Geometry and the Imagination

Based on materials from the course taught
at the University of Minnesota Geometry Center
in June 1991

by John Conway, Peter Doyle, Jane Gilman, and Bill Thurston

Version 0.93, Fall 1999
Derived from works

Contents

1 Bicycle tracks

C. Dennis Thron has called attention to the following passage from *The Adventure of the Priory School*, by Sir Arthur Conan Doyle:

> 'This track, as you perceive, was made by a rider who was going from the direction of the school.'

‘Or towards it?’

‘No, no, my dear Watson. The more deeply sunk impression is, of course, the hind wheel, upon which the weight rests. You perceive several places where it has passed across and obliterated the more shallow mark of the front one. It was undoubtedly heading away from the school.’

Problems

1. Discuss this passage. Does Holmes know what he’s talking about?

2. Try to come up with a method for telling which way a bike has gone by looking at the track it has left. There are all kinds of possibilities here. Which methods do you honestly think will work, and under what conditions? For example, does your method only work if the bike has passed through a patch of wet cement? Would it work for tracks on the beach? Tracks on a patch of dry sidewalk between puddles? Tracks through short, dewy grass? Tracks along a thirty-foot length of brown package-wrapping paper, made by a bike whose tires have been carefully coated with mud, and which has been just ridden long enough before reaching the paper so that the tracks are not appreciably darker at one end of the paper than the other?

3. Try to determine the direction of travel for the idealized bike tracks in Figure 1.

4. Try to sketch some idealized bicycle tracks of your own. You don’t need a computer for this; just an idea of what the relationship is between the track of the front wheel and the track of the back wheel. How good do you think your simulated tracks are?

5. Go out and observe some bicycle tracks in the wild. Can you tell what way the bike was going? Keep your eye out for bike tracks, and practice until you can determine the direction of travel quickly and accurately.

6. Tell your friends about this problem. Don’t give them the answers; just the questions. Let them think for themselves. Let them stew a little. See what they come up with. You can give them a little help, if need be, once they’ve spent enough time thinking about it for themselves to

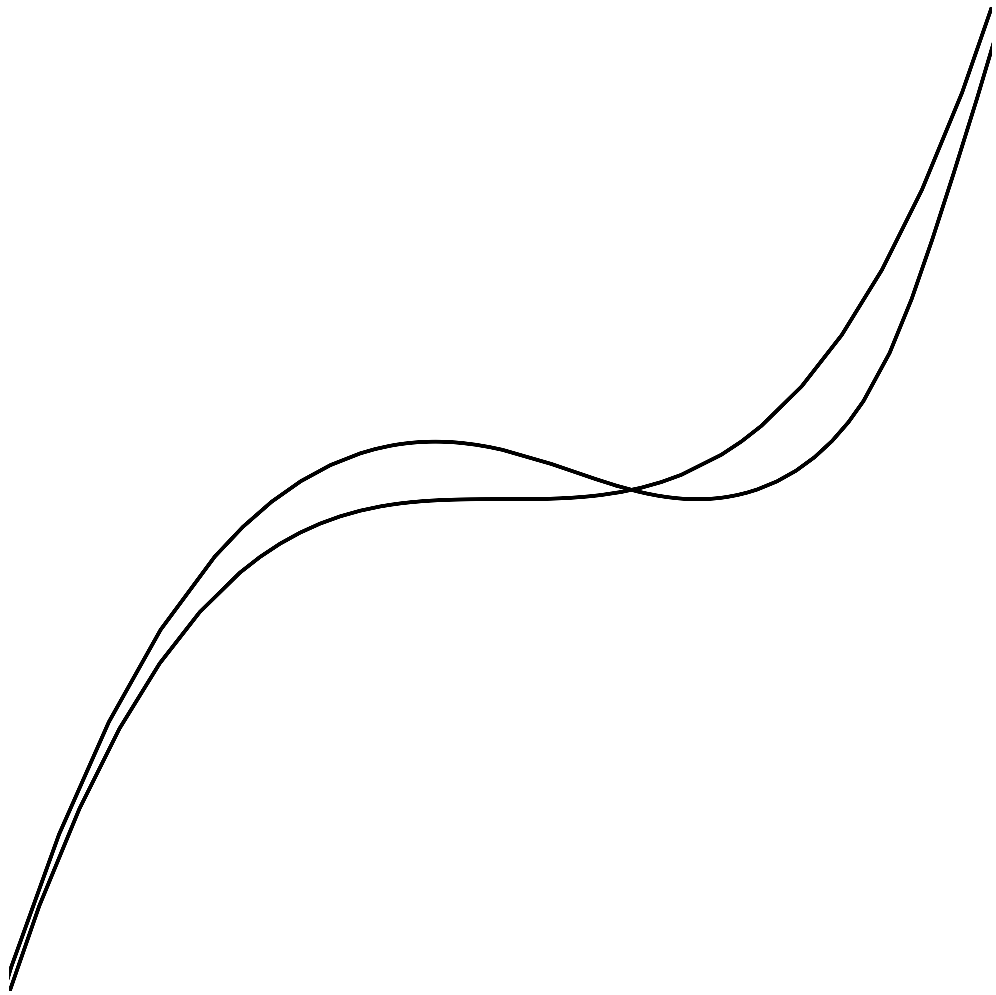

Figure 1: Which way did the bicycle go?

appreciate the problem. Make sure that they, too, finally master this useful trick of bike direction-finding.

2 Pulling back on a pedal

Imagine that I am steadying a bicycle to keep it from falling over, but without preventing it from moving forward or back if it decides that it wants to. The reason it might want to move is that there is a string tied to the right-hand pedal (which is to say, the right-foot pedal), which is at its lowest point, so that the right-hand crank is vertical. You are squatting behind the bike, a couple of feet back, holding the string so that it runs (nearly) horizontally from your hand forward to where it is tied to the pedal.

Problems

1. Suppose you now pull gently but firmly back on the string. Does the bicycle go forward, or backward? Remember that I am only steadying it, so that it can move if it has a mind to. No, this isn't a trick; the bike really does move one way or the other. Can you reason it out? Can you imagine it clearly enough so that you can *feel* the answer intuitively?

2. Try it and see.

John Conway makes the following outrageous claim. Say that you have a group of six or more people, none of whom have thought about this problem before. You tell them the problem, and get them all to agree to the following proposal. They will each take out a dollar bill, and announce which way they think the bike will go. They will be allowed to change their minds as often as they like. When everyone has stopped waffling, you will take the dollars from those who were wrong, give some of the dollars to those who were right, and pocket the rest of the dollars yourself. You might worry that you stand to lose money if there are more right answers than wrong answers, but Conway claims that in his experience this *never* happens. There are always more wrong answers than right answers, and this despite the fact that you tell them in advance that there are going to be more wrong answers than right answers, and allow them to bear this in mind during the waffling process.

(Or is it *because* you tell them that there will be more wrong answers than right answers?)

Problem

- Try this trick for yourself, and see if it works. Of course the fact that this trick invariably works for Conway doesn't mean it will *necessarily* work for you; mileage may vary. In other words, do not expect anyone to cover your losses if you mess this trick up.

3 Bicycle pedals

It is fairly commonly known, at least among bicyclists, that there is something funny about the way that the pedals of a bicycle screw into the cranks. One of the pedals has a normal 'right-hand thread', so that you screw it in clockwise—the usual way—like a normal screw or lightbulb, and you unscrew it counter-clockwise. The other pedal has a 'left-hand thread', so that it works exactly backwards: You screw it in counter-clockwise, and you unscrew it clockwise.

This 'asymmetry' between the two pedals—actually it's a surfeit of symmetry we have here, rather than a dearth—is not just some whimsical notion on the part of bike manufacturers. If the pedals both had normal threads, one of them would fall out before you got to the end of the block.

If you're assembling a new bike out of the box, the fact that one of the pedals screws in the wrong way may cause some momentary confusion, but you can easily figure out what to do by looking at the threads. The real confusion comes when for one reason or another you need to unscrew one of your pedals, and you can't remember whether this pedal is the normal one or the screwy one, and the pedal is on so tightly that a modest torque in either direction fails to budge it. You get set to give it a major twist, only which way do you turn it? You worry that if you to turn it the wrong way you'll get it on so tight that you'll never get it off.

If you try to figure out which pedal is the normal one using common sense, the chances are overwhelming that you will figure it out exactly wrong. If you remember this, then you're all set: Just figure it out by common sense, and then go for the opposite answer. Another good strategy is to remember that ‘right is right; left is wrong.’

Problems

1. What is the difference between a screw and a bolt?

2. Do all barber poles spiral the same way? What about candy canes? What other things spiral? Do they always spiral the same way?

3. Take two identical bolts or screws or lightbulbs or barber poles, and place them tip to tip. Describe how the two spirals meet.

4. Take a bolt or a screw or a lightbulb or a barberpole and hold it perpendicular to a mirror so that its tip appears to touch the tip of its mirror image. Describe how the two spirals meet.

5. When you hold something up to an ordinary mirror you can’t quite get it to appear to touch its mirror image. Why not? How close can you come? What if you use a different kind of mirror?

6. Why is a right-hand thread called a ‘right-hand thread’? What is the ‘right-hand rule’?

7. Which way do tornados and hurricanes rotate in the northern hemisphere? Why?

8. Which way does water spiral down the drain in the southern hemisphere, and how do you know?

9. Use common sense to figure out which pedal on a bike has the normal, right-hand thread. If you come up with the correct answer that ‘right is right; left is wrong’ then we offer you our humblest apologies.

10. Now see if you can figure out the correct explanation.

11. You can simulate what is going on here by curling your fingers loosely around the eraser end of a nice long pencil (a long thin stick works even better), so that there’s a little extra room for the pencil to roll around inside your grip. Get someone else to press down gently on the business end of the pencil, to simulate the weight of the rider’s foot on

the pedal, and see what happens when you rotate your arm like the crank of a bicycle.

12. The best thing is to make a wooden model. Drill a block through a block of wood to represent the hole in the crank that the pedal screws into, and use a dowel just a little smaller in diameter than the hole to represent the pedal.

13. If your pedal is on really, really tight, you might be tempted to use a 'cheater', which is a pipe that you slip over the end of a wrench to increase the effective length of the handle. If it takes a force of 150 pounds on the end of a 9-inch long adjustable wrench to loosen the pedal, how much force will be required on the end of an 18-inch long cheater?

14. Wrench manufacturers, pipe manufacturers, bicycle manufacturers, your insurance underwriter, your family doctor, and your geometry teacher all maintain that using a cheater is a bad idea. Do you understand why? Discuss some of the dangers to which you might expose yourself by using a cheater. If despite all this well-meaning advice you should go ahead and use a cheater, and some harm should come to you, who will be to blame?

4 Bicycle chains

Sometimes, when you come to put the rear wheel back on your bike after fixing a flat, or when you are fooling around trying to get the chain back onto the sprockets after it has slipped off (what? you don't always keep your gears adjusted perfectly?), you may find that the chain is in the peculiar kinked configuration shown in Figure 2.

Problems

1. Since you haven't removed a link of the chain or anything like that, you know it must be possible to get the chain unkinked, but how? Play around with a bike chain (a pair of rubber gloves is handy), and figure out how to introduce and remove kinks of this kind.

Figure 2: Kinked bicycle chain.

2. Draw a sequence of diagrams showing intermediate stages that you go through to get from the kinked to the kinked configuration.

3. Now close your eyes, and see if you can visualize the process of kinking and unkinking the chain. This can be tricky at first; practice it until you get good at it. Visualization is a skill that you can learn. The more you practice it, the better you get at it.

4. Keep a sharp eye out for opportunities to impress people with your ability to disentangle bicycle chains. Such opportunities are not common, but the satisfaction attendant on straightening out someone's chain is hard to describe.

5. Take a look at the bicycle chains shown in Figure 3. Some of these chain are not in configurations that the chain can get into from the normal configuration without removing a link. To disentangle these recalcitrant chain, you would need to remove one of the links using a tool called a 'chain-puller', mess around with the open-ended chain, and then do the link back up again. Can you tell which chains require a chain-puller?

6. Some of the chains in Figure 3 that require a chain-puller can be untangled without one if you know how to perform Chain Magic, which is a magical spell that will convert between an overcrossing and an undercrossing, as shown in Figure 4. Which?

7. Try to formulate a general rule that will tell you which chains can be untangled with Chain Magic, but without the aid of a chain-puller.

8. Now how about a rule to tell which chains can be untangled without Chain Magic?

9. The theory of straightening out bicycle chains using Chain Magic is called 'regular homotopy theory'. A higher-dimensional version of the theory explains how you can turn a sphere in three dimensional space 'inside out'. What this means and how it is done is explained in the video 'Inside Out', produced by the Minnesota Geometry Center. Keep your eye out for an opportunity to watch this amazing video.

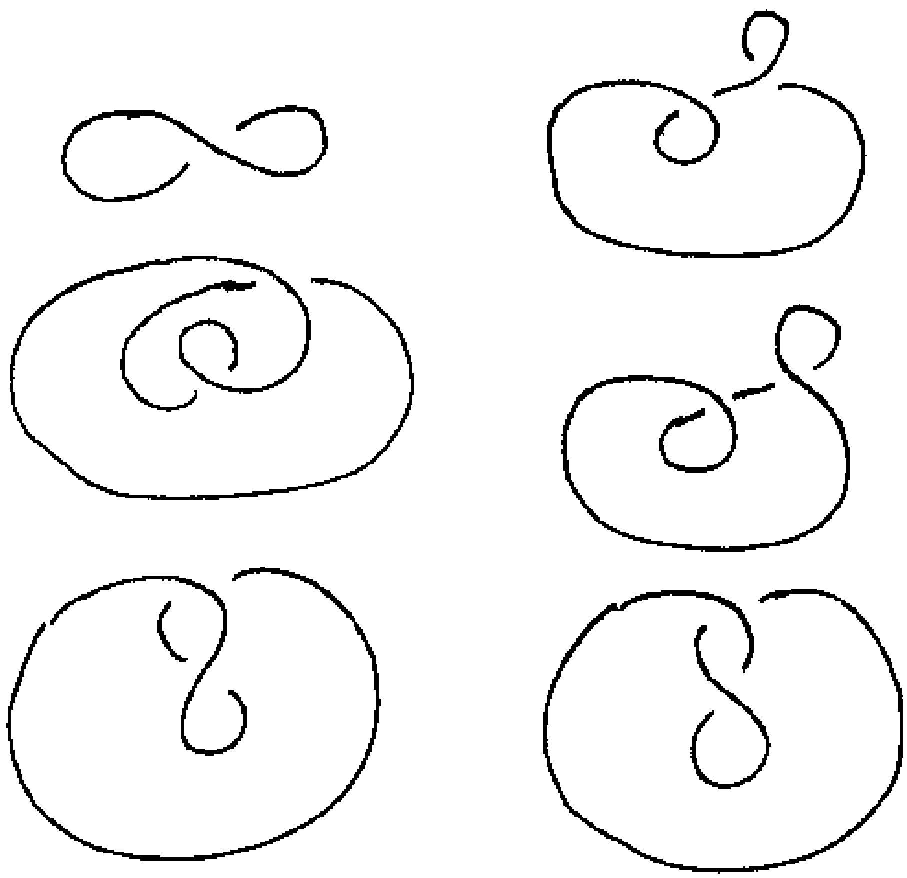

Figure 3: More kinked bicycle chains.

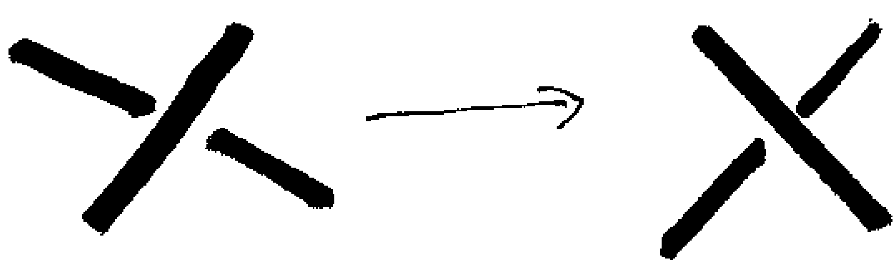

Figure 4: Chain Magic.

5 Push left, go left

Motorcycle riders have a saying:‘Push left, go left’.

Problems

1. What does this saying mean?

2. Would this saying apply to bicycles? tricycles?

6 Knots

A mathematical knot is a knotted loop. For example, you might take an extension cord from a drawer and plug one end into the other: this makes a mathematical knot.

It might or might not be possible to unknot it without unplugging the cord. A knot which can be unknotted is called an *unknot*.

Two knots are considered equivalent if it is possible to rearrange one to the form of the other, without cutting the loop and without allowing it to pass through itself. The reason for using loops of string in the mathematical definition is that knots in a length of string can always be undone, so any two lengths of string are equivalent in this sense.

If you drop a knotted loop of string on a table, it crosses over itself in a certain number of places. Possibly, there are ways to rearrange it with fewer crossings—the minimum possible number of crossings is the *crossing number* of the knot.

Make drawings and use short lengths of string to investigate the following problems.

