

Computing fundamental domains for Fuchsian groups

par JOHN VOIGHT

RÉSUMÉ. Nous exposons un algorithme pour calculer un domaine de Dirichlet pour un Fuchsian groupe Γ avec aire cofinis. En conséquence, nous calculons les invariants de Γ et une présentation explicite finis pour Γ .

ABSTRACT. We exhibit an algorithm to compute a Dirichlet domain for a Fuchsian group Γ with cofinite area. As a consequence, we compute the invariants of Γ , including an explicit finite presentation for Γ .

Let $\Gamma \subset \mathrm{PSL}_2(\mathbb{R})$ be a *Fuchsian group*, a discrete group of orientation-preserving isometries of the upper half-plane \mathfrak{H} with hyperbolic metric d . A *fundamental domain* for Γ is a closed domain $D \subset \mathfrak{H}$ such that:

- (i) $\Gamma D = \mathfrak{H}$, and
- (ii) $gD^\circ \cap D^\circ = \emptyset$ for all $g \in \Gamma \setminus \{1\}$, where $^\circ$ denotes the interior.

Assume further that Γ has cofinite area, i.e., the coset space $X = \Gamma \backslash \mathfrak{H}$ has finite hyperbolic area $\mu(X) < \infty$; then it follows that Γ is finitely generated.

In this article, we exhibit an algorithm to compute a fundamental domain for Γ ; we assume that Γ is specified by a finite set of generators $G \subset \mathrm{SL}_2(K)$ with $K \hookrightarrow \mathbb{R} \cap \overline{\mathbb{Q}}$ a number field, and we call Γ *exact*. Suppose that $p \in \mathfrak{H}$ has trivial stabilizer $\Gamma_p = \{1\}$. Then the set

$$D(p) = \{z \in \mathfrak{H} : d(z, p) \leq d(gz, p) \text{ for all } g \in \Gamma\},$$

known as a *Dirichlet domain*, is a hyperbolically convex fundamental domain for Γ . The boundary of $D(p)$ consists of finitely many geodesic segments or *sides*. We specify $D(p)$ by a sequence of vertices, oriented counterclockwise around p . The domain $D(p)$ has a natural *side pairing*: For each side s of $D(p)$, there exists a unique side s^* and $g \in \Gamma \setminus \{1\}$ such that $s^* = gs$, and the set of such g comprises a set of generators for Γ .

Our main theorem is as follows.

Theorem. *There exists an algorithm which, given an exact Fuchsian group Γ with cofinite area and a point $p \in \mathfrak{H}$ with $\Gamma_p = \{1\}$, returns the Dirichlet domain $D(p)$, a side pairing for $D(p)$, and a finite presentation for Γ with a minimal set of generators.*

This algorithm also provides a solution to the word problem for the computed presentation of Γ .

Of particular and relevant interest is the class of *arithmetic Fuchsian groups*, those groups commensurable with the group of units \mathcal{O}_1^* of reduced norm 1 in a maximal order \mathcal{O} of a quaternion algebra B defined over a totally real field and split at exactly one real place. Alsina-Bayer [1] and Kohel-Verrill [18] give several examples of fundamental domains for arithmetic Fuchsian groups with $F = \mathbb{Q}$. Our work generalizes that of Johansson [15], who first made use of a Dirichlet domain for algorithmic purposes: he restricts to the case of arithmetic Fuchsian groups, and we improve on his methods in several respects (see the discussion preceding Algorithm 2.5 and the reduction algorithms in §4).

The algorithm described in the above theorem has the following applications. The first is a noncommutative generalization of the problem of computing generators for the unit group of a number field.

Corollary. *There exists an algorithm which, given an order $\mathcal{O} \subset B$ of a quaternion algebra B defined over a totally real field and split at exactly one real place, returns a finite presentation for \mathcal{O}_1^* with a minimal set of generators.*

We may also use the presentation for Γ to compute invariants. The group Γ has finitely many orbits with nontrivial stabilizer, known as *elliptic cycles* or *parabolic cycles* according as the stabilizer is finite or infinite. The coset space $X = \Gamma \backslash \mathfrak{H}$ can be given the structure of a Riemann surface, and we say that Γ has *signature* $(g; m_1, \dots, m_t; s)$ if X has genus g and Γ has exactly t elliptic cycles of orders $m_1, \dots, m_t \in \mathbb{Z}_{\geq 2}$ and s parabolic cycles.

Corollary. *There exists an algorithm which, given Γ , returns the signature of Γ and a set of representatives for the elliptic and parabolic cycles in Γ .*

Finally, we mention a corollary which is useful for the evaluation of automorphic forms.

Corollary. *There exists an algorithm which, given Γ and $z, p \in \mathfrak{H}$ with $\Gamma_p = \{1\}$, returns a point $z' \in D(p)$ and $g \in \Gamma$ such that $z' = g(z)$.*

The article is organized as follows. We begin by fixing notation and discussing the necessary background from the theory of Fuchsian groups (§1–2). We then treat arithmetic Fuchsian groups and give methods for enumerating “small” elements of the group \mathcal{O}_1^* , with $\mathcal{O} \subset B$ a quaternion order as above (§3). Next, we describe the basic algorithm to reduce an element $g \in \Gamma$ with respect to a finite set $G \subset \Gamma$ (§4). We then prove the main theorem (§5) and conclude by giving two examples (§6).

The author would like to thank the Magma group at the University of Sydney for their hospitality, Steve Donnelly and David Kohel for their

helpful input, Stefan Lemurell for his careful reading of the paper, and Gebhard Böckle, Aurel Page, Jeroen Sijsling, and Charles Stibitz for finding mistakes corrected here.

1. Fuchsian groups

In this section, we present the relevant background from the theory of Fuchsian groups; suggested references include Katok [16, Chapters 3–4] and Beardon [2, Chapter 9]. Throughout, we let $\Gamma \subset \mathrm{PSL}_2(\mathbb{R})$ denote a Fuchsian group with cofinite area, which is finitely generated by a result of Siegel [16, Theorem 4.1.1], [12, §1]. To simplify, we will identify a matrix $g \in \mathrm{SL}_2(\mathbb{R})$ with its image in $\mathrm{PSL}_2(\mathbb{R})$.

Throughout this section, let $p \in \mathfrak{H}$ be a point with trivial stabilizer $\Gamma_p = \{1\}$. Almost all points p satisfy this property: there exist only finitely many p with $\Gamma_p \neq \{1\}$ in any compact subdomain of \mathfrak{H} , and in particular, the set of $p \in \mathfrak{H}$ with $\Gamma_p \neq \{1\}$ have area zero. In practice, with probability 1 a “random” choice of p will suffice.

We define the *Dirichlet domain* centered at p to be

$$D(p) = \{z \in \mathfrak{H} : d(z, p) \leq d(gz, p) \text{ for all } g \in \Gamma\}.$$

The set $D(p)$ is a fundamental domain for Γ , and is a *hyperbolic polygon*. More generally, we define a *generalized hyperbolic polygon* to be a closed, connected, and hyperbolically convex domain whose boundary consists of finitely many geodesic segments, called *sides*, so that a hyperbolic polygon is a generalized hyperbolic polygon with finite area.

Let $D \subset \mathfrak{H}$ be a hyperbolic polygon. Let $S = S(D)$ denote the set of sides of D , with the following convention: if $g \in \Gamma$ is an element of order 2 which fixes a side s of D , and s contains the fixed point of g , we instead consider s to be the union of two sides meeting at the fixed point of g . We define a labeled equivalence relation on S by

$$P = \{(g, s, s^*) : s^* = g(s)\} \subset \Gamma \times (S \times S).$$

We say that P is a *side pairing* for D if P induces a partition of S into pairs, and we denote by $G(P)$ the projection of P to Γ .

Let $D \subset \mathfrak{H}$ be a hyperbolic polygon. For a vertex v of D , we denote by $\vartheta_D(v)$ the interior angle of D at v . Let P be a side pairing for D . We say that P *satisfies the cycle condition* if for every cycle C of vertices in D under P there exists $e \in \mathbb{Z}_{>0}$ such that

$$\sum_{v \in C} \vartheta_D(v) = \frac{2\pi}{e}.$$

Proposition 1.1. *The Dirichlet domain $D(p)$ has a side pairing P , and the set $G(P)$ generates Γ . Conversely, let $D \subset \mathfrak{H}$ be a hyperbolic polygon*

and let P be a side pairing for D which satisfies the cycle condition. Then D is a fundamental domain for the group generated by $G(P)$.

