Junior Seminar: Hyperbolic Geometry
Lecture Notes

Tim Campion

January 20, 2010

1 Motivation

Our first construction is very similar in spirit to an analogous one in Euclidean space. The group of isometries of Euclidean $\mathbb{R}^n$ is given by $O(n) \oplus \mathbb{R}^n$ or $SO(n) \oplus \mathbb{R}^n$ if we want to preserve orientation; the first factor is made up of rotations and reflections (in the non-orientation-preserving case) about the origin, while the second factor gives translations. Now, it turns out that $O(n)$ is generated by reflections. What is a reflection?

Definition 1 (Euclidean Reflection)**.** A reflection $r_V : \mathbb{R}^n \to \mathbb{R}^n$ in Euclidean space has a $n-1$-dimensional space V (the "mirror") on which it acts as the identity; it multiplies the orthogonal complement $V^\perp$ of V by -1.

$$r_V(w) = -\operatorname{pr}_{V^\perp}(w) + \operatorname{pr}_V(w)$$

where $\operatorname{pr}_V, \operatorname{pr}_{V^\perp}$ are the orthogonal projections onto $V, V^\perp$ repectively. Note that since $V^\perp$ is 1-dimensional, the reflection can alternatively be uniquely specified by a unit vector $v \in V^\perp$, and using a parallel notation, we can call it r_v.

Notice that a reflection is always orientation-reversing. If we compose two reflections, though, we get a rotation. Try visualizing it in two dimensions: the composition $r_v \circ r_{v'}$ rotates through an angle which is twice the angle between v and v'. Now, if we allow ourselves to conjugate by translations, i.e. using such transformations as $r_v^T = T \circ r_v \circ T^{-1}$ where T is an arbitrary translation, then the composition of two such isometries $r_v^T \circ r_{v'}^{T'}$ will just be a rotation-plus-translation, and in fact, we can create any isometry of $\mathbb{R}^n$ in this way. You can prove this if you want: start with a Euclidean isometry, and successively make it fix more vectors by subtracting off the image of a vector it doesn't fix, using a reflection. In this way, we can express any Euclidean $\mathbb{R}^n$ isometry as the composition of at most n reflections, subject to the constraint that we need an even number of them to have an orientation-preserving isometry. If we choose some clever normalization, then we can even choose a canonical form for Euclidean isometries in terms of reflections.
We will do just that in the hyperbolic plane.

2 Hyperbolic Analog

If we want to imitate this Euclidean decomposition, we're going to need some notion of a reflection in hyperbolic space. In two dimensions, this means we need some analog of a $2-1=1$-dimensional vector space, the fixed space of the reflection. Of course, this is just a line, and we already have a notion of a line – a geodesic. This leads us to the following notion.

Definition 2 (Inversion in a Circle)**.** Let C be a circle of radius R with center O. If P is inside the circle C and Q is outside of C, let X be the intersetction between L and C which lies between P and Q. We say that P and Q are inverse to each other in C, if

1. the line PQ passes through O, and
2. $(PX)(QX) = R^2$

Furthermore, we say that P is inverse in C to itself if P lies on C.

Now we have

Proposition 1. *Let C be a circle with center O and radius R. Then for every $P \neq O, \infty$, there is a unique point Q such that P is inverse in C to Q. Moreover, let i_C be the map sending a point to its inverse in C. Then i_C is anticonformal, when extended to exchange 0 and ∞.*
Furthermore, if C is a geodesic circle in $\mathbb{H}$ (resp. $\mathbb{U}$), then i_C maps $\mathbb{H}$ (resp. $\mathbb{U}$) bijectively to itself.

Proof. For a unit circle centered at the origin, the equation saying that $z, z' \neq 0, \infty$ lie on the same radial line is $z = re^{i\theta}, z' = r'e^{i\theta}$, and then the equation saying they are inverse is $1 = |z||z'| = rr'$. There is a unique solution: $z = 1/\bar{z}'$. Any other circle inversion can be attained by the conformal transformations of dilation and translation. When extended to exhange 0 and ∞, this is just a reflection composed with a Möbius transformation, which is anticonformal.
Inversion in the unit circle takes $\mathbb{H}$ to itself because $\Im(\frac{1}{\bar{z}}) = \Im(\frac{z}{|z|^2})$ which is > 0 if $\Im z > 0$. Inversion in any other circle is just a translation along $\mathbb{R}$ and a dilation composed with this map, so it also preserves $\mathbb{H}$. The map is bijective because it is its own inverse. Since the isomorphism between $\mathbb{H}$ and $\mathbb{U}$ preserves geodesics and bijectively maps between $\mathbb{H}$ and $\mathbb{U}$, it follows that the same holds for $\mathbb{U}$. □

Now we need some more-or-less canonically defined circle to invert in, if we want to find a more-or-less unique decomposition of a generic isometry. The following construction depends on what model of hyperbolic space we use, but behaves similarly in $\mathbb{H}$ and $\mathbb{U}$.

