in \mathbb{P}(\mathcal{O})$ is an oppositely projected pair. Moreover, if $\{[p, q], [q, p]\} \in U$ then $[p^-, [p, q]]$ and $[[p, q], p^-]$ are defined and equal q and q^-, respectively.*

Proof Definition 3.6 gives that

$$\boldsymbol{\pi}([p, q]) = \boldsymbol{\pi}(q) - \boldsymbol{\pi}(p) = -\boldsymbol{\pi}([q, p]).$$

Since p and q belong to the same π-sector with $\arg(p) \neq \arg(q)$, we have $\boldsymbol{\pi}(q) - \boldsymbol{\pi}(p) \neq 0$, so $\{[p, q], [q, p]\} \in \mathbb{P}(\mathcal{O})$.

Condition (ii) of Definition 3.6 guarantees that p^- and $[p, q]$ belong to the same open π-sector of the same component. Since p and q belong to the same open π-sector with different arguments, condition (i) guarantees that $\arg(p^-) \neq \arg([p, q])$, so $[p^-, [p, q]]$ and $[[p, q], p^-]$ are defined. Next we show $[p^-, [p, q]] = q$. Condition (ii) of Definition 3.6 (applied to $[p^-, [p, q]]$) gives that p and $[p^-, [p, q]]$ belong to the same open π-sector of the same component of $\mathcal{O}$; by assumption, the same is true of p and q, so we find that q and $[p^-, [p, q]]$ belong to the same open 2π-sector. As $\boldsymbol{\pi}(\cdot)$ is injective on open 2π-sectors, it suffices to show that these latter two points are sent under this map to the same point in the plane. We compute

$$\boldsymbol{\pi}([p^-, [p, q]]) = \boldsymbol{\pi}([p, q]) - \boldsymbol{\pi}(p^-) = \boldsymbol{\pi}(q) - \boldsymbol{\pi}(p) + \boldsymbol{\pi}(p) = \boldsymbol{\pi}(q)$$

as desired. The proof that $[[p, q], p^-]$ equals q^- is similar. $\square$

The following definition and subsequent results will prove to be essential in Sects. 4.4 and 4.5 below.

Definition 3.7 Let P, Q and U be distinctive, limit-point free subsets of $\mathbb{P}(\mathcal{O})$ and $(\{p, p^-\}, \{q, q^-\}, \{u, u^-\})$ a triple of oppositely projected pairs in $P \times Q \times U$. We call $(p, q, u) \in P_F \times Q_F \times U_F$ a *permissible triple* if there exist *permissible scalars* $(r, s, t) \in \mathbb{R}_+^3$ such that, for $(\{rp, rp^-\}, \{sq, sq^-\}, \{tu, tu^-\}) \in rP \times sQ \times tU$,

(i) $[rp, sq]$ and $[sq, rp]$ are defined and equal tu and tu^-, respectively, and
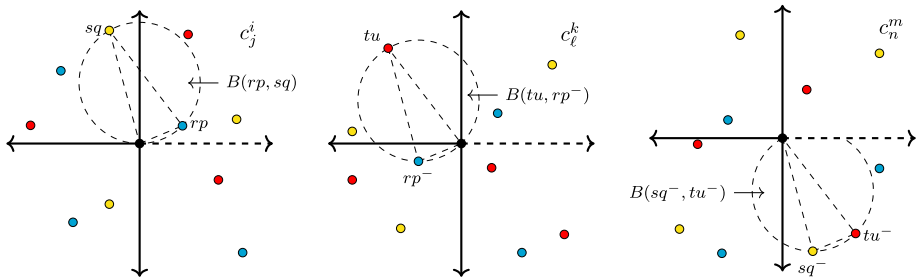
Fig. 4 A permissible triple (p, q, u) with permissible scalars (r, s, t). Points of $r P_F$, $s Q_F$ and $t U_F$ are in cyan, yellow, and red, respectively. Condition (i) of Definition 3.7 is met, as $[rp, sq] = tu$ and $[sq, rp] = tu^-$. Condition (ii) is met since none of the open balls $B(rp, sq)$, $B(tu, rp^-)$ or $B(sq^-, tu^-)$ contain any points of $r P_F$, $s Q_F$ or $t U_F$

(ii) $\left(B(rp, sq) \cup B(tu, rp^-) \cup B(sq^-, tu^-) \right) \cap (r P_F \cup s Q_F \cup t U_F) = \varnothing.$

(See Fig. 4.) For a collection $\mathbf{P} = \{P_i\}_{i \in I}$ of distinctive, limit-point free subsets of $\mathbb{P}(\mathcal{O})$, we let $\mathcal{P}(\mathbf{P})$ denote the *set of all permissible triples* arising from $\mathbf{P}$; that is,

$$\mathcal{P}(\mathbf{P}) := \{(p_i, p_j, p_k) \in (P_i)_F \times (P_j)_F \times (P_k)_F \mid i, j, k \in I \text{ and } (p_i, p_j, p_k)$$
$$\text{is a permissible triple}\}.$$

One immediate, but important, observation regarding permissible triples and permissible scalars is the following:

Proposition 3.5 *Let (p, q, u) be a permissible triple with permissible scalars $(r, s, t) \in \mathbb{R}_+^3$. The scalars s and t are uniquely determined by $p, q, u \in \mathcal{O}$ and $r \in \mathbb{R}_+$. In particular, the set of all permissible scalars for (p, q, u) is precisely the ray $\{(ar, as, at) \mid a \in \mathbb{R}_+\}$.*

Proof By Definitions 3.6 and 3.7, we have

$$t\boldsymbol{\pi}(u) = \boldsymbol{\pi}([rp, sq]) = s\boldsymbol{\pi}(q) - r\boldsymbol{\pi}(p). \tag{1}$$

Rearranging, we find that s and t are solutions to

$$\begin{pmatrix} \boldsymbol{\pi}(q)_x & -\boldsymbol{\pi}(u)_x \\ \boldsymbol{\pi}(q)_y & -\boldsymbol{\pi}(u)_y \end{pmatrix} \begin{pmatrix} s \\ t \end{pmatrix} = \begin{pmatrix} r\boldsymbol{\pi}(p)_x \\ r\boldsymbol{\pi}(p)_y \end{pmatrix}.$$

The fact that rp and sq belong to the same open π-sector with different arguments implies that $\boldsymbol{\pi}(p)$ and $\boldsymbol{\pi}(q)$ are $\mathbb{R}$-linearly independent. This, together with Equation (1), implies that also $\boldsymbol{\pi}(u)$ and $\boldsymbol{\pi}(q)$ are $\mathbb{R}$-linearly independent. Hence the matrix on the left side of the previous equation is invertible, so s and t are uniquely determined.

For the second statement, let P_1, $P_2 \subset \mathbb{P}(O)$ be distinctive and $a \in \mathbb{R}_+$. From Definition 3.6, one finds for $(\{p_1, p_1^-\}, \{p_2, p_2^-\}) \in P_1 \times P_2$ and $(\{ap_1, ap_1^-\}, \{ap_2, ap_2^-\}) \in a P_1 \times a P_2$ that if $[p_1, p_2]$ is defined then $[ap_1, ap_2]$ is defined and equals $a[p_1, p_2]$. It follows then from Definition 3.7 that if (r, s, t) are permissible scalars for (p, q, u), then so are (ar, as, at) for any $a \in \mathbb{R}_+$. On the other hand, if (r_1, s_1, t_1) and (r_2, s_2, t_2) are both permissible scalars for (p, q, u), then so are $(1, s_1/r_1, t_1/r_1)$ and $(1, s_2/r_2, t_2/r_2)$. The first statement of this proposition implies that (r_1, s_1, t_1) and (r_2, s_2, t_2) belong to a common ray. $\quad\square$

Note that the order of entries of a permissible triple (p, q, u) is important in Definition 3.7; for instance, the first two entries p and q must necessarily belong to the same open π-sector

of the same component of $\mathcal{O}$ for condition (i) to hold. Nevertheless, this ordering does admit some flexibility.

Proposition 3.6 *With notation as in Definition* 3.7, *the following statements are equivalent:*

(a) (p, q, u) *is a permissible triple with permissible scalars* (r, s, t),
(b) (q^-, u^-, p) *is a permissible triple with permissible scalars* (s, t, r), *and*
(c) (u, p^-, q^-) *is a permissible triple with permissible scalars* (t, r, s).

Proof We need only show (a) implies (b): the same argument will give (b) implies (c) and (c) implies (a). Let (p, q, u) be a permissible triple with permissible scalars (r, s, t). We claim that (q^-, u^-, p) is a permissible triple with permissible scalars (s, t, r).

(i) We must show that $[sq^-, tu^-]$ and $[tu^-, sq^-]$ are defined and equal rp and rp^-, respectively. We have by assumption that $[sq, rp]$ is defined and equals tu^-. By Lemma 3.4, both $[sq^-, [sq, rp]] = [sq^-, tu^-]$ and $[[sq, rp], sq^-] = [tu^-, sq^-]$ are defined. By the same Lemma, these equal rp and rp^-, respectively.
(ii) We have

$$\left(B(sq^-, tu^-) \cup B(rp, sq) \cup B(tu, rp^-)\right) \cap (sQ_F \cup tU_F \cup rP_F) = \varnothing$$

by assumption.

$\square$

One might suspect from Definition 3.7 that it is difficult for some $(p, q, u) \in P_F \times Q_F \times U_F$ to be a permissible triple. Indeed, this suspicion is confirmed by the following:

Lemma 3.7 *Let* $\mathbf{P} = \{P_i\}_{i=1}^n$ *be a finite collection of distinctive, limit-point free subsets of* $\mathbb{P}(\mathcal{O})$. *Then the set* $\mathcal{P}(\mathbf{P})$ *of all permissible triples arising from* $\mathbf{P}$ *is finite.*

Proof It suffices to show that for any distinctive, limit-point free $P, Q, U \subset \mathbb{P}(\mathcal{O})$, the set of permissible triples $(p, q, u) \in P_F \times Q_F \times U_F$ is finite. Suppose on the contrary that $\{(p_k, q_k, u_k)\}_{k\in\mathbb{N}} \subset P_F \times Q_F \times U_F$ is an infinite set of permissible triples with corresponding permissible scalars $\{(r_k, s_k, t_k)\}_{k\in\mathbb{N}} \subset \mathbb{R}_+^3$. At least one of the sets $\{p_k\}_{k\in\mathbb{N}}, \{q_k\}_{k\in\mathbb{N}}$ or $\{u_k\}_{k\in\mathbb{N}}$ is infinite; by Proposition 3.6, we may assume that either $\{p_k\}_{k\in\mathbb{R}}$ or $\{p_k^-\}_{k\in\mathbb{N}}$ is infinite. Since P is distinctive, these sets have the same cardinality and hence are both infinite. As the set of possible directions of points in each $\mathcal{O}_i$ is compact and there are only finitely many components $\mathcal{O}_i$ of $\mathcal{O}$, we may also assume that each p_k (resp. p_k^-) belongs to some small sector of a fixed component of $\mathcal{O}$ and that $\lim_k \arg(\boldsymbol{\pi}(p_k)) = \theta_P$ (resp. $\lim_k \arg(\boldsymbol{\pi}(p_k^-)) = \theta_P^-$) exists. Rotating each of P, Q and U by $-\theta_P$, assume for simplicity that $\theta_P = 0$ (and hence $\theta_P^- = \pi$).

Rescaling each (r_k, s_k, t_k) as necessary (see Proposition 3.5), we further assume that $|r_k p_k| = 1$ for all k. Under these assumptions, we find that $\{r_k \boldsymbol{\pi}(p_k)\}_{k\in\mathbb{N}} \subset S^1$ with $r_k \boldsymbol{\pi}(p_k) \to 1 \in \mathbb{C}$. Also note that since P is limit-point free, $|p_k| \to \infty$ and hence $r_k = 1/|p_k| \to 0$.

Condition (ii) of Definition 3.7 (together with the fact that $\boldsymbol{\pi}(\cdot)$ is an isometry on π-sectors) guarantees that in the plane, we have for each $k \in \mathbb{N}$ both
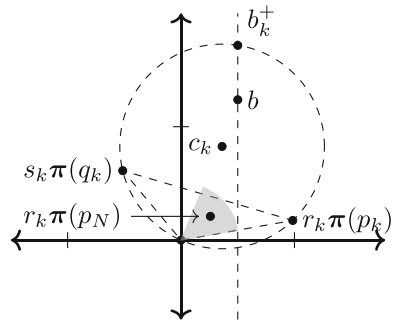
$$B(r_k\boldsymbol{\pi}(p_k), s_k\boldsymbol{\pi}(q_k)) \cap \{r_k\boldsymbol{\pi}(p_\ell)\}_{\ell\in\mathbb{N}} = \varnothing \tag{2}$$

and

$$B(t_k\boldsymbol{\pi}(u_k), r_k\boldsymbol{\pi}(p_k^-)) \cap \{r_k\boldsymbol{\pi}(p_\ell^-)\}_{\ell\in\mathbb{N}} = \varnothing.$$

Fig. 5 Illustration of the proof of Lemma 3.7. The vertical dashed line represents $x = 1/2$, and the shaded region is the subset S_k of the ball $B_k = B(r_k \boldsymbol{\pi}(p_k), s_k \boldsymbol{\pi}(q_k))$

This latter intersection may be rewritten

$$B(t_k \boldsymbol{\pi}(u_k), -r_k \boldsymbol{\pi}(p_k)) \cap \{-r_k \boldsymbol{\pi}(p_\ell)\}_{\ell \in \mathbb{N}} = \varnothing. \tag{3}$$

For each k, let $c_k := c(r_k \boldsymbol{\pi}(p_k), s_k \boldsymbol{\pi}(q_k))$ and $B_k := B(r_k \boldsymbol{\pi}(p_k), s_k \boldsymbol{\pi}(q_k))$ be the circumcenter and ball determined by $r_k \boldsymbol{\pi}(p_k)$ and $s_k \boldsymbol{\pi}(q_k)$. The circumcenter c_k is the intersection of the perpendicular bisectors of the straight line segments from the origin to $r_k \boldsymbol{\pi}(p_k)$ and from the origin to $s_k \boldsymbol{\pi}(q_k)$. Since $\arg(\boldsymbol{\pi}(p_k)) \to 0$, for large k the former perpendicular bisector does not intersect the negative real axis, so for all large k we have $\arg(c_k) \in [0, \pi) \cup (\pi, 2\pi)$. Passing to a subsequence, assume without loss of generality that $\arg(c_k) \in [0, \pi)$ for all k.

We consider two cases (see Fig. 5):

(i) Suppose there is a subsequence for which $\boldsymbol{\pi}(p_k)_y \geq 0$ for all k. Since $\arg(\boldsymbol{\pi}(p_k)) \to 0$, we may pass to a subsequence to assume $\arg(\boldsymbol{\pi}(p_{k+1})) \leq \arg(\boldsymbol{\pi}(p_k))$ for all k. For large enough k, the x-coordinate of $r_k \boldsymbol{\pi}(p_k) \in S^1$ is greater than $1/2$. Since both 0 and $r_k \boldsymbol{\pi}(p_k)$ belong to $C_k := \partial B_k$, we find that the vertical line $x = 1/2$ intersects C_k at two distinct points; let b_k^+ denote the point of intersection on the upper-semicircle of C_k. Note that $b_y := \inf_k \{(b_k^+)_y\} > 0$; otherwise $\arg(c_k) \in (\pi, 2\pi)$ contrary to our assumption. Set $b := (1/2, b_y)$ and note that for large enough k, we have $\arg(\boldsymbol{\pi}(p_k)) < \arg(b)$. Let

$$S_k = \{p \in \mathbb{C} \mid |p| < 1/2, \ \arg(p) \in (\arg(\boldsymbol{\pi}(p_k)), \arg(b))\}.$$

Note that S_k is contained in the triangle with vertices 0, $r_k \boldsymbol{\pi}(p_k)$ and b_k^+, which is contained in B_k; hence $S_k \subset B_k$ for each k. Choose some $N \in \mathbb{N}$ large enough that $\arg(\boldsymbol{\pi}(p_N)) \in [0, \arg(b))$. For each $k > N$, we have $\arg(\boldsymbol{\pi}(p_k)) \leq \arg(\boldsymbol{\pi}(p_N))$ by assumption, so $\arg(\boldsymbol{\pi}(p_N)) \in [\arg(\boldsymbol{\pi}(p_k)), \arg(b))$. Taking k large enough, we also have $|r_k \boldsymbol{\pi}(p_N)| = |\boldsymbol{\pi}(p_N)|/|\boldsymbol{\pi}(p_k)| < 1/2$. Hence $r_k \boldsymbol{\pi}(p_N) \in S_k \subset B_k = B(r_k \boldsymbol{\pi}(p_k), s_k \boldsymbol{\pi}(q_k))$, contradicting Eq. 2.