Proof. The first statement is well-known [2, Theorem 9.3.3], [16, Theorem 3.5.4]. For the second statement, we refer to Beardon [2, Theorem 9.8.4] and the accompanying exercises: the condition that $\mu(D) < \infty$ ensures that any vertex which lies on the circle at infinity is fixed by a hyperbolic element [12, §1]. One must verify Beardon's condition (A6) [2, p. 246] or (A6)' [2, p. 249], which formalizes the equivalent angle condition (g) given by Maskit [19, p. 223]. \square

Remark 1.2. The second statement of Proposition 1.1 extends to a larger class of polygons (see [2, §9.8]), and therefore conceivably our results extend to the class of finitely generated non-elementary Fuchsian groups of the first kind. For simplicity, we restrict to the case of groups with cofinite area.

We can define an analogous equivalence relation on the set of vertices of D , and we say that a vertex v of D is *paired* if each side s containing v is paired to a side s^* via an element $g \in G$ such that gv is a vertex of D .

We now consider the corresponding notions in the hyperbolic unit disc \mathfrak{D} , which will prove more convenient for algorithmic purposes. The maps

$$(1.1) \quad \begin{array}{ccc} \phi : \mathfrak{H} & \rightarrow & \mathfrak{D} & & \phi^{-1} : \mathfrak{D} & \rightarrow & \mathfrak{H} \\ & & z \mapsto \frac{z-p}{z-\bar{p}} & & & & w \mapsto \frac{\bar{p}w-p}{w-1} \end{array}$$

define a conformal equivalence between \mathfrak{H} and \mathfrak{D} with $p \mapsto \phi(p) = 0$. Via the map ϕ , the group Γ acts on \mathfrak{D} as

$$\Gamma^\phi = \phi\Gamma\phi^{-1} \subset \text{PSU}(1,1) = \left\{ \pm \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{PSL}_2(\mathbb{C}) : a = \bar{d}, b = \bar{c} \right\}.$$

We may analogously define a Dirichlet domain $D(q)$ for $q \in \mathfrak{D}$ with $\Gamma_q = \{1\}$, and we have $\phi(D(p)) = D(0) \subset \mathfrak{D}$. To ease notation, we identify Γ with Γ^ϕ by $g \mapsto g^\phi = \phi g \phi^{-1}$ when no confusion can result.

Any matrix $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SU}(1,1)$ acts on \mathfrak{D} , multiplying lengths by $|g'(z)| = |cz + d|^{-2}$, and therefore Euclidean lengths (and areas) are preserved if and only if $|cz + d| = 1$. We define the *isometric circle* of g to be

$$I(g) = \{z \in \mathbb{C} : |cz + d| = 1\};$$

if $c \neq 0$, then $I(g)$ is a circle with radius $1/|c|$ and center $-d/c$, and if $c = 0$ then $I(g) = \mathbb{C}$. We denote by

$$\text{int}(I(g)) = \{z \in \mathbb{C} : |cz + d| < 1\}, \quad \text{ext}(I(g)) = \{z \in \mathbb{C} : |cz + d| > 1\}$$

the *interior* and *exterior* of $I(g)$, respectively.

With these notations, we now find the following alternative description of the Dirichlet domain $D(0) \subset \mathfrak{D}$.

Proposition 1.3.

(a) The domain $D(0)$ is the closure in \mathfrak{D} of

$$\bigcap_{g \in \Gamma \setminus \{1\}} \text{ext}(I(g)).$$

(b) For any $g \in \text{SU}(1, 1)$, we have

$$d(z, 0) \begin{cases} < \\ = \\ > \end{cases} d(gz, 0) \text{ according as } \begin{cases} z \in \text{ext}(I(g)), \\ z \in I(g), \\ z \in \text{int}(I(g)). \end{cases}$$

Proof. See Katok [16, Theorem 3.3.5]; we note that if $g \in \Gamma$ and $c = 0$, then $q = 0$ is a fixed point of g , so by hypothesis $g = 1$, and hence $\text{ext}(I(g)) \neq \emptyset$ for all $g \neq 1$. In particular, since Γ has cofinite area we note that the intersection in (a) is nonempty. \square

Corollary 1.4. For any $g \in \text{SU}(1, 1)$, we have $gI(g) = I(g^{-1})$.

Proof. By Proposition 1.3(b), we have

$$w = gz \in I(g^{-1}) \Leftrightarrow d(g^{-1}w, 0) = d(w, 0) \Leftrightarrow d(z, 0) = d(gz, 0) \Leftrightarrow z \in I(g)$$

and the result follows. \square

Remark 1.5. One can similarly define isometric circles $I(g)$ for $g \in \text{PSL}_2(\mathbb{R})$ acting on \mathfrak{H} . One warning is due, however: although $\phi^{-1}(D(0)) = D(p) \subset \mathfrak{H}$ is again a Dirichlet domain, its sides need not be contained in isometric circles (as the map ϕ is a hyperbolic isometry, whereas isometric circles are defined by a Euclidean condition). Instead, we see easily that

$$\phi^{-1}I(g^\phi) = \{z \in \mathfrak{H} : d(z, p) = d(gz, p)\},$$

i.e., the isometric circle $I(g^\phi)$ corresponds in \mathfrak{H} to the perpendicular bisector of the geodesic between p and $g(p)$. In particular, if $p = i$ then

a somewhat lengthy calculation reveals that for $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(\mathbb{R})$,

this perpendicular bisector is the half-circle of radius $\frac{a^2 + b^2 + c^2 + d^2 - 2}{(a^2 + c^2 - 1)^2}$

centered at $\frac{ab + cd}{a^2 + c^2 - 1} \in \mathbb{R}$.

The domain $D(0)$ is also known as a *Ford domain*, since Proposition 1.3 is originally attributed to Ford [11, Theorem 7, §20]. The heart of our algorithm (as provided in the main theorem) will be to algorithmically construct a Ford domain.

2. Algorithms for the upper half-plane and unit disc

We represent points $p \in \mathfrak{H}, \mathfrak{D}$ using exact complex arithmetic: see Pour-El–Richards [20], Weihrauch [25] for theoretical foundations (the subject of computable analysis) and e.g. Boehm [3], Gowland-Lester [13] for a discussion of practical implementations. Alternatively, our algorithms can be interpreted using fixed and sufficiently large precision; even though one cannot predict in advance the precision required to guarantee correct output, it is likely that an error due to round-off will only very rarely occur in practice; see also Remark 2.6. The induced action on \mathfrak{D} has $\Gamma \leftrightarrow \Gamma^\phi \subset \mathrm{SU}(1, 1)$, represented as matrices with exact complex entries.

A Fuchsian group Γ is *exact* if it has a finite set of generators $G \subset \mathrm{SL}_2(K)$ with $K \hookrightarrow \overline{\mathbb{Q}} \cap \mathbb{R}$ a number field; from now on, we assume that the group Γ is exact. Even up to conjugation in $\mathrm{PSL}_2(\mathbb{R})$, not every finitely generated Fuchsian group is exact; our methods conceivably extend to the case where the set of generators $G \subset \mathrm{SL}_2(\mathbb{R})$ are specified with (exact) real entries, but we will not discuss this case any further. Algorithms for efficiently computing with algebraic number fields are well-known (see e.g. Cohen [6]).

We now discuss some elementary methods for working with generalized hyperbolic polygons in \mathfrak{D} , which are defined analogously as those in \mathfrak{H} .

Let $\overline{\mathfrak{D}} = \{z \in \mathbb{C} : |z| \leq 1\}$ denote the closure of \mathfrak{D} and let $\partial\mathfrak{D} = \{z \in \mathbb{C} : |z| = 1\}$ be the *circle at infinity*. We represent a geodesic L in \mathfrak{D} in bits by four pieces of data:

- the center $c = \mathrm{ctr}(L) \in \mathbb{C} \cup \{\infty\}$,
- the radius $r = \mathrm{rad}(L) \in \mathbb{R} \cup \{\infty\}$ of L , and
- the *initial point* $z = \mathrm{in}(L) \in \overline{\mathfrak{D}}$ and the *terminal point* $w \in \overline{\mathfrak{D}}$;

the initial and terminal points are normalized so that the path along L follows a counterclockwise orientation. Although this data is redundant, it will be more efficient in practice to store all values rather than, say, to recompute c and r when needed.

If $L_1, L_2 \subset \mathfrak{D}$ are geodesics which intersect at a point $v \in \mathfrak{D} \setminus \{0\}$, then we define $\angle(L_1, L_2)$ to be the counterclockwise-oriented angle at v from the geodesics L_1 to L_2 for the wedge directed toward the origin, so that in particular we have $\angle(L_2, L_1) = -\angle(L_1, L_2)$.

