Introduction to Knot Theory

George Voutsadakis[1]

[1]Mathematics and Computer Science
Lake Superior State University

LSSU Math 500

Subsection 1

Surfaces without Boundary

Surfaces or Two-Manifolds

- A **surface** or **two-manifold** is defined to be any object such that every point in the object has a neighborhood in the object that is a (possibly nonflat) disk.

  Examples: The **sphere** (surface of a ball) and the **torus** (surface of a doughnut) are two-manifolds.

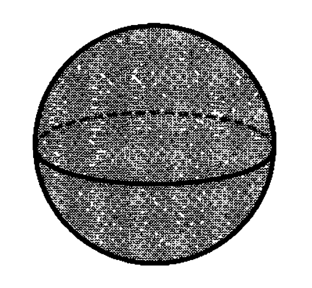

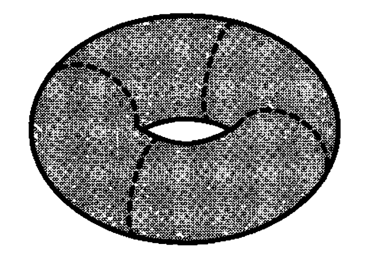

Examples

- Surfaces:

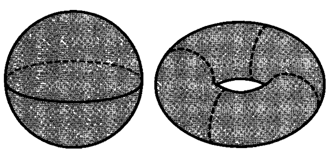

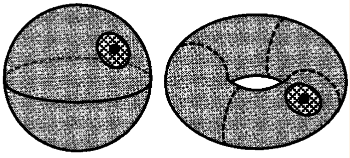

- Non-Surfaces:

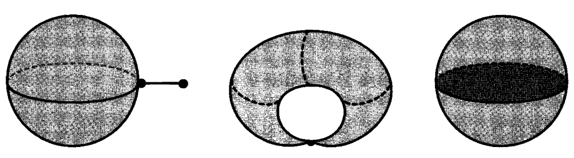

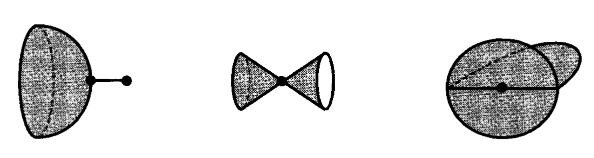

Deformability of Surfaces

- We think of all surfaces as being made of rubber, and hence deformable.

  Example: We consider a sphere and a cube to be equivalent surfaces, since we could pullout eight points on a rubber sphere to make it look like a cube, without having to do any cutting and pasting.

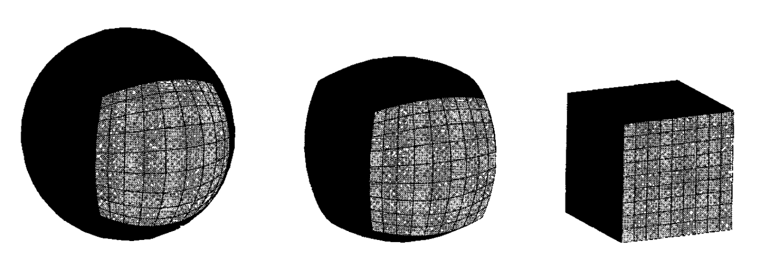

Isotopies

Example: Consider the three surfaces shown below.

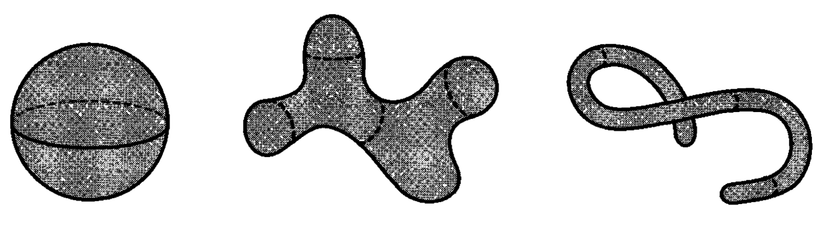

They are equivalently placed in space because we could get from anyone to any other by a rubber deformation.

- We call such a rubber deformation an **isotopy**.
- Two surfaces in space that are equivalent under a rubber deformation are called **isotopic surfaces**.

Ttriangulations

- In order to better work with surfaces, we divide them into triangles.
- The triangles have to fit together nicely along their edges so that they cover the entire surface.
- They cannot intersect each other in any of the ways pictured:

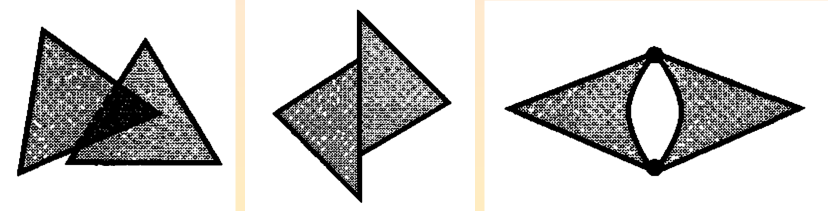

- The triangles need not be flat with straight edges, since they are deformable.
- We call such a division of a surface into triangles a **triangulation**.

Examples

- Examples of triangulations of the sphere and the torus.

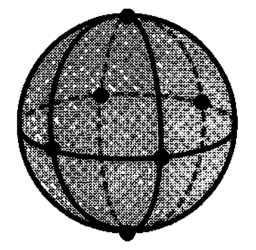

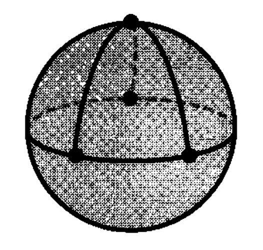

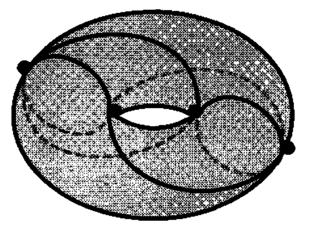

Cutting the Surface into Triangles

- Given a surface with a triangulation, we can cut it into the individual triangles, keeping track of the original surface by:
  - Labeling the edges that should be glued back together;
  - Placing matching arrows on the pairs of edges that are to be glued.