Definition 3 (Isometric Circle)**.** Let T be a Möbius transformation acting on the model M, with $M = \mathbb{H}$ or $\mathbb{U}$. The isometric circle $I(T)$ is defined to be the set of points where T locally fixes Euclidean length, i.e. $I(T) = \{z \in M \mid |T'(z)| = 1\}$.

Proposition 2. *In either model, $I(T) = T(I(T^{-1}))$.*

Proof. Let u be a tangent vector of $I(T)$ at the point P. Then the induced map $T_*(u) = T'(p)(u)$ satisfies $|T_*(u)| = |u|$. Since $T_*^{-1} = (T_*)^{-1}$, this gives $|v| = T_*^{-1}(v)$ for $v = T(u)$. Since T is a diffeomorphism, every tangent vector v of the upper half plane can be written uniquely as the image of a tangent vector of $\mathbb{H}$ under T. This shows that $T(I(T)) \subset I(T^{-1})$. By similar reasoning, we have $T^{-1}(I(T^{-1})) \subset I(T)$, and taking the T-image, $I(T^{-1}) \subset T(I(T))$, completing the proof. □

We would like to justify the name "isometric circle," not only because it's a little silly to call something a circle when you don't know it's a circle, but also because being a geodesic circle associated to an isometry is practically the whole point of the isometric circle.

Proposition 3. *For any non-parabolic T acting on $\mathbb{H}$, the set $I_{\mathbb{H}}(T)$ is a geodesic circle. For any non-elliptic T acting on $\mathbb{U}$, the set $I_{\mathbb{U}}(T)$ is a geodesic circle.*
Moreover, in either model, T locally shortens Euclidean lengths at points inside $I(T)$, and lengthens them outside of $I(T)$.
If T is parabolic, then $I_{\mathbb{H}}(T)$ is either $\emptyset$ or $\mathbb{H}$; if T is elliptic, then $I_{\mathbb{U}}(T)$ is $\mathbb{U}$.

Proof. In the upper half plane, let T be the fractional linear transformation associated to the matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in SL_2(\mathbb{R})$. Then the condition of membership in $I(T)$ is $1 = |T'(z)| = \frac{1}{|cz+d|^2}$.
If T is nonparabolic, then $c \neq 0$ so we can rewrite this as $|z + d/c| = 1/|c|$. This is manifestly a circle of radius $1/|c|$ and center $-d/c$; since d, c are real, the center of the circle is in $\mathbb{R}$, and so it is a geodesic circle.
Similarly, in the Poincaré disc, let T be the fractional linear transformation associated to the matrix $\begin{pmatrix} \bar{a} & \bar{c} \\ c & a \end{pmatrix} \in SL_2(\mathbb{C})$. Then the condition of membership in $I_{\mathbb{U}}(T)$ is $1 = |T'(z)| = \frac{1}{|cz+a|}$. This is a circle of radius $1/|c|$ centered at c/a. So the triangle formed by the origin, the center of $I_{\mathbb{U}}(T)$, and intersection of the unit disc and $I(T)$ has side lengths $1, |a/c|, 1/|c|$. We have $1 + 1/|c|^2 = |a/c|^2$ since T has determinant 1, i.e. the triangle is a right triangle, with a right angle at the intersection of the circles. Since both triangle sides are perpendicular to the circles, this means the circles are orthogonal, and hence $I(T)$ is a geodesic.
In either model, T shortens Euclidean lengths where $|T'(z)| > 1$, and this inequality says precisely that z is on the outside of $I(T)$; similarly, the condition $|T'(z)| < 1$ says taht z is outside of $I(T)$.
If T is parabolic, then in the action on $\mathbb{H}$ we have $c = 0$, and so $z \in I(T)$ iff $|d| = 1$. Geometrically, T is a dilation and a translation; $I(T)$ is $\mathbb{H}$ if the dilation is the identity, and $\emptyset$ otherwise.
If T is elliptic, then in the action on $\mathbb{U}$ we have $c = 0$, and so T acts as a rotation; in this case $I(T) = \mathbb{U}$. □

