

Geometry Seminar Notes

Muye

19 May 2024

1 Dense geodesics: an outline

Before discussing Section 5.4 of the textbook, let us summarize the objective. The main line of the book constructs surfaces by quotienting group actions. Here, as in Section 2.6, we instead study a property of the resulting surface. We want a hyperbolic analogue of a dense straight line on a Euclidean cylinder. The hyperbolic statement is stronger than pointwise density: we seek a geodesic L such that, for every geodesic segment σ and every $\varepsilon > 0$, some segment $L_{\sigma,\varepsilon} \subset L$ remains within perpendicular distance ε of σ along its entire length.

Why should this be possible? Near the cusp of a complete pseudosphere, or punctured sphere, a geodesic can exhibit behavior for which Euclidean intuition is poor. A geodesic on $\mathbb{H}^2/\Gamma$ must therefore be lifted to $\mathbb{H}^2$. On a Euclidean torus, two lifts represent the same line precisely when they differ by an integral horizontal or vertical translation. The group used for the punctured hyperbolic sphere is less arithmetically transparent, so an equally explicit coordinate description seems unlikely.

For other groups, however, arithmetic descriptions become remarkably effective. Artin's first theorem about dense geodesics on a modular quotient appeared in 1924 [1]. The account below also uses Series's later explanations and extensions [3, 2, 4].

For the punctured sphere, the right replacement for coordinates is a symbolic description. We label the sides of a tessellation by generators of Γ and record, in order, the sides crossed by a

lift. Two bi-infinite sequences are regarded as the same if one is a shift of the other: there is no distinguished zeroth crossing.

2 Labelled geodesics and frontier sequences

The frontier sequence does more than distinguish geodesics modulo Γ : every admissible sequence is realized by a geodesic.

Def. 2.1 (Frontier sequence). Fix oriented side labels g, b on the ideal quadrilateral tessellation of $\mathbb{H}^2$ associated with the punctured sphere. The *frontier sequence* of an oriented geodesic $L \subset \mathbb{H}^2 / \Gamma$ (other than the two labelled geodesics themselves) is the bi-infinite sequence

$$\dots, c_{-1}, c_0, c_1, \dots, \quad c_i \in \{g, b, g^{-1}, b^{-1}\},$$

that records the labelled sides crossed by a lift of L . Crossing against the labelled orientation contributes the inverse symbol. The sequence is *reduced* if no adjacent terms are inverses.

Lem. 2.2. Every nonconstant reduced bi-infinite sequence in the alphabet $\{g, b, g^{-1}, b^{-1}\}$ is the frontier sequence of an oriented geodesic on the punctured sphere. Two frontier sequences that differ by a shift describe the same oriented geodesic on the quotient.

We first explain the labels. In the upper half-plane, the two distinguished geodesics lift to the lines $\operatorname{Re} z = -\frac{1}{2}$ and $|z - \frac{1}{2}| = 1$, together with all their Γ -translates. Label the four sides of a fundamental region R as follows: if $x \in \{g, b, g^{-1}, b^{-1}\}$ maps R across a side ℓ to the neighboring region R' , label ℓ by x and direct its transverse arrow from R to R' . Reversing arrows allows the visible labels to be written as g and b . Because elements of $\operatorname{PSL}(2, \mathbb{R})$ preserve orientation, the induced side orientations are consistent throughout the tessellation.

Proof of Lemma 2.2. Choose a region C_0 , say the fundamental region. Reading the positive indices of the sequence determines successive adjacent regions $C_1, C_2, \dots$; reading the negative indices determines $C_{-1}, C_{-2}, \dots$. Reducedness prevents immediate backtracking.




FIGURE 5.8.

Figure 1: A schematic labelled tessellation. Every ideal quadrilateral has one incoming and one outgoing g -side, and similarly for h .

Pass to the Klein projective disk by the isometry $z \mapsto (iz + 1)/(z + i)$. Every C_i becomes a convex ideal quadrilateral. The vertices of these quadrilaterals lie on the boundary circle, and the Euclidean lengths of the successive crossed sides tend to zero as $i \rightarrow \pm\infty$. Thus the positive and negative chains converge to boundary points A_+ and A_- .

If $A_- \neq A_+$, convexity shows that the chord joining them lies in the union of the C_i and crosses precisely the prescribed sides. The nonconstant hypothesis guarantees $A_- \neq A_+$: only an eventually constant choice at both ends can force both chains toward the same ideal vertex. Therefore the chord is the required geodesic.