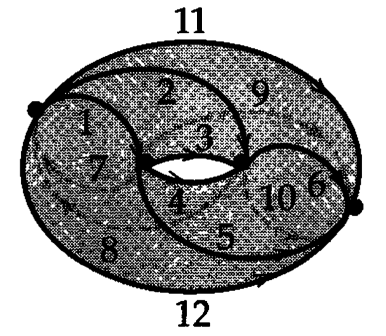

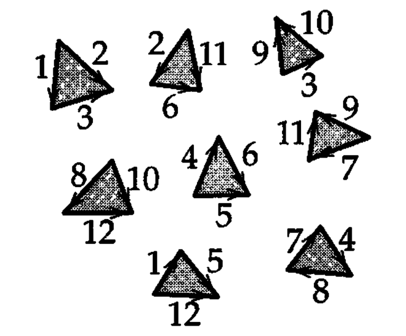

Homeomorphic Surfaces

- We say that two surfaces are **homeomorphic** if from one of them we can obtain the other via the following process:
  - Triangulate;
  - Cut along a subset of the edges into pieces;
  - Glue back together along the edges according to the instructions given by the orientations and labels on the edges.

Example: Two homeomorphic copies of the torus.

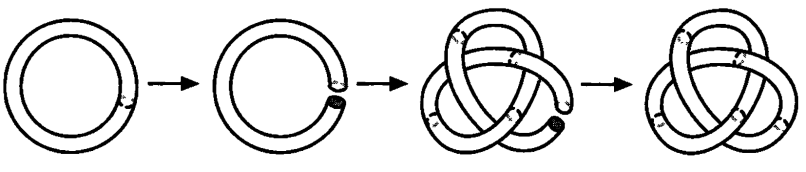

- We cut along a subset of the edges of a triangulation that form a circle.
- We knot the resulting cylinder;
- We glue the two circles back together.

Non-Homeomorphic Surfaces

- A sphere and a torus are not homeomorphic.

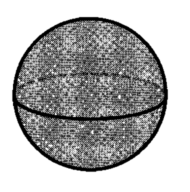

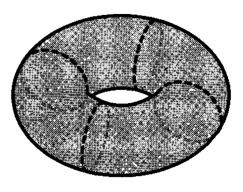

- We could also have the surface of a two-hole doughnut or a three-hole doughnut.

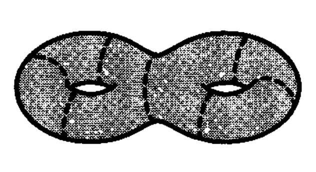

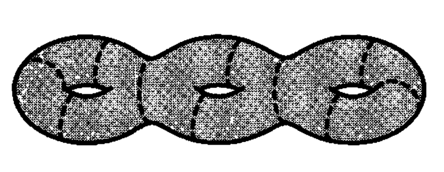

- None of these four examples are homeomorphic to one another.

Genus of a Surface

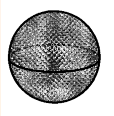 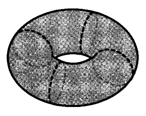 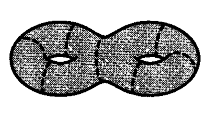 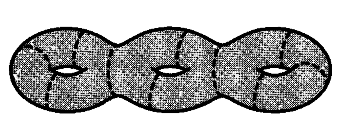

- Since we could just keep increasing the number of holes in our doughnuts, there are an infinite number of distinct (nonhomeomorphic) surfaces.
- We call the number of holes in the doughnut the **genus** of the surface.

  Example:
  - The sphere has genus 0;
  - The torus has genus 1;
  - The 2- and 3-holed doughnut surfaces have genera 2 and 3.

Embeddings of a Surface in Space

- Each of the surfaces considered can be placed in space in different ways.

  Example: We saw two ways to put a torus in space.

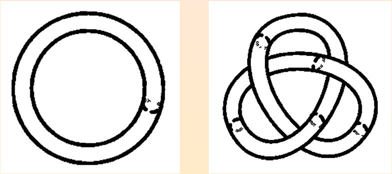

- Even though both of those surfaces were tori, they were not isotopic, since there was no rubber deformation that would take us from the one to the other.
- However, they were still homeomorphic surfaces, just placed in space in two different ways.
- A specific placement of a surface in space is an **embedding** of the surface.

Example: Embeddings of a Genus 3 Surface in Space

- The figure shows three distinct embeddings of a genus 3 surface in space.

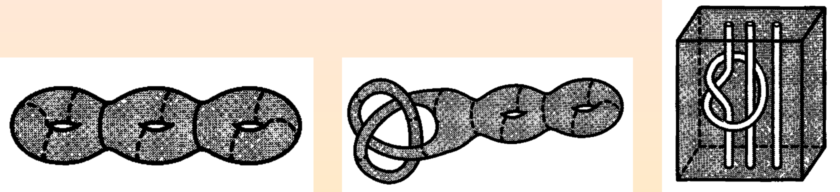

- Although they are all homeomorphic to one another, only the first and the third are isotopic to one another.
- On the third surface, we can slide the end of one of the tubes along another tube to un- knot the knotting.
- We call the third surface the surface of a **cube-with-holes**.

  It is the surface of the solid object obtained by drilling three wormholes out of a cube.

Homeomorphism Type of a Surface

- Given a surface in space, we would like to be able to tell what surface it is, i.e., what is its **homeomorphism type**.
- It might be a sphere or a torus, but so mangled that we do not recognize it.
- One option is to cut and paste to simplify the appearance of our surface until we can identify it.
- This technique requires us to make a clever choice of how to cut up the surface and rearrange the pieces before regluing.
- It would be better if there were a method for recognizing surfaces that did not require the cut-and-paste technique.

The Euler Characteristic

- Consider a triangulation of the surface and let:
  - V be the number of vertices;
  - E be the number of edges;
  - F be the number of triangles (F stands for faces).
- We define the **Euler characteristic** of the triangulation to be

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

Example: Consider the triangulation of the sphere in the figure.

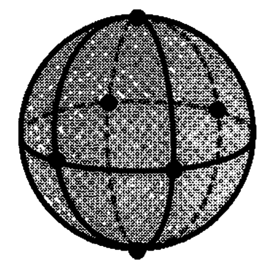

We have

$$V = 6, \quad E = 12, \quad F = 8.$$

So

$$\chi = 6 - 12 + 8 = 2$$

Independence from Triangulation

- The Euler characteristic depends only on the surface, not on the particular triangulation of the surface that we use.

  We describe the idea behind the proof avoiding the technicalities.

  Let T_1 and T_2 be two different triangulations of surface S.

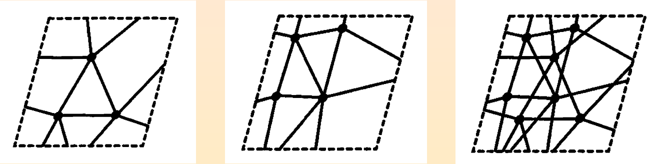

  Place them on the surface so that they are overlapping.