Example 2.1. In Figure 2.1, we depict a geodesic and the angle $\angle(L_1, L_2) \approx 3\pi/8$ between geodesics.

We leave it to the reader to show that one can compute using elementary formulae the following quantities: for geodesics L_1, L_2 , the intersection $L_1 \cap L_2$ and (if nonzero) the angle $\angle(L_1, L_2)$, as well as the area of a hyperbolic polygon.




Figure 2.1: Geodesics and angles

Definition 2.2. Let $G \subset \Gamma \setminus \{1\}$. The *exterior domain* of G , denoted $E = \text{ext}(G)$, is the closure in $\overline{\mathfrak{D}}$ of the set $\bigcap_{g \in G} \text{ext}(I(g)) \cap \mathfrak{D}$.

With this definition, Proposition 1.3(a) becomes simply the statement that $\text{ext}(\Gamma \setminus \{1\})$ is the closure of $D(0)$.

Let $G \subset \Gamma$ be a finite subset and let $E = \text{ext}(G)$ be its exterior domain. Then E is a generalized hyperbolic polygon whose sides are contained in isometric circles $I(g)$ with $g \in G$. A *proper vertex* of E is a point of intersection $v \in I(g) \cap I(g')$ between two sides (with $g \neq g' \in G$); a *vertex at infinity* of E is a point of intersection $v \in I(g) \cap \partial\mathfrak{D}$ between a side and the circle of infinity. A *vertex* of E is either a proper vertex or a vertex at infinity.

Definition 2.3. Let $E = \text{ext}(G)$ be an exterior domain. A sequence $U = g_1, \dots, g_n$ is a *normalized boundary* for E if:

- (i) $E = \text{ext}(U)$;
- (ii) $I(g_1), \dots, I(g_n)$ contain the counterclockwise consecutive sides of D ;
and
- (iii) the vertex $v \in E$ with minimal $\arg(v) \in (0, 2\pi)$ is either a proper vertex with $v \in I(g_1) \cap I(g_2)$ or a vertex at infinity with $v \in I(g_1)$.

It is clear that for each exterior domain E , there exists a unique normalized boundary G for E : in (i) and (ii) we order exactly those g_i for which $I(g_i)$ are sides of E and in (iii) we choose a consistent place to start.

Example 2.4. In the following figure, we exhibit a normalized boundary $G = \{g_1, g_2, g_3, g_4\}$; the vertices v_1, v_2 are on the circle at infinity whereas v_3, v_4, v_5 are proper.




Figure 2.4: Normalized boundary of a generalized hyperbolic polygon

We now detail an algorithm which computes a normalized boundary for a given exterior domain.

Algorithm 2.5. Let $G \subset \Gamma$ be a finite subset. This algorithm returns the normalized boundary U of the exterior domain $E = \text{ext}(G)$.

1. Initialize $\theta := 0$, $U := \emptyset$, and $L := [0, 1]$.
2. Let

$$H := \{g \in G : \arg(I(g) \cap L) \geq \theta\}.$$

- a. If $H = \emptyset$, let $g \in G$ be such that

$$\theta := \arg(\text{in}(I(g))) \in \theta + [0, 2\pi)$$

is minimal. If more than one such g exists, let g be the one which maximizes $\text{rad } I(g)$.

- b. Otherwise $H \neq \emptyset$. If $L = [0, 1]$, then let $g \in H$ be such that $I(g) \cap L$ is minimal. If $L \neq [0, 1]$, let $g \in H$ be such that

$$\theta := \arg(I(g) \cap L) \in \theta + [0, 2\pi)$$


is minimal. In either case, if more than one such g exists, let g be the one that minimizes $\angle(L, I(g))$.

Let $U := U \cup \{g\}$ and let $L := I(g) \cap \overline{\mathcal{D}}$.

3. If $U = \{g_1, \dots, g_n\}$ and $g_n = g_1$, return $U := \{g_1, \dots, g_{n-1}\}$. Otherwise, return to Step 2.

Proof of correctness. By definition $\text{ext}(U)$ is a generalized hyperbolic polygon. Suppose that $E \neq \text{ext}(U)$. Then there exists $g \in G$ such that $L = I(g) \cap \text{ext}(U)$ is not just a vertex of $\text{ext}(U)$. Consider the initial point $z = \text{in}(L)$: either z lies on a side $I(g_i)$ of $\text{ext}(U)$ or $z \in \partial\mathcal{D}$.

Suppose that $z \in I(g_i)$. Let v_i be the initial vertex of the side $s_i \subset I(g_i)$. Then in the i th iteration of Step 2 of the algorithm we have $g \in H$, so the terminal vertex v_{i+1} of s_i is proper and we are in case (b). But by assumption we have $d(v_i, z) \leq d(v_i, v_{i+1})$ since $I(g_i)$ is a geodesic, and \arg increases along s_i with the distance, thus according to the stipulations of the algorithm we must have $z = v_{i+1}$. But then in order for the interior of $I(g)$ to intersect $\text{ext}(U)$ nontrivially, we must have $\angle(L, I(g_i)) < \angle(I(g_{i+1}), I(g))$, a contradiction.



So suppose that $z \in \partial\mathfrak{D}$. Then there exists i such that z lies on the principal circle between the terminal point of $I(g_i)$ and the initial point of $I(g_{i+1})$. But then $\arg(\text{in}(I(g))) < \arg(\text{in}(I(g_{i+1})))$, contradicting (a). This proves that (i) holds in Definition 2.3.

It is obvious that (ii) holds, and condition (iii) holds by initialization: if the vertex $v \in E$ with minimal $\arg(v) \in (0, 2\pi)$ is a vertex at infinity then it is found in the first iteration of the algorithm in stage (a), and if it is a proper vertex then it is found in the second iteration in stage (b). \square

A Ford domain $D(0)$ is specified in bits by a normalized boundary G for $D(0)$. We can similarly specify a Dirichlet domain $D(p)$ by an analogously defined normalized boundary of perpendicular bisectors, as in Remark 1.5; for many purposes, it will be sufficient to represent $D(p)$ by a sequence of vertices (ordered in a counterclockwise orientation around p).

Remark 2.6. Although the intermediate computations as above are of a numerical sort, an algorithm to compute a Dirichlet domain accepts exact input and produces exact output.

3. Element enumeration in arithmetic Fuchsian groups

In this section, we treat arithmetic Fuchsian groups, and in particular we exhibit methods for enumerating “small” elements of these groups. See Vigneras [23] for background material and Voight [24, Chapter 4] for a discussion of algorithms for quaternion algebras.

Let F be a number field with $[F : \mathbb{Q}] = n$ and discriminant d_F . A *quaternion algebra* B over F is an F -algebra with generators $\alpha, \beta \in B$ such

that

$$\alpha^2 = h, \quad \beta^2 = k, \quad \beta\alpha = -\alpha\beta$$

with $h, k \in F^*$; such an algebra is denoted $B = \left(\frac{h, k}{F}\right)$ and is specified in bits by $h, k \in F^*$. An element $\gamma \in B$ is represented by $\gamma = x + y\alpha + z\beta + w\alpha\beta$ with $x, y, z, w \in F$, and we define the *reduced trace* and *reduced norm* of γ by $\text{trd}(\gamma) = 2x$ and $\text{nrd}(\gamma) = x^2 - hy^2 - kz^2 + hkw^2$, respectively.

Let B be a quaternion algebra over F and let \mathbb{Z}_F denote the ring of integers of F . An *order* $\mathcal{O} \subset B$ is a finitely generated \mathbb{Z}_F -submodule with $F\mathcal{O} = B$ which is also subring; an order is *maximal* if it is not properly contained in any other order. We represent an order by a *pseudobasis* over \mathbb{Z}_F ; see Cohen [7, §1] for methods of computing with finitely generated modules over Dedekind domains using pseudobases.

A place v of F is *split* or *ramified* according as $B_v = B \otimes_F F_v \cong M_2(F_v)$ or not, where F_v denotes the completion at v . The set S of ramified places of B is finite and of even cardinality, and the ideal $\mathfrak{d} = \prod_{v \in S, v \neq \infty} \mathfrak{p}_v$ of \mathbb{Z}_F is called the *discriminant* of B .

