ON THE FUNDAMENTAL GROUP OF SURFACES

MATTHEW ACTIPES

Abstract. This paper begins by reviewing the idea of homotopy while motivating the idea of the fundamental group. We then define free groups and the free product operation in order to state the van Kampen Theorem, allowing us to calculate the fundamental group of many different topological spaces. We focus on surfaces in particular and are able to classify them by defining the genus as the number of surgical cuts required to bring the surface into isomorphism with S^2. After defining cell complexes we are able to combine van Kampen's Theorem with the notion of genus in order to provide an explicit formula for the fundamental group of any closed, oriented surface of genus g.

Contents

1. Homotopy 1
2. Homotopy and the Fundamental Group 3
3. Free Groups 6
3.1. Free Product 7
4. Van Kampen's Theorem 8
5. Surfaces 9
5.1. Classification of Surfaces 9
6. Calculating the Fundamental Group of Surfaces 13
6.1. Cell Complexes 13
6.2. Calculation of π_1 14
Acknowledgments 15
References 15

1. Homotopy

In algebraic topology one of the basic ways of classifying spaces is by using the notion of homotopy. One reason we use this is to classify when spaces have a "similar shape," i.e. when there exist continuous functions between the spaces subject to a certain set of restrictions. To do so we shall begin with the following definitions:

Definition 1.1. For any two spaces X and Y, a *homotopy* is any family of maps $f_t : X \to Y$, $t \in I$ such that the associated map $F : X \times I \to Y$ given by $F(x, t) = f_t(x)$ is continuous.

Definition 1.2. We say that two maps f_0 and f_1 are *homotopic* if there exists a homotopy f_t connecting them. We denote this as $f_0 \simeq f_1$.

Date: August 30th, 2013.

For now consider the case of a subspace $A\subset X$, and define a retraction on A as follows:

Definition 1.3. Given a subspace $A\subset X$, a *retraction* of X onto A is a map $r : X\to X$ such that $r(X) = A$ and $r|A = \mathbb{1}$.

Our goal is to "shrink" our space X onto its subspace A by sending each point of X to a point in A via a continuous function in finite time. Note that any point already in A does not move. Formally, this is defined as a deformation retract:

Definition 1.4. A *deformation retraction* of space X onto a subspace A is a family of maps $f_t : X\to X$, where $t\in I = [0, 1]$ such that

(1) $f_0 = \mathbb{1}$,
(2) $f_1(X) = A$,
(3) $f_t|A = \mathbb{1}$ for all t.

We require that the family f_t be continuous in the sense that the associated map $X\times I\to X$, $(x, t)\to f_t(x)$ is continuous.

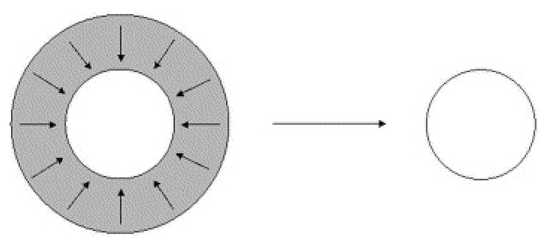

FIGURE 1. A deformation retract of a once-punctured disk onto its core circle [e]

Using this, we see that a deformation retraction of a space X onto a subspace A is a homotopy from the identity map on X to a retraction of X onto A. Now suppose $f_t : X\to X$ is the given deformation retraction of X onto A. If $r : X\to A$ denotes the resulting retraction map and $i : A\hookrightarrow X$ the standard inclusion map, then we have that $r\circ i = \mathbb{1}_A$ and $i\circ r\simeq\mathbb{1}_X$ with homotopy given by the family f_t, with $f_t|A = \mathbb{1}_A$. This can be generalized to the following:

Definition 1.5. For spaces X and Y, a map $f : X\to Y$ is called a *homotopy equivalence* if there exists a map $g : Y\to X$ such that $f\circ g\simeq\mathbb{1}_Y$ and $g\circ f\simeq\mathbb{1}_X$. The spaces X and Y are said to be *homotopy equivalent* or to have the same *homotopy type*. We denote this by $X\simeq Y$.

