Math 181 Handout 13

Rich Schwartz

November 28, 2005

The purpose of this handout is to talk about flat cone surfaces, the combinatorial Gauss-Bonnet Theorem, and polygonal billiards.

1 Sectors and Euclidean Cones

A *sector* in $\boldsymbol{R}^2$ is the closure of one of the two components of $\boldsymbol{R}^2 - \rho_1 - \rho_2$, where ρ_1 and ρ_2 are two distinct rays emanating from the origin. For example non-negative quadrant is a sector. The *angle* of the sector is defined as the angle between ρ_1 and ρ_2 as measured from inside the cone. For instance, the angle of the non-negative quadrant is $\pi/2$.

Two sectors in $\boldsymbol{R}^2$ can be glued together isometrically along one of their edges. A *Euclidean cone* is a space obtained by gluing together, in a cyclic pattern, a finite number of sectors. The *angle* of the Euclidean cone is the sum of the angles of the sectors. The *cone point* is the equivalence class of the origin(s) under the gluing. The cone point is the only point which potentially does not have a neighborhood locally isometric to $\boldsymbol{R}^2$.

Note that two isometric Euclidean cones might have different descriptions. For instance, $\boldsymbol{R}^2$ can be broken into 4 quadrants or 8 sectors of angle $\pi/4$.

Exercise 1: Prove that two Euclidean cones are isometric if and only if they have the same angle.

Exercise 2: Define the unit circle in a Euclidean cone to be the set of points which are 1 unit away from the cone point. On the cone of angle 4Π find the shortest path between every pair of points on the unit circle. (This breaks into finitely many cases.)

2 Euclidean Cone Surfaces

Say that a compact surface Σ is a *Euclidean cone surface* if it has the following two properties:

- Every point $p \in \Sigma$ has a neighborhood which is isometric to a neighborhood of the cone point in a Euclidean cone of angle $\theta(p)$.
- We have $\theta(p) = 2\pi$ for all but finitely many points.

The points p where $\theta(p) \neq \pi$ are called the *cone points*. The quantity

$$\delta(p) = 2\pi - \theta(p)$$

is called the *angle deficit*. So, there are only finitely many points with nonzero angle deficit, and these deficits could be positive or negative.

Here are two examples:

- Let P be a convex polyhedron in $\boldsymbol{R}^3$. Then ∂P is a Euclidean cone surface. The metric on ∂P is the intrinsic one: The distance between two points is the length of the shortest curve which remains on ∂P and joins the points.
- Let $P_1, ..., P_n$ be a finite union of polygons. Suppose that these polygons can be glued together, isometrically along their edges, so that the result is a surface. Then the surface in question is a Euclidean cone surface if it is given its intrinsic metric−i.e. the shortest path metric.

Amazingly, every example of type 2 is also an example of type 1 provided that the underlying surface is a sphere and all the angle deficits are positive. This result is known as the Alexandrov Theorem. (To make this strictly true we have to allow for the possibility that P is contained in a plane in $\boldsymbol{R}^3$.) One interesting open problem is to determine the combinatorics of the convex polyhedron you get, based on the intrinsic geometry of the cone surface.

Exercise 3: Prove Alexandrov's theorem in case there are just 3 cone points (and the underlying space is a sphere.)

3 The Gauss-Bonnet Theorem

A *Euclidean triangle* on a Euclidean cone surface S is a region isometric to (you guessed it) a Euclidean triangle. For instance, on the boundary of a tetrahedron, there are 4 obvious Euclidean triangles. Two triangles on a cone surface *intersect* normally if they are either disjoint, or share a vertex, or share an edge. A *triangulation* of S is a decomposition of S into finitely many triangles, such that each pair of triangles intersects normally.

Exercise 4: Prove that every Euclidean cone surface has a triangulation. Hint: First draw lines between the cone points.

If S has been triangulated, then we define $\chi(S) = F + V - E$, where F is the number of triangles, E is the number of edges, and V is the number of vertices. This famous formula is known as the *Euler Characteristic*. It turns out that $\chi(S)$ only depends on S and not on the chosen triangulation. Here is a combinatorial version of the Gauss-Bonnet Theorem:

Theorem 3.1

$$\sum_p \delta(p) = 4\pi\chi(S).$$

Here the sum is taken over all angle deficits.

