Snapshots of modern mathematics from Oberwolfach
№ 1/2015

Billiards and flat surfaces

Diana Davis

Billiards, the study of a ball bouncing around on a table, is a rich area of current mathematical research. We discuss questions and results on billiards, and on the related topic of flat surfaces.

1 Let's start playing!

Wherever you are sitting to read this, molecules of oxygen and other gas particles are bouncing against each other and against the walls of the room. Understanding the movements of all those particles would be interesting, but very complicated. A common strategy that mathematicians use is to simplify a problem as much as possible, and try to understand the simpler system. In this case, we could simplify it so that there are no collisions or other interactions between molecules, by studying just one particle bouncing around the room and ignoring all the others. In fact, we could simplify it even further, to two dimensions instead of three, a particle bouncing around on a surface, just like a ball on a billiards or pool table. It turns out that this is already very interesting.

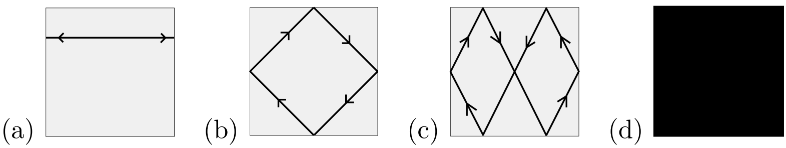

Figure 1: Different paths on a square billiard table.

Consider the simplest case, when the billiard table is a square. We will assume that the ball (or particle) is just a point, moving with no friction (so it goes forever), and that when it hits the edge of the table, the angle of reflection is equal to the angle of incidence (just as in real life).

Is it possible to hit the ball so that it repeats its path? Yes: if we hit it vertically or horizontally, it will bounce back and forth between two points on parallel edges, see Figure 1 (a). We say that this *trajectory* (path) is *periodic*, with *period* 2 (it bounces off the edge of the table two times before retracing its own steps). Other examples, with periods 4 and 6, respectively, are in Figure 1 (b) and Figure 1 (c).

Is it possible to hit the ball so that it *never* repeats its path? It is more difficult to draw a picture of an example of such a *non-periodic* trajectory because the trajectory never repeats. In fact, it will gradually fill up the table until the picture is a black square (Figure 1 (d)). However, if we don't restrict ourselves to just the table, we can draw a picture of the trajectory by *unfolding* the table:

Consider the simple trajectory in Figure 2 (a) below. When the ball hits the top edge, instead of having it bounce and go downwards, we *unfold* the table upward, creating another copy of the table in which the ball can keep going straight (Figure 2 (b)).[1] Now when the trajectory hits the right edge, we do the same thing: we unfold the table to the right, creating another copy of the table, in which the ball can keep going straight (Figure 2 (c)). We can keep doing this, creating a new square every time the trajectory crosses an edge. In this way, a trajectory on the square table is represented as a line on a piece of graph paper.

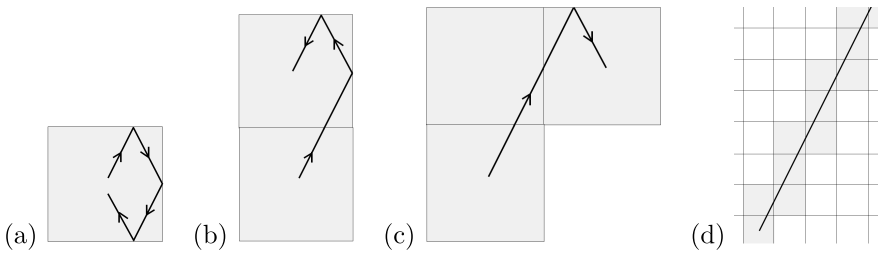

Figure 2: Unfolding the billiard table.

By thinking of the trajectory as a line on graph paper, we can easily find an example of a non-periodic trajectory. Suppose that we draw a line with an *irrational* slope. Then it will never cross two different horizontal (or vertical) edges at the same point: if it did, then the slope between those corresponding

[1] It is best to think of a transparent table, so that you don't have to worry about whether the ball is now on the "upside" or the "downside" of the table.

points would be a ratio of two integers. But we chose the slope to be irrational, so this can't happen. Hence, if we hit the ball with *any* irrational slope, its trajectory on the table will be non-periodic. By a similar argument, if we hit the ball with any *rational* slope, its trajectory on the table will be periodic.