Remark 1.6. Notice that this is an equivalence relation:

(1) We see that $X\simeq X$ since the identity map satisfies the above definition.
(2) If $X\simeq Y$ then for the homotopy equivalence map f there exists a map g satisfying $f\circ g\simeq\mathbb{1}_Y$ and $g\circ f\simeq\mathbb{1}_X$. Thus g is also a homotopy equivalence since picking the map f gives the same consequences, showing that $Y\simeq X$.
(3) If $X\simeq Y$ via the map f and $Y\simeq Z$ via the map g, then $X\simeq Z$ via the map

$$h(x,t) = \begin{cases} f(x,2t), t\leq .5 \\ g(x,2t-1), t > .5 \end{cases}$$

This function travels each homotopy "twice as fast" in order to travel their piecewise composition in the same amount of time.

Using this result, we can prove the following claim concerning homeomorphic spaces:

Proposition 1.7. *If X is homeomorphic to Y, then $X \simeq Y$.*

Proof. If $h : X \to Y$ is a homeomorphism, then $h \circ h^{-1} = \mathbb{1}_Y$ and $h^{-1} \circ h = \mathbb{1}_X$, so that homotopy equivalence is an immediate consequence of reflexivity. □

Of course, the converse is not necessarily true. For example, the disk D^1 is homotopy equivalent to a point, but the two are not homeomorphic.
Homotopy equivalence between spaces is an intuitive way to understand spaces via "stretching" and "compressing" even when finding explicit formulas for the homotopies between them may be dificult. For example, consider figure 1. Intuitively the picture of the once-puctured disk looks like it could be squeezed together continuously until it resembles the circle. However if there was no puncture in the disk, i.e. we were only working with a copy of D^1, then we would not be able to deform the disk into the circle since there would be no logical place on the circle to send the center point. Therefore we see that $D^1 \not\simeq S^1$ but $D^1 - [\text{point}] \simeq S^1$.

2. Homotopy and the Fundamental Group

In this section we will use the idea of homotopy to motivate what we will call the fundamental group of a space X. Before proceeding, we must first introduce some new terminology.

Definition 2.1. Given a space X, a *path* is a continuous map $f : I \to X$.

Definition 2.2. Fix $x, y \in X$. A *homotopy* of paths in X is a family $f_t : I \to X$ such that the endpoints $f_t(0) = x$ and $f_t(1) = y$ are constant for each t and the associated map $H : I \times I \to X$ defined by $H(t, s) = f_t(s)$ is continuous.

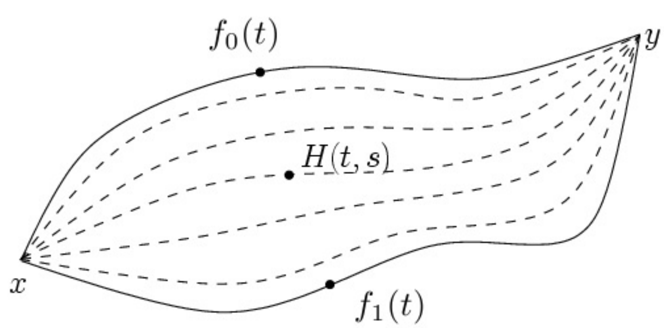

Figure 2. An example of a homotopy [f]

If two paths are connected in such a way then we say that they are *homotopic*, noted by $f_0 \simeq f_1$.

Definition 2.3. For a path f, the *homotopy class* of f, denoted $[f]$, is the equivalence class of f under the equivalence relation of homotopy.

To motivate the fundamental group, we will look at the paths where the starting point and ending point are the same. In this case, we call these paths *loops* fixed at a *basepoint* $x_0 \in X$.

Definition 2.4 (Fundamental Group)**.** The set of all homotopy classes $[f]$ of loops $f : I \to X$ at a basepoint x_0 is denoted by $\pi_1(X, x_0)$.

Remark 2.5. Notice that $\pi_1(X, x_0)$ is a group with respect to the product $[f][g]=[f \cdot g]$ (given without proof).