Proof: Let $T_1, ..., T_F$ be the list of triangles in the triangulation. Each T_i has associated to it three angles a_i, b_i, c_i, with $a_i + b_i + c_i = \pi$. The cone points are all at vertices of the triangles, and so

$$\sum_p \delta(p) = 2\pi V - (\sum_{i=1}^F a_i + \sum_{i=1}^F b_i + \sum_{i=1}^F c_i).$$

In other words, we add up all the angles and see hows the total sum differs from the expected $2\pi V$. Given that $a_i + b_i + c_i = \pi$ we have

$$\sum_p \delta(p) = 2\pi V - \pi F = 2\pi(V - F/2) =^* 2\pi(V + F - E) = 2\pi\chi(S).$$

The starred equality can be explained like this: Each triangle contributes 3/2 edges to the total number of edges. That is, $E = 3F/2 = F + F/2$. Hence $-F/2 = F - E$. ♠

4 Billiards and Translation Surfaces

Let P be a Euclidean polygon. A *billiard path* in P is the motion taken by an infinitesimal frictionless ball as it rolls around inside P, bouncing off the walls according to the laws of inelastic collisions: The angle of incidence equals the angle of reflection. The billiard path is *periodic* if it eventually repeats itself. Geometrically, a periodic billiard path corresponds to a polygonal path Q with the following properties:

- $Q \subset P$ (I mean the solid planar region.)
- The vertices of Q are contained in the interiors of the edges of P.
- A neighborhood of $P \cup Q$ in each vertex is isometric to a perfect K, with the crooked part being a subset of P and the straight part being a subset of Q.

Exercise 5: Find (with proof) all the examples of periodic billiard paths in a square which do not have self intersections. So, the path Q has to be embedded.

The polygon P is called *rational* if all its angles are rational multiples of π. For instance, the equilateral triangle is a rational polygon. A Euclidean cone surface is a *translation surface* if all the cone angles are integer multiples of π. In this section I'll explain how to associated a translation surface to a rational polygon. This is a classical construction, attributed by some people to Anatoly Katok. The geometry of the translation surface encodes a lot of the features of billiards in the polygon.

For each edge e of P there is a reflection R_e in the line through the origin parallel to e. Like all reflections, R_e has order 2. That is, $R_e \circ R_e$ is the identity map. Let G be the group generated by the elements $R_1, ..., R_n$. Here R_j stands for R_{e_j} and $e_1, ..., e_n$ is the complete list of edges. If e_i and e_j are parallel then $R_i = R_j$. If P is a rational polygon then there is some N such that e_j is parallel to some Nth root of unity. But then G is a group of order at most $2N$. In particular, G is a finite group.

Let $X = P \times G$. We think of X as a disjoint union of copies of P, one per element of G. We isometrically glue the edge (e_i, g) of (P, g) to the edge (e_i, h) of (P, h) if and only if $g = R_i h$. Note that $g = R_i h$ if and only if $h = R_i g$, because R_i has order 2. We call the resulting space $\widehat{P}$. Note that each edge of $P \times G$ is glued to exactly one other edge, because G is a group. Hence $\widehat{P}$ is a closed surface. By construction $\widehat{P}$ is a Euclidean cone surface. The only potential cone points are at the (equivalence classes) of vertices of X.

Lemma 4.1 *$\widehat{P}$ is a translation surface.*

Proof: To analyze the vertices of $\widehat{P}$, let (v, g) be a vertex of (P, g). Then p is incident to two consecutive edges of P, say e_1 and e_2. The polygon (P, g) is glued along (e_1, g) to the polygon $(P, R_1 g)$, which is then glued along $(e_2, R_1 g)$ to $(P, R_2 R_1 g)$, which is then glued along $(e_1, R_2 R_1 g)$ to $(P, R_1 R_2 R_1 g)$, and so on. This process continues until we reach the smallest k such that $(R_1 R_2)^k$ is the identity. If the angle at v is $\theta/2$ then the element $R_1 R_2$ is rotation by θ. We have $2k$ copies glued together around v and so the total cone angle around v is $k\theta$. But then $k\theta$ is a multiple of 2π because $(R_1 R_2)^k$ is the identity rotation. ♠

A path $\gamma \in \widehat{P}$ is called *straight* if every point $p \in \gamma$ has a neighborhood U with the following property: Any isometry between U and a subset of $\boldsymbol{R}^2$ maps $\gamma \cap U$ to a straight line segment. (For concreteness we can always take U to be a little Euclidean ball centered at p.) There is an obvious map $\pi : X \to P$. We just forget the group element involved. This forgetting respects the way we have done the gluing and so π is a well defined continuous map from $\widehat{P}$ to P. The map π is somewhat like a covering map, except that it is not locally a homeomorphism around points on the edges or vertices.