  We build a new triangulation T_3 of S that "contains" T_1 and T_2.

  We show that it has the same Euler characteristic as T_1.

  The same argument shows T_3 has the same Euler characteristic as T_2.

  So it shows that T_1 and T_2 have the same Euler characteristic.

Independence from Triangulation (Adding Vertices)

- We will assume that:
  - Each edge of T_1 intersects each of the edges of T_2 a finite number of times. There is a technical proof that the edges of T_1 can be moved just slightly to make sure that this is the case.
  - The vertices of T_2 do not lie on top of a vertex or edge from T_1. This can be made true by moving T_1 slightly.

We begin to build the new triangulation T_3 by starting with T_1.

One at a time, we add to the vertices of T_1 a new set of vertices corresponding to where the edges of T_2 cross the edges of T_1.

Each new vertex also cuts an edge into two edges.

When computing χ, the number of vertices is added and the number of edges is subtracted.

So the Euler characteristic is unchanged by this operation.

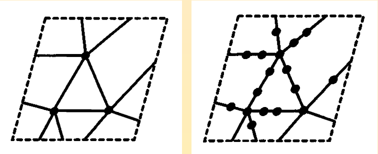

Independence from Triangulation (Cont'd)

- We also add each vertex in the second triangulation T_2 to T_3, together with one edge that runs from that vertex to one of the vertices that is already in T_3.

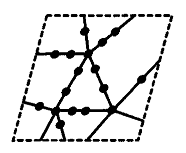
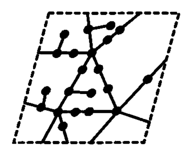

We choose each of these new edges to be a subset of one of the edges from T_2.
The number of faces does not change.
The numbers of vertices and edges each goes up by one.
So this addition does not change the Euler characteristic.
We may need to add a chain of edges to connect a vertex of T_2 and T_3.
However, the Euler characteristic remains unchanged.

Independence from Triangulation (Cont'd)

- Now we add all of the pieces of edges from T_2 that have not been added yet, each of which becomes a separate edge in T_3.

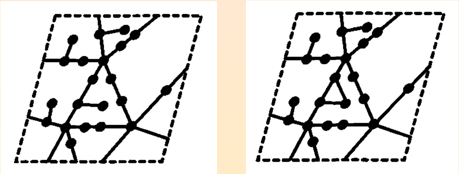

Note that as we add one of these edges, we cut a face in two.

Hence, the numbers of edges and of faces each goes up by one.

This leaves the Euler characteristic unchanged.

Independence from Triangulation (Conclusion)

- At this point, we may not have a triangulation. Some of the faces may not be triangles. So now we just add edges to cut the faces into triangles.

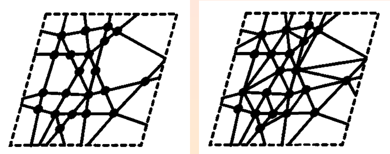

When we add such an edge, it cuts an existing face into two pieces.

So both numbers of edges and faces go up by one.

This leaves the Euler characteristic unchanged.

We showed that there exists a third triangulation, T_3, with the same Euler characteristic as T_1, such that it "contains" both T_1 and T_2.

Note that we could have built T_3 by starting with T_2.

So it also has the same Euler characteristic as T_2.

Hence, we have shown that T_1 and T_2 must have the same Euler characteristic.

Euler Characteristic of Sphere and Torus

- The Euler characteristic only depends on the type of surface that we have, and not on the particular triangulation.

  Example: Any triangulation of the sphere has Euler characteristic 2.

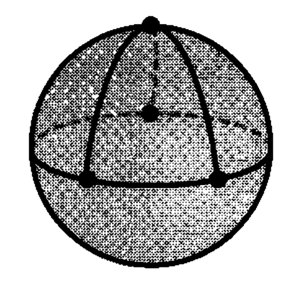

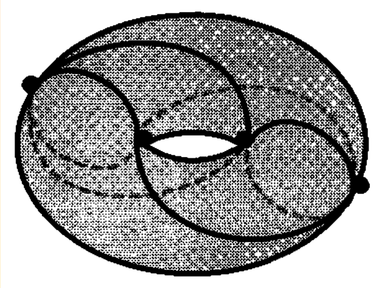

Any triangulation of the torus has Euler characteristic 0.

Genus 2 Surface: Connected Sums

- To compute the Euler characteristic of a genus 2 surface S, we could:
  - Take a triangulation of the surface;
  - Compute its Euler characteristic.
- Instead, we notice that, to obtain a genus 2 surface, we may:
  - Remove a disk from each of two tori T_1 and T_2;
  - Glue the tori together along the resulting circle boundaries.

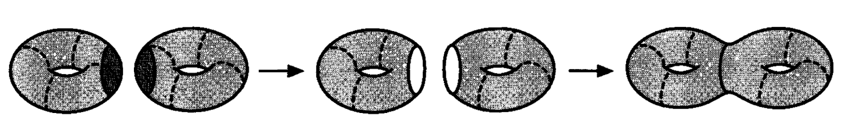

- This operation is called the **connected sum** of the tori.

Connected Sum of Triangulated Surfaces

- Suppose that we already have triangulations of the two tori.
- Then we can think of taking their connected sum as:
  - Removing the interior of a triangle from each torus;
  - Gluing together the boundaries of the two missing triangles by pairing up the vertices and edges.

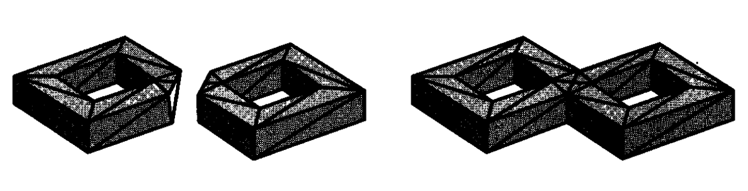

- The result is a triangulated genus 2 surface.
- Now we have a triangulation for it.
- So we can figure out what the Euler characteristic is.

Characteristic of Connected Sum

- The total number of vertices, edges, and faces in the triangulation of S is just the total number of vertices, edges, and faces in T_1 and T_2 with:
  - Three fewer vertices, since we identified three vertices in T_1 with three vertices in T_2;
  - Three fewer edges, since we identified three edges in T_1 with three edges in T_2;
  - Two fewer faces, since we threw away the interiors of two triangles in order to construct the connected sum.

  In the formula for χ, V is added and E is subtracted.

  So the loss of three vertices and three edges has no net effect on χ.

  Hence, the only effect is the loss of two faces.