Since we are interested in classifying spaces X, we will next compare $\pi_1(X, x_0)$ to $\pi_1(X, x_1)$ in order to determine when our choice of basepoint is arbitrary. Let $g : I \to X$ be a path from x_0 to x_1 and let $g^*(t) = g(1-t)$ be the inverse path from x_1 back to x_0. Then to each loop f centered at x_1 we know that $g \cdot f \cdot g^*$ is a loop based at x_0, leading to our next proposition:

Proposition 2.6. *The map* $F_g : \pi_1(X, x_1) \to \pi_1(X, x_0)$ *given by* $F_g[f] = [g \cdot f \cdot g^*]$ *is an isomorphism.*

Proof. First, F_g is well defined by our construction above. Now $F_g[f \cdot h] = [g \cdot f \cdot h \cdot g^*] = [g \cdot f \cdot g^* \cdot g \cdot h \cdot g^*] = F_g[f]F_g[h]$. Hence F_g is a homomorphism. Also, F_g is an isomorphism with inverse F_{g^*} since $F_g F_{g^*}[f] = F_g[g^* \cdot f \cdot g] = [g \cdot g^* \cdot f \cdot g \cdot g^*] = [f]$, and $F_{g^*} F_g[f] = F_{g^*}[g \cdot f \cdot g^*] = [g^* \cdot g \cdot f \cdot g^* \cdot g] = [f]$. □

Remark 2.7. The only thing we assumed in this proof is that x_0 and x_1 lie in the same path component of X. Thus we see that if X is path-connected, the group $\pi_1(X, x_0)$ is independent of the choice of basepoint x_0 up to isomorphism. In cases such as this we can simply denote the fundamental group by $\pi_1(X)$.

Example 2.8. Consider any loop in $\mathbb{R}^n$. This loop can be continuously deformed into a loop traveling arbitrarily close to its basepoint, meaning that there is only one homotopy class of loops. In this case we say that $\mathbb{R}^n$ has trivial fundamental group, denoted by $\pi_1(\mathbb{R}^n) = 0$. We say that a space is *simply-connected* if it is path-connected and has trivial fundamental group.

Example 2.9. Now let's try to intuitively understand what $\pi_1(S^1)$ is without any hard calculations. First let's impose an orientation on our circle: we'll denote the positive direction as loops traveling clockwise. Then each homotopy class consists of all loops that travel around the circle a given number of times. If one loop travels around the circle a times and another loop travels around the circle b times, then their product travels the circle $a + b$ times. Therefore the set of all possible loops around the circle is isomorphic to the group of integers with group operation of addition, meaning $\pi_1(S^1) \cong \mathbb{Z}$.

We'd also like to know how to find the fundamental group of a product space. Thankfully the relation is very intuitive.

Proposition 2.10. *If* X *and* Y *are path-connected, then* $\pi_1(X \times Y)$ *is isomorphic to* $\pi_1(X) \times \pi_1(Y)$*.*

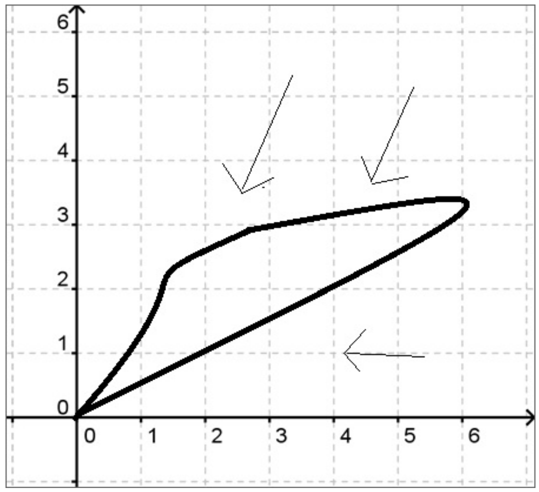

FIGURE 3. Any loop in $\mathbb{R}^2$ is equivalent to the identity map [j]

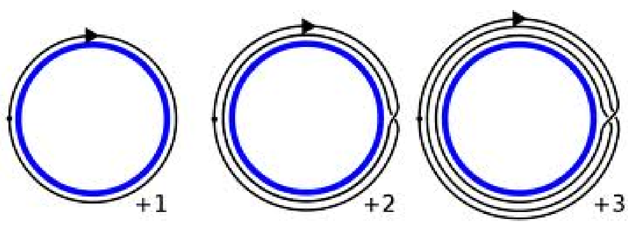

FIGURE 4. Intuition behind $\pi_1(S^1)$ [g]

Proof. A map $f : Z \to X \times Y$ is continuous if and only if the maps $g : Z \to X$ and $h : Z \to Y$ defined by $f(z) = (g(z), h(z))$ are continuous. Therefore a loop f in $X \times Y$ at a basepoint (x_0, y_0) is equivalent to a pair of loops g in X and h in Y based at x_0 and y_0 respectively. Also, a homotopy f_t of a loop in $X \times Y$ is equivalent to apair of homotopies g_t in X and h_t in Y. This gives a bijection $[f] \to ([g], [h])$, where $\pi_1(X \times Y, (x_0, y_0))$ is isomorphic to $\pi_1(X, x_0) \times \pi_1(Y, y_0)$. This is a group homomorphism, so it is an isomorphism as well, completing the proof. □

