8. Billiards in curved domains

Billiards makes sense in regions with curved boundaries. Some of the things we've talked about (like phase space) also work in this context. However, we like to focus on what's new.

- We discuss the law of reflection for curved boundaries, and playing billiards in an ellipse.
- We use the extremal-length interpretation of the law of reflection to give an abstract existence criterion for periodic orbits.
- We return to phase space, and talk about how that can be used to understand optical devices consisting of curved mirrors.

(8a) The law of reflection. When light bounces off a curved mirror, the law of reflection says that the incoming and outgoing angles between the ray of light and the tangent line to the mirror must be equal. This is what we will choose as the behaviour for billiards in curved billiards tables.

One can derive the equal-angle reflection law from an extremal (shortest path) principle. Namely, suppose that we consider a curved mirror parametrized as $c(t) \in \mathbb{R}^2$. Fix two points p and q that do not lie on our mirror. Let's suppose that we go from p to a point $c(t)$ on the curve, and then from there to q, each time along a straight line.

(8.1)

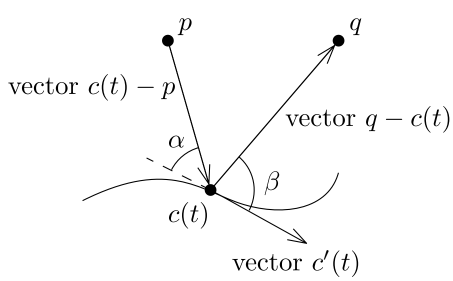

The total length of this path is

$$S(t) = \|c(t) - p\| + \|c(t) - q\| = \sqrt{(c(t) - p) \cdot (c(t) - p)} + \sqrt{(c(t) - q) \cdot (c(t) - q)}. \tag{8.2}$$

We differentiate this, and use the product rule for the scalar product:

$$S'(t) = \frac{(c(t) - p) \cdot c'(t)}{\|c(t) - p\|} + \frac{(c(t) - q) \cdot c'(t)}{\|c(t) - q\|}. \tag{8.3}$$

Let's divide this by $\|c'(t)\|$ and rewrite it as

$$\frac{S'(t)}{\|c'(t)\|} = \frac{c(t) - p}{\|c(t) - p\|} \cdot \frac{c'(t)}{\|c'(t)\|} - \frac{q - c(t)}{\|q - c(t)\|} \cdot \frac{c'(t)}{\|c'(t)\|}. \tag{8.4}$$

Now, all the vectors in the scalar products have length one, so the scalar product is the cosine of the angles α and β. In particular, $S'(t) = 0$ if and only if the angles are equal. Let's remember this:

PROPOSITION 8.1. *$S'(t) = 0$ if and only if the trajectory from p to $c(t)$ to q satisfies the equal-angle law.*

Let's look at the special case of an ellipse. Take two points p and q, which will be the foci of the ellipse, and some number s which is bigger than the distance between them. The ellipse is then defined as the set of points $v \in \mathbb{R}^2$ such that

$$\|v - p\| + \|v - q\| = s. \tag{8.5}$$

Consider paths from p to a point on the ellipse and then back to q. By definition, the length of such path is then independent of which point on the ellipse we pick. In other words, if we parametrized the ellipse in some way, then $R(t) = r$ would be a constant function. This means that the assumption from Proposition 8.1 is always satisfied:

Fact 8.2. *If a billiards trajectory in the ellipse starts at one focus, then after bouncing off once it will reach the other focus.*

The trajectories that keep hitting foci have a characteristic behaviour. Namely, as we extend them far into the future, they keep getting closer and closer to the major axis, which is the line between the two foci. (An extreme special case is the periodic trajectory that just keeps bouncing back and forth along that axis.)

(8.6)

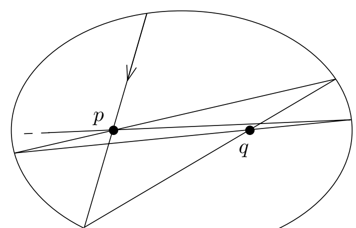

Of course, these are not all the trajectories in the ellipse: there are lots of others, which never hit a focus.

Now we leave the ellipse behind, and return to the general question of periodic orbits, for a large class of billiards tables.

Theorem 8.3. *Suppose that our billiards region has no corners, and is strictly convex. Then, for every $n \geq 2$, there is (at least) one periodic billiards trajectory with n bounces.*

No corners means that the boundary of the region is smooth everywhere. The other condition is:

Definition 8.4. Strict convexity *means that if you have two points on the boundary of our region, the line segment connecting them stays inside our region, and doesn't touch the boundary anywhere except at its endpoints (so, a polygon can never be strictly convex, but a curved region can be).*

The proof of the theorem is remarkably easy: consider all possible polygonal loops with n vertices, with the vertices lying on the boundary of our region. Clearly, there's an upper bound on the length of such loops. Analysis tells us that the maximum is achieved by some loop. When we derived the equal-angle law, we used only vanishing of the derivative of the length. Hence, that reasoning applies not just to the shortest path, but also the longest one. This shows that the longest path must satisfy the equal-angle law, hence is a billiards trajectory. You can see where the argument would go wrong for polygonal billiards: the longest loop can be one that goes through corners, so it's not well-defined as a billiards trajectory.