(ii) Suppose there is no subsequence for which $\boldsymbol{\pi}(p_k)_y \geq 0$ for all k. Then there is some subsequence for which $\boldsymbol{\pi}(p_k)_y < 0$—and hence $-\boldsymbol{\pi}(p_k)_y > 0$—for all k. Notice that $c(t_k \boldsymbol{\pi}(u_k), -r_k \boldsymbol{\pi}(p_k)) = c_k - r_k \boldsymbol{\pi}(p_k)$, so the argument of this circumcenter belongs to $(0, \pi)$. An analogous proof (reflecting about the x-axis) to that of case (i) implies that for some fixed N and large enough k, $-r_k \boldsymbol{\pi}(p_N) \in B(t_k \boldsymbol{\pi}(u_k), -r_k \boldsymbol{\pi}(p_k))$, which contradicts Eq. 3.

We conclude that the set of permissible triples must be finite. $\qquad \square$

3.3 Marked segments and Voronoi staples

In this subsection we fix a stratum $\mathcal{H}(d_1, \ldots, d_\kappa)$, its corresponding canonical surface $\mathcal{O} = \mathcal{O}(d_1, \ldots, d_\kappa)$ and a translation surface $(X, \omega) \in \mathcal{H}(d_1, \ldots, d_\kappa)$. Label the singularities of (X, ω) as $\sigma_1, \ldots, \sigma_\kappa$ so that σ_i has cone angle $2\pi(d_i+1)$. In a sufficiently small neighborhood of σ_i define generalized polar coordinates in a fashion analogous to that on $\mathcal{O}_i$ (see Sect. 3.1 and the proof of Lemma 5 of [7]).

Let s be a separatrix emanating from σ_i of length $|s| > 0$ and corresponding angle $\theta \in [0, 2\pi(d_i + 1))$, and denote by $\hat{s}$ the point of $\mathcal{O}_i$ with polar coordinates $(|s|, \theta)$. If $|s| = 0$, set $\hat{s} := 0 \in \mathcal{O}_i$. If s is in fact a saddle connection, we call $\hat{s}$ the *marked segment* determined by s; note that in this case, $\mathrm{hol}(s) = \boldsymbol{\pi}(\hat{s}) \in \mathbb{C}$. Denote by s' the identical, but oppositely-oriented saddle connection to s. Let $\mathcal{M}(X, \omega)$ be the set of all pairs $\{s, s'\}$ of oppositely-oriented saddle connections of (X, ω) and

$$\mathcal{M}_F(X, \omega) := \bigcup_{\{s, s'\} \in \mathcal{M}(X, \omega)} \{s, s'\}$$

the set of all oriented saddle connections on (X, ω). Let $\widehat{\mathcal{M}}(X, \omega)$ denote the set of all *orientation-paired marked segments* $\{\hat{s}, \hat{s}'\}$ determined by oppositely oriented saddle connections $\{s, s'\} \in \mathcal{M}(X, \omega)$.

Proposition 3.8 *The set $\widehat{\mathcal{M}}(X, \omega)$ is a limit-point free subset of the set $\mathbb{P}(\mathcal{O})$ of oppositely projected pairs. Moreover, if $\hat{s}$ and $\hat{t}$ are marked segments with identical arguments belonging to the same component of $\mathcal{O}$, then in fact $\{\hat{s}, \hat{s}'\} = \{\hat{t}, \hat{t}'\}$. In particular, $\widehat{\mathcal{M}}(X, \omega)$ is distinctive.*

Proof Suppose $\{\hat{s}, \hat{s}'\} \in \widehat{\mathcal{M}}(X, \omega)$. Since the saddle connections s and s' are identical but oppositely-oriented, it is clear that $\boldsymbol{\pi}(\hat{s}') = -\boldsymbol{\pi}(\hat{s})$, so $\{\hat{s}, \hat{s}'\} \in \mathbb{P}(\mathcal{O})$ and $\widehat{\mathcal{M}}(X, \omega) \subset \mathbb{P}(\mathcal{O})$. The image of $(\widehat{\mathcal{M}}(X, \omega))_F$ under $\boldsymbol{\pi}(\cdot)$ equals the image of $\mathcal{M}_F(X, \omega)$ under $\mathrm{hol}(\cdot)$. Since the set of holonomy vectors has no limit points (Sect. 2.1.2) and $\boldsymbol{\pi}(\cdot)$ is a homeomorphism on sufficiently small neighborhoods of regular points, we have that $(\widehat{\mathcal{M}}(X, \omega))_F$ has no limit points and hence $\widehat{\mathcal{M}}(X, \omega)$ is limit-point free.

Now suppose $\hat{s}$ and $\hat{t}$ are marked segments with identical arguments in the same component $\mathcal{O}_i \subset \mathcal{O}$. Then the underlying saddle connections s and t both emanate in the same direction from $\sigma_i \in \Sigma$. If their lengths differ, say $|\hat{s}| < |\hat{t}|$, then $|s| < |t|$. This implies that t has a singularity in its interior, which is impossible. If $|\hat{s}| = |\hat{t}|$, then in fact $s = t$, and so $\{\hat{s}, \hat{s}'\} = \{\hat{t}, \hat{t}'\}$. It follows that $\widehat{\mathcal{M}}(X, \omega)$ is distinctive. $\qquad\square$

Let $\widehat{\mathcal{M}}_F(X, \omega) := (\widehat{\mathcal{M}}(X, \omega))_F$ denote the set of all marked segments determined by saddle connections on (X, ω). The *star domain* for $\sigma_i \in \Sigma$ is

$$\mathrm{star}_i(X, \omega) := \{\hat{s} \mid s \text{ is a separatrix on } (X, \omega) \text{ emanating from } \sigma_i\} \subset \mathcal{O}_i.$$

Note that $\mathrm{star}_i(X, \omega)$ consists of the union of closed rays emanating from $0 \in \mathcal{O}_i$ which stop only when meeting a marked segment (and thus almost every such ray is infinite). The *star domain* for (X, ω) is

$$\mathrm{star}(X, \omega) = \bigsqcup_{i=1}^{\kappa} \mathrm{star}_i(X, \omega) \subset \mathcal{O}.$$

Define a map $\eta : \mathrm{star}(X, \omega) \to (X, \omega)$, where for each point $p \in \mathrm{star}_i(X, \omega)$, if s is the separatrix for which $p = \hat{s}$, then $\eta(p)$ is the endpoint of s on (X, ω). In other words, if
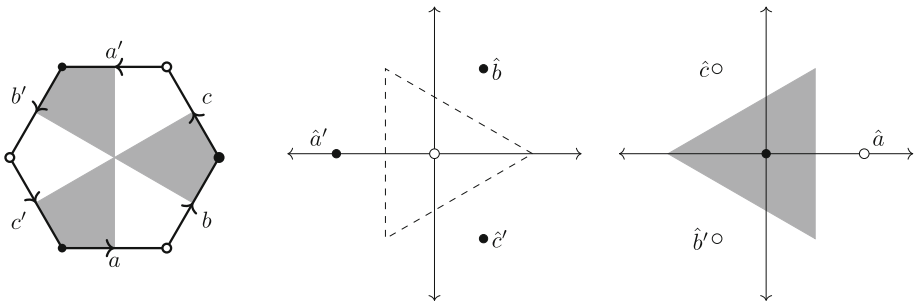
Fig. 6 Left: Voronoi decomposition of a translation surface subordinate to two removable singularities, σ_1 (white) and σ_2 (black). The 2-cell $C_1 := C_{\sigma_1}$ (resp. $C_2 := C_{\sigma_2}$) is the open region in white (resp. gray). Middle: The image $\iota(C_1)$ in $\mathcal{O}_1$, along with three marked segments. Right: The image $\iota(C_2)$ in $\mathcal{O}_2$, along with three marked segments

$p \in \text{star}(X, \omega)$ has polar coordinates $(|p|, \theta)$, then $\eta(p)$ is the endpoint of the separatrix of length $|p|$ emanating from $\sigma_i \in \Sigma$ with angle θ.

For each x in the Voronoi 2-cell $C_i := C_{\sigma_i}$, let s_x be the unique length-minimizing separatrix from σ_i to x. Note that η is injective—and thus invertible—on the set

$$\{\hat{s}_x \mid x \in C_i\} \subset \text{star}_i(X, \omega) :$$

if $\eta(\hat{s}_x) = \eta(\hat{s}_y)$ for $x, y \in C_i$, then $x = y$ and $s_x = s_y$ by the aforementioned uniqueness of these separatrices. Hence $\hat{s}_x = \hat{s}_y$. Let $\iota_i : C_i \to \{\hat{s}_x \mid x \in C_i\}$ denote the corresponding inverse, namely $x \mapsto \hat{s}_x$ for each $x \in C_i$, and define $\iota : \bigsqcup_{i=1}^{K} C_i \to \mathcal{O}$ by setting $\iota|_{C_i} = \iota_i$ for each i; see Fig. 6.

The translation surface (X, ω) is isometric to the quotient space of $\bigsqcup_{i=1}^{K} \overline{C_i}$ under the equivalence relation defined by identifying shared edges of Voronoi 2-cells. Proposition 7 of [7] shows that in a similar fashion, (X, ω) may be recovered from the closure of the image under ι of its Voronoi 2-cells by identifying appropriate edges of the various $\overline{\iota(C_i)}$. We provide a brief overview of the method by which these edge identifications are made. Recall that the half-space $H(p)$ determined by $p \in \mathcal{O}_i$ is convex, and thus so is any intersection of such half-spaces (see Definition 3.2).

Definition 3.8 For S a subset of $\mathcal{O}$ with no limit points, the *convex body* of $\mathcal{O}_i$ subordinate to S is defined by

$$\Omega_i(S) := \bigcap_{p \in S \cap \mathcal{O}_i} H(p).$$

The set of *essential points* of $\Omega_i(S)$ is the (unique) minimal subset $\mathcal{E}_i(S) \subset S$ for which

$$\Omega_i(S) = \bigcap_{p \in \mathcal{E}_i(S)} H(p).$$

When $S = \widehat{\mathcal{M}}_F(X, \omega)$ is the set of all marked segments of (X, ω), we use the suppressed notation $\Omega_i := \Omega_i(\widehat{\mathcal{M}}_F(X, \omega))$ and $\mathcal{E}_i := \mathcal{E}_i(\widehat{\mathcal{M}}_F(X, \omega))$. Call distinct points $p, q \in S$ *adjacent* within S if $p, q \neq 0$ belong to the same component $\mathcal{O}_i$, and either of the two open sectors centered at 0 between p and q contains no points of S.

Proposition 14 of [7] shows that the convex body Ω_i is precisely the set $\overline{\iota(C_i)}$. Furthermore, Proposition 14 and Definition 15 of [7] show that there is a subset $\widehat{\mathcal{S}}(X, \omega)$ of $\widehat{\mathcal{M}}(X, \omega)$ for

which the union of essential points $\mathcal{E} := \sqcup_{i=1}^{\kappa} \mathcal{E}_i$ equals $\widehat{\mathcal{S}}_F(X, \omega) := (\widehat{\mathcal{S}}(X, \omega))_F$, i.e. the essential points come equipped with a natural pairing. Elements $\{\hat{s}, \hat{s}'\}$ of $\widehat{\mathcal{S}}(X, \omega)$ are called *marked Voronoi staples* and their underlying pairs of saddle connections $\{s, s'\}$ in $\mathcal{M}(X, \omega)$ are called *Voronoi staples*. It follows from the definition of $\mathcal{E}_i$ that each edge on the boundary of $\Omega_i = \overline{\iota(\mathcal{C}_i)}$ belongs to the boundary of a half-space $H(\hat{s})$ for some $\hat{s} \in \mathcal{E}_i \subset \widehat{\mathcal{S}}_F(X, \omega)$; conversely, the boundary of the half-space determined by each $\hat{s} \in \widehat{\mathcal{S}}_F(X, \omega)$ contains an edge on the boundary of some $\Omega_i = \overline{\iota(\mathcal{C}_i)}$. Propositions 7 and 14 of [7] show that the edges of the convex bodies of $\sqcup_{i=1}^{\kappa} \Omega_i$ corresponding to $\hat{s}$ and its orientation-paired $\hat{s}'$ are equal length, and that (X, ω) is isometric to the quotient space of $\sqcup_{i=1}^{\kappa} \Omega_i$ under the equivalence relation given by identifying these edges via translation.

Of import for our purposes is the observation that a translation surface is uniquely determined by its marked Voronoi staples $\widehat{\mathcal{S}}(X, \omega)$, and that $\widehat{\mathcal{S}}_F(X, \omega)$ is precisely the union of essential points $\mathcal{E} = \sqcup_{i=1}^{\kappa} \mathcal{E}_i$.

Remark Under this identification of edges, the vertices of the various Ω_i are regarded as regular points on the resulting translation surface (they correspond to the Voronoi 0-cells of (X, ω)). The origins $0 \in \Omega_i \subset \mathcal{O}_i$ become singularities of the resulting translation surface; we require this also if $\mathcal{O}_i = (\mathbb{C}, dz)$, in which case the singularity is a marked point.

Example 3.9 In Fig. 6, the marked Voronoi staples are $\{\hat{a}, \hat{a}'\}$, $\{\hat{b}, \hat{b}'\}$ and $\{\hat{c}, \hat{c}'\}$. The hexagonal translation surface on the left is recovered by identifying the edges of the convex bodies corresponding to these respective pairs.

We conclude this subsection a brief observation:

Proposition 3.10 *Any two adjacent elements $\hat{s}$ and $\hat{t}$ of $\widehat{\mathcal{S}}_F(X, \omega)$ belong to the same open π-sector.*

Proof If not, then $\partial H(\hat{s}) \cap \partial H(\hat{t}) = \varnothing$, and the corresponding convex body has infinite area. This would imply that the closed translation surface (X, ω) has infinite area, which is a contradiction. $\qquad\square$

# 4 Fanning groups, simulations and finiteness of lattices in strata

4.1 Fanning groups and lattices

In this subsection we consider a set of directions $\Theta_\Gamma \subset S^1$ announced by a discrete subset $\Gamma \subset \mathrm{SL}_2\mathbb{R}$ and give an equivalent definition of lattice groups in terms of this set of directions (Lemma 4.3). Abusing notation, we write Γ for both a discrete subset of $\mathrm{SL}_2\mathbb{R}$ whose elements act as linear transformations on $\mathbb{R}^2$ and for its image as a discrete subset of $\mathrm{PSL}_2\mathbb{R}$ whose elements act as Möbius transformations on the closure of the upper-half plane $\mathbb{H}$. Similarly, we denote elements of Γ in both of these settings by the same notation $A \in \Gamma$. When there is risk of confusion, we shall explicitly mention the setting in which to consider Γ.