Problems

1. Are there any knots with one or two crossings? Why?

2. How many inequivalent knots are there with three crossings?

3. How many knots are there with four crossings?

4. How many knots can you find with five crossings?

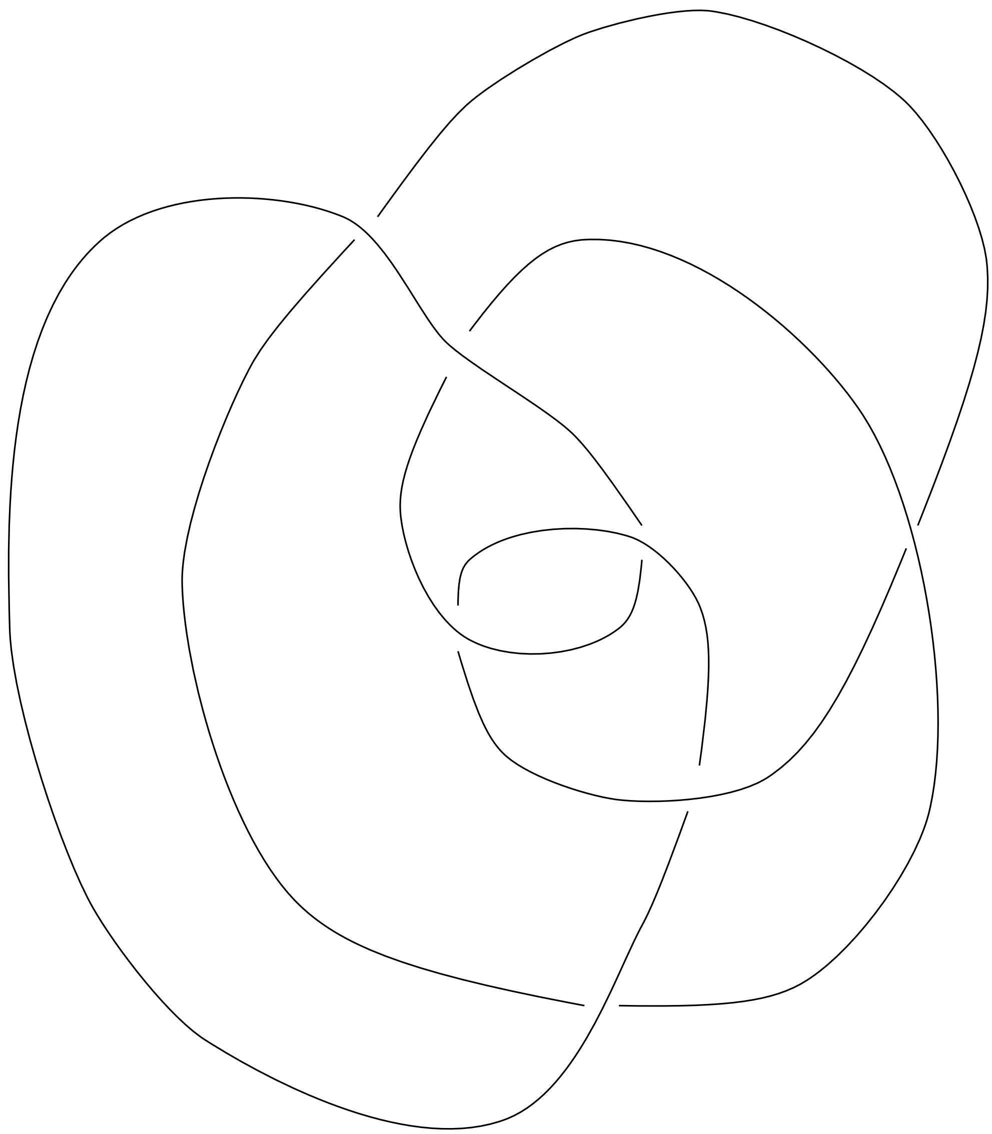

Figure 5: This is drawing of a knot with 7 crossings. Is it possible to rearrange it to have fewer crossings?

5. How many knots can you find with six crossings?

7 Reidemeister moves

Problems

1. Are the two trefoils the same? How do you know? How could you convince someone else of your answer?

2. Are the two figures-of-eight the same? How do you know? How could you convince someone else of your answer?

3. What is involved in showing that two knot diagrams represent the same knot, and in showing that two knot diagrams do not represent the same knot?

4. Make a Reidemeister movie to show that the figure-of-eight is amphicheiral.

8 Notation for some knots

It is a hard mathematical question to completely codify all possible knots. Given two knots, it is hard to tell whether they are the same. It is harder still to tell for sure that they are different.

Many simple knots can be arranged in a certain form, as illustrated below, which is described by a string of positive integers along with a sign.

9 Knots diagrams and maps

A knot diagram gives a map on the plane, where there are four edges coming together at each vertex. Actually, it is better to think of the diagram as a map on the sphere, with a polygon on the outside. It sometimes helps in recognizing when diagrams are topologically identical to label the regions with how many edges they have.

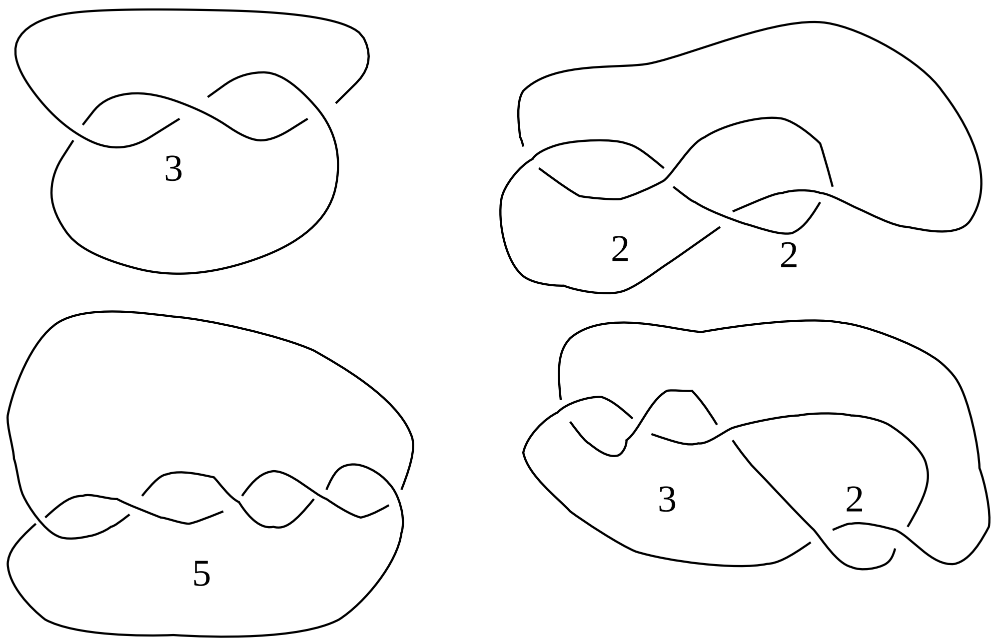

Figure 6: Here are drawings of some examples of knots that Conway 'names' by a string of positive integers. The drawings use the convention that when one strand crosses under another strand, it is broken. Notice that as you run along the knot, the strand alternates going over and under at its crossings. Knots with this property are called *alternating knots*. Can you find any examples of knots with more than one name of this type?

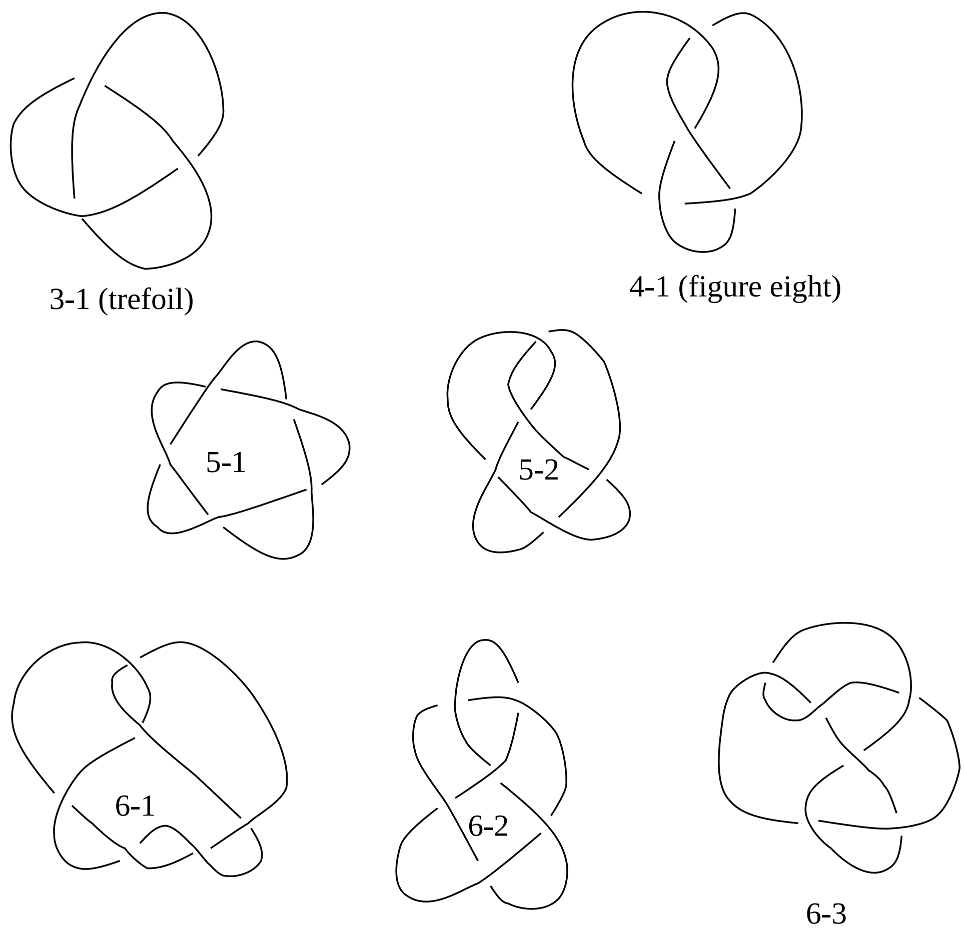

Figure 7: Here are the prime knots with up to six crossings. The names follow an old system, used widely in knot tables, where the kth knot with n crossings is called $n - k$. Mirror images are not included: some of these knots are equivalent to their mirror images, and some are not. Can you tell which are which?

10 Unicursal curves and knot diagrams

A unicursal curve in the plane is a curve that you get when you put down your pencil, and draw until you get back to the starting point. As you draw, your pencil mark can intersect itself, but you're not supposed to have any triple intersections. You could say that you pencil is allowed to pass over an point of the plane at most twice. This property of not having any triple intersections is *generic*: If you scribble the curve with your eyes closed (and somehow magically manage to make the curve finish off exactly where it began), the curve won't have any triple intersections.

A unicursal curve differs from the curves shown in knot diagrams in that there is no sense of the curve's crossing over or under itself at an intersection. You can convert a unicursal curve into a knot diagram by indicating (probably with the aid of an eraser), which strand crosses over and which strand crosses under at each of the intersections.

A unicursal curve with 5 intersections can be converted into a knot diagram in 2^5 ways, because each intersection can be converted into a crossing in two ways. These 32 diagrams will not represent 32 different knots, however.

Problems

1. Draw the 32 knot diagrams that arise from the unicursal curve underlying the diagram of knot 5-2, and identify the knots that these diagrams represent.

2. Show that any unicursal curve can be converted into a diagram of the unknot.

3. Show that any unicursal curve can be converted into the diagram of an alternating knot in precisely two ways. These two diagrams may or may not represent different knots. Give an example where the two knots are the same, and another where the two knots are different.

4. Show that any unicursal curve gives a map of the plane whose regions can be colored black and white in such a way that adjacent regions have different colors. In how many ways can this coloring be done? Give examples.

11 Exercises in imagining

How do you imagine geometric figures in your head? Most people talk about their three-dimensional imagination as 'visualization', but that isn't exactly right. The image you form in your head is more conceptual than a picture—you locate things in more of a three-dimensional model than in a picture. In fact, it is not easy to go from a mental image to a two-dimensional visual picture. Three-dimensional mental images are connected with your visual sense, but they are also connected with your sense of place and motion. In forming an image, it often helps to imagine moving around it, or tracing it out with your hands.

Geometric imagery is not just something that either you are born with or you are not. Like any other skill, it is something that needs to be developed with practice.

Below are some images to practice with. Some are two-dimensional, some are three-dimensional. Some are easy, some are hard, but not necessarily in numerical order. Work through these exercises in pairs. Evoke the images by talking about them, not by drawing them. It will probably help to close your eyes, although sometimes gestures and drawings in the air will help. Skip around to try to find exercises that are the right level for you.

Problems

1. Picture your first name, and read off the letters backwards. If you can't see your whole name at once, do it by groups of three letters. Try the same for your partner's name, and for a few other words. Make sure to do it by sight, not by sound.

2. Cut off each corner of a square, as far as the midpoints of the edges. What shape is left over? How can you re-assemble the four corners to make another square?

3. Mark the sides of an equilateral triangle into thirds. Cut off each corner of the triangle, as far as the marks. What do you get?

4. Take two squares. Place the second square centered over the first square but at a forty-five degree angle. What is the intersection of the two squares?

5. Mark the sides of a square into thirds, and cut off each of its corners back to the marks. What does it look like?
6. How many edges does a cube have?
7. Take a wire frame which forms the edges of a cube. Trace out a closed path which goes exactly once through each corner.
8. Take a 3×4 rectangular array of dots in the plane, and connect the dots vertically and horizontally. How many squares are enclosed?
9. Find a closed path along the edges of the diagram above which visits each vertex exactly once? Can you do it for a 3×3 array of dots?
10. How many different colors are required to color the faces of a cube so that no two adjacent faces have the same color?
11. A tetrahedron is a pyramid with a triangular base. How many faces does it have? How many edges? How many vertices?
12. Rest a tetrahedron on its base, and cut it halfway up. What shape is the smaller piece? What shapes are the faces of the larger pieces?
13. Rest a tetrahedron so that it is balanced on one edge, and slice it horizontally halfway between its lowest edge and its highest edge. What shape is the slice?
14. Cut off the corners of an equilateral triangle as far as the midpoints of its edges. What is left over?
15. Cut off the corners of a tetrahedron as far as the midpoints of the edges. What shape is left over?
16. You see the silhouette of a cube, viewed from the corner. What does it look like?
17. How many colors are required to color the faces of an octahedron so that faces which share an edge have different colors?
18. Imagine a wire is shaped to go up one inch, right one inch, back one inch, up one inch, right one inch, back one inch, ... What does it look like, viewed from different perspectives?

19. The game of tetris has pieces whose shapes are all the possible ways that four squares can be glued together along edges. Left-handed and right-handed forms are distinguished. What are the shapes, and how many are there?

20. Someone is designing a three-dimensional tetris, and wants to use all possible shapes formed by gluing four cubes together. What are the shapes, and how many are there?

21. An octahedron is the shape formed by gluing together equilateral triangles four to a vertex. Balance it on a corner, and slice it halfway up. What shape is the slice?

22. Rest an octahedron on a face, so that another face is on top. Slice it halfway up. What shape is the slice?

23. Take a $3 \times 3 \times 3$ array of dots in space, and connect them by edges up-and-down, left-and-right, and forward-and-back. Can you find a closed path which visits every dot but one exactly once? Every dot?

24. Do the same for a $4 \times 4 \times 4$ array of dots, finding a closed path that visits every dot exactly once.

25. What three-dimensional solid has circular profile viewed from above, a square profile viewed from the front, and a triangular profile viewed from the side? Do these three profiles determine the three-dimensional shape?

26. Find a path through edges of the dodecahedron which visits each vertex exactly once.

12 Pizza

Problems

1. How much more pizza does a 16-inch pie contain than a 14-inch pie?

2. How much more water does a 10-inch tall pitcher hold than an 8-inch tall pitcher?

3. How much more work does it take to build a 200-foot pyramid than a 100-foot pyramid?

4. What causes the phases of the moon?

5. Which way does water swirl down the drain in the southern hemisphere, and how do you know?

13 Ideas for projects

Here are some ideas for projects. Be creative—don't feel limited by these ideas.

- Make sets of tiles which exhibit various kinds of symmetry and which tile the plane in various symmetrical patterns.

- Write a computer program that replicates three-dimensional objects according to a three-dimensional pattern, as in the tetrahedron, octahedron, and icosahedron.

- The Archimidean solids are solids whose faces are regular polygons (but not necessarily all the same) such that every vertex is symmetric with every other vertex. Make models of the the Archimedean solids

- Write a computer program for visualizing four-dimensional space.

- Make stick models of the regular four-dimensional solids.

- Make models of three-dimensional cross-sections of regular four-dimensional solids.

- Design and implement three-dimensional tetris.

- Make models of the regular star polyhedra (Kepler-Poinsot polyhedron).

- Knit a Klein bottle, or a projective plane.

- Make some hyperbolic cloth.

- Sew topological surfaces and maps.
- Infinite Euclidean polyhedra.
- Hyperbolic polyhedra.
- Make a (possibly computational) orrery.
- Design and make a sundial.
- Astrolabe (Like a primitive sextant).
- Calendars: perpetual, lunar, eclipse.
- Cubic surface with 27 lines.
- Spherical Trigonometry or Geometry: Explore spherical trigonometry or geometry. What is the analog on the sphere of a circle in the plane? Does every spherical triangle have a unique inscribed and circumscribed circle? Answer these and other similar questions.
- Hyperbolic Trigonometry or Geometry: Explore hyperbolic trigonometry or geometry. What is the analog in the hyperbolic plane of a circle in the Euclidean plane? Does every hyperbolic triangle have a unique inscribed and circumscribed circle? Answer these and other similar questions.
- Make a convincing model showing how a torus can be filled with circular circles in four different ways.
- Turning the sphere inside out.
- Stereographic lamp.
- Flexible polyhedra.
- Models of ruled surfaces.
- Models of the projective plane.
- Puzzles and models illustrating extrinsic topology.

- Folding ellipsoids, hyperboloids, and other figures.
- Optical models: elliptical mirrors, *etc.*
- Mechanical devices for angle trisection, *etc.*
- Panoramic polyhedron (similar to an astronomical globe) made from faces which are photographs.
- Write a computer program for drawing tilings of the hyperbolic plane, using one or two of the possible hyperbolic symmetry groups.

14 Polyhedra

A *polyhedron* is the three-dimensional version of a polygon: it is a chunk of space with flat walls. In other words, it is a three-dimensional figure made by gluing polygons together. The word is Greek in origin, meaning many-seated. The plural is polyhedra. The polygonal sides of a polyhedron are called its *faces*.

Collect some equilateral triangles, either the snap-together plastic polydrons or paper triangles. Try gluing them together in various ways to form polyhedra.

Problems

1. Fasten three triangles together at a vertex. Complete the figure by adding one more triangle. Notice how there are three triangles at *every* vertex. This figure is called a *tetrahedron* because it has four faces (see the table of Greek number prefixes.)
2. Fasten triangles together so there are four at every vertex. How many faces does it have? From the table of prefixes below, deduce its name.
3. Do the same, with five at each vertex.
4. What happens when you fasten triangles six per vertex?
5. What happens when you fasten triangles seven per vertex?

1	mono
2	di
3	tri
4	tetra
5	penta
6	hexa
7	hepta
8	octa
9	ennia
10	deca
11	hendeca
12	dodeca
13	triskaideca
14	tetrakaideca
15	pentakaideca
16	hexakaideca
17	heptakaideca
18	octakaideca
19	enniakaideca
20	icosa

Table 1: The first 20 Greek number prefixes

A *regular polygon* is a polygon with all its edges equal and all angles equal. A *regular polyhedron* is whose faces are regular polygons, all congruent, and with the same number of polygons at each vertex.

Problem

- Construct models of all possible regular polyhedra, by trying what happens when you fasten together regular polygons with 3, 4, 5, 6, 7, *etc* sides so the same number come together at each vertex. Make a table listing the number of faces, vertices, and edges of each. What should they be called?

15 Maps

A *map* in the plane is a collection of vertices and edges (possibly curved) joining the vertices such that if you cut along the edges the plane falls apart into polygons. These polygons are called the faces. A map on the sphere or any other surface is defined similarly. Two maps are considered to be the same if you can get from one to the other by a continouous motion of the whole plane. Thus the two maps in figure 8 are considered to be the same.