Example 2.11. This proposition gives us a way to calculate the fundamental group of many different spaces. For instance, the torus T^2 is represented by the product space $S^1 \times S^1$. We know that $\pi_1(S^1) = \mathbb{Z}$, so by proposition 2.10 we have

$$\pi_1(T^2) = \pi_1(S^1 \times S^1) \cong \mathbb{Z} \times \mathbb{Z} = \mathbb{Z}^2.$$

In fact, when we notice that the n-torus T^n is a product space of n copies of S^1, we see that, in general,

$$\pi_1(T^n) \cong \mathbb{Z}^n.$$

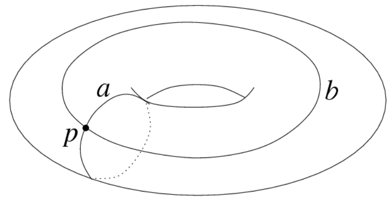

Figure 5. A torus decomposed to its product space $S^1{\times}S^1$ (denoted by generators a and b). [h]

3. Free Groups

Now that we have spent some time working with the fundamental group, we will establish the basics behind a particular type of group that we will come in frequent contact with.

Definition 3.1. A group G is *free* if there exists a subset $T{\subset}G$ such that every element of G can be written uniquely as a product of finitely many elements of T along with their inverses. We call T the *free generating set* of G, and we denote our resulting free group by F_T.

Remark 3.2. Notice that F_T is a group under the operation of concatenation, followed by reduction if necessary. Define T^{-1} such that for every $t{\in}T$ there exists $t^{-1}{\in}T^{-1}$ and let $S = T{\cup}T^{-1}$. A *word* in T is any written product of elements of S. A word is called *reduced* if no element of the word is directly next to its inverse on either side. For example, the word $a^2b^{-1}a^{-1}a$ is not reduced but ab^2ab^{-1} is reduced. Then the free group F_T is defined to be the group of all reduced words in T.

Example 3.3. The group $(\mathbb{Z}, +)$ is free with generating set $T = [1]$. This is because any integer $z{\in}\mathbb{Z}$ is equal to a finite summation

$$\sum_{i=1}^{m}(+1) + \sum_{k=1}^{n}(-1),$$

where $m, n{\in}\mathbb{N}$ and $m - n = z$.

We would like to find a way to characterize free groups. Given a generating set T of a group H, we'd like to be able to categorize the free group F_T by functions $f : T{\to}G$ for any group G. We can do so using the universal property of free groups.

Definition 3.4 (Universal Property)**.** Let $T{\subset}H$ generate the free group F_T. For any group G and any function $f : T{\to}G$, there exists a unique homomorphism $\varphi : F_T{\to}G$ that makes the following diagram commute:

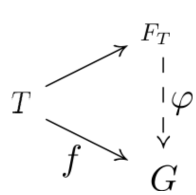

This shows that homomorphisms between F_T and G are in one-to-one correspondence with functions between T and G, giving us the ability to characterize free groups up to isomorphism.
We're also interested in a way to classify free groups based on their generating sets. We can do so using the following definition.

Definition 3.5. The *rank* of a free group F_T is equal to the cardinality of the generating set T.

Proposition 3.6. *Two free froups F_S and F_T are isomorphic if and only if their generating sets S and T have the same cardinality.*

Using this with the universal property shows that for every $n \in \mathbb{N}$ there is exactly one free group of rank n, up to isomorphism.

3.1. **Free Product.** We will now introduce a way for free groups to interact with each other, but we must first state the following definition:

Definition 3.7. Let G and H be groups. A *word* in G and H is a product of the form $a_1a_2a_3 \cdots a_n$ where $a_i \in G$ or $a_i \in H$. If a_j and a_{j+1} are from the same group then perform the group operation on them, removing all identity elements $\mathbb{1}_G$ and $\mathbb{1}_H$ in the process. This gives a reduced word in G and H of the form $g_1h_1g_2h_2 \cdots g_kh_k$.

With this in mind, we are going to define an operation between groups, denoted $G * H$, creating a group containing the reduced words in G and H under the group operation of concatenation and reduction.