  We conclude that

$$\chi(S) = \chi(T_1) + \chi(T_2) - 2.$$

  Since the Euler characteristic of a torus is 0, $\chi(S) = -2$.

Generalizing Triangulation

- We make the computation of Euler characteristic even easier.
- We subdivide the surface into vertices, edges, and faces.
- A face is simply a disk with its boundary made up of a sequence of edges connecting the vertices (better known as a polygon).
- A face must be a single piece that has no holes in it.

Example

- We could subdivide the torus into a single face, with one vertex and two edges. This gives

$$\chi = V - E + F = 1 - 2 + 1 = 0.$$

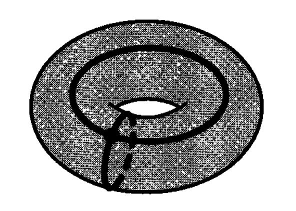
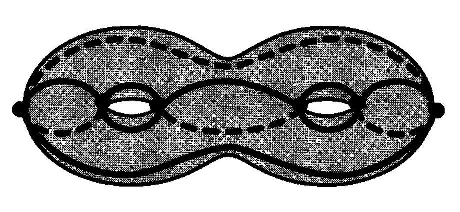

- We could cut the genus 2 surface up into 4 faces, with 6 vertices and 12 edges. This yields

$$\chi = V - E + F = 6 - 12 + 4 = -2.$$

Compact Surfaces

- A hard technical fact is that every surface has a triangulation.
- However, not every surface has one with a finite number of triangles.
- We say that a surface is **compact** if it has a triangulation with a finite number of triangles.

  Example: The sphere and torus are compact surfaces.

  But neither the plane nor a torus minus a disk is compact, as neither can be triangulated with finitely many triangles.

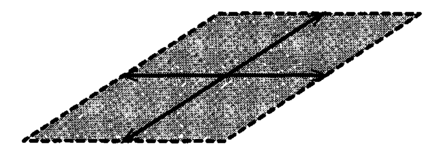

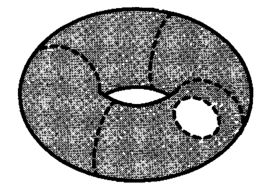

- Note that both the plane and the torus minus a disk do satisfy the definition of a surface.

The Complement of a Knot

- Surfaces appear in knot theory in the space around the knot.
- Let $\mathbb{R}^3$ be the three-dimensional space that the knot K sits in.
- The space M around the knot is everything but the knot.
- This set $M = \mathbb{R}^3 - K$ is called the **complement of the knot**.
- The complement M is what is left over if we drill K out of space.

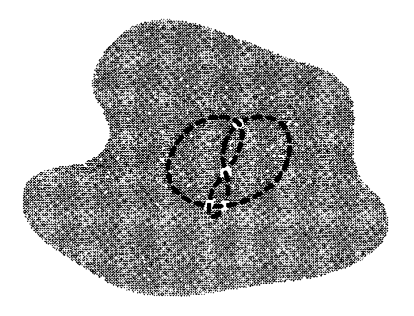

Complements and Splittable Links

- The figure shows an example of a surface in the complement of a link when the link is splittable.

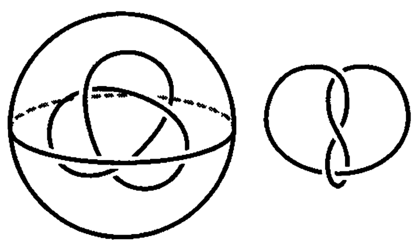

- Since we can pull the components of the link apart, we can think of there being a sphere that separates the components from one another.
- An alternative way to define a *splittable link* is simply to say that it is a link such that there is a sphere in the link complement that has components of the link on either side of it.

Additional Examples

- Every knot is contained in a torus.

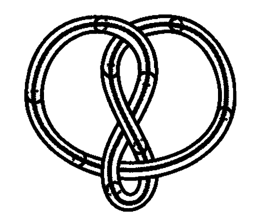

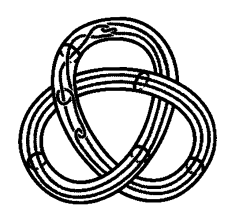

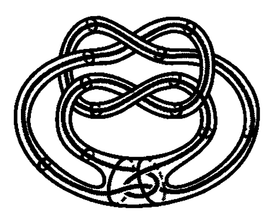

- The second figure contains a torus that surrounds a knot in a more unusual way.
- The last figure shows a genus 2 surface in the complement of a knot.

Compressible Surfaces

- Let L be a link in $\mathbb{R}^3$.
- Let F be a surface in the complement $\mathbb{R}^3 - L$.
- We say that F is **compressible** if there is a disk D in $\mathbb{R}^3 - L$, such that:
  - D intersects F exactly in its boundary;
  - The boundary of D does not bound another disk on F.
- Example: The surface F in the figure is compressible.

  The disk D is a disk in $\mathbb{R}^3$, such that:
  - D does not intersect the link L;
  - D intersects F exactly in its boundary;
  - The boundary of D does not bound a disk on F.

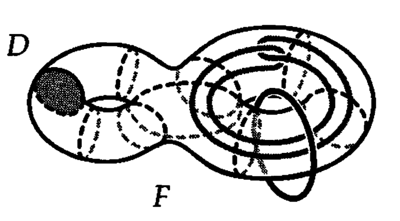

Compression

- A compressible surface can be simplified by:
  - Cutting it open along the boundary of the disk;
  - Gluing two copies of the disk to the two curves that result.
- We obtain a simpler surface (or sometimes a pair of surfaces) that still lies in the link complement.

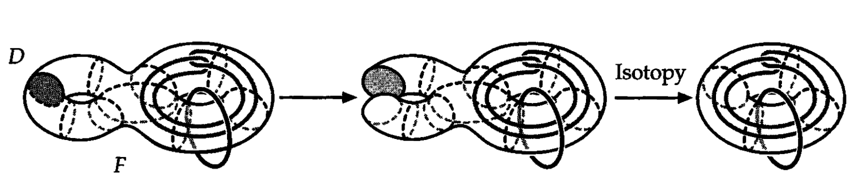

- This simplifying operation is called a **compression** of the original surface.

Incompressible Surfaces

- If a surface is not compressible, we say that it is **incompressible**.

  Example: The torus shown is incompressible, although proving it is somewhat difficult.

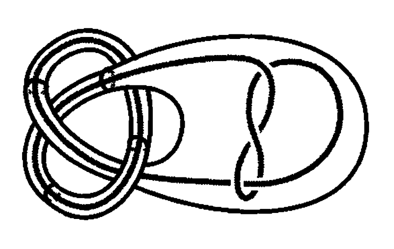

- Notice that any disk that intersects the torus in its boundary looks like it must satisfy one of the following:
  - It must intersect the link L;
  - Its boundary must cut a disk off of the torus.