Lemma 4.2 *Suppose $\widehat{\gamma}$ is a straight path on $\widehat{P}$ which does not go through any vertices of $\widehat{P}$. Then $\gamma = \pi(\widehat{\gamma})$ is a billiard path on P.*

Proof: By construction γ is a polygonal path whose only vertices are contained in the interiors of edges of P. We just have to check the perfect K condition at each vertex. You can see why this works by building a physical model: Take a piece of paper and make a crease in it by folding it in half (and then unfolding it.) Now draw a straight line on the paper which crosses the crease. This straight line corresponds to a piece of $\widehat{\gamma}$ which crosses an edge. When you fold the paper in half you see the straight line turn back at the crease and form a perfect K. This folded path corresponds to γ. ♠

The converse is also true:

Lemma 4.3 *Suppose that γ is a billiard path on P. Then there is a straight path $\widehat{\gamma}$ on $\widehat{P}$ such that $\pi(\widehat{\gamma}) = \gamma$.*

Proof: We use the fact that the map π is almost a covering map. Think of γ as a parameterized path $\gamma : \boldsymbol{R} \to P$, with $\gamma(0)$ contained in the interior of P. We define $\widehat{\gamma}(0)$ to be the corresponding interior point of (P, g), where $g \in G$ is any initial element of G we like. We can define $\widehat{\gamma}(t)$ until the first value $t_1 > 0$ such that $\gamma(t_1)$ lies on an edge, say e_1, of P. But then we can define $\widehat{\gamma}$ in a neighborhood of t_1 in such a way that $\widehat{\gamma}(t_1 - s) \in (P, g)$ and $\widehat{\gamma}(t + s) \in (P, e_1 g)$ for $s > 0$ small. If you think about the folding construction described in the previous lemma, you will see that the straight path $\gamma(t_1 - \epsilon, t_1 + \epsilon)$ projects to $\widehat{\gamma}(t_1 - \epsilon, t_1 + \epsilon)$. Here ϵ is some small value which depends on the location of $\gamma(t_1)$. We can define $\widehat{\gamma}$ for $t > t_1$ until we reach the next time t_2 such that $\gamma(t_2)$ lies in an edge of P. Then we repeat the above construction for parameter values in a neighborhood of t_2. And so on. This process continues indefinitely, and defined $\widehat{\gamma}$ for all $t \geq 0$. Now we go in the other direction and define $\widehat{\gamma}$ for all $t < 0$. ♠

Note that $\widehat{\gamma}$ is a closed loop in $\widehat{P}$ if and only if γ is a periodic billiard path. Thus, the closed straight loops in $\widehat{P}$ correspond, via π, to periodic billiard paths in P.

Exercise 6: Suppose that P is the regular 7-gon. What is the genus of $\widehat{P}$?

Exercise 7: The same construction can be made when P has some irrational angles. What do you get if P is a right triangle with the two small angles irrational multiples of π?

5 Area Preserving Maps

We would like to understand how straight lines move around on P. In order to do this, we have to go a bit into some concepts from measure theory and dynamics. A subset $S \subset \boldsymbol{R}^2$ is called *null* if, for any $\epsilon > 0$, one can find an open subset $U \subset \boldsymbol{R}^n$ such that $S \subset U$ and U has volume less than ϵ. For instance, any countable set is null. In fact

Exercise 8: Prove that the countable union of null sets is null.

We would like to define null sets on $\widehat{P}$ but one technical problem is that $\widehat{P}$ isn't a subset of Euclidean space. The easiest approach is to recall that $\widehat{P}$ is the finite union of Euclidean polygons. Say that $S \subset \widehat{P}$ is null if its intersection with each of the polygons is null.

A subset $S \subset \widehat{P}$ is called *full* if $\widehat{P} - S$ is null. A *full map* on $\widehat{P}$ is a map $f : S \to \widehat{P}$, where both S and $f(S)$ are full. We say that f is *area preserving* if $f(V)$ has the same area as V provided that f is defined on all of V.

Remark: This definition presupposes that we have a well defined notion of area on P. The easiest way to do this is to recall that $\widehat{P}$ is the finite union of Euclidean polygons. We can take a set $S \subset \widehat{P}$, intersect it with each polygon, compute the area, and then add up the result.

