Some problems

The ICERM 2018 Team

1 A relation between self-intersection number of geodesics on a surface

Problem: Fix a hyperbolic surface S. For each $k \geq 0$, define $l_k(S)$ as the length of the shortest closed geodesic in S with at least k self-intersection points. Prove that l_k is increasing.

An **orientable closed surface** can be thought of as the outer layer of a solid (for instance, the solid can be a ball, or a donut or a pretzel...). An **orientable surface with boundary** is an orientable closed surface with some disks removed. From now on, we'll refer to both orientable surfaces closed or with boundary as surfaces.

A **pair of pants** is the following surface with boundary: a sphere with three disks removed.

A **closed curve on a surface** is a map from the circle to the surface (a useful way to visualize a closed curve is picture it as a very thin rubber band winding around the surface).

A **metric** is a certain way of measuring distance between points on a space. We are also going to assume that one can measure angles between straight lines (or whatever a "straight line" is on a surface).

The **hyperbolic plane** is two dimensional space with a **hyperbolic metric**. A hyperbolic metric is... this is hard to explain in a few words, so whatever I'll say has to be taken with a grain of salt (or a full salt shaker). On a very small scale, a hyperbolic metric is very similar to the Euclidean metric you know. On a larger scale, it is a metric such that given a straight line and a point not on that line, there are infinitely many lines through that point and disjoint from the given line.

A surface is **hyperbolic** if it has a hyperbolic metric, that is, it is made by gluing pairs of edges of polygonal pieces of the hyperbolic plane, with the appropriate angles at the vertices. Each (not too kinky, for instance, smooth) curve on a hyperbolic surface has a **length**.

A **geodesic on a surface** is a closed curve that cannot be continuously deformed (without leaving the surface) to a shorter curve.

A **self-intersection point of a geodesic** is a point through which two or more branches of the geodesic pass.

2 Hausdorff Dimension and Hyperbolic Pair of Pants

Problem: Determine (or find bounds on) the Hausdorff dimension of the limit set of a pair of pants in terms of the length of the boundary components. Possible start: Make two boundary components a fixed length C and find the Hausdorff dimension varying the third length.

For background, see that of the problem **A relation between self-intersection number of geodesics on a surface** and add the following:

A hyperbolic pair of pants is a pair of pants with a hyperbolic metric for which the three boundary components are closed geodesics. We will see that the hyperbolic metric is determined by the length of the three boundary components. Thus each hyperbolic pair of pants is "labeled" by a triple of positive numbers.

A **Cantor set** can be obtained as follows:

- Step 0: Start with the closed interval $[0, 1]$. Set $C_0 = [0, 1]$.
- Step 1: Remove from $[0, 1]$ the open interval $(1/3, 2/3)$, the "open middle third" of $[0, 1]$. Note that you obtain two closed intervals $[0, 1/3]$ and $[2/3, 1]$. Denote by C_1 the union of these two closed intervals.
- Step 2: Remove the open middle thirds of each of the two closed intervals forming C_1. Now you have four closed intervals. Denote by C_2 the union of these four closed intervals.
- Step n: You have 2^{n-1} closed intervals and remove the open middle third of each of them. Denote by C_n the union of the 2^n intervals you obtained.

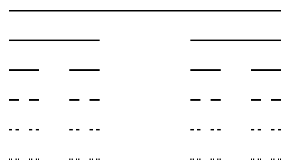

Figure 1: The sequence of C_n's

The Cantor set is the intersection of all the C_n's. (for more about Cantor sets see https://www.math.hmc.edu/funfacts/ffiles/20004.3.shtml#).

We will see how a hyperbolic pair of pants determines a Cantor set (a bit more general than the description we just gave.) This Cantor set is the **limit set** of the hyperbolic pair of pants.

Cantor sets are not one-dimensional (we defined them as a subset of an interval, but they have "too many holes" to be one-dimensional). But Cantor sets are (dimension-wise) more than "just points", thus they are more than 0-dimensional. There is a way to assign to each Cantor set a **Hausdorff dimension**, which in this case is a number between 0 and 1.

The problem consists of studying the relation between the Hausdorff dimension of a hyperbolic pair of pants, and the length of its three boundary components.