2 Unusual billiard tables

These are fundamental questions to ask about any dynamical system: Does it have periodic behavior? Does it have non-periodic behavior? We have seen that both behaviors can occur in the simplest case, that of the square table. Stretching the square table horizontally or vertically preserves the rules of the system, so the same results are true for all rectangular tables. Now perhaps we are ready to make our system more complicated. We can ask the same questions for any other table shape − for example, a triangular table, a pentagonal table, or a table made from several rectangles glued together. A trajectory on a pentagonal table with many bounces is shown in Figure 3. (Can you figure out what a trajectory on a circular table would look like?)

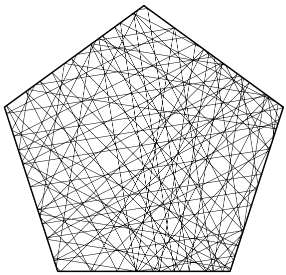

Figure 3: A trajectory on a pentagonal table.

For the square table, we mentioned above that every trajectory is either periodic (we can draw a picture), or its path eventually covers the table (the picture is black). Does this always happen? Surprisingly, the answer is no! There are tables where some trajectory completely covers one region of the table, but never goes to another region of the table: An illustration of this phenomenon, constructed by Curtis McMullen, is shown in Figure 4. The trajectory gets "trapped" in the rectangle and a portion of the square, and never visits the other two corners of the square. If we let the ball keep going, the shaded region would become completely black, but the corners would stay white.

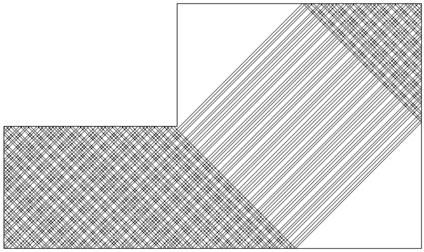

Figure 4: A non-periodic trajectory that never enters the white corners.

This raises the question: which table shapes have the property that every non-periodic trajectory covers the whole table? It turns out that most of them don't: this is only true of tables that have a lot of symmetry, such as regular polygons, tables made from multiple squares glued together, and some simple triangles. People are trying to find more examples of table shapes that have this property, but so far, only these few families of examples are known [1, 5, 6].

3 Playing billiards on a donut

Let's return to the unfolded billiard table. We unfolded the top edge of the table, creating another copy in which the ball can keep going straight. The edge that the ball would have bounced off of no longer acts as a wall, so we make it a dotted line (Figure 5). The new top edge is just a copy of the bottom edge, so we now label them both A to remember that they are the same. Similarly, we unfolded the right edge of the table, creating another copy of the unfolded table, which gives us four copies of the original table. The new right edge is a copy of the left edge, so we now label them both B. When the trajectory hits the top edge A, it just reappears in the same place on the bottom edge A and keeps going. Similarly, when the trajectory hits the right edge B, it reappears on the left edge B. This is called *identifying* the top and bottom edges, and *identifying* the left and right edges.

You may be familiar with this idea of entering the right wall and re-emerging from the left wall from video games like "Pac-Man", "Snake", or "Portal". However, you may not have realized that with these edge identifications, you are no longer on the flat plane, but on an entirely different flat surface! It turns out that this surface is actually the surface of a bagel or a donut, which mathematicians call a *torus*.

It is a fun exercise to see why this is the case. You can see that if you glue both copies of edge A to each other, you get a cylinder, whose ends are both

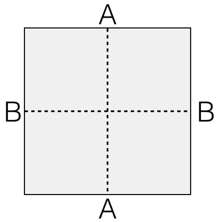

Figure 5: Identifying the edges of a square.

edge B (Figure 6). When you wrap the cylinder around to glue both copies of edge B to each other, you get the torus surface (you will have to stretch and compress the cylinder considerably, but no worries – you may think of it as being very elastic).

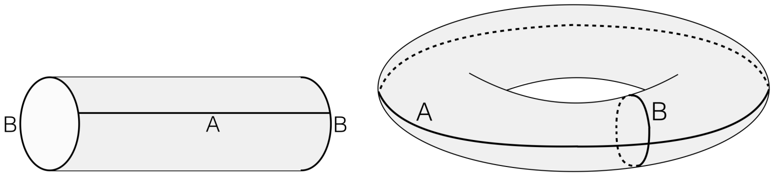

Figure 6: The gluing of a torus.