Swallow-Follow Tori

- An incompressible torus like the one in the figure

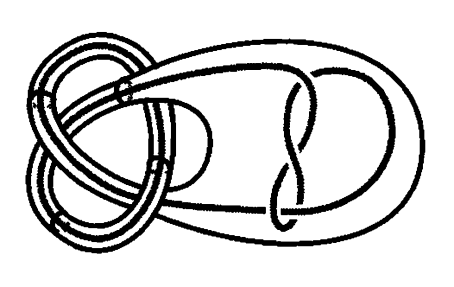

exists any time that we have a composite knot.

It is called a **swallow-follow torus** because:

- It swallows one of the two factor knots;
- It follows the other one.

Subsection 2

Surfaces with Boundary

Surfaces with Boundary

- In order to obtain **surfaces with boundary**, we can just remove the interiors of disks from the surfaces that we already have.

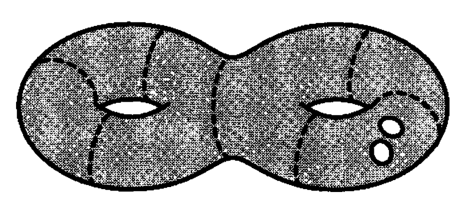

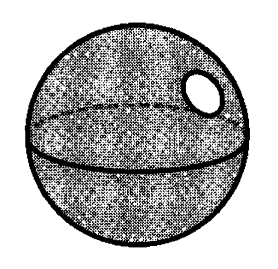

- We leave the boundaries of the disks in the surfaces, which become the boundaries of the surfaces.
- All of the resulting boundaries are circles.
- These circles are called **boundary components**.

Deformations of Surfaces with Boundary

- Since all of our surfaces are made of rubber, they can look very different when we deform them.

  Example: The figure shows two different pictures of a torus with one boundary component.

  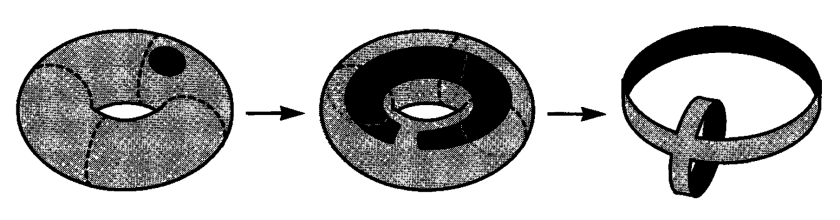

  It also shows the deformation for getting from the one picture to the other.

Surfaces with Boundary and Euler Characteristic

- When we remove a disk from a surface without boundary, we can think of it as removing the interior of one triangle in a triangulation of the surface.
- Hence the Euler characteristic goes down by one.

  Example: Consider a surface with three boundary components.
  It has an Euler characteristic three less than the Euler characteristic of the surface obtained by filling in the three boundaries with disks.

  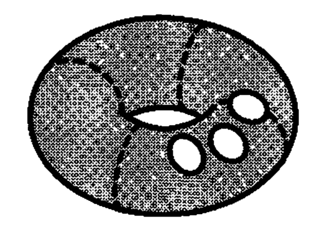

- Filling in boundary components by attaching disks is called **capping off** a surface with boundary.

Example

- We can find the Euler characteristics of the surface in the figure without triangulating it.

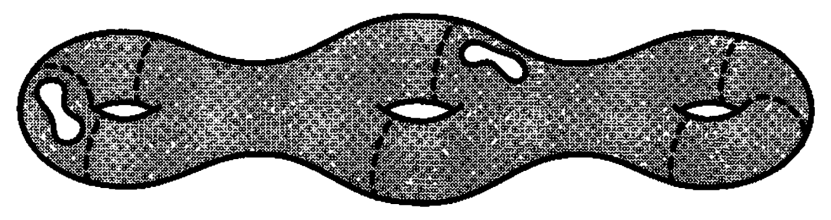

- A genus 3 surface has Euler characteristic -4.
- Since it has two boundary components it Euler characteristic is 2 less than -4.
- Therefore, the surface in the figure has Euler characteristic -6.

Euler Characteristic and Distinguishability

- Unlike surfaces without boundary in three-space, surfaces with boundary cannot all be distinguished by Euler characteristic.

  Example: The figure contains two surfaces with boundary that have the same Euler characteristic, but that are not homeomorphic.

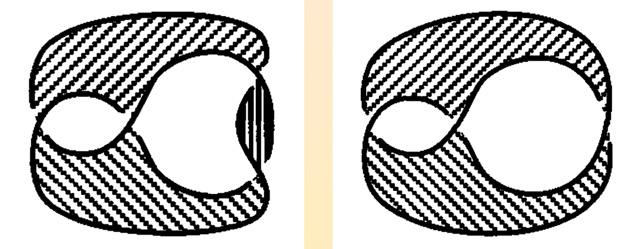

- It might help to picture these surfaces by thinking of their boundaries as wire frames and the surfaces as soap films spanning the wires.

Constructing Surfaces Using Paper

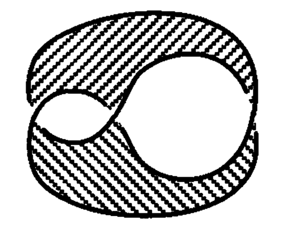

- To construct the surface on the right from paper:
  - Cut out two larger disks and three thin strips of paper;
  - At one end of each of the disks, tape two of the strips running from one disk to the other, each with a half twist in it;
  - Tape the last strip from the one disk to the other with a full twist in it, making sure that the direction of the twist matches the direction of the twist in the figure.

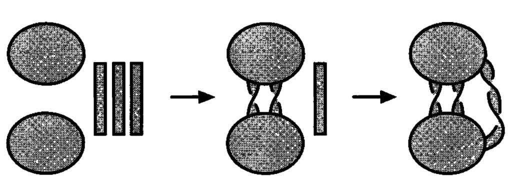

Orientability

- Consider again the two surfaces in the picture.
- A trait (other than Euler characteristic) that distinguishes between the two is orientability.

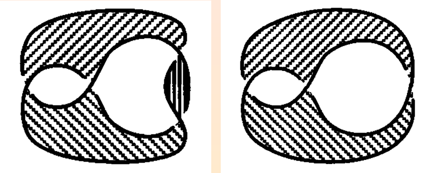

- Suppose we start painting one side of the first surface gray.

