1 Translation Surfaces

These notes are a strict subset of Chapter 17 of my book, Mostly Surfaces, lightly edited to make them self-contained. I'm not going to talk about the Gauss Bonnet Theorem or billiards in class, but I kept this stuff in because I thought that some people might like it.

1.1 Sectors and Euclidean Cones

A *sector* in $\boldsymbol{R}^2$ is the closure of one of the 2 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, the nonnegative 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 sector. For instance, the angle of the nonnegative 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 – e.g., $\boldsymbol{R}^2$ 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 problem breaks down into finitely many cases, depending on where the points are located.

Exercise 3. Let C be a Euclidean cone, with cone point x. Say that a vector field on $C - x$ is parallel if an isometry carrying any open set of $C - x$ into $\boldsymbol{R}^2$ carries the vector field to a constant vector field. Prove that $C - x$ has a parallel vector field in a neighborhood of x if and only if the cone angle of C is a multiple of 2π. (*Hint*: Unroll C into the plane and watch the vector field as you go once around the cone point.)

1.2 Euclidean Cone Surfaces

A surface is oriented if it does not contain any Möbius bands. For ease of exposition, we only consider oriented surfaces.

Say that a compact oriented 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 $\theta(p) \neq 2\pi$, 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 $P_1, \ldots, P_n$ 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.

1.3 The Gauss–Bonnet Theorem

Here is combinatorial version of the Gauss–Bonnet Theorem:

Theorem 1.1 *If S is a compact cone surface, then*

$$\sum_p \delta(p) = 2\pi\chi(S).$$

Here the sum is taken over all angle deficits.

Proof: 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 maximal 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.

Choose a triangulation of S such that the triangles in the triangulation are Euclidean. Let $T_1, \ldots, T_F$ be the list of triangles in the triangulation. Each T_i has associated to it three angles a_i, b_i, c_i, with $a_i + b_i + c_i = \pi$. Since we picked a Euclidean triangulation, 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 how 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 has the following explanation. For each triangle we get 3 edges, but each edge is shared between exactly 2 triangles. That is, $E = 3F/2 = F + F/2$. Hence $-F/2 = F - E$. ♠

For comparison, we mention that the differential geometric version of the Gauss–Bonnet Theorem says that the total curvature of a surface S is $2\pi\chi(S)$, where χ is the Euler characteristic of S. One can view each version as the limit of the other.

1.4 Translation Surfaces

A *translation surface* is a Euclidean cone surface which admits a parallel vector field which is defined everywhere except at the cone points. By Exercise 3 above, the cone angles of a translation surface are all integer multiples of 2π.

At first it might seem that a Euclidean surface whose cone angles are all integer multiples of 2π must admit a parallel vector field, but this is not so. As Rick Kenyon pointed out to me, M. Troyanov constructed some counter examples. See "Les surfaces euclidienne a singularites coniques", by M. Troyanov, published in Enseign. Math (2) 32 (1986), 76-94. You might like to try to find some examples yourself without looking up Troyanov's article.

Recall that a *gluing diagram* for a surface is a list of finitely many polygons, together with a recipe for gluing together the sides of the polygon in pairs.

Lemma 1.2 *Suppose that S is a flat cone surface obtained from a gluing diagram in which the two sides in each glued pair are parallel. Then S is a translation surface.*

Proof: Once we show that S is orientable, we will know that S is a cone surface. On each polygon, we consider the standard pair of vector fields V_1 and V_2. Here V_j consists of vectors parallel to the basis vector e_j. Given the nature of the gluing maps, the vector fields piece together across the edges to give parallel vector fields V_1 and V_2 defined on the complement of finitely many points.

We first show that S is orientable. If S is not orientable, then S contains a Möbius band M. By shrinking M if necessary, we can arrange that M lies entirely in the region where both V_1 and V_2 are defined. But then we can define a continuous pair of linearly independent vector fields on a Möbius band. This is easily seen to be impossible. Hence S is oriented.

It now follows from definition that S is a translation surface. ♠

In light of Lemma 1.2, the surface obtained by gluing (with translations) the opposite sides of a regular $2n$-gon is a translation surface.

Translation Principle. Whenever we consider gluing diagrams for translation surfaces, in which more than one polygon is involved, we always think of the polygons in the plane as being pairwise disjoint. How the polygons sit in the plane is really not so important, in the following sense. Suppose that $P_1, \dots, P_n$ are the polygons involved in a gluing diagram for some surface. Suppose that $Q_1, \dots, Q_n$ are new polygons, such that Q_k is a translation of P_k for all k, and the pattern of gluing for the Q's is the same as the pattern of gluing for the Ps. Then the two resulting surfaces are canonically isometric. The canonical isometry is obtained by piecing together the translations that carry each P_k to Q_k. We mention this rather obvious principle because it guarantees that certain constructions, which seem based on arbitrary choices, are actually well defined independent of these choices.