If f is an area preserving full map then $f(S)$ and $f^2(S)$, etc. all are full. Hence $\bigcap f^k(S)$, the countable intersection of full sets, is also full. (Here we are using the complement of the result that the countable union of null sets is null.) Thus f and all its iterates are defined on a full set. We say that f is *invertible* if there is an area preserving full map g such that $f \circ g$ and $g \circ f$ are the identity map on a full set.

Here is a toy version of the famous Poincare Recurrence Theorem:

Theorem 5.1 *Let f be an area preserving full map on $\widehat{P}$. Let $p \in \widehat{P}$ be any point and let $\epsilon > 0$ be arbitrary. Then there is some $q \in \widehat{p}$ and some n such that $d(p, q) < \epsilon$ and $d(p, f^n(q)) < \epsilon$.*

Proof: Let Δ be the disk of radius ϵ about p. Let S be a full set on which f and all its iterates are defined. The infinite union of (almost) disks $D_i = f^n(\Delta \cap S)$ all have the same area and the total area of $\widehat{P}$ is finite. Hence there are two disks, say D_a and D_b, which intersect. But then $D_{a-1} = g(D_a)$ and $D_{b-1} = g(D_b)$ also intersect. But then $D_0 = \Delta \cap S$ intersects some D_n. Any $q \in D_0 \cap D_n$ satisfies the conclusions of the lemma. ♠

6 Existence of Periodic Billiard Paths

It is a theorem of Howie Masur that every rational polygon has a periodic billiard path. In fact, Masur gives bounds on the number of such billiard paths of length at most L. He proves that there are at least $L^2/C - C$ of them, and at most $CL + C$ of them, for some constant C which depends on the polygon. In some cases, it is possible to get sharper results. For instance:

Exercise 9: Prove that there is a constant C such that

$$\lim_{L \to \infty} N(L)/L^2 = C,$$

where $N(L)$ is the number of periodic billiard paths of length less than L on the unit square. What is C?

In this section, I'll sketch an elementary proof, due to Boshernitsyn, that every rational polygon has at least one periodic billiard path. You will see that the proof actually gives the existence of many periodic billiard paths, but no bounds like the ones mentioned above.

Choose a vector V on P which is perpendicular to one of the sides of P. We are going to use V to put a vector field on $\widehat{P}$ (minus the vertices.) Recall that $\widehat{P}$ is constructed from $X = P \times G$ by making some gluings. We define our vectorfield on (P, g) so that it is everywhere parallel to $g(V)$. This construction is such that the vectorfield patches together across the glued polygons. Also, the trajectories of the vector field are straight paths.

We would like to define a map $f : \widehat{P} \to \widehat{P}$ as follows: Given a point $p \in \widehat{P}$ we simple move 1 unit along the vector field. The problem with this definition is that the paths we take might go through the vertices, where the vectorfield is not defined. However, there is really just a countable set of bad straight paths, and the union of these is null. Thus f is really a full map. If we start with a little disk on $\widehat{P}$ and apply f to it, we are just translating this disk by 1 unit, all in the same direction. Hence f is area preserving. Clearly f is invertible: We can just flow one unit in the other direction.

Exercise 10: Prove that a countable union of straight paths on $\widehat{P}$ really is null.

Let $p \in \widehat{P}$ be some point. We think of p as the lift of $\gamma(0)$, where $\gamma : \boldsymbol{R} \to P$ is a billiard path which, at time 0, is travelling perpendicular to a side of P. That is, γ is travelling parallel to V at time 0. By the Poincare Recurrence Theorem we can find some q very close to p and some n so that $q\widehat{\beta}(0)$ and $f^n(q) = \widehat{\beta}(n)$ are very close together, and $\widehat{\beta}(0)$ is very close to $\widehat{\gamma}(0)$. Here $\widehat{\beta}(0)$ is a straight path in $\widehat{P}$ which goes through q at time 0.

If $\widehat{\beta}(0)$ and $\widehat{\gamma}(0)$ are sufficiently close then these two points are on the same polygon of $\widehat{P}$. Hence β and γ are travelling in the same direction at time 0. Likewise β is travelling in the same direction at times 0 and n. In short, β travels perpendicular to a side of P at time 0 and also at some much later time n. This means that β hits the same side of P twice, and both times at right angles. But then β is periodic. Each time it hits P perpendicularly, β just reverses itself and retraces its path.