We can create many other surfaces in the same way (see Figure 7). For example, if we take a regular octagon and identify each pair of parallel edges, we get another surface, actually a torus with two holes [4]. We could also take two pentagons, one of which has been rotated by 180°, and identify the five pairs of parallel edges.[2] It turns out that this surface is also the two-holed torus.

Can you visualize how to roll up and glue the identified edges of the octagon, in the same way that we did it for the square? It is much more difficult – you have to squeeze the octagon considerably when fitting its corners together. Still, it is possible to do it; there is a nice picture in Allen Hatcher's *Algebraic Topology* [3, p. 5].

[2] I made a video that uses dance to rigorously explain my research on the double pentagon: http://vimeo.com/47049144.

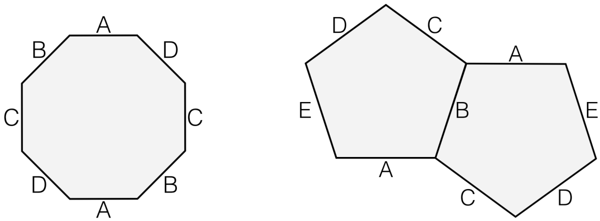

Figure 7: An octagon or two pentagons can be glued into a two-holed torus.

4 Twisting the surface

My research, on the double pentagon surface and other related surfaces, is on the following problem: Suppose that you have a trajectory on a surface. Then you cut the surface, twist it, and put it back together in such a way that the cut edges are glued back together again. What happens to the trajectory?

On the torus, this amounts to slicing the surface open, twisting it one, two, three, or more times (it gets twisted once in Figure 8), and then gluing the ends back together. We can see that if we start with a very simple trajectory (an equator of the torus), twisting the surface makes the trajectory go around the center hole.

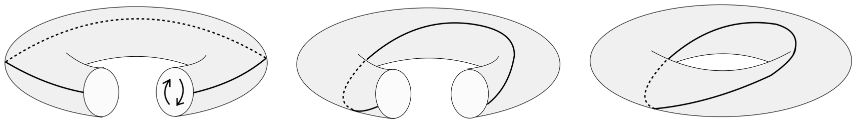

Figure 8: Cutting, twisting, and gluing back together the torus.

Even with this simplest example, it is a bit difficult to see the trajectory on the torus surface. You can see in Figure 9 that it is much easier to draw it on the square, both the original trajectory in Figure 9 (a) and the twisted trajectory in Figure 9 (b). For more complicated surfaces and more complicated trajectories (see Figure 9 (c)), it is far easier to study this twisting action on the polygons than it would be to try to draw a three-dimensional representation of the surface.

Still, we need a way to refer to a trajectory on a surface. It is convenient to do this by writing down the edge labels that the trajectory crosses: the trajectory in Figure 9 (a) is $\ldots BBB \ldots$ or $\overline{B}$ for short, the trajectory in Figure 9 (b) is $\ldots ABABAB \ldots$ or $\overline{AB}$, and the trajectory in Figure 9 (c) is $\overline{EBECDC}$. These are called the *cutting sequences* corresponding to each of the trajectories.

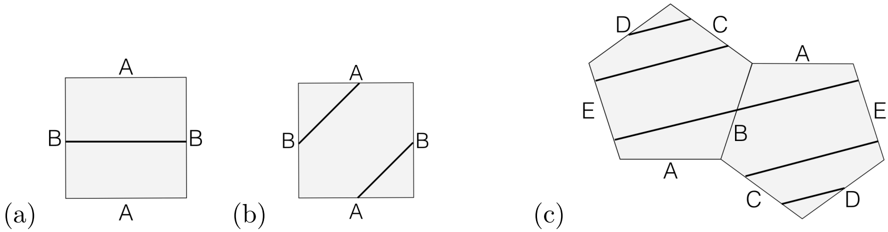

Figure 9: Twisting a trajectory on a square (a) - (b) and a trajectory on two pentagons (c).

