Introduction to Rational Billiards III

Talk by John Smillie

August 22, 2007

●Affine Maps of Translation Surfaces

Let S and S' be translation surfaces and $f : S \to S'$ a homeomorphism that is smooth away from singular points. At each nonsingular point we have the map $Df_p : T_p \to T_{f(p)}$, and as $T_p \cong T_{f(p)} \cong \mathbb{R}^2$, we can view Df_p as an element of $GL_2(\mathbb{R})$.

If Df_p does not depend on p, then f is said to be affine. If Df_p is the identity for all p then f is an equivalence of translation surfaces. If $Df_p \in O(2)$ then f is an isometry. A general affine map takes directional flows to directional flows, but as f does not necessarily preserve lengths it may change the parametrization of the these flows.

When studying maps between surfaces in the topological or conformal category, it is useful to consider a class of objects called half-translation surfaces, which are slightly more general than the translations surfaces we have introduced so far. A half-translation surface is a surface given by gluing polygons, but here the gluing maps are allowed to be maps of the form $v \mapsto -v + c$. These maps reverse the directions of the horizontal and vertical foliations that the surface inherits from the plane, so directions on these surfaces are only defined up to sign. Translation surfaces and half-translation surfaces are often referred to as flat structures. An example of such a surface is shown below.
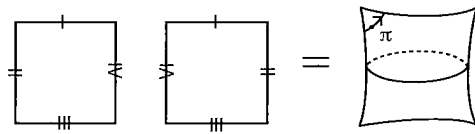

Notice that the cone angle at one corner of the pillowcase is π, whereas for a translation surface cone angles always come in multiples of 2π. From the definition one can deduce that the allowable cone angles for a half-translation surface are of the form πn where $n \geq 1$. For half-translation surfaces Df is well defined modulo $\pm Id$, so we can define an affine maps to

1

be one in which Df is constant as an element of $GL_2(R)/\pm Id$. Recall that a translation structure induces a conformal structure. The same is true for half-translation surfaces. To go from a conformal structure to a flat structure, one needs an abelian differential in the case of translation surfaces and a quadratic differential in the case of half-translation surfaces.

Let $Aff(S)$ be the group of affine transformations from a (half)-translation surface to itself. Then we have a homeomorphism from $Aff(S) \to PSL_2(\mathbb{R})$ given by $f \mapsto [Df]$. The image of this map is called the Veech group of S, and its image in $PSL_2(\mathbb{R})$ is discrete. As an example we can consider the Veech group of the square torus. The affine maps preserving the torus are exactly the group $SL_2(\mathbb{Z})$, so its Veech group is $PSL_2(\mathbb{Z})$.

Because the Veech group is a subgroup of $PSL_2(\mathbb{R})$, the group of isometries of the hyperbolic plane, we can use some of the language of hyperbolic geometry in this setting. We say an affine automorphism is elliptic, parabolic or hyperbolic depending on the type of its image in $PSL_2(\mathbb{R})$. Elliptic affine automorphisms are finite order, for example the automorphism given by rotating a glued regular octagon by an eighth of a turn. Hyperbolic automorphisms offer a nice class of examples of pseudo-Anosov transformations. In fact, the first examples of pseudo-Anosov homeomorphisms were constructed by Thurston using flat structures. Furthermore, every pseudo-Anosov homeomorphism is homotopic to an affine automorphism of a flat surface.

For parabolic diffeomorphisms, we can think of Dehn twists on cylinders. For a single cylinder, a Dehn twist can be realized as an affine automorphism of the cylinder. Unrolling the cylinder, we see a glued rectangle with width c and height h. The Dehn twist takes the vector $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ to $\begin{pmatrix} 0 \\ h \end{pmatrix}$ and takes horizontal vectors to themselves, so is given by the matrix $\begin{pmatrix} 1 & \frac{c}{h} \\ 0 & 1 \end{pmatrix}$.
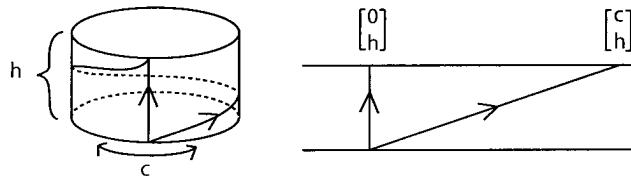
We can also do this simultaneously on a collection of cylinders. In order for this map to be affine, the derivative of a Dehn twist on one cylinder has to equal the derivative of a Dehn twist on any other cylinder, which will be the case provided that each cylinder has the same modulus. One can also get a map with constant derivative by doing different numbers of Dehn twists on different cylinders, provided that the moduli of these cylinders are rationally related. The following illustration shows an example of a flat surface that