(8b) The limitations of optical devices. The phase space picture also applies to billiards in curved domains. Suppose we have a billiards table whose boundary consists of curved segments (corners are allowed). One defines phase space in $(s, t = -cos(\theta))$ coordinates again as

$$\Omega = \bigsqcup_e (0, \text{length}(e)) \times (-1, 1). \tag{8.7}$$

where e are the segments of the boundary, and length(e) is the arclength. The second coordinate is $t = -\cos(\theta)$, where $\theta \in (0, \pi)$ measures the angle between with respect to the tangent line at the appropriate boundary point. Everything we have said, including the billiards transformation T and its area-preserving property (Liouville's theorem), still applies.

There are implications of this in optics (in our discussion, the world is two-dimensional, but the three-dimensional world also satisfies a version of that). Take a box with two holes ("input" and "output") of length l and m, respectively. Inside the box, we arrange N curved mirrors. What we want is that *all* the light coming through the input hole, at an angle of at most α from the perpendicular direction, bounces off our mirrors in fixed order $1, 2, \dots, N$, and then leaves through the output hole, at an angle of at most β from the perpendicular direction:

(8.8)

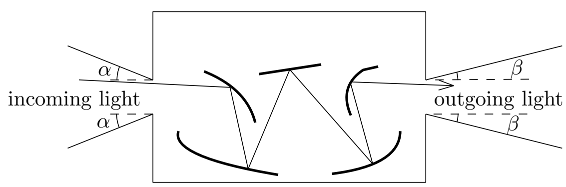

You can take any n, and make the mirrors of arbitrarily curved shapes. When is this possible? The first insight is this:

Proposition 8.5. *If the input and output lengths are equal, $l = m$, then the angles must satisfy $\beta \geq \alpha$; you can't squeeze the light into a smaller angle!*

Proposition 8.6. *If the input and output angles are equal, $\alpha = \beta$, then the lengths must satisfy $l \leq m$: you can't make the output hole smaller than the input!*

The idea is to think of the whole thing as a curved billards table, by joining the mirrors with arbitrary other pieces, and where the input and output holes are straight-line parts of the boundary.

(8.9)

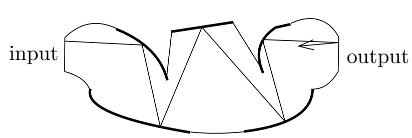

The incoming light can be thought of as a billiards trajectory starting at an input line, and whose angular coordinate in phase space is constrained to

$$\begin{aligned} &\theta \in (\pi/2 - \alpha, \pi/2 + \alpha), \text{ or equivalently,} \\ &t \in (-\cos(\pi/2 - \alpha), -\cos(\pi/2 + \alpha)) = (-\sin(\alpha), \sin(\alpha)). \end{aligned} \tag{8.10}$$

In other words, those starting points form a region in phase space,

$$U = (0, l) \times (-\sin(\alpha), \sin(\alpha)). \tag{8.11}$$

After $N + 1$ applications of the billiards map, we end up in a similar rectangle describing how one would bounce off the straight line corresponding to the output end,

$$V = (0, m) \times (-\sin(\beta), -\sin(\beta)). \tag{8.12}$$

Our requirement that this should work for all light rays then means that the billiards transformation T should satisfy

$$T^{N+1}(U) \subset V. \tag{8.13}$$

Since areas are preserved under T, this can only happen if

$$\mathrm{area}(U) = 2l\sin(\alpha) \leq \mathrm{area}(V) = 2m\sin(\beta). \tag{8.14}$$

Special cases of this ($l = m$, or $\alpha = \beta$) explain the two Propositions above. We have actually proved more:

PROPOSITION 8.7. *In general, for the optical contraption to be theoretically possible, we need*

$$\frac{l}{m} \leq \frac{\sin(\beta)}{\sin(\alpha)}. \tag{8.15}$$

In words, you can squeeze the angles ($\beta < \alpha$) but only by making the output opening larger than the input one ($l < m$); or you can make the output smaller than the input, but have to tolerate a larger output angle; and there are precise quantitative bounds on that. One could check by computation see that a specific system (say, a single circular-piece mirror) satisfies those constraints. The remarkable thing is that they apply to systems of mirrors of any shape.