3 Symmetries of curve graph variants

The *curve graph* of a surface S is a very big graph that keeps track of how loops on the surface cross one another. Loosely speaking, each vertex of the curve graph represents a loop on the surface S which if it were made of string, you wouldn't be able to pull it tightly to a point– before the loop got that small, it would get caught on some feature of the surface. For instance, if S is the surface of a donut, a good example of a loop we might consider is one which goes around the "tube" of the donut. An example of a loop we wouldn't consider would be one that bounds a very tiny disk (see Figure 2). The technical term is that we want only loops that are **homotopically non-trivial**.

Whenever one loop can be obtained from another by sliding it along the surface, we think of them as being the same. In Figure 3, the white, green, and pink loops are all equivalent to each other, but the orange loop is different from all of them because there is no way to slide the orange loop on the surface so that it coincides with any of the other three. The technical way to say this is that the white, green, and pink loops are all **homotopic to each other**.

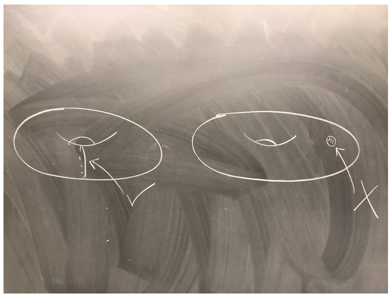

Figure 2: One of these loops is homotopically non-trivial

The curve graph is a graph whose vertices represent the different loops on a surface, S, and where we connect two vertices by an edge when the corresponding loops cross the minimum number of times possible for any two different loops. The figure below illustrates a little piece of the curve graph of a surface of genus 3 (3 "holes"):

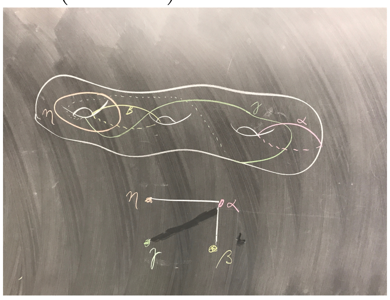

The goal of this project is to understand the **symmetries** of a sort of generalized curve graph. The symmetries of the curve graph are known and well-understood, so the strategy will be to try to deduce something about the symmetries of this less understood and more general graph, called the k-**curve graph**, by using what is already known about the curve graph.

If k is any given natural number, the k-curve graph has the same vertices as the curve graph, but now we draw an edge between a pair of vertices whenever those loops can be made to cross at most k times.

Problem: Characterize the symmetries of the k-curve graph.

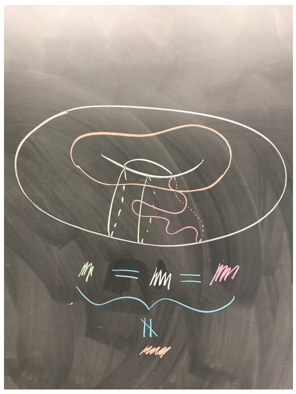

Figure 3: The orange loop is not homotopic to any of the others

4 Square-tilings

A **square-tiling** (also called an **origami**) is a region of the plane tiled by squares, so that every boundary edge has a "decoration". The decorations come in pairs, so that every "top" edge has the same decoration as some "bottom" edge, and similarly every "left" boundary edge is paired off with some "right" boundary edge. See Figure 4 on the right for an example.

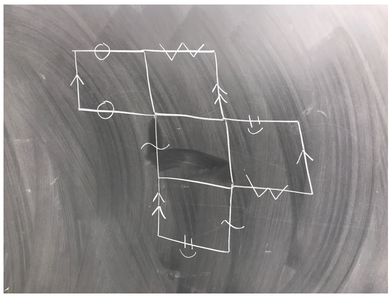

Figure 4: Every right edge is paired to a left edge and visa versa, and similarly for top and bottom edges.

A square-tiling gives rise to a surface when we **glue** the boundary edges up in pairs: each boundary edge glues to the edge with the same decorative pattern. The square-tiling becomes a fascinating graph on the surface whose faces are all squares.