Definition 3.8. Let T_i be the set of generators for group i and R_i be the set of relations, so that $G =< T_G|R_G >$ and $H =< T_H|R_H >$. Then the *free product* of G and H is defined by

$$G * H =< T_G \cup T_H | R_G \cup R_H > .$$

Remark 3.9. One property of the free product is that a pair of homomorphisms $\phi_G : G \to F$ and $\phi_H : H \to F$ extends uniquely to a homomorphism $\phi : G * H \to F$.

Remark 3.10. There are no relations on free groups, so the free product between free groups is always itself a free group. Hence for $m, n \in \mathbb{N}$, we have that

$$F_m * F_n \cong F_{m+n},$$

where F_n denotes the free group on n generators.

Notice that all of these definitions can apply to more than two groups: $*_\alpha G_\alpha$ refers to the free product of a collection of groups G_α defined just as one would expect. So a collection of homomorphisms $\phi_\alpha : G_\alpha \to H$ extends uniquely to a homomorphism $\phi : *_\alpha G_\alpha \to H$ with its action determined on the basis, i.e. $\phi(g_1g_1 \cdots g_n)$ with $g_i \in G_{\alpha_i}$ is equal to $\phi_{\alpha_1}(g_1) \cdots \phi_{\alpha_n}(g_n)$. This generalizes to the universal property of the free product.

Definition 3.11 (Universal Property)**.** The *free product* of a family of groups $[B_\alpha]$ is a group B for which there exist homomorphisms $\phi_\alpha : B_\alpha \to B$ such that for any group G and any family of homomorphisms $f_\alpha : B_\alpha \to G$ there exists a unique homomorphism $f : B \to G$ such that $f \circ \phi_\alpha = f_\alpha$ for all α.

4. Van Kampen's Theorem

The tools we have developed so far now allow us to state the van Kampen Theorem, which is helpful in calculating the fundamental group of a space composed of spaces that are familiar to us. To motivate this, suppose we have a space $X = \bigcap_\alpha B_\alpha$ where each B_α is open and contains some basepoint $x_0 \in X$. We know by properties of the free product that the homomorphisms $j_\alpha : \pi_1(B_\alpha) \to \pi_1(X)$ induced by the inclusions $B_\alpha \hookrightarrow X$ extend to a homomorphism $\phi : *_\alpha \pi_1(B_\alpha) \to \pi_1(X)$. Now if $i_{\alpha\beta} : \pi_1(B_\alpha \cup B_\beta) \to \pi_1(B_\alpha)$ is the homomorphism induced by the inclusion $B_\alpha \cap B_\beta \hookrightarrow B_\alpha$, then this combined with the inclusion $B_\alpha \cap B_\beta \hookrightarrow X$ gives us that $j_\alpha i_{\alpha\beta} = j_\beta i_{\beta\alpha}$. This shows us that the kernel of ϕ contains all elements of the form $i_{\alpha\beta}(r) i_{\beta\alpha}(r)^{-1}$ for $r \in \pi_1(B_\alpha \cap B_\beta)$, meaning that the kernel of ϕ is generally non-trivial. Van Kampen's Theorem combines all this to give an explcit way to calculate the fundamental group of such spaces.

Theorem 4.1 (The van Kampen Theorem)**.** *If X is the union of path-connected open sets B_α each containing a basepoint $x_0 \in X$ and if each intersection $B_\alpha \cap B_\beta$ is path-connected, then the homomorphism $\phi : *_\alpha \pi_1(B_\alpha) \to \pi_1(X)$ is surjective. Furthermore, if each intersection $B_\alpha \cap B_\beta \cap B_\gamma$ is path-connected, then the kernel of ϕ is the normal subgroup N generated by all elements of the form $i_{\alpha\beta}(r) i_{\beta\alpha}(r)^{-1}$, meaning that ϕ induces an isomorphism $\pi_1(X) \approx *_\alpha \pi_1(B_\alpha)/N$.*

Let's compute some examples to better our understanding of this theorem.