Proposition 4. *If T is a non-parabolic Möbius transformation acting on $\mathbb{H}$ (respectively, a non-elliptic Möbius transformation acting on $\mathbb{U}$), then T can be uniquely decomposed as inversion in $I_{\mathbb{H}}(T)$ (respectively, in $I_{\mathbb{U}}(T)$) followed by reflection in a line L which is perpendicular to $\mathbb{R}$ (the line, of course, depends on model you take the isometric circle in).*

Proof. Most of what we say here works in either model (even though we are not talking about the same circle and line); when this is the case, will try to use model-neutral language. Since T is non-parabolic (resp. non-elliptic), so is T^{-1}, and hence both these sets are geodesic circles. Since T extends to a homeomorphism on the closure of the hyperbolic plane, this extended map sends boundary points to boundary points. In particular, it maps the endpoints of $I(T)$ to the endpoints of $I(T^{-1})$. Since Euclidean lengths are preserved locally on $I(T)$, we can pull back any measurement of Euclidean length along $I(T^{-1})$ to a measurement of length between preimages under T along $I(T)$. Since $I(T^{-1})$ is fully paramaterized by Euclidean arc length measured from one of the endpoints, this ability to pull back fixes the entire transformation T; i.e. any other isometry that sends $I(T)$ to $I(T^{-1})$ and sends the endpoints to the same places will agree with T at at least three points, and hence be equal to T. So first we will find where T sends the endpoints, and then we will construct an isometry that sends $I(T)$ to $I(T^{-1})$ and the endpoints to the same places. By pulling back measurements of length along $I(T)$, we see that $I(T)$ and $I(T^{-1})$ have the same Euclidean circumference, and hence the same radius. In fact, the Euclidean distance between the endpoints is the same or each circle (in the case of $\mathbb{H}$, this is obvious because the center of each circle lies on $\mathbb{R}$. In $\mathbb{U}$, one must note not only that the circumferences are equal, but also the portion of the circumference that lies in $\mathbb{U}$). So if T sends the lefthand endpoint of $I(T)$ to the lefthand endpoint of $I(T^{-1})$, then T is just a translation (resp. a rotation), i.e. it is parabolic (resp. elliptic), contrary to our hypothesis. Hence T sends the left endpoint of $I(T)$ to the right endpoint of $I(T^{-1})$.
Now, we can of course find a unique Euclidean line L which is also a geodesic such that reflection in this line takes the left-hand endpoint of $I(T^{-1})$ to the right-hand endpoint of $I(T)$; by the preceeding sentence (and the fact that the endpoints are on the boundary of the hyperbolic plane), this reflection also takes the right-hand endpoint of $I(T^{-1})$ to the left-hand endpoint of $I(T)$. This reflection anticonformal, and sends everything on $I(T)$ to the same place as T itself does. So if we precompose by inversion in $I(T)$, we get a conformal map which agrees with T on $I(T)$, and hence everywhere (because inversion in $I(T)$ fixes $I(T)$). This completes the proof. □

3 The Ford Fundamental Region

Given a Fuchsian group Γ, we can use isometric circles to find a very nice fundamental region for Γ, called the *Ford Fundamental Region*.

Definition 4. Let Γ be a Fuchsian group. Then the Ford Fundamental Region

F_0 for Γ is defined as

$$F_0 = \bigcap_{T\in\Gamma} I^-_{\mathbb{U}}(T) \cap \mathbb{U}$$

where $I^-(T)$ is meant to indicate the inside of $I(T)$. Of course, we could give an entirely analogous definition using isometric circles in $\mathbb{H}$.

We would like to show that this is a fundamental region.

Theorem 1. *For any isometry T, of $\mathbb{U}$, we have*

1. $\rho(Tz, 0) = \rho(z, 0)$ *if* $z \in I(T)$
2. $\rho(Tz, 0) > \rho(z, 0)$ *if z is outside of* $I(T)$
3. $\rho(Tz, 0) < \rho(z, 0)$ *if z is inside of* $I(T)$