Now suppose that F is a totally real field, and there is a unique split real place $v \notin S$ corresponding to $\iota_\infty : B \hookrightarrow M_2(\mathbb{R})$. Let $\mathcal{O} \subset B$ be an order and let \mathcal{O}_1^* denote the group of units of reduced norm 1 in \mathcal{O} . Then the group $\Gamma(\mathcal{O}) = \iota_\infty(\mathcal{O}_1^*/\{\pm 1\}) \subset \text{PSL}_2(\mathbb{R})$ is a Fuchsian group [16, §§5.2–5.3]. If \mathcal{O} is maximal, we denote $\Gamma^B(1) = \Gamma(\mathcal{O})$. An *arithmetic Fuchsian group* Γ is a Fuchsian group commensurable with $\Gamma^B(1)$ for some choice of B . One can, for instance, recover the usual modular groups in this way, taking $F = \mathbb{Q}$, $\mathcal{O} = M_2(\mathbb{Z}) \subset M_2(\mathbb{Q}) = B$, and $\Gamma \subset \text{PSL}_2(\mathbb{Z})$ a subgroup of finite index.

An arithmetic Fuchsian group Γ has cofinite area; indeed, by a formula of Shimizu [21, Appendix], the area $A = \mu(X) = \mu(\Gamma \backslash \mathfrak{H})$ is given by

$$(3.1) \quad A = \frac{4}{(2\pi)^{2n}} d_F^{3/2} \zeta_F(2) \Phi(\mathfrak{d}) [\Gamma^B(1) : \Gamma],$$

where $\zeta_F(s)$ denotes the Dedekind zeta function of F , and

$$\Phi(\mathfrak{d}) = \#(\mathbb{Z}_F/\mathfrak{d}\mathbb{Z}_F)^* = N(\mathfrak{d}) \prod_{\mathfrak{p}|\mathfrak{d}} \left(1 - \frac{1}{N(\mathfrak{p})}\right);$$

here the hyperbolic area is normalized so that

$$\mu(\Omega) = \frac{1}{2\pi} \iint_{\Omega} \frac{dx dy}{y^2}$$

and hence an ideal triangle has area $1/2$.

Remark 3.1. The area A is effectively computable from the formula (3.1). By the Riemann-Hurwitz formula, we have

$$(3.2) \quad A = 2g - 2 + \sum_q e_q \left(1 - \frac{1}{q}\right) + e_\infty$$

where e_q is the number of elliptic cycles of order $q \in \mathbb{Z}_{\geq 2}$ in Γ and e_∞ the number of parabolic cycles. In particular, $A \in \mathbb{Q}$; and since $e_q > 0$ implies $F(\zeta_{2q}) \hookrightarrow B$, the denominator of A is bounded by the least common multiple of all q such that $[F(\zeta_{2q}) : F] = 2$ (which in particular requires that F contains the totally real subfield $\mathbb{Q}(\zeta_{2q})^+$ of $\mathbb{Q}(\zeta_{2q})$). Therefore, it suffices to compute the usual Dirichlet series or Euler product expansion for $\zeta_F(2)$ with the required precision; see also Dokchitser [9].

We now relate isometric circles to the arithmetic of B . Let $p \in \mathfrak{H}$ have $\Gamma_p = \{1\}$. A short calculation with the maps defined in (1.1) shows that if $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$, then $g^\phi = \phi g \phi^{-1} \in \mathrm{SU}(1, 1)$ has radius

$$\mathrm{rad}(I(g^\phi)) = \frac{2 \mathrm{Im}(p)}{|f_g(p)|},$$

where $f_g(t) = ct^2 + (d - a)t - b$, a polynomial whose roots are the fixed points of g in \mathbb{C} . We will abbreviate $\mathrm{rad}(g) = \mathrm{rad}(I(g^\phi))$. The map

$$(3.3) \quad \begin{aligned} \mathrm{inrad} : M_2(\mathbb{R}) &\rightarrow \mathbb{R} \\ g &\mapsto |f_g(p)|^2 + 2y^2 \det(g) \end{aligned}$$

yields a quadratic form on $M_2(\mathbb{R})$: explicitly, if $p = x + yi$, we have

$$\begin{aligned} \mathrm{inrad} \begin{pmatrix} a & b \\ c & d \end{pmatrix} &= (xa + b - (x^2 - y^2)c - xd)^2 + y^2(a - 2xc - d)^2 + 2y^2(ad - bc) \\ &= y^2(a - xc)^2 + (xa + b - x^2c - xd)^2 + y^4c^2 + y^2(xc + d)^2, \end{aligned}$$

and hence the form inrad is positive definite and via ι_∞ induces a positive definite form $\mathrm{inrad} : B \rightarrow \mathbb{R}$. For $g \in B$, we note that $\det \iota_\infty(g) = v(\mathrm{nrd}(g))$, where v is the unique split real place of B .

Suppose that $p = i$. Then we have simply $\mathrm{inrad} \begin{pmatrix} a & b \\ c & d \end{pmatrix} = a^2 + b^2 + c^2 + d^2$. Let $B = \left(\frac{h, k}{F}\right)$. Identify F with its image $F \hookrightarrow \mathbb{R}$ under the unique split real place of B ; without loss of generality, we may assume that $h > 0$. We may therefore embed $\iota_\infty : B \hookrightarrow M_2(\mathbb{R})$ by letting

$$(3.4) \quad \alpha \mapsto \begin{pmatrix} \sqrt{h} & 0 \\ 0 & -\sqrt{h} \end{pmatrix}, \quad \beta \mapsto \begin{pmatrix} 0 & \sqrt{|k|} \\ \mathrm{sgn}(k)\sqrt{|k|} & 0 \end{pmatrix}$$

where sgn denotes the sign. Therefore if $g = x + y\alpha + z\beta + w\alpha\beta \in B$, then we see directly that

$$\text{invrad}(g) = x^2 + hy^2 + |k|z^2 + h|k|w^2.$$

For the ramified real places v of F , corresponding to $B \hookrightarrow B \otimes_F \mathbb{R} \cong \mathbb{H}$, the reduced norm form $\text{nrd}_v : B \rightarrow \mathbb{R}$ by $g \mapsto v(\text{nrd}(g))$ is positive definite. Putting these together, we find that the *absolute reduced norm*

$$N : B \rightarrow \mathbb{R}$$

$$g \mapsto 2y^2 \sum_{v \in S, v|_\infty} \text{nrd}_v(g) + \text{invrad}(g) = |f_g(p)|^2 + 2y^2 \text{Tr}_{F/\mathbb{Q}} \text{nrd}(g)$$

is positive definite and gives \mathcal{O} the structure of a lattice of rank $4n$.

The elements $g \in \mathcal{O}$ with small absolute reduced norm N are those such that $|f_g(p)|$ and $\text{Tr}_{F/\mathbb{Q}} \text{nrd}(g)$ are both small—in particular, this will include the elements of \mathcal{O}_1^* with small invrad (with respect to $p \in \mathfrak{H}$), which correspond to elements $g \in \Gamma$ whose isometric circle in \mathfrak{D} (centered at p) has large radius. Since the Dirichlet domain $D(p)$ has only finitely many sides, those $g \in \Gamma$ with $\text{rad}(g)$ sufficiently small radius cannot contribute to the boundary of $D(p)$.

Hence, one simple idea to construct $D(p)$ would be to enumerate all elements of \mathcal{O}_1^* by increasing absolute reduced norm N until the exterior domain of these elements has area equal to $\mu(\Gamma \backslash \mathfrak{H})$. This method shows that $D(p)$ is indeed computable, and may have been known to Klein; it is mentioned by Katok [17] when $F = \mathbb{Q}$ and sees further explication by Johansson [15]. Using the above framework, we can immediately improve upon this method by enumerating such elements efficiently using lattice reduction, as follows.

Algorithm 3.2. Let $\mathcal{O} \subset B$ be a quaternion order. This algorithm returns a Dirichlet domain for $\Gamma(\mathcal{O})$.

1. Compute $A = \mu(\Gamma(\mathcal{O}) \backslash \mathfrak{H})$ by Remark 3.1.
2. Embed $\mathcal{O} \hookrightarrow \mathbb{R}^{4n}$ as a lattice using the absolute reduced norm form N , and choose $C \in \mathbb{R}_{>0}$.
3. Using the Fincke-Pohst algorithm [10], compute the set

$$G(C) = \{ \iota_\infty(g/u) : \pm g \in \mathcal{O}, N(g) \leq C, \text{nrd}(g) = u^2 \in \mathbb{Z}_F^{*2} \} \subset \Gamma.$$

4. From Algorithm 2.5, compute $E = \text{ext}(G(C))$. If $\mu(E) = A < \infty$, then return E ; otherwise, increase C and return to Step 2.