  If we continue to paint that side, eventually we will end up painting the entire surface gray on both sides.

  In essence, the surface does not have two distinct sides.
- On the second surface, the two sides could be painted black and white so that nowhere would any black paint touch any white paint.

  There really are two distinct sides of the surface.
- We say that the second surface is **orientable**.

Orientable Surfaces

- A surface sitting in three-dimensional space is **orientable** if it has two sides that can be painted black and white, so that the black paint never meets the white paint except along the boundary of the surface.

  Example: A torus is an orientable surface, because we could always paint the outer side black and the inner side white.

  Also, the following surfaces with boundary are all orientable.

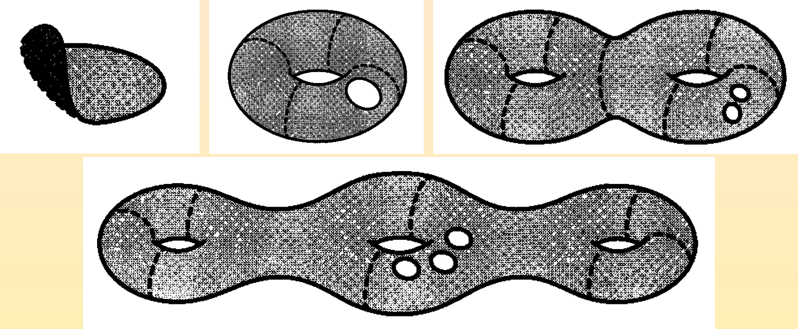

Non-Orientability: The Möbius Band

- One of the simplest surfaces that are not orientable is the Möbius band.

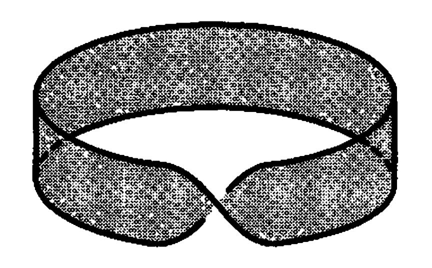

- This surface is not orientable because if we started painting one side of it black and continued working on that side, we would find that when we were done, we had painted all of it black.
- A surface that has only one side is called **nonorientable**.

Non-Orientability: The Klein Bottle

- Another nonorientable surface is shown.
- A surface is nonorientable if and only if it contains a Möbius band within it.
- The Möbius band may have an odd number of half-twists in it rather than just one half-twist, since it would be homeomorphic to the usual Möbius band.
- We can easily imagine a copy of the Möbius band on the Klein bottle pictured.

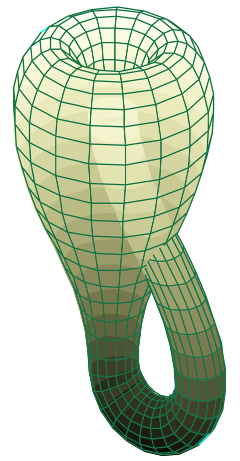

Determining the Homeomorphism Type of a Surface

- Suppose that we have a surface with boundary and we want to figure out what surface it is.
- To do so, we need to know three facts.
  1. Whether it is orientable or nonorientable;
  2. How many boundary components it has;
  3. Its Euler characteristic.

  These three pieces of information will completely determine the homeomorphism type of the surface.

Example

- Consider the surface shown in the figure.

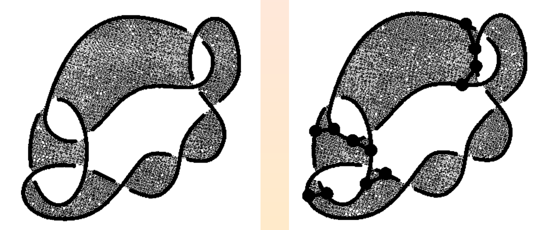

  - The surface is orientable.
  - It has three boundary components.
  - It can be subdivided, as in the second figure, in order to determine that its Euler characteristic is -3.
- Therefore, if we cap off its boundary components with three disks, the resulting surface without boundary will have $\chi = 0$.
- So the resulting surface without boundary is a torus.
- Hence, our surface is simply a torus with three disks removed.

Genus of Surfaces with Boundary

- If a surface has boundary, we define its **genus** to be the genus of the corresponding surface without boundary obtained by capping off each of its boundary components with a disk.

  Example: Consider again the following surface with boundary.

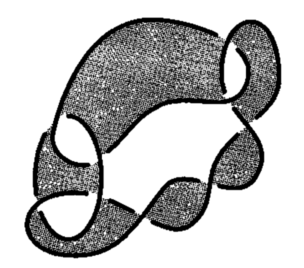

  It results from the torus by removing three disks.

  So its genus must be 1.

Surfaces with Boundary and the Unknot

- One way to define the unknot is to say that it is the only knot that forms the boundary of a disk.

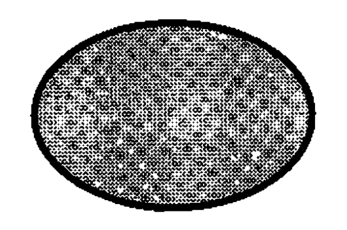

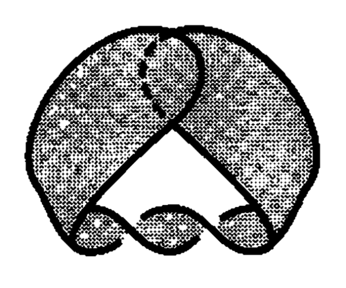

- In some projections of the unknot, the disk is not at all obvious, but it is always there.

Surfaces with Boundary and Composite Knots

- If we have a composite knot, there is a sphere with two boundary components that lies outside the knot.
- This surface is also called an **annulus**.

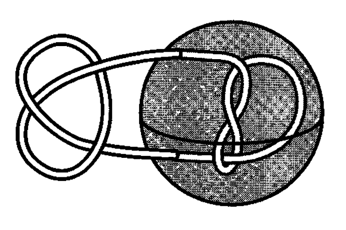

- Note that we thickened the knot up a little in this picture.
- If we had left the knot infinitely thin, we would have said the surface was a sphere with two punctures, the punctures occurring where the knot passed through the sphere.
- An alternative definition of a *composite knot* is a knot such that there is a sphere in space punctured twice by the knot, such that the knot is nontrivial both inside and outside the sphere.

Surfaces with Boundary and Tangles

- We looked at tangles as regions in the projection plane with four outgoing strands.
- We can also think of a *tangle* as a portion of the knot surrounded by a sphere with four punctures, the punctures occurring where the knot passes through the sphere.
- Such a sphere is called a **Conway sphere**.

