3. SURGERY ON SURFACES

§3.1. Atlases and Surfaces

We've discussed a number of specific examples of surfaces so our next task is to decide on a suitable definition of 'surface'.

The most important requirement is that a surface should be locally like a disk. That is, if we lived on any surface, be it a sphere or a torus or even a Klein Bottle, our immediate surroundings should be exactly the same as if we were living on a disk. What distinguishes one surface from another is not the local environment but the global structure, arising from the way these disk-like local environments connect together.

So we'll start by insisting that in order for a topological space to be called a 'surface' it must have the property that every point P has a neighbourhood that is homeomorphic to either an open disk, with P at the centre, or the interior of a semicircle together with the diameter, with P on the diameter.

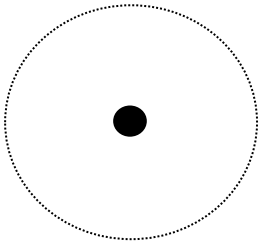

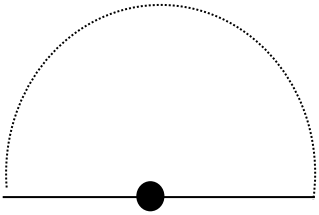

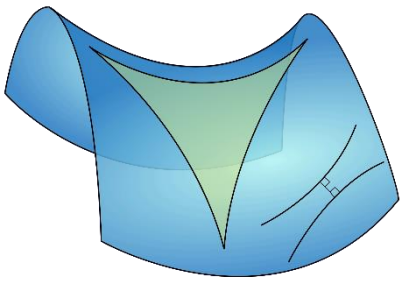

With this definition we allow surfaces to have boundaries – edges with the surface on one side only. The standard definition of a surface doesn't allow this – it insists that every point should have a neighbourhood that's homeomorphic to an open disk so that at every point you can walk right around the point. But this would rule out cylinders and Möbius Bands, which would be a pity.

We'd like to require more, before a space can be called a surface. For a start we'd like surfaces to be *pathwise connected*. The disjoint union of a sphere and a torus will certainly satisfy our first requirement but we'd rather call it a union of two surfaces than a single surface. We'd know everything about such a disjoint union if we knew everything about its connected components, so it's purely for convenience that we limit ourselves to connected spaces.

Finally we'll insist that our surfaces be *compact*. This means that the surface is the union of some finite collection of neighbourhoods. Loosely speaking this means that we can cut the surface into a finite number of pieces, each of which is homeomorphic to a disk.

If we actually cut some paper with scissors it's not clear what happens to the edges along which we cut. For

the purpose of this discussion we'll suppose that a cut edge splits into two with one copy going with each piece and with the edges being identified, that is with both edges being considered the same. With the freedom that topology allows us we can consider the edges to be straight and so the pieces can be polygons.

We can think of these polygons with identified edges as being maps of parts of the whole space. The collection of all these maps could be thought of as an *atlas*. Now an atlas of the world is a collection of maps, each describing a portion of the earth's surface. We know that the world is topologically a sphere but each map is a rectangle (topologically a disk). To sail around the world you'd need to know how these maps fit together.

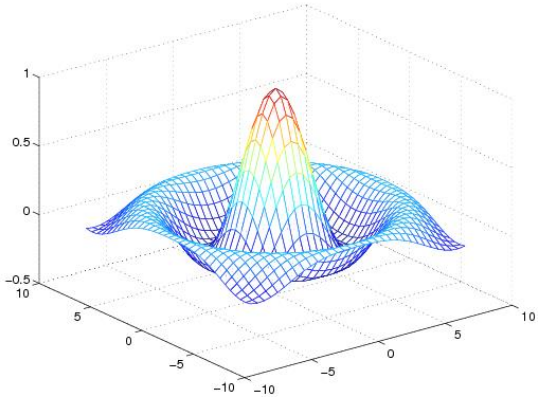

This discussion is rather less rigorous than we would like it to be. Consider it merely as a way of motivating our formal and precise definition of a surface.

A **surface** is any pathwise connected topological space that is homeomorphic to a finite collection of polygons with identified edges.

In a road atlas the edges of each map have the words “MAP XXX adjoins” and so we know how they join up. Some maps have along one or two edges the words “LIMIT OF MAPS”, indicating a boundary. Once this information is provided on individual maps, these maps can be pulled apart and assembled into a book in any order. There’s usually a map of the whole area, showing where each map sits in relation to the others, but this is purely for convenience. The information at the edges of the maps is all that’s really needed.

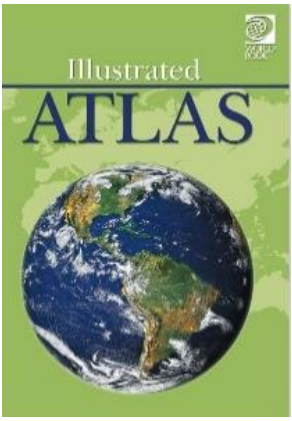

Example 1: The following is an atlas consisting of four maps.

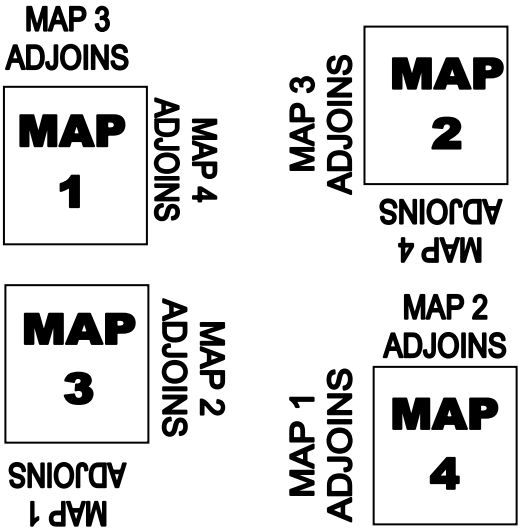

These maps can be joined together to make one large map representing a disk.