Remark 3.3. In choosing C , we note that

$$\{g \in \mathcal{O}_1^* : \text{rad}(g) \geq R\} = \left\{ g \in \mathcal{O} : N(g) \leq 2y^2 \left(n + \frac{2}{R^2} \right) \right\} \cap \mathcal{O}_1^*;$$

in practice, we would like to take C large enough so that $G(C) \neq \emptyset$ but not too large. It is not immediately clear how to choose C (and a strategy

for its incrementation) optimally in general, unless one knows something about the radii of the sides of the Dirichlet domain.

Our final algorithm (Algorithm 4.9) significantly improves on Algorithm 3.2 by the use of a reduction algorithm, which we introduce in the next section.

4. Reduction algorithm

In this section, we introduce the reduction algorithm (Algorithm 4.3) which forms the heart of the paper. This algorithm will allow us to find a normalized basis for the group Γ (Algorithm 4.7), yielding a fundamental domain.

Throughout this section, let $G = \{g_1, \dots, g_t\} \subset \Gamma \setminus \{1\}$ be an (ordered) finite subset of a Fuchsian group Γ , and denote by $\langle G \rangle$ the group generated by G . For any $z \in \mathfrak{D}$, we have a map

$$\begin{aligned} \rho : \Gamma &\rightarrow \mathbb{R}_{\geq 0} \\ \gamma &\mapsto \rho(\gamma; z) = d(\gamma z, 0) \end{aligned}$$

where d denotes hyperbolic distance. We abbreviate $\rho(\gamma; 0) = \rho(\gamma)$.

Definition 4.1. Let $z \in \mathfrak{D}$. An element $\gamma \in \Gamma$ is (G, z) -reduced if for all $g \in G$, we have $\rho(\gamma; z) \leq \rho(g\gamma; z)$, and γ is G -reduced if it is $(G, 0)$ -reduced.

Remark 4.2. By Proposition 1.3, we note that γ is (G, z) -reduced if and only if $\gamma z \in \text{ext}(G)$.

We arrive at the following straightforward algorithm to perform (G, z) -reduction.

Algorithm 4.3. Let $\gamma \in \Gamma$ and let $z \in \mathfrak{D}$. This algorithm returns elements $\bar{\gamma} \in \Gamma$ and $\delta \in \langle G \rangle$ such that $\bar{\gamma}$ is (G, z) -reduced and $\bar{\gamma} = \delta\gamma$.

1. Initialize $\bar{\gamma} := \gamma$ and $\delta := 1$.
2. If $\rho(\bar{\gamma}; z) \leq \rho(g\bar{\gamma}; z)$ for all $g \in G$, return $\bar{\gamma}, \delta$. Otherwise, let $g \in G$ be the first element in G such that

$$\rho(g\bar{\gamma}; z) = \min_i \rho(g_i\bar{\gamma}; z).$$

Let $\bar{\gamma} := g_i\bar{\gamma}$ and $\delta := g_i\delta$, and return to Step 2.

We denote the output of the above algorithm $\bar{\gamma} = \text{red}_G(\gamma; z)$ and abbreviate $\text{red}_G(\gamma; 0) = \text{red}_G(\gamma)$.

Proof of correctness. The output of the algorithm $\bar{\gamma}$ is by definition G -reduced. The algorithm terminates because if $\bar{\gamma}_1, \bar{\gamma}_2, \dots$ are the elements that arise in the iteration of Step 2, then $\rho(\bar{\gamma}_1; z) > \rho(\bar{\gamma}_2; z) > \dots$; however, the action of Γ is discrete, so among the points $\{\bar{\gamma}_i(z)\}_i$, only finitely many are distinct. \square

A priori, Step 2 in Algorithm 4.3 depends on the ordering of the set G and therefore the output $\bar{\gamma}$ will depend on this ordering. This is analogous to the situation of the reduction theory of polynomials, as follows. Let k be a field, let $R = k[x_1, \dots, x_n]$ be the polynomial ring over k in n variables with a choice of term order, and let $G = g_1, \dots, g_t \in R$ be not all zero. Applying the generalized division algorithm, one can reduce a polynomial $f \in R$ with respect to G , and the result is unique (i.e., independent of the ordering of the g_i) for all f if G is a Gröbner basis of the ideal $I = \langle g_1, \dots, g_t \rangle$. Moreover, if G is a Gröbner basis, then $f \in I$ if and only if the remainder on division of f by G is zero. (See e.g. Cox-Little-O’Shea [8, Chapter 2].) We can prove analogous statements, replacing the ring R by the group Γ , as follows.

Proposition 4.4. *Suppose that $\text{ext}(G)$ is a fundamental domain for $\langle G \rangle$. Then for almost all $z \in \mathfrak{D}$, $\text{red}_G(\gamma; z)$ as an element of Γ is independent of the ordering of G for all $\gamma \in \langle G \rangle$. Moreover, for all $\gamma \in \Gamma$, we have $\text{red}_G(\gamma) = 1$ if and only if $\gamma \in \langle G \rangle$.*

Here, “almost all” means for all z outside of a set of measure zero: it suffices to take z in the Γ -orbit of the interior of $\text{ext}(G)$.

Proof. Suppose that $\text{ext}(G)$ is a fundamental domain for $\langle G \rangle$. Let z be in the Γ -orbit of $z_0 \in \text{int}(\text{ext}(G))$, let $\gamma \in \langle G \rangle$, and let $\bar{\gamma} = \text{red}_G(\gamma; z)$. Then by Remark 4.2, we have $\bar{\gamma}z \in \text{ext}(G)$, and since $\text{ext}(G)$ is a fundamental domain and $\Gamma z = \Gamma z_0$ with $z_0 \in \text{int}(\text{ext}(G))$, we must have $\bar{\gamma}z = z_0$; in particular, $\bar{\gamma}$ is unique and independent of the ordering of G . The second statement follows similarly: we have that $0, \bar{\gamma}(0) \in \text{int}(\text{ext}(G))$, so if $\bar{\gamma} \neq 1$ then $\gamma \notin \langle G \rangle$. \square

Inspired by the preceding proposition, we make the following definition.

Definition 4.5. A set G is a *basis* for Γ if $\text{ext}(G)$ is a fundamental domain for $\langle G \rangle = \Gamma$. If G is a basis that forms a normalized boundary for Γ , then we say that G is a *normalized basis*.

Remark 4.6. It follows from Proposition 4.4 that if one can compute a normalized basis G for Γ , then one also has a solution to the word problem: given any element $\gamma \in \Gamma$, we compute $\bar{\gamma} = \text{red}_G \gamma$, which by Proposition 4.4 must satisfy $\bar{\gamma} = 1$, so we have explicitly written γ as a word from G .

We construct a normalized basis for as follows.

Algorithm 4.7. Let $G \subset \Gamma$. This algorithm returns a normalized basis for $\langle G \rangle$ for all points $p \in \mathfrak{H}$ outside of a set of measure zero.

1. Let $G := \{g_1, \dots, g_t, g_1^{-1}, \dots, g_t^{-1}\}$.
2. Compute the normalized boundary U of $\text{ext}(G)$ by Algorithm 2.5.

3. Let $G' := U$. For each $g \in G$, compute $\bar{g} = \text{red}_U(g)$ using Algorithm 4.3. If $\bar{g} \neq 1$, set $G' := G' \cup \{\bar{g}\}$.
4. Compute the normalized boundary U' of $\text{ext}(G')$. Let $G := G'$. If $U' = U$, proceed to Step 5; otherwise set $U := U'$ and return to Step 3.
5. If all vertices of $E = \text{ext}(U)$ are paired, return U . Otherwise, for each $g \in G$ with a vertex $v \in I(g)$ which is not paired, compute $\bar{g} := \text{red}_G(g; v)$, where if v is a vertex at infinity we replace v by a nearby point in $I(g^{-1}) \setminus E \subset \mathfrak{D}$. Add the reductions \bar{g} for each unpaired vertex v to G and return to Step 2.

Proof of correctness. First, note that if v be a vertex of $E = \text{ext}(G)$, then by Corollary 1.4, v is a paired vertex if and only if for every side $s \subset I(g)$ containing v , we have that $gv \in I(g^{-1})$ is a vertex of E .

Next, we prove that if the algorithm terminates it does so correctly. We construct a side pairing as in §1. A side s of E pairs up with $gs \subset I(g^{-1})$ if and only if its vertices are paired, necessarily with the vertices of $I(g^{-1})$ by Corollary 1.4. Therefore if we terminate in Step 5, we have in fact paired all sides of $\text{ext}(U)$ and by Proposition 1.1, $\text{ext}(U)$ is a Dirichlet domain and U is a basis.