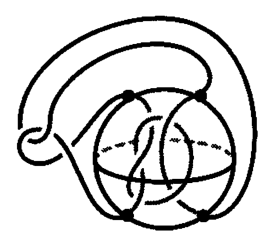

- If we thicken up the knot, the punctures become holes and we have a sphere with four boundary components.

The Möbius Band and the Trefoil Knot

- In the figure is pictured a Möbius band with boundary the trefoil knot.

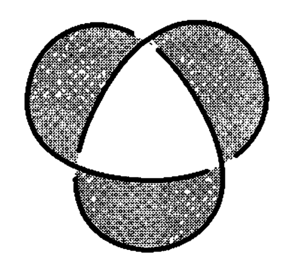

- Even though the band has three twists instead of one, it is still (homeomorphic to) a Möbius band:
  - We can cut this band open along an arc;
  - Untwist one full twist;
  - Reidentify the points we first cut along, obtaining the Möbius band.

Torus with Boundary and Trefoil Knot

- We will be particularly interested in orientable surfaces with one boundary component such that the boundary component is a knot.

  Example: The picture shows a torus with one boundary component where that boundary component is a trefoil knot.

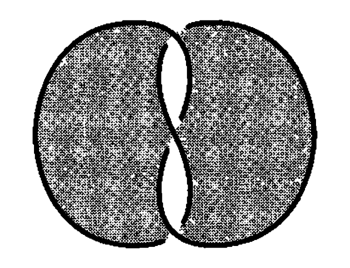

- The surface pictured does not look like a torus with one boundary.
- We can use the Euler characteristic to show that it actually is.

Subsection 3

Genus and Seifert Surfaces

Seifert's Algorithm

- Seifert's algorithm:
  - Takes as input any knot;
  - Creates an orientable surface with one boundary component, such that the boundary circle is that knot.
- Suppose we want to construct such a surface for a knot K.
- Starting with a projection of the knot, choose an orientation on K.

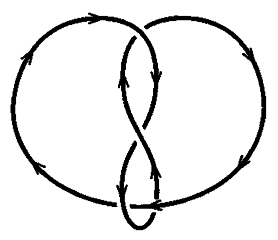

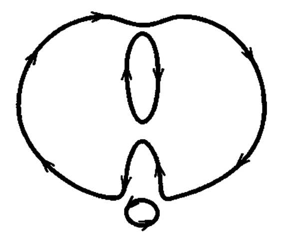

- At each crossing, two strands come in and two strands go out.
- Eliminate the crossing by connecting each of the strands coming into the crossing to the adjacent strand leaving the crossing.

Sefert's Algorithm (Seifert Circles)

- Now all of the resultant strands will no longer cross.
- The result will be a set of (topological) circles in the plane, called **Seifert circles**.

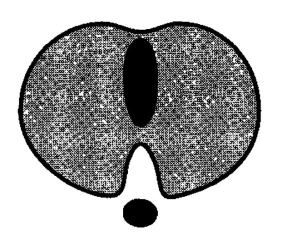

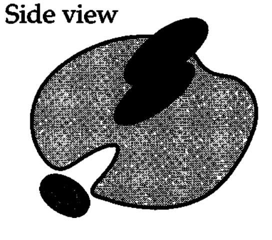

- Each circle will bound a disk in the plane.
- Since we do not want the disks to intersect one another, we will choose them to be at different heights rather than having them all in the same plane.

Sefert's Algorithm (Connecting the Circles)

- Now we would like to connect the disks to one another at the crossings of the knot by twisted bands.

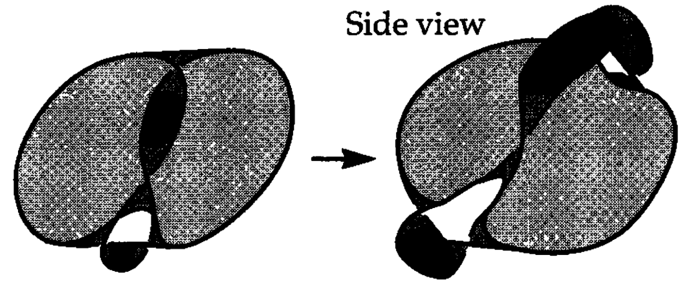

- The result is a surface with one boundary component such that the boundary component is the knot K.

A Second Example

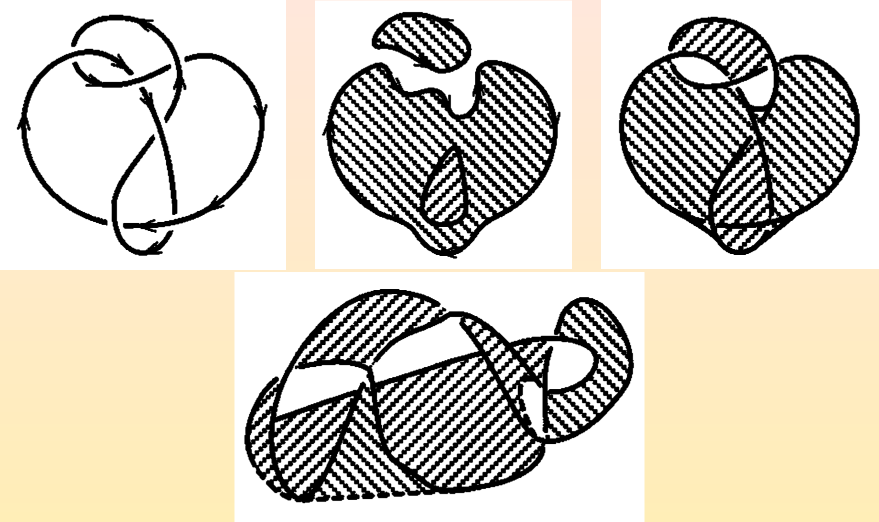

Orientability of the Generated Surfaces

- The surfaces that we are generating are always orientable.

  To see this, we need to show that each surface has two distinct sides, which can be painted two different colors.

  Let's give each Seifert circle the orientation that it inherits from the knot, either clockwise or counterclockwise.

  - For each disk that has a clockwise orientation on its bounding Seifert circle, we paint its upward pointing face white and its downward pointing face black.
  - For each disk that has a counterclockwise orientation on its bounding Seifert circle, we paint its upward pointing face black and its downward pointing face white.

Orientability of the Generated Surfaces (Cont'd)

- At each crossing in the knot, we connect two of the disks bounded by the Seifert circles by a band containing a half-twist.
  - If one of the two disks is adjacent to the other, the two disks have opposite orientations on their boundaries.
    Hence, the twist in the band allows us to extend the black and white paint across the two faces of the band so that they match up consistently on the disks.
  - If one of the two disks is on top of the other, the two disks have the same orientation on their boundaries.
    Again, the twist in the band allows us to extend the paint consistently across the band.