A map on the sphere can be represented by a map in the plane by removing a point from the sphere and then stretching the rest of the sphere out to cover the plane. (Imagine popping a balloon and stretching the rubber out onto on the plane, making sure to stretch the material near the puncture all the way out to infinity.)

Depending on which point you remove from the sphere, you can get different maps in the plane. For instance, figure 9 shows three ways of representing the map depicting the edges and vertices of the cube in the plane; these three different pictures arise according to whether the point you remove lies in the middle of a face, lies on an edge, or coincides with one of the vertices of the cube.

16 Euler numbers

For the regular polyhedra, the *Euler number* $V - E + F$ takes on the value 2. The Euler number is also called the *Euler characteristic*, and it is commonly

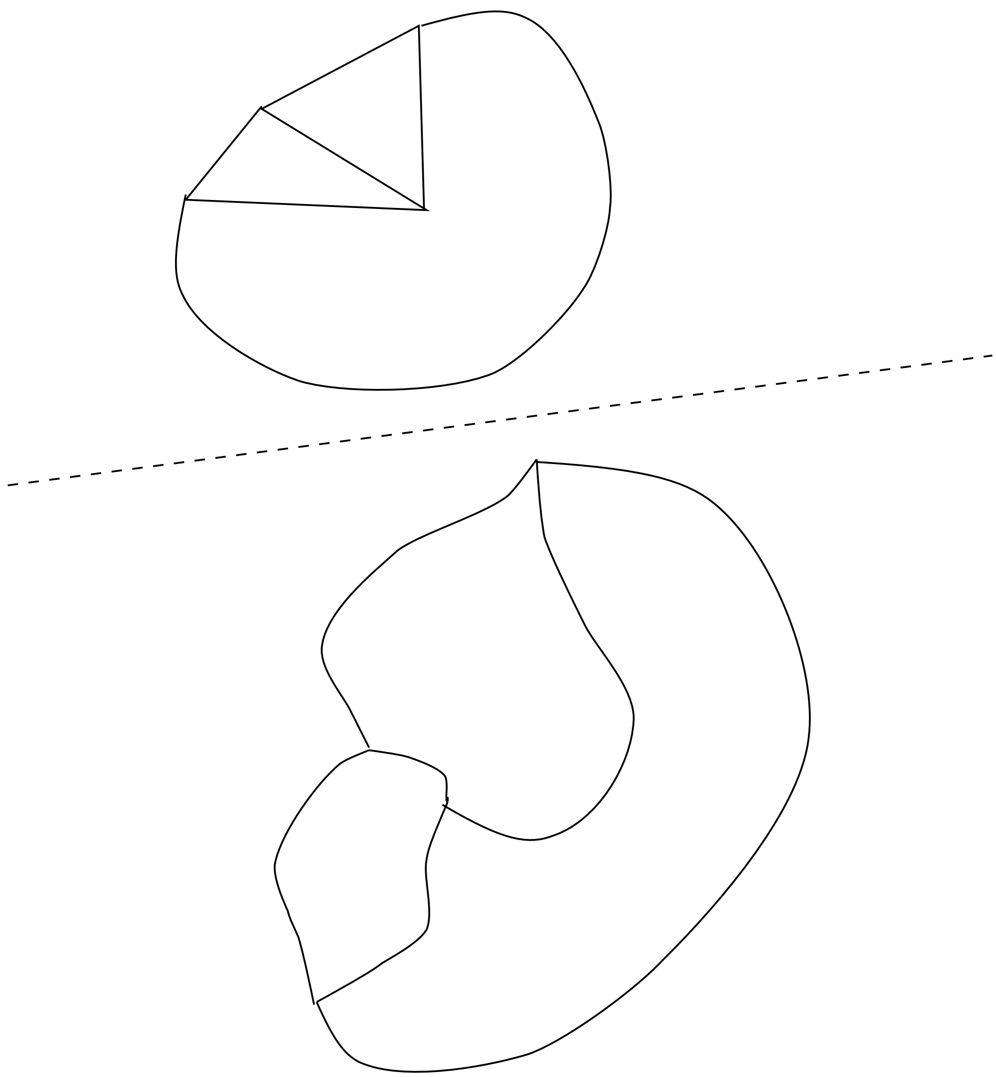

Figure 8: These two maps are considered the same (topologically equivalent), because it is possible to continuously move one to obtain the other.

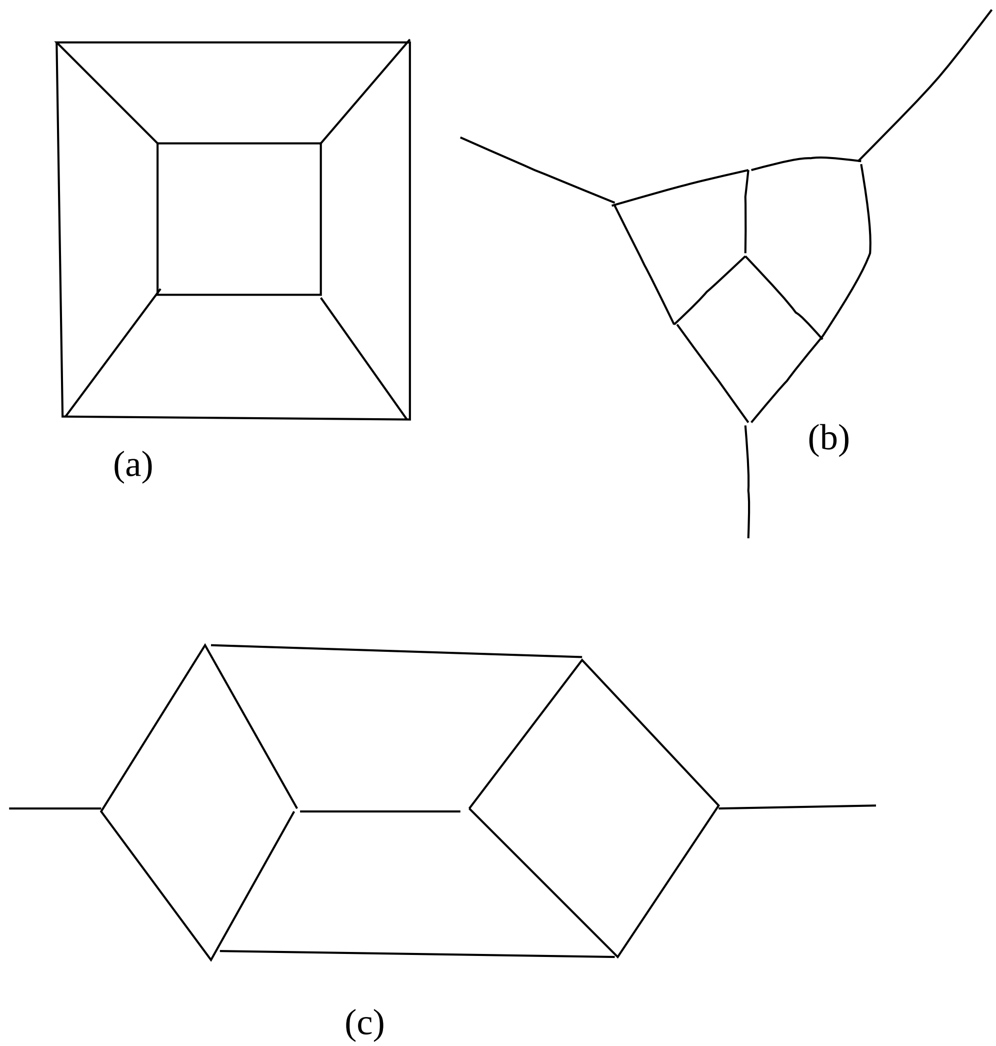

Figure 9: These three diagrams are maps of the cube, stretched out in the plane. In (a), a point has been removed from a face in order to stretch it out. In (b), a vertex has been removed. In (c), a point has been removed from an edge.

denoted by the Greek letter χ (pronounced 'kai', to rhyme with 'sky'):

$$\chi = V - E + F.$$

We propose to investigate the extent to which it is true that the Euler number of a polyhedron is always equal to 2. In the course of this investigation, you will gain some experience with representing polyhedra in the plane using maps, and with drawing dual maps.

Collect, or have someone else collect, a whole bunch of polyhedra, including among them some with 'holes' in them.

Problems

1. For as many of the polyhedra as you can, determine the values of V, E, F, and the Euler number χ.

2. When you are counting the vertices and so forth, see if you can think of more than one way to count them, so that you can check your answers. Can you make use of symmetry to simplify counting?

3. The number χ is frequently very small compared with V, E, and F, Can you think of ways to find the value of χ without having to compute V, E, and F, by 'cancelling out' vertices or faces with edges? This gives another way to check your work.

The *dual* of a map is a map you get by putting a vertex in the each face, connecting the neighboring faces by new edges which cross the old edges, and removing all the old vertices and edges.

Problem

- To the extent feasible, draw a maps in the plane of the polyhedra you've been investigating, draw (in a different color) the dual maps.

17 Gas, water, electricity

The diagram below shows three houses, each connected up to three utilities.

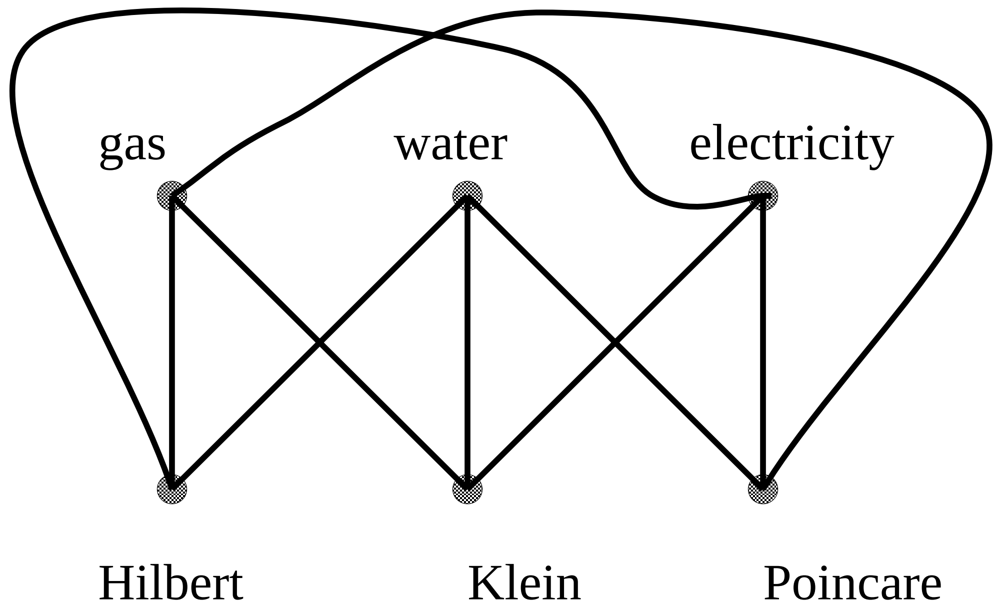

Figure 10: This is no good because we don't want the lines to intersect.

Problems

1. Show that it isn't possible to rearrange the connections so that they don't intersect each other.

2. Could you do it if the earth were a not a sphere but some other surface?

18 Topology

Topology is the theory of shapes which are allowed to stretch, compress, flex and bend, but without tearing or gluing. For example, a square is topologically equivalent to a circle, since a square can be continously deformed into a circle. As another example, a doughnut and a coffee cup with a handle for are topologically equivalent, since a doughnut can be reshaped into a coffee cup without tearing or gluing.

18.1 Letters

As a starting exercise in topology, let's look at the letters of the alphabet. We think of the letters as figures made from lines and curves, without fancy doodads such as serifs.

Problem

- Which of the capital letters are topologically the same, and which are topologically different? How many topologically different capital letters are there?

18.2 Surfaces

A *surface*, or 2-*manifold*, is a shape any small enough neighborhood of which is topologically equivalent to a neighborhood of a point in the plane. For instance, a the surface of a cube is a surface topologically equivalent to the surface of a sphere. On the other hand, if we put an extra wall inside a cube dividing it into two rooms, we no longer have a surface, because there are points at which three sheets come together. No small neighborhood of those points is topologically equivalent to a small neighborhood in the plane.

Here are some pictures of surfaces. The pictures are intended to indicate things like doughnuts and pretzels rather than flat strips of paper.

Problem

- Can you identify these surfaces, topologically? Which ones are topologically the same intrinsically, and which extrinsically?

19 Surfaces

Recall that you get a torus by identifying the sides of a rectangle as in Figure 2.10 of *SS* (The Shape of Space). If you identify the sides slightly differently, as in Figure 4.3, you get a surface called a *Klein bottle*, shown in Figure 4.9.

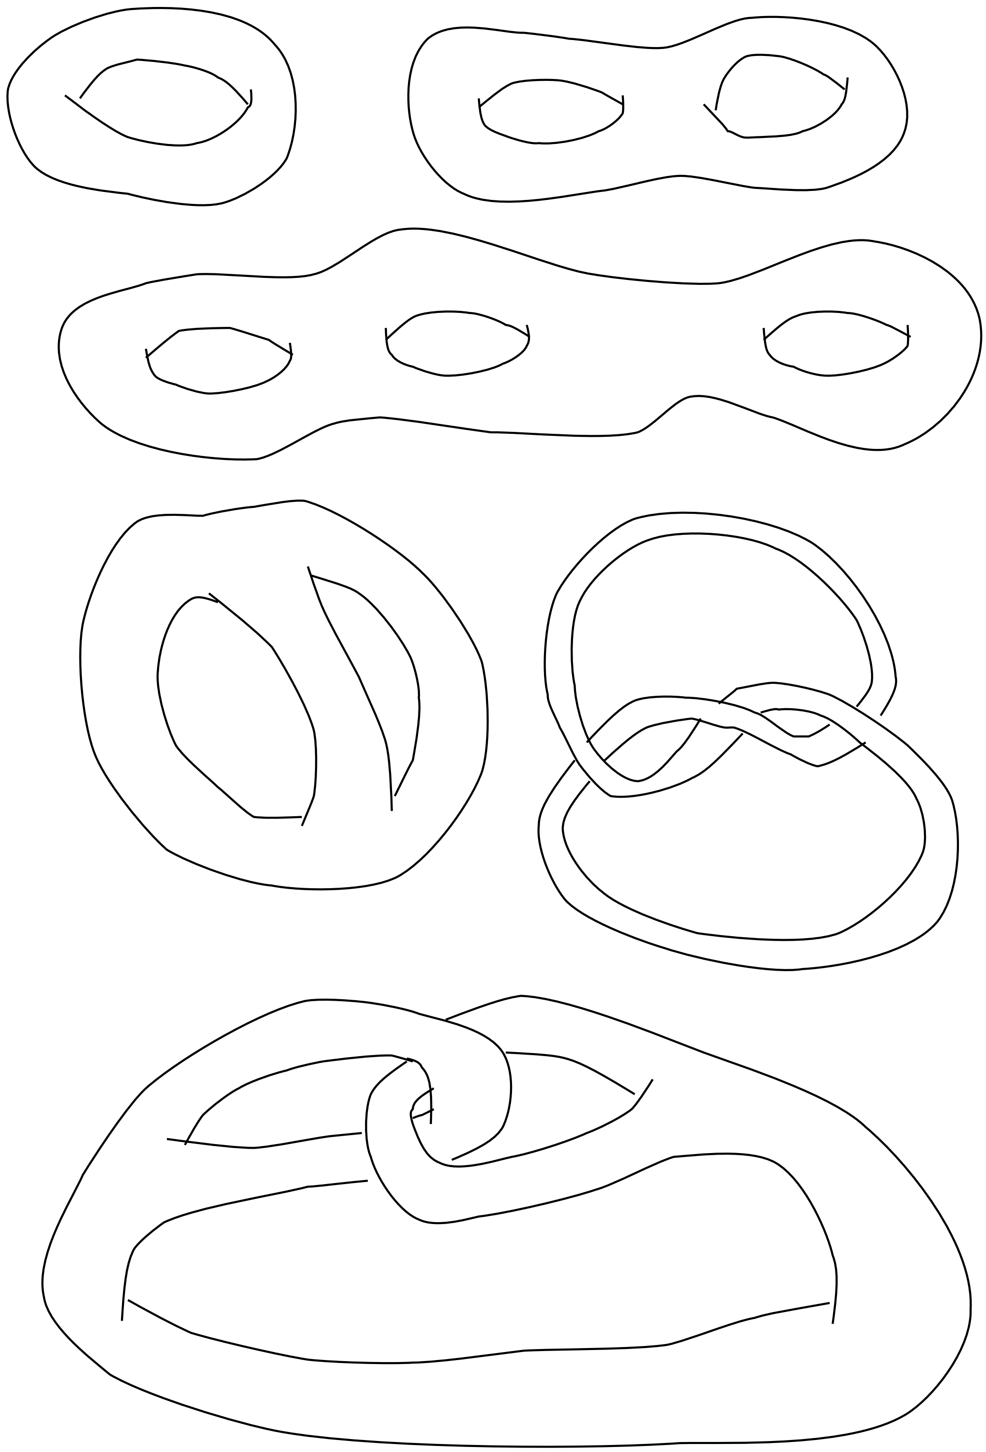

Figure 11: Some surfaces

Problems

1. Take some strips and join the opposite ends of each strip together as follows:

   with no twists;

   with one twist (half-turn); this is called a *Möbius strip*;

   with two twists;

   with three twists.

2. Imagine that you are a two-dimensional being who lives in one of these four surfaces. To what extent can you tell exactly which one it is?

3. Now cut each of the above along the midline of the original strip. Describe what you get. Can you explain why?

4. What is the Euler number of a disk? A Möbius strip? A torus with a circular hole cut from it? A Klein bottle? A Klein bottle with a circular hole cut from it?

5. What is the maximum number of points in the plane such that you can draw non-intersecting segments joining each pair of points? What about on a sphere? On a torus?

20 How to knit a Möbius Band

Start with a different color from the one you want to make the band in. Call this the spare color. With the spare color and normal knitting needles cast on 90 stitches.

Change to your main color yarn. Knit your row of 90 stitches onto a circular needle. Your work now lies on about 2/3 of the needle. One end of the work is near the tip of the needle and has the yarn attached. This is the working end. Bend the working end around to the other end of your work, and begin to knit those stitches onto the working end, but do not slip them off the other end of the needle as you normally would. When you have knitted all 90 stitches in this way, the needle loops the work twice.