Once C_0 is fixed, this construction is unique. Replacing C_0 by another region applies an element of Γ and changes only the index origin. Hence shifted sequences describe the same quotient geodesic. $\square$




FIGURE 5.9.

Figure 2: The projective-disk construction of the realizing geodesic.

3 Existence of a dense geodesic

Thm. 3.1. On the punctured sphere there is an oriented geodesic L such that, for every geodesic segment σ and every $\varepsilon > 0$, a segment $L_{\sigma, \varepsilon} \subset L$ has perpendicular distance less than ε from σ at every corresponding point.

Proof. Lift σ to a segment $\tilde{\sigma} \subset \mathbb{H}^2$ and extend it to a complete geodesic. Let $(c_i)_{i \in \mathbb{Z}}$ be its frontier sequence. In the projective disk, the crossed sides shrink toward the two endpoints. Consequently, for sufficiently large n , every geodesic joining the two boundary sides of C_{-n} and C_n remains within perpendicular distance ε of $\tilde{\sigma}$ along the intervening segment. Thus it is enough that another frontier sequence contain the block

$$c_{-n}, c_{-n+1}, \dots, c_n.$$

For each length N , there are only finitely many reduced nonconstant words of length N in our four-letter alphabet. The union over N is countable. Concatenate an enumeration of all these

finite words, inserting short reduced separators when necessary, and extend in both directions. The resulting nonconstant reduced bi-infinite sequence contains every finite reduced word. By Lemma 2.2, it determines a geodesic L , and the preceding paragraph proves the required density property. $\square$

4 The modular group and its quotient

Artin's original setting uses the modular group

$$\mathcal{G} = \mathrm{PSL}(2, \mathbb{Z}) \subset \mathrm{PSL}(2, \mathbb{R}).$$

This is the most natural arithmetic discrete subgroup of the orientation-preserving isometry group of $\mathbb{H}^2$. The quotient $\mathbb{H}^2/\mathcal{G}$ is an orbifold rather than a hyperbolic surface, because $\mathcal{G}$ contains elliptic elements with fixed points. Thus Artin's theorem does not literally produce a dense line on a smooth hyperbolic surface, although its symbolic and arithmetic structure is exceptionally clear.

Two standard fundamental regions are the Dirichlet and Ford regions. The following familiar generating result will be used repeatedly.

Thm. 4.1. The group $\mathrm{SL}(2, \mathbb{Z})$ is generated by

$$T = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \quad (1)$$

In fractional-linear notation, $T(z) = z + 1$ and $J(z) = -1/z$.

5 Continued fractions

We now describe the $\mathcal{G}$ -orbits of boundary points. Write

$$x = [n_0; n_1, n_2, \dots]$$

for the regular continued fraction of an irrational real number, using the usual harmless convention for negative numbers after applying $x \mapsto -1/x$.




FIGURE 7.19.

Figure 3: The standard Dirichlet region for the modular group.

Thm. 5.1. For irrational real numbers x, y , there exists $g \in \mathcal{G}$ with $g(x) = y$ if and only if their continued-fraction expansions have a common tail with the appropriate parity: if

$$x = [n_0; n_1, n_2, \dots], \quad y = [m_0; m_1, m_2, \dots],$$

then there are $k, \ell \in \mathbb{N}$ with $k + \ell$ even and

$$n_{k+r} = m_{\ell+r} \quad (r \geq 0).$$

The idea is visible from the generators. Applying T^{-n_0} removes the integer part, and applying J replaces the remainder by its negative reciprocal. Iteration removes successive continued-fraction entries. The determinant-one condition accounts for the parity. Conversely, writing an element of the modular group as a word in J and powers of T retraces exactly these operations.

A geodesic in $\mathbb{H}^2$ is determined by its two endpoints. After applying an element of $\mathcal{G}$, arrange them as $\eta \in (-1, 0]$ and $\xi \in [1, \infty)$. If

$$-1/\eta = [b_0; b_1, b_2, \dots], \quad \xi = [a_0; a_1, a_2, \dots],$$

associate the bi-infinite sequence

$$\dots, b_2, b_1, b_0, a_0, a_1, a_2, \dots \tag{2}$$




Figure 7. The tessellation Γ subdivided into fundamental regions for $SL(2, \mathbb{Z})$.

Figure 4: A Ford region inside the modular tessellation.

Thm. 5.2 (Artin). Two geodesics $L, L' \subset \mathbb{H}^2$ project to the same oriented geodesic on $\mathbb{H}^2/\mathcal{G}$ if and only if their sequences (2) agree up to a shift.