1.5 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; see Figure 1 below. We make a convention that a path stops if it lands precisely at a vertex. (The infinitesimal ball falls into the infinitesimal pocket.)

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:

- $Q \subset P$ (that is, the solid planar region).
- The vertices of Q are contained in the interiors of the edges of P.
- Q obeys the angle of incidence rule discussed above.

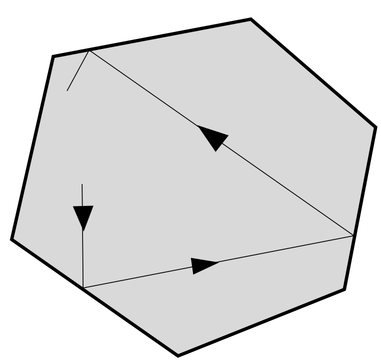

Figure 1. Polygonal billiards

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.

In this section I will explain how to associate a translation surface to a rational polygon. This is a classical construction, attributed by some people to A. Katok and A.N. Zemylakov. The geometry of the translation surface encodes many 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 $R_1, \ldots, R_n$. Here R_j stands for R_{e_j} and $e_1, \ldots, e_n$ 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, after we suitably rotate P, 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.

For each $g \in G$, we define a polygon

$$P_g = g(P) + V_g. \tag{1}$$

Here V_g is a vector included so that all the polygons $\{P_g |\ g \in G\}$ are disjoint. Thanks to the Translation Principle, the surface we will produce is independent of the choices of the translation vectors.

To form a gluing diagram, we declare that every two edges of the form

$$e_1 = g(e) + V_g, \qquad e_2 = gr(e) + V_{gr}, \qquad r = R_e. \tag{2}$$

are glued together by a translation. Here e is an arbitrary edge of P. Since $gr(e) = g(e)$, the edges e_1 and e_2 are parallel. Hence, it makes sense to glue them by a translation. Note also that $(gr)r = g$. So, our instructions tell us to glue e_1 to e_2 if and only if they tell us to glue e_2 to e_1. Let $\widehat{P}$ be the space obtained from the gluing diagram. Since the edges are glued in pairs, $\widehat{P}$ is a surface. By Lemma 1.2, $\widehat{P}$ is a translation surface.

Here we work out the example where P is an isosceles triangle with small angles $\pi/8$. In this case, the group G has order 16 and our surface will be made from 16 isometric copies of P.

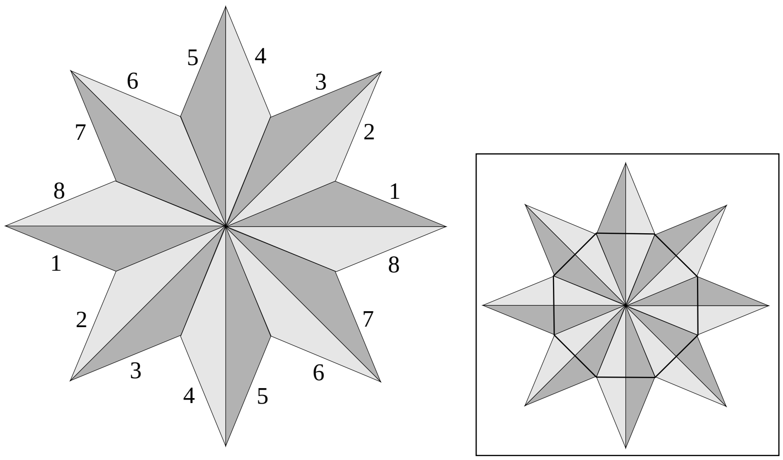

Figure 2. Gluing diagram for a translation surface

Figure 2 shows the resulting gluing diagram. We have chosen the translations so that all the long sides have already been glued together. Also, we have colored the triangles alternately light and dark so as to better show the pattern. The numbers around the outside of the figure indicate the gluing pattern for the short edges.

The gluing pattern in Figure 2 has an alternate description. Take two regular Euclidean octagons and glue each side of one to the opposite side of the other. The smaller inset picture in Figure 2 shows one of the two octagons. The other octagon is splayed open, and made by gluing together the pieces that are outside the octagon shown.

Let $\widehat{P}$ be the translation surface constructed above. 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 $\pi : \widehat{P} \to P$. 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 1.3 *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 angle incidence 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 bounce like a billiard path. This folded path corresponds to γ. ♠

The converse is also true:

Lemma 1.4 *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 parametrized 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 $\widehat{\gamma}(t + s) \in P_{rg}$ for $s > 0$ small, where r is reflection over side e_1. If you think about the folding construction described in the previous lemma, you will see that the straight path $\widehat{\gamma}(t_1 - \epsilon, t_1 + \epsilon)$ projects to $\gamma(t_1 - \epsilon, t_1 + \epsilon)$. 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 defines $\widehat{\gamma}$ 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 Euler characteristic of $\widehat{P}$? As a much harder problem, can you find a formula for the Euler characteristic of $\widehat{P}$ as a function of the angles of 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?