Example 4.2. We want to compute $\pi_1(\bigvee_\alpha S^1_\alpha)$. From van Kampen's Theorem we know that this is isomorphic to $*_\alpha \pi_1(S^1)/N$, which by our previous calculation is equal to $*_\alpha \mathbb{Z}/N$. Now we need to compute N. Remember that N is equal to the kernel of the homomorphism ϕ, so if we define a_{α_i} to be the positively-oriented generator for the circle $S^1_{\alpha_i}$, then elements of the form $a_{\alpha_i} a_{\alpha_i}^{-1} = \mathbb{1}_{S^1_{\alpha_i}}$ are in the kernel of ϕ. Therefore the kernel consists only of one point since the wedge sum gives the relation that the base point of each copy of S^1 is equivalent. Therefore we see that $\pi_1(\bigvee_\alpha S^1_\alpha) = *_\alpha \mathbb{Z}$.

Example 4.3. Now let's try a less-trivial example: computing the fundamental group of the klein bottle. Below is a polygonal representation of the klein bottle. To apply van Kampen's Theorem we will cut it into two strips, the blue and pink strips in figure 6 (Note that the pink area is only one strip since we can connect the top and bottom edge to each other). Then the yellow strips represent their intersection. Notice that the pink and blue strips give representations of the mobius strip, which has a fundamental group equal to $\mathbb{Z}$ since the mobius strip can be homotoped to a circle. Therefore we have that $\pi_1(K) \approx \mathbb{Z} * \mathbb{Z}/N$. To find N, let's look at our intersection. It also gives a copy of the mobius strip, which needs to be traversed twice in order to get the identity element. So if a and b are our generators for our two copies of $\mathbb{Z}$, we need $i_{\alpha\beta}(r) i_{\beta\alpha}(r)^{-1} = a^2 b^{-2} = \mathbb{1}$, or $a^2 = b^2$. Therefore, we arrive at the conclusion that $\pi_1(K) = <\mathbb{Z} * \mathbb{Z} | a^2 = b^2> = <a, b | a^2 = b^2>$.

Now that we are equipped with the tools we need, we will switch our discussion of general spaces X to that of surfaces. Van Kampen's Theorem will be used later in showing how to compute the fundamental group of any closed, orientable surface.

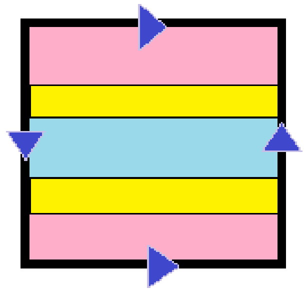

FIGURE 6. A representation of the klein bottle [j]

5. Surfaces

For the purposes of this paper we will define a *surface* as a two-dimensional topological manifold. Intuitively this means that every point on a surface is locally homeomorphic to $\mathbb{R}^2$. For now, we will define the *genus* of a surface to be the number of punctures throughout the surface.

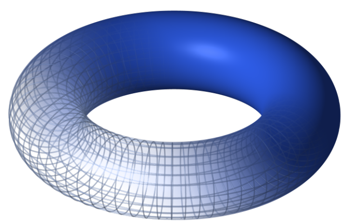

FIGURE 7. An example of a genus-1 surface (the torus) [a]

5.1. **Classification of Surfaces.** The purpose of this section is to motivate a way to quantify surfaces. One natural way of doing so is to look at the components that make up surfaces and to extraxt information from them. This leads into our first theorem:

Theorem 5.1. *Every Surface is Triangulable.*

The intuition behind this result is easily realizable if we can relate surfaces down to the 2-dimensional level. Consider a square with boundary $aba^{-1}b^{-1}$, where a^{-1} and b^{-1} are the oppositely-oriented sides across from a and b. From here we can obtain a cylinder with two copies of S^1 as its boundary by gluing one pair of oppositely-oriented sides together. Then we glue the two copies of S^1 together (Note that they are oppositely-oriented by design) to obtain a torus. Notice that if we were to start with the surface it would take two cuts along a and b in order to recover our original polygon.
We can generalize this process to all closed surfaces:

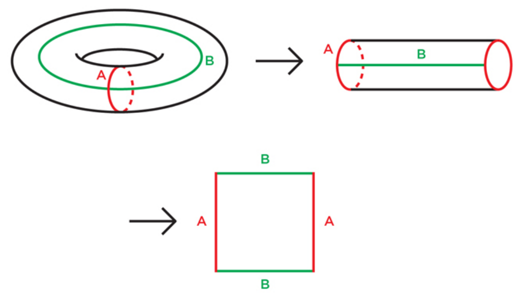

Figure 8. Polygonal construction of the torus [c]

Lemma 5.2. *Any orientable surface of genus g can be cut open using 2g cuts and can be represented by a regular 4g-gon with boundary*

$$a_1b_1a_1^{-1}b_1^{-1}a_2b_2a_2^{-1}b_2^{-1}\cdots a_gb_ga_g^{-1}b_g^{-1}.$$

Remark 5.3. Note that any orientable surface of genus 0 is isomorphic to S^2, meaning that it can be represented by an "empty polygon" with boundary aa^{-1}.