Later continued-fraction entries have decreasing influence on the endpoint. Therefore a bi-infinite sequence containing every finite sequence of positive integers codes a geodesic that shadows every finite geodesic segment. This is the arithmetic counterpart of Theorem 3.1.

6 The Farey tessellation

There is a geometric development parallel to the continued-fraction argument.

Def. 6.1. Two reduced fractions p/q and r/s are *Farey neighbors* if

$$|ps - rq| = 1.$$

Their *Farey sum* is

$$\frac{p}{q} \oplus \frac{r}{s} = \frac{p+r}{q+s}.$$

The Farey tessellation is obtained by joining every pair of Farey neighbors by the hyperbolic geodesic with those endpoints.

Begin with the vertical geodesics through consecutive integers. Insert the mediants $(2n+1)/2$ between n and $n+1$, join each median to its two parents, and continue recursively. Equivalently, this tessellation is generated by a subgroup of the modular group acting on one ideal triangle.

Lem. 6.2. The Farey tessellation has the following properties.

- (1) Every pair of Farey neighbors is joined by an edge.
- (2) Every rational boundary point is a vertex of an ideal triangle.
- (3) Every modular-group element sends each Farey triangle to another Farey triangle.




Figure 2: Fundamental domain for the action of $SL(2, \mathbb{Z})$ on $\mathbb{H}$.

Figure 5: Farey tessellation with the beginning of a cutting sequence.

7 Cutting sequences

Def. 7.1. Let γ be an oriented curve beginning on the imaginary axis and ending at a boundary point $x > 0$. Whenever γ crosses an ideal triangle of the Farey tessellation, record L if the common vertex of the two crossed sides lies to the left of γ , and R if it lies to the right. The resulting word

$$L^{n_0} R^{n_1} L^{n_2} \dots$$

is the *cutting sequence* of γ , or equivalently of x .

The starting point on the imaginary axis does not affect the sequence.

Thm. 7.2. For $x > 0$, the cutting sequence is

$$L^{n_0} R^{n_1} L^{n_2} \dots$$

if and only if

$$x = [n_0; n_1, n_2, \dots],$$

where $n_0 \geq 0$ and $n_i > 0$ for $i \geq 1$.

Proof. If $n_0 \leq x < n_0 + 1$, the first n_0 triangles lie on the same side, giving $L^{n_0} R$. Apply T^{-n_0} and then J . Modular transformations preserve the tessellation and its orientation; J fixes the imaginary axis while T^k changes only the starting position by k triangles. The transformed endpoint is the negative reciprocal of $x - n_0$, so the same argument reads off n_1 , then n_2 , and so on. This is exactly the continued-fraction algorithm, and the converse retraces the construction. $\square$

This theorem connects an arithmetic expansion to a purely geometric crossing pattern. Frontier sequences are more flexible because their construction uses little beyond an ideal polygonal fundamental region. There is, however, an important restriction: all vertices of the relevant quadrilaterals must lie at infinity. In the Klein disk, the regions must be inscribed in the boundary circle; otherwise a reduced nonconstant symbolic sequence need not converge to two distinct boundary points. For example, the standard Ford tessellation of the modular group contains interior elliptic vertices and admits formal alternating words that no geodesic realizes.

One can pass to a suitable finite-index subgroup whose ideal quadrilaterals do reach the boundary, but its frontier sequence need not be invariant under the full modular group. Tracking how the larger group changes adjacent symbols leads precisely to the L/R cutting sequence; see [4]. The punctured torus provides another fundamental ideal-quadrilateral example, and its refined frontier sequences lead to remarkable Diophantine phenomena.




Figure 11. Curves containing $AB^3\bar{A}$ cut H .

Figure 6: A cutting-pattern model associated with the punctured torus.

8 A glimpse of Diophantine approximation

We briefly indicate one application unrelated to the textbook's main line. Consider the horizontal Euclidean line, a horocycle,

$$\ell_b = \{z \in \mathbb{H}^2 : \text{Im } z = b\}.$$

Let

$$\phi(z) = \frac{az + b}{cz + d}, \quad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}(2, \mathbb{R}), \quad c \neq 0.$$

Writing $w = \phi^{-1}(z) = (dz - b)/(-cz + a)$ and imposing $\text{Im } w = b$ yields

$$\left| z - \left(\frac{a}{c} + \frac{i}{2bc^2} \right) \right| = \frac{1}{2bc^2}. \quad (3)$$

Thus the image is a Euclidean circle tangent to the real axis, centered at

$$\frac{a}{c} + \frac{i}{2bc^2}.$$

The appearance of the rational a/c and the scale c^{-2} suggests Diophantine approximation.

Thm. 8.1 (Dirichlet). For every irrational θ there are infinitely many pairs $(p, q) \in \mathbb{Z} \times \mathbb{N}$ such that

$$\left| \theta - \frac{p}{q} \right| < \frac{1}{q^2}. \quad (4)$$

Define the optimal recurrent constant

$$\nu(\theta) = \inf \left\{ c : \left| \theta - \frac{p}{q} \right| < \frac{c}{q^2} \text{ for infinitely many } (p, q) \right\}. \quad (5)$$

Hurwitz's theorem gives the universal constant $1/\sqrt{5}$, attained by the golden ratio $(1 + \sqrt{5})/2$. In continued-fraction language, equality occurs for numbers whose tails eventually agree with the all-ones tail. This exceptional set is small. Relaxing the extremal requirement produces the Markoff spectrum, whose geometry is encoded by geodesics and cutting sequences on the punctured torus; see [2].

9 Constructing surfaces: the logical framework

The chapter aims to understand hyperbolic surfaces, that is, surfaces locally isometric to $\mathbb{H}^2$. One route uses the Killing–Hopf theorem and writes a complete surface as $\mathbb{H}^2/\Gamma$, where Γ acts freely and properly discontinuously. Direct classification of such groups is complicated in hyperbolic geometry, and completeness must first be established.

A second route returns to the intuitive construction by gluing sides of polygons. The chapter develops these routes in parallel:

$$\text{polygon gluing} \quad \longleftrightarrow \quad \text{quotient by a discrete group.}$$

For compact polygons, completeness lets one pass from the gluing construction back to a quotient via Killing–Hopf. A more elementary connection holds whenever the polygon itself is a fundamental region.

Thm. 9.1. Let $\Pi \subset \mathbb{H}^2$ be a polygon whose paired sides are identified by isometries $t_1, \dots, t_s$, and put $\Gamma = \langle t_1, \dots, t_s \rangle$. If the interior of Π is a fundamental region for Γ , then the metric gluing space S_Π is isometric to $\mathbb{H}^2/\Gamma$.

Li Linjie. The gluing metric defines the distance between two classes by taking the infimum of lengths of polygonal paths. A representative path from A to B may be written

$$AW_1, W_1'W_2, \dots, W_{n-1}'W_n, W_n'B,$$

where W_i and W_i' are paired boundary points. After suitable indexing, $t_i W_i = W_i'$. Apply t_1 to the second segment, $t_1 t_2$ to the third, and so forth. The transformed segments meet end to end in $\mathbb{H}^2$ and have the same total length L . Hence

$$L \geq d_{\mathbb{H}^2}(A, t_1 \cdots t_n B) \geq d_{\mathbb{H}^2/\Gamma}(\Gamma A, \Gamma B). \quad (6)$$

Taking the infimum over gluing paths gives one inequality. Conversely, a geodesic from A to a translate γB can be cut at the sides of the tessellation and folded back into Π , producing a gluing path of the same length. Thus equality holds. $\square$

This theorem also makes the group constructive: it is generated by the side pairings. Compactness of Π is not necessary; for example, the punctured sphere has a noncompact ideal fundamental region but is complete.

For a discrete subgroup of $\mathrm{PSL}(2, \mathbb{R})$, called a *Fuchsian group*, the quotient need not be a surface; it may be an orbifold. Finite generation is the natural condition governing polygonal fundamental regions: a finitely generated Fuchsian group has finite-sided fundamental polygons, subject to the usual technical qualifications on sides. Conversely, the Poincaré polygon theorem gives conditions under which a side-paired hyperbolic polygon is a fundamental polygon for the discrete group generated by its pairings, and it reads the group relations from the vertex cycles [7, Chs. 9–10].

10 Glued hyperbolic surfaces

Suppose several polygon vertices are identified in a vertex cycle. If the total angle there is θ , a hyperbolic circle of radius r around the glued point has length

$$\theta \sinh r. \tag{7}$$

Unless $\theta = 2\pi$, this differs from the circumference $2\pi \sinh r$ of a genuine hyperbolic disk. The glued point is a cone point rather than a smooth surface point.

It is useful to distinguish a *side*, a segment actually paired by a gluing isometry, from an *edge*, a maximal geodesic segment in the polygonal boundary. A single infinite edge may contain infinitely many paired sides, so “finitely many edges” need not imply that the side-pairing group is finitely generated.