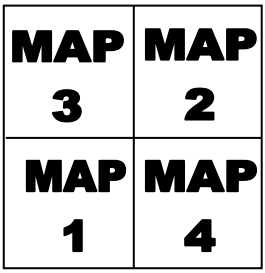

Example 2: The following atlas has four maps.

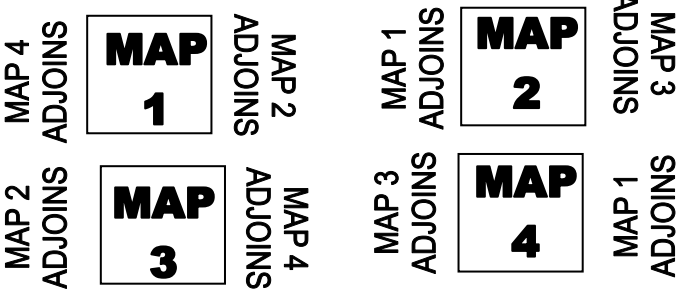

We can join them into a single map:

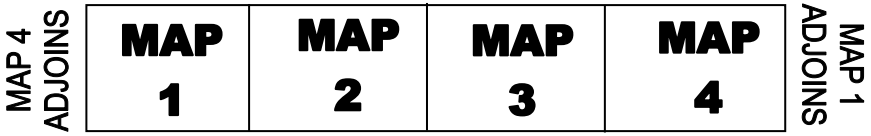

But there's a pair of edges that are yet to be joined. Joining them we get a cylinder.

It should be clear from these two examples that, if we have a finite number of maps in such an atlas, and if the space is pathwise connected, we can join the maps to form a single map. This map will be a single polygon, but perhaps some of its edges remain to be joined. Rather than actually join them we simply *identify* them, that is, mark them in such a way that they are treated as if they were the same edge. What results is a **Polygon with Identified Edges**, or simply as a **PIE**.

Instead of numbering the maps and indicating which maps adjoin we'll follow the practice that we introduced in the previous chapter – that of labelling corresponding edges and using arrows to indicate the appropriate orientation for the edges to be joined.

If there were infinitely many maps in an atlas, things could get very complicated so we'll restrict our attention to what are usually called **compact surfaces** – those which can be described by an atlas with *finitely* many maps. In fact we'll build the finiteness into the definition of an atlas.

An **atlas** is a finite set of disjoint polygons, in which certain edges (possibly none at all) are labelled and given a direction. Each such label must occur exactly twice among the edges of the maps in the atlas. Any remaining edges are called **unidentified edges**.

Every atlas represents a topological space where neighbourhoods of points on identified edges take in

points that are close to the corresponding point on the other member of the identified pair. (Neighbourhoods of all other points are just as they would be in a normal polygon.)

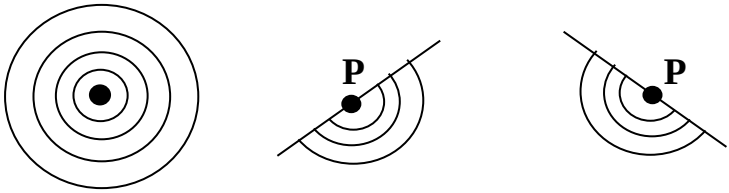

An atlas with just one polygon is called a **polygon-with-identified-edges** (**PIE**). Such a PIE is clearly pathwise-connected and, if we have a pathwise-connected atlas, we can join pairs of identified edges to convert it to a PIE.

Example 3: The following atlas represents the disjoint union of a torus and a cylinder.

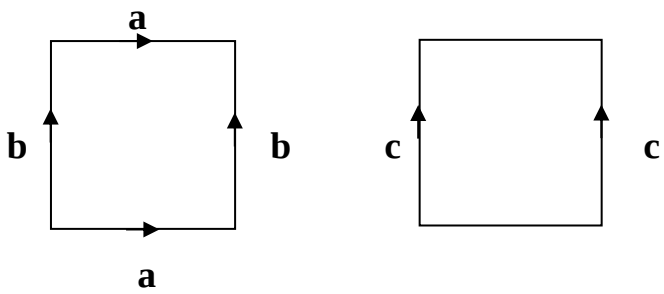

Since it isn't pathwise-connected there's no way we can convert it to a PIE.

We are now ready to give our definition of a surface.

A **surface** is any topological space that is homeomorphic to a polygon with identified edges.

You need to be aware that this definition differs from the standard one you might find in some other texts in two ways:

(1) Our definition includes only compact surfaces, since atlases are finite.

(2) We include the possibility of boundaries.

Every atlas therefore represents the disjoint union of one or more surfaces.

§3.2. Surgery

Surgery on an atlas is the process of cutting and re-stitching. When we cut along an edge to separate a map into two separate maps we must identify the edges of the cut. We do this by placing a label along the cut edge in each map, representing the direction of the cut by arrows. The actual direction of the cut doesn't matter. We could reverse the direction of one arrow, as long as we reversed the direction of the arrow with the same label. What's important is that these new arrows in the two maps

must reflect the relative orientation of the identified edges. When we join two maps to form one, by bringing together the two members of an identified pair with the appropriate orientation, we can then remove that edge completely as it no longer serves any purpose. (When we say "remove" we don't mean that we eliminate that edge. It's still there. It's just that it no longer needs to be given a label.)

Surgery is a way of transforming one atlas into another. What's important is that the process doesn't change the topological nature of the surface. The old atlas is homeomorphic to the new one. But it may be simpler and that may make it easier to recognise.

Example 4: Here is a mystery atlas (with just one map). Do you recognise it?

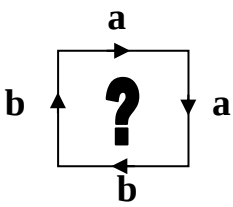

By carrying out the following surgery we recognise it as a Klein Bottle.

