A PULLBACK ALGORITHM FOR GENERAL (COFINITE) FUCHSIAN GROUPS

ANDREAS STRÖMBERGSSON

In this short note, we give an algorithm for computing the pullback of any given point in the hyperbolic plane into a Dirichlet fundamental domain, for a given, arbitrary cofinite Fuchsian group.

$\mathcal{H}$ denotes the hyperbolic plane, and Γ is our cofinite Fuchsian group. We use $d(\cdot,\cdot)$ to denote hyperbolic distance.

Take some point p in $\mathcal{H}$, not an elliptic fixpoint.

Let $\mathcal{F}$ be the corresponding Dirichlet fundamental domain:

$$\mathcal{F} = \{z \in \mathcal{H} \mid d(p,z) \leqq d(T(p),z) \text{ for all } T \in \Gamma\}.$$

Then we know from Siegel's theorem (as it is proved in Katok, p.80-), that $\mathcal{F}$ has a finite number of sides. We also know that the sides of $\mathcal{F}$ fall into congruent pairs, let us denote them s_1, s_{n+1}; s_2, s_{n+2}; ...; s_n, s_{2n}, and let $T_1, T_2, ..., T_n \in \Gamma$ be the identification maps (so that $T_j(s_j) = s_{j+n}$). (Recall the convention an elliptic fixpoint of order two is considered as a vertex separating *two distinct* sides of $\mathcal{F}$.)

We know that $T_1, T_2, ..., T_n$ *generates* Γ (Katok, p.73-74).

Our simple pullback algorithm is as follows:

---

1. Start with an arbitrary point z.

2. Compute the $2n$ points $T_1(z), T_1^{-1}(z), T_2(z), T_2^{-1}(z), ..., T_n^{-1}(z)$. Let z' be the one of these which has the *smallest* [hyperbolic] distance to p.

3. If $d(p,z') < d(p,z)$, then *replace z by z'*, and repeat from 2.

4. In the other case, $d(p,z') \geqq d(p,z)$, we *know that z lies in $\mathcal{F}$*, hence z is the searched-for pullback!

We now prove that this will always work (in theory).

Let us call the sequence obtained by working through the algorithm $z_0 = z$ (the initial point), z_1, z_2, ... We first show that the process must terminate after a finite number of steps.

Let K be the closed ball with center p and hyperbolic radius $= d(p,z_0)$. Then K is a compact set. We have $d(p,z_0) > d(p,z_1) > d(p,z_2) > ...$, hence all the points $z_0, z_1, ...$ lie in K and are distinct. They all belong to the

Γ-orbit of $z = z_0$. Hence, since Γ acts properly discontinuously, they can only be finitely many, and the process must terminate at some z_m.

It now remains to show that $z_m \in \mathcal{F}$. The proof of this depends on a certain relation between the sides $s_1, s_2, ..., s_{2n}$ and the corresponding identification maps $T_1, T_2, ..., T_n$, which we will now describe. (I do not find this made explicit in Katok's book.)

Each side s_j is a segment of a geodesic, and this geodesic splits $\mathcal{H}$ into two hyperbolic half-planes. Let H_j be the one of these two half-planes which contains $\mathcal{F}$. (Recall that $\mathcal{F}$ is known to be convex.)

For any $T \in \mathrm{PSL}(2,\mathbb{R})$ with $T(p) \neq p$, we define

$$B_p(T) := \{z \in \mathcal{H} \mid d(p,z) = d(T(p),z)\}$$
$$H_p(T) := \{z \in \mathcal{H} \mid d(p,z) \leqq d(T(p),z)\}.$$

Then $B_p(T)$ is a geodesic; it is the perpendicular bisector to the geodesic segment joining p and $T(p)$. Also, $H_p(T)$ is a hyperbolic half-plane with boundary $= B_p(T)$. Cf. Katok, p.54.

The relation mentioned above is the following:

$$\begin{aligned} H_j &= H_p(T_j^{-1}) \qquad \text{for } j \in \{1,2,...,n\}; \\ H_j &= H_p(T_{j-n}) \qquad \text{for } j \in \{n+1, n+2, ..., 2n\}. \end{aligned} \tag{1}$$

We prove this for $j \in \{1,...,n\}$; the case $j > n$ is very similar. Let us consider any point w on s_j. Then $T_j(w) \in s_{j+n}$, and hence both w and $T_1(w)$ belong to $\mathcal{F}$. Hence from the definition of $\mathcal{F}$ it follows that $d(p,w) \leq d(T_j^{-1}(p),w)$ and $d(p,T_j(w)) \leq d(T_j(p),T_j(w))$. Using the invariance of $d(\cdot,\cdot)$, the second inequality translates into $d(T_j^{-1}(p),w) \leq d(p,w)$. Hence $d(p,w) = d(T_j^{-1}(p),w)$, i.e. $w \in B_p(T_j^{-1})$. Since this is true for all $w \in s_j$, we actually have $s_1 \subset B_p(T_j^{-1})$. Hence $H_p(T_j^{-1})$ is either equal to the half-plane H_j, or to the half-plane on the opposite side of s_j. But we know that p belongs to the interior of $H_p(T_j^{-1})$, and to the interior of H_j, hence $H_j = H_p(T_j^{-1})$, and the claim is proved.

The proof of $z_m \in \mathcal{F}$ is now very simple. Since $\mathcal{F}$ is a convex polygon with sides $s_1, ..., s_{2n}$, it is sufficient to show that z_m belongs to each of the half-planes $H_1, ..., H_{2n}$. But from the fact that z_m was the point that caused the algorithm to terminate it follows that $d(p,T_j(z_m)) \geq d(p,z_m)$ and $d(p,T_j^{-1}(z_m)) \geq d(p,z_m)$ for all $j \in \{1,2,...,n\}$. In other words, $d(T_j^{-1}(p),z_m) \geq d(p,z_m)$ and $d(T_j(p),(z_m)) \geq d(p,z_m)$ for all $j \in \{1,2,...,n\}$, i.e. z_m belongs to all the sets $H_p(T_j)$ and $H_p(T_j^{-1})$. Hence our relation (1) shows that z_m belongs to $H_1, H_2, ..., H_{2n}$, and hence $z_m \in \mathcal{F}$.