Definition 4.1 Let Γ be a discrete subset of $\mathrm{SL}_2\mathbb{R}$ and Δ the open unit disk. Set

$$\Theta_\Gamma := S^1 \setminus \bigcup_{A \in \Gamma} A \cdot \Delta,$$

where $A \cdot \Delta$ denotes the usual action of a matrix on $\mathbb{R}^2$ as a linear transformation. A group Γ for which Θ_Γ is finite is called a *fanning group*. We call (X, ω) a *fanning surface* if its Veech group $\mathrm{SL}(X, \omega)$ is fanning.
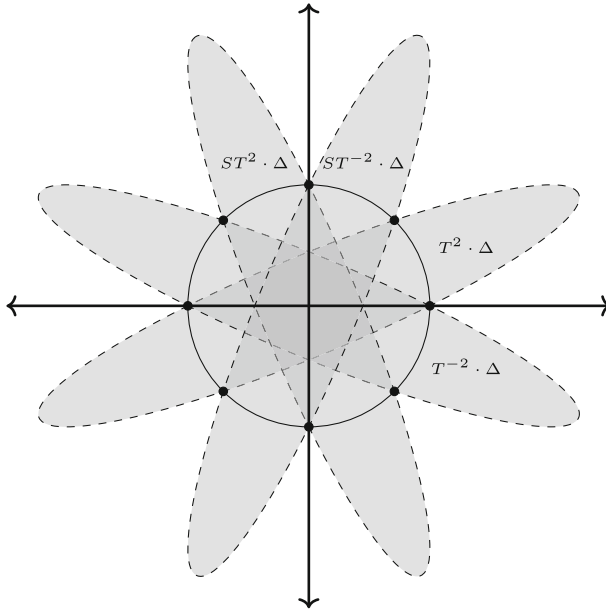
Fig. 7 The set of eighth roots of unity contains Θ_Γ with $\Gamma = \langle S, T^2 \rangle$ and S and T defined as in Example 4.1

As we shall see in Corollary 4.6 below, Θ_Γ contains the directions of certain short marked segments of any translation surface with Veech group Γ. These directions will be instrumental in our construction of such surfaces.

Example 4.1 Let $S := \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, $T := \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, and set $\Gamma := \langle S, T^2 \rangle$. Figure 7 shows that Θ_Γ is contained in the set of eighth roots of unity and thus Γ a fanning group. One verifies that in fact Θ_Γ equals this set.

We introduce notation to help us relate lattices to fanning groups. For $A \in \mathrm{PSL}_2 \mathbb{R}$, set

$$H(A) := \{ z \in \mathbb{H} \mid d_\mathbb{H}(i, z) \leq d_\mathbb{H}(Ai, z) \},$$

where $d_\mathbb{H}$ denotes the hyperbolic metric on $\mathbb{H}$, and for a discrete subset $\Gamma \subset \mathrm{PSL}_2 \mathbb{R}$, define

$$D(\Gamma) := \bigcap_{A \in \Gamma} H(A)$$

(compare with Definitions 3.2 and 3.8). Recall that if Γ is a Fuchsian group which trivially stabilizes $i \in \mathbb{H}$—or, equivalently, $\Gamma \cap \mathrm{PSO}_2 \mathbb{R} = \{\mathrm{Id}\}$—then $D(\Gamma)$ is the Dirichlet region centered at i for Γ. In this case $D(\Gamma)$ is a convex fundamental polygon for Γ; see, say, [1] and [15]. For any subset $S \subset \mathbb{H}$, let $\overline{S}$ and ∂S denote the closure and boundary, respectively, of S in $\hat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$. Define

$$S(A) := \{ x/y \mid (x, y)^T \in S^1 \backslash A \cdot \Delta \}$$

and

$$S_\Gamma := \bigcap_{A \in \Gamma} S(A) = \{ x/y \mid (x, y)^T \in \Theta_\Gamma \},$$

so that $S(A)$ and S_Γ belong to $\partial\mathbb{H} = \mathbb{R} \cup \{\infty\}$. The following proposition states that the reciprocals of the slopes of directions in $\Theta_\Gamma \subset S^1$ coincide with the points of $\overline{D(\Gamma)}$ in $\partial\mathbb{H}$.

Proposition 4.2 *For any discrete $\Gamma \subset \mathrm{SL}_2\mathbb{R}$,*

$$S_\Gamma = \overline{D(\Gamma)} \cap \partial\mathbb{H}.$$

Proof We first show that for any $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{SL}_2\mathbb{R}$,

$$S(A) = \overline{H(A)} \cap \partial\mathbb{H}.$$

If $A \in \mathrm{SO}_2\mathbb{R}$, then both sets are all of $\partial\mathbb{H}$, so suppose $A \notin \mathrm{SO}_2\mathbb{R}$. Assume that A satisfies $Ai = i + 2t_0$ for some $t_0 \in \mathbb{R}_+$. Then $H(A)$ consists of all points $\tau \in \mathbb{H}$ with $\mathrm{Re}(\tau) \le t_0$, and

$$\overline{H(A)} \cap \partial\mathbb{H} = \{t \in \mathbb{R} \mid t \le t_0\} \cup \{\infty\}.$$

From $(ai + b)/(ci + d) = i + 2t_0$ we find that

$$a = 2ct_0 + d \quad \text{and} \quad b = 2dt_0 - c,$$

and, moreover,

$$1 = \det(A) = ad - bc = c^2 + d^2.$$

Notice that $(x, y)^T \in S^1 \backslash A \cdot \Delta$ if and only if $(x, y)^T \in S^1$ and $A^{-1}(x, y)^T \notin \Delta$. The latter requirement may be rewritten as $|A^{-1}(x, y)^T| \ge |(x, y)^T|$. We claim that $(1, 0)^T$ and $\frac{1}{\sqrt{t_0^2+1}}(t_0, 1)^T$ satisfy this with equality. We compute

$$|A^{-1}(1, 0)^T|^2 = \left| \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = \left| \begin{pmatrix} d \\ -c \end{pmatrix} \right|^2 = c^2 + d^2 = 1 = |(1, 0)^T|^2$$

and

$$\begin{aligned}
|A^{-1}(t_0, 1)^T|^2 &= \left| \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} \begin{pmatrix} t_0 \\ 1 \end{pmatrix} \right|^2 \\
&= \left| \begin{pmatrix} dt_0 - b \\ -ct_0 + a \end{pmatrix} \right|^2 \\
&= (dt_0 - b)^2 + (-ct_0 + a)^2 \\
&= (-dt_0 + c)^2 + (ct_0 + d)^2 \\
&= d^2t_0^2 - 2cdt_0 + c^2 + c^2t_0^2 + 2cdt_0 + d^2 \\
&= t_0^2 + 1 \\
&= |(t_0, 1)^T|^2,
\end{aligned}$$

so the claim holds. Note that $A \cdot \Delta$ is the open region whose boundary is the ellipse $A \cdot S^1$ centered at the origin, so $S^1 \cap A \cdot \Delta$ consists of all points in S^1 with slope strictly between 0 and $1/t_0$ (see, for instance, Fig. 7, setting $A = T^2$ and $t_0 = 1$). It follows that

$$S(A) = \{t \in \mathbb{R} \mid t \le t_0\} \cup \{\infty\} = \overline{H(A)} \cap \partial\mathbb{H}.$$

Now consider general $A \in \Gamma \backslash \mathrm{SO}_2\mathbb{R}$, and fix $B \in \mathrm{SO}_2\mathbb{R}$ such that $BAi = i + 2t_0$ for some $t_0 \in \mathbb{R}_+$. From the argument above, we have

$$S(BA) = \overline{H(BA)} \cap \partial\mathbb{H}.$$

We claim that $S(BA) = BS(A)$ and $H(BA) = BH(A)$. By definition, $t \in S(BA)$ if and only if there is some $(x, y)^T \in S^1 \backslash BA \cdot \Delta$ with $x/y = t$. The latter inclusion is equivalent to $B^{-1}(x, y)^T \in B^{-1}(S^1 \backslash BA \cdot \Delta) = S^1 \backslash A \cdot \Delta$ since $B^{-1} \in SO_2\mathbb{R}$. Writing $B^{-1} = \left(\begin{smallmatrix} e & f \\ g & h \end{smallmatrix} \right)$, we find $B^{-1}(x, y)^T = (ex + fy, gx + hy)^T$. But $B^{-1}t = B^{-1}(x/y) = (ex + fy)/(gx + hy)$ is the reciprocal of the slope of $B^{-1}(x, y)^T \in S^1 \backslash A \cdot \Delta$, so $t \in S(BA)$ if and only if $B^{-1}t \in S(A)$, as desired. For the second claim, note that $\tau \in H(BA)$ if and only if $d_{\mathbb{H}}(i, \tau) \leq d_{\mathbb{H}}(BAi, \tau)$. Since B^{-1} is an isometry fixing i, this inequality is equivalent to $d_{\mathbb{H}}(i, B^{-1}\tau) \leq d_{\mathbb{H}}(Ai, B^{-1}\tau)$, which is true if and only if $B^{-1}\tau \in H(A)$.

Next we show that

$$\overline{D(\Gamma)} = \bigcap_{A \in \Gamma} \overline{H(A)}.$$

The forward inclusion holds as the closure of an intersection of sets is always contained in the intersection of the closures of the sets. For the reverse inclusion, suppose $\tau \in \bigcap_{A \in \Gamma} \overline{H(A)} \backslash \{i\}$, and let γ be a geodesic segment in $\mathbb{H}$ without endpoints, for which $\overline{\gamma}$ has endpoints i and τ. Since $\tau \in \overline{H(A)}$ for each $A \in \Gamma$ and each $H(A)$ is convex with i in its interior, we have $\gamma \subset \bigcap_{A \in \Gamma} H(A) = D(\Gamma)$. Hence $\tau \in \overline{\gamma} \subset \overline{D(\Gamma)}$, so the claim holds.

It follows that

$$S_\Gamma = \bigcap_{A \in \Gamma} S(A) = \bigcap_{A \in \Gamma} \left(\overline{H(A)} \cap \partial\mathbb{H} \right) = \left(\bigcap_{A \in \Gamma} \overline{H(A)} \right) \cap \partial\mathbb{H} = \overline{D(\Gamma)} \cap \partial\mathbb{H}.$$

<div align="right">□</div>

With the aid of Proposition 4.2, lattices may be characterized as the finitely generated discrete groups Γ for which the set of directions Θ_Γ is finite. Furthermore, for a lattice Γ, the set of directions Θ_Γ may be computed in finite time:

Lemma 4.3 *A discrete subgroup $\Gamma \subset SL_2\mathbb{R}$ is a lattice if and only if it is a finitely generated fanning group. In this case there is some finite subset $\Gamma_n \subset \Gamma$ for which $\Theta_\Gamma = \Theta_{\Gamma_n}$.*

Proof Let Λ denote the image in $PSL_2\mathbb{R}$ of $\Gamma \cap SO_2\mathbb{R}$. As Λ contains only elliptic elements, it is a finite cyclic group of order, say, n (Corollary 2.4.2 of [15]). Let $A_0 \in \Lambda$ be a generator of this group. Recall that if $n = 1$, then $D(\Gamma)$ is the Dirichlet region centered at i for Γ and is thus a convex fundamental polygon. If $n > 1$, then $D(\Gamma)$ is no longer a fundamental domain: one finds that $A_0 D(\Gamma) = D(\Gamma)$, and the interior of $D(\Gamma)$ contains n points from each orbit $\Gamma\tau$, $\tau \in \mathbb{H}$. Let γ_0 denote a geodesic segment beginning at i and ending on $\partial D(\Gamma)$. For each $0 \leq j < n$, the geodesic segment $\gamma_j := A_0^j \gamma_0$ also begins and ends at i and $\partial D(\Gamma)$. Let $\mathcal{F}_0 \subset D(\Gamma)$ denote the convex polygon bounded by γ_0, γ_1 and $\partial D(\Gamma)$. A slight generalization of the proof of Theorem 3.2.2 of [15] shows that $\mathcal{F}_0$ is a convex fundamental polygon for Γ. In particular,

$$\mu_{\mathbb{H}}(D(\Gamma)) = n\mu_{\mathbb{H}}(\mathcal{F}_0) = n\mu_{\mathbb{H}}(\Gamma\backslash\mathbb{H}),$$

where $\mu_{\mathbb{H}}$ denotes hyperbolic area. Hence Γ is a lattice if and only if $\mu_{\mathbb{H}}(D(\Gamma))$ is finite.

Suppose Γ is a lattice. Then Γ is geometrically finite and hence finitely generated. If Γ is not a fanning group, then Θ_Γ—and, consequently, S_Γ—is infinite. By Proposition 4.2, $\overline{D(\Gamma)} \cap \partial\mathbb{H}$ is infinite, and by Gauss-Bonnet, $D(\Gamma)$ has infinite hyperbolic area. This is a contradiction, so Γ is a finitely generated fanning group.

Now assume Γ is a finitely generated fanning group. Again by Proposition 4.2, this implies that $\overline{D(\Gamma)} \cap \partial\mathbb{H}$ is finite, so $\mathcal{F}_0 \subset D(\Gamma)$ has no free sides and at most finitely many cusps.

Since Γ is finitely generated, it is geometrically finite and thus every convex fundamental polygon for Γ has finitely many sides (Theorem 10.1.2 of [1]). By Gauss-Bonnet, Γ is a lattice, and the first statement is proven.

For the second statement, suppose Γ is a lattice. From the previous paragraph, we know that $\mathcal{F}_0$ has finitely many sides, and thus the same is true of $D(\Gamma)$. But then there is some finite subset $\Gamma_n \subset \Gamma$ for which $D(\Gamma_n) = D(\Gamma)$, and by Proposition 4.2, $S_\Gamma = S_{\Gamma_n}$, which implies $\Theta_\Gamma = \Theta_{\Gamma_n}$.   □

4.2 The group $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ and its action on marked segments

Let $(X, \omega) \in \mathcal{H}(d_1, \ldots, d_\kappa)$ and $\mathcal{O} = \mathcal{O}(d_1, \ldots, d_\kappa)$. A central result of [6] (Theorem 18) and [7] (Proposition 17) is that a matrix $A \in \mathrm{SL}_2\mathbb{R}$ belongs to $\mathrm{SL}(X, \omega)$ if and only if there is some $f \in \mathrm{Aff}^+(\mathcal{O})$ with $\mathrm{der}(f) = A$ satisfying $f(\widehat{\mathcal{S}}(X, \omega)) \subset \widehat{\mathcal{M}}(X, \omega)$. Theorem 17 of [6] shows that the same statement holds with the latter subset inclusion replaced by the equality $f(\widehat{\mathcal{M}}(X, \omega)) = \widehat{\mathcal{M}}(X, \omega)$. Note, in particular, that these statements require the affine automorphism f to respect the orientation-pairing of marked segments—it is not enough that the set of marked segments $\widehat{\mathcal{M}}_F(X, \omega)$ is invariant under f to conclude that $\mathrm{der}(f) \in \mathrm{SL}(X, \omega)$ (see Example 21 of [7]). We introduce the following notation for the collection of such affine automorphisms of $\mathcal{O}$:

Definition 4.2 Let

$$\mathrm{Aff}^+_\mathcal{O}(X, \omega) := \{f \in \mathrm{Aff}^+(\mathcal{O}) \mid f(\widehat{\mathcal{M}}(X, \omega)) = \widehat{\mathcal{M}}(X, \omega)\}.$$

Note that $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ is a subgroup of $\mathrm{Aff}^+(\mathcal{O})$. From the comments above we see that the image of $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ under $\mathrm{der}(\cdot)$ is precisely the Veech group $\Gamma = \mathrm{SL}(X, \omega)$ (though $\mathrm{der}(\cdot)$ restricted to $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ need not be injective, so $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ and $\mathrm{SL}(X, \omega)$ are in general non-isomorphic). The group $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ acts on $\widehat{\mathcal{M}}(X, \omega)$ via

$$\mathrm{Aff}^+_\mathcal{O}(X, \omega) \times \widehat{\mathcal{M}}(X, \omega) \to \widehat{\mathcal{M}}(X, \omega)$$
$$(f, \{\hat{s}, \hat{s}'\}) \mapsto \{f(\hat{s}), f(\hat{s}')\}.$$

Note in particular that $f(\hat{s}') = f(\hat{s})'$. For any $\{\hat{s}, \hat{s}'\} \in \widehat{\mathcal{M}}(X, \omega)$, let $[\{\hat{s}, \hat{s}'\}]$ denote the orbit of $\{\hat{s}, \hat{s}'\}$ under this action.