We must argue that the output $D = \text{ext}(U)$ of Algorithm 4.7 satisfies the cycle condition. Let C be a cycle of vertices in D . Consider small neighborhoods of each vertex in C in D . If these neighborhoods are disjoint under the action of Γ , then they glue to give a neighborhood in the quotient $\Gamma \backslash \mathfrak{H}$, hence the cycle condition holds for C . Making these neighborhoods smaller, we may assume that each $v \in C$ is an elliptic fixed point. But then by Proposition 5.4 (which applies equally well to exterior domains) and the accompanying discussion, we may assume that the elliptic cycle has length 1, and consequently the cycle condition is trivially satisfied.

Otherwise, by Step 5 we have $v \in s$ such that $gv \notin \text{ext}(G)$. We now compute $\bar{g} = \text{red}_G(g; v)$, and refer to Proposition 1.3. Since $v \in I(g)$, we have $d(v, 0) = d(gv, 0)$, and since $gv \notin \text{ext}(G)$, we have $d(gv, 0) > d(\bar{g}v, 0)$. Putting these together, we find that $v \in \text{int}(I(\bar{g}))$ and hence $\text{ext}(G \cup \{\bar{g}\}) \subsetneq \text{ext}(G)$.

Consider now the limit of the sets $G_\infty = \lim G$ and $U_\infty = \lim U$ as we let the algorithm run forever. Accordingly, every vertex v of $\text{ext}(U_\infty)$ must be paired, otherwise it would be caught in some step of the algorithm. Therefore by the above, U_∞ is a basis for $\langle G_\infty \rangle$. But at each step of the algorithm, the group $\langle G \rangle$ remains the same, even as G changes: indeed, in Step 3, if $\bar{g} = 1$ then $g \in \langle G \setminus \{g\} \rangle$. Therefore $\langle G_\infty \rangle = \langle G \rangle$, and since $\langle G \rangle$ is finitely generated we know that U is finite, and hence the algorithm terminates after finitely many steps. \square

Remark 4.8. If one desires Algorithm 4.7 to work for every point $p \in \mathfrak{H}$, then modify the last two steps as follows.

5. If all vertices of $E = \text{ext}(U)$ are paired, proceed to Step 6. Otherwise, for each $g \in G$ with a vertex $v \in I(g)$ which is not paired, compute $\bar{g} := \text{red}_G(g; v)$, where if v is a vertex at infinity we replace v by a nearby point in $I(g^{-1}) \setminus E \subset \mathfrak{D}$. Add the reductions \bar{g} for each nonpaired vertex v to G and return to Step 2.
6. Run Algorithm 5.2, and let S be the set of minimal cycles $g \neq 1$ with a fixed point in \mathfrak{D} . If $S \subseteq U$, return U ; otherwise, set $U := U \cup S \cup S^{-1}$ and return to Step 2.

Following as in the proof above, in the third paragraph we have checked explicitly that the elliptic fixed point $v \in C$ has a neighborhood of the right size, with edges contained in $I(g)$ and $I(g^{-1})$ where g is the minimal cycle fixing v .

We now extend this in the natural way to an arithmetic Fuchsian group $\Gamma(\mathcal{O})$.

Algorithm 4.9. Let \mathcal{O} be a quaternion order. This algorithm returns a basis G for $\Gamma = \Gamma(\mathcal{O})$.

1. Choose $C \in \mathbb{R}_{>0}$, initialize $G := \emptyset$, and compute $A = \mu(\Gamma \setminus \mathfrak{H})$.
2. Using Steps 1–2 in Algorithm 3.2, compute the set $G(C) \subset \Gamma$.
3. Call Algorithm 4.7 with input $G \cup G(C)$ and let G be the output. If $\mu(\text{ext}(G)) = A < \infty$, then return G ; otherwise, increase C and return to Step 2.

A fundamental domain for an arithmetic Fuchsian group $\Gamma \subset \Gamma(\mathcal{O})$ can easily be computed from this by first running Algorithm 4.9 and then computing a coset decomposition of Γ in $\Gamma(\mathcal{O})$; and for that reason, one may even restrict consideration to the case where \mathcal{O} is maximal.

Remark 4.10. In practice, in some cases we can improve Step 5 of Algorithm 4.7 for arithmetic Fuchsian groups as follows. For each nonpaired vertex v , we can consider those elements with small absolute reduced norm N relative to $p \in \mathfrak{D}$ taken to be a point along the geodesic between 0 and v : indeed, by continuity if $g \in \mathcal{O}_1^*$ has $v \in \text{int}(I(g))$, then $\text{rad}(g)$ increases as the center p moves towards v and thus $N(g)$ decreases, so using lattice reduction we are likely to find a small such g .

5. Proof of the main theorem

We are now ready to prove the main theorem of this paper.

Theorem 5.1. *There exists an algorithm which, given a finitely generated Fuchsian group Γ and a point $p \in \mathfrak{H}$ with $\Gamma_p = \{1\}$, returns the Dirichlet*

domain $D(p)$, a side pairing for $D(p)$, and a finite presentation for Γ with a minimal set of generators.

To prove the theorem, we need to show how the output of Algorithm 4.7 yields a finite presentation for Γ with a minimal set of generators. Indeed, Algorithm 4.7 terminates only if it has computed a side pairing P (which we may assume meets the convention in §1) for the Dirichlet domain D . Such a side pairing P gives a set G of generators for Γ by Proposition 1.1.

We now consider the induced relation on the set of vertices. A *cycle* of D is a sequence $v_1, \dots, v_n = v_1$ which is the (ordered) intersection of the Γ -orbit of $v = v_1$ with D . To each cycle, we associate the word $g = g_n g_{n-1} \cdots g_2 g_1$ where $g_i(v_i) = v_{i+1}$ and the indices are taken modulo n . We say that a cycle is a *pairing cycle* if $g_i \in G$ for all i , and without further mention we shall assume from now on that a cycle is a pairing cycle.

A cycle is *minimal* if $v_i \neq v_j$ for all $i \neq j$. Every vertex v of D is contained in a unique minimal cycle (up to reversion and cyclic permutation). Indeed, by the uniqueness of the side pairing, a vertex $v \in I(g) \cap I(g')$ either has $v = gv = g'v$, in which case v has nontrivial stabilizer and one has the singleton cycle v , or v has trivial stabilizer and is paired with the distinct elements $gv \in I(g^{-1})$ and $g'v \in I(g'^{-1})$, each of which also has trivial stabilizer, and then continuing in this way one constructs a (unique minimal) cycle. This analysis gives rise to the following algorithm.

Algorithm 5.2. Let P be a side pairing for a Dirichlet domain D for Γ . This algorithm returns a set of minimal cycles for D .

1. Initialize V to be the set of vertices of D and $M := \emptyset$.
2. If $V = \emptyset$, terminate. Otherwise, choose $v \in V$ with $v \in I(g) \cap I(g')$ for $g, g' \in G(P)$. If $gv = v$, add the cycle v to M and remove v from V , and return to Step 2. Otherwise, let $i := 1$ and $v_1 := v$.
3. Let $v_{i+1} := gv_i \in I(g^{-1}) \cap I(g')$. If $v_{i+1} = v_1$, add the cycle v_1, \dots, v_i to M , remove these elements from V , and return to Step 2; otherwise, increment $i := i + 1$, let $g := g'$ and return to Step 3.

The relations associated to minimal cycles have the following important property.

Lemma 5.3. *Let $g \in G$ be a side-pairing element. Then g appears at most once in any word associated to a minimal cycle. Moreover, g and its inverse appears in exactly two such words.*

Proof. By definition, a side-pairing element g pairs a unique set of sides: in particular, g pairs the vertices of one side s with the vertices of another. Suppose that g occurs twice in a word associated to a minimal cycle. Then by minimality, the vertices of s are in the same Γ orbit. But this implies that g maps $I(g)$ to itself, so g has order 2 and therefore one of the vertices of s is fixed by g , a contradiction.

In a similar way, we see that g and its inverse can appear in at most two words since each vertex belongs to exactly one minimal cycle. \square

Now, to each cycle, associated to the word g , we further associate a relation in Γ as follows. By definition, we have $g \in \Gamma_v$, and therefore we have one of three possibilities. If $\#\Gamma_v = 1$, then we have the relation $g = 1$; we call g an *accidental cycle*. If $1 < \#\Gamma_v < \infty$, then we associate the relation $g^k = 1$ where k is the order of g , and we call g an *elliptic cycle*. Otherwise, if $\#\Gamma_v = \infty$, then we associate the empty relation, a *parabolic cycle*. We note that the latter occurs if and only if g has infinite order if and only if $\text{trd}(g) = \pm 2$, so the relation g is computable.