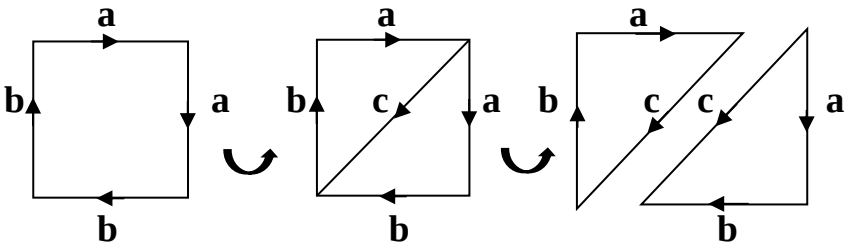

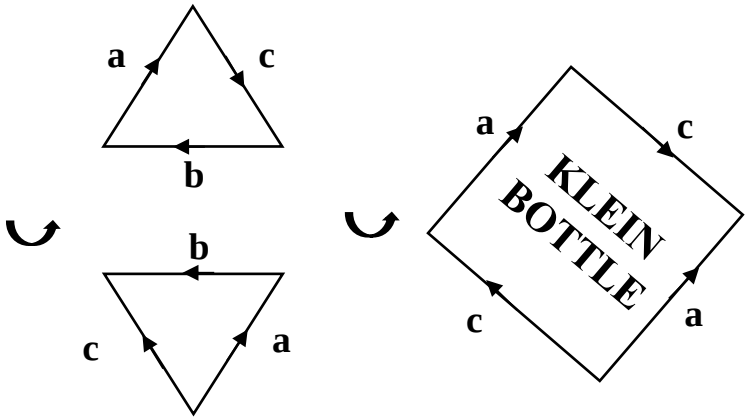

§3.3. The Sum of Two Connected Surfaces

The **sum**, S + T, of two surfaces S and T, is the surface obtained from them in the following way. Cut a hole in each of S and T and join (identify) the resulting boundaries.

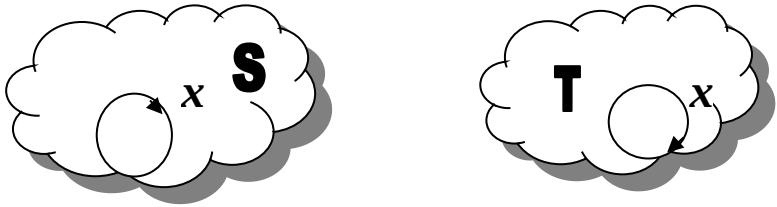

Sometimes, to make a clearer picture, we join the holes by a cylinder. But it amounts to the same thing because the surface of the cylindrical tube can be drawn back into the surfaces themselves by suitable homeomorphisms.

Example 5: The following is the sum of two toruses:

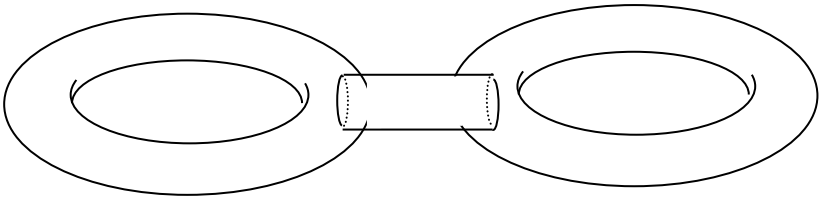

which is homeomorphic to what could be called “Siamese toruses”:

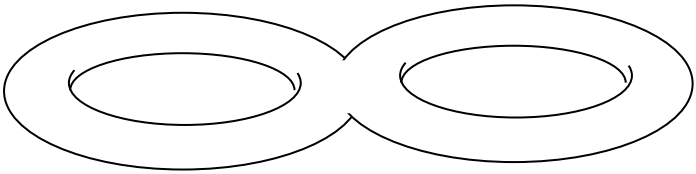

It’s clear that S + T doesn’t depend on where the holes are cut since any hole can be moved into any other position by an appropriate homeomorphism. Moreover the sum is independent of the orientation of the holes in the joining-up process. You could put T inside S and join up the holes with the opposite orientation and you’d get the same sum.

Addition of surfaces is commutative: S + T is homeomorphic to T + S. It’s also easy to see (though rather messy to prove in a formal way) that the operation is associative:

S + (T + U) is homeomorphic to (S + T) + U.

So we have a sort of arithmetic on the set of connected surfaces. This leads us to ask whether there's an identity for this operation.

Theorem 1: The sphere is the identity for connected sums. If S is any surface and **O** is the sphere then

$$S + \mathbf{O} \approx S.$$

Proof: Cut a hole in each of S and **O**. Cutting a hole in **O** results in a disk. By a suitable homeomorphism the boundary of this disk can be deformed so that it fits the hole in S. The resulting surface is the original surface, in which a hole has been cut and then patched up – in other words we return to the original surface.

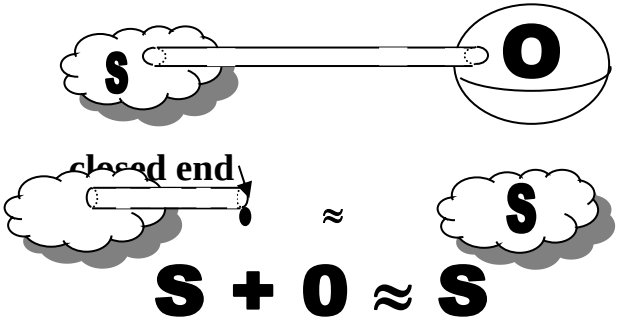

Theorem 2: If S is a surface and **D** is a disk then S + **D** is homeomorphic to the original surface with a hole cut out of it.

Proof: Cut a hole in each of S and **D**. Then **D** becomes a disk with one hole, that is, a sphere with two holes or, more simply a cylinder. We now join this cylinder to the hole that we've made in S. The area around what was the

hole can be deformed and the cylinder can be contracted back so that it becomes part of the surface. What's left is the original surface with one hole.