One way to visualize this result is to notice that a genus-g surface is homeomorphic to a glued collection of g toruses. This surface would require 2 cuts to open each torus completely and each resulting polygon would have 4 sides, thus leading to a 4g-gon and a total of 2g cuts. Using a well-known result in geometry that any polygon can be triangulated using a finite number of cuts, we see that orientable surfaces can be triangulated as well. With this in mind we are now able to quantify a surface based on its triangulation:

Definition 5.4. For any triangulation T of a surface S, the *Euler Characteristic* of S, denoted $\chi(S)$, is defined by

$$\chi(S) = Vertices(T) - Edges(T) + Faces(T).$$

Although we will leave it without proof, using homology one can prove that for a surface S, $\chi(S)$ is invariant under *any* triangulation of S. Therefore we see that if $T \simeq S$, then $\chi(S) = \chi(T)$.

Our next goal will be to introduce a way to classify surfaces. To do this we must find an upper bound on the Euler Characteristic which will require the use of some basic graph theory (we know such a bound exists since our surface S is compact).

Definition 5.5. A *graph* $G = (V, E)$ is a collection of vertices and edges between these vertices. Every edge must connect two vertices, but it is not necessary that every vertex be connected to an edge. A *connected graph* is a graph in which for every vertex there exist paths connecting every other vertex to the initial vertex. A *tree* P is a graph on which any two vertices are connected by at most one edge, i.e. a connected graph with no cycles. A *spanning tree* K of a graph G is a tree which contains every vertex of G and contains the maximum number of edges such that no cycles are found in K.

With this in mind, we will be able to bound the euler characteristic by relating surfaces to their triangulations. But first we will need the following result:

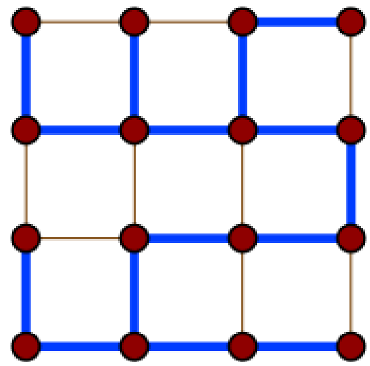

FIGURE 9. A spanning tree of the 4 by 4 grid [d]

Lemma 5.6. *Let G be a finite connected graph. Then $\chi(G) \leq 1$ with equality if and only if G is a tree.*

Proof. First notice that for a graph G, $\chi(G) = V(G) - E(G)$. Suppose G contains n vertices. G is connected, so there must be a path connecting any two vertices of G. This implies that there must be at least $n-1$ edges in G since any less would leave a vertex with no edges leading into it, contradicting the connectedness of G. Thus $\chi(G)$ is bounded above by

$$\chi(G) = V(G) - E(G) \leq n - (n-1) = 1$$

for any finite connected graph G.
Now suppose $\chi(G) = 1$, where G is a finite connected graph. Then we know that there is one more vertex than edge in G. G is connected, so to show that G is a tree we must show that it contains no cycles. Suppose it does. Then we can create a new graph G' by removing one edge from this cycle while still leaving the graph connected, meaning that $\chi(G') = n - (n-2) = 2$, contradicting our earlier proof that $\chi(G') \leq 1$. Therefore G is connected and contains no cycles, thus implying that G is a tree.
Now suppose we have a finite connected graph G that is a tree. We will prove by induction that $\chi(G) = 1$. Suppose G has 2 vertices. Then the addition of one edge connecting the two vertices would result in a tree, where $\chi(G) = 2 - 1 = 1$. Now suppose G is a tree with n vertices and $n-1$ edges. If we add one more vertex to G then it is no longer connected. However, if we add an edge connecting the new vertex to any other vertex of G, then G is connected and is a tree since the new edge does not create any cycles in G. Now $\chi(G) = (n+1) - n = 1$, so by induction $\chi(G) = 1$ for any tree G. □