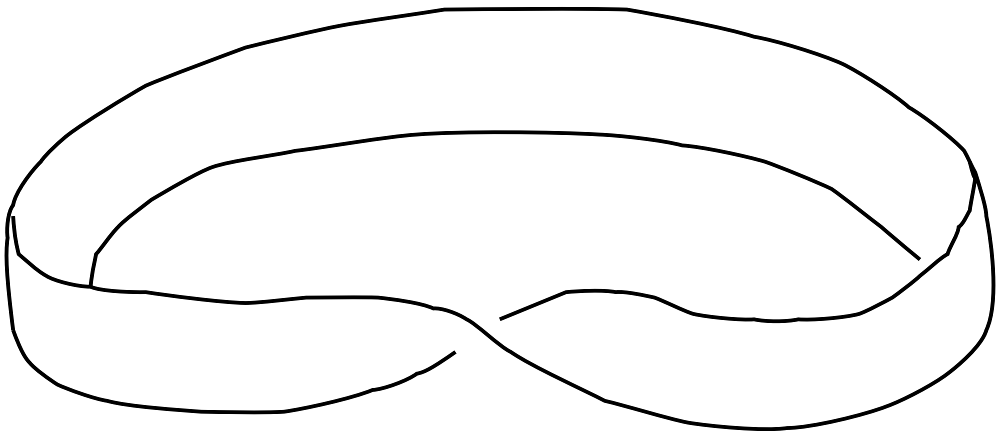

Figure 12: A Mobius band.

Carry on knitting in the same direction, slipping stitches off the needle when you knit them, as normal. The needle will remain looped around the work twice. Knit five 'rows' (that is 5×90 stitches) in this way.

Cast off. You now have a Mobius band with a row of your spare color running around the middle. Cut out and remove the spare colored yarn. You will be left with one loose stitch in your main color which needs to be secured.

(Expanded by Maria Iano-Fletcher from an original recipe by Miles Reid.)

21 Classification of surfaces

You can identify the topological type of a surface either by cutting and pasting, or by computing its invariants: Euler characteristic; orientability; number of boundary components. Use both of these methods in addressing the following problems.

Problems

1. What do you get when you cut a hole in a projective plane P^2?
2. Show that gluing two Mobius strips together along their boundary gives a Klein bottle. Can you see the two Mobius strips in the Klein bottle?

3. What do you get gluing opposite sides of a regular hexagon via translation? What about an octagon? a decagon?

4. Show that the connected sum of two projective planes is a Klein bottle.

5. Cut the globe along the equator and join the southern hemisphere to the northern by strips with a half twist. Is the result orientable? What is its boundary? What is its topological type?

6. Consider the great dodecahedron with self-intersections removed. Is it orientable? What is its topological type?

22 Mirrors

Problems

1. How do you hold two mirrors so as to get an integral number of images of yourself? Discuss the handedness of the images.

2. Set up two mirrors so as to make perfect kaleidoscopic patterns. How can you use them to make a snowflake?

3. Fold and cut hearts out of paper. Then make paper dolls. Then honest snowflakes.

4. Set up three or more mirrors so as to make perfect kaleidoscopic patterns. Fold and cut such patterns out of paper.

5. Why does a mirror reverse right and left rather than up and down?

23 More paper-cutting patterns

Experiment with the constructions below. Put the best examples into your journal, along with comments that describe and explain what is going on. Be careful to make your examples large enough to illustrate clearly the symmetries that are present. Also make sure that your cuts are interesting enough so that extra symmetries do not creep in. Concentrate on creating a collection of examples that will get across clearly what is going on, and include enough written commentary to make a connected narrative.

Problems

1. **Conical patterns.** Many rotationally-symmetric designs, like the twin blades of a food processor, cannot be made by folding and cutting. However, they can be formed by wrapping paper into a conical shape.

   Fold a sheet of paper in half, and then unfold. Cut along the fold to the center of the paper. Now wrap the paper into a conical shape, so that the cut edge lines up with the uncut half of the fold. Continue wrapping, so that the two cut edges line up and the original sheet of paper wraps two full turns around a cone. Now cut out any pattern you like from the cone. Unwrap and lay it out flat. The resulting pattern should have two-fold rotational symmetry.

   Try other examples of this technique, and also try experimenting with rolling the paper more than twice around a cone.

2. **Cylindrical patterns.** Similarly, it is possible to make repeating designs on strips. If you roll a strip of paper into a cylindrical shape, cut it, and unroll it, you should get a repeating pattern on the edge. Try it.

3. **Möbius patterns.** A Möbius band is formed by taking a strip of paper, and joining one end to the other with a twist so that the left edge of the strip continues to the right.

   Make or round up a strip of paper which is long compared to its width (perhaps made from ribbon, computer paper, adding-machine rolls, or formed by joining several shorter strips together end-to-end). Coil it around several times around in a Möbius band pattern. Cut out a pattern along the edge of the Möbius band, and unroll.

4. **Other patterns.** Can you come up with any other creative ideas for forming symmetrical patterns?

24 Symmetry and orbifolds

Given a symmetric pattern, what happens when you identify equivalent points? It gives an object with interesting topological and geometrical properties, called an orbifold.

The first instance of this is an object with bilateral symmetry, such as a (stylized) heart. Children learn to cut out a heart by folding a sheet of paper in half, and cutting out half of the pattern. When you identify equivalent points, you get half a heart.

A second instance is the paper doll pattern. Here, there are two different fold lines. You make paper dolls by folding a strip of paper zig-zag, and then cutting out half a person. The half-person is enough to reconstruct the whole pattern. The quotient orbifold is a half-person, with two mirror lines.

A wave pattern is the next example. This pattern repeats horizontally, with no reflections or rotations. The wave pattern can be rolled up into a cylinder. It can be constructed by rolling up a strip of paper around a cylinder, and cutting a single wave, through several layers, with a sharp knife. When it is unrolled, the bottom part will be like the waves.

When a pattern repeats both horizontally and vertically, but without reflections or rotations, the quotient orbifold is a torus. You can think of it by first rolling up the pattern in one direction, matching up equivalent points, to get a long cylinder. The cylinder has a pattern which still repeats vertically. Now coil the cylinder in the other direction to match up equivalent points on the cylinder. This gives a torus.

25 Names for features of symmetrical patterns

We begin by introducing names for certain features that may occur in symmetrical patterns. To each such feature of the pattern, there is a corresponding feature of the quotient orbifold, which we will discuss later.

25.1 Mirrors and mirror strings

A *mirror* is a line about which the pattern has mirror symmetry. Mirrors are perhaps the easiest features to pick out by eye.

At a *crossing point*, where two or more mirrors cross, the pattern will necessarily also have rotational symmetry. An n-way crossing point is one where precisely n mirrors meet. At an n-way crossing point, adjacent mirrors meet at an angle of π/n. (Beware: at a 2-way crossing point, where two mirrors meet at right angles, there will be 4 slices of pie coming together.)

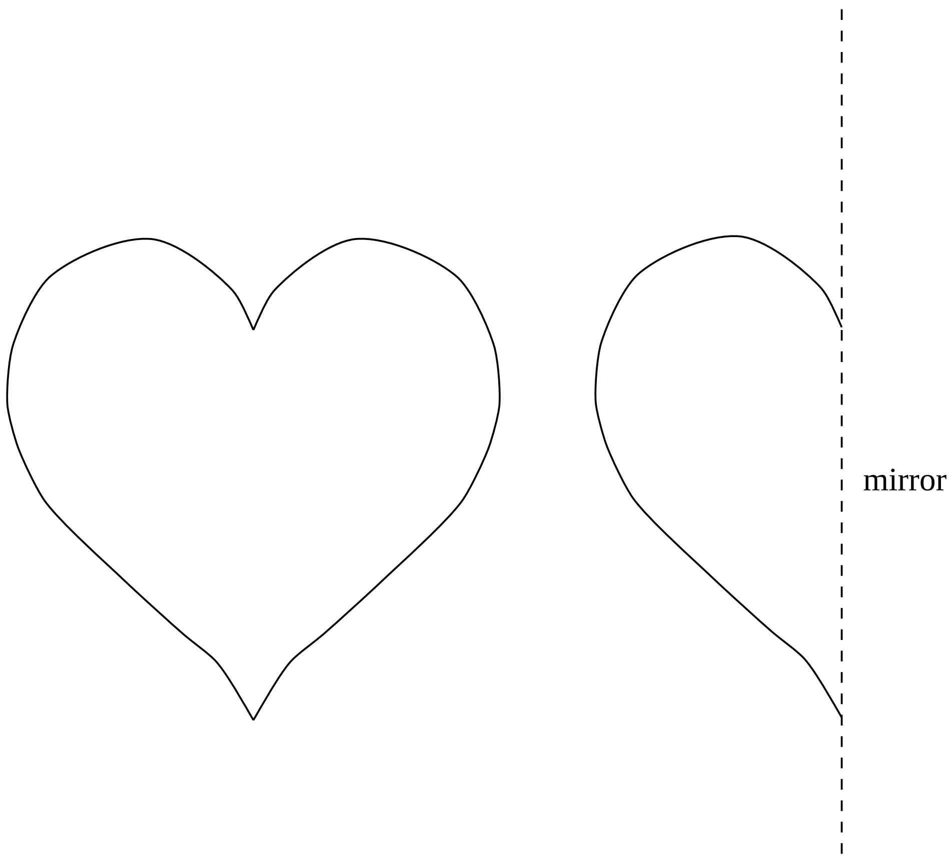

Figure 13: A heart is obtained by folding a sheet of paper in half, and cutting out half a heart. The half-heart is the orbifold for the pattern. A heart can also be recreated from a half-heart by holding it up to a mirror.

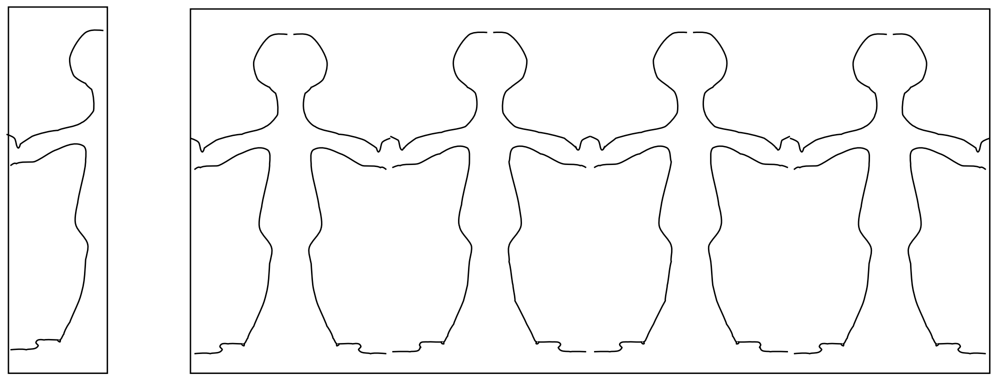

Figure 14: A string of paper dolls

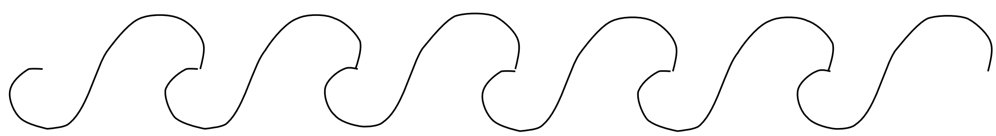

Figure 15: This wave pattern repeats horizontally, with no reflections or rotations. The quotient orbifold is a cylinder.

We obtain a *mirror string* by starting somewhere on a mirror and walking along the mirror to the next crossing point, turning as far right as we can so as to walk along another mirror, walking to the next crossing point on it, and so on. (See figure 22.)

Suppose that you walk along a mirror string until you first reach a point exactly like the one you started from. If the crossings you turned at were (say) a 6-way, then a 3-way, and then a 2-way crossing, then the mirror string would be of *type* $*632$, etc. As a special case, the notation $*$ denotes a mirror that meets no others.

For example, look at a standard brick wall. There are horizontal mirrors that each bisect a whole row of bricks, and vertical mirrors that pass through bricks and cement alternately. The crossing points, all 2-way, are of two kinds: one at the center of a brick, one between bricks. The mirror strings have four corners, and you might expect that their type would be $*2222$. However, the correct type is $*22$. The reason is that after going only half way round, we come to a point exactly like our starting point.

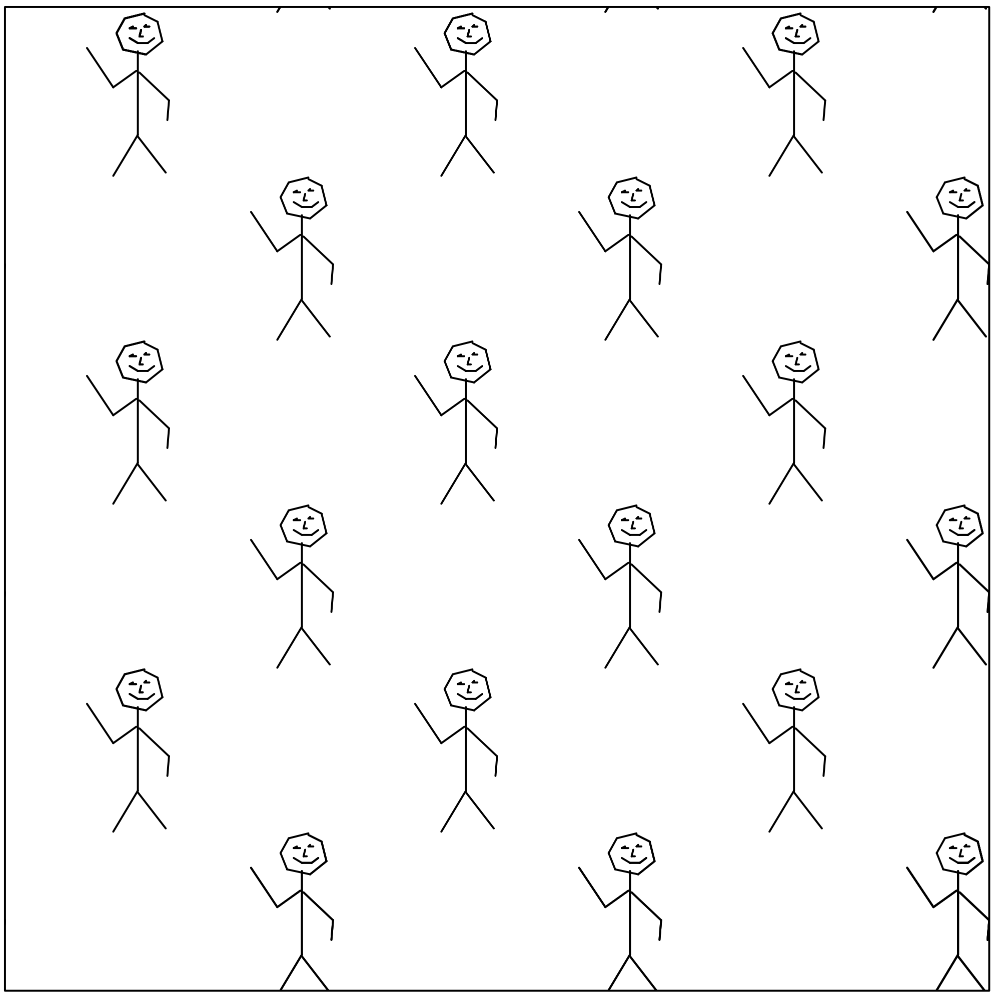

Figure 16: This pattern has quotient orbifold a torus. It repeats both horizontally and vertically, but without any reflections or rotations. It can be rolled up horizontally to form a cylinder, and then vertically (with a twist) to form a torus.

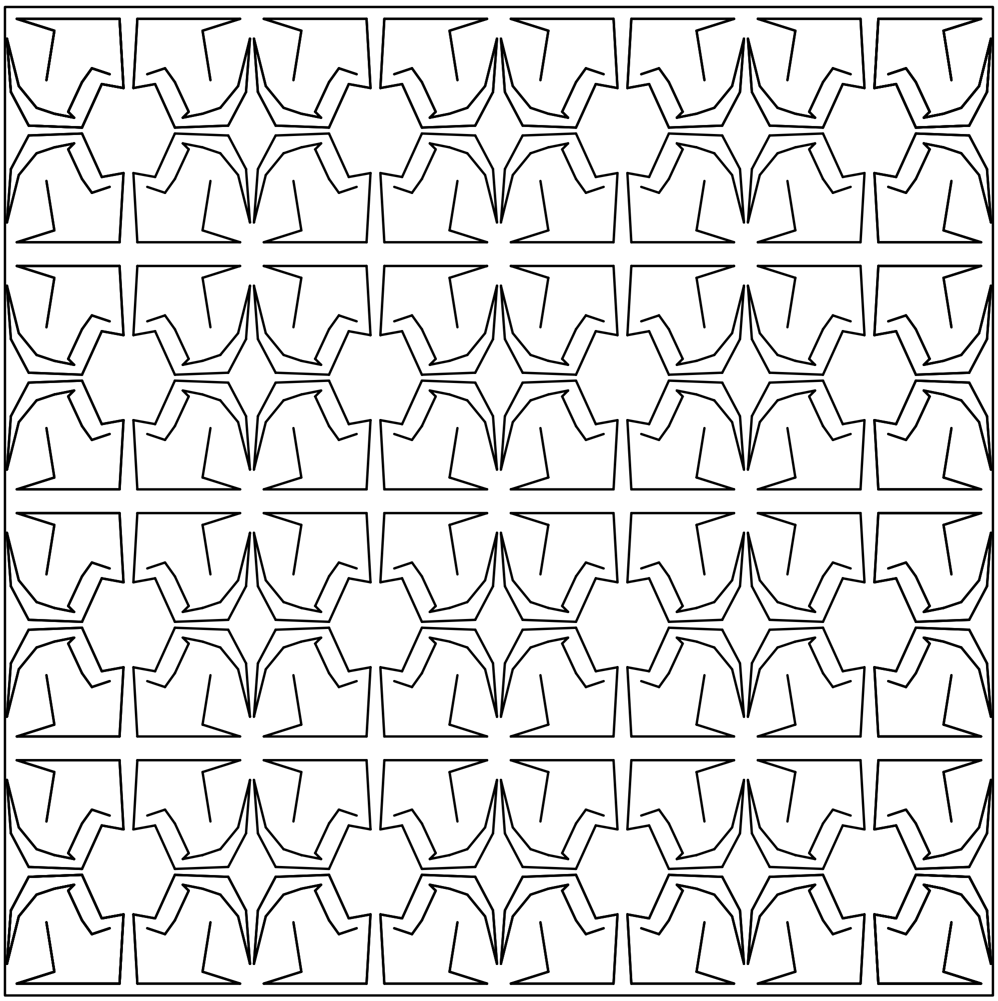

Figure 17: The quotient orbifold is a rectangle, with four mirrors around it.