Proof. We know that the hyperbolic distance to 0 in $\mathbb{U}$ increases monotonically with the Euclidean distance, so it suffices to prove the theorem for Euclidean distance.
Any isometry T is composed of inversion in $I(T)$ followed by reflection in a line through 0. The reflection doesn't change the distance to 0, so we need only consider inversion in the circle. We can perform a translation and a dilation, which preserves the proportions of Euclidean distances, to place the circle to be inverted in at the origin and with unit radius; call this circle $I = I(T)$. Then the circle U represeting $\mathbb{U}$ is centered at a point c, which by a rotation we can take to be on the real axis, and has radius R.
Now, for any point z in U, its inverse z' lies along the line $L = \{te^{i\theta} : t \in \mathbb{R}\}$ which passes through z. For a given line L, there are at most two obvious pairs of points which are mapped to points the same distance from 0: two identical points which lie at the intersection of L and I (this obvious observation proves the first point of the theorem), and the two (possibly identical) points which lie at the intersection of L and U.
Then the inversion map is $i_I(z) = \frac{1}{\bar{z}}$. Let $f(z)$ be the squared distance z to c i.e. $f(z) = |z - c|$. The theorem statement (note that the remaining two parts are equivalent) is that $f(z) < f(i_I(z))$ for $z \in U - I$. When restricted to L, we can write $f : te^{i\theta} \mapsto |te^{i\theta} - c|^2$, and $i_I : te^{i\theta} \mapsto t^{-1}e^{-i\theta}$. We can parameterize L by the map $h : t \mapsto te^{i\theta}$. We compute the derivatives:

$$\begin{aligned} f'(te^{i\theta}) &= e^{-i\theta}\frac{d}{dt}\left((te^{i\theta} - c)(te^{-i\theta} - c)\right) \\ &= e^{-i\theta}\left(2\cos\theta(t\cos\theta - c) + 2t\sin^2\theta\right) \\ i_I'(te^{i\theta}) &= -e^{-2i\theta}/t^2 \\ h'(t) &= e^{i\theta} \end{aligned}$$

To prove the theorem, it will suffice to show that $(f \circ i_I \circ h)'(t) < (f \circ h)'(t)$ for $t > 1$ and $te^{i\theta}$ inside the circle U, because $f \circ i_I = f$ on I. So let us calculate:

$$\begin{aligned} (f \circ h)'(t) &= 2\cos\theta(t\cos\theta - c) + 2t\sin^2\theta \\ (f \circ i_I \circ h)'(t) &= \left(2\cos\theta(t^{-1}\cos\theta - c) + 2t^{-1}\sin^2\theta\right)/t^2 \end{aligned}$$

Now, for $t > 1$, the deired inequality holds. $\square$

Recall that given a Fuchsian group Γ and a point p in the hyperbolic plaen, the Dirichlet region $D_p(\Gamma)$ can be equivalently defined in two ways:

$$D_p(\Gamma) = \{z \in \mathbb{H} \mid \forall T \in \Gamma\ \rho(z,p) \leq \rho(z,Tz)\} = \{z \in \mathbb{H} \mid \forall T \in \Gamma\ \rho(z,p) \leq \rho(Tz,p)\}$$

where ρ is the hyperbolic metric; this definition is independent of the choice of model. As Marli proved, a Dirichlet region is always a fundamental region for its Fuchsian group (up to the usual boundary ambiguities). The following theorem should have an analog in $\mathbb{H}$, but we will be satisfied with

Corollary 1. *For any Fuchsian group Γ, the Ford Fundamental Region in $\mathbb{U}$ is a Dirichlet Region. Specifically, $F_0 = D_0(\Gamma)$.*

Proof. It is immediate, since the intersections in the definitions of the two regions agree term-by-term. $\square$

In particular, the Ford fundamental region is in fact a fundamental region.

4 Limit Points

We're interested in what points on the boundary can be approached from within by the action of a Fuchsian group.

Definition 5. Given a Fuchsian group Γ acting on $\mathbb{U}$, define the limit set

$$\Lambda(\Gamma) = \bigcup_{T \in \Gamma, z \in \mathbb{U}} \text{Limit points}(Tz)$$

It is immediately obvious that $\Lambda(\Gamma)$ is Γ-invariant. That is, for each $\alpha \in \Lambda$, we have $z \in \mathbb{U}$ and $T_1, T_2, \dots \in \Gamma$ such that $T_1 z, T_2 z, \dots$ approaches α. So $ST_nS^{-1}(Sz)$ will approach $S\alpha$ for any $S \in \Gamma$, by continuity.
There is a certain genericity to this set Λ; our first toehold in demonstrating this is by showing a relationship to isometric circles.