Recall from Sect. 3.3 that (X, ω) may be reconstructed from its marked Voronoi staples $\widehat{\mathcal{S}}(X, \omega) \subset \widehat{\mathcal{M}}(X, \omega)$. The goal of Algorithm 1.1 is to construct—using only $\Gamma = \mathrm{SL}(X, \omega)$ and $d_1 \leq \cdots \leq d_\kappa$—increasing subsets of $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$-orbits $[\{\hat{s}, \hat{s}'\}]$ whose union eventually contains $\widehat{\mathcal{S}}(X, \omega)$. To this end, the two main obstacles we face are

(i) determining $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$, and
(ii) determining a representative $\{\hat{s}, \hat{s}'\}$ of each of the $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$-orbits $[\{\hat{s}, \hat{s}'\}]$ intersecting $\widehat{\mathcal{S}}(X, \omega)$

using only the data Γ and $d_1 \leq \cdots \leq d_\kappa$. The first challenge is addressed by the following:

Lemma 4.4 *Let G be a set of generators of $\Gamma \leq \mathrm{SL}_2\mathbb{R}$. Then for any $(X, \omega) \in \mathcal{H}(d_1, \ldots, d_\kappa)$ with $\mathrm{SL}(X, \omega) = \Gamma$, the group $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ is generated by a subset of*

$$G_\mathcal{O} := \{\tau \circ f_A \mid \tau \in \mathit{Trans}(\mathcal{O}), \ f_A \in \mathit{Aff}^+_C(\mathcal{O}), \ A \in G\}.$$

In particular, if G is finite then $\mathrm{Aff}^+_\mathcal{O}(X, \omega)$ is generated by a subset of the finite set $G_\mathcal{O}$.

Proof For each $A \in G$, fix $g_A \in \text{Aff}_{\mathcal{O}}^+(X, \omega)$ with $\text{der}(g_A) = A$. Recall that g_A may be written uniquely as $g_A = \tau_A \circ f_A$, where $\tau_A \in \text{Trans}(\mathcal{O})$ and $f_A \in \text{Aff}_C^+(\mathcal{O})$, so the set of such g_A is contained in $G_{\mathcal{O}}$. Let Λ be the subgroup of $\text{Aff}_{\mathcal{O}}^+(X, \omega)$ generated by these g_A. Since $\text{Trans}(\mathcal{O})$ is closed under composition, it suffices to show that $\text{Aff}_{\mathcal{O}}^+(X, \omega) = \cup \tau \Lambda$, where the union is over some collection of $\tau \in \text{Trans}(\mathcal{O})$.

Now let $f \in \text{Aff}_{\mathcal{O}}^+(X, \omega)$ be arbitrary and $A := \text{der}(f) \in \text{SL}(X, \omega)$. The matrix A may be written as a product $A = A_1^{\delta_1} \cdots A_m^{\delta_m}$, $\delta_j \in \{\pm 1\}$, of elements A_j in the generating set G and their inverses. Set $g := g_{A_1}^{\delta_1} \ldots g_{A_m}^{\delta_m} \in \Lambda$, with each g_{A_j} a chosen generator of Λ as above. Then $\text{der}(f \circ g^{-1}) = \text{Id}$, so $\tau := f \circ g^{-1} \in \text{Trans}(\mathcal{O})$ and $f \in \tau \Lambda$ as desired.

The final statement of the Lemma follows immediately from finiteness of both $\text{Trans}(\mathcal{O})$ (Lemma 3.1) and G. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$

For challenge (ii) mentioned above, we wish to determine the generalized polar coordinates of representatives of $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits. In the proof of the following Lemma—which is also crucial for Algorithm 1.1—it is shown that the directions of the shortest representatives $\{\hat{s}, \hat{s}'\}$ of each orbit $[\{\hat{s}, \hat{s}'\}]$ are announced by the set $\Theta_{\text{SL}(X, \omega)}$. The more delicate procedure of determining the lengths of such representatives is addressed in Sect. 4.5.

Lemma 4.5 *If (X, ω) is a fanning surface, then the orbit space*

$$\widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega) = \left\{ [\{\hat{s}, \hat{s}'\}] \,\big|\, \{\hat{s}, \hat{s}'\} \in \widehat{\mathcal{M}}(X, \omega) \right\}$$

is finite.

Proof Let $[\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega)$. The image under $\pi(\cdot)$ of $[\{\hat{s}, \hat{s}'\}]$ consists of pairs $\{\pm v\}$ of holonomy vectors in $\mathbb{C}$, among which there is some pair of minimal length. The preimage under $\pi(\cdot)$ of this pair is a finite set in $\mathcal{O}$ containing a pair of orientation-paired marked segments $\{\hat{s}, \hat{s}'\}$ of minimal length in $[\{\hat{s}, \hat{s}'\}]$. Rescaling (X, ω) if necessary, we may assume $|\hat{s}| = |\hat{s}'| = 1$. We claim that

$$\left\{ \pi(\hat{s}), \pi(\hat{s}') \right\} \subset \Theta_{\text{SL}(X, \omega)}.$$

Note for each $A \in \text{SL}(X, \omega)$ that $\pi(\hat{s}) \notin S^1 \cap A \cdot \Delta$: otherwise $A^{-1} \cdot \pi(\hat{s}) \in \Delta$, but the comments following Definition 4.2 together with Lemma 3.2 imply $A^{-1} \cdot \pi(\hat{s}) = \pi \circ f(\hat{s})$ for some $f \in \text{Aff}_{\mathcal{O}}^+(X, \omega)$ with $\text{der}(f) = A^{-1}$. Since $1 > |\pi \circ f(\hat{s})| = |f(\hat{s})|$, this contradicts the assumption that $\{\hat{s}, \hat{s}'\}$ is a pair of minimal length in $[\{\hat{s}, \hat{s}'\}]$. The same is true for $\pi(\hat{s}')$, so we find that

$$\left\{ \pi(\hat{s}), \pi(\hat{s}') \right\} \subset \bigcap_{A \in \text{SL}(X, \omega)} (S^1 \backslash A \cdot \Delta) = S^1 \backslash \bigcup_{A \in \text{SL}(X, \omega)} A \cdot \Delta = \Theta_{\text{SL}(X, \omega)}$$

as claimed.

For each $v \in S^1$, let r_v denote the infinite open ray emanating from $0 \in \mathbb{C}$ in the direction of v. Note that the preimage $\pi^{-1}(r_v) \subset \mathcal{O}$ consists of $\sum_{i=1}^{\kappa} (d_i + 1) = 2g - 2 + \kappa$ infinite open rays emanating from the origins of the various components of $\mathcal{O}$, and each of these rays contains at most one marked segment of (X, ω) (Proposition 3.8). From the claim above, the minimal-length representatives of each $[\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega)$ belong to one of the $2g - 2 + \kappa$ open rays of $\pi^{-1}(r_v)$ for some $v \in \Theta_{\text{SL}(X, \omega)}$. Hence

$$\left| \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega) \right| \leq (2g - 2 + \kappa) \cdot |\Theta_{\text{SL}(X, \omega)}|,$$

where $|\cdot|$ denotes cardinality. Since (X, ω) is a fanning surface, the right-hand side is finite. $\qquad\square$

The proof of Lemma 4.5 also gives the following:

Corollary 4.6 *Let (X, ω) be a fanning surface and $[\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega)$. If $\{\hat{s}, \hat{s}'\}$ are orientation-paired marked segments of minimal length in $[\{\hat{s}, \hat{s}'\}]$, then $\pi(\hat{s})/|\pi(\hat{s})|$ and $\pi(\hat{s}')/|\pi(\hat{s}')|$ belong to the finite set $\Theta_{\text{SL}(X,\omega)}$.*

4.3 Simulating normalized $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits

The set of orientation-paired marked segments $\widehat{\mathcal{M}}(X, \omega)$ (and the set of $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits into which it partitions) depends intrinsically on (X, ω) and its geometry. The results of the previous subsection suggest that in the case of a fanning surface, much of this information is encoded in the Veech group $\text{SL}(X, \omega)$. This subsection further explores these ideas by introducing *simulations* of $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits which are constructed via generators of a fanning group. We again fix a stratum $\mathcal{H}(d_1, \dots, d_\kappa)$.

Definition 4.3 Let $\Gamma \leq \text{SL}_2\mathbb{R}$ be a fanning group generated by $G \subset \Gamma$ and set

$$\mathfrak{S}_G := \{(H, \{p, p^-\}) \in 2^{G_{\mathcal{O}}} \setminus \{\varnothing\} \times \mathbb{P}(\mathcal{O}) \mid \pi(p), \pi(p^-) \in \Theta_\Gamma\},$$

where $G_{\mathcal{O}}$ is defined as in Lemma 4.4 and $2^{G_{\mathcal{O}}}$ is the power set of $G_{\mathcal{O}}$. For each $\mathfrak{s} := (H, \{p, p^-\}) \in \mathfrak{S}_G$, define the *stage 0 simulation* determined by $\mathfrak{s}$ by $\text{sim}^0(\mathfrak{s}) := \{\{p, p^-\}\} \subset \mathbb{P}(\mathcal{O})$, and for $n \geq 1$ define the *stage n simulation* determined by $\mathfrak{s}$ recursively as

$$\text{sim}^n(\mathfrak{s}) := \bigcup_{f \in H \cup \{\text{Id}\}} f^{\pm 1}(\text{sim}^{n-1}(\mathfrak{s})) \subset \mathbb{P}(\mathcal{O}).$$

The *simulation* determined by $\mathfrak{s}$ is the union $\text{sim}(\mathfrak{s}) := \bigcup_{n \geq 0} \text{sim}^n(\mathfrak{s})$ of all stage n simulations determined by $\mathfrak{s}$. Denote by

$$\text{Sims}_G := \{\text{sim}(\mathfrak{s}) \mid \text{sim}(\mathfrak{s}) \text{ is distinctive}, \ \mathfrak{s} \in \mathfrak{S}_G\} \quad \text{and}$$
$$\text{Sims}_G^n := \{\text{sim}^n(\mathfrak{s}) \mid \text{sim}^n(\mathfrak{s}) \text{ is distinctive}, \ \mathfrak{s} \in \mathfrak{S}_G\}$$

the *set of all distinctive simulations* and *set of all stage n distinctive simulations*, respectively, determined by G. For any $r \in \mathbb{R}_+$, the *r-scaled (stage n)* simulation is $r\,\text{sim}(\mathfrak{s})$ ($r\,\text{sim}^n(\mathfrak{s})$). Finally, set $\text{sim}_F(\mathfrak{s}) := (\text{sim}(\mathfrak{s}))_F \subset \mathcal{O}$ and $\text{sim}_F^n(\mathfrak{s}) := (\text{sim}^n(\mathfrak{s}))_F \subset \mathcal{O}$ (recall Definition 3.4).

Note that by construction, the stage n simulation $\text{sim}^n(\mathfrak{s})$ is the set of images in $\mathcal{O}$ of the oppositely-projected pair $\{p, p^-\}$ under all compositions of at most n elements of H and their inverses, and the simulation $\text{sim}(\mathfrak{s})$ is simply the $\langle H \rangle$-orbit of $\{p, p^-\}$. Furthermore, the distinctive simulations in Sims_G and Sims_G^n are limit-point free since Γ is discrete.

Recall S, T and Γ from Example 4.1.

Example 4.7 Let $G := \{S, T^2\}$, $\mathcal{O} = \mathcal{O}(2)$, and set

$$\mathfrak{s}_0 := (H, \{p_0, p_0^-\}) \in \mathfrak{S}_G,$$

where $H = \{\rho_1^2 \circ f_S, f_{T^2}\}$, and $p_0, p_0^- \in \mathcal{O}_1 = \mathcal{O}$ have unit length (this is required since $\pi(p_0), \pi(p_0^-) \in \Theta_\Gamma \subset S^1$) and arguments $\arg(p_0) = 0$ and $\arg(p_0^-) = 5\pi$. The map $\rho_1^2 \circ f_S$ acts as a counterclockwise rotation of $\mathcal{O}$ by an angle of $\pi/2 + 4\pi = 9\pi/2$, and f_{T^2} is a bijective horizontal shear on each c_j^1. A subset of the simulation $\text{sim}(\mathfrak{s}_0)$ is shown in cyan in Fig. 8. In red and yellow are subsets of $\text{sim}(\mathfrak{s}_1)$ and $\text{sim}(\mathfrak{s}_2)$, respectively, for
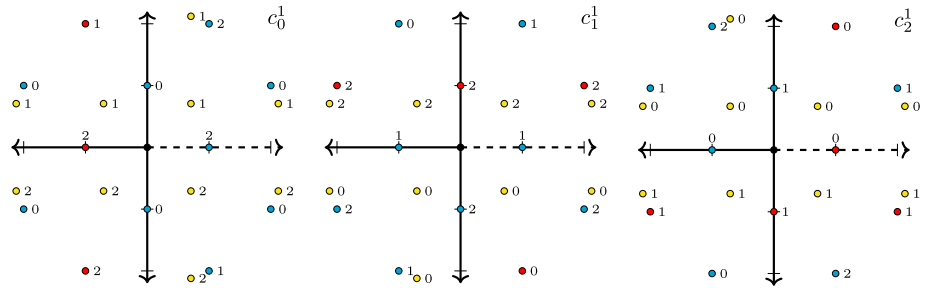
Fig. 8 Three (subsets) of the simulations $\text{sim}(\mathfrak{s}_0)$, $\text{sim}(\mathfrak{s}_1)$, $\text{sim}(\mathfrak{s}_2) \in \text{Sims}_G$ from Example 4.7 are shown in cyan, red and yellow, respectively. A label j next to a point p indicates to which c_j^1 the point p^- belongs

$$\mathfrak{s}_1 := (H, \{p_1, p_1^-\}) \in \mathfrak{S}_G$$

and

$$\mathfrak{s}_2 := (H, \{p_2, p_2^-\}) \in \mathfrak{S}_G,$$

where $\arg(p_1) = 4\pi$, $\arg(p_1^-) = \pi$, $\arg(p_2) = \pi/4$ and $\arg(p_2^-) = 13\pi/4$.

Observe that Definition 4.3 and Lemmas 4.3 and 4.4 immediately imply:

Lemma 4.8 *If Γ is a lattice, then there exists a set of generators G of Γ for which the set of simulations Sims_G is finite.*

Remark Recall from Lemma 4.3 that for a lattice Γ, the set of directions Θ_Γ may be computed in finite time. Indeed, letting Γ_n be the subset of Γ consisting of all words of length no greater than n in the generators G and their inverses, there is some $n \in \mathbb{N}$ for which $\Theta_{\Gamma_n} = \Theta_\Gamma$. In general, $\Theta_{\Gamma_n} \supset \Theta_\Gamma$, and it may be difficult to determine when equality holds. Thus in practice, the set Θ_Γ in the definition of $\mathfrak{S}_G$ (Definition 4.3) is replaced by some finite Θ_{Γ_n} containing Θ_Γ. Such a replacement does not affect the statements of the results in the remainder of the paper, though it could increase the size of the—still finite—set of simulations Sims_G.