We have the following characterization of the minimal cycles.

Proposition 5.4 (Beardon [2, Theorem 9.4.5]). *For all $p \in \mathfrak{H}$ outside of a set of area zero, the following statements hold:*

- (i) *Every elliptic cycle has length 1;*
- (ii) *Every accidental cycle has length 3; and*
- (iii) *Every parabolic cycle has length 1.*

Remark 5.5. The exceptional set of p is contained in the union

$$E_2 = \bigcup_{f,g,h \in \Gamma} \{z : R(z) \in \mathbb{R}\}$$

over all triples $f, g, h \in \Gamma$ such that

$$R(z) = \frac{(z - gz)(fz - hz)}{(z - fz)(gz - hz)}$$

is not constant. It is easy to see that the set E_2 has area zero.

For the purposes of computing a minimal set of generators and relations, we may and do assume that p does not lie in the exceptional set; indeed, a sufficiently general choice of p will suffice, and so in practice the conditions of Proposition 5.4 always hold. In particular, every elliptic cycle is represented by a minimal cycle (whose fixed point is a vertex of D).

We now appeal to the structure theory for Fuchsian groups with cofinite area [16, §4.3]. Suppose that Γ has exactly t elliptic cycles of orders $m_1, \dots, m_t \in \mathbb{Z}_{\geq 2}$ and s parabolic cycles, and that $X = \Gamma \backslash \mathfrak{H}$ has genus g . We say then that Γ has signature $(g; m_1, \dots, m_t; s)$. Moreover, Γ is generated by elements

$$(5.1) \quad \alpha_1, \dots, \alpha_g, \beta_1, \dots, \beta_g, \gamma_1, \dots, \gamma_t, \gamma_{t+1}, \dots, \gamma_{t+s}$$

subject to the relations

$$(5.2) \quad \gamma_1^{m_1} = \dots = \gamma_t^{m_t} = [\alpha_1, \beta_1] \cdots [\alpha_g, \beta_g] \gamma_1 \cdots \gamma_{t+s} = 1,$$

where $[\alpha, \beta] = \alpha\beta\alpha^{-1}\beta^{-1}$ is the commutator. (One obtains a minimal set of generators from this presentation by eliminating γ_{t+s} whenever $t + s > 0$.)

From the set of generators coming from the side-pairing elements and the set of relations coming from the minimal cycles, we can build a minimal set of generators and relations by “back substitution”. First, we prove a lemma.

Lemma 5.6. *Suppose $\Gamma \cong \Gamma_1 * \Gamma_2$ is a free product, and that $\gamma_i \in \Gamma_1$ or Γ_2 for $i = 1, \dots, s + t$. Then either Γ_1 or Γ_2 is isomorphic to the free product of cyclic groups.*

Proof. Let $\phi : \Gamma \xrightarrow{\sim} \Gamma_1 * \Gamma_2$ be an isomorphism. Passing to the quotient by the γ_i , for $i = 1, \dots, t + s$, we may assume that $s = t = 0$. But then the homology groups $H_i(\Gamma, \mathbb{Z})$ (coming from group homology) coincide with the homology groups $H_i(Y, \mathbb{Z})$ (coming from topology) where Y is the orientable surface of genus g [5, §II.4]; in particular, we have $H_0(\Gamma, \mathbb{Z}) = \mathbb{Z}$. By the Mayer-Vietoris sequence [5, Corollary II.7.7], we have

$$\mathbb{Z} \cong H_2(\Gamma, \mathbb{Z}) \cong H_2(\Gamma_1 * \Gamma_2, \mathbb{Z}) \cong H_2(\Gamma_1, \mathbb{Z}) \oplus H_2(\Gamma_2, \mathbb{Z})$$

so say $H_2(\Gamma_2, \mathbb{Z}) = 0$; but this immediately implies Γ_2 is trivial as well, and the result now follows. \square

Algorithm 5.7. Let P be a side pairing for D and let M be a set of minimal cycles for D . This algorithm returns a minimal set of generators and relations for Γ .

1. Let $H \subset G(P)$ be such that $g \in H$ implies either $g = g^{-1}$ or $g^{-1} \notin H$.
2. Let R be the set of elliptic cycles in M and let A be the set of accidental cycles. Initialize r to be an element of A and remove r from A .
3. If $A = \emptyset$, add r to R and return the generators H and the relations R . Otherwise, choose an element $g \in A$ such that g and r have an element $g_i \in H$ in common; then solve for g_i , substitute this expression in for g_i in the relation r , and remove g_i from A . Return to Step 3.

Proof of correctness. If in Step 3 there is always an element $g \in A$ such that g and r have an element in common, then the algorithm terminates correctly: in the notation of (5.1–5.2), there are exactly $t + 1$ relations, and hence the set of generators must also be minimal.

So suppose otherwise. Let H_1 be the set of $g \in H$ such that g or g^{-1} occurs in the relation r and let $H_2 = H \setminus H_1$. Let Γ_1, Γ_2 be the groups generated by H_1, H_2 . Then by assumption, Γ is the free product of Γ_1 and Γ_2 . By Lemma 5.6, since the relation in Γ_1 is nontrivial, it follows that Γ_2 is the free product of finite cyclic groups, and hence cannot contain any accidental cycles, which is a contradiction. \square

The minimal presentation resulting from Algorithm 5.7 is not necessarily of the form (5.1)–(5.2); we refer to the methods of Imbert [14] for an

alternative approach using fat graphs which computes such a canonical presentation.

This completes the proof of the theorem and the accompanying corollaries in the introduction.

Remark 5.8. If in the first corollary, one wants the structure of \mathcal{O}^* , we use the exact sequence

$$1 \rightarrow \mathbb{Z}_F^{*2} \mathcal{O}_1^* \rightarrow \mathcal{O}^* \xrightarrow{\text{nrd}} \mathbb{Z}_{F,+}^* / \mathbb{Z}_F^{*2} \rightarrow 1$$

where $\mathbb{Z}_{F,+}^* = \{u \in \mathbb{Z}_F^* : v(u) > 0 \text{ for all ramified places } v \mid \infty\}$. From the solution to the word problem, it then suffices to find elements $\gamma \in \mathcal{O}^*$ such that $\text{nrd}(\gamma) = u$ generates the finite group $\mathbb{Z}_{F,+}^* / \mathbb{Z}_F^{*2}$, and these can be found using the methods of §3.

6. Examples

We have implemented a variant of the above algorithm in the computer system **Magma** [4]. In this section, we provide two examples of the output of this algorithm.

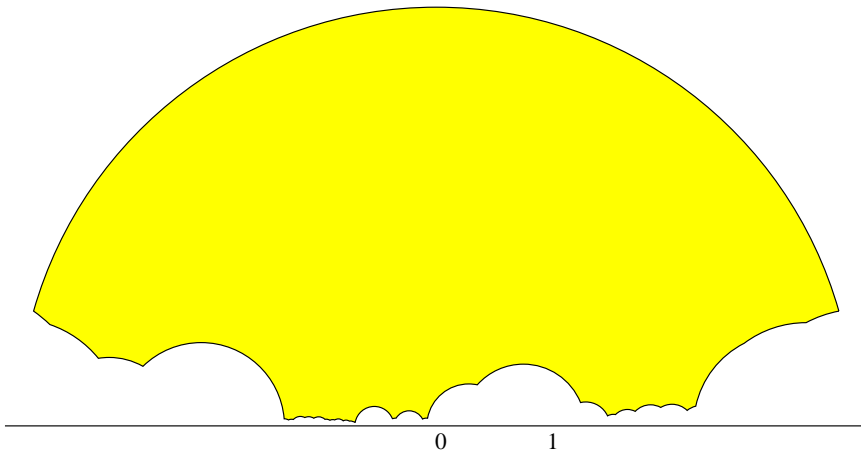


Figure 6.1: A Dirichlet domain for the arithmetic Fuchsian group $\Gamma_0^6(13)$

First, we consider the quaternion algebra $B = \left(\frac{3, -1}{\mathbb{Q}} \right)$ of discriminant 6. A maximal order \mathcal{O} is given by

$$\mathcal{O} = \mathbb{Z} \oplus \mathbb{Z}\alpha \oplus \mathbb{Z}\beta \oplus \mathbb{Z} \frac{1 + \alpha + \beta + \alpha\beta}{2}.$$

We consider the Eichler order contained in \mathcal{O} of level 13, given by

$$\begin{aligned} \mathcal{O}(13) = \mathbb{Z} \oplus \mathbb{Z} \frac{3 - 5\alpha - 5\beta + 3\alpha\beta}{2} \oplus \mathbb{Z}(2 - 2\alpha - \beta + \alpha\beta) \\ \oplus \mathbb{Z} \frac{13 - 13\alpha - 13\beta + 13\alpha\beta}{2}. \end{aligned}$$

We denote $\Gamma(\mathcal{O}) = \Gamma_0^6(13)$. We embed $B \hookrightarrow M_2(\mathbb{R})$ by the embedding (3.4), and take $p = 9i/10 \in \mathfrak{H}$. By (3.1), we compute that the Fuchsian group $\Gamma_0^6(13)$ has coarea $14/3$.

