Certifying Maximal Stretch Ratios for Fuchsian Free Group Representations

Joseph Hummel, Kaius Noble, Prajith Velicheti

University of Texas at Austin

Background

Geometries: In 2 dimensions there exist 3 types of geometries; Spherical, Euclidean and Hyperbolic. They are differentiated by the number of lines parallel to a given line through a point.

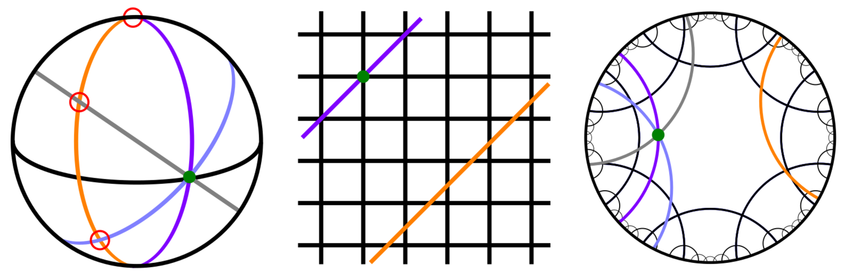

It is possible to define geometry on a surface by looking at the geometry of a covering space. The 1-holed torus is covered by the real plane $\mathbb{R}^2$, which allows it to be given Euclidean geometry. Similarly, the 2-holed torus is covered by the hyperbolic disk $\mathbb{H}^2$, which allows it to be given hyperbolic geometry.

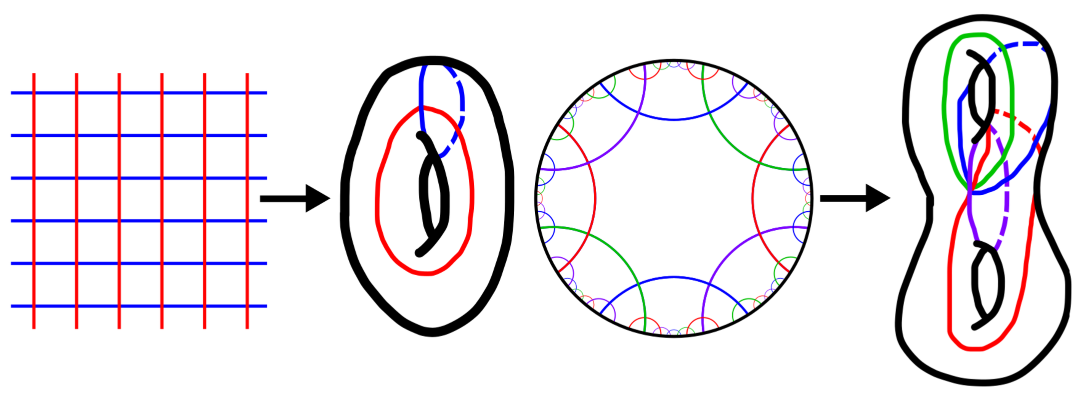

Wall-Generator Correspondence: The fundamental group of the surface acts upon the covering space via deck transformations. From this action, we get a correspondence between the walls of a fundamental domain and the generators of the fundamental group.