Thm. 10.1. Let Π be a hyperbolic polygon with side identifications. If the sum of the interior angles in every vertex cycle is 2π , then the gluing space S_Π carries a metric locally equal to the standard metric on $\mathbb{H}^2$; with this metric it is a hyperbolic surface.

The proof constructs charts in polygon interiors, across paired sides, and around vertex cycles. The 2π condition makes the last charts genuine disks instead of cones.

11 Fundamental regions of groups

Thm. 11.1. For every compact hyperbolic surface $\mathbb{H}^2/\Gamma$, the group Γ admits a polygonal fundamental region.

The textbook proves this using a Dirichlet region. In the surface case, Γ acts freely. This matters: the purported bisector

$$L_g = \{z \in \mathbb{H}^2 : d(z, c) = d(z, gc)\} \tag{8}$$

is a genuine geodesic only when $gc \neq c$. For an orbifold group, one must choose the center c outside every nontrivial fixed-point set.

We finish by returning to the modular group and proving the standard region directly.

Thm. 11.2. Let $c = 2i$ (more generally $c = ki$ with $k > 0, k \neq 1$). The region

$$R = \left\{ z \in \mathbb{H}^2 : |\operatorname{Re} z| < \frac{1}{2}, |z| > 1 \right\} \quad (9)$$

is a Dirichlet fundamental region for $\mathcal{G} = \operatorname{PSL}(2, \mathbb{Z})$.

Proof. The bisectors associated with J, T, T^{-1} are respectively

$$|z| = 1, \quad \operatorname{Re} z = -\frac{1}{2}, \quad \operatorname{Re} z = \frac{1}{2},$$

so the Dirichlet region is contained in R .

It remains to show $R \cap gR = \emptyset$ for $g \neq 1$. Suppose $z \in R$ and $g^{-1}z \in R$, and write

$$g^{-1} = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

If $c \neq 0$, then

$$\begin{aligned} |cz + d|^2 &= c^2 |z|^2 + 2cd \operatorname{Re} z + d^2 \\ &> c^2 - |cd| + d^2 = (|c| - |d|)^2 + |cd| \geq 1. \end{aligned} \quad (10)$$

Hence

$$\operatorname{Im}(g^{-1}z) = \frac{\operatorname{Im} z}{|cz + d|^2} < \operatorname{Im} z.$$

Applying the same argument to $g^{-1}z \in R$ and g gives the reverse strict inequality, a contradiction.

If $c = 0$, then g is an integral translation, and distinct translates of the strip $|\operatorname{Re} z| < 1/2$ are disjoint. Thus R is the required region. $\square$

Finally, consider the Ford region. Put

$$A = JT^{-1} = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}.$$

In $\operatorname{PSL}(2, \mathbb{Z})$, J has order 2 and A has order 3. The Poincaré polygon theorem shows that these are the only relations required:

$$\operatorname{PSL}(2, \mathbb{Z}) \cong \langle g, b \mid g^2 = b^3 = 1 \rangle. \quad (11)$$

Geometrically, the orders record the vertex cycles of angles π and $2\pi/3$. They also explain why the quotient is an orbifold rather than a smooth hyperbolic surface. This completes the discussion of dense geodesics, glued surfaces, fundamental regions, and the modular-group supplement.

References

- [1] E. Artin, Ein mechanisches System mit quasi-ergodischen Bahnen, *Abh. Math. Sem. Univ. Hamburg* **3** (1924), 170–175, doi:10.1007/BF02954622.
- [2] C. Series, The geometry of Markoff numbers, *Math. Intelligencer* **7** (1985), no. 3, 20–29, doi:10.1007/BF03025802.
- [3] C. Series, Non-Euclidean geometry, continued fractions and ergodic theory, *Math. Intelligencer* **4** (1982), no. 1, 24–31.
- [4] C. Series, *Continued Fractions and Hyperbolic Geometry*, lecture notes for the London Mathematical Society Summer School, Loughborough University.
- [5] J. Stillwell, *Geometry of Surfaces*, Universitext, Springer, New York, 1992, doi:10.1007/978-1-4612-0929-4.
- [6] C. Series, *Hyperbolic Geometry*, MA448 lecture notes, University of Warwick.
- [7] A. F. Beardon, *The Geometry of Discrete Groups*, Graduate Texts in Mathematics 91, Springer, New York, 1983, doi:10.1007/978-1-4612-1146-4.