Given a square-tiling, we can transform it by a 45-degree **shear**, meaning that we replace each square with a parallelogram of the same area and with acute angles equal to 45 degrees:

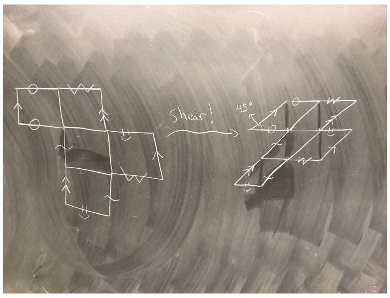

After the shearing, we are then allowed to chop off right triangles in the right places and re-attach them somewhere else **so long as we only ever glue edges together with the same decoration**. It turns out that you can always reshuffle after a shear to get a new square-tiling! Example:

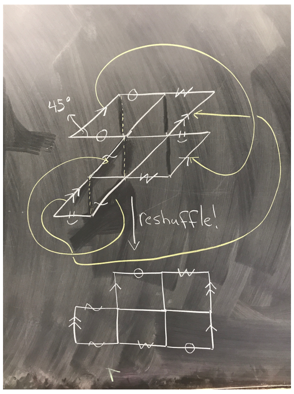

Something to be a bit careful about: when we chop off a triangle and place it somewhere else, we may create new boundary edges which used to be in the middle of the tiling somewhere. We need to decorate these new edges, and the rule is that we assign two of these new edges the same decoration when they used to be the same edge before the chopping.

There are other ways we can get a new square-tiling from an old one, e.g. we could rotate a square-tiling by 90 degrees. We can also chain these two transformations together, e.g., first shear, then rotate the resulting square-tiling by 90 degrees, then shear twice in a row, then rotate again, etc. etc.

Problem: Which square-tiled surfaces can be obtained from which others by applying a sequence of the shear & rotation moves? The goal of this project is to think about this problem for a certain family of square-tilings that arise naturally when thinking about pairs of loops on a surface (called **filling pairs**). This project will have both theoretical and computer-experimental components– participants can choose to work on one or both of these components.

5 Coloring problems for arc and curve graphs

This problem is about investigating graph theoretic concepts in the topology/geometry context of curves and surfaces. In fact, this problem could be about the *curve graph*, the graph introduced in a problem above, but for technical reasons this is a little more approachable if we work with a close relative of the curve graph. Let's introduce this related graph first.

Let S be a sphere with n marked points. A *simple arc* on S is an injective map from an interval $I = [0,1]$ to S that sends the endpoints of I to marked points, with one additional complication: we think of two simple arcs $\alpha : I \to S$ and $\alpha' : I \to S$ as 'the same' when one of these intervals, say $\alpha(I)$, can be pushed along the surface, avoiding any of the marked points, until it coincides with $\alpha'(I)$. In this case, we view α and α' as **representatives** of the same simple arc. For instance, in Figure 5 to the right, the arc α on the left side of the figure is the 'same' as the arc α' on the right side.

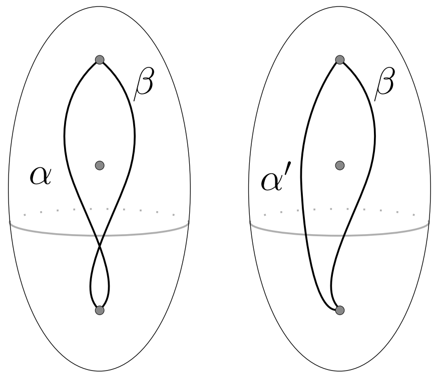

Figure 5: Because α and α' are representatives of the same simple arc, and α' and β are disjoint, there is an edge in the arc graph between these two arcs.

The **arc graph** of S, denoted by $\mathcal{A}(S)$, has one vertex for each simple arc on S. The arc graph has an edge between the vertices corresponding to α and β when the intervals $\alpha(I)$ and $\beta(I)$ are disjoint. (Ok, there's a technicality here if you are reading closely: It's possible that α and β are *not* disjoint–i.e. they intersect–but there might be a different representative α' of the simple arc represented by α so that $\alpha'(I)$ and $\beta(I)$ are in fact disjoint—see Figure 5 for exactly this situation. To deal with this annoyance, we should really say that we have an edge between two simple arcs when they have disjoint represntatives.)

Problem: Compute the chromatic number of the graph $\mathcal{A}(S)$.