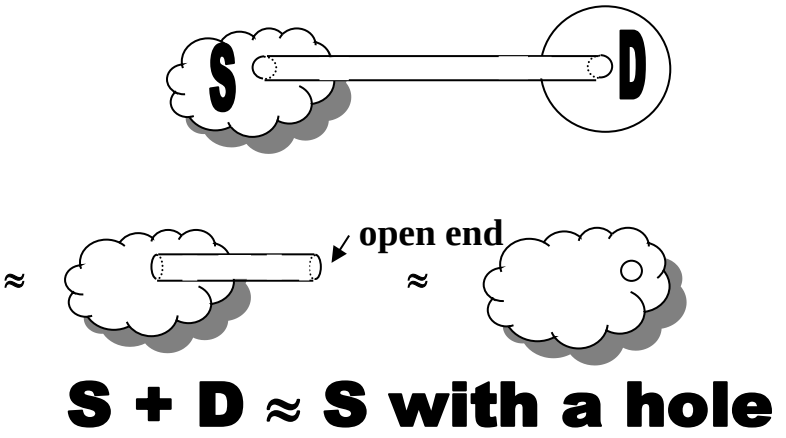

Example 6: A cylinder can be written as **D** + **D**, where **D** is a disk. This is because a cylinder is homeomorphic to a disk with an extra hole (the boundary of the disk is the other hole).

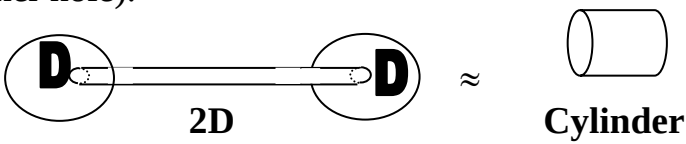

As in ordinary algebra we write **D** + **D** as 2**D** and, more generally, where S is any surface we write S + S + … + S, with *n* copies, as *n*S.

We can perform addition of connected surfaces at the level of atlases. Suppose S and T are atlases in which

the label *x* isn't used. An atlas for S + T is obtained by the following process:

Break open the polygons surrounding one of the maps in each atlas. Let the end vertices that result in S be S_1 and S_2 and let the end vertices in T be T_1 and T_2.

Now join S_1 to T_1 by an edge labelled *x* (with an arrow from S_1 to T_1) and join S_2 to T_2 by an edge, also labelled *x*, in the direction from S_2 to T_2.

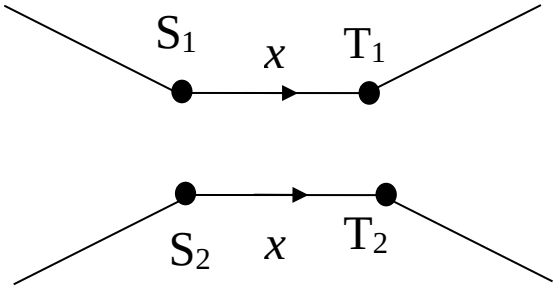

The identification of the edges marked *x* will identify S_1 with S_2 and T_1 with T_2 and the edges marked *x* can be rolled up to form a cylinder. This cylinder can be absorbed into the surface at each end and so may be removed provided we remember that S_1 and S_2 are to be identified. The removal of the *x* edge no longer guarantees this. But if the remaining identified edges identify S_1 with S_2 we are OK.

Theorem 3: Suppose we have two PIEs S and T and we cut open each at some point and join the two PIEs to form a single one. Suppose that the vertices where the join took place are identified, as a result of the identification of the edges in this combined PIE. Then the combined PIE represents S + T.

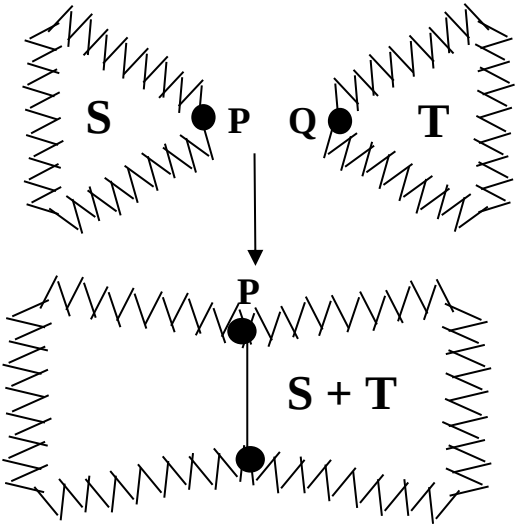

Example 7: The following PIEs represent a torus and a cylinder respectively.

We can represent their sum as a PIE as follows:

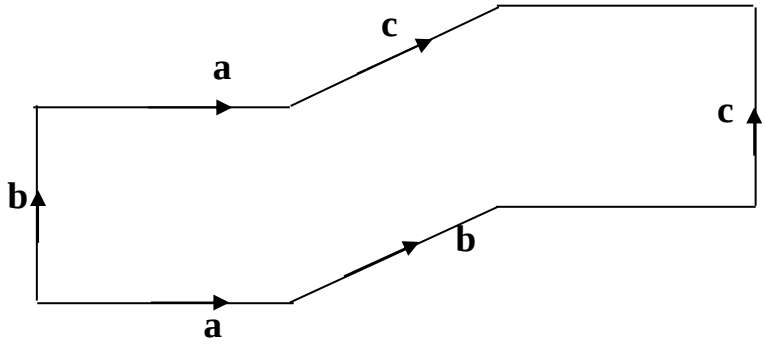

Example 8: A Möbius Band is homeomorphic to **P** + **D** where **P** is a projective plane and **D** is a disk.