2

admits two different cylinder decomposition in which all cylinders have the same modulus.
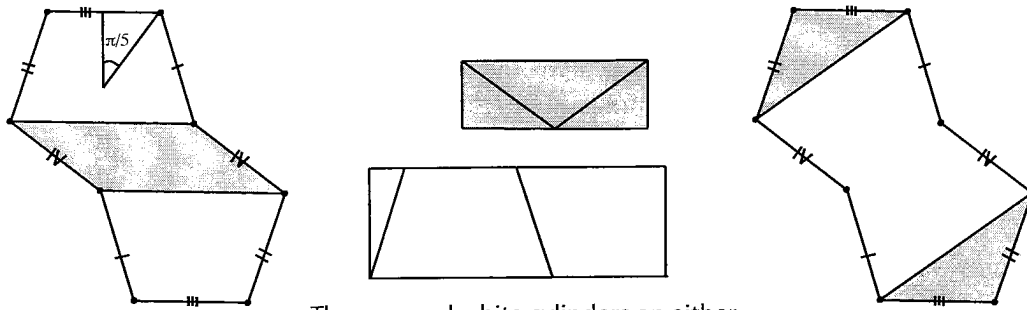

The gray and white cylinders on either side may be cut up and reglued to form these cylinders with equal moduli

The above example shows that affine automorphisms are not always easy to read off from looking at the surface.

●THE CONCEPT OF RENORMALIZATION

As a final topic, we will describe a billiards problem whose solution employs a technique that is generally useful in dynamics. In the previous lecture, we discussed how billiard trajectories are distributed in a topological sense. In this section we will discuss how billiard trajectories are distributed in a different sense, namely how often they hit a given side of the table.

Consider a billiard trajectory in the square torus whose initial slope is irrational. We would like to describe the sequence of edges that this trajectory hits. Labeling the vertical sides with a 0 and the horizontal sides with a 1 (without regard to top and bottom, left and right), we can associate a sequence of zeros and ones to a billiard trajectory.


0   0100101...

1
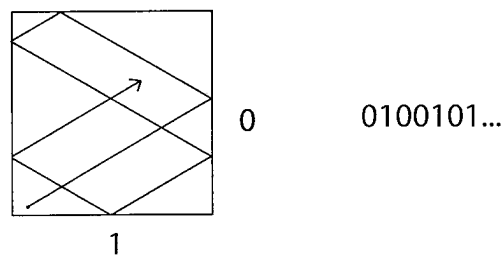
For concreteness we will study the trajectory whose initial slope is $\phi^{-1} = \frac{-1+\sqrt{5}}{2}$. Unwrapping the torus, this trajectory becomes a straight line in the plane and the sequence of
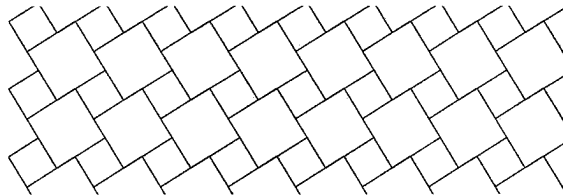
3

ones and zeros records the order in which our trajectory hits vertical and horizontal lines between lattice points.

ϕ^{-1} is the slope of an eigenvector for the matrix $L = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, so L preserves this trajectory but scales it by the eigenvalue ϕ. Thus the billiard flow in this direction is taken to itself by L but with a change of parameter. Using iteration of the map L to rescale the system exhibits one instance of a renormalization technique.