Step 2 in Algorithm 4.9 finds the units $(1 - \alpha - 3\beta + \alpha\beta)/2, \alpha - 2\beta, \dots$, and following the algorithm, reduction and further enumeration automatically yields the fundamental domain as in Figure 6.1. (The methods in Magma for producing the postscript graphic are due to Helena Verrill [22].)

This domain already exhibits significant complexity: it has 38 sides and hence 19 side-pairing elements, which yields a set of 10 minimal generators $\gamma_1, \dots, \gamma_{10}$ for $\Gamma_0^6(13)$, namely

$$\begin{aligned} 12 - 7\alpha + 4\beta + 2\alpha\beta, (1 - \alpha - 33\beta - 19\alpha\beta)/2, 2\alpha + 16\beta + 9\alpha\beta, \\ (37 - 19\alpha + 9\beta + 11\alpha\beta)/2, 2\alpha + 4\beta + \alpha\beta, (1 - \alpha - 3\beta + \alpha\beta)/2, \\ \alpha - 2\beta, (1 + 7\alpha - 15\beta - 5\alpha\beta)/2, (1 + 7\alpha - 45\beta - 25\alpha\beta)/2, \alpha - 14\beta - 8\alpha\beta, \end{aligned}$$

subject to the relations

$$\begin{aligned} \gamma_3^2 = \gamma_5^2 = \gamma_7^2 = \gamma_{10}^2 = \gamma_2^3 = \gamma_6^3 = \gamma_8^3 = \gamma_9^3 = 1 \\ \gamma_1^{-1} \gamma_4 \gamma_5 \gamma_6^{-1} \gamma_1 \gamma_2^{-1} \gamma_3 \gamma_4^{-1} \gamma_7 \gamma_8^{-1} \gamma_9^{-1} \gamma_{10}^{-1} = 1. \end{aligned}$$

We deduce that $\Gamma_0^6(13)$ has signature $(1; 2, 2, 2, 2, 3, 3, 3, 3; 0)$, a fact which can be independently verified by well-known formulae [1].

Second, we consider the totally real number field F generated by a root t of the polynomial $x^7 - x^6 - 6x^5 + 4x^4 + 10x^3 - 4x^2 - 4x + 1$; it is the minimal septic totally real field, having discriminant $d_F = 20134393 = 71 \cdot 283583$. We consider the quaternion algebra B which is ramified at 6 of the 7 real places of F and no finite place: explicitly, $B = \left(\frac{h, k}{F} \right)$ where $h = -t^6 + 6t^4 + t^3 - 9t^2 - 3t + 1$ and $k = -t^2 + 2t - 1$, and in fact $h, k \in \mathbb{Z}_F^*$. We compute a maximal order \mathcal{O} of B . Letting $\Gamma = \Gamma(\mathcal{O})$, we see that Γ has coarea $5/2$. The output of Algorithm 4.9 in this case is given in Figure 6.2; we find that Γ has signature $(0; 2, 2, 2, 2, 2, 3, 3, 3; 0)$.

We conclude by noting that it would be interesting to extend the methods in this paper to other arithmetic groups; this would allow the computation of unit groups for a wider range of quaternion algebras over number fields and would have further consequences for the algorithmic theory of Shimura varieties.




Figure 6.2: A Dirichlet domain for the arithmetic Fuchsian group Γ associated to a quaternion algebra over the minimal septic totally real field

References

- [1] M. ALSINA and P. BAYER, *Quaternion orders, quadratic forms, and Shimura curves*, CRM monograph series, vol. 22, AMS, Providence, 2004.
- [2] A. BEARDON, *The geometry of discrete groups*, Grad. Texts in Math., vol. 91, Springer-Verlag, New York, 1995.
- [3] H.-J. BOEHM, *The constructive reals as a Java library*, J. Log. Algebr. Program. **64** (2005), 3–11.
- [4] W. BOSMA, J. CANNON, and C. PLAYOUST, *The Magma algebra system. I. The user language.*, J. Symbolic Comput., **24** (3–4), 1997, 235–265.
- [5] K. S. BROWN, *Cohomology of groups*, Grad. Texts in Math., vol. 87, Springer-Verlag, New York, 1982.
- [6] H. COHEN, *A course in computational algebraic number theory*, Grad. Texts in Math., vol. 138, Springer-Verlag, New York, 1993.
- [7] H. COHEN, *Advanced topics in computational algebraic number theory*, Grad. Texts in Math., vol. 193, Springer-Verlag, Berlin, 2000.
- [8] D. COX, J. LITTLE, and D. O’SHEA, *Ideals, varieties, and algorithms: An introduction to computational algebraic geometry and commutative algebra*, 2nd ed., Undergrad. Texts in Math., Springer-Verlag, New York, 1997.
- [9] T. DOKCHITSER, *Computing special values of motivic L-functions*, Experiment. Math. **13** (2004), no. 2, 137–149.
- [10] U. FINCKE and M. POHST, *Improved methods for calculating vectors of short length in a lattice, including a complexity analysis*, Math. Comp. **44** (1985), no. 170, 463–471.
- [11] L. R. FORD, *Automorphic functions*, 2nd. ed., Chelsea, New York, 1972.
- [12] I.M. GEL’FAND, M.I. GRAEV, and I.I. PYATETSKII-SHAPIRO, *Representation theory and automorphic functions*, trans. K.A. Hirsch, Generalized Functions, vol. 6, Academic Press, Boston, 1990.
- [13] P. GOWLAND and D. LESTER, *A survey of exact computer arithmetic*, in Computability and Complexity in Analysis, Lecture Notes in Computer Science, eds. Blanck et al., vol. 2064, Springer, 2001, 30–47.
- [14] M. IMBERT, *Calculs de présentations de groupes fuchsien via les graphes rubanés*, Expo. Math. **19** (2001), no. 3, 213–227.

- [15] S. JOHANSSON, *On fundamental domains of arithmetic Fuchsian groups*, Math. Comp **69** (2000), no. 229, 339–349.
- [16] S. KATOK, *Fuchsian groups*, Chicago Lect. in Math., U. of Chicago Press, Chicago, 1992.
- [17] S. KATOK, *Reduction theory for Fuchsian groups*, Math. Ann. **273** (1986), no. 3, 461–470.
- [18] D. R. KOHEL and H. A. VERRILL, *Fundamental domains for Shimura curves*, Les XXIIèmes Journées Arithmétiques (Lille, 2001), J. Théor. Nombres Bordeaux **15** (2003), no. 1, 205–222.
- [19] B. MASKIT, *On Poincaré’s theorem for fundamental polygons*, Advances in Math. **7** (1971), 219–230.
- [20] M.B. POUR-EL and J.I. RICHARDS, *Computability in analysis and physics*, Perspect. in Math. Logic, Springer, Berlin, 1989.
- [21] H. SHIMIZU, *On zeta functions of quaternion algebras*, Ann. of Math. (2) **81**, 1965, 166–193.
- [22] H. VERRILL, *Subgroups of $\mathrm{PSL}_2(\mathbb{R})$* , Handbook of Magma Functions, eds. John Cannon and Wieb Bosma, Edition 2.14 (2007).
- [23] M.-F. VIGNÉRAS, *Arithmétique des algèbres de quaternions*, Lect. Notes in Math., vol. 800, Springer, Berlin, 1980.
- [24] J. VOIGHT, *Quadratic forms and quaternion algebras: algorithms and arithmetic*, Ph.D. Thesis, University of California, Berkeley, 2005.
- [25] K. WEIHRAUCH, *An introduction to computable analysis*, Springer-Verlag, New York, 2000.

John VOIGHT
Department of Mathematics and Statistics
16 Colchester Avenue
University of Vermont
Burlington, Vermont 05401-1455
USA
E-mail: jvoight@gmail.com
URL: <http://www.cems.uvm.edu/~voight/>