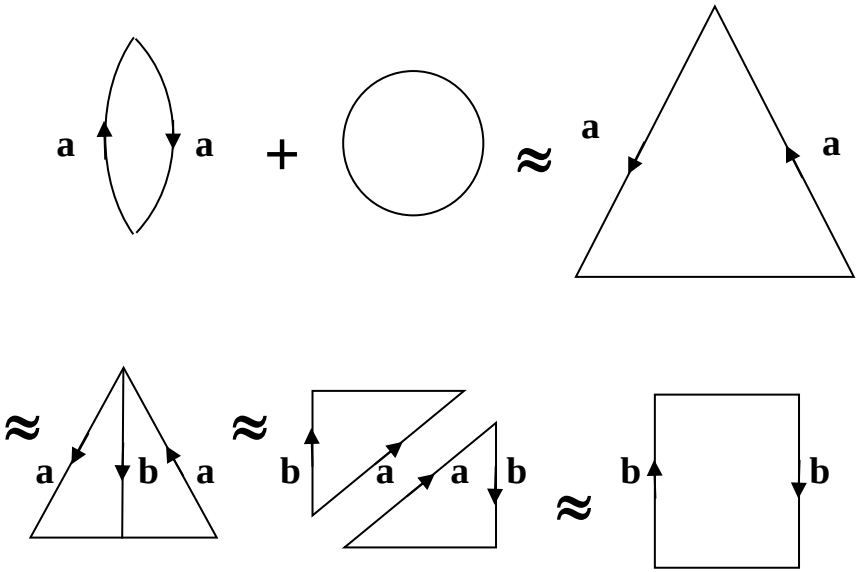

Example 9: A Klein Bottle is homeomorphic to **2P** where **P** is a projective plane.

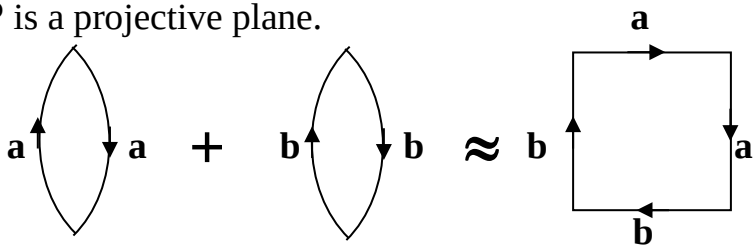

P + P = Klein Bottle

All seven of our fundamental surfaces can be expressed as a sum of disks, projective planes and toruses. Letting these be represented by **D**, **P** and **T** respectively we have:

Sphere	**O** (a sum of no terms)
Disk	**D**
Projective Plane	**P**
Torus	**T**
Cylinder	**2D**
Möbius Band	**P** + **D**
Klein Bottle	**2P**

Suppose we have two PIEs and we cut open each at some point and join the two PIEs to form a single one.

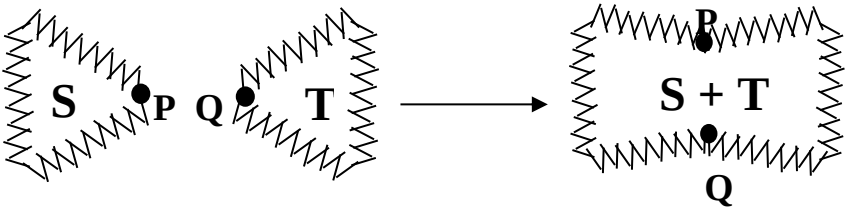

Do we get the sum of the two PIEs?

The answer is no, not always. Consider the following example.

Example 10:

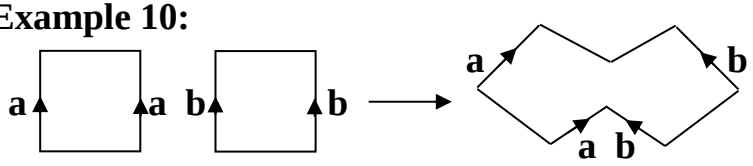

The PIE on the right clearly has 3 holes (closed boundaries). However each of the two PIEs at the left represents a cylinder and the sum of two cylinders has 4 holes.

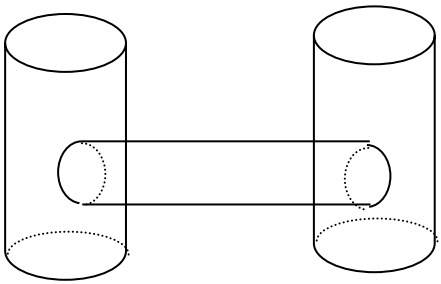

The problem is that in the combined PIE the two vertices where the join took place aren't identified as a result of the identification of the edges. Yet in the sum they must be the same point.

§3.4. The Sum of a Torus and a Projective Plane

Theorem 4: $\mathbf{T} + \mathbf{P} \approx 3\mathbf{P}$.

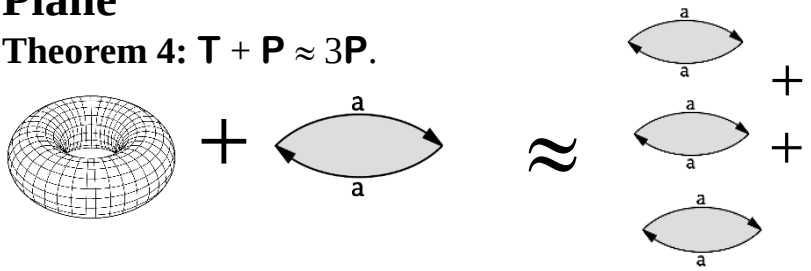

Proof: Again we can prove this by a little bit of surgery.