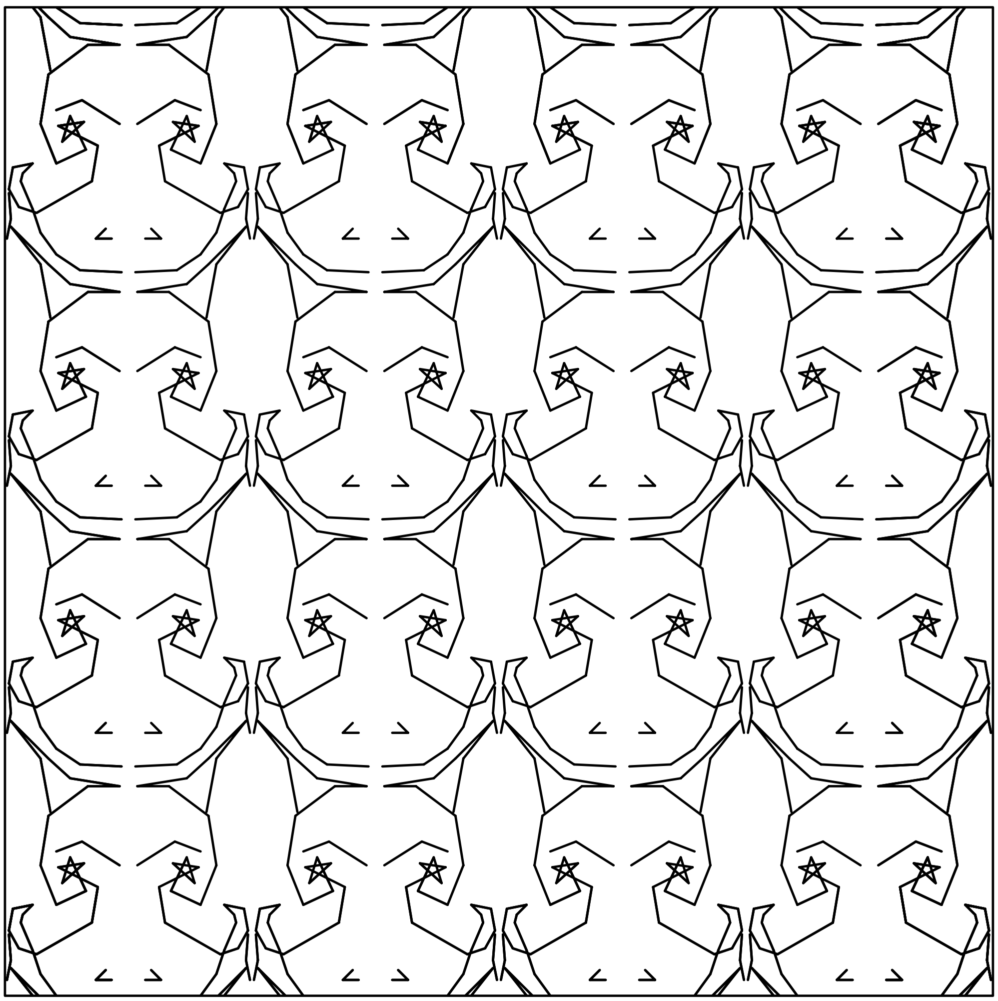

Figure 18: The quotient orbifold is an annulus, with two mirrors, one on each boundary.

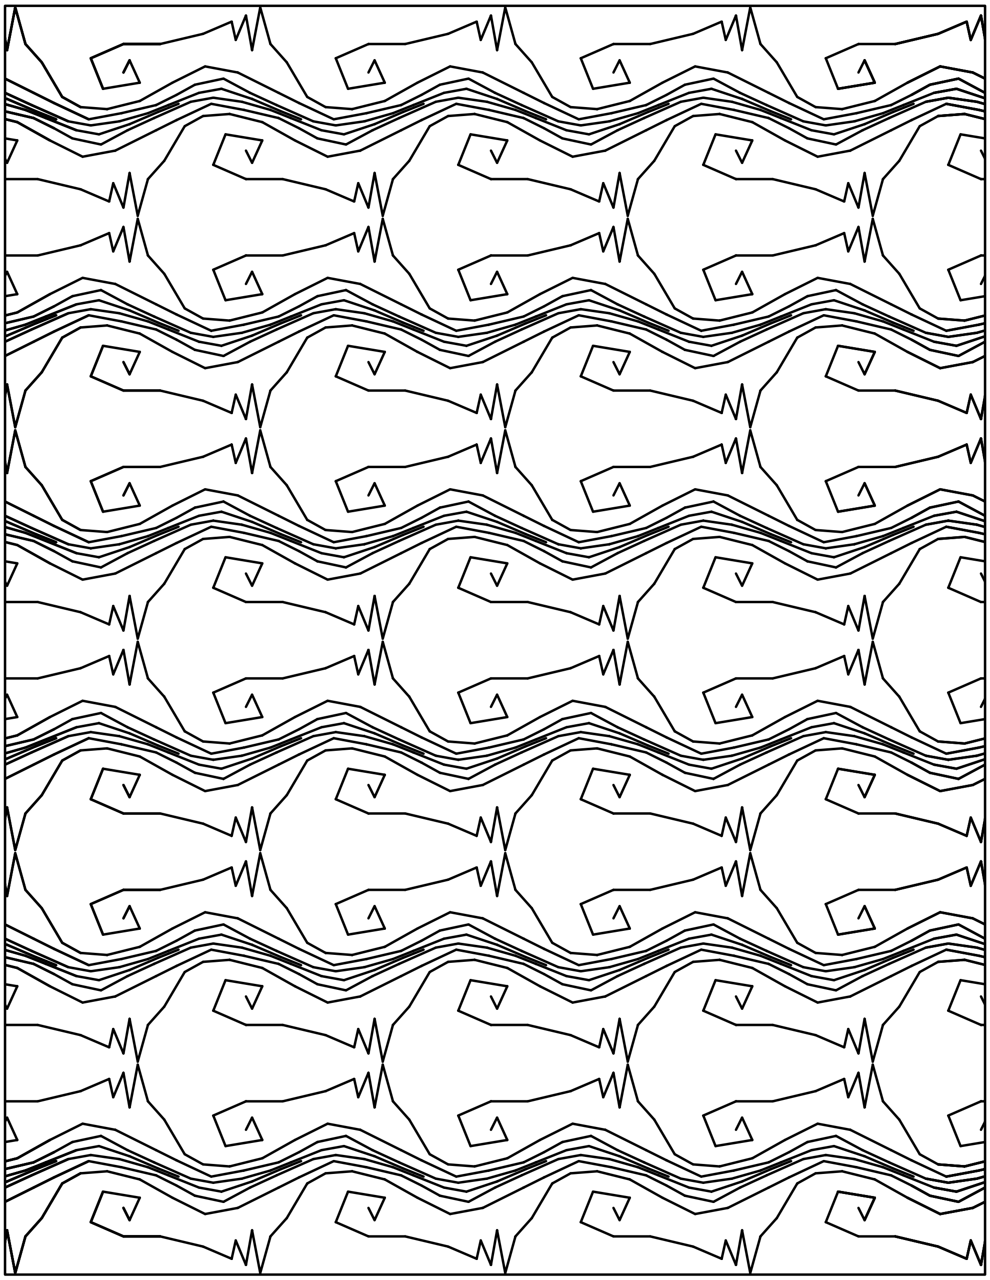

Figure 19: The quotient orbifold is a Moebius band, with a single mirror on its single boundary.

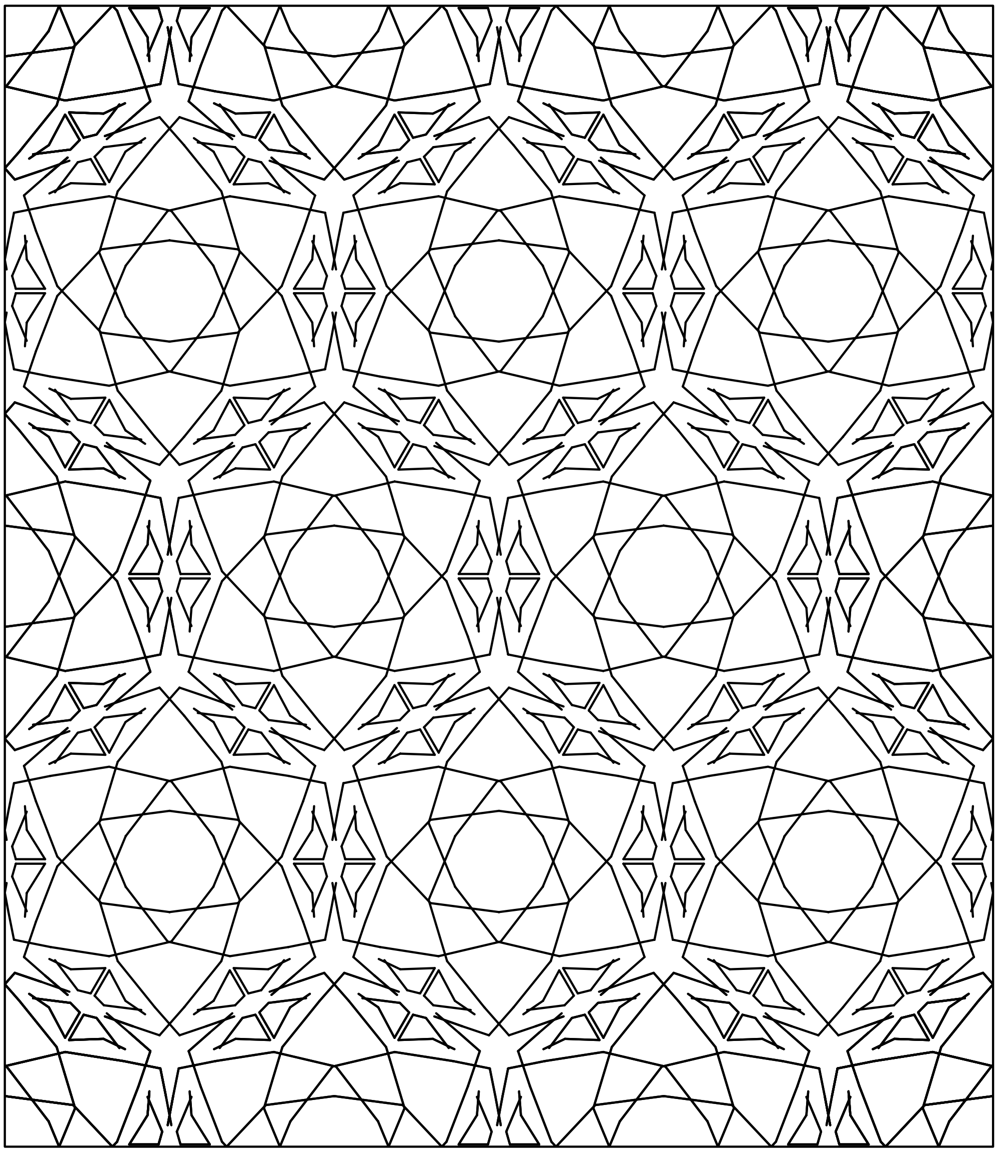

Figure 20: The quotient orbifold is a 60°, 30°, 90° triangle, with three mirrors from sides.

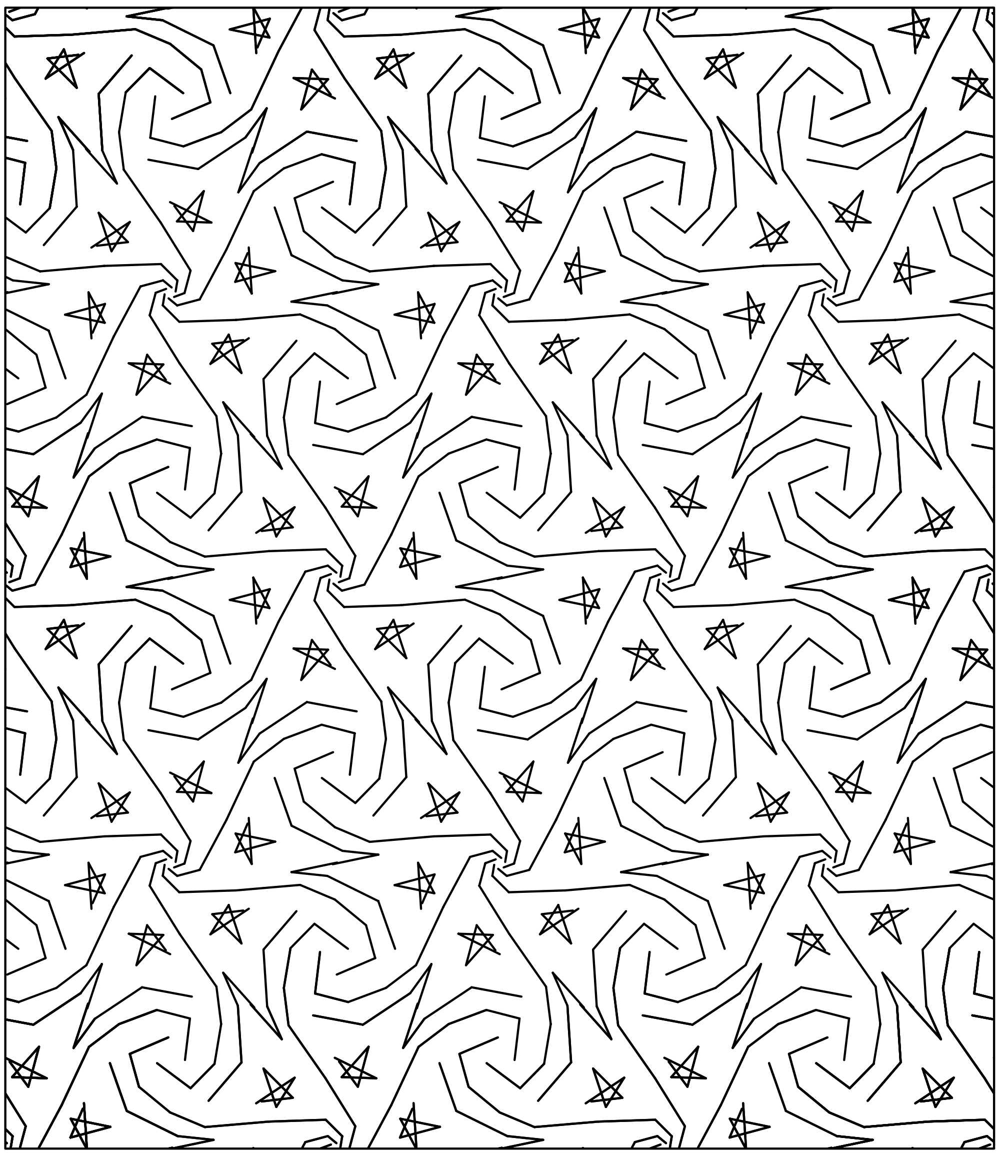

Figure 21: This pattern has rotational symmetry about various points, but no reflections. The rotations are of order 6, 3 and 2. The quotient orbifold is a triangular pillow, with three cone points.

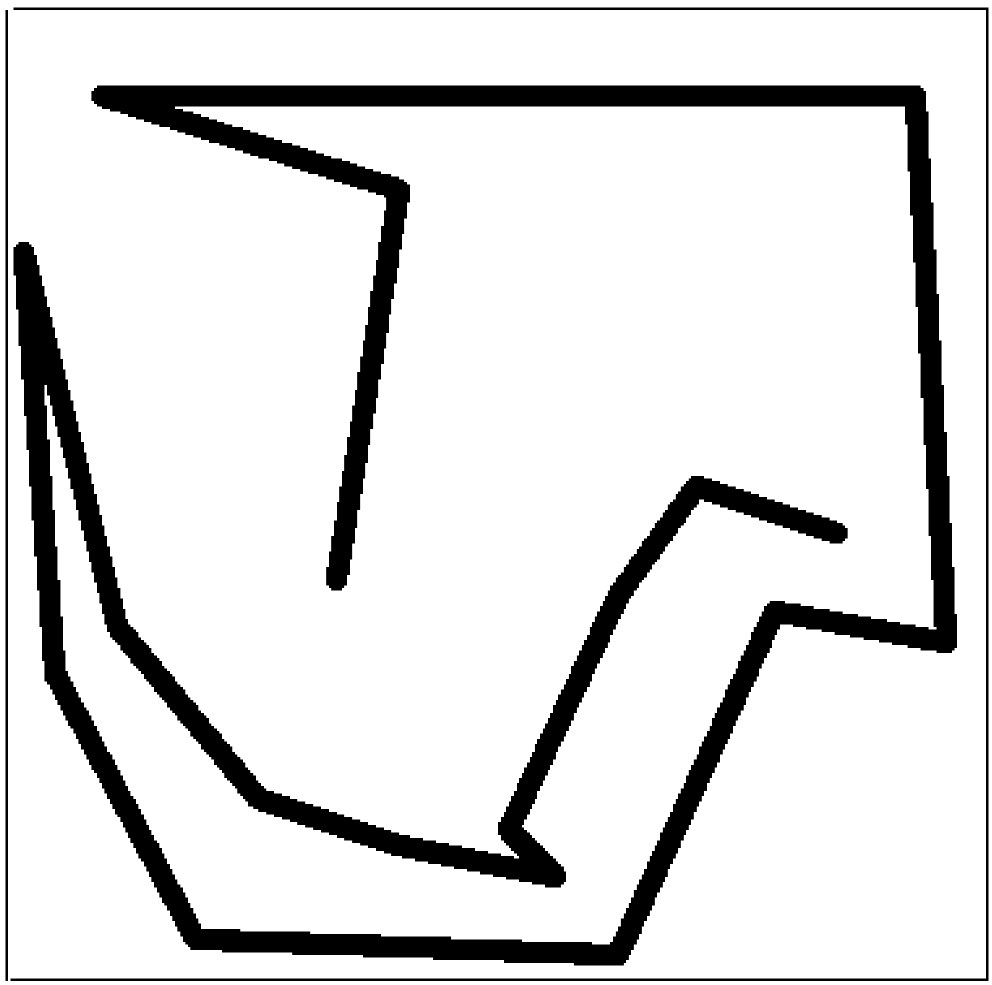

Figure 22: The quotient billiard orbifold.

25.2 Mirror boundaries

In the quotient orbifold, a mirror string of type $*abc$ becomes a boundary wall, along which there are corners of angles $\pi/a, \pi/b, \pi/c$. We call this a *mirror boundary* of type $*abc$. For example, a mirror boundary with no corners at all has type $*$. The quotient orbifold of a brick wall has a mirror boundary with just two right-angled corners, type $*22$.

25.3 Gyration points

Any point around which a pattern has rotational symmetry is called a *rotation point*. Crossing points are rotation points, but there may also be others. A rotation point that does NOT lie on a mirror is called a *gyration point*. A gyration point has *order* n if the smallest angle of any rotation about it is $2\pi/n$.

For example, on our brick wall there is an order 2 gyration point in the middle of the rectangle outlined by any mirror string.

25.4 Cone points

In the quotient orbifold, a gyration point of order n becomes a cone point with cone angle $2\pi/n$.

26 Names for symmetry groups and orbifolds

A symmetry group is the collection of all symmetry operations of a pattern. We give the same names to symmetry groups as to the corresponding quotient orbifolds.

We regard every orbifold as obtained from a sphere by adding cone-points, mirror boundaries, handles, and cross-caps. The major part of the notation enumerates the orders of the distinct cone points, and then the types of all the different mirror boundaries. An initial black spot $\bullet$ indicates the addition of a handle; a final circle $\circ$ the addition of a cross cap.

For example, our brick wall gives $2*22$, corresponding to its gyration point of order 2, and its mirror string with two 2-way corners.

Here are the types of some of the patterns shown in section 31:

Figure 16: $\bullet$; Figure 17: $*2222$; Figure 18: $**$; Figure 19: $*\circ$. Figure 20: $*632$. Figure 21: 632.

Apart from the spots and circles, these can be read directly from the pictures. The important thing to remember is that if two things are equivalent by a symmetry, then you only record one of them. A dodecahedron is very like a sphere. The orbifold corresponding to its symmetry group is a spherical triangle having angles $\pi/5, \pi/3, \pi/2$; so its symmetry group is $*532$.

You, the topologically spherical reader, approximately have symmetry group $*$, because the quotient orbifold of a sphere by a single reflection is a hemisphere whose mirror boundary has no corners.

27 The orbifold shop

The Orbifold Shop has gone into the business of installing orbifold parts. They offer a special promotional deal: a free coupon for \$2.00 worth of parts, installation included, to anyone acquiring a new orbifold.

There are only a few kinds of features for two-dimensional orbifolds, but they can be used in interesting combinations.

- Handle: \$2.00.
- Mirror: \$1.00.
- Cross-cap: \$1.00.
- Order n cone point: $\$1.00 \times (n-1)/n$.
- Order n corner reflector: $.50 \times (n-1)/n$. Prerequisite: at least one mirror. Must specify in mirror and position in mirror to be installed.

With the \$2.00 coupon, for example, you could order an orbifold with four order 2 cone points, costing \$.50 each. Or, you could order an order 3 cone point costing \$.66..., a mirror costing \$1.00, and an order 3 corner reflector costing \$.33....

Theorem. If you exactly spend your coupon at the Orbifold Shop, you will have a quotient orbifold coming from a symmetrically repeating pattern in the Euclidean plane with a bounded fundamental domain. There are exactly 17 different ways to do this, and corresponding to the 17 different symmetrically repeating patterns with bounded fundamental domain in the Euclidean plane.

Problem

- What combinations of parts can you find that cost exactly \$2.00?

28 The Euler characteristic of an orbifold

Suppose we have a symmetric pattern in the plane. We can make a symmetric map by subdividing the quotient orbifold into polygons, and then 'unrolling it' or 'unfolding it' to get a map in the plane.

If we look at a large area A in the plane, made up from N copies of a fundamental domain, then each face in the map on the quotient orbifold contributes N faces to the region. An edge which is not on a mirror also contributes approximately N copies — approximately, because when it is on the boundary of A, we don't quite know how to match it with a fundametnal region.

In general, if an edge or point has order k symmetry which which preserves it, it contributes approximately N/k copies of itself to A, since each time it occurs, as long as it is not on the boundary of A, it is counted in k copies of the fundamental domain.

Thus,

- If an edge is on a mirror, it contributes only approximately $N/2$ copies.
- If a vertex is not on a mirror and not on a cone point, it contributes approximately N vertices to A.
- If a vertex is on a cone point of order m it contributes approximately N/m vertices.
- If a vertex is on a mirror but not on a corner reflector, it contributes approximately $N/2$.
- If a vertex is on an order m corner reflector, it contributes approximately $N/2m$

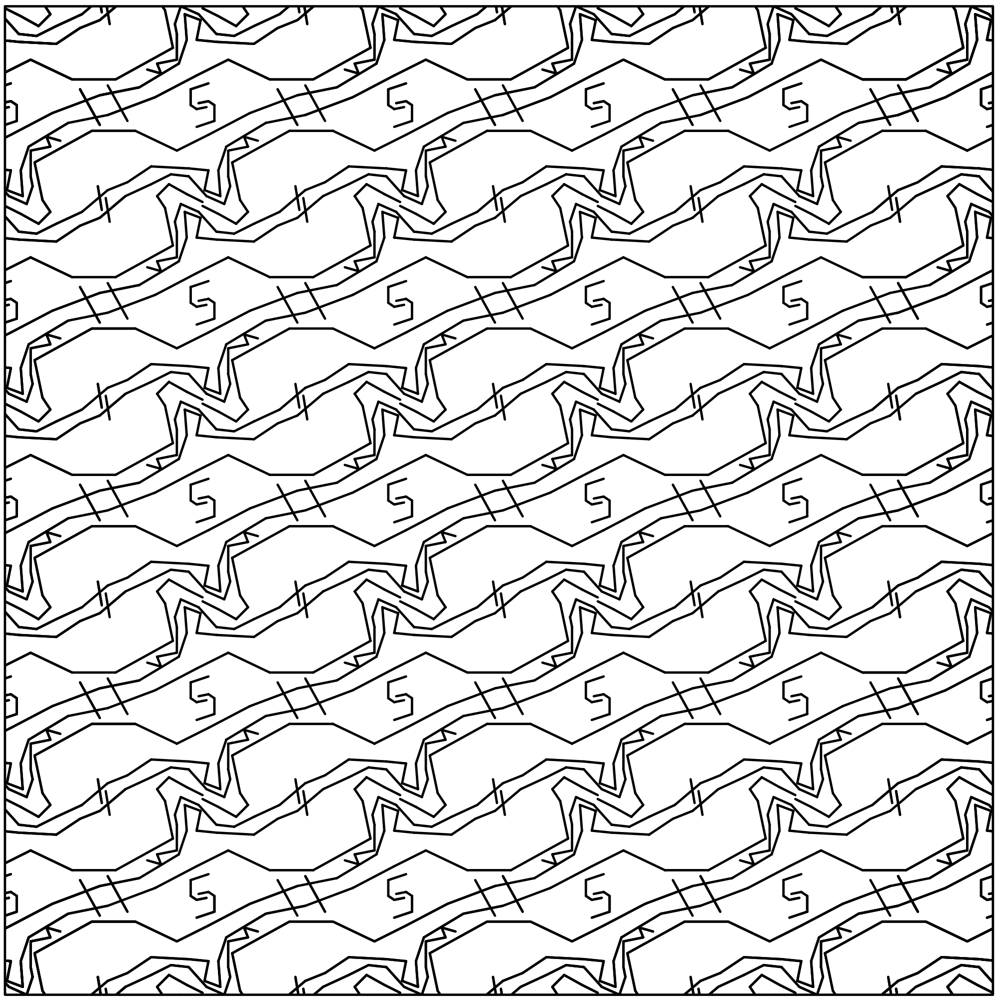

Figure 23: This is the pattern obtained when you buy four order 2 cone points for $.50 each.

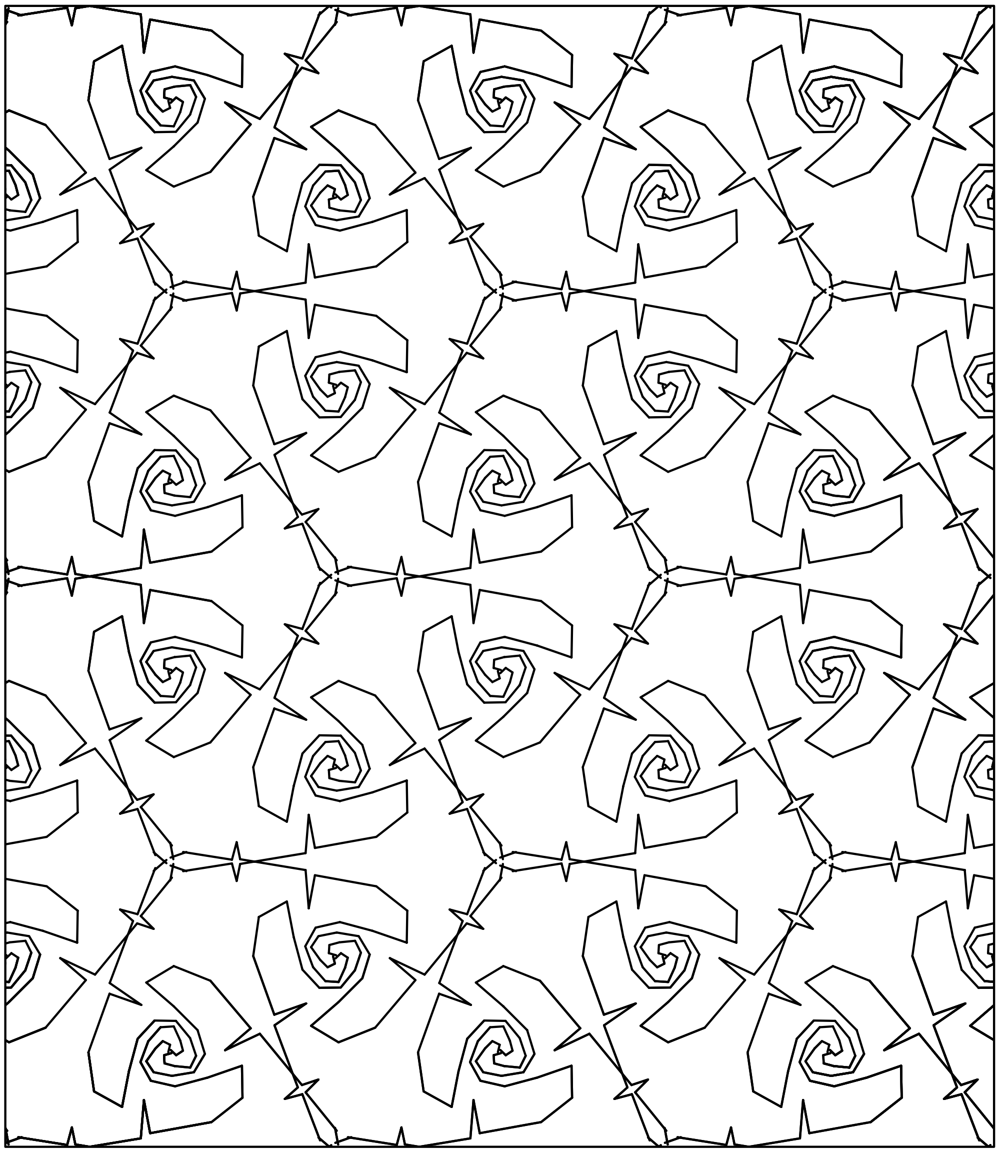

Figure 24: This is the pattern obtained by buying an order 3 cone point, a mirror, and an order 3 corner reflector.

Problem

- Can you justify the use of 'approximately' in the list above? Take the area A_R to be the union of all vertices, edges, and faces that intersect a disk of radius R in the plane, along with all edges of any face that intersects and all vertices of any edge that intersects. Can you show that the ratio of the true number to the estimated number is arbitrarily close to 1, for R high enough?

Definition. The *orbifold Euler characteristic* is $V - E + F$, where each vertex and edge is given weight $1/k$, where k is the order of symmetry which preserves it.

It is important to keep in mind the distinction between the topological Euler characteristic and the orbifold Euler characteristic. For instance, consider the billiard table orbifold, which is just a rectangle. In the orbifold Euler characteristic, the four corners each count $1/4$, the four edges count $-1/2$, and the face counts 1, for a total of 0. In contrast, the topological Euler characteristic is $4 - 4 + 1 = 1$.

Theorem. The quotient orbifold of any symmetry pattern in the Euclidean plane which has a bounded fundamental region has orbifold Euler number 0.

Sketch of proof: take a large area in the plane that is topologically a disk. Its Euler characteristic is 1. This is approximately equal to N times the orbifold Euler characteristic, for some large N, so the orbifold Euler characteristic must be 0.

How do the people at The Orbifold Shop figure its prices? The cost is based on the orbifold Euler characteristic: it costs \$1.00 to lower the orbifold Euler characteristic by 1. When they install a fancy new part, they calculate the difference between the new part and the part that was traded in.

For instance, to install a cone point, they remove an ordinary point. An ordinary point counts 1, while an order k cone point counts $1/k$, so the difference is $(k-1)/k$.

To install a handle, they arrange a map on the original orbifold so that it has a square face. They remove the square, and identify opposite edges of it. This identifies all four vertices to a single vertex. The net effect is to remove 1 face, remove 2 edges (since 4 are reduced to 2), and to remove 3 vertices.

The effect on the orbifold Euler characteristic is to subtract $1 - 2 + 3 = 2$, so the cost is \$2.00.

Problem

- Check the validity of the costs charged by The Orbifold Shop for the other parts of an orbifold.

To complete the connection between orbifold Euler characteristic and symmetry patterns, we would have to verify that each of the possible configurations of parts with orbifold Euler characteristic 0 actually does come from a symmetry pattern in the plane. This can be done in a straightforward way by explicit constructions. It is illuminating to see a few representative examples, but it is not very illuminating to see the entire exercise unless you go through it yourself.

29 Positive and negative Euler characteristic

A symmetry pattern on the sphere always gives rise to a quotient orbifold with positive Euler characteristic. In fact, if the order of symmetry is k, then the Euler characteristic of the quotient orbifold is $2/k$, since the Euler characteristic of the sphere is 2.

However, the converse is not true. Not every collection of parts costing less than \$2.00 can be put together to make a viable pattern for symmetry on the sphere. Fortunately, the experts at The Orbifold Shop know the four bad configurations which are too skimpy to be viable:

- A single cone point, with no other part, is bad.
- Two cone points, with no other parts, is a bad configuration unless they have the same order.
- A mirror with a single corner reflector, and no other parts, is bad.
- A mirror with only two corner reflectors, and no other parts, is bad unless they have the same order.

All other configurations are good. If they form an orbifold with positive orbifold Euler characteristic, they come from a pattern of symmetry on the sphere.

The situation for negative orbifold Euler characteristic is straightforward, but we will not prove it:

Theorem. Every orbifold with negative orbifold Euler characteristic comes from a pattern of symmetry in the hyperbolic plane with bounded fundamental domain. Every pattern of symmetry in the hyperbolic plane with compact fundamental domain gives rise to a quotient orbifold with negative orbifold Euler characteristic.

Since you can spend as much as you want, there are an infinite number of these.

30 A field guide to the orbifolds

The number 17 is *just right* for the number of types of symmetry patterns in the Euclidean plane: neither too large nor too small. It's large enough to make learning to recognize them a challenge, but not so large that this is an impossible task. It is by no means necessary to learn to distinguish the 17 types of patterns quickly, but if you learn to do it, it will give you a real feeling of accomplishment, and it is a great way to amaze and overawe your friends, at least if they're a bunch of nerds and geeks.

In this section, we will give some hints about how to learn to classify the patterns. However, we want to emphasize that this is a tricky business, and the only way to learn it is by hard work. As usual, when you analyze a pattern, you should look first for the mirror strings. The information in this section is meant as a way that you can learn to become more familiar with the 17 types of patterns, in a way that will help you to distinguish between them more quickly, and perhaps in some cases to be able to classify some of the more complicated patterns without seeing clearly and precisely what the quotient is. This kind of superficial knowledge is no substitute for a real visceral understanding of what the quotient orbifold is, and in every case you should go on and try to understand why the pattern is what you say it is while your friends are busy admiring your cleverness.

This information presented in this section has been gleaned from a cryptic manuscript discovered among the personal papers of John Conway after his death. For each of the 17 types of patterns, the manuscript shows a small piece of the pattern, the notation for the quotient orbifold, and Conway's idiosyncratic pidgin-Greek name for the corresponding pattern. These names are far from standard, and while they are unlikely ever to enter common use, we have found from our own experience that they are not wholly useless as a method for recognizing the patterns.

We will begin by discussing Conway's names for the orbifolds. A reproduction of Conway's manuscript appears at the end of the section. You should refer to the reproduction as you try to understand the basis for the names.

30.1 Conway's names

Each of Conway's 17 names consists of two parts, a *prefix* and a *descriptor*.

30.1.1 The prefix

The prefix tells the number of directions from which you can view the pattern without noticing any difference. The possibilities for the prefix are: *hexa-; tetra-; tri-; di-; mono-*.

For example, if you are looking at a standard brick wall, it will look essentially the same whether you stand on your feet or on your head. This will be true even if the courses of bricks in the wall do not run parallel to the ground, as they invariably do. Thus you can recognize right away that the brick-wall pattern is *di*-something-or-other In fact, it is *dirhombic*.

Another way to think about this is that if you could manage to turn the brick wall upside down, you wouldn't notice the difference. Again, this would be true even if you kept your head tilted to one side. More to the point, try looking at a dirhombic pattern drawn on a sheet of paper. Place the paper at an arbitrary angle, note what the pattern looks like in the large, and rotate the pattern around until it looks in the large like it did to begin with. When this happens, you will have turned the paper through half a rev. No matter how the pattern is tilted originally, there is always one and only one other direction from which it appear the same in the large.

This 'in the large' business means that you are not supposed to notice if, after twisting the paper around, the pattern appears to have been shifted by a translation. You don't have to go grubbing around looking for some pesky little point about which to rotate the pattern. Just take the wide, relaxed view.

30.1.2 The descriptor

The descriptor represents an attempt on Conway's part to unite patterns that seem more like each other than they do like the other patterns. The possibilities for the descriptor are: *scopic; tropic; gyro; glide; rhombic*.

The *scopic* patterns are those that emerge from kaleidoscopes: $*632 =$ *hexascopic*; $*442 =$ *tetrascopic*; $*333 =$ *triscopic*; $*2222 =$ *discopic*; $** =$ *monoscopic*;

Their $*$-less counterparts are the *tropic* patterns (from the Greek for 'turn'): $632 =$ *hexatropic*; $442 =$ *tetratropic*; $333 =$ *tritropic*; $2222 =$ *ditropic*; $\bullet =$ *monotropic*.

With the scopic patterns, it's all done with mirrors, while with the tropic patterns, it's all done with gyration points. The two exceptions are: $** =$ *monoscopic*; $\bullet =$ *monotropic*. There is evidence that Conway did not consider these to be exceptions, on the grounds that 'with the scopics it's all done with mirrors and translations, while with the tropics, it's all done with turnings and translations'.

The *gyro* patterns contain both mirrors and gyration points: $4*2 =$ *tetragyro*; $3*3 =$ *trigyro*; $22* =$ *digyro*.

Since both tropic and gyro patterns involve gyration points, there is a real possibility of confusing the names. Strangely, it is the tropic patterns that are the more closely connected to gyration points. In practice, it seems to be easy enough to draw this distinction correctly, probably because the tropics correspond closely to the scopics, and 'tropic' rhymes with 'scopic'. Conway's view appears to have been that a gyration point, which is a point of rotational symmetry that does NOT lie on a mirror, becomes ever so much more of a gyration point when there are mirrors around that it might have been tempted to lie on, and that therefore patterns that contain both gyration points and mirrors are more gyro than patterns with gyration points but no mirrors.

The *glide* patterns involve glide-reflections: $22\circ =$ *diglide*; $\circ\circ =$ *monoglide*.

The glide patterns are the hardest to recognize. The quotient orbifold of the diglide pattern is a projective plane with two cone points; the quotient of the monoglide patterns is a Klein bottle. When you run up against one of these patterns, you just have to sweat it out. One trick is that when you meet something that has glide-reflections but not much else, then you decide that it must be either a diglide or a monoglide, and you can distinguish between them by deciding whether it's a *di-* or a *mono-* pattern, which is a distinction that is relatively easy to make. Another clue to help distinguish these two cases is that a diglide pattern has glides in two different directions, while a monoglide has glides in only one direction. Yet another clue is that in a monoglide you can often spot two disjoint Möbius strips within the quotient orbifold, corresponding to the fact that the quotient orbifold for a monoglide pattern is a Klein bottle, which can be pieced together from two Möbius strips. These two disjoint Möbius strips arise from the action of glide-reflections along parallel but inequivalent axes.

The *rhombic* patterns often give a feeling of rhombosity: $2{*}22$ = *dirhombic*; $*\circ$ = *monorhombic*.

An ordinary brick wall is dirhombic; it can be made monorhombic by breaking the gyrational symmetry. The quotient of a monorhombic pattern is a Möbius strip. Like the two glide quotients, it is non-orientable, but it is much easier to identify because of the presence of the mirrors.

30.2 How to learn to recognize the patterns

As you will see, Conway's manuscript shows only a small portion of each of the patterns. A very worthwhile way of becoming acquainted with the patterns is to draw larger portions of the patterns, and then go through and analyze each one, to see why it has the stated notation and name. You may wish to make flashcards to practice with. When you use these flashcards, you should make sure that you can not only remember the correct notation and name, but also that you can analyze the pattern quickly, locating the distinguishing features. This is important because the patterns you will see in the real world won't be precisely these ones.

Another hint is to keep your eyes open for symmetrical patterns in the world around you. When you see a pattern, copy it onto a flashcard, even if you cannot analyze it immediately. When you have determined the correct analysis, write it on the back and add it to your deck.

30.3 The manuscript

What follows is an exact reproduction of Conway's manuscript. In addition to the 17 types of repeating patterns, Conway's manuscript also gives tables of the 7 types of frieze patterns, and of the 14 types of symmetrical patterns on the sphere. These parts of the manuscript appear to be mainly gibberish. We reproduce these tables here in the hope that they may someday come to the attention of a scholar who will be able to make sense of them.

31 Geometry on the sphere

We want to explore some aspects of geometry on the surface of the sphere. This is an interesting subject in itself, and it will come in handy later on when we discuss Descartes's angle-defect formula.

31.1 Discussion

Great circles on the sphere are the analogs of straight lines in the plane. Such curves are often called *geodesics*. A *spherical triangle* is a region of the sphere bounded by three arcs of geodesics.

Problems

1. Do any two distinct points on the sphere determine a unique geodesic? Do two distinct geodesics intersect in at most one point?

2. Do any three 'non-collinear' points on the sphere determine a unique triangle? Does the sum of the angles of a spherical triangle always equal π? Well, no. What values can the sum of the angles take on?

The area of a spherical triangle is the amount by which the sum of its angles exceeds the sum of the angles (π) of a Euclidean triangle. In fact, for any spherical polygon, the sum of its angles minus the sum of the angles of a Euclidean polygon with the same number of sides is equal to its area.

A proof of the area formula can be found in Chapter 9 of Weeks, *The Shape of Space*.

Conway 47:1

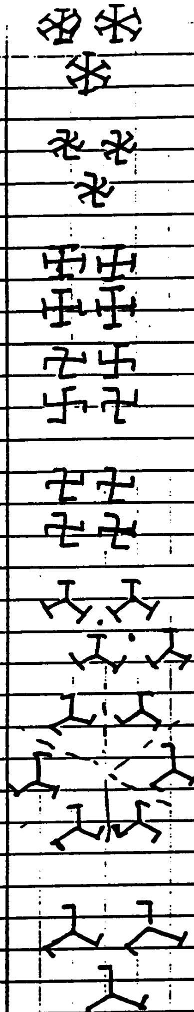

*632 hexascopic

632 hexatropic

*442 tetrascopic

4*2 tetragyro

442 tetratropic

discopic

*333 triscopic

dirhombic

3*3 trigyro

digyro

333 tritropic

diglide

ditropic

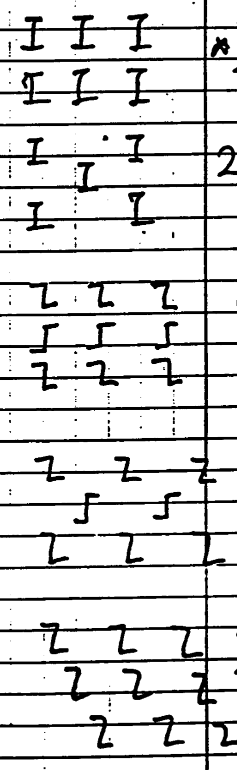

*2222

2*22

22*

22°

2222

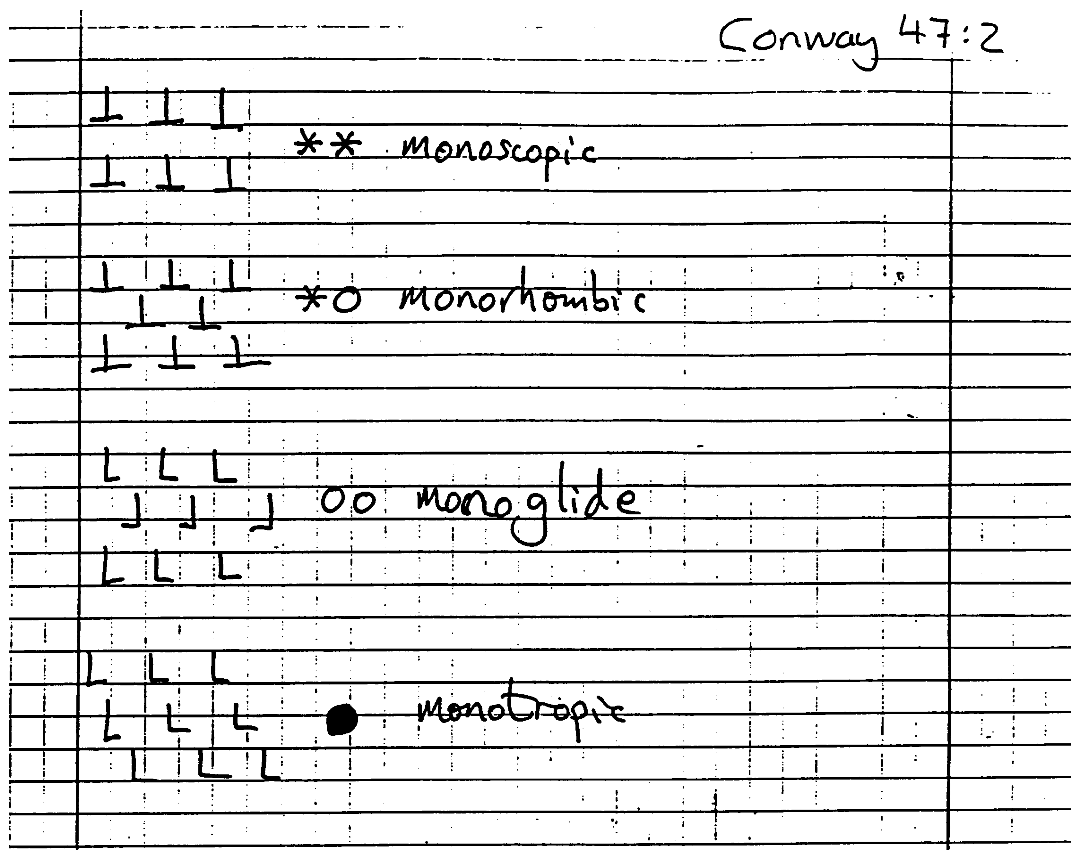
Conway 47:2
** monoscopic
*o monorhombic
oo monoglide
● monotropic

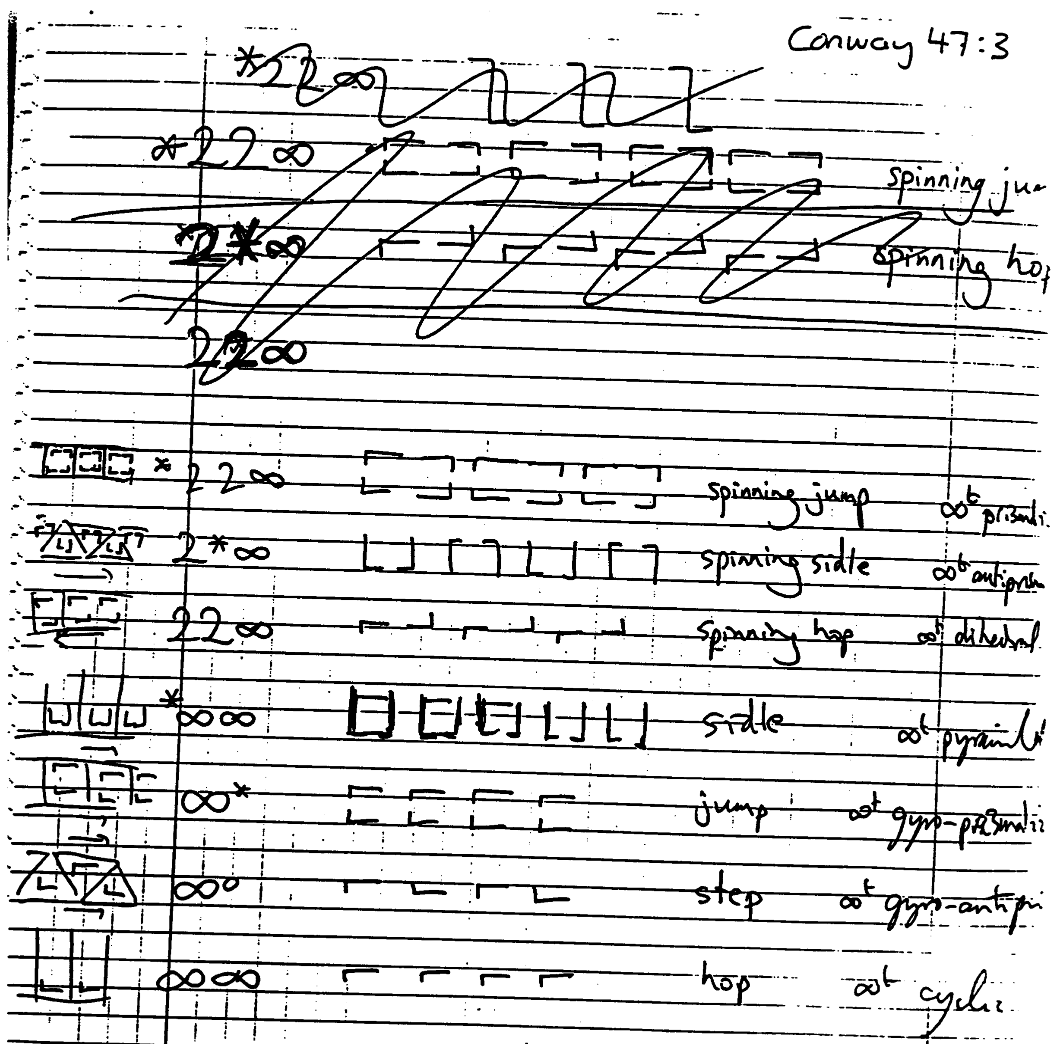

Conway 47:3
*22∞
*22∞
spinning ju
2*∞
spinning ho
22∞
*22∞
spinning jump
∞t prismati
2*∞
spinning sidle
∞t antipri
22∞
spinning hop
∞t dihedral
*∞∞
sidle
∞t pyramid
∞*
jump
∞t gyro-prismati
∞∞
step
∞t gyro-antipri
∞∞
hop
∞t cyclic

Conway 47:4

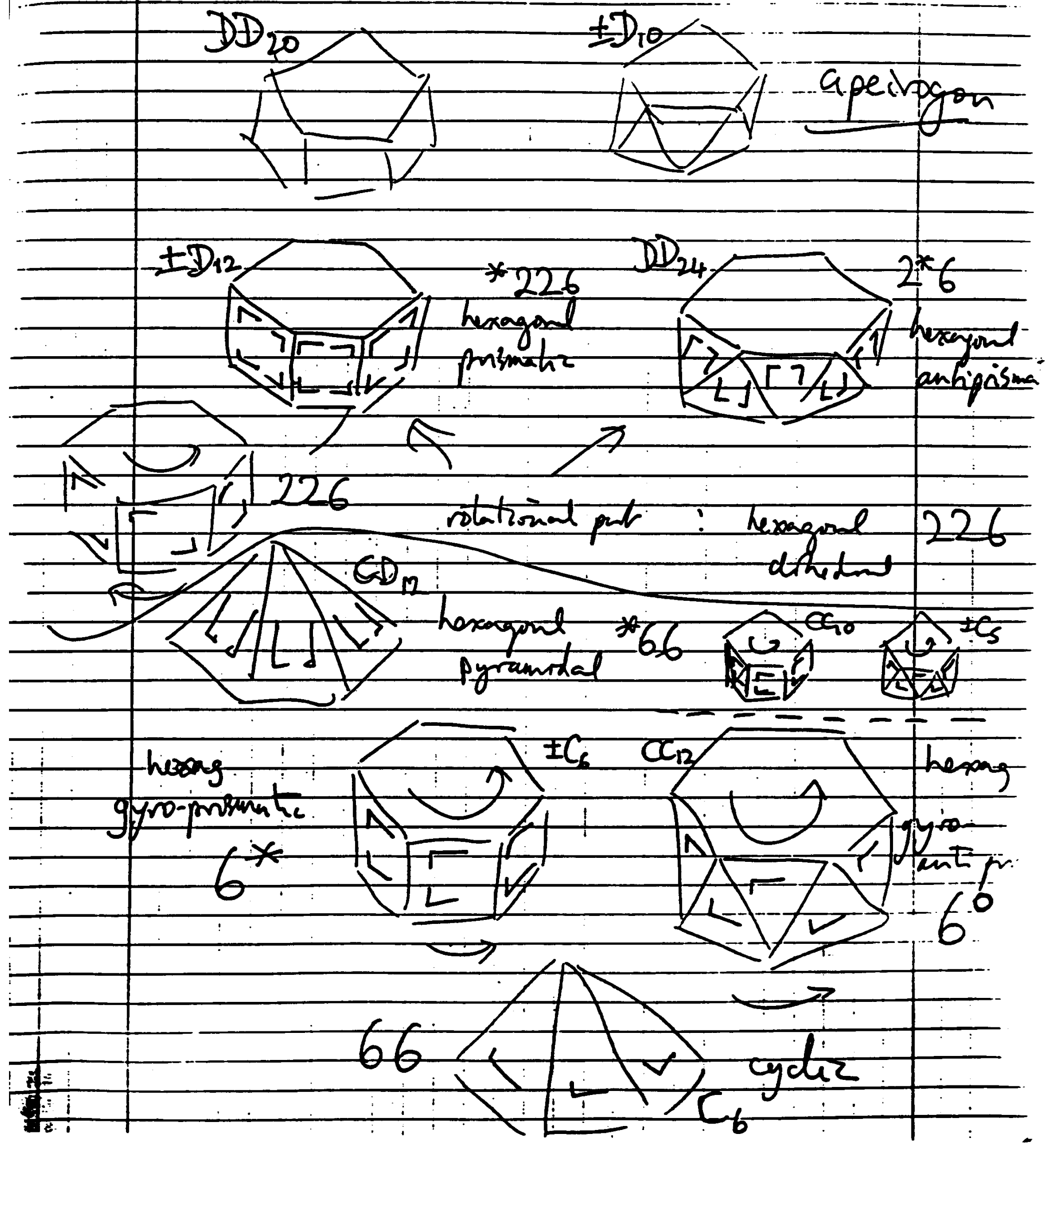

DD20
±D10
apeirogon
±D12
*226
hexagonal prismatic
DD24
2*6
hexagonal antiprismatic
226
rotational part : hexagonal dihedral
226
CD12
hexagonal pyramidal
*66
CC10
±C5
hexag gyro-prismatic
6*
±C6
CC12
hexag gyro-anti pr
6°
66
cyclic
C6

Conway ~~47~~:5

geom name	orbifold	alg not^n	abstract struc	algebraic name
icosahedral ref^n	*532	$\pm I$	$2\times A_5$	diplo-icosahedral
icosahedral rot^n	532	I	A_5	icosahedral (rotation)
octahedral ref^n	*432	$\pm O$	$2\times S_4$	diplo-octahedral
octahedral rot^n	432	O	S_4	octahedral (rotation)
tetrahedral ref^n	*332	TO	S_4	tetra-octahedral
gyro-octahedral	3*2	$\pm T$	$2\times A_4$	diplo-tetrahedral
tetrahedral rot^n	332	T	A_4	tetrahedral (rotation)

abs. struc	geom name	orbifold	alg not^n	abstract struc	algebraic name
$2\times D_{2n}$	prismatic (ref^n)	*22n	$\pm D_{2n}$	$2\times D_{2n}$	diplo-dihedral
D_{4n}	antiprismatic (ref^n)	2*n	DD_{4n}	D_{4n}	dihedro-dihedral
D_{2n}	Prism rot^n or Antiprism rot^n dihedral (rot^n)	22n	D_{2n}	D_{2n}	dihedral (rotation)
D_{2n}	pyramidal (ref^n)	*nn	CD_{2n}	D_{2n}	cyclo-dihedral
$2\times C_n$	gyro-prismatic	n*	$\pm C_n$	$2\times C_n$	diplo-cyclic
C_{2n}	gyro-antiprismatic	n^{o}	CC_{2n}	C_{2n}	cyclo-cyclic
C_n	(Pyr. rot^n) cyclic (rot^n)	nn	C_n	C_n	cyclic (rotation)

⧗ indicates = when n is even, X when n is odd.

Thus $6^* = \pm C_6$, $6^{o} = CC_{12}$ but $5^* = CC_{10}$, $5^{o} = \pm C_5$

$*226 = \pm D_{12}$, $2*6 = DD_{24}$ but $*225 = DD_{20}$, $2*6 = \pm D_{10}$

32 The angle defect of a polyhedron

The *angle defect* at a vertex of a polygon is defined to be 2π minus the sum of the angles at the corners of the faces at that vertex. For instance, at any vertex of a cube there are three angles of $\pi/2$, so the angle defect is $\pi/2$. You can visualize the angle defect by cutting along an edge at that vertex, and then flattening out a neighborhood of the vertex into the plane. A little gap will form where the slit is: the angle by which it opens up is the angle defect.

The *total angle defect* of the polyhedron is gotten by adding up the angle defects at all the vertices of the polyhedron. For a cube, the total angle defect is $8 \times \pi/2 = 4\pi$.

Problems

1. What is the angle sum for a polygon (in the plane) with n sides?

2. Determine the total angle defect for each of the 5 regular polyhedra, and for various other polyhedra.

33 Descartes's Formula.

The *angle defect* at a vertex of a polygon was defined to be the amount by which the sum of the angles at the corners of the faces at that vertex falls short of 2π and the *total angle defect* of the polyhedron was defined to be what one got when one added up the angle defects at all the vertices of the polyhedron. We call the total defect T. Descartes discovered that there is a connection between the total defect, T, and the Euler Number $E - V - F$. Namely,

$$T = 2\pi(V - E + F). \tag{1}$$

Here are two proofs. They both use the fact that the sum of the angles of a polygon with n sides is $(n-2)\pi$.

33.1 First proof

Think of $2\pi(V - E + F)$ as putting $+2\pi$ at each vertex, -2π on each edge, and $+2\pi$ on each face.

We will try to cancel out the terms as much as possible, by grouping within polygons.

For each edge, there is -2π to allocate. An edge has a polygon on each side: put $-\pi$ on one side, and $-\pi$ on the other.

For each vertex, there is $+2\pi$ to allocate: we will do it according to the angles of polygons at that vertex. If the angle of a polygon at the vertex is a, allocate a of the 2π to that polygon. This leaves something at the vertex: the angle defect.

In each polygon, we now have a total of the sum of its angles minus $n\pi$ (where n is the number of sides) plus 2π. Since the sum of the angles of any polygon is $(n-2)\pi$, this is 0. Therefore,

$$2\pi(V - E + F) = T.$$

33.2 Second proof

We begin to compute:

$$T = \sum_{\text{Vertices}} \text{the angle defect at the vertex.}$$

$$= \sum_{\text{Vertices}} (2\pi - \text{the sum of the angles at the corners of those faces that meet at the vertex}).$$

$$= 2\pi V - \sum_{\text{Vertices}} (\text{the sum of the angles at the corners of those faces that meet at the vertex}).$$

$$= 2\pi V - \sum_{\text{Faces}} \text{the sum of the interior angles of the face.}$$

$$= 2\pi V - \sum_{\text{Faces}} (n_f - 2)\pi.$$

Here n_f denotes the number of edges on the face f.

$$T = 2\pi V - \sum_{\text{Faces}} n_f\pi + \sum_{\text{Each face}} 2\pi.$$

Thus

$$T = 2\pi V - (\sum_{\text{Faces}} \text{the number of edges on the face} \cdot \pi) + 2\pi F.$$

If we sum the number of edges on each face over all of the faces, we will have counted each edge **twice**. Thus

$$T = 2\pi V - 2E\pi + 2\pi F.$$

Whence,

$$T = 2\pi(V - E + F).$$

Problems

1. Discuss both proofs with the aim of understanding them.
2. Draw a sketch of the first proof.
3. Discuss the differences between the two proofs. Can you describe the ways in which they are different? Which is easier to understand? Which is more pleasing? Which is more conceptual?

34 The celestial image of a polyhedron

We want now to discuss the celestial image of a polyhedron, and use it to get yet another proof of Descartes's angle-defect formula.

Problems

1. What pattern is traced out on the celestial sphere when you move a flashlight around on the surface of a cube, keeping its tail as flat as possible on the surface? What is the celestial pattern for a dodecahedron?
2. On a convex polyhedron, the celestial image of a region containing a solitary vertex v where three faces meet is a triangle. Show that the three angles of this celestial triangle are the supplements of the angles of the three faces that meet at v.
3. Show that the area of this celestial triangle is the angle defect at v.
4. Show that the total angle defect of a convex polyhedron is 4π.

35 Curvature of surfaces

If you take a flat piece of paper and bend it gently, it bends in only one direction at a time. At any point on the paper, you can find at least one direction through which there is a straight line on the surface. You can bend it into a cylinder, or into a cone, but you can never bend it without crumpling or distorting to the get a portion of the surface of a sphere.

If you take the skin of a sphere, it cannot be flattened out into the plane without distortion or crumpling. This phenomenon is familiar from orange peels or apple peels. Not even a small area of the skin of a sphere can be flattened out without some distortion, although the distortion is very small for a small piece of the sphere. That's why rectangular maps of small areas of the earth work pretty well, but maps of larger areas are forced to have considerable distortion.

The physical descriptions of what happens as you bend various surfaces without distortion do not have to do with the topological properties of the surfaces. Rather, they have to do with the *intrinsic geometry* of the surfaces. The intrinsic geometry has to do with geometric properties which can be detected by measurements along the surface, without considering the space around it.

There is a mathematical way to explain the intrinsic geometric property of a surface that tells when one surface can or cannot be bent into another. The mathematical concept is called the *Gaussian curvature* of a surface, or often simply the *curvature* of a surface. This kind of curvature is not to be confused with the curvature of a curve. The curvature of a curve is an extrinsic geometric property, telling how it is bent in the plane, or bent in space. Gaussian curvature is an intrinsic geometric property: it stays the same no matter how a surface is bent, as long as it is not distorted, neither stretched or compressed.

To get a first qualitative idea of how curvature works, here are some examples.

A surface which bulges out in all directions, such as the surface of a sphere, is *positively curved*. A rough test for positive curvature is that if you take any point on the surface, there is some plane touching the surface at that point so that the surface lies all on one side except at that point. No matter how you (gently) bend the surface, that property remains.

A flat piece of paper, or the surface of a cylinder or cone, has 0 curvature.

A saddle-shaped surface has negative curvature: every plane through a point on the saddle actually cuts the saddle surface in two or more pieces.

Problem

- What surfaces can you think of that have positive, zero, or negative curvature.

Gaussian curvature is a numerical quantity associated to an area of a surface, very closely related to angle defect. Recall that the angle defect of a polyhedron at a vertex is the angle by which a small neighborhood of a vertex opens up, when it is slit along one of the edges going into the vertex.

The total Gaussian curvature of a region on a surface is the angle by which its boundary opens up, when laid out in the plane. To actually measure Gaussian curvature of a region bounded by a curve, you can cut out a narrow strip on the surface in neighborhood of the bounding curve. You also need to cut open the curve, so it will be free to flatten out. Apply it to a flat surface, being careful to distort it as little as possible. If the surface is positively curved in the region inside the curve, when you flatten it out, the curve will open up. The angle between the tangents to the curve at the two sides of the cut is the total Gaussian curvature. This is like angle defect: in fact, the total curvature of a region of a polyhedron containing exactly one vertex is the angle defect at that vertex. You must pay attention not just to the angle between the ends of the strip, but how the strip curled around, keeping in mind that the standard for zero curvature is a strip which comes back and meets itself. Pay attention to π's and 2π's.

If the total curvature inside the region is negative, the strip will curl around further than necessary to close. The curvature is negative, and is measured by the angle by which the curve overshoots.

A less destructive way to measure total Gaussian curvature of a region is to apply narrow strips of paper to the surface, e.g., masking tape. They can be then be removed and flattened out in the plane to measure the curvature.

Problems

1. Measure the total Gaussian curvature of

   (a) a cabbage leaf.

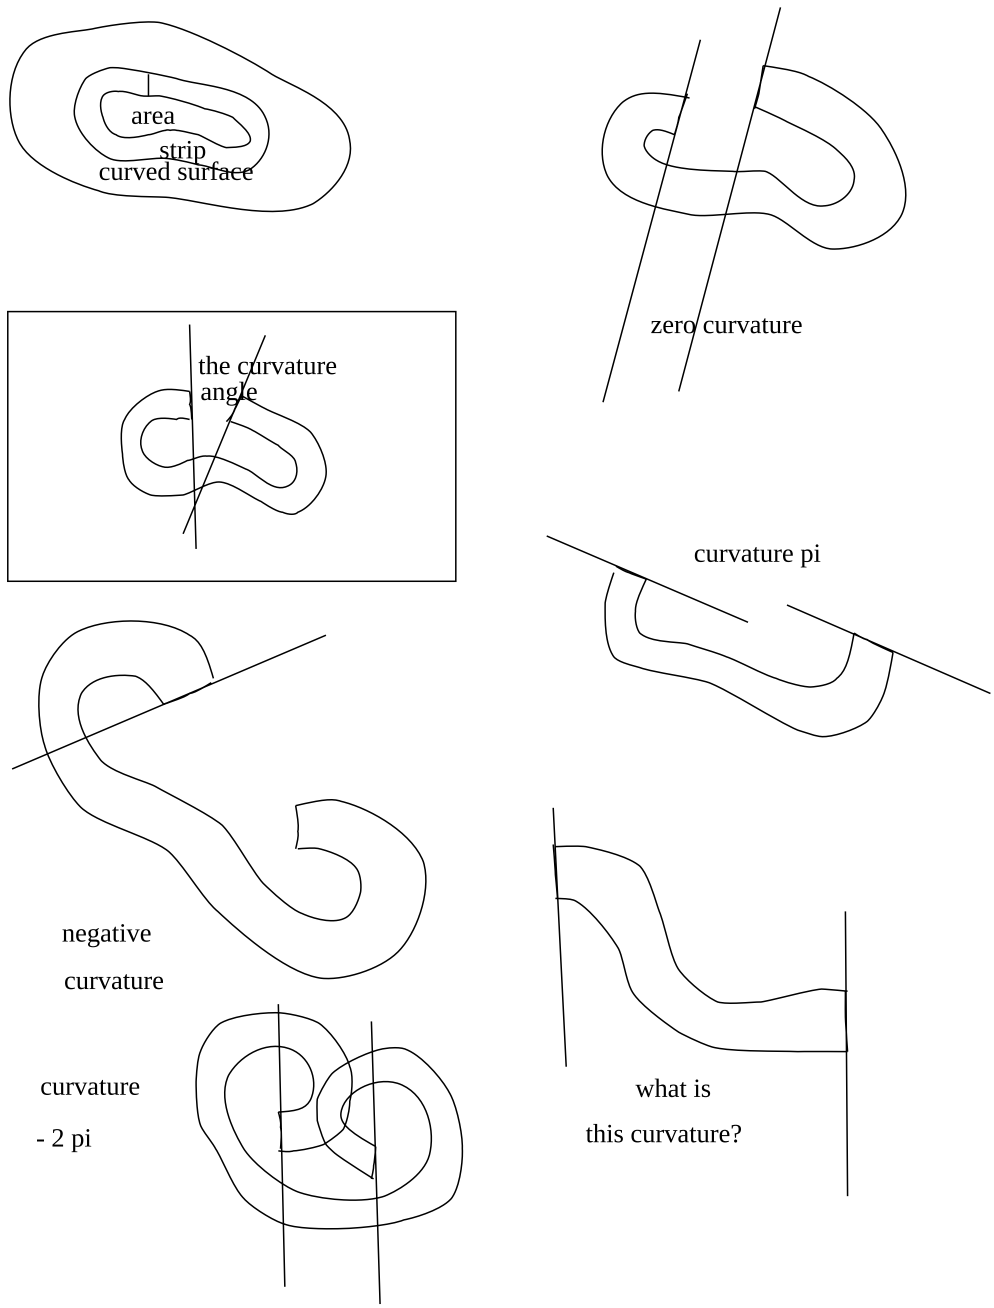

Figure 25: This diagram illustrates how to measure the total Gaussian curvature of a patch by cutting out a strip which bounds the patch, and laying it out on a flat surface. The angle by which the strip 'opens up' is the total Gaussian curvature. You must pay attention not just to the angle between the lines on the paper, but how it got there, keeping in mind that the standard for zero curvature is a strip which comes back and meets itself. Pay attention to π's and 2π's.

(b) a lettuce leaf

(c) a piece of banana peel

(d) a piece of potato skin

If you take two adjacent regions, bounded by a θ-shape, is the total curvature in the whole equal to the sum of the total curvature in the parts? Why?

2. The angle defect of a convex polyhedron at one of its vertices can be measured by rolling the polyhedron in a circle around its vertex. Mark one of the edges, and rest it on a sheet of paper. Mark the line on which it contacts the paper. Now roll the polyhedron, keeping the vertex in contact with the paper. When the given edge first touches the paper again, draw another line. The angle between the two lines (in the area where the polyhedron did not touch) is the angle defect. In fact, the area where the polyhedron did touch the paper can be rolled up to form a paper model of a neighborhood of the vertex in question.

3. A polyhedron can also be rolled in a more general way. Mark some closed path on the surface of the polyhedron, avoiding vertices. Lay the polyhedron on a sheet of paper so that part of the curve is in contact. Mark the position of one of the edges in contact with the paper. now roll the polyhedron, along the curve, until the original face is in contact again, and mark the new position of the same edge. What is the angle between the original position of the line, and the new position of the line?

4. On a polyhedron, what is the curvature inside a region containing a single vertex? two vertices? all but one vertex? all the vertices?

5. What is the curvature inside the region on a sphere *exterior* to a tiny circle?

36 Clocks and curvature

The total curvature of any surface topologically equivalent to the sphere is 4π. This can be seen very simply from the definition of the curvature of a region in terms of the angle of rotation when the surface is rolled around on the plane; the only problem is the perennial one of keeping proper track of multiples of π when measuring the angle of rotation. Since are trying to show that the total curvature is a specific multiple of π, this problem is crucial. So to begin with let's think carefully about how to reckon these angles correctly

36.1 Clocks

Suppose we have a number of clocks on the wall. These clocks are good mathematician's clocks, with a 0 up at the top where the 12 usually is. (If you think about it, 0 o'clock makes a lot more sense than 12 o'clock: With the 12 o'clock system, a half hour into the new millennium on 1 Jan 2001, the time will be 12:30 AM, the 12 being some kind of hold-over from the departed millennium.)

Let the clocks be labelled A, B, C, ... To start off, we set all the clocks to 0 o'clock. (little hand on the 0; big hand on the 0), Now we set clock B ahead half an hour so that it now the time it tells is 0:30 (little hand on the 0 (as they say); big hand on the 6). What angle does its big hand make with that of clock A? Or rather, through what angle has its big hand moved relative to that of clock A? The angle is π. If instead of degrees or radians, we measure our angles in *revs* (short for *revolutions*), then the angle is 1/2 rev. We could also say that the angle is 1/2 hour: as far as the big hand of a clock is concerned, an hour is the same as a rev.

Now take clock C and set it to 1:00. Relative to the big hand of clock A, the big hand of C has moved through an angle of 2π, or 1 rev, or 1 hour. Relative to the big hand of B, the big hand of C has moved through an angle of π, or 1/2 rev. Relative to the big hand of C, the big hand of A has moved through an angle of -2π, or -1 rev, and the big hand of B has moved $-\pi$, or -1 rev.

36.2 Curvature

Now let's describe how to find the curvature inside a disk-like region R on a surface S, i.e. a region topologically equivalent to a disk. What we do is cut a small circular band running around the boundary of the region, cut the band open to form a thin strip, lay the thin strip flat on the plane, and measure the angle between the lines at the two end of the strip. In order to keep the π's straight, let us go through this process very slowly and carefully.

To begin with, let's designate the two ends of the strip as the *left end* and the *right end* in such a way that traversing the strip from the left end to the right end corresponds to circling *clockwise* around the region. We begin by fixing the left-hand end of the strip to the wall so that the straight edge of the cut at the left end of the strip—the cut that we made to convert the band into a strip—runs straight up and down, parallel to the big hand of clock A, and so that the strip runs off toward the right. Now we move from left to right along the strip, i.e. clockwise around the boundary of the region, fixing the strip so that it lies as flat as possible, until we come to the right end of the strip. Then we look at the cut bounding the right-hand end of the strip, and see how far it has turned relative to the left-hand end of the strip. Since we were so careful in laying out the left-hand end of the strip, our task in reckoning the angle of the right-hand end of the strip amounts to deciding what time you get if you think of the right-hand end of the strip as the big hand of a clock. The curvature inside the region will correspond to the amount by which the time told by the right-hand end of the strip falls short of 1:00.

For instance, say the region R is a tiny disk in the Euclidean plane. When we cut a strip from its boundary and lay it out as described above, the time told by its right hand end will be precisely 1:00, so the curvature of R will be exactly 0. If R is a tiny disk on the sphere, then when the strip is laid out the time told will be just shy of 1:00, say 0:59, and the curvature of the region will be $\frac{1}{60}$ rev, or $\frac{\pi}{30}$.

When the region R is the lower hemisphere of a round sphere, the strip you get will be laid out in a straight line, and the time told by the right-hand end will be 0:00, so the total curvature will be 1 rev, i.e. 2π. The total curvature of the upper hemisphere is 2π as well, so that the total curvature of the sphere is 4π.

Another way to see that the total curvature of the sphere is 4π is to take as the region R the *outside* of a small circle on the sphere. When we lay out a strip following the prescription above, being sure to traverse the boundary of the region R in the clockwise sense as viewed from the point of view of the region R, we see that the time told by the right hand end of the strip is very nearly -1 o'clock! The precise time will be just shy of this, say -1:59, and the total curvature of the region will then be $1\frac{59}{60}$ revs. Taking the limit, the total curvature of the sphere is 2 revs, or 4π.

But this last argument will work equally well on any surface topologically equivalent to a sphere, so any such surface has total curvature 4π.

36.3 Where's the beef?

This proof that the total curvature of a topological sphere is 4π gives the definite feeling of being some sort of trick. How can we get away without doing any work at all? And why doesn't the argument work equally well on a torus, which as we know should have total curvature 0? What gives?

What gives is the lemma that states that if you take a disklike region R and divide it into two disklike subregions R_1 and R_2, then the curvature inside R when measured by laying out its boundary is the sum of the curvatures inside R_1 and R_2 measured in this way. This lemma might seem like a tautology. Why should there be anything to prove here? How could it fail to be the case that the curvature inside the whole is the sum of the curvatures inside the parts? The answer is, it could fail to be the case by virtue of our having given a faulty definition. When we define the curvature inside a region, we have to make sure that the quantity we're defining has the additivity property, or the definition is no good. Simply calling some quantity the curvature inside the region will not make it have this additivity property. For instance, what if we had defined the curvature inside a region to be 4π, no matter what the region? More to the point, what if in the definition of the curvature inside a region we had forgotten the proviso that the region R be disklike? Think about it.