Definition 6. Given a Fuchsian group Γ acting on $\mathbb{U}$, define

$$\Lambda_0(\Gamma) = \text{Limit points}\left(\bigcup_{T \in \Gamma} C(I(T))\right)$$

where $C(I(T))$ is meant to denote the center of $I(T)$.

Since $\Lambda_0(\Gamma)$ is defined as a limit set in a metric space, it is closed. For if we have a sequence of limit points $x_1, x_2, \dots$ of a set S approaching a point z, then let $s_{n1}, s_{n2}, \dots$ be a sequence in S approaching x_n. By the triangle inequality, $s_{11}, s_{22}, \dots$ approaches z.
Remarkably, these two sets coincide, even though the centers of the isometric circles li *outside* of the closed unit disc! We first need a lemma:

Lemma 1. *For any Fuchsian group Γ acting on $\mathbb{U}$, and for any $\epsilon > 0$, there are at most finitely many $T \in \Gamma$ such that the radius of $I(T)$ is $\geq \epsilon$.*

Proof. If Γ is finite then the lemma is trivial, so assume Γ is infinite.
For any $T \in \Gamma$, we can write T as a matrix $\begin{pmatrix} a & c \\ \bar{c} & \bar{a} \end{pmatrix}$, with $|a|^2 - |c|^2 = 1$. The radius of the isometric circle is $1/|c|$. Suppose that there were infinitely many T with $1/|c| \geq \epsilon$, i.e. $|c| \leq \epsilon$. Then there would be some c_0, and a sequence $T_1, T_2, \dots \in \Gamma$ with "$|c|$-values" $|c_1|, |c_2|, \dots$ approaching c_0. Hence their $|a|$-values would approach $1 + |c_0|$. This would give an infinite set contained in a compact subset of $PSL_2(\mathbb{C})$, which must have a limit point, contradicting the discreteness of Γ. □

This fact resolves the initial puzzlement over the definition of $\Lambda_0(\Gamma)$, which lay in the fact that the centers of the $I(T)$ are not in $\overline{\mathbb{U}}$ at all.

Corollary 2. *For any Fuchsian group Γ, $\Lambda_0(\Gamma) \subset \partial\,\mathbb{U}$*

Proof. Suppose $\Lambda_0(\Gamma)$ has a limit point $p \notin \partial\,\mathbb{U}$. Then there is a sequence of isometries $T_1, T_2, \dots \in \Gamma$ such that $C(I(T_n)) \to p$. But by the lemma, the radii of the $I(T_n)$ go to zero. In particular, all but finitely many of the $I(T_n)$ do not even intersect $\mathbb{U}$. But we know that every $I(T)$ intersects $\mathbb{U}$, a gross contradition. □

Now we can prove our theorem.

Theorem 2. *For a Fuchsian group Γ acting on $\mathbb{U}$,*

$$\Lambda(\Gamma) = \Lambda_0(\Gamma)$$

Proof. Suppose p is a limit point of Λ_0. Then there is a sequence $T_1, T_2, \dots \in \Gamma$ such that $C(I(T)) \to p$. Consider the circles $I(T_n^{-1})$ Since the radii of the $I(T_n^{-1})$ must go to zero, we can pick a z which is outside of all but finitely many $I(T_n^{-1})$. The action of T_n^{-1} is to first invert in $I(T_n^{-1})$, mapping z to the inside of $I(T_n^{-1})$, and then reflect, sending z inside of $I(T_n)$. So $d(T_n^{-1}z, p) \leq d(T_n^{-1}z, C(I(T_n))) + d(C(I(T)), p)$, and both terms go to zero, the first by the lemma, and the second by hypothesis. Hence $T_n z \to p$.
Conversely, suppose p is a limit point of Λ. Then there is a $z \in \mathbb{U}$ and a sequence $T_1, T_2, \dots \in \Gamma$ such that $T_n z \to p$. As above means that $T_n z$ is inside of $I(T_n^{-1})$ for all but finitely many n, because z is on the outside of $I(T_n^{-1})$. So $d(C(I(T_n)), p) \leq d(C(I(T_n)), T_n z) + d(T_n z, p)$, and both go to zero, the first by the lemma, the second by hypothesis. □

We get for free that $\Lambda(\Gamma)$ is closed (a nontrivial fact), that $\Lambda(\Gamma) \subset \partial\,\mathbb{U}$ (though that is pretty obvious anyway) and that $\Lambda_0(\Gamma)$ is Γ-invariant.
We can use the similar reasoning to prove a number of facts.