You might ask, given a string of As and Bs, how do you know if it is the cutting sequence of some trajectory on the square torus? Can you think of a string of As and Bs that is *not* a cutting sequence on the square torus?[3]

These cutting sequences of As and Bs on the square torus are beautiful and interesting. They are related to continued fractions: you can use the cutting sequence of a trajectory to find the continued fraction expansion for the slope of that trajectory. Even if the slope is irrational, there is an algorithm to give closer and closer approximations to the slope. If the trajectory is periodic, can you figure out how to use the number of As and Bs in one period of its cutting sequence to determine the (rational) slope of the trajectory?[4]

The subject of my research is to describe all possible cutting sequences for the double pentagon surface and for related surfaces: which sequences of As, Bs, Cs, Ds and Es represent actual trajectories on the double pentagon? Twisting the surface, as described above, helps to answer this question, because we start with a trajectory that we already know, and then twisting the surface gives us a new trajectory on the surface [2]. John Smillie and Corinna Ulcigrai have completely described all possible cutting sequences on the regular octagon surface, and even found an analog of the continued fraction expansion for the direction of the trajectory [4]. Unfortunately, the answer is much more difficult to describe than the result for sequences on the square torus, but that is not surprising; the simplest examples are often the most elegant ones.

The study of flat surfaces[5] is a rich area of current research. Maryam Mirzakhani recently received the Fields Medal, the highest honor in mathematics, for her work on areas related to billiards and flat surfaces. However, Mirzakhani

---

[3] Solution: any sequence containing both AA and BB, for example.

[4] Solution: slope $= \frac{\text{rise}}{\text{run}} = \frac{\text{number of } As}{\text{number of } Bs}$.

[5] Surfaces as the ones we have seen above are often called flat because they arise from gluing edges in polygons that are part of the flat plane, in such a way that a trajectory is always parallel to itself.

does not just study the dynamics of a ball on one billiard table or surface, or even on a group of surfaces; she essentially studies the space of *all possible* surfaces. The study of such spaces, and the study of billiards, flat surfaces, and other related questions, are collectively called *dynamics*. This is a relatively new field in mathematics, in which many people are doing research.

Image credits

Figures 3 and 4 were provided by Curtis McMullen.

References

[1] Irene Bouw and Martin Möller, *Teichmüller curves, triangle groups, and Lyapunov exponents*, Annals of Mathematics **172** (2010), no. 1, 139–185.

[2] Diana Davis, *Cutting sequences, regular polygons, and the Veech group*, Geometriae Dedicata **162** (2013), no. 1, 231–261.

[3] Allen Hatcher, *Algebraic topology*, Cambridge University Press, 2002.

[4] John Smillie and Corinna Ulcigrai, *Beyond Sturmian sequences: coding linear trajectories in the regular octagon*, Proceedings of the London Mathematical Society **102** (2011), no. 2, 291–340.

[5] William Veech, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Inventiones Mathematicae **87** (1989), 553–583.

[6] Clayton Ward, *Calculation of Fuchsian groups associated to billiards in a rational triangle*, Ergodic Theory and Dynamical Systems **18** (1998), no. 4, 1019–1042, http://journals.cambridge.org/article_S0143385798117479.

Diana Davis *is a postdoc at Northwestern University (IL).*

Mathematical subjects
Geometry and Topology

Connections to other fields
Physics

DOI
10.14760/SNAP-2015-001-EN

Snapshots of modern mathematics from Oberwolfach are written by participants in the scientific program of the Mathematisches Forschungsinstitut Oberwolfach (MFO). The snapshot project is designed to promote the understanding and appreciation of modern mathematics and mathematical research in the general public worldwide. It is part of the mathematics communication project "Oberwolfach meets IMAGINARY" funded by the Klaus Tschira Foundation and the Oberwolfach Foundation. All snapshots can be found on www.imaginary.org and on www.mfo.de/snapshots.

Mathematisches Forschungsinstitut
Oberwolfach gGmbH
Schwarzwaldstr. 9 – 11
77709 Oberwolfach
Germany

Director
Gerhard Huisken

Mathematisches Forschungsinstitut Oberwolfach
Member of the Leibniz Association
Klaus Tschira Stiftung gemeinnützige GmbH
KTS
oberwolfach FOUNDATION
IMAGINARY open mathematics