The *chromatic number* of a graph is defined as follows: A *proper coloring* of a graph is an assignment of colors to each of the vertices, in such a way that any vertices connected by an edge get different colors. The **chromatic number** of a graph is the minimal number of colors needed to construct a proper coloring of it.

Some comments about this problem: It is not at all obvious that the chromatic number of $\mathcal{A}(S)$ is even finite! In fact, it's known that the chromatic number of $\mathcal{A}(S)$ lies somewhere between $n \log n$ and n^3. A precise answer to this problem is likely too hard, but there is definitely room for improved bounds. For example, here are some more approachable versions of this problem:

1. Compute the chromatic number of the subgraph spanned by arcs incident to a fixed puncture.
2. Compute the chromatic number of the subgraph of arcs that go from a fixed puncture to either of two other punctures.
3. Compute the chromatic number of the four-punctured sphere arc graph.

6 Curves intersecting each other twice

This problem is secretly about the following question: Suppose that a collection of simple closed curves on a surface pairwise intersect each other r times. How many such curves might there be? When $r = 1$, this problem has a very concrete answer, and the answer is obtained by finding a reinterpretation using linear algebra. However, when $r = 2$ no answer is known, and the same kind of linear algebra reinterpretation needs some critical thought. That reinterpretation is how the problem below is phrased.

Let V be a vector space of dimension $2n$ over the finite field $\mathbb{Z}/p\mathbb{Z}$. In case this sounds intimidating, it is all a lot simpler than it might sound. All you do is pretend that the constants that you are allowed to multiply vectors in V by are in the set $0, 1, 2, \ldots, p-1$, and also $p \equiv 0$. Moreover, suppose that V is equipped with a **symplectic form** $\omega : V \times V \to \mathbb{Z}/p\mathbb{Z}$. The latter object is a bilinear form that is *nondegenerate*—in the sense that $\omega(u, v) = 0$ for all $v \in V$ implies that $u = 0$—and *alternating*—in the sense that $\omega(u, u) = 0$ for all u. Think of ω as a funky-inverted-inner-product (that is, it's not an inner product, but it shares some features with one). When $\omega(u, v) = 0$ for vectors $u, v \in V$, we say that u and v are *ω-orthogonal.*

Problem: Suppose that that $\Gamma \subset V$ is a collection of vectors that are pairwise *not* ω-orthogonal, i.e. $\omega(u, v) \neq 0$ for every pair of vectors $u, v \in V$ with $u \neq v$. How big may Γ be?

Any bounds for an answer to this problem that depend on n will be interesting—the tighter the bounds the more interesting the answer will be! In order to inform the the curve question above, one should take $p \geq 3$. When $p = 2$, it turns out that the answer is $2n + 1$.

7 Navigating the space of $\mathsf{SL}(2, \mathbb{C})$-characters

The group of all 2×2 matrices with determinant one and complex numbers as entries is denoted by $SL(2, \mathbb{C})$. Given a group G, the $SL(2, \mathbb{C})$-**character variety** of G is, very loosely speaking, a space which keeps track of all the different ways that G can be mapped into $SL(2, \mathbb{C})$.

A classic result Vogt-Fricke asserts that the $SL_2(\mathbb{C})$ character variety of a two-generator free group $\langle A, B \rangle$ is isomorphic to $\mathbb{C}^3$; the three coordinates are $a := \mathsf{tr}(A), b := \mathsf{tr}(B)$ and $c := \mathsf{tr}(AB)$. An important role is played by the trace of the commutator

$$\mathsf{tr}(ABA^{-1}B^{-1}) = \kappa(a, b, c) := a^2 + b^2 + c^2 - abc - 2.$$

Points on the level sets of κ correspond to various hyperbolic surfaces (possibly with singularities), subject to natural geometric constraints. For example, when $a, b, c \in \mathbb{R}$ and $k < -2$, the level sets $\kappa^{-1}(k)$ parametrize hyperbolic structures on a torus with boundary of fixed length.

Problem: Use and develop software packages to interactively explore the geometry of these cubic surfaces in terms of the corresponding polygons and tilings of the hyperbolic plane. Devise and try to prove some conjectures based on these explorations.