Definition 4.4 Let $r \in \mathbb{R}_+$ be the minimal length of orientation-paired marked segments in $[\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega)$. For any marked segment $\hat{s}$ belonging to a pair in the orbit $[\{\hat{s}, \hat{s}'\}]$, the *normalized marked segment* corresponding to $\hat{s}$ is $n(\hat{s}) := (1/r)\hat{s}$. Set $n(\hat{s})' := n(\hat{s}')$, and define the *normalized* $\text{Aff}_{\mathcal{O}}^+(X, \omega)$*-orbit* corresponding to $[\{\hat{s}, \hat{s}'\}]$ as

$$n([\{\hat{s}, \hat{s}'\}]) := \left\{\{n(\hat{s}), n(\hat{s})'\} \mid \{\hat{s}, \hat{s}'\} \in [\{\hat{s}, \hat{s}'\}]\right\}.$$

Thus the normalized $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbit corresponding to $[\{\hat{s}, \hat{s}'\}]$ takes all points of the orbit $[\{\hat{s}, \hat{s}'\}]$ and rescales them (by a constant scalar $1/r$) so that the minimal length of a pair in the orbit is one.

Theorem 4.9 *If (X, ω) is a fanning surface, then $\widehat{\mathcal{M}}(X, \omega)$ is a finite union of scaled simulations. In particular, if $\Gamma \leq \text{SL}_2\mathbb{R}$ is a fanning group generated by $G \subset \Gamma$, then*

$$\left\{n([\{\hat{s}, \hat{s}'\}]) \mid [\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega) \text{ for some } (X, \omega) \text{ with } \text{SL}(X, \omega) = \Gamma\right\} \subset \text{Sims}_G.$$

Proof We will prove the second statement; the first follows from this together with Lemma 4.5. Let (X, ω) be a translation surface with $\text{SL}(X, \omega) = \Gamma$, and let G be a (finite or

infinite) set of generators for Γ. By Lemma 4.4, there is some subset $H \subset G_{\mathcal{O}}$ which generates $\text{Aff}_{\mathcal{O}}^+(X, \omega)$. Now let $[\{\hat{s}, \hat{s}'\}] \in \widehat{\mathcal{M}}(X, \omega)/\text{Aff}_{\mathcal{O}}^+(X, \omega)$ with $\{\hat{s}, \hat{s}'\}$ a pair of minimal length in $[\{\hat{s}, \hat{s}'\}]$. Set $p := n(\hat{s})$ and $p^- := n(\hat{s})'$. By Corollary 4.6, $\boldsymbol{\pi}(p) = \boldsymbol{\pi}(\hat{s})/|\boldsymbol{\pi}(\hat{s})|$ and $\boldsymbol{\pi}(p^-) = \boldsymbol{\pi}(\hat{s}')/|\boldsymbol{\pi}(\hat{s}')|$ belong to the finite set Θ_Γ, so $\mathfrak{s} := (H, \{p, p^-\}) \in \mathfrak{S}_G$. Furthermore, $\text{sim}(\mathfrak{s})$ is the $\langle H \rangle = \text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbit of $\{p, p^-\} = \{n(\hat{s}), n(\hat{s})'\}$, and this is precisely the set $n([\{\hat{s}, \hat{s}'\}])$. By Proposition 3.8, $[\{\hat{s}, \hat{s}'\}]$—and thus also $n([\{\hat{s}, \hat{s}'\}])$—is distinctive; hence $n([\{\hat{s}, \hat{s}'\}]) \in \text{Sims}_G$.   □

Theorem 4.9 implies that for any fanning group Γ, the orientation-paired marked segments $\widehat{\mathcal{M}}(X, \omega)$ of any (X, ω) with Veech group Γ are given by a finite collection of simulations, each of which is appropriately scaled. Recall that the marked Voronoi staples $\widehat{S}(X, \omega)$—from which (X, ω) may be reconstructed—are determined by $\widehat{\mathcal{M}}(X, \omega)$ (Sect. 3.3). In Sect. 4.5, we show that for any finite collection of simulations, there are at most finitely many scalars for which the union of scaled simulations agrees with the marked segments of a unit-area translation surface. Together with Lemma 4.8, this will prove Theorem 1.2.

4.4 Marked Voronoi staples determine permissible triples

Let $(X, \omega) \in \mathcal{H}(d_1, \dots, d_\kappa)$ and $\mathcal{O} = \mathcal{O}(d_1, \dots, d_\kappa)$. In this subsection we apply the results of Sect. 3.2 to orientation-paired marked segments and their $\text{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits. In particular, we shall see that adjacent elements of $\widehat{S}_F(X, \omega)$ naturally determine permissible triples (Proposition 4.11). To this end, we begin with the following:

Lemma 4.10 *For any two adjacent $\hat{s}, \hat{t} \in \widehat{S}_F(X, \omega)$, we have*

$$B(\hat{s}, \hat{t}) \cap \widehat{\mathcal{M}}_F(X, \omega) = \varnothing.$$

Moreover, $[\hat{s}, \hat{t}] \in \widehat{\mathcal{M}}_F(X, \omega)$ with $[\hat{s}, \hat{t}]' = [\hat{t}, \hat{s}]$.

Proof Let i be the index for which $\hat{s}, \hat{t} \in \mathcal{O}_i$, and recall from Definition 3.3 that the circumcenter $c := c(\hat{s}, \hat{t})$ determined by $\hat{s}, \hat{t} \in \mathcal{O}_i$ is the intersection of $\partial H(\hat{s})$ and $\partial H(\hat{t})$. Since $\hat{s}, \hat{t} \in \widehat{S}_F(X, \omega)$ are essential points of Ω_i, there are two (adjacent) edges of Ω_i contained in $\partial H(\hat{s})$ and $\partial H(\hat{t})$, respectively. The intersection of these edges is precisely the circumcenter c, so $c \in \Omega_i$. Suppose for the sake of contradiction that there is some marked segment $\hat{u} \in B(\hat{s}, \hat{t})$. Then

$$d(\hat{u}, c) < |c| = d(0, c),$$

and hence $c \notin H(\hat{u})$. This contradicts the fact that $c \in \Omega_i$, so we have

$$B(\hat{s}, \hat{t}) \cap \widehat{\mathcal{M}}_F(X, \omega) = \varnothing.$$

This result—together with the convexity of $B(\hat{s}, \hat{t})$ and the fact that $0 \in \partial B(\hat{s}, \hat{t})$—gives that $B(\hat{s}, \hat{t}) \subset \text{star}_i(X, \omega)$, so η is defined on all of $B(\hat{s}, \hat{t})$. Let γ be the straight line segment in $\mathcal{O}_i$ from $\hat{s}$ to $\hat{t}$, and let $\sigma_j, \sigma_k \in \Sigma$ be the singularities at which the oriented saddle connections s and t end, respectively. Then, recycling notation, $u := \eta(\gamma)$ is a saddle connection from σ_j to σ_k, and $\hat{u} \in \widehat{\mathcal{M}}_F(X, \omega)$. We claim that $\hat{u} = [\hat{s}, \hat{t}]$.

By Proposition 3.10, $\hat{s}$ and $\hat{t}$ belong to the same open π-sector of $\mathcal{O}_i$. We also have by Proposition 3.8 that $\arg(\hat{s}) \neq \arg(\hat{t})$, so $[\hat{s}, \hat{t}]$ is in fact defined. It is clear from the definition of $\hat{u}$ that $\hat{s}'$ and $\hat{u}$ belong to the same open π-sector of the same component $\mathcal{O}_j$. Furthermore, $\boldsymbol{\pi}(\hat{u}) = \boldsymbol{\pi}(\hat{t}) - \boldsymbol{\pi}(\hat{s})$; see Fig. 9. By Definition 3.6, we have $\hat{u} = [\hat{s}, \hat{t}]$. Reversing the orientation of γ in the argument above gives that $\hat{u}' = [\hat{t}, \hat{s}]$.   □
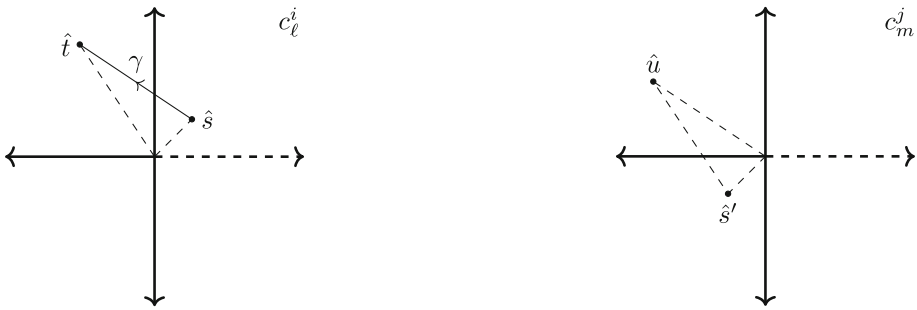
Fig. 9 Illustration of the proof of Lemma 4.10, showing that $\hat{u} = [\hat{s}, \hat{t}]$ is a marked segment for adjacent $\hat{s}, \hat{t} \in \widehat{\mathcal{S}}_F(X, \omega)$

Recall that permissible triples are defined in terms of three distinctive, limit-point free subsets of $\mathbb{P}(\mathcal{O})$ (Definition 3.7). We introduce notation for a set of—a priori non-permissible—triples determined by adjacent $\hat{s}, \hat{t} \in \widehat{\mathcal{S}}_F(X, \omega)$ and subsequently prove that these are in fact permissible triples belonging to the set $\mathcal{P}(\mathrm{Sim}_G)$, where $\mathrm{SL}(X, \omega) = \langle G \rangle$.

Definition 4.5 For $(X, \omega) \in \mathcal{H}(d_1, \dots, d_\kappa)$, let

$$\mathcal{P}(X, \omega) := \{(n(\hat{s}), n(\hat{t}), n([\hat{s}, \hat{t}])) \mid \hat{s}, \ \hat{t} \text{ are adjacent elements of } \widehat{\mathcal{S}}_F(X, \omega)\}.$$

Proposition 4.11 *For any fanning group* $\Gamma \leq \mathrm{SL}_2\mathbb{R}$ *generated by* $G \subset \Gamma$ *and any* $(X, \omega) \in \mathcal{H}(d_1, \dots, d_\kappa)$ *with Veech group* $\mathrm{SL}(X, \omega) = \Gamma$*, we have* $\mathcal{P}(X, \omega) \subset \mathcal{P}(\mathrm{Sim}_G)$*. Moreover, if* G *is finite (i.e.* Γ *is a lattice), then there exists some* $N \in \mathbb{N}$ *such that for any* $(X, \omega) \in \mathcal{H}(d_1, \dots, d_\kappa)$ *with Veech group* $\mathrm{SL}(X, \omega) = \Gamma$*, we have* $\mathcal{P}(X, \omega) \subset \mathcal{P}(\mathrm{Sim}_G) \subset \mathcal{P}(\mathrm{Sim}_G^N)$*.*

Proof Let $(n(\hat{s}), n(\hat{t}), n([\hat{s}, \hat{t}])) \in \mathcal{P}(X, \omega)$ for some (X, ω) with Veech group Γ. By Theorem 4.9, there are $\mathfrak{s}_{\hat{s}}, \mathfrak{s}_{\hat{t}}, \mathfrak{s}_{[\hat{s}, \hat{t}]} \in \mathfrak{S}_G$ for which

$$n([\{\hat{s}, \hat{s}'\}]) = \mathrm{sim}(\mathfrak{s}_{\hat{s}}), \quad n([\{\hat{t}, \hat{t}'\}]) = \mathrm{sim}(\mathfrak{s}_{\hat{t}}), \quad \text{and} \quad n([\{[\hat{s}, \hat{t}], [\hat{s}, \hat{t}]'\}]) = \mathrm{sim}(\mathfrak{s}_{[\hat{s}, \hat{t}]}).$$

Hence

$$(n(\hat{s}), n(\hat{t}), n([\hat{s}, \hat{t}])) \in \mathrm{sim}_F(\mathfrak{s}_{\hat{s}}) \times \mathrm{sim}_F(\mathfrak{s}_{\hat{t}}) \times \mathrm{sim}_F(\mathfrak{s}_{[\hat{s}, \hat{t}]}),$$

and we must show that $(n(\hat{s}), n(\hat{t}), n([\hat{s}, \hat{t}]))$ is a permissible triple. The corresponding permissible scalars we shall use are the minimal lengths $(r_{\hat{s}}, r_{\hat{t}}, r_{[\hat{s}, \hat{t}]}) \in \mathbb{R}_+^3$ of pairs in the respective $\mathrm{Aff}_{\mathcal{O}}^+(X, \omega)$-orbits $[\{\hat{s}, \hat{s}'\}]$, $[\{\hat{t}, \hat{t}'\}]$ and $[\{[\hat{s}, \hat{t}], [\hat{s}, \hat{t}]'\}]$. In particular, $r_{\hat{s}} n(\hat{s}) = \hat{s}$, $r_{\hat{t}} n(\hat{t}) = \hat{t}$ and $r_{[\hat{s}, \hat{t}]} n([\hat{s}, \hat{t}]) = [\hat{s}, \hat{t}]$. Moreover, $r_{\hat{s}} \mathrm{sim}(\mathfrak{s}_{\hat{s}}) = [\{\hat{s}, \hat{s}'\}]$, $r_{\hat{t}} \mathrm{sim}(\mathfrak{s}_{\hat{t}}) = [\{\hat{t}, \hat{t}'\}]$ and $r_{[\hat{s}, \hat{t}]} \mathrm{sim}(\mathfrak{s}_{[\hat{s}, \hat{t}]}) = [\{[\hat{s}, \hat{t}], [\hat{s}, \hat{t}]'\}]$. We now verify the two criteria of Definition 3.7 on the triple $(n(\hat{s}), n(\hat{t}), n([\hat{s}, \hat{t}]))$ and the scalars $(r_{\hat{s}}, r_{\hat{t}}, r_{[\hat{s}, \hat{t}]})$.