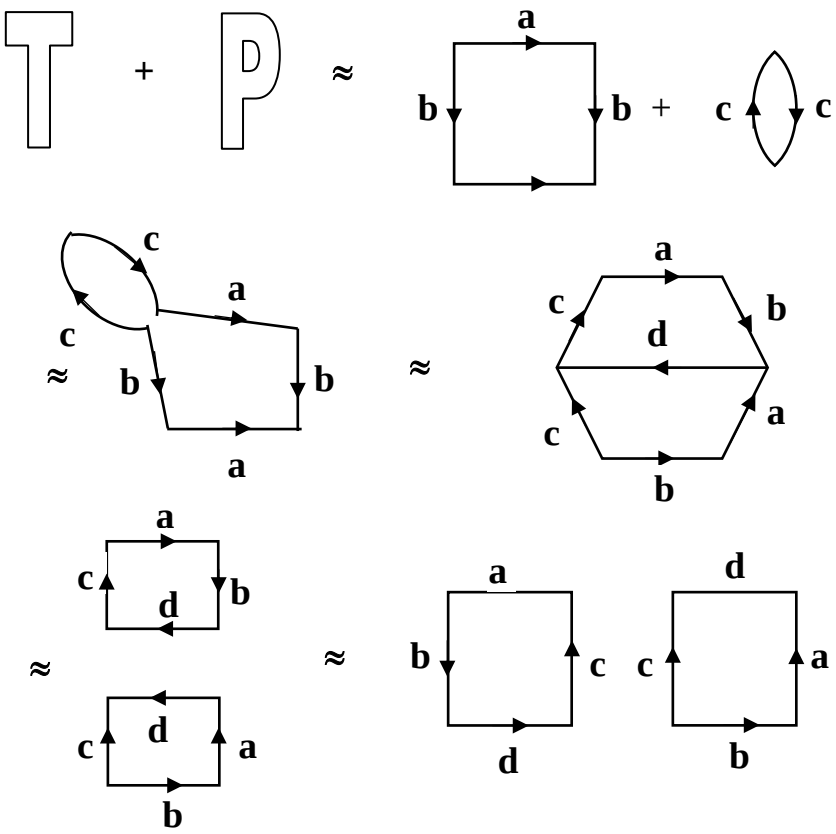

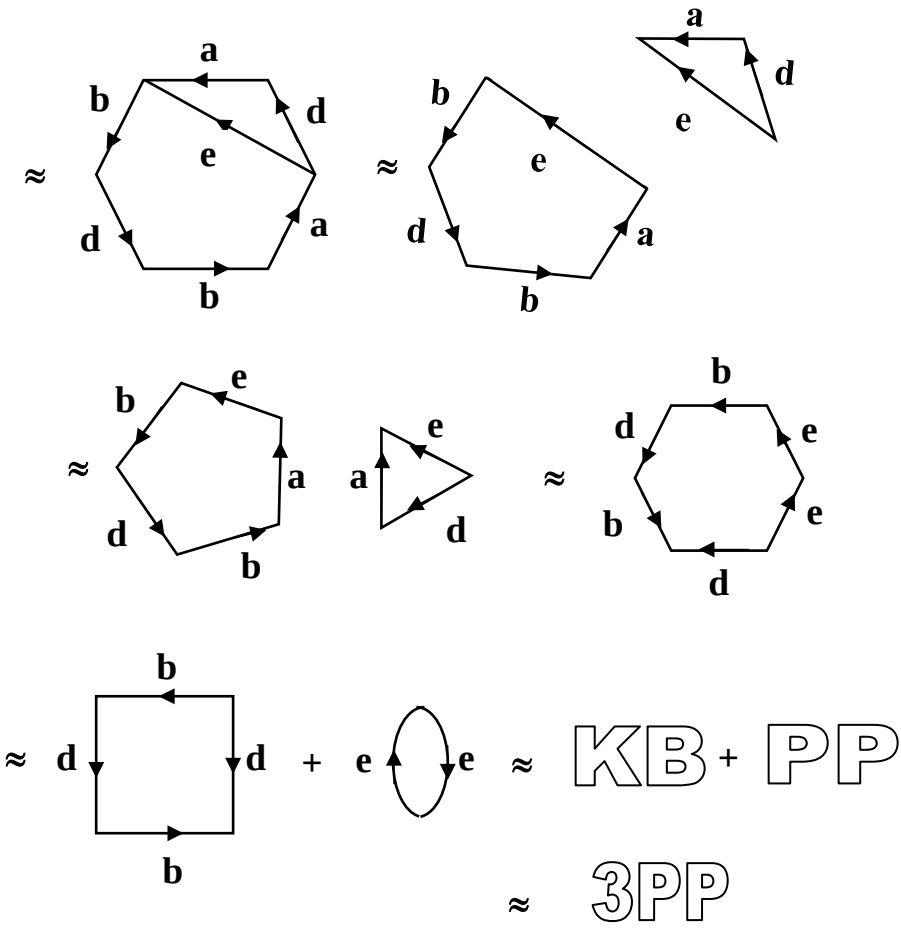

It's very tempting to say that, since $\mathbf{T} + \mathbf{P} \approx 3\mathbf{P}$ then, by subtraction, $\mathbf{T} \approx 2\mathbf{P}$. However this conclusion is false. The torus is not homeomorphic to the Klein Bottle, as we'll show later. This shows that subtraction isn't valid in the arithmetic of surfaces.

§3.5. The Fundamental Theorem of Surfaces

We define an identified pair of edges to be a **like** pair if they have the same direction around the polygon. An **unlike** pair is a pair where they have opposite directions. We define a pair of identified edges to be **adjacent** if they're consecutive edges around the polygon.

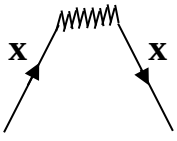	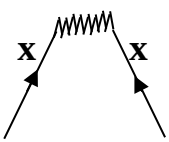	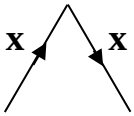	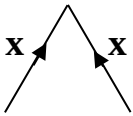
LIKE PAIR	**UNLIKE PAIR**	**ADJACENT LIKE PAIR**	**ADJACENT UNLIKE PAIR**

Theorem 5: Every PIE, S, is homeomorphic to a one of the form $m\mathbf{D} + n\mathbf{T} + r\mathbf{P}$ for $m, n, r \geq 0$.

Proof: We prove this by induction on the number of edges of S. With one edge the PIE must be a disk. Suppose we have a PIE with n edges and that the theorem holds for PIEs with fewer than n edges.

Case I: S has at least one like pair.

Cut from the tip of one arrow in the like pair to the tip of the other and join those two edges. The resulting PIE is P + S′ where S′ has fewer edges than S.