Thus, the entire surface can be painted black and white so that no black paint touches any white paint.

This shows that the surface is orientable.

Genus of the Generated Surface

- Consider a knot.
- Construct the Seifert surface generated by the knot.
- Suppose that:
  - c is the number of crossings;
  - s is the number of Seifert circles.
- Then its characteristic is given by

$$\chi = v - e + f = 2c - (2c + c) + s = s - c.$$

- On the other hand, we know that a surface of genus g, with b boundary components, has characteristic

$$\chi = 2 - 2g - b.$$

- Thus, for our Seifert surface, we get

$$s - c = 2 - 2g - 1 \ \Rightarrow \ g = \frac{1 + c - s}{2}.$$

Different Surfaces from the Same Knot

- Notice that Seifert's algorithm can be used to generate lots of different surfaces for the same knot.
- We could alter the projection of the knot and then obtain a surface that at least looks different.

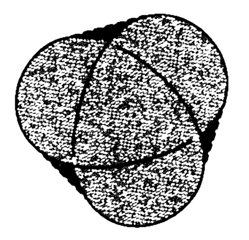
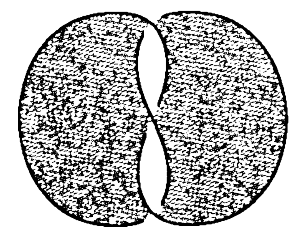
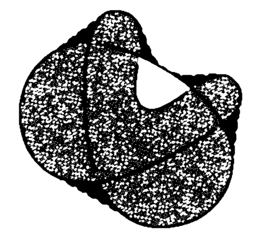

Seifert Surfaces for a Knot; Genus of a Knot

- Given a knot K, a **Seifert surface for** K is an orientable surface with one boundary component such that the boundary component of the surface is the knot K.
- We have described one way to obtain a Seifert surface for a knot.
- However, there may be other Seifert surfaces for the same knot.
- We define the **genus** of a knot to be the least genus of any Seifert surface for that knot.

Genus of the Unknot

- The unknot bounds a disk.
- When we cap off the disk, we get a sphere, which has genus 0.

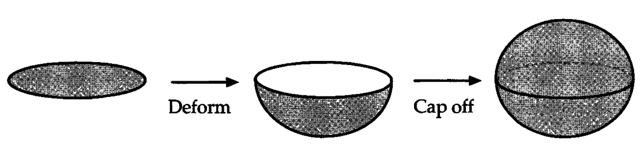

- Therefore the unknot has genus 0.
- Note that the unknot is the only knot with genus 0.

Genus of the Figure-Eight Knot

- Consider the figure-eight knot.
- By Seifert's algorithm, we obtain a Seifert surface with genus 1.

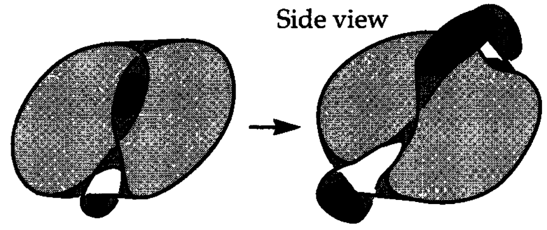

- Since the figure-eight knot is not trivial, it cannot bound a surface of genus 0.
- So 1 is the least genus of a Seifert surface for the figure-eight knot.
- Thus, the genus of the figure-eight knot is 1.

Calculating the Genus of an Alternating Knot

Theorem

Applying Seifert's algorithm to an alternating projection of an alternating knot or link yields a Seifert surface of minimal genus.

- We use the theorem to calculate the genus of the knot in the figure.

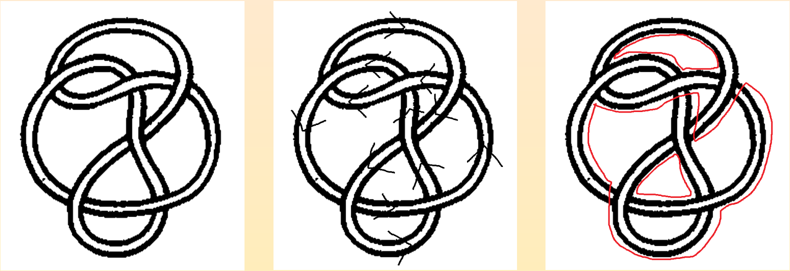

Therefore, $g = \frac{1+c-s}{2} = \frac{1+6-3}{2} = 2$.

The Genus of a Composite Knot

- Let $g(K)$ denote the genus of knot K.

Theorem

$g(J\#K) = g(J) + g(K)$.

- It's easy to see that $g(J\#K) \leq g(J) + g(K)$.
  Take a Seifert surface Q of genus $g(J)$ for J.
  Take a Seifert surface R of genus $g(K)$ for K.

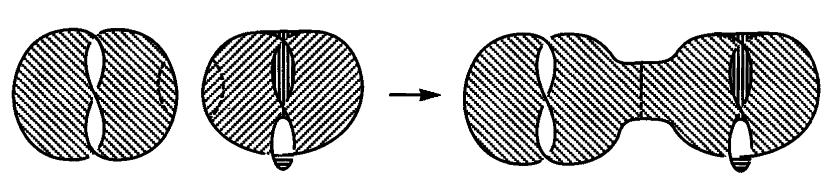

  Remove a little piece of each along their boundaries, and sew them together to obtain a Seifert surface of genus $g(J) + g(K)$ for $J\#K$.
- The proof of the converse is much more technical and we omit it.

Indecomposability of the Trivial Knot

Claim: The trivial knot cannot be the composition of two nontrivial knots.

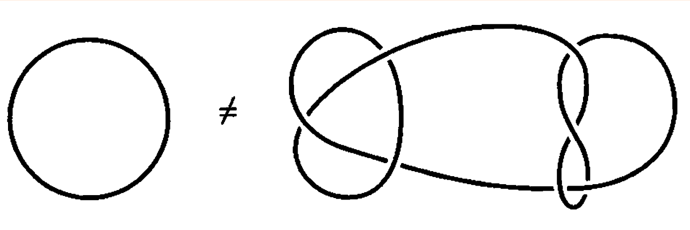

Genus 0 means the knot bounds a disk and is therefore trivial.

So any nontrivial knot has genus at least 1.

So the composition of two nontrivial knots has genus at least 2.

Therefore it cannot be the trivial knot.