Lemma 2. *Let Γ be a Fuchsian group, and suppose $\alpha, \beta, \delta \in \partial\,\mathbb{U}$ are distinct, with $\alpha \in \Lambda(\Gamma)$. Then α is a limit point of $\Gamma\beta$ or of $\Gamma\delta$.*

Proof. Choose $T_1, T_2, \dots \in \Gamma$ such that the centers of the $I(T_n^{-1})$ go to α. Then the centers of the $I(T_n)$ approach some $\gamma \in \partial \mathbb{U}$. Now it is possible that γ coincides with either β or δ, but not both; so without loss of generality assume $\gamma \neq \beta$. Then, for all but finitely many n, we know that β is outside of $I(T_n)$. It follows that T_n maps β inside of $I(T_n^{-1})$, and hence $T_n\beta$ approaches α, by Lemma 1. $\square$

This allows us to prove another characterization of $\Lambda(\Gamma)$.

Theorem 3. *If Γ is a Fuchsian group, and if $\Lambda(\Gamma)$ contains more than one point, then $\Lambda(\Gamma)$ is the closure of the fixed points of the hyperbolic transformations of Γ.*

Proof. First we show that Γ has a hyperbolic element. Γ is not purely elliptical, for then it is finite and $\Lambda(\Gamma)$ is empty. If it doesn't have a hyperbolic element, then it at least has a parabolic one. In $\mathbb{H}$, take a conjugate of the group so that we have a parabolic element $T = \begin{pmatrix} 1 & k \\ 0 & 1 \end{pmatrix}$. Then there must be another element with $S = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $c \neq 0$, for otherwise ∞ would be the only limit point. Then we have $|\operatorname{Tr} T^n S| = |a + d + nkc|$, which is > 2 for n large enough, giving us a hyperbolic element.
Now, if T is hyperbolic, then we can conjugate it to a transformation of the form $T = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}$, $\lambda \neq 1$, with fixed points $0, \infty$. Then for any $z \in \mathbb{H}$, $T^n z$ approaches ∞ for $\lambda > 1$, while T^{-n} approaches 0 (and vice versa for $\lambda < 1$). Since conjugation is continuous, it follows that every hyperbolic fixed point is in $\Lambda(\Gamma)$. Since $\Lambda(\Gamma)$ is closed, it in fact contains the closure of the set of hyperbolic fixed points.
Conversely, if $\alpha \in \Lambda$, then let T be a hyperbolic element of Γ, with distinct fixed points β_1, β_2. Then by the preceeding lemma, α is in the closure of $\Gamma\beta_1$ or $\Gamma\beta_2$. Say $T_1\beta_1, T_2\beta_2, \dots$ approaches α. These are fixed points for $T_1 T T_1^{-1}, T_2 T T_2^{-1}, \dots$, so that α is a limit of hypberbolic fixed points, completing the proof. $\square$

This characterization is useful partly because it gives $\Lambda(\Gamma)$ as the closure of a set, rather than just the limit points of a set, which is a little more flexible. It also seems to be an example of how $\Lambda(\Gamma)$ is in some sense a generic attractor for almost any set associated with Γ.
Our study of limit points culminates in a sharp classification theorem:

Theorem 4. *If Γ is a Fuchsian group with more than one fixed point, then $\Lambda(\Gamma)$ is either*

1. *all of $\partial \mathbb{U}$ or*

2. *a perfect nowhere dense subset of $\partial \mathbb{U}$*

These are called Fuchsian groups of the first and second kinds, respectively.

You may need to be reminded, as I did, that a perfect set is a closed set with no isolated points, and that the classic example of a perfect nowhere dense set is the Cantor set.

Proof. First, $\Lambda(\Gamma)$ is perfect, since it is closed and if $\alpha \in \Lambda$, then we have a sequence of hyperbolic fixed points approaching α. These hyperbolic fixed points are in Λ.
Now, if Λ is dense in $\partial\mathbb{U}$, then it is equal to $\partial\mathbb{U}$, since Λ is closed. If it is not dense, then pick $\alpha \in \Lambda$ with a neighborhood containing no other points of Λ. Suppose there is a $\beta \in \Lambda$ which has $\delta_1, \delta_2, \cdots \in \Lambda$ approaching it. Then at least one of the δ_n is not equal to α. So by the earlier lemma, either $\Gamma\beta$ or $\Gamma\delta_n$ has α as a limit point. Since Λ is Γ-invariant, this contradicts the hypothesis that α had a neighborhood free of other elements of Λ. So if Λ is not dense, then it is nowhere dense. □