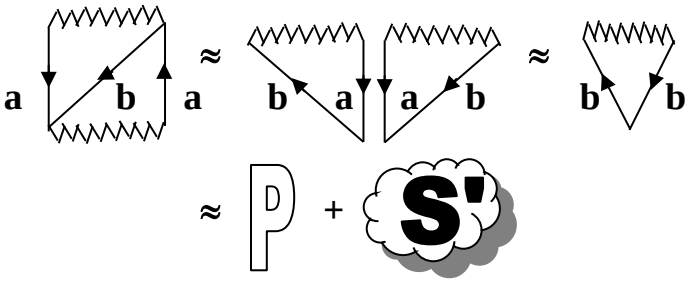

The result now follows by induction.

Case 2: All pairs are unlike and one pair is adjacent. These edges can be brought together, resulting in a PIE with fewer edges.

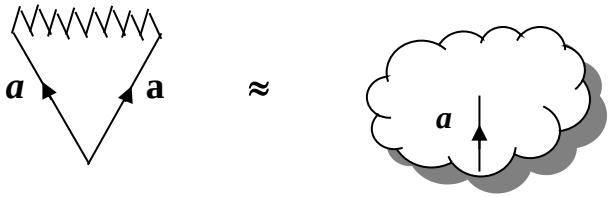

Case 3: All pairs are unlike and no pair is adjacent. Removing such a pair would divide the remaining edges into two subsets. If the number of edges (counting identified edges separately) in these two subsets are *m* and *n*, define the *distance* between the members of the pair to be the minimum of *m* and *n*. Choose a pair of identified edges, *x*, with the smallest distance apart on the diagram. Then every pair of identified edges is separated by each copy of ***x***.

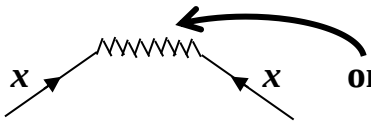

one member of every other identified pair here

Case 3A: The members of the pair of *x*'s are separated (on one side of the polygon) by unidentified edges only. These unidentified edges can be combined to form a single unidentified edge and the resulting PIE is **D** + S′ where S′ has fewer edges. The result now follows by induction.

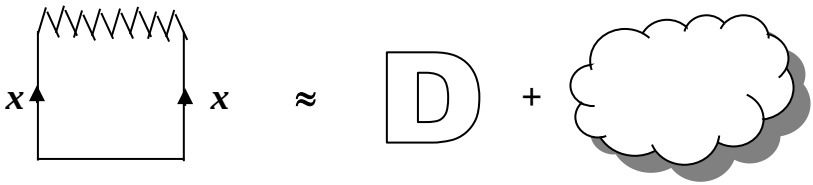

Case 3B: The pair of ***x***'s separates an identified pair. The following piece of surgery shows that S is homeomorphic to **T** + S′ where S′ has fewer edges. The result now follows by induction.

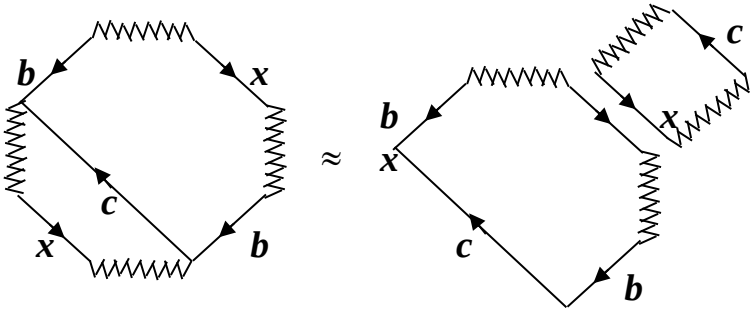

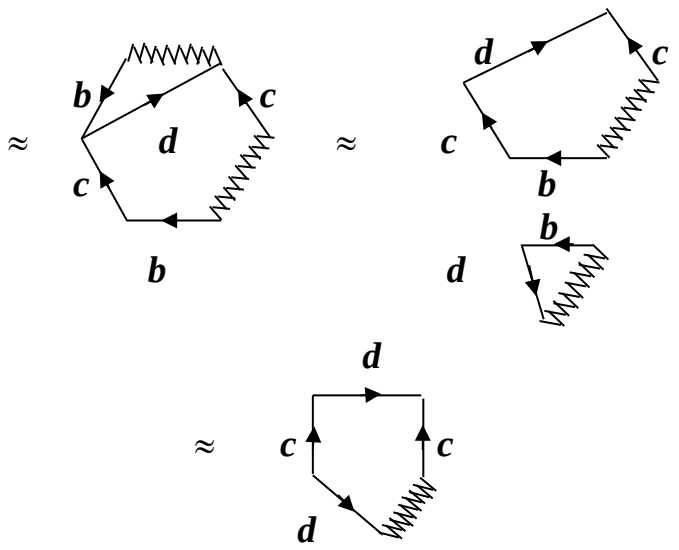

Corollary: Every surface is homeomorphic to

$m\mathbf{D} + n\mathbf{T}$ for some $m, n \geq 0$ or

$m\mathbf{D} + n\mathbf{P}$ for some $n > 0$.

Proof: Since $\mathbf{T} + \mathbf{P} \approx 3\mathbf{P}$ it follows that if $r \geq 1$ then $m\mathbf{D} + n\mathbf{T} + r\mathbf{P} \approx m\mathbf{D} + (2n + r)\mathbf{P}$.

EXERCISES FOR CHAPTER 3

Exercise 1: For each of the following determine whether it is true or false.

(1) The following is a polygon with identified edges that has one like pair and one unlike pair.

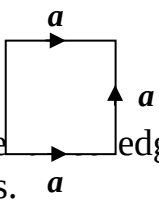

(2) If a hexagon is a polygon with identified edges and has no unidentified edges it has 3 edges.

(3) The following represents a cylinder.

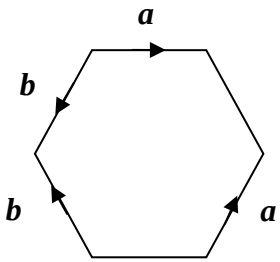

(4) The following represents a Klein Bottle.