(i) From the observations above, $[r_{\hat{s}} n(\hat{s}), r_{\hat{t}} n(\hat{t})] = [\hat{s}, \hat{t}]$ and $[r_{\hat{t}} n(\hat{t}), r_{\hat{s}} n(\hat{s})] = [\hat{t}, \hat{s}]$. Since $\hat{s}$ and $\hat{t}$ are adjacent elements of $\widehat{\mathcal{S}}_F(X, \omega)$, Lemma 4.10 guarantees these are defined. Moreover,

$$[r_{\hat{s}} n(\hat{s}), r_{\hat{t}} n(\hat{t})] = [\hat{s}, \hat{t}] = r_{[\hat{s}, \hat{t}]} n([\hat{s}, \hat{t}])$$

and

$$[r_{\hat{t}} n(\hat{t}), r_{\hat{s}} n(\hat{s})] = [\hat{t}, \hat{s}] = [\hat{s}, \hat{t}]' = (r_{[\hat{s}, \hat{t}]} n([\hat{s}, \hat{t}]))' = r_{[\hat{s}, \hat{t}]} n([\hat{s}, \hat{t}])'.$$
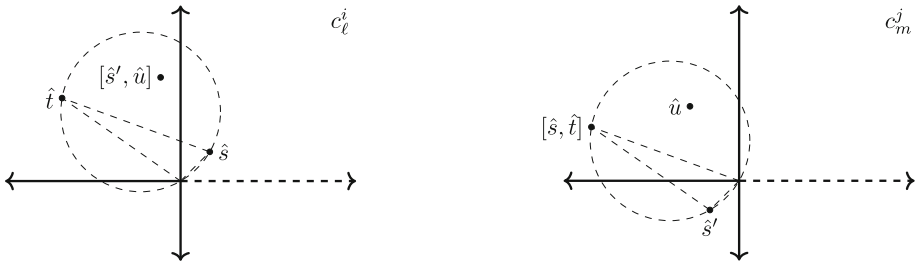
Fig. 10 Illustration of the proof of Proposition 4.11

(ii) We must show that

$$B(r_{\hat{s}}n(\hat{s}), r_{\hat{t}}n(\hat{t})) \cup B(r_{[\hat{s},\hat{t}]}n([\hat{s},\hat{t}]), r_{\hat{s}}n(\hat{s})') \cup B(r_{\hat{t}}n(\hat{t})', r_{[\hat{s},\hat{t}]}n([\hat{s},\hat{t}])')$$

does not intersect

$$r_{\hat{s}}\mathrm{sim}_F(\mathfrak{s}_{\hat{s}}) \cup r_{\hat{t}}\mathrm{sim}_F(\mathfrak{s}_{\hat{t}}) \cup r_{[\hat{s},\hat{t}]}\mathrm{sim}_F(\mathfrak{s}_{[\hat{s},\hat{t}]}),$$

or, equivalently,

$$\left(B(\hat{s},\hat{t}) \cup B([\hat{s},\hat{t}],\hat{s}') \cup B(\hat{t}',[\hat{s},\hat{t}]')\right) \cap \left([\{\hat{s},\hat{s}'\}]_F \cup [\{\hat{t},\hat{t}'\}]_F \cup [\{[\hat{s},\hat{t}],[\hat{s},\hat{t}]'\}]_F\right) = \varnothing.$$

Since each of the $\mathrm{Aff}_{\mathcal{O}}^+(X,\omega)$-orbits are contained in $\widehat{\mathcal{M}}(X,\omega)$, it suffices to show

$$\left(B(\hat{s},\hat{t}) \cup B([\hat{s},\hat{t}],\hat{s}') \cup B(\hat{t}',[\hat{s},\hat{t}]')\right) \cap \widehat{\mathcal{M}}_F(X,\omega) = \varnothing.$$

Lemma 4.10 gives that

$$B(\hat{s},\hat{t}) \cap \widehat{\mathcal{M}}_F(X,\omega) = \varnothing.$$

Now suppose for the sake of contradiction that

$$B([\hat{s},\hat{t}],\hat{s}') \cap \widehat{\mathcal{M}}_F(X,\omega) \neq \varnothing,$$

and let j be the index for which $\hat{s}' \in \mathcal{O}_j$. Since $\widehat{\mathcal{M}}_F(X,\omega)$ has no limit points, there are at most finitely many marked segments belonging to this intersection, and by Proposition 3.8, none of these marked segments have the same argument as $\hat{s}'$. Choose some $\hat{u}$ in the intersection with argument nearest that of $\hat{s}'$. Then the straight-line segment from $\hat{s}'$ to $\hat{u}$ belongs to $\mathrm{star}_j(X,\omega)$, and a similar argument to that in the proof of Lemma 4.10 gives that $[\hat{s}',\hat{u}] \in \widehat{\mathcal{M}}(X,\omega)$. By definition of $[\hat{s}',\hat{u}]$, this marked segment belongs to the same open π-sector as $\hat{s}$ and hence to $B(\hat{s},\hat{t})$, contradicting Lemma 4.10 (see Fig. 10). A similar argument gives

$$B(\hat{t}',[\hat{s},\hat{t}]') \cap \widehat{\mathcal{M}}_F(X,\omega) = \varnothing.$$

Thus $(n(\hat{s}), n(\hat{t}), n([\hat{s},\hat{t}]))$ is a permissible triple, and $\mathcal{P}(X,\omega) \subset \mathcal{P}(\mathrm{Sims}_G)$.

Now suppose that G is finite. Then the set Sims_G of all simulations is finite, and by Lemma 3.7, the set $\mathcal{P}(\mathrm{Sims}_G)$ is finite. Hence for each $(p,q,u) \in \mathcal{P}(\mathrm{Sims}_G)$ there is some n for which $(p,q,u) \in \mathcal{P}(\mathrm{Sims}_G^n)$. Taking N to be the maximum of these, we have $\mathcal{P}(X,\omega) \subset \mathcal{P}(\mathrm{Sims}_G) \subset \mathcal{P}(\mathrm{Sims}_G^N)$. $\qquad\square$
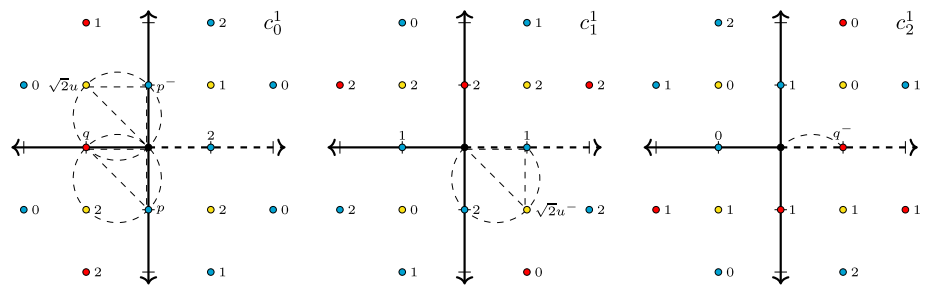
Fig. 11 Subsets of the scaled simulations $\text{sim}(\mathfrak{s}_0)$, $\text{sim}(\mathfrak{s}_1)$ and $\sqrt{2}\text{sim}(\mathfrak{s}_2)$, with $\mathfrak{s}_0$, $\mathfrak{s}_1$ $\mathfrak{s}_2 \in \mathfrak{S}_G$ as defined in Example 4.7, and balls showing that $(p, q, u) \in \mathcal{P}(\text{Sims}_G)$ is a permissible triple with permissible scalars $(1, 1, \sqrt{2})$

4.5 Finiteness of lattice surfaces with given Veech groups

Throughout this subsection, let G be a set of generators of a fanning group Γ. As noted at the end of Sect. 4.3, to prove Theorem 1.2 it suffices to show that for any finite collection of simulations determined by G, there are at most finitely many scalars for which the union of scaled simulations coincides with the marked segments of a unit-area translation surface. Propositions 3.5, 3.6 and 4.11 suggest that for adjacent elements $\hat{s}, \hat{t} \in \widehat{\mathcal{S}}_F(X, \omega)$ of a translation surface (X, ω) with Veech group Γ, the desired scalars for the simulations corresponding to $\hat{s}, \hat{t}$ and $[\hat{s}, \hat{t}]$ are uniquely determined by one another. Hence if the elements of $\widehat{\mathcal{S}}_F(X, \omega)$ are all—loosely speaking—'transitively adjacent' to one another (or their orientation-paired inverses), then an arbitrary choice of one scalar uniquely determines all other scalars.

We make these vague notions precise by introducing a *permissible triples graph* corresponding to a set P of permissible triples whose vertices are the distinctive, limit-point free sets corresponding to the triples and whose edges are defined between the vertices corresponding to entries of the same permissible triple in P. The 'transitive adjacency' alluded to above refers to the connectedness of such a graph.

Definition 4.6 Let $\mathbf{Q} = \{Q_i\}_{i \in I}$ be a collection of distinctive, limit-point free subsets of $\mathbb{P}(\mathcal{O})$. For a subset $P \subset \mathcal{P}(\mathbf{Q})$ of permissible triples, define the *permissible triples graph* determined by P to be the graph $\mathfrak{G}_P = (\mathfrak{V}_P, \mathfrak{E}_P)$ with vertex set

$$\mathfrak{V}_P = \{Q_i, Q_j, Q_k \mid \text{there exists some } (q_i, q_j, q_k) \in P \cap ((Q_i)_F \times (Q_j)_F \times (Q_k)_F)\}$$

consisting of the sets in $\mathbf{Q}$ corresponding to the permissible triples in P, and edge set

$$\mathfrak{E}_P = \{\{Q_\ell, Q_m\} \mid \ell, m \in \{i, j, k\} \text{ for some } (q_i, q_j, q_k) \in P \cap ((Q_i)_F \times (Q_j)_F \times (Q_k)_F)\}$$

where an edge connects vertices Q_ℓ and Q_m if there is some permissible triple in P to which both Q_ℓ and Q_m correspond. We call $(r_i)_{i \in I} \in \mathbb{R}_+^{|I|}$ *consistent scalars* for the graph $\mathfrak{G}_P$ if (r_i, r_j, r_k) are permissible scalars for each permissible triple $(q_i, q_j, q_k) \in P \cap ((Q_i)_F \times (Q_j)_F \times (Q_k)_F)$. We call the graph $\mathfrak{G}_P$ *consistent* if it admits consistent scalars.

Example 4.12 Consider again Example 4.7 and Fig. 8; let $p \in \text{sim}(\mathfrak{s}_0)$ be the point with argument $3\pi/2$, $q \in \text{sim}(\mathfrak{s}_1)$ the point with argument π and $u \in \text{sim}(\mathfrak{s}_2)$ the point with argument $3\pi/4$. Figure 11 shows that $(p, q, u) \in \mathcal{P}(\text{Sims}_G)$ is a permissible triple with permissible scalars $(1, 1, \sqrt{2})$. Let $P := \{(p, q, u)\}$. Then $\mathfrak{G}_P$ is the graph with vertex set $\mathfrak{V}_P = \{\text{sim}(\mathfrak{s}_0), \text{sim}(\mathfrak{s}_1), \text{sim}(\mathfrak{s}_2)\}$ and edge set

$$\mathfrak{E}_P = \{\{\text{sim}(\mathfrak{s}_0), \text{sim}(\mathfrak{s}_1)\}, \{\text{sim}(\mathfrak{s}_1), \text{sim}(\mathfrak{s}_2)\}, \{\text{sim}(\mathfrak{s}_2), \text{sim}(\mathfrak{s}_0)\}\},$$

and $(1, 1, \sqrt{2})$ are consistent scalars for $\mathfrak{G}_P$.

Proposition 4.13 *Let (X, ω) be a fanning surface with Veech group $\mathrm{SL}(X, \omega) = \Gamma$, and let $P := \mathcal{P}(X, \omega) \subset \mathcal{P}(\mathrm{Sims}_G)$. Then the permissible triples graph $\mathfrak{G}_P$ is connected and consistent, and if $(r_{\mathfrak{s}})_{sim(\mathfrak{s}) \in \mathfrak{V}_P}$ are consistent scalars for $\mathfrak{G}_P$, then any other consistent scalars are of the form $(ar_{\mathfrak{s}})_{sim(\mathfrak{s}) \in \mathfrak{V}_P}$ with $a \in \mathbb{R}_+$.*

Proof Theorem 4.9 implies that the vertex set $\mathfrak{V}_P$ is finite, say $\mathfrak{V}_P = \{sim(\mathfrak{s}_i)\}_{i=1}^n$. The same theorem also gives that

$$\bigcup_{i=1}^n r_i sim(\mathfrak{s}_i) \subset \widehat{\mathcal{M}}(X, \omega)$$

for some $(r_i)_{i=1}^n \in \mathbb{R}_+^n$. In fact, each r_i is the minimal length of a pair in the $\mathrm{Aff}_{\mathcal{O}}^+(X, \omega)$-orbit for which $r_i sim(\mathfrak{s}_i) = [\{\hat{s}, \hat{s}'\}]$. The proof of Proposition 4.11 shows that the scalars $(r_i)_{i=1}^n \in \mathbb{R}_+^n$ are consistent, so the graph $\mathfrak{G}_P$ is consistent.

Now suppose that $(r_i')_{i=1}^n \in \mathbb{R}_+^n$ are also consistent scalars for $\mathfrak{G}_P$, and recall the definition of the edge set $\mathfrak{E}_P$. Induction with Propositions 3.5 and 3.6 implies that any of the scalars r_i' corresponding to a simulation of a connected component of $\mathfrak{G}_P$ uniquely determines the remaining scalars r_j' of the simulations of the component. In particular, if $r_i' = ar_i$ for some $a \in \mathbb{R}_+$, then $r_j' = ar_j$ for each of these scalars. Hence it suffices to show that $\mathfrak{G}_P$ is connected.

Consider the subgraph (V, E) of $(\mathfrak{V}_P, \mathfrak{E}_P)$ with vertex set V consisting of all simulations in $\mathfrak{V}_P$ which contain a pair of normalized marked Voronoi staples and edge set E consisting of edges $\{sim(\mathfrak{s}_i), sim(\mathfrak{s}_j)\}$, where $sim(\mathfrak{s}_i)$ and $sim(\mathfrak{s}_j)$ contain $\{n(\hat{s}), n(\hat{s})'\}$ and $\{n(\hat{t}), n(\hat{t})'\}$, respectively, for adjacent $\hat{s}, \hat{t} \in \widehat{\mathcal{S}}_F(X, \omega)$. Recall that (X, ω) is isometric to the quotient space of $\bigsqcup_{i=1}^K \Omega_i$ under the equivalence relation given by identifying the two edges of convex bodies which correspond to a pair of marked Voronoi staples (Sect. 3.3 above, and also Sect. 3 of [7]). If (V, E) were disconnected, then after this identification of edges of convex bodies the resulting surface—and hence (X, ω)—would be disconnected. This is a contradiction, so the subgraph (V, E) is connected. By definition of $P = \mathcal{P}(X, \omega)$ and $\mathfrak{V}_P$, any vertex $sim(\mathfrak{s}_i) \in \mathfrak{V}_P \backslash V$ must contain $\{n([\hat{s}, \hat{t}]), n([\hat{s}, \hat{t}])'\}$ for some adjacent $\hat{s}, \hat{t} \in \widehat{\mathcal{S}}_F(X, \omega)$. By definition of $\mathfrak{E}_P$, there are edges of $\mathfrak{G}_P$ from $sim(\mathfrak{s}_i)$ to the simulations containing $\{n(\hat{s}), n(\hat{s})'\}$ and $\{n(\hat{t}), n(\hat{t})'\}$, both of which belong to $V \subset \mathfrak{V}_P$. Thus $\mathfrak{G}_P$ is connected, and the result follows. $\square$

We are now ready to prove Theorem 1.2:

Proof of Theorem 1.2 Let $\Gamma \leq \mathrm{SL}_2\mathbb{R}$ be a lattice. By Lemma 4.3, Γ is a finitely generated fanning group. Let G be a finite set of generators of Γ and fix $\mathcal{H}_1(d_1, \ldots, d_\kappa)$. We must show that the set of all $(X, \omega) \in \mathcal{H}_1(d_1, \ldots, d_\kappa)$ for which $\mathrm{SL}(X, \omega) = \Gamma$ is finite. By Proposition 4.11, each of these translation surfaces belongs to the set

$$\{(X, \omega) \in \mathcal{H}_1(d_1, \ldots, d_\kappa) \mid \mathcal{P}(X, \omega) \subset \mathcal{P}(\mathrm{Sims}_G)\},$$

so it suffices to show that this latter set is finite. By Lemmas 3.7 and 4.8, the set $\mathcal{P}(\mathrm{Sims}_G)$ of all permissible triples is finite, so it suffices to show for any $P \subset \mathcal{P}(\mathrm{Sims}_G)$ that the set

$$\{(X, \omega) \in \mathcal{H}_1(d_1, \ldots, d_\kappa) \mid \mathcal{P}(X, \omega) = P\}$$

is finite. Let $\mathfrak{V}_P = \{sim(\mathfrak{s}_i)\}_{i=1}^n$ be the vertex set of the permissible triples graph $\mathfrak{G}_P$. As in the proof of Proposition 4.13, for each (X, ω) with $\mathcal{P}(X, \omega) = P$, there are consistent

scalars $(r_i)_{i=1}^n \in \mathbb{R}_+^n$ for which

$$\bigcup_{i=1}^n r_i \operatorname{sim}(\mathfrak{s}_i) \subset \widehat{\mathcal{M}}(X, \omega).$$

By definition of $\mathfrak{G}_P$ and its vertex set $\mathfrak{V}_P$, we have that the set of marked Voronoi staples, $\widehat{\mathcal{S}}(X, \omega)$—from which (X, ω) may be recovered—is contained in this union of scaled simulations. Again from Proposition 4.13, if $(Y, \eta) \in \mathcal{H}_1(d_1, \dots, d_\kappa)$ also satisfies $\mathcal{P}(Y, \eta) = P$, then its corresponding consistent scalars are of the form $(a r_i)_{i=1}^n$ for some $a \in \mathbb{R}_+^n$. From Sect. 3.3, we have

$$\bigsqcup_{j=1}^\kappa \Omega_j(\widehat{\mathcal{S}}_F(X, \omega)) = \bigsqcup_{j=1}^\kappa \Omega_j \left(\bigcup_{i=1}^n r_i \operatorname{sim}_F(\mathfrak{s}_i) \right)$$

and

$$\bigsqcup_{j=1}^\kappa \Omega_j(\widehat{\mathcal{S}}_F(Y, \eta)) = \bigsqcup_{j=1}^\kappa \Omega_j \left(\bigcup_{i=1}^n a r_i \operatorname{sim}_F(\mathfrak{s}_i) \right) = a \left(\bigsqcup_{j=1}^\kappa \Omega_j \left(\bigcup_{i=1}^n r_i \operatorname{sim}_F(\mathfrak{s}_i) \right) \right).$$

Since (X, ω) and (Y, η) are isometric to the quotient spaces of these respective sets under an equivalence relation identifying edges, and since both translation surfaces have area one, we conclude that $a = 1$ and thus $(X, \omega) = (Y, \eta)$. $\qquad\square$

5 The algorithm

The constructive results of the previous sections suggest an algorithm (Algorithm 1.1) which returns all translation surfaces with a given lattice Veech group Γ in any given stratum $\mathcal{H}_1(d_1, \dots, d_\kappa)$. We highlight the main steps of this algorithm here, providing necessary references and further commentary in italics. Recall that G is a given finite set of generators of the lattice $\Gamma \le \mathrm{SL}_2\mathbb{R}$.

(1) Compute a finite set $\Theta_{\Gamma_m} \subset S^1$ containing Θ_Γ, and let

$$\mathfrak{S}_G := \{(H, \{p, p^-\}) \in 2^{G\mathcal{O}} \backslash \{\varnothing\} \times \mathbb{P}(\mathcal{O}) \mid \pi(p), \pi(p^-) \in \Theta_{\Gamma_m}\}.$$

   (*See Definition 4.3 and the remark following Lemma 4.8.*)
(2) For each $n \in \mathbb{N}$:

   (a) Compute the set $\mathrm{Sims}_G^n = \{\operatorname{sim}^n(\mathfrak{s}) \mid \mathfrak{s} \in \mathfrak{S}_G\}$ of all stage n distinctive simulations determined by G.
   (*See Defintion 4.3.*)
   (b) Compute the set $\mathcal{P}(\mathrm{Sims}_G^n)$ of all permissible triples arising from Sims_G^n along with their respective rays of permissible scalars.
   (*See Definition 3.7 and Proposition 3.5. Recall from Definition 4.5 and Proposition 4.11 that there is some $N \in \mathbb{N}$ so that for any $(Y, \eta) \in \mathcal{H}_1(d_1, \dots, d_\kappa)$ with $\mathrm{SL}(Y, \eta) = \Gamma$, the set $\mathcal{P}(Y, \eta)$ determined by the marked Voronoi staples $\widehat{\mathcal{S}}(Y, \eta)$ is contained in the finite set of permissible triples $\mathcal{P}(\mathrm{Sims}_G^N)$.*)
   (c) For each $P \subset \mathcal{P}(\mathrm{Sims}_G^n) \backslash \{\varnothing\}$:
   (*For $n = N$ and any (Y, η) as above, this loop will set $P = \mathcal{P}(Y, \eta)$ for some P.*)
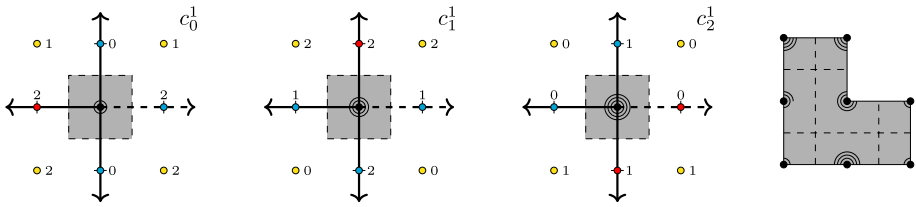
Fig. 12 The convex body and resulting translation surface (X, ω) (with opposite sides identified) using the scaled simulations $\text{sim}(\mathfrak{s}_0)$, $\text{sim}(\mathfrak{s}_1)$ and $\sqrt{2}\text{sim}(\mathfrak{s}_2)$ from Examples 4.7 and 4.12. The arcs about the singularity of $\mathcal{O}$ and of (X, ω) indicate arguments of generalized polar coordinates

(i) Construct the permissible triples graph $\mathfrak{G}_P = (\mathfrak{V}_P, \mathfrak{E}_P)$. If $\mathfrak{G}_P$ is connected and consistent with consistent scalars $(r_\mathfrak{s})_{\text{sim}^n(\mathfrak{s}) \in \mathfrak{V}_P^n}$, then set

$$S := \bigcup_{\text{sim}^n(\mathfrak{s}) \in \mathfrak{V}_P} r_\mathfrak{s} \text{sim}^n(\mathfrak{s}),$$

construct the convex bodies $\Omega_i(S_F)$ and determine the corresponding essential points $\mathcal{E}_i(S_F)$ for each $i \in \{1, \ldots, \kappa\}$.
(*See Definition 4.6 and Sect. 3.3.*)

(ii) If each $\Omega_i(S_F)$ is compact, the essential points $\sqcup_{i=1}^\kappa \mathcal{E}_i(S_F)$ come in pairs $\{p, p^-\}$ determined by the simulations $\text{sim}^n(\mathfrak{s}) \in \mathfrak{V}_P$, and the edges of the convex bodies determined by p and p^- are equal length, then construct the translation surface (X, ω) by identifying these edges via translation.
(*See Sect. 3.3. When $P = \mathcal{P}(Y, \eta)$, the proof of Proposition 4.13 shows that for some $a \in \mathbb{R}_+$, the set aS contains the set of marked Voronoi staples $\widehat{\mathcal{S}}(Y, \eta)$, and thus the union of scaled essential points $\sqcup_{i=1}^\kappa a\mathcal{E}_i(S_F)$ is precisely $\widehat{\mathcal{S}}_F(Y, \eta)$. Hence (X, ω) is a scaled version of (Y, η).*)

(iii) Rescale (X, ω) if necessary so that it has area one, and verify that $\text{SL}(X, \omega) = \Gamma$.
(*Computing $\text{SL}(X, \omega)$ is done using one of the algorithms mentioned in Sect. 1.*)

Example 5.1 Examples 4.7 and 4.12 illustrate (parts of) Steps 1, 2.a, 2.b and 2.c.i of the algorithm. Figure 12 illustrates the construction of convex bodies and corresponding essential points in Step 2.c.i as well as the resulting translation surface (X, ω) constructed in Step 2.c.ii. Notice that the essential points come only from $\text{sim}(\mathfrak{s}_0)$ and $\text{sim}(\mathfrak{s}_1)$; however, $\text{sim}(\mathfrak{s}_2)$ was necessary to determine the permissible scalars $(1, 1, \sqrt{2})$ (Example 4.12). Rescaling to unit-area, one verifies that indeed $\text{SL}(X, \omega) = \langle S, T^2 \rangle$.

A number of improvements to Algorithm 1.1 can be made for computational efficiency. We mention three in particular:

- It is optimal to choose a minimal set of generators G for Γ, as the size of $\mathfrak{G}_G$—and hence of Sims_G—depends on $|G|$.
- Recall from Proposition 3.8 that if $\hat{s}, \hat{t} \in \widehat{\mathcal{M}}_F(X, \omega)$ are two marked segments with identical arguments in the same component $\mathcal{O}_i \subset \mathcal{O}$, then in fact $\{\hat{s}, \hat{s}'\} = \{\hat{t}, \hat{t}'\}$. This requirement restricts the permissible triples graphs that one needs to consider in step 2.c.
- In Step 2.b, one need only consider permissible triples $(p, q, u) \in \text{sim}_F^n(\mathfrak{r}) \times \text{sim}_F^n(\mathfrak{s}) \times \text{sim}_F^n(\mathfrak{t})$ where $\mathfrak{r}$, $\mathfrak{s}$ and $\mathfrak{t}$ are defined using the same subset $H \subset G_\mathcal{O}$. A similar restriction applies to the subsets $P \subset \mathcal{P}(\text{Sims}_G^n)$ considered in Step 2.c.

While Algorithm 1.1 does produce the finite set of lattice surfaces asserted by Theorem 1.2 in finite time, it does not give a halting criterion to determine when the set of

returned translation surfaces is exhaustive. An explicit test to determine the $N \in \mathbb{N}$ for which $\mathcal{P}(\mathrm{Sims}_G) \subset \mathcal{P}(\mathrm{Sims}_G^N)$ would give such a halting criterion (see Proposition 4.11). In particular, Algorithm 1.1 together with this halting criterion would give a general procedure to determine whether or not any given lattice group is realized as a Veech group in any given stratum. However, even without this criterion, the ideas of this paper may still be used in special cases to show that certain lattice groups are not realized in certain strata; see Sect. 6.

As the set of all strata is countable, Algorithm 1.1—if allowed to run indefinitely—is easily adapted to produce all unit-area translation surfaces with a given lattice Veech group in *any* stratum. A full implementation could also provide experimental evidence for lattices which are never realized as Veech groups, though—even with the aforementioned stratum-wise halting criterion—a criterion to definitively assert that a lattice is not a Veech group in any stratum seems to the author to be elusive.

We also mention that as the set of generators for any fanning group is at most countable, Algorithm 1.1 may be adapted to construct surfaces (X, ω) with given fanning Veech groups in given strata. However, for countably infinite generating sets G, the set $\mathfrak{S}_G$ is uncountable, so a construction of such surfaces may not be exhaustive.

6 The modular group in minimal strata

Here we give an example showing how the ideas of the previous sections may be used—in certain cases—to give obstructions for lattices being realized as Veech groups in strata.[3] In particular, we show that the square torus is the only translation surface with Veech group $\mathrm{SL}_2\mathbb{Z}$ in the collection of all minimal strata $\mathcal{H}(2g - 2)$, $g > 0$. Throughout this section, fix $g > 0$ and $(X, \omega) \in \mathcal{H}(2g - 2)$. Let $\sigma := \sigma_1$ be the sole singularity of (X, ω), $\Omega := \Omega_1$ the convex body subordinate to $\widehat{\mathcal{M}}_F(X, \omega)$ and $\mathcal{O} = \mathcal{O}(2g - 2)$ the canonical surface associated to $\mathcal{H}(2g - 2)$. We begin with two lemmas.

Lemma 6.1 *Let* $\widehat{\mathcal{S}}(X, \omega) = \{\{\hat{s}_i, \hat{s}_i'\}\}_{i=1}^n$ *be the marked Voronoi staples of* (X, ω). *If the counterclockwise angle measured from* $\hat{s}_i$ *to* $\hat{s}_i'$ *equals that measured from* $\hat{s}_i'$ *to* $\hat{s}_i$ *for each* i, *then* $n = 2g$ *for* n *even and* $n = 2g + 1$ *for* n *odd.*

Proof Since $\widehat{\mathcal{S}}_F(X, \omega)$ is the set of all essential points of Ω, the convex body Ω has $2n$ edges. Decompose Ω into $2n$ closed triangular regions with disjoint interiors, where each triangle has one vertex at the origin $0 \in \mathcal{O}$, and remaining two vertices at the endpoints of an edge of Ω. The sum of the interior angles of all triangles in this decomposition is $2n\pi$. The origin $0 \in \mathcal{O}$ is a singularity of cone angle $2(2g - 1)\pi$, so the sum of the interior angles of Ω is

$$2n\pi - 2(2g - 1)\pi = 2(n - 2g + 1)\pi.$$

Recall that (X, ω) is isometric to the quotient space of Ω under the equivalence relation given by identifying edges determined by the marked Voronoi staples $\widehat{\mathcal{S}}(X, \omega)$, and the vertices of Ω under this identification correspond to Voronoi 0-cells of (X, ω) (Sect. 3.3). In particular, the sum of interior angles of Ω which are identified as a single point must equal 2π. Our assumptions on the counterclockwise angles between $\hat{s}_i$ and $\hat{s}_i'$ and between $\hat{s}_i'$ and $\hat{s}_i$ imply that for n even, the vertices of Ω become a single point under the identification, while for n odd, the vertices of Ω become two distinct points under the identification (the argument

---

[3] Due to time constraints, we have only considered the group $\mathrm{SL}_2\mathbb{Z}$ in minimal strata. However, the author suspects that similar methods can be applied more generally to obtain further restrictions for different groups and strata.

here is analogous to that which shows the translation surface constructed by identifying opposite edges of a regular $2m$-gon has one singularity when m is even and two singularities when m is odd). In particular, for n even,

$$2(n - 2g + 1)\pi = 2\pi$$

implies $n = 2g$, while for n odd,

$$2(n - 2g + 1)\pi = 4\pi$$

implies $n = 2g + 1$. ☐

Lemma 6.2 *Suppose that the minimal length of a saddle connection on (X, ω) is one. If the set*

$$P := \{p \in \mathcal{O} \mid |p| = 1, \ arg(\boldsymbol{\pi}(p)) \in \{0, \pi/2, \pi, 3\pi/2\}\}$$

belongs to the set of marked segments $\widehat{\mathcal{M}}_F(X, \omega)$, then in fact $P = \widehat{\mathcal{S}}_F(X, \omega)$.

Proof It suffices to show that P is the set of essential points of the convex body Ω, i.e. that (i) $\Omega = \cap_{p \in P} H(p)$ and (ii) for any $q \in P$, $\Omega \subsetneq \cap_{p \in P \setminus \{q\}} H(p)$. Condition (ii) is immediate from the fact that Ω is compact while $\cap_{p \in P \setminus \{q\}} H(p)$ is non-compact for any $q \in P$. Now

$$\Omega = \bigcap_{\hat{s} \in \widehat{\mathcal{M}}_F(X,\omega)} H(\hat{s}) = \left(\bigcap_{p \in P} H(p) \right) \cap \left(\bigcap_{\hat{s} \in \widehat{\mathcal{M}}_F(X,\omega) \setminus P} H(\hat{s}) \right),$$

so it suffices to show for each marked segment $\hat{s} \in \widehat{\mathcal{M}}_F(X, \omega) \setminus P$, that $\cap_{p \in P} H(p) \subset H(\hat{s})$.

We claim that for each $p \in P$, the ball centered at p of radius one contains no marked segments other than p, i.e.

$$B_1(p) \cap \widehat{\mathcal{M}}_F(X, \omega) \setminus \{p\} = \varnothing.$$

Suppose on the contrary that the intersection is non-empty. As in the proof of Proposition 4.11, we can choose some $\hat{u}$ in the intersection with argument nearest that of p and find that $[p, \hat{u}]$ is a marked segment of length strictly less than one. Its corresponding saddle connection also has length strictly less than one, contrary to our assumptions, so the claim holds.

Next, we claim that the ball centered at the origin of radius $\sqrt{2}$ is contained in the union of balls of radius one centered at the various $p \in P$,

$$B_{\sqrt{2}}(0) \subset \bigcup_{p \in P} B_1(p).$$

This is clear from elementary Euclidean geometry, restricting to each 2π-sector c_j^1; see Fig. 13. In particular, the previous two claims imply that

$$B_{\sqrt{2}}(0) \cap \left(\widehat{\mathcal{M}}_F(X, \omega) \setminus P \right) = \varnothing. \tag{4}$$

Notice that the maximum distance from the origin to any point in $\cap_{p \in P} H(p)$ is $\sqrt{2}/2$ (this distance being realized at the vertices of this intersection). Let $\hat{s} \in \widehat{\mathcal{M}}_F(X, \omega) \setminus P$ and let q belong to the complement of $H(\hat{s})$. Certainly $|q| > |\hat{s}|/2$. Equation 4 implies $|\hat{s}| \geq \sqrt{2}$, so $|q| > \sqrt{2}/2$. Hence $q \notin \cap_{p \in P} H(p)$, so the complement of $H(\hat{s})$ does not intersect $\cap_{p \in P} H(p)$. Thus $\cap_{p \in P} H(p) \subset H(\hat{s})$ as desired. ☐

With these two lemmas, we now prove the main result of this section.

Fig. 13 The ball $B_{\sqrt{2}}(0)$ is contained in the union of balls $B_1(p)$, $p \in P$. Note that this figure shows only a portion of the ball $B_{\sqrt{2}}(0)$, as this ball is centered at the singularity $0 \in \mathcal{O}$. Similarly, only the upper half of $B_1(p_0)$ is shown
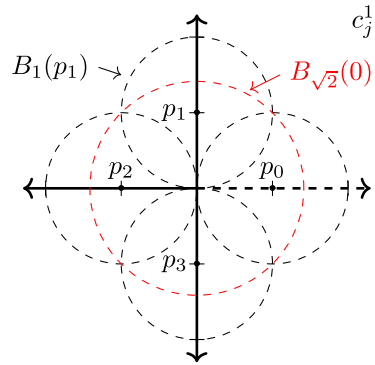
Fig. 14 Labelling of the points $p_0, \ldots, p_{16g-9}$ of $P \cup P'$ for $j \in \mathbb{Z}_{2g-1}$
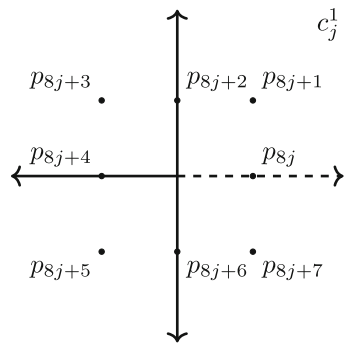
Proof of Theorem 1.3 Let $S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ and $T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, fix $\tau_S, \tau_T \in \mathrm{Trans}(\mathcal{O})$, and set $g_S = \tau_S \circ f_S$ and $g_T = \tau_T \circ f_T$ where $f_S, f_T \in \mathrm{Aff}_{\mathbb{C}}^+(\mathcal{O})$. We first show that all elements of

$$P := \{p \in \mathcal{O} \mid |p| = 1, \ \arg(\pi(p)) \in \{0, \pi/2, \pi, 3\pi/2\}\}$$

belong to the same $\langle g_S, g_T \rangle$-orbit. Note that $\mathrm{Trans}(\mathcal{O}) = \langle \rho_1 \rangle$; let $n \in \mathbb{Z}_{2g-1}$ be such that $\tau_T = \rho_1^n$. Set

$$P' := \{p \in \mathcal{O} \mid |p| = \sqrt{2}, \ \arg(\pi(p)) \in \{\pi/4, 3\pi/4, 5\pi/4, 7\pi/4\}\}.$$

Beginning with $p_0 \in P \cap c_0^1$, $\arg(p_0) = 0$ and sweeping counterclockwise, label the elements of $P \cup P'$ as $p_0, p_1, \ldots, p_{16g-9}$; see Fig. 14. For each $k \in \mathbb{Z}_{16g-8}$, we have

$$f_S(p_k) = p_{k+2},$$

and for $k \not\equiv 1 \pmod 4$,

$$f_T(p_k) = \begin{cases} p_k & \text{if } k \equiv 0 \pmod 4 \\ p_{k-1} & \text{if } k \equiv 2, 3 \pmod 4 \end{cases}.$$

Moreover, for each $k \in \mathbb{Z}_{16g-8}$,

$$\tau_T(p_k) = p_{k+8n}.$$

Set $f := g_T^{-1} \circ g_S^{-1} \circ g_T^{-1} \circ g_S$ and $h := g_S^{-1} \circ g_T^{-1} \circ g_S \circ g_T^{-1}$. Using the relations above—together with Lemma 3.3 and the fact that $\mathrm{Trans}(\mathcal{O}) = \langle \rho_1 \rangle$ is abelian—we find for each

$k \in \mathbb{Z}_{16g-8}$ that

$$
\begin{aligned}
f(p_{4k}) &= g_T^{-1} \circ g_S^{-1} \circ g_T^{-1} \circ g_S(p_{4k}) \\
&= (f_T^{-1} \circ \tau_T^{-1}) \circ (f_S^{-1} \circ \tau_S^{-1}) \circ (f_T^{-1} \circ \tau_T^{-1}) \circ (\tau_S \circ f_S)(p_{4k}) \\
&= \tau_T^{-2} \circ f_T^{-1} \circ f_S^{-1} \circ f_T^{-1} \circ f_S(p_{4k}) \\
&= \tau_T^{-2} \circ f_T^{-1} \circ f_S^{-1} \circ f_T^{-1}(p_{4k+2}) \\
&= \tau_T^{-2} \circ f_T^{-1} \circ f_S^{-1}(p_{4k+3}) \\
&= \tau_T^{-2} \circ f_T^{-1}(p_{4k+1}) \\
&= \tau_T^{-2}(p_{4k+2}) \\
&= p_{4k-16n+2}
\end{aligned}
$$

and

$$
\begin{aligned}
h(p_{4k+2}) &= g_S^{-1} \circ g_T^{-1} \circ g_S \circ g_T^{-1}(p_{4k+2}) \\
&= (f_S^{-1} \circ \tau_S^{-1}) \circ (f_T^{-1} \circ \tau_T^{-1}) \circ (\tau_S \circ f_S) \circ (f_T^{-1} \circ \tau_T^{-1})(p_{4k+2}) \\
&= \tau_T^{-2} \circ f_S^{-1} \circ f_T^{-1} \circ f_S \circ f_T^{-1}(p_{4k+2}) \\
&= \tau_T^{-2} \circ f_S^{-1} \circ f_T^{-1} \circ f_S(p_{4k+3}) \\
&= \tau_T^{-2} \circ f_S^{-1} \circ f_T^{-1}(p_{4k+5}) \\
&= \tau_T^{-2} \circ f_S^{-1}(p_{4k+6}) \\
&= \tau_T^{-2}(p_{4k+4}) \\
&= p_{4k-16n+4}.
\end{aligned}
$$

Then

$$
\begin{aligned}
g_T^8 \circ h \circ f \circ h \circ f(p_0) &= g_T^8 \circ h \circ f \circ h(p_{-16n+2}) \\
&= g_T^8 \circ h \circ f(p_{-32n+4}) \\
&= g_T^8 \circ h(p_{-48n+6}) \\
&= g_T^8(p_{-64n+8}) \\
&= \tau_T^8 \circ f_T^8(p_{-64n+8}) \\
&= \tau_T^8(p_{-64n+8}) \\
&= p_8.
\end{aligned}
$$

Iterating this latter composition shows that each p_{8k} belongs to the same $\langle g_S, g_T \rangle$-orbit. Injectivity of g_S implies that the union of the images under g_S, g_S^2 and g_S^3 of these p_{8k} is all of P, as claimed.

Now suppose $(X, \omega) \in \mathcal{H}(2g - 2)$ with $\mathrm{SL}(X, \omega) = \mathrm{SL}_2\mathbb{Z} = \langle S, T \rangle$ for some $g > 0$. Normalize (X, ω) if necessary so that its shortest saddle connection has unit length. One computes

$$
S^1 \backslash (T \cdot \Delta \cup T^{-1} \cdot \Delta \cup ST \cdot \Delta \cup ST^{-1} \cdot \Delta) = \{(\pm 1, 0), (0, \pm 1)\}.
$$

This set contains $\Theta_{\mathrm{SL}_2\mathbb{Z}} = \Theta_{\mathrm{SL}(X,\omega)}$, so Corollary 4.6 implies that the shortest saddle connection of (X, ω) is horizontal or vertical. Since the orthogonal matrix S belongs to $\mathrm{SL}(X, \omega)$ and any affine automorphism of (X, ω) sends saddle connections to saddle connections, (X, ω) has both horizontal and vertical saddle connections of minimal length one.

Recycling notation, let $f \in \mathrm{Aff}^+(\mathcal{O})$ and $\{p, p^-\} \in \mathbb{P}(\mathcal{O})$. Lemma 3.2 together with the fact that f is a self-homeomorphism of $\mathcal{O} = \mathcal{O}(2g - 2)$ implies that the counterclockwise angle measured from p to p^- equals that measured from $f(p)$ to $f(p^-)$. In particular, this is true of any $f \in \mathrm{Aff}_{\mathcal{O}}^+(X, \omega)$ and horizontal, unit-length $\{\hat{s}, \hat{s}'\} \in \widehat{\mathcal{M}}(X, \omega)$. Since $S, T \in \mathrm{SL}(X, \omega)$, there exist $g_S, g_T \in \mathrm{Aff}_{\mathcal{O}}^+(X, \omega)$ with $\mathrm{der}(g_S) = S$, $\mathrm{der}(g_T) = T$ (Sect. 4.2). The argument above shows that the set P belongs to the $\langle g_S, g_T \rangle$-orbit of $\hat{s}$. Since this orbit belongs to $\widehat{\mathcal{M}}_F(X, \omega)$, Lemma 6.2 implies that $P = \widehat{\mathcal{S}}_F(X, \omega)$. Moreover, $\hat{s}'$ belongs to this orbit, so the observation above implies that the counterclockwise angle measured from $\hat{s}$ to $\hat{s}'$ equals that measured from $\hat{s}'$ to $\hat{s}$, and the same is true of any pair of marked Voronoi staples in $\widehat{\mathcal{S}}(X, \omega)$. Since

$$|\widehat{\mathcal{S}}(X, \omega)| = |\widehat{\mathcal{S}}_F(X, \omega)|/2 = |P|/2 = 2(2g - 1)$$

is even, Lemma 6.1 implies that $2(2g - 1) = 2g$, which is only true for $g = 1$. When $g = 1$, the convex body Ω is the unit-square centered at $0 \in \mathcal{O} = (\mathbb{C}, dz)$, and the reconstruction of (X, ω) from Ω gives the square torus. $\square$

Declarations

References

1. Beardon, A.F.: The Geometry of Discrete Groups. Springer, New York (1983)
2. Bouw, I.I., Möller, M.: Teichmüller curves, triangle groups, and Lyapunov exponents. Ann. Math. **2**(1), 139–185 (2010)
3. Bowman, J.P.: Teichmüller geodesics, Delaunay triangulations, and Veech groups. Teichmüller Theory and Moduli Problems, Ramanujan Math. Society Lecture Notes Series, Vol 10, pp. 113–129 (2010)
4. Broughton, S.A., Judge, C.: Ellipses in translation surfaces. Geom. Dedicata. **157**(1), 111–151 (2012)

5. Davis, D., Pasquinelli, I., Ulcigrai, C.: Cutting sequences on Bouw–Möller surfaces: an $\mathcal{S}$-adic characterization. Ann. Sci. de l'Éc. Norm. Supér. **52**(4), 927–1023 (2019)

6. Edwards, B.: A new algorithm for computing the Veech group of a translation surface. Oregon State University, PhD Dissertation (2017)

7. Edwards, B., Sanderson, S., Schmidt, T. A.: Canonical translation surfaces for computing Veech groups. Geom. Dedicata, Vol. 216 No. 5 Paper No. 60 (2022)

8. Eskin, A., Mirzakhani, M., Mohammadi, A.: Isolation, equidistribution, and orbit closures for the SL $(2,\mathbb{R})$ action on moduli space. Ann. Math. **182**(2), 673–721 (2015)

9. Freidinger, M.: Stabilisatorgruppen in Aut(F_2) und Veechgruppen von Überlagerungen. Diplome Thesis, Universität Karlsruhe (2008)

10. Gutkin, E., Judge, C.: Affine mappings on translation surfaces: geometry and arithmetic. Duke Math. J. **103**(2), 191–213 (2000)

11. Hooper, W.P.: Grid graphs and lattice surfaces. Int. Math. Res. Not. IMRN **12**, 2657–2698 (2013)

12. Hubert, P., Lanneau, E.: Veech groups without parabolic elements. Duke Math. J. **133**(2), 335–346 (2006)

13. Hubert, P., Masur, H., Schmidt, T. A., Zorich, A.: Problems on billiards, flat surfaces and translation surfaces. In Farb B (ed)., Proceedings of symposium pure mathematical problems on mapping class groups and related topics vol 74, AMS, pp 233–243 (2006)

14. Hubert, P., Schmidt, A.T.: Infinitely generated Veech groups. Duke Math. J. **123**(1), 49–69 (2004)

15. Katok, S.: Fuchsian Groups. University of Chicago Press (1992)

16. Kenyon, R., Smillie, J.: Billiards on rational-angled triangles. Comment. Math. Helv. **75**(1), 65–108 (2000)

17. Masur, H., Smillie, J.: Hausdorff dimension of sets of nonergodic measured foliations. Ann. Math. **134**(3), 455–543 (1991)

18. McMullen, C.: Teichmüller geodesics of infinite complexity. Acta Math. **191**(2), 191–223 (2003)

19. McMullen, C.: Rigidity of Teichmüller curves. Math. Res. Lett. **16**(4), 647–649 (2009)

20. Mukamel, R.: Fundamental domains and generators for lattice Veech groups. Comment. Math. Helv. **92**(1), 57–83 (2017)

21. Ratner, M.: Raghunathan's topological conjecture and distributions of unipotent flows. Duke Math. J. **63**(1), 235–280 (1991)

22. The Sage Developers. SageMath, the Sage Mathematics Software System (Version 9.4). 2021, https://www.sagemath.org

23. Schmithüsen, G.: An algorithm for finding the Veech group of an origami. Exper. Math. **13**(4), 459–472 (2004)

24. Smillie, J., Weiss, B.: Minimal sets for flows on moduli space. Israel J. Math. **142**(1), 249–260 (2004)

25. Smillie, J., Weiss, B.: Characterizations of lattice surfaces. Invent. Math. **180**(3), 535–557 (2010)

26. Smillie, J., Weiss, B.: Finiteness results for flat surfaces: large cusps and short geodesics. Comment. Math. Helv. **85**(2), 313–336 (2010)

27. Veech, W.A.: Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards. Invent. Math. **97**, 553–583 (1989)

28. Veech, W.A.: Geometric realizations of hyperelliptic curves. Algorithms, fractals, and dynamics (Okayama/Kyoto, 1992), pp. 217–226 (1995)

29. Veech, W.A.: Bicuspid F-structures and Hecke groups. Proc. London Math. Soc. **103**(4), 710–745 (2011)

30. Vorobets, Y.B.: Planar structures and billiards in rational polygons: the Veech alternative. Russian Math. Surv. **51**(5), 779–817 (1996)

31. Wright, A.: Translation surfaces and their orbit closures: an introduction for a broad audience. EMS Surv. Math. Sci. **2**(1), 63–108 (2015)

32. Zorich, A.: Flat Surfaces. Frontiers in number theory, physics, and geometry Vol. 1, pp. 439–586 (2006)