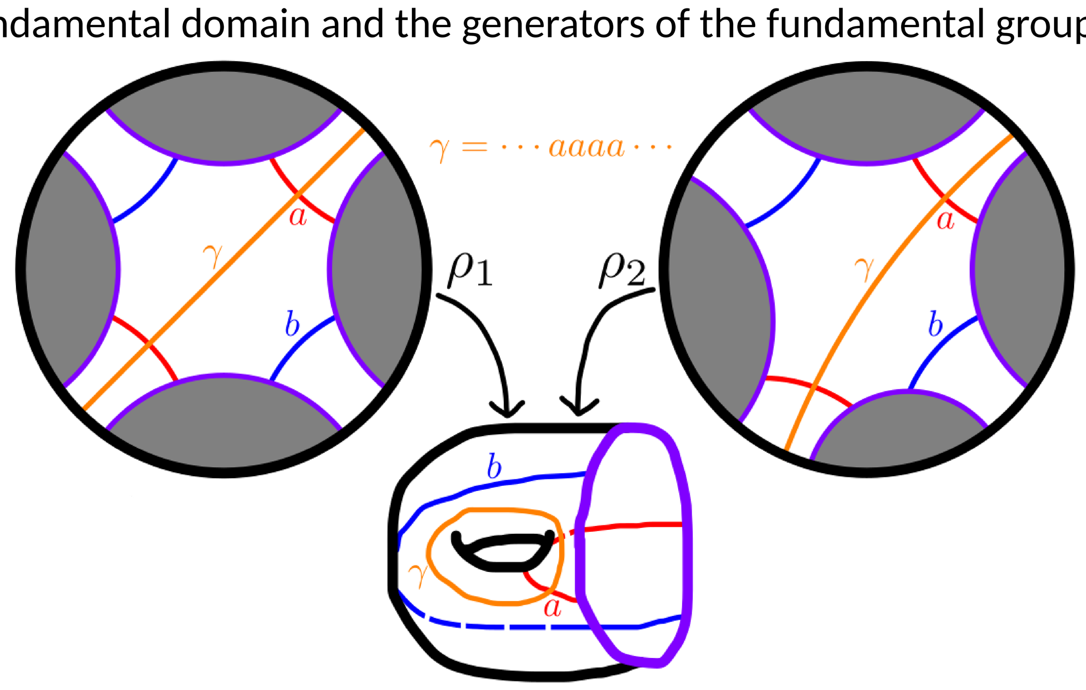

Question: By changing the walls in the cover, we can change the geometry on the surface. A natural question then becomes how do we compare two of these geometries. Thurston defined a way of doing this by looking at the maximal stretch ratio of all periodic geodesics between the two structures. This is called the Thurston metric and is given by:

$$d_{Th}(\rho_1, \rho_2) = \log\left(\sup_{\gamma \in \pi_1(S)} \frac{\ell_2(\gamma)}{\ell_1(\gamma)}\right)$$

Problem

With probability 1, this supremum is attained by an element γ of the fundamental group, which we think of as a periodic geodesic in the universal cover.

Purpose: Our project is to determine, with certainty, which periodic geodesic realizes the Thurston metric for a given pair of representations.

Preprocessed Words - Approximating Sections of Geodesics

The combinatorics of a word tells us geometric information about any geodesic representative as shown below.

Definition: A preprocessed word of length n is a reduced word $p = p_0 \cdots p_n$. Let $W_p \subset \pi_1(S)$ be the set of all words of the form $w_1 \dots w_k p w_{k+1} \dots w_{2k}$. We define:

$$L_r(p) := \left[\inf_{w \in W_p} \ell_r^C(w),\ \sup_{w \in W_p} \ell_r^C(w)\right]$$

where $\ell_r^C(w)$ is the length of the geodesic representative of w in the central fundamental domain.

Example: Considering the preprocessed word $p = baab$, we see in the picture a geodesic representative of some $w \in W_p$. The value $\ell_r^C(w)$ is the length of the geodesic between the middle red walls.

Why? There is a finite set of preprocessed words of length n so calculating $L_r(p)$ for each of them gives us length bounds for any geodesic in any particular fundamental domain.