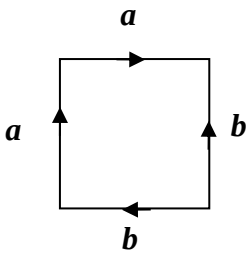

(5) A Klein Bottle plus a Cylinder is homeomorphic to a sum of two Möbius Bands.

(6) Since **T** + **P** is homeomorphic to a sum of 3 Projective Planes a torus is homeomorphic to the sum of 2 projective planes.

(7) {**D**, **T**, **P**} is a basis for the space of all polygons with identified edges.

(8) The following PIE has 2 like pairs, 1 adjacent like pairs, 1 unlike pair and 0 adjacent like pairs.

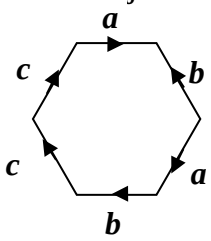

(9) Every PIE is homeomorphic to one with no adjacent unlike pairs.

(10) An adjacent like pair can be removed from a PIE.

Exercise 2:
(a) Join the four maps in the following atlas along identified edges so that it becomes a single polygon with identified edges.

(b) Identify this surface.

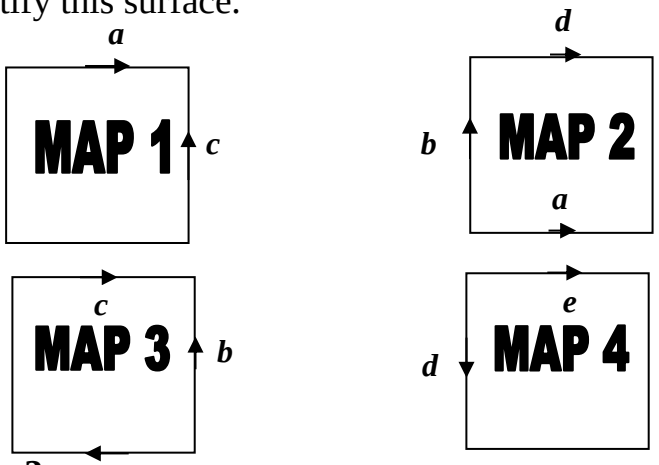

Exercise 3:

Carry out surgery (by showing a sequence of cutting and pasting) on either one of the following surfaces to obtain the other.

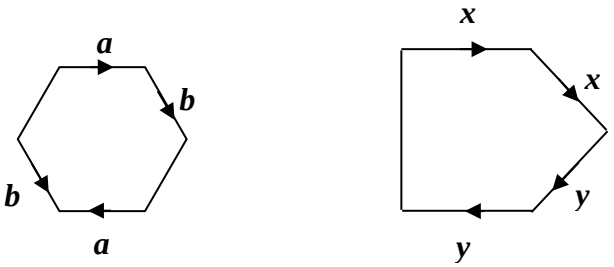

SOLUTIONS FOR CHAPTER 3

Exercise 1:
(1) FALSE (Identified edges can only come in pairs.)
(2) TRUE (Each identified pair is a single edge.)
(3) TRUE (The **b**'s may be stitched up.)
(4) FALSE (After stitching up the **b**'s we see that it is a projective plane.)
(5) TRUE since KB + Cylinder = 2**P** + 2**D** = 2(**P** + **D**)
= 2MB
(6) FALSE (Subtraction is not valid for connected sums.)
(7) FALSE (For a start PIEs don't form a vector space over any field as inverses don't exist.

Secondly **T** + **P** = 3**P** would prevent them from being linearly independent.)
(8) TRUE
(9) TRUE
(10) FALSE

Exercise 2: Rearranging the maps we get:

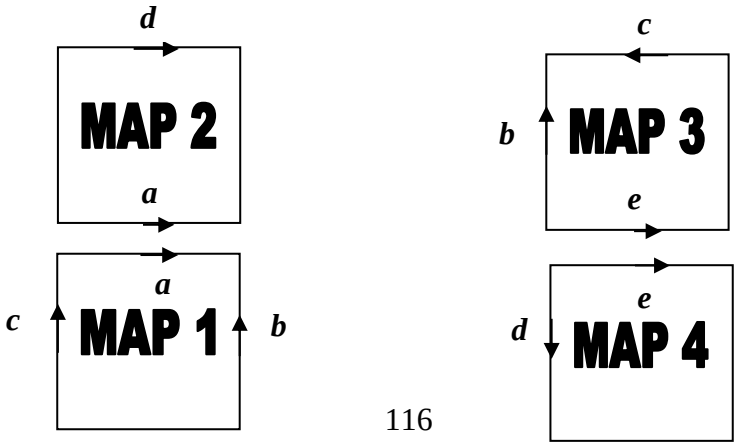

Joining maps 1, 2 and maps 3, 4 we get:

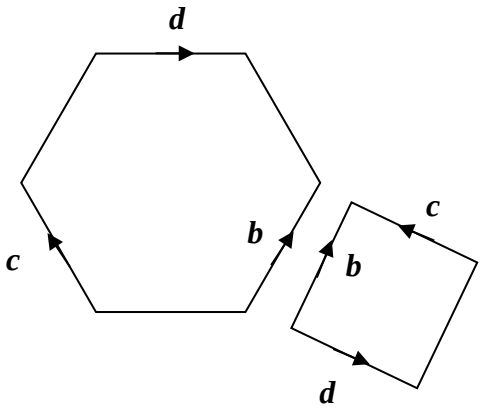

Joining these into a single polygon with identified edges we get:

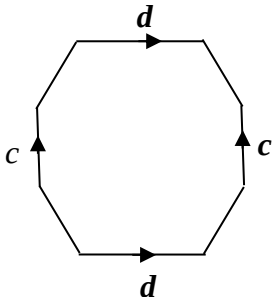

Without the unidentified edges we would get a torus. With the unidentified edges it is a torus with one or more holes. The unidentified edges form the edge around these holes. It is not easy to see that here there's just one hole. In the next chapter we'll develop techniques that will make such questions routine.

Exercise 3: