2

Surfaces

2.1 The shape of the world. How do we know the world is round?

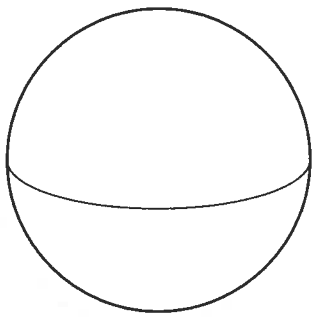

A sphere

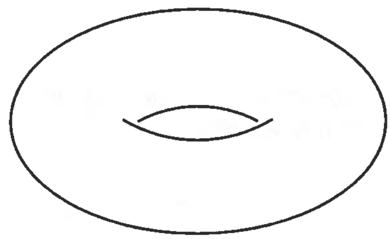

A torus

Is it possible that we live on a torus, as opposed to a sphere? Pictures taken from space give uncontestable proof that the world is not a big donut, but people knew that the surface of the Earth was a sphere long before humans developed spaceflight. What was the evidence?

Task 2.1.1: List the evidence for why the surface of the Earth is a sphere.

Task 2.1.2: If the world was perpetually shrouded in dense fog and the terrain was extremely lumpy, would the evidence you gave in Task 2.1.1 still work? Do your methods permit the inhabitants to distinguish between the two worlds below?

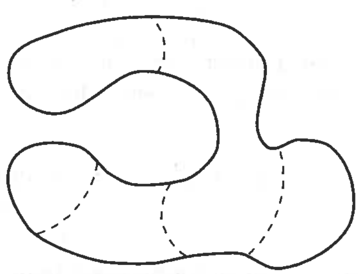

A lumpy sphere

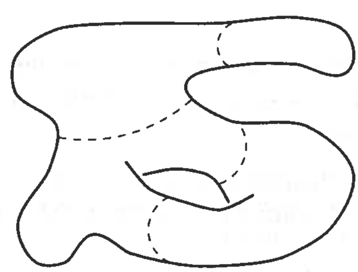

A lumpy torus

If you confine yourself to just a small bit of your world then it is not possible to determine whether you live on a sphere, or a torus, or on some other surface. You must travel over the whole surface in order to truly understand it. The inhabitants of a world perpetually shrouded in dense fog would not have an easy task to determine the shape of their world. Keep that picture in mind as you read the rest of the chapter.

Note. In this chapter we are studying surfaces, not 'solid' objects. When we say *sphere* you should have in mind the outer skin of a ball. Similarly, the torus can be thought of as the outer skin of a donut.

One way to explore the whole world is to form a loop of people all facing the same way, and have them walk across the surface while holding hands.

A large group of people hold hands to form a big loop on the Earth.

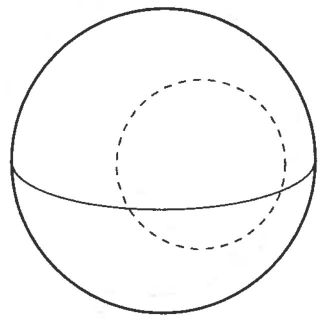

As the people walk forward they all bunch together in one spot.

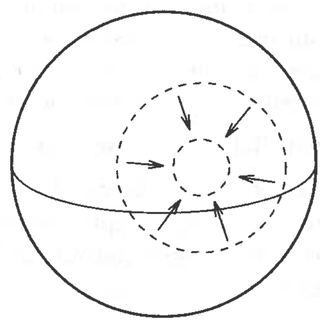

Task 2.1.3: Explain why any loop of people on a sphere will end up in one spot as they walk forward. Note: the loop may first get bigger before it gets smaller.

Definition. A loop on a surface is called **contractible** if the loop can be shrunk to a point without leaving the surface.

The result of Task 2.1.3 can be rephrased as, "On the sphere, all loops are contractible." The situation is quite different on the torus.

Task 2.1.4: Convince yourself that neither of these loops on the torus is contractible.

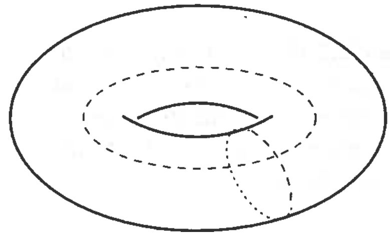

Keep in mind that the loops in Task 2.1.4 are on the *surface.* The loops must stay on the surface; they are not allowed to 'dig through the donut.'

Loops on a torus can have amazing properties unlike anything on a sphere:

A large loop of people are holding hands. They begin to walk forward.

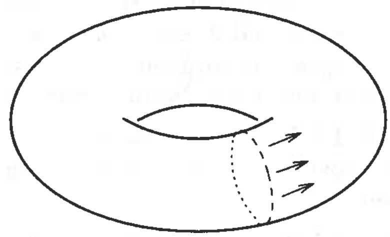

After a while they will have walked halfway around the world.

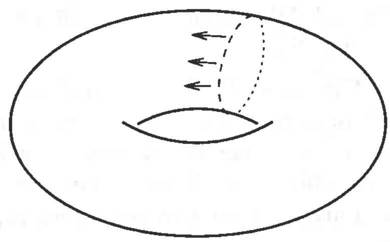

Eventually everybody gets back where they started.

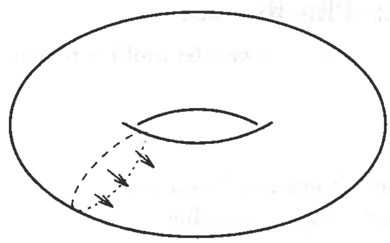

Task 2.1.5: Invent a story about the inhabitants of a torus planet and their attempts to determine the shape of their world. Would the members of the 'human chain' described above be surprised when they all returned to their original spot?

Task 2.1.6: What will happen as this loop of people starts walking forward? Is this different from what happens with the loop shown above?

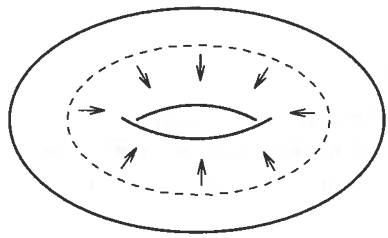

Task 2.1.7: How would your answers to the last two Tasks change if the loops were on the lumpy torus shown at the beginning of the chapter, and everything was shrouded in dense fog? If any of your answers involved measuring distances or looking outside the torus, then those reasons are no longer valid.

Part of Tasks 2.1.6 and 2.1.7 asks you to determine how the inhabitants of a torus can tell the difference between the two loops shown in Task 2.1.4. If you are unable to measure distances properly, as would happen if the torus were very lumpy, and if there was no way to look at the torus from the outside, then it is *impossible* to distinguish between those two loops. This is an important concept which may become more clear in the next section.

Task 2.1.8: The two loops in Task 2.1.4 cross at one point. Explain why, on the plane, or on the sphere, it is impossible for two loops to cross at only one point.

Task 2.1.9: Draw a loop on the torus which crosses each of the loops in Task 2.1.4 exactly once.

Task 2.1.10: Draw a different pair of loops on the torus which cross at only one point.

The book *The Shape of Space* [SoS], by Jeff Weeks, contains an amusing story of a creature who travels around its world drawing lines on the ground. Two of its longer trips result in a huge blue loop and a huge red loop drawn around the world. It turns out that these loops cross each other only once. This was quite confusing to the creature, but all it means is that its world is not a sphere.

2.2 The flat torus

Here is a way to build a torus. Start with a square:

The arrows show you which edges to glue together.

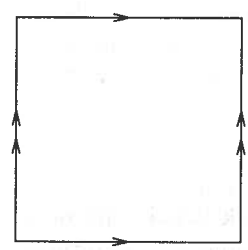

continued...

First glue the top edge to the bottom to make a cylinder.

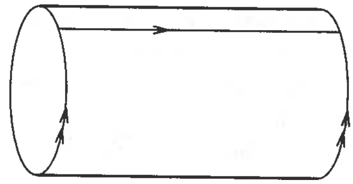

Then bend the cylinder...

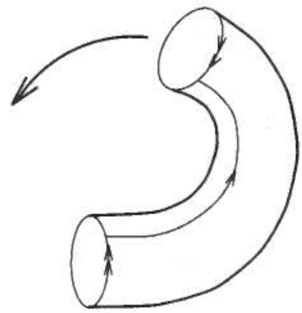

and glue the other edges to complete the torus. The glue lines become the loops shown in Task 2.1.4.

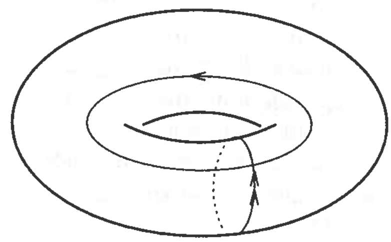

It is easy to draw a square, so we will use a square to represent a torus whenever possible. We call this the **flat torus**. Putting 'arrows' on the sides of a square shows that we mean for the opposite edges to be glued, with the understanding that the glue lines become the two curves shown directly above. Here is an example. Both figures represent the same loop on the torus:

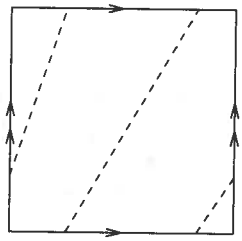

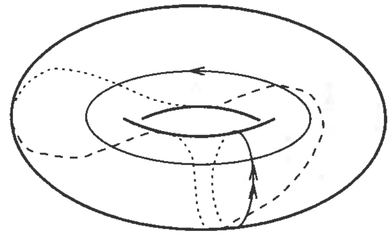

It is worth spending some time looking at that example. On the square, the places where the dotted line hits opposite edges must 'match up.' This ensures that, when the edges are glued, the dotted line will form one continuous loop. To verify that the two pictures represent the same thing, it is easiest to look at where the loop intersects the glue lines. The left picture shows the loop broken into three segments. You need only check that each segment in the right picture is drawn properly.

Here is another example. The segments are marked so that it is easier to see how things correspond.

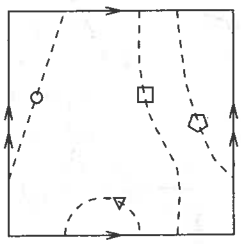

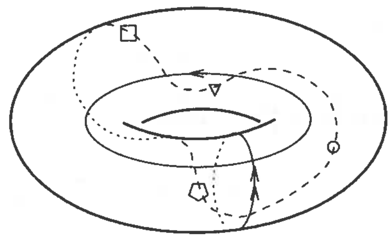

Task 2.2.1: Draw your answers to Tasks 2.1.9 and 2.1.10 on the flat torus.

The next section will provide plenty of practice drawing on the torus.

We mentioned earlier that it is not possible, in general, to distinguish between the two curves shown in Task 2.1.4. Thinking in terms of the flat torus will make this clear. If you repeat the procedure of gluing the square to give a torus, but first glue the left edge to the right edge to make a 'vertical' cylinder, the result will be a torus with glue lines the reverse of those shown above. Since it is impossible to distinguish between the edges of the flat torus, it is impossible to distinguish between the glue lines.

2.3 Graphs on the torus

In the previous chapter we showed that it is impossible to draw the graphs K_5 and $K_{3,3}$ without crossing edges. Actually, that statement isn't quite right. It is impossible to draw those graphs on the *plane* or the *sphere* without crossing edges. We will see that both K_5 and $K_{3,3}$ can be drawn on the torus. Here is one way to draw K_5, given in both representations:

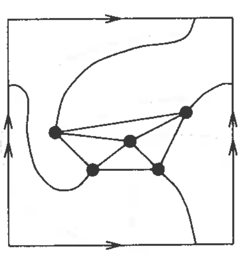

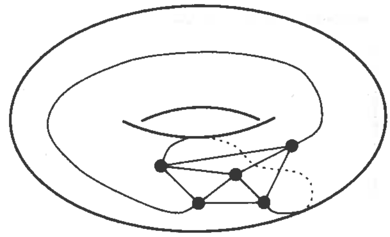

Usually it is easiest to first draw on the square, and then transfer everything to the other picture. When possible, show both pictures.

Task 2.3.1: Draw $K_{3,3}$ on the torus. Do the same for $K_{3,4}$ and $K_{4,4}$. There are several nice representations for $K_{4,4}$.

Task 2.3.2: Draw K_6 and K_7 on the torus.

The graphs K_8 and $K_{4,5}$ cannot be drawn on the torus. This is discussed in the next section.

Task 2.3.3: Draw the Petersen graph on the torus.

Task 2.3.4: Gluing opposite sides of a hexagon produces a torus. Use this representation to give a nice way to draw K_7.

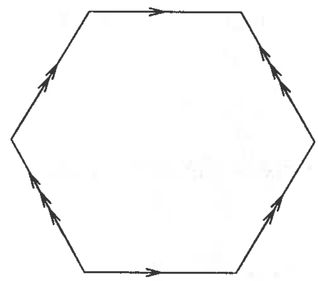

In Section 2.10 we will see why gluing opposite sides of a hexagon gives a torus.

2.4 Euler's formula, again

In the previous chapter we established Euler's formula $v - e + f = 2$ for any graph drawn on the plane. This formula also holds for any graph on the sphere.

Task 2.4.1: Explain why $v - e + f = 2$ holds for any connected graph drawn on the sphere.

Let's investigate $v - e + f$ for the torus. In these examples it is important to keep track of what is being glued together when counting edges and regions.

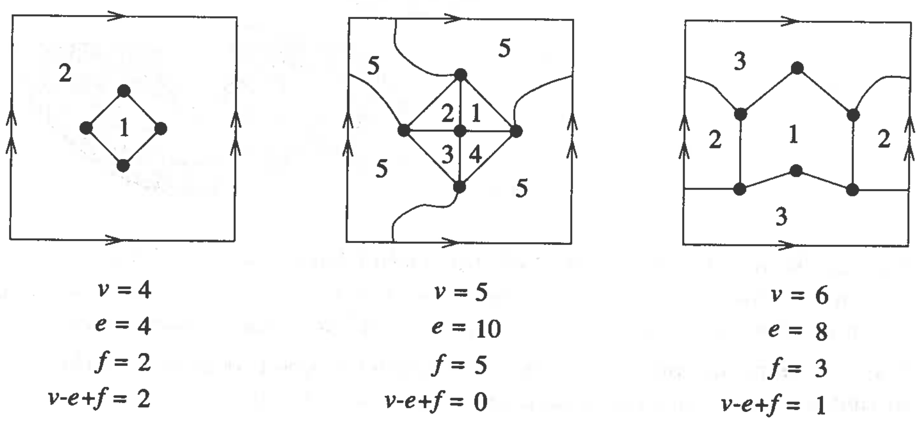

This looks like bad news. The value of $v - e + f$ doesn't appear to always be the same. Fortunately, the discrepancy is just an illusion. The key lies in the separate regions of the graph. To get a workable formula, we must only use graphs whose separate regions are *cells*. A region of a surface is called a **cell** if all loops in that region are contractible.

This is region 2 of the first graph given above. The loop shown is not contractible, so the region is not a cell.

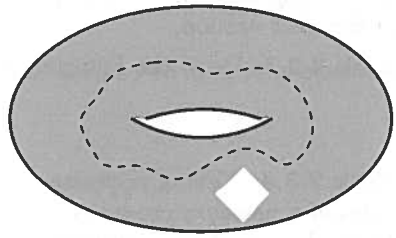

Here are more examples.

These are cells:

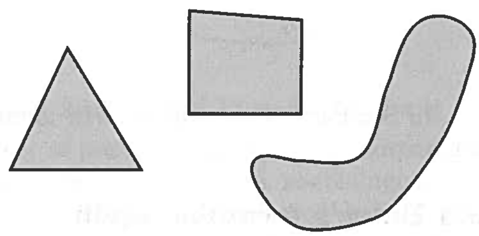

These are **not** cells:

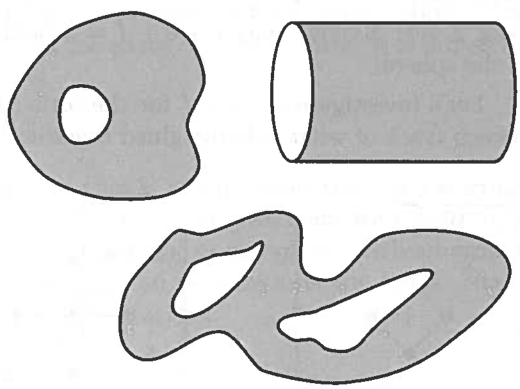

Roughly speaking, a region is a cell if it doesn't have any 'holes' in it.

In the examples shown previously, we found $v - e + f = 0$ for the graph which divides the torus into cells. This is Euler's formula for the torus.

Euler's formula for the torus. If a connected graph is drawn on the torus so that the separate regions are cells, then $v - e + f = 0$.

The next two Tasks suggest methods of seeing why $v - e + f = 0$ is true for the torus.

Task 2.4.2: Refer to Task 1.5.1. The graphs you drew in that Task can all be drawn on the torus, but the large outside region will not be a cell. It takes 2 more edges to turn that region into a cell. Putting in those edges adds 2 to e and C, and leaves all other quantities unchanged. Replace e by $e+2$ and C by $C+2$ in all of your formulas for the sphere, and you will wind up with $v-e+f=0$ for the torus.

Task 2.4.3: A torus can be built from a sphere by cutting out two triangles and gluing the cut edges together:

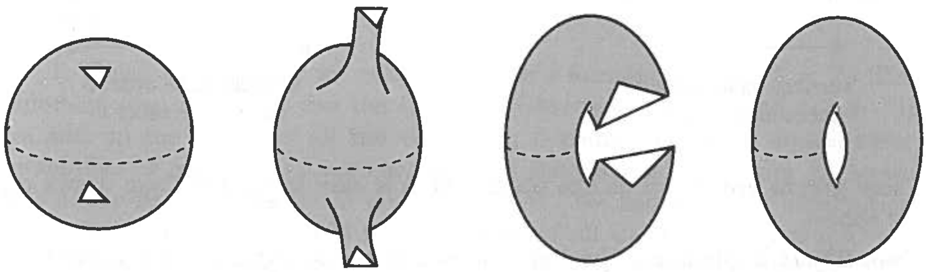

Altogether the triangles had 6 edges, 6 vertices, and 2 faces. The 2 faces were thrown away, and after the gluing, the two triangles became one triangle. So in going from a sphere to a torus we 'lost' 3 vertices, 3 edges, and 2 faces. Check that if you start with Euler's formula for the sphere, replace v by $v-3$, replace e by $e-3$, and replace f by $f-2$, you get Euler's formula for the torus.

Task 2.4.4: What happens if you modify Task 2.4.3 by cutting out squares instead of triangles? Do you still get Euler's formula for the torus?

Task 2.4.5: Use $v-e+f=0$ and $3f \leq 2e$ to show that K_8 cannot be drawn on the torus.

Task 2.4.6: Use $v-e+f=0$ and $4f \leq 2e$ to show that $K_{4,5}$ cannot be drawn on the torus. Can $K_{3,6}$ be drawn on the torus? Either do it or explain why it is impossible.

Task 2.4.7: Suppose that removing one edge from a graph results in a graph which can be drawn on the sphere. Does it follow that the original graph can be drawn on the torus?

In Section 1.10 we discussed maps on the plane (or the sphere), and we mentioned that all planar maps can be 4–colored. One can also look at maps on the torus, and in this case the result is that all maps on the torus can be 7–colored. The proof is described in the Notes at the end of the chapter.

2.5 Regular graphs

In this section we study a special kind of graph.

Definition. A graph diagram is called **regular** if every vertex has the same order and every face is a cell of the same order.

We will often refer to regular graphs, although technically it is only correct to speak of regular graph diagrams. First we study regular graphs on the sphere.

Here are two regular graphs:

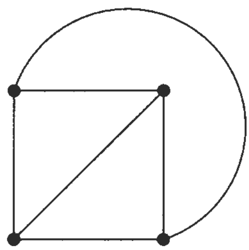

vertices have order 3
faces have order 3

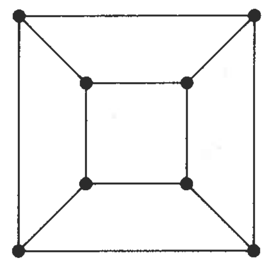

vertices have order 3
faces have order 4

Those graphs are drawn on the plane, but it is easy to picture them drawn on the sphere.

Task 2.5.1: Explain why the dual of a regular graph diagram is also regular.

Our goal is to find all of the regular graphs on the sphere. There is an infinite list of uninteresting ones, three of which are:

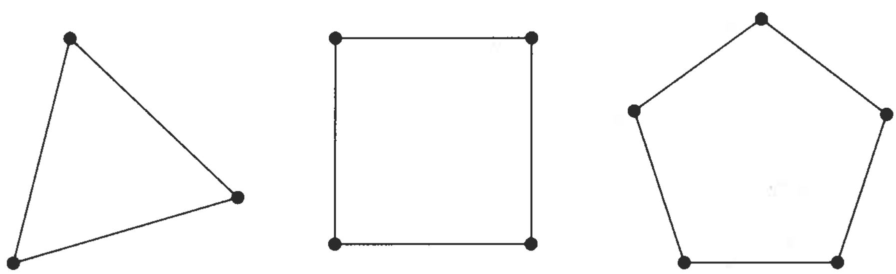

Those graphs are commonly referred to as *cyclic graphs.*

The dual of a cyclic graph has multiple edges, so it is not actually a graph in our strict sense. You can decide for yourself if it should be included on our list of regular graphs.

Task 2.5.2: Find as many regular graphs as possible. Try to find at least three regular graphs in addition to the cyclic graphs and the two at the beginning of this section. For each one that you find, list the number of vertices, edges, and faces, and the orders of the vertices and the faces.

Now that you have a list of regular graphs, we want to check if the list is complete. Our secret weapon, as usual, will be Euler's formula $v - e + f = 2$.

Let M be the order of each vertex in a regular graph, and let J be the order of each face. First we will determine the possibilities for M and J. A little cleverness will save us lots of work. Each face of a graph must have order at least 3, so $J \geq 3$. By Task 2.5.1, the dual of a regular graph is regular. Taking duals switches vertices and faces, so taking duals switches J and M. Therefore, we also have $M \geq 3$.

In Section 1.11 we showed that any planar graph must have a vertex of order less than or equal to 5. Since all the vertices in a regular graph have the same order, we conclude that $M \le 5$. By the same argument as above, we also have $J \le 5$.

We have shown that J and M must both be between 3 and 5. It remains to be determined which possibilities can actually occur. You already listed some of these in Task 2.5.2.

Task 2.5.3: The cyclic graphs have $M = 2$. This contradicts $M \ge 3$. What is the problem?

To finish our analysis we must use Euler's formula $v - e + f = 2$. The important step here is to use the Counting Observation from Section 1.6: "If you add up the orders of all the vertices in a graph, the result equals twice the number of edges." In a regular graph all the vertices have order M, so we have $Mv = 2e$. Dividing by M gives $v = 2e/M$. We had a similar observation concerning faces: "If you add up the orders of all the faces in a graph, the result equals twice the number of edges." In a regular graph, this says $Jf = 2e$. Dividing by J gives $f = 2e/J$. Plugging our new formulas into Euler's formula $v - e + f = 2$ gives

$$\frac{2e}{M} - e + \frac{2e}{J} = 2.$$

For each of the possible values for M and J, we plug into that formula and solve for e. For example, if $J = 3$ and $M = 3$ then

$$\frac{2e}{3} - e + \frac{2e}{3} = 2,$$

which tells us $e = 6$. Sure enough, $J = 3$, $M = 3$, $e = 6$ describes one of the regular graphs we found earlier.

Another possibility is $J = 4$ and $M = 5$. This gives

$$\frac{2e}{5} - e + \frac{2e}{4} = 2,$$

which reduces to $e = -20$. That is nonsense, because the number of edges must be positive. We conclude that there is no regular graph with $J = 4$ and $M = 5$. The next Task finishes our study of regular graphs on the sphere.

Task 2.5.4: For each possible value of J and M, use the formula above to find the value of e. If the value of e is sensible, then find a regular graph with those values of J, M, and e. Note: there are 9 combinations of J and M to be checked. Two of them were already done in the text above.

The regular graphs on the sphere were first studied by the ancient Greeks in terms of the **regular solids**:

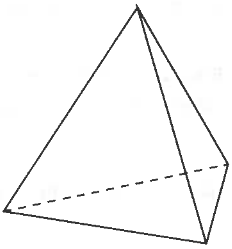

Tetrahedron

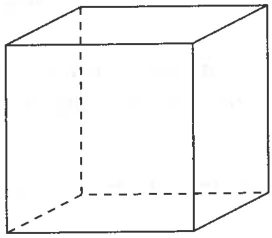

Cube

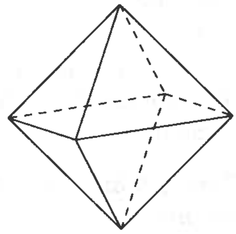

Octahedron

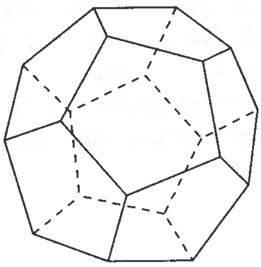

Dodecahedron

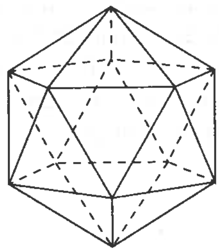

Icosahedron

For more information on the regular solids, see the book *Shapes, Space, and Symmetry* [SSS] by Holden.

Now we turn our attention to the torus. Part of our previous work can be reused: the relations $Mv = 2e$ and $Jf = 2e$ hold for a regular graph on any surface. Plugging into Euler's formula for the torus gives

$$\frac{2e}{M} - e + \frac{2e}{J} = 0.$$

We can factor the left side to get

$$e\left(\frac{2}{M} + \frac{2}{J} - 1\right) = 0.$$

We may assume $e \neq 0$, otherwise our graph would have no edges. So we must have

$$\frac{2}{M} + \frac{2}{J} - 1 = 0.$$

When solving that equation we must keep in mind the meaning of J and M. For example, $M = \frac{14}{3}$ and $J = \frac{7}{2}$ is a solution, but this is a nonsense solution because it is impossible for $\frac{14}{3}$ to be the order of a vertex. We must restrict ourselves to whole number values of J and M.

Task 2.5.5: Check that these are the only positive integer solutions to the above equation:

$$J = 3 \qquad \text{and} \qquad M = 6$$
$$J = 4 \qquad \text{and} \qquad M = 4$$
$$J = 6 \qquad \text{and} \qquad M = 3$$

Task 2.5.6: Draw regular graphs on the torus corresponding to each possibility in Task 2.5.5. You have already done some of these in Section 2.3.

2.6 More surfaces: holes

So far we have studied the sphere and the torus. Now we look at some other surfaces.

One way to get a new surface is to cut holes in a surface you already have.

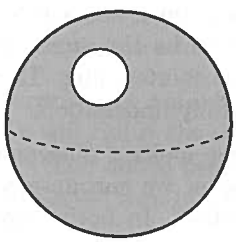

A sphere with one hole

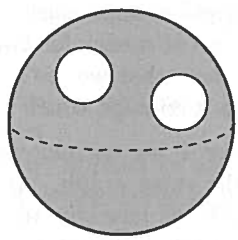

A sphere with two holes

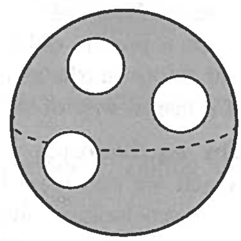

A sphere with three holes

Surfaces can appear in disguised form. Each of these surfaces is 'the same' as the corresponding sphere with holes shown above:

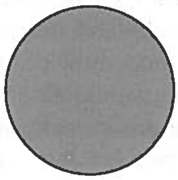

A disk

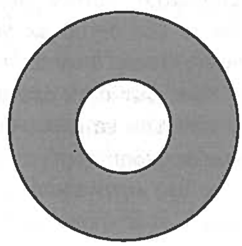

An annulus

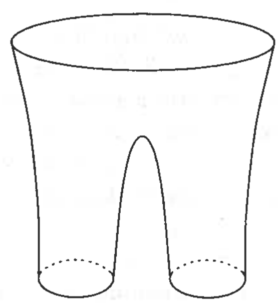

A pair of pants

In a later section we will discuss what it means for two surfaces to be 'the same.' For now, it is sufficient to think of two surfaces as 'the same' if one can be smoothly deformed to give the other. For example, an annulus is 'the same' as a cylinder:

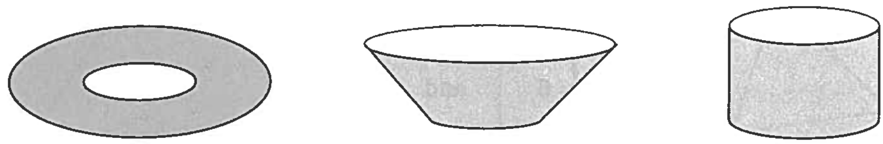

Task 2.6.1: Convince yourself that a disk is 'the same' as a sphere with one hole, an annulus is 'the same' as a sphere with two holes, and a pair of pants is 'the same' as a sphere with three holes.

Task 2.6.2: What common article of clothing is 'the same' as a sphere with four holes?

On 'quotes'. In the last page we have used the phrase 'the same' several times. Each time we were careful to put 'quote marks' around it. The quote marks are used to show that we are using words imprecisely. Specifically, since we have not given a precise definition of what it means for two surfaces to be *the same*, we put quote marks around it to show that we are aware of this shortcoming. This is a useful way of working with a concept which is not yet fully understood.

On equivalence. In Chapter 1 we encountered different–looking diagrams which we considered to be 'the same graph.' In this chapter we encountered different–looking surfaces which we consider to be 'the same.' In both cases there was the underlying idea that, although we had two things which weren't 'really' the same, we were going to think of them as 'the same' for our current purposes. This situation occurs so frequently that mathematicians have a special word to describe it. Instead of using saying 'the same,' we use the word **equivalent**. An example sentence is, "An annulus is equivalent to a sphere with two holes." Now, we still haven't given a precise definition of what *equivalent* means in this context, but at least we no longer have to use those annoying quote marks.

In everyday life the words "the same" rarely means "absolutely exactly the same." Take the phrase, "my car is the same as your car." That statement is perfectly sensible, and we know it doesn't mean that we both own the very same automobile! It could mean that our cars are the same make and model, or maybe it means that our cars are the same make, model, year, and color. If you heard that phrase in a conversation, you would be able to understand what it meant. In the same way, we use *equivalent* to mean "the same, as far as our current purposes are concerned." It is everybody's job to keep track of the current meaning of *equivalent*.

2.7 More surfaces: connected sums

We created new surfaces by cutting holes in a surface we already had. Another way to get a new surface is to 'glue together' two surfaces. The official name for this is **taking the connected sum** of the two surfaces. Here is an example where we take the connected sum of two tori:

Remove a disk from each surface.

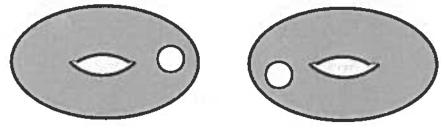

Glue the surfaces together along the cut edges.

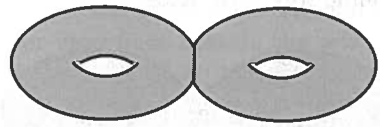

Smooth out the seam.

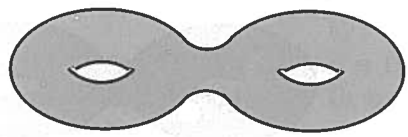

Surprisingly, nobody has invented a good name for the surface we just created. We will call it the **double torus**.

The above procedure can be applied to any pair of surfaces. Taking the connected sum of a torus and a double torus produces this surface:

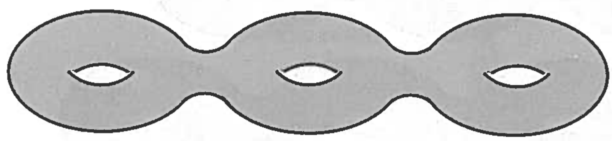

We will call this the **triple torus**.

Task 2.7.1: Draw a sequence of pictures to show that if you take the connected sum of a surface with a sphere, the result is equivalent to the original surface.

Another way to prove the fact mentioned in Task 2.7.1 is: A sphere minus a disk is the same as a disk, so taking the connected sum with a sphere just replaces the disk which was removed from the original surface.

Combining our two methods of producing surfaces gives this catalog:
Spheres with holes:

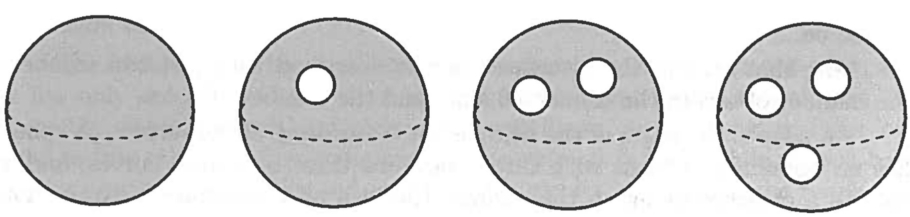

Tori with holes:

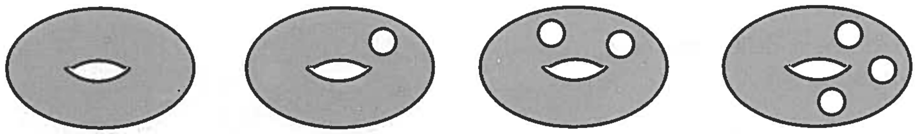

Double tori with holes:

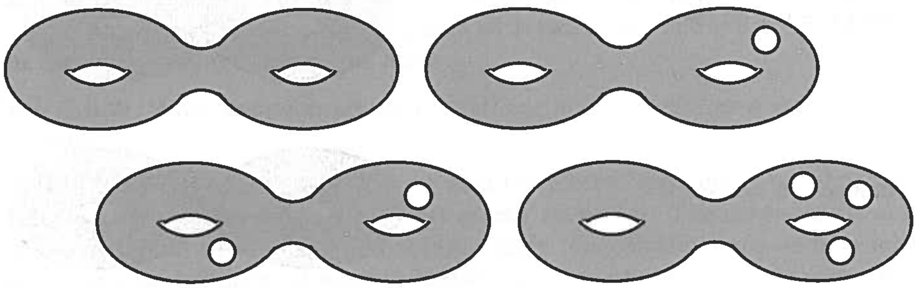

Triple tori with holes:

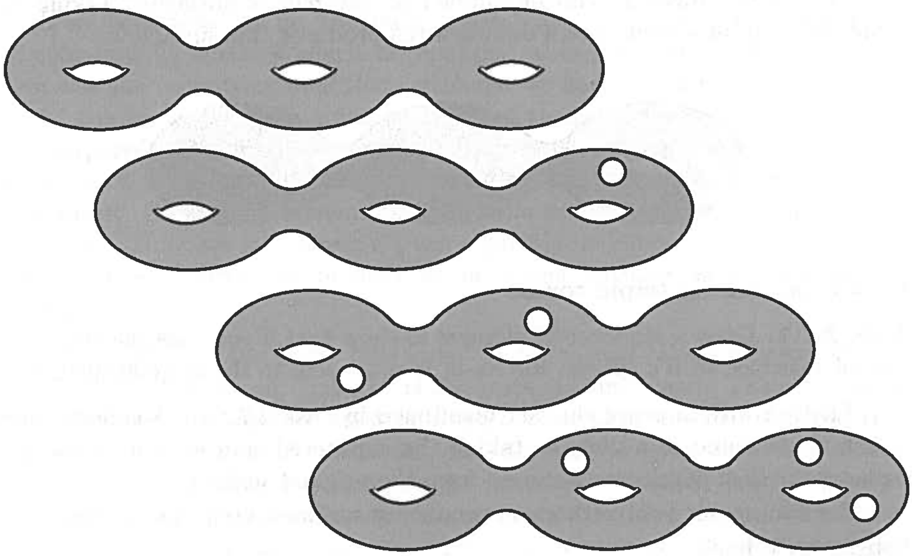

and so on....

Note that each of these surfaces can be described with just two numbers: the number of tori in the connected sum, and the number of holes.

We refer to the edges of the holes as the **boundary** of the surface. A sphere has no boundary, a torus with three holes has three boundary curves, and so on. To a creature living on the surface, the idea of a boundary curve is more natural than the idea of a hole. As the inhabitants move about on the surface they occasionally encounter an 'edge' to their world. From the outside we can see that they have walked up to the one of the holes, but on the surface it merely looks like the world ends. An inhabitant can walk along the edge until it gets back to where it started. To the creature, it just seems like there is a line on the ground which marks the boundary of the world. If the surface has more than one 'hole' then the inhabitant can walk along the surface to another boundary curve and walk around it until it gets back to where it started. By making marks on the ground it can count how many boundaries the world has. We can see that the creature is just counting holes, but our view from outside the world is unnatural. The natural perspective is that of the creature on the surface, and so boundary curves, not holes, are the preferred object of study. We will look at this further in a later section.

2.8 One–sided surfaces

Our list of surfaces is not complete. One surface we are missing is the *Möbius strip*, named after the German mathematician August Möbius. To make a Möbius strip, glue the ends of a strip of paper with a half–twist.

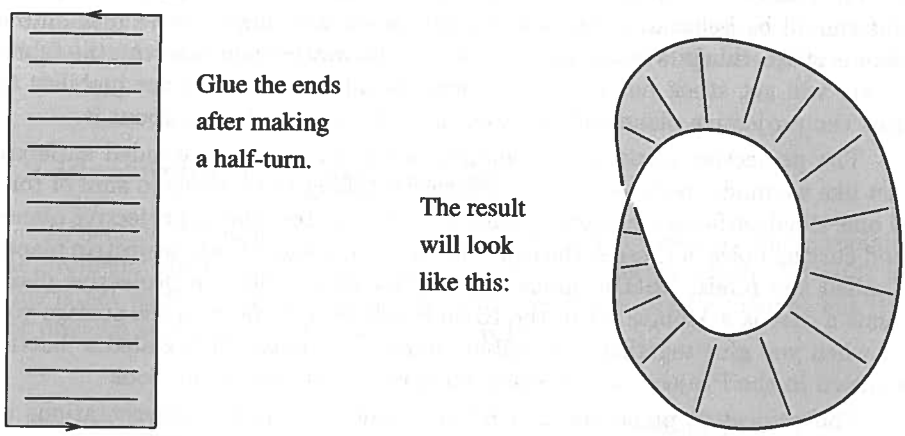

Task 2.8.1: Imagine that you take a Möbius strip made of paper and you cut it in half down the center. What will be the result? What if you cut the resulting surface in half again? What if you cut the original Möbius strip into thirds? Make detailed predictions as to the result of each operation.

Task 2.8.2: Make some Möbius strips out of paper and check your predictions from the previous Task. If your predictions were wrong, try to figure out where you made an error.

The Möbius strip has two interesting properties: it has just one edge, and it has only one side. Check those properties with your paper model: you can trace the whole edge without picking up your finger, and you can go from 'one side' to 'the other side' by going around the strip. It is worthwhile to look back at the previous two Tasks and express your results in terms of the number of edges and sides of the surfaces.

Since the Möbius strip has only one side, it cannot be in the large catalog of surfaces we produced previously. If you look back on that list, you can easily convince yourself that all those surfaces have two sides.

A Möbius strip has one curve for its boundary. A disk also has one boundary curve. We can imagine gluing a disk onto a Möbius strip, producing a surface with no boundary. The resulting surface is called the **projective plane**.

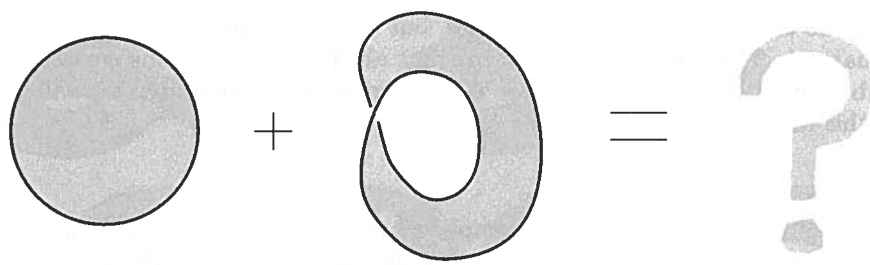

We put a fuzzy ? because it is impossible to accurately draw a projective plane. This should be believable, because you can imagine trying to sew a disk onto a Möbius strip: things will be easy at first, but no matter how stretchy the fabric is, you will get stuck before you are able to finish. However, our inability to draw the projective plane will not stop us from learning things about it.

The projective plane is the building block for all the one–sided surfaces. Just like we made more two–sided surfaces by taking the connected sum of tori, all one–sided surfaces are made by taking the connected sum of projective planes (and cutting holes in those surfaces). The connected sum of two projective planes is called the **Klein bottle**, named after Felix Klein. Since a projective plane minus a disk is a Möbius strip, the Klein bottle is equivalent to the surface you get when you glue together two Möbius strips. The reason it is called a 'bottle' is shown in the Project on one–sided surfaces, at the end of the book.

The projective plane and the Klein bottle have useful representations as squares with opposite sides glued:

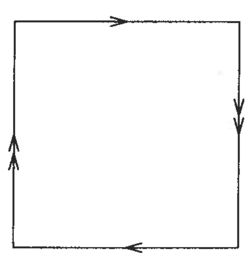

The projective plane

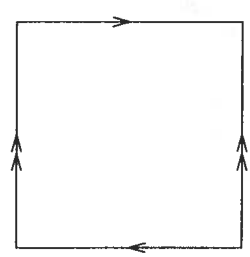

The Klein bottle

Note that the arrows do not all face the same way, sometimes we 'make a twist' when gluing. These representations are useful for drawing graphs on the surfaces:

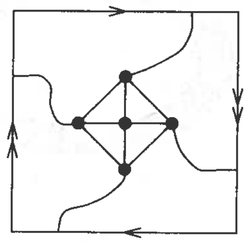

K_5 on the projective plane

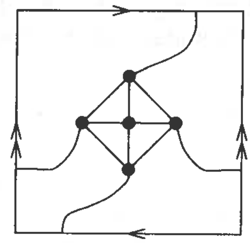

K_5 on the Klein bottle

It is important that edges of the graph match up properly after we glue the sides of the square. The presence of a 'twist' makes things a bit more tricky than when we were dealing with the torus. Check that the above examples are drawn correctly.

This ends our discussion of one–sided surfaces. To learn more about these interesting surfaces see the Project on one–sided surfaces at the end of the book. In the next section we return to two–sided surfaces.

2.9 Identifying two–sided surfaces

In Section 2.7 we gave a complete list of two sided surfaces. There remains the problem: given a two–sided surface, how do we decide which surface it is? This is not as trivial as it sounds. The surfaces below are two–sided, but it is not obvious where each fits on our previous list.

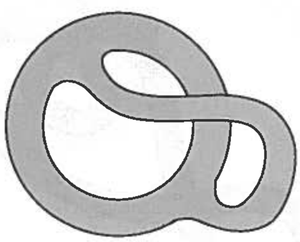

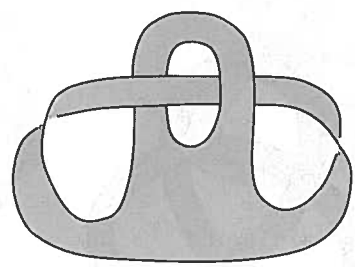

Before identifying these surfaces in terms of our previous list, we must develop our concept of what it means for two surfaces to be equivalent. Recall that we take the perspective of someone living on the surface. If you are confined to the surface then you can't distinguish many things which are evident from the 'outside.' For example, suppose you live on an annulus. You would be oblivious to the following procedure:

Choose any line going from one boundary to the other.

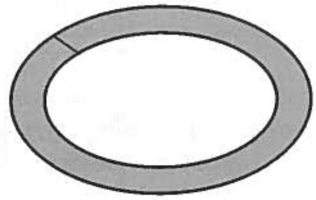

Cut along the line and twist the surface in various ways.

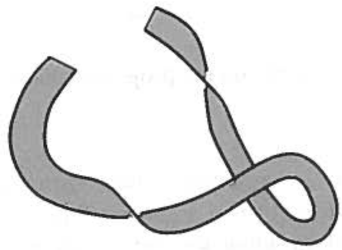

Glue the cut edges together the way they were originally.

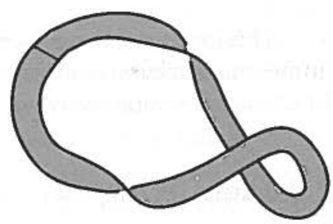

The cut edge was put back exactly the way it started, so if you are confined to the surface it is impossible to distinguish between the original annulus and the twisted annulus.

Task 2.9.1: Show that each of these is equivalent to an annulus.

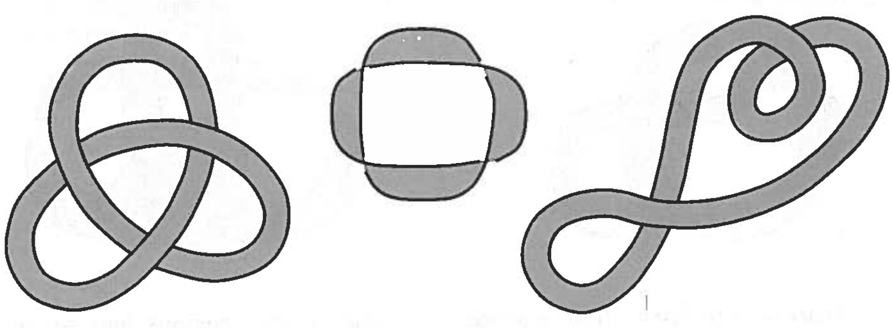

Also, the figure on the cover of this book is equivalent to an annulus.

It is difficult to write down a useful and completely rigorous definition of what it means for two surfaces to be equivalent. For our purposes it is sufficient to think of it as "indistinguishable by people who are confined to the surface, have poor eyesight, and are bad at measuring distances." Two important ways of manipulating a surface to get an equivalent surface are: bend and stretch without tearing it; and cut it apart and twist it around, reassembling the pieces so they fit together the same as originally.

On topology. The above discussion of a mystery world inhabited by people who have poor eyesight and are bad at measuring distances might remind you of the insect world from Chapter 1. This is a main theme of this book: we ignore size and distance, and just look at how things are connected. For example, when we studied graphs the only thing we cared about was which vertices were connected to which other vertices. And as we study surfaces, we only care how the surface is 'connected to itself,' the outside appearance of the surface being of only minor consequence. Mathematicians call this area of mathematics **topology**. Note: do not confuse topology, a branch of mathematics, with topography, the study of mapmaking.

Now we develop an organized way to distinguish surfaces. The surfaces on our list can each be described as the connected sum of some number of tori, each with some number of holes. Counting holes is the easier part. Each hole has one boundary curve, so we count holes by counting boundary curves. To do this, start at one point on the edge of the surface, and trace that edge until you end up back where you started. If there is a part of the boundary which hasn't been traced, then start over on an untraced part. Keep going until you trace the entire edge. The number of times you started tracing is equal to the number of boundary curves. An easy way to keep track is to trace each boundary curve with a different color pen.

Task 2.9.2: Check that these are labeled correctly:

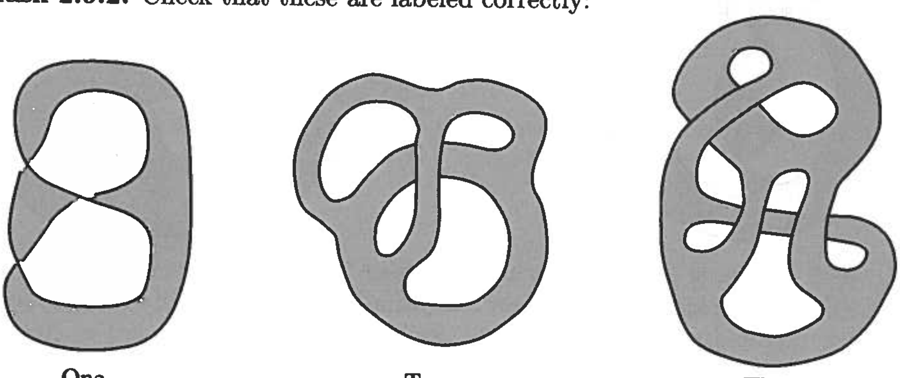

One boundary curve | Two boundary curves | Three boundary curves

Task 2.9.3: Count the number of boundary curves on the two surfaces at the beginning of this section.

If a two–sided surface has no boundary then it is a sphere, or a torus, or a double–torus, or.... If a surface has one boundary curve then it is a *something* with one hole. The problem is to determine the *something*. The key to this is our friend $v - e + f$. In the case of the sphere we found $v - e + f = 2$ and for the torus we found $v - e + f = 0$, so we can use $v - e + f$ to distinguish between the sphere and the torus. This suggests that it will be useful to find $v - e + f$ for our other surfaces.

Since f stands for the number of separate regions...

Task 2.9.4: Explain why cutting a hole in a surface decreases $v - e + f$ by 1.

So, a sphere with one hole has $v - e + f = 1$, a sphere with 5 holes has $v - e + f = -3$, a torus with 2 holes has $v - e + f = -2$, and so on.

If we knew $v - e + f$ for the double torus, triple torus, and so on, then we could determine $v - e + f$ for all the 2–sided surfaces.

Task 2.9.5: Explain why taking the connected sum of a surface with a torus decreases $v - e + f$ by 2. So, for the double torus $v - e + f = -2$, for the triple torus $v - e + f = -4$, etc. The reasoning in Task 2.4.3 may be helpful here.

Task 2.9.6: Make a chart of all the 2–sided surfaces and write the value of $v - e + f$ next to each one.

Now we have an infallible way to identify any 2–sided surface: first count the number of boundary curves, then find $v - e + f$ and look up the answer in your chart from Task 2.9.6. To find $v - e + f$, draw a graph on the surface which divides the surface into cells. The usual method is to decide how many edges need to cut across the surface to divide it into cells, put vertices at the ends of those edges, and connect the vertices by putting graph edges along the entire boundary of the surface. Here are two examples:

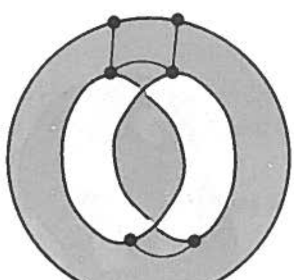

$v = 6$

$e = 11$

$f = 3$

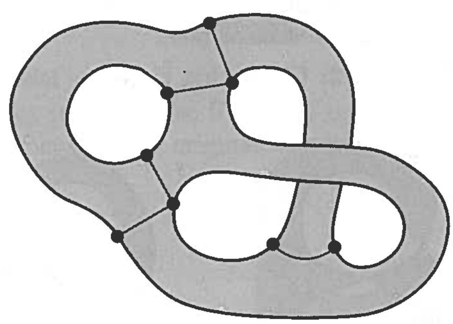

$v = 8$

$e = 13$

$f = 3$

The first surface has three boundary curves and $v - e + f = -1$, so it is a sphere with three holes. The second has two boundary curves and $v - e + f = -2$, so it is a torus with two holes. In both examples we used more vertices and edges than were absolutely necessary to cut the surface into cells. Adding too many vertices and edges makes it difficult to correctly count v, e, and f, but if you fail to cut the surface into cells then you won't get the right answer.

Task 2.9.7: Identify all the surfaces which have appeared in this section.

2.10 Cell complexes

This section finishes our study of two–sided surfaces.

Below are assembly instructions for building a surface. Edges labeled with the same number get glued together, and the arrows show which way to match the edges.

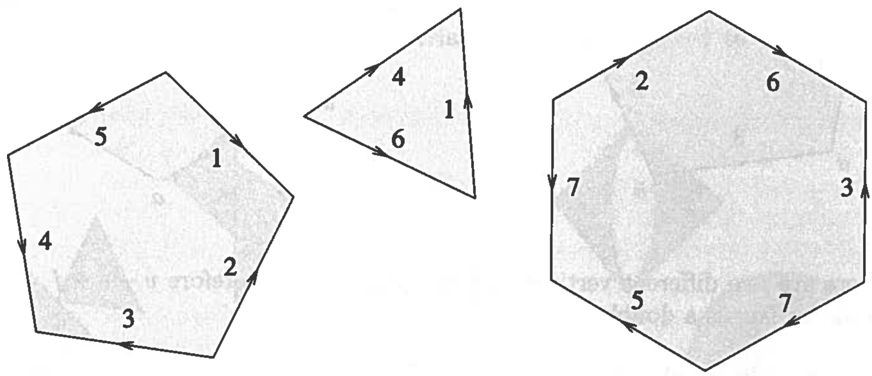

This is an example of a **cell complex**: a bunch of cells which are glued together to make a surface. We first encountered this idea back when we began studying the flat torus. Gluing the sides of a square to get a torus is an example of a cell complex with just one cell.

Task 2.10.1: Is it possible to make the above surface out of actual pieces of paper?

We use our usual method to determine what surface is formed from the cell complex. First we count boundary curves. In this example, every edge gets glued, so the surface has no boundary, so the surface is a *something* with no holes. Next we find $v - e + f$. The glue lines form the edges of a graph, so we will use that graph to find v, e, and f. The cell complex has 3 cells, so $f = 3$. There are 7 glue lines, labeled 1 through 7, so $e = 7$. The corners of the cells will become vertices, but many corners will be glued together at the same vertex. When a pair of edges get glued, that also glues together the corners at the respective ends of the edges. To count vertices, label a corner and then trace through all the glued edges and give the same label to all the corners which get glued to the original corner. If some corner hasn't been labeled, then repeat the process. The number of labels you need is equal to the number of vertices. Here is the result for the above example:

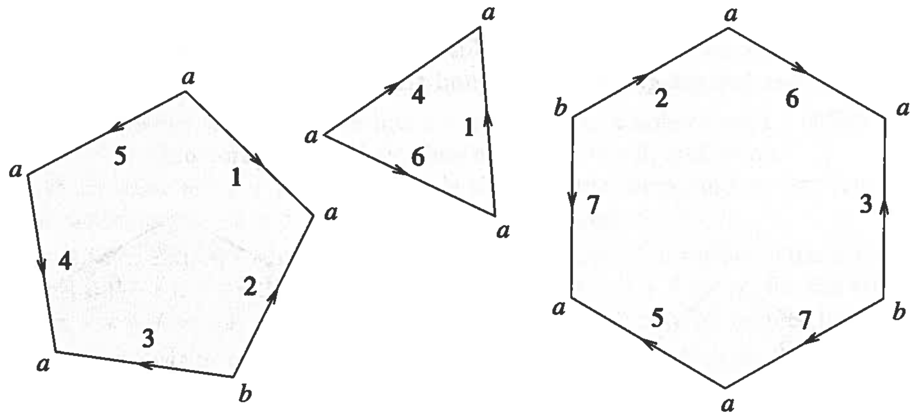

There are two different vertices: a and b, so $v = 2$. Therefore $v - e + f = -2$, so the surface is a double torus.

Task 2.10.2: Identify this surface:

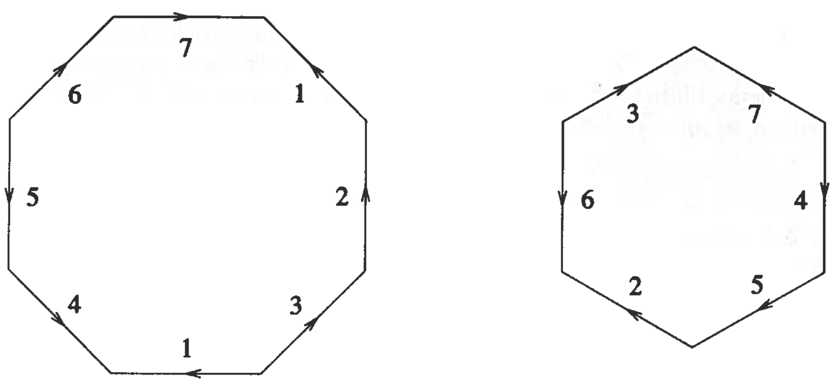

Task 2.10.3: In Task 2.3.4 we stated that gluing opposite sides of a hexagon gives a torus. Verify that this is true.

Task 2.10.4: What surface do you get when you glue opposite sides of an octagon? decagon? dodecagon? (Those polygons have 8, 10, and 12 sides, respectively.)

We have repeatedly stated that our list of two–sided surfaces is complete. The easiest proof of this fact uses cell complexes. This is briefly discussed in the Notes at the end of the Chapter.

One type of cell complex is of particular importance:

Definition. A **triangulation** is a cell complex where:

- All the cells are triangles.
- Two different triangles in the cell complex share either exactly one vertex, or exactly one edge (and the vertices at the ends of that edge), or else they don't meet at all.
- No triangle shares a vertex or edge with itself.

The triangles can look bent or twisted, but they must still be cells whose border has three vertices and three edges.

Here are some of the ways that triangles are *not* allowed to border in a triangulation:

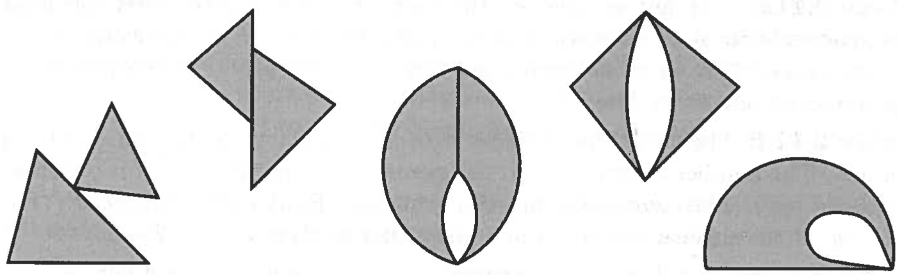

Task 2.10.5: Explain why the number of triangles in a triangulation must be even. Hint: this is related to Task 1.6.5.

The answers to the next Task have already appeared in this Chapter.

Task 2.10.6: Find a triangulation of the sphere using 4 triangles. Find a triangulation of the torus using 14 triangles.

It is impossible to triangulate the sphere with fewer than 4 triangles. We must use an even number of triangles, so the only possibility would be to use 2 triangles. You can easily make a sphere by gluing together two triangles, but the result is not a triangulation.

It is also impossible to triangulate the torus with fewer than 14 triangles. You can show this by combining $v - e + f = 0$, the idea behind Task 2.10.5, and the observation that a graph with n vertices has at most as many edges as K_n. There is a simple way to triangulate the torus with 18 triangles: cut the flat torus into 9 squares, then divide each square into two triangles.

2.11 Notes

Note 2.11.a: The statement "The Earth is a sphere" is inaccurate in several ways. It is more precise to say "The surface of the Earth is a sphere." That statement is also inaccurate because certain geographic features, such as arches and tunnels, cause the surface of the Earth to be quite complicated. It would be impossible to determine exactly which two–sided surface the surface of the Earth is, because the answer depends on how fine of a scale you use to measure.

If you measure things on a very large scale then the surface of the Earth is a good approximation to a sphere. Outside of pure mathematics, that is the best you could hope for.

Note 2.11.b: The plural of *torus* is **tori**.

Note 2.11.c: Although the word *equivalent* will have different meanings in different contexts, in each case we require the following:

- Everything is equivalent to itself.
- If A is equivalent to B, then B is equivalent to A.
- If A is equivalent to B, and B is equivalent to C, then A is equivalent to C.

Note 2.11.d: It is traditional to use the letter F to represent a surface. This is because the German word for "surface" is "fläche."

Note 2.11.e: For any surface F, the number $v - e + f$ is called the **Euler characteristic** of the surface, and it is denoted by $\chi(F)$. For example, if T is the torus, then we would write $\chi(T) = 0$. Note: χ is the Greek letter chi, pronounced 'kī,' as in 'kite.'

Note 2.11.f: The two–sided surfaces with no boundary are all connected sums of tori. The number of tori is called the **genus** of the surface, and it is commonly denoted by g. The connection between genus and Euler characteristic is $\chi(F) = 2 - 2g$. If the surface also has b boundary curves, then $\chi(F) = 2 - 2g - b$.

Note 2.11.g: The distinction between two–sided and one–sided surfaces may seem natural, but this actually reflects our prejudice as 3–dimensional beings. First, you should not think of the inhabitants of a surface as people walking on top it, but rather as 2–dimensional beings embedded in it. These beings do not think in terms of 'sides' of the surface. They can only conceive of two directions of motion, so the possibility of a third direction of motion towards the 'top' is beyond their comprehension. Fortunately, they can still understand our distinction between one– and two–sided surfaces by using the related concept of orientability. We will see that the two–sided surfaces are orientable and the one–sided surfaces are nonorientable.

It is easiest to explain by an example. The following illustrates that the Möbius strip is nonorientable. Keep in mind that the objects 'on' the surface are actually 2–dimensional objects 'in' the surface.

Take a right–hand glove and slide it around a Möbius strip.

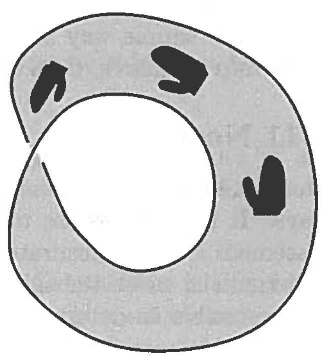

Keep sliding until it returns to its original place. It is now a left–hand glove!

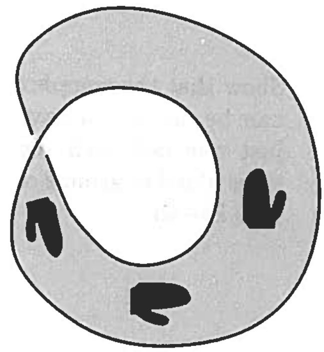

A surface is **nonorientable** if it is possible to turn any shape into its mirror–image by moving it along some path in the surface. Such a path is called an **orientation reversing path**. A surface is **orientable** if it isn't nonorientable. You should convince yourself that the two–sided surfaces are orientable, meaning that it is impossible to turn a right–hand glove into a left–hand glove by moving around on those surfaces.

An alternate definition of orientability uses cell complexes. To **orient a cell** means to specify a direction of travel around its boundary.

An oriented cell.

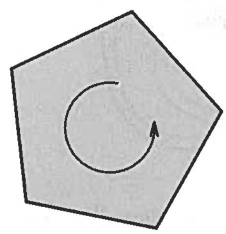

A cell complex is **orientable** if we can orient each cell, such that at the edge where two cells meet, the orientations from each of the two cells are in opposite directions. We say that a surface is orientable if it is given by a cell complex that is orientable. This is equivalent to the previous definition of orientable. Note: it is best to think of an orientation of a cell not as a cyclic ordering of its edges, but rather as a cyclic ordering of its vertices. In two dimensions the distinction is irrelevant, but in higher dimensions the latter view is better.

Note 2.11.h: Here is one way to prove that our list of two–sided surfaces is complete:

- Show that all surfaces can be written as a cell complex.
- Show that the cell complex can be chosen to have just one cell.
- Show that the cell complex can be chosen to have just one cell and one vertex.

- Show that the complex can be chosen to have just one cell with the sides glued in groups of four like so:

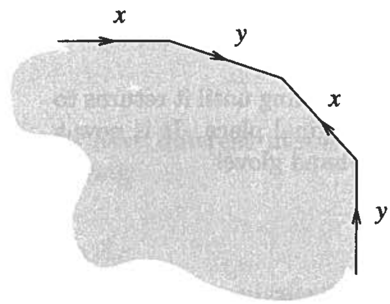

- Finally, conclude that the cell complex gives one of the surfaces on our list.

The representation of a surface as one cell with edges glued in groups of four (as above) is called the **standard form** for the surface.

Standard form of the double torus:

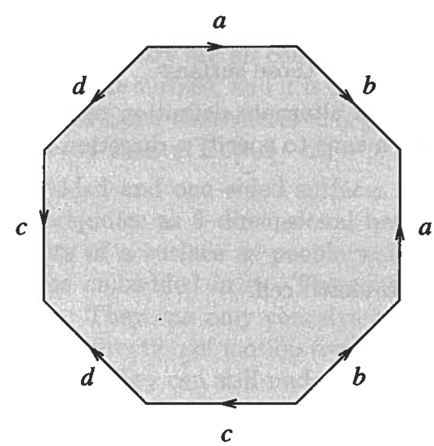

Note 2.11.i: On the plane, or the sphere, any map can be 4–colored, and some maps cannot be done with fewer than 4 colors. The proof of this is very difficult. On the torus, any map can be 7–colored, and some maps require 7 colors. This can be proven by an easy modification of the 6–coloring method from Chapter 1. Here are the steps:

- Use $v - e + f = 0$ to show that any graph on the torus has a vertex of order 6 or less.
- Modify the 6–coloring method of Chapter 1 to show that any graph on the torus can be 7–colored.
- Give an example of a graph on the torus which requires 7 colors.

The first step requires a small bit of calculation, the second is nearly identical to the planar method of Chapter 1, and the third has already been done in Task 2.3.2.

It is interesting that the Four–Color Theorem for the sphere is very difficult, while the Seven–Color Theorem for the torus is very easy.