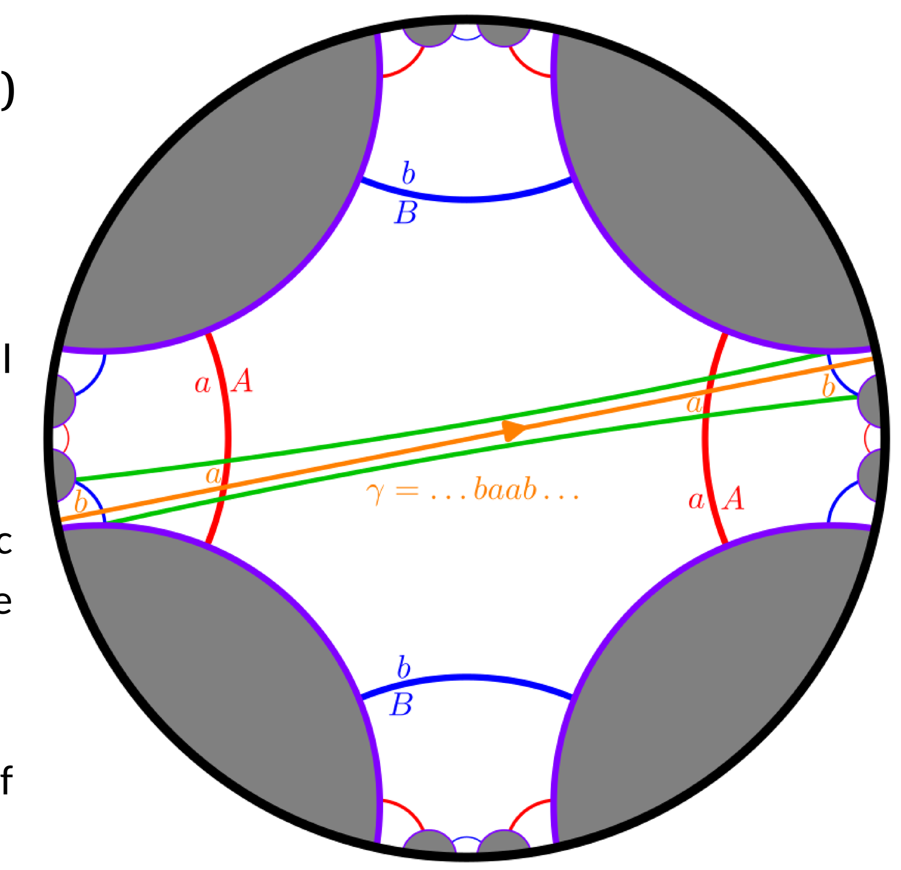

Groupings - Approximating Longer Geodesic Sections

We can group together preprocessed words to get length bounds over multiple fundamental domains.

Definition: A grouping g of length m is a reduced word of length $m + n - 1$ which is composed of m consecutive preprocessed words. Below is a particular example with $m = 3$ and $n = 8$.

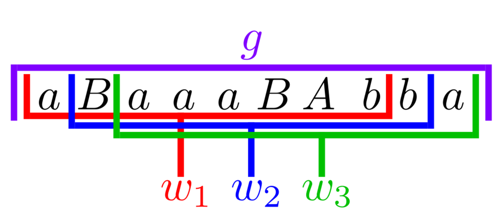

We define $L_r(g) = \sum_{p \in g} L_r(p)$ which will approximate the lengths of the middle m fundamental domains.

Application: We can push this idea even further and approximate the periodic length of any reduced word w using our preprocessed words. This is how we will use our preprocessed words to prove statements about general reduced words.

Beta - Comparing Stretch Ratios

Fix a potential optimizer $u \in \pi_1(S)$. For a word $z \in \pi_1(S)$, we define:

$$\beta(z) := \begin{bmatrix} \ell_1(z) \\ \ell_2(z) \end{bmatrix} \cdot \begin{bmatrix} -\ell_2(u) \\ \ell_1(u) \end{bmatrix}$$

Why? This measures how much a geodesic representative is stretched when compared to the optimizer. Positive beta means more stretched while negative beta means less stretched.

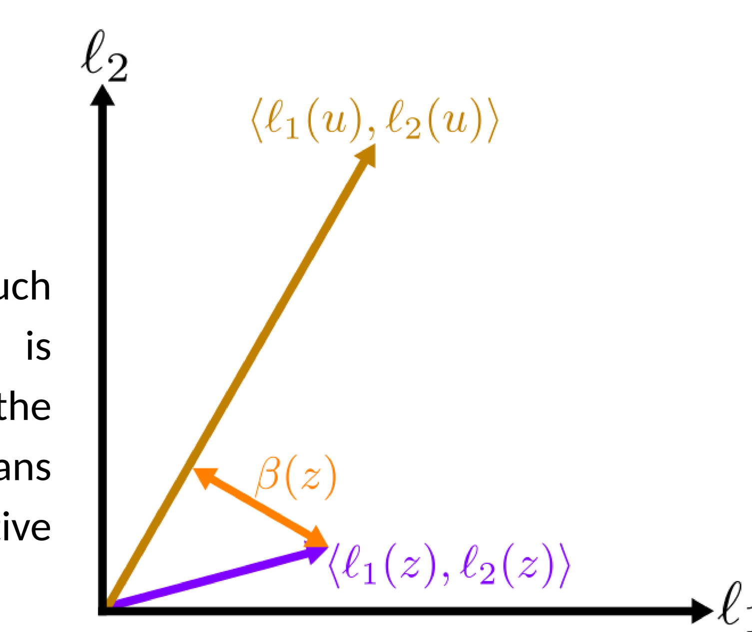

Extension: We can define for preprocessed words and groupings:

$$\beta^P(p) := \begin{bmatrix} \min(L_1(p)) \\ \max(L_2(p)) \end{bmatrix} \cdot \begin{bmatrix} -\ell_2(u) \\ \ell_1(u) \end{bmatrix} \qquad \beta^G(g) := \sum_{p \in g} \beta^P(p)$$

These values give an upper bound for how stretched a geodesic representative can be in the particular fundamental domains when compared to the optimizer.

Progress Check - Example

Preprocessed Word Distribution: For a particular example, the distribution of $\beta^P(p)$ for $n = 6$ is shown below, computed with respect to the real optimizer ab.

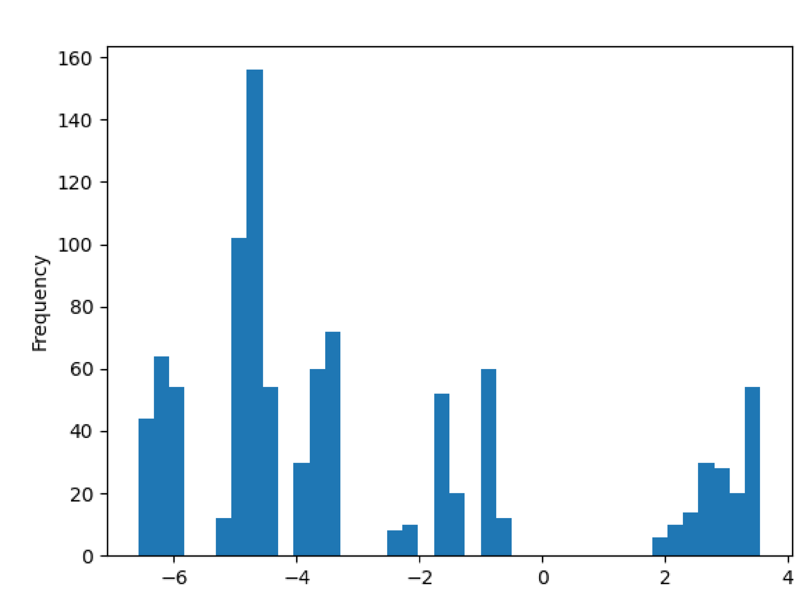

We find 162/972 $\sim$ 16.67% of the preprocessed word betas are positive.

What does this tell us? Even though the optimizer is periodically the most stretched, sections of geodesics can still be more stretched than the optimizer itself.

Grouping Motivation: Many positive-beta preprocessed words must be followed by negative-beta preprocessed words. Thus, considering multiple fundamental domains may yield a better distribution.

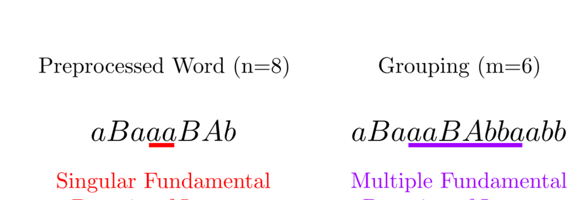

Grouping Distribution: The distribution of groupings for $m = 7$ and $n = 6$ is shown below.

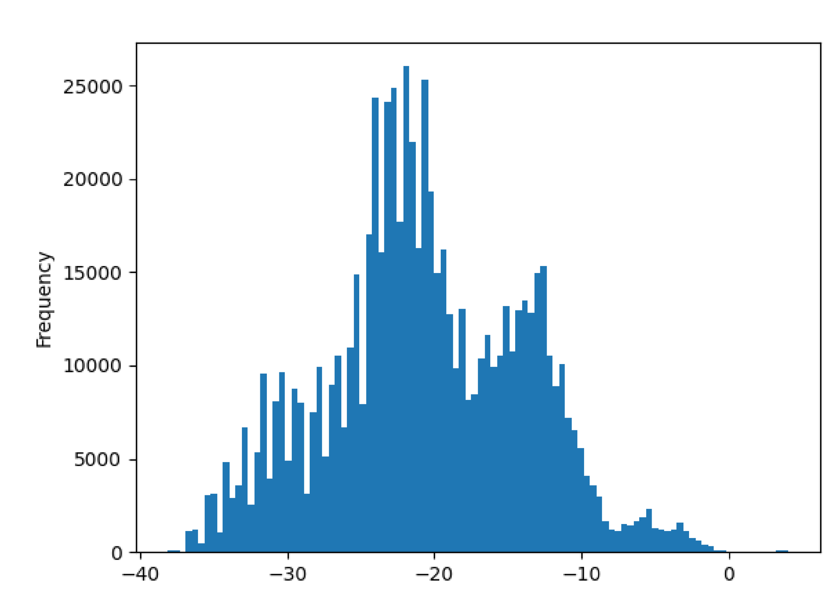

We find 162/708588 $\sim$ 0.023% have positive beta. Additionally, every positive-beta grouping contains one of:

$$(ab)^4 \quad (ba)^4 \quad (AB)^4 \quad (BA)^4$$

Intuition

Sections of geodesics can outperform the optimizer on small intervals. However, the longer the section of a non-optimizer geodesic we consider, the less likely it is to outperform the optimizer.

Key Insight: Groupings approximate the central fundamental domains of a word. If our approximations are sufficient, then the middle portion of positive groupings should converge to the optimizer as we increase the grouping length. For example, in the previous case where the optimizer is the periodic word ab, we observe:

m	Most Positive Grouping	Beta
2	ABabaBA	1.365
4	ABababaBA	1.683
6	ABabababaBA	2.000
7	BAbabababaBA	4.089

General Argument: Certification involves preprocessing to sufficient accuracy, and then increasing grouping length m until we observe that all positive groupings contain a significant portion of the optimizer. Then we can use some analytic tools similar to the Anosov Closing Lemma to effectively ignore groupings with a significant portion of the optimizer and then leverage this tool to certify the optimizer.

Theorems

We have proved the following results.

Certification Theorem: The algorithm we have informally outlined can certify that a particular geodesic is the optimizer.

Converse Theorem: Suppose there exists some $\delta > 0$ such that for every word w that is not the optimizer, we have $\beta(w) \leq -\delta$. Then the algorithm we have outlined will certify the optimizer.

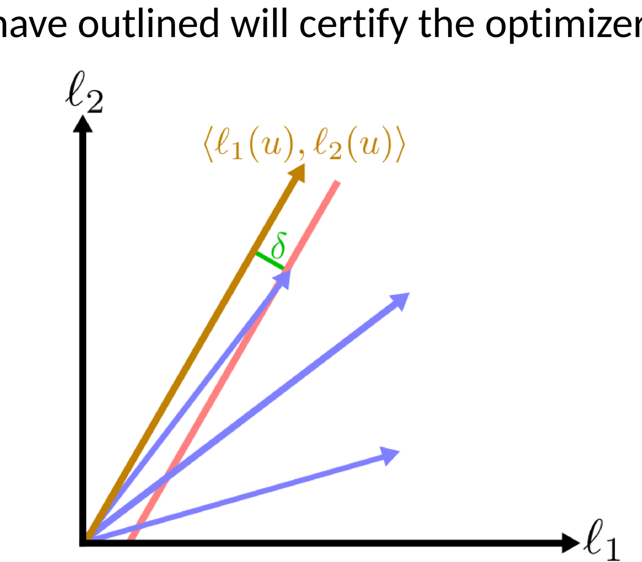

The details can be found in our paper entitled "Certifying Maximal Stretch Ratios for Fuchsian Free Group Representations".

Acknowledgments

This project was completed in the Texas Experimental Geometry Lab (TXGL) at the University of Texas at Austin under the guidance of Jeff Danciger and Aaron Benda.

Contact Information

- Web: https://texas-experimental-geometry-lab.gitlab.io/
- Email: knoble6782@gmail.com, jhummel343@utexas.edu

Bibliography

- William Thurston. "Minimal Stretch Maps between Hyperbolic Surfaces". en. In: (1986)
- A. Benda, J. Hummel, K. Noble and P. Velicheti. "Certifying Maximal Stretch Ratios for Fuchsian Free Group Representations". (In Progress)