Because L maps the billiard trajectory to itself with a change of parameter, to find the position of the billiard ball at a later time we can either continue the flow or apply the map L. More explicitly, if at time t we are at the point $\psi(t)$, then $L\psi(t) = \psi(\phi t)$. One way in which we can use this fact is to compute the symbol sequence given by a segment of a billiard trajectory. We may color the plane gray and white according to whether the trajectory has just hit a horizontal or vertical side. The symbol sequence of a point in the gray region will have just added a zero to its symbol sequence, and a point in the white region will have just added a one. If we add a transversal to the flow as shown below, we see a pattern of L-shaped tiles.
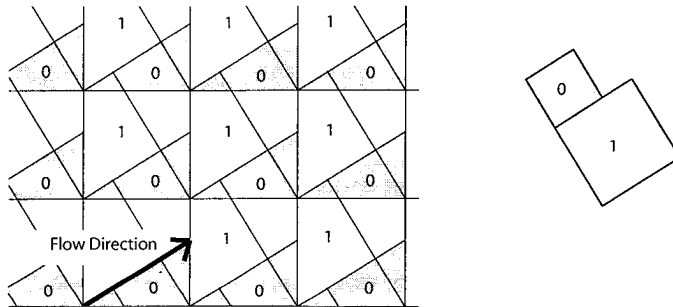
It's not hard to see that if we look at the L-shaped tiles, the part shown in gray on the right hand side of the above diagram corresponds to a collision with a horizontal side (as any time you cross a horizontal edge must enter this box) and likewise the white box corresponds to a collision with a vertical side. We can therefore recolor the plane as indicated on the right of the above diagram, and adding a one to the symbol sequence each time the trajectory enters a white box and a zero each time it enters a gray box we get the same symbol sequence.

4

The following illustration shows how this tiling transforms under the map L.

The L-shaped tiles have been flipped and stretched, but remarkably return to their original shape, with the gray and white parts reconfigured. The following illustration shows the image under L of one of these tiles as it sits in the original tiling.

The above diagram shows that a segment of the billiard trajectory that traverses a white box is mapped under L to a trajectory that traverses both a white and a gray box in the original tiling, and a segment of the trajectory that traverses a gray box is mapped to a trajectory that traverses a white box in the original tiling. Given the symbol sequence of a segment of a billiard trajectory, we can get the symbol sequence for the image of this segment under L (which is some extension of the segment by the billiard flow) by replacing digits via the rules $0 \mapsto 1$ and $1 \mapsto 10$. For example, starting with a segment l of a trajectory traversing a white box, we derive the symbol sequences of its iterates under L as follows:

$(l; 1)$, $(L(l); 10)$, $(L^2(l); 101)$, $(L^3(l); 10110)$,$(L^4(l); 10110101)$,$(L^5(l); 1011010110110)$, $(L^6(l); 1011010110110101101101110101)$,$(L^7(l); 1011010110110101101011011010110110)$...

Using this process we can very efficiently give the symbol sequence to a specified number of digits, as the number of digits grows exponentially in the number of times we apply L. The rule for replacing zeros and ones also tells us the frequency with which a given digit appears in this sequence, and moreover one can deduce from this rule the frequency with

which any finite sequence will appear.

In the above discussion, we only looked at a single direction on the torus, but this analysis will also work for any trajectory on the torus. Given a direction and a perpendicular transversal, we can find a matrix L that preserves these directions. By rotating the entire picture, we can ensure that the billiard trajectory is vertical, at the expense of changing the lattice that gives us our torus. This will have the effect of making the matrix L a diagonal matrix.

We can now think of how the matrix $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$ acts on our particular lattice. Note that for some t, $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix} = L$, and the fact L preserves the lattice says precisely that the orbit of our lattice under the flow $\begin{pmatrix} e^t & 0 \\ 0 & e^{-t} \end{pmatrix}$ is periodic. This flow is called the Teichmüller flow.

The Teichmüller flow as a flow on the space of lattices, which can be thought of as the moduli space of translation structures on the torus or as $SL_2(\mathbb{R})/SL_2(\mathbb{Z})$. This space is the unit tangent bundle of the moduli space of Riemann surfaces of genus 1, and the Teichmüller flow is precisely the geodesic flow on this space. The fact that L preserves the lattice shows that L^2 corresponds to a closed geodesic on moduli space (we square to get an orientation preserving transformation).

It turns out that a concrete analysis of billiard orbits is possible whenever the Teichmüller flow corresponding to our trajectory returns to a compact set in moduli space. For details see

Arnoux, Pierre. Le codage du flot godsique sur la surface modulaire. *Enseign. Math. (2)* 40 (1994), no. 1-2, 29–48.

There are other billiard tables on which a complete analysis of billiard orbits is possible. This analysis relies on being able to understand the $SL_2(\mathbb{R})$ orbit of the flat surface corresponding to the billiard table. Such an understanding is possible in the case when the Veech group is a lattice. An example of a surface whose Veech group is a lattice is the one shown above, corresponding to billiards in a right angled triangle with an angle of $\pi/5$.