Theorem 5.7. *For any surface S, $\chi(S) \leq 2$.*

Proof. Let T be a triangulation of S. T is a graph, so pick K to be a spanning tree in T. Now let $R \subset T$ be a connected graph that has vertices on every face of T and every edge in T that is not in K. Then we have

$$\begin{aligned}\chi(S) &= V(T) - E(T) + F(T)\\ &= V(K) - (E(K) + E(R)) + V(R)\\ &= (V(K) - E(K)) + (V(R) - E(R))\\ &= \chi(K) + \chi(R)\\ &= 1 + \chi(R)\\ &\leq 2,\end{aligned}$$

since K is a tree and R is a connected graph. □

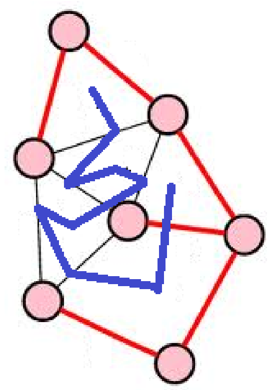

FIGURE 10. Example triangulation. R is in blue, K is in red [i]

Remark 5.8. Notice that $\chi(S) = 2$ only when R is a tree. This implies that there are no cycles in R, meaning there does not exist a loop $\gamma : S^1 \to S$ that does not separate S. Therefore, we see that $S \cong S^2$ by the Jordan Separation Theorem.

We've now seen how surfaces without boundary can be classified by their Euler characteristic. However, for more complex surfaces it is often undesirable to find a triangulation of the surface and actually compute χ. Our next theorem will give a much simpler way to calculate χ that depends only on the genus of the surface. But first we must state the following lemma:

Lemma 5.9. *If M and N are topological spaces, then*

$$\chi(M \amalg N) = \chi(M) + \chi(N).$$

Proof. This lemma follows from the fact that their triangulations are disjoint, meaning that $\chi(M \amalg N) = V_M - E_M + F_M + V_N - E_N + F_N = \chi(M) + \chi(N)$. □

Our next theorem gives us the tool we need to classify surfaces by formally defining the notion of *genus*.

Theorem 5.10 (Classifaction of Surfaces)**.** *For any closed, orientable surface S there exists $g \in \mathbb{N}$ such that $\chi(S) = 2 - 2g$.*

Proof. We will prove this theorem using a procedure called topological surgery. Suppose we have a surface N. If $\chi(N) = 2$, then we know that $N \cong S^2$ and we define it to have genus 0 by Remark 5.8. If $\chi(N) \neq 2$, then N is not isomorphic to S^2, so there exists a loop γ that does not separate N. We can cut the surface along this loop to create two boundary components in N, call them γ_1 and γ_2.

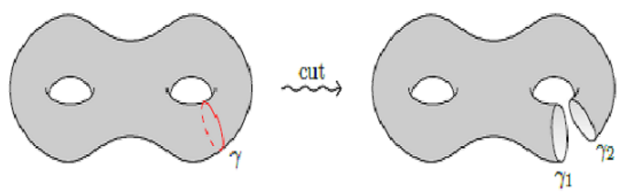

Figure 11. Surgery on a surface of genus 2 [b]

Next we glue one copy of D^2 on to each boundary component, making sure that they are properly oriented. Call this new surface N'. Notice that N' is closed and has no boundary components. Furthermore,

$$\begin{aligned}\chi(N') &= \chi(N \amalg D^2 \amalg D^2)\\ &= \chi(N) + \chi(D^2) + \chi(D^2)\\ &= \chi(N) + 2.\end{aligned}$$

Here, we repeat the previous process: If $\chi(N') = 2$ then we terminate the process because $N' \cong S^2$. In this case we say that N has genus 1. If $\chi(N') \neq 2$ then we know that we can repeat the process over again, arriving at a surface N'' such that $\chi(N'') = \chi(N') + 2 = \chi(N) + 4$. We keep doing this until we find a surface N^k such that $\chi(N^k) = 2$. Then the process terminates and we say that N has genus k. Thus we have that

$$\chi(N) = \chi(N^k) - 2k = 2 - 2g,$$

which is our desired result. □

This shows that the Euler Characteristic of any closed, oriented surface is uniquely determined by its genus, which is defined above as the number of surgical cuts required to make the surface isomorphic to the sphere.

6. Calculating the Fundamental Group of Surfaces

In this section we will combine the power of van Kampen's Theorem with the notion of genus in order to give an explicit formula for the fundamental group of a closed, oriented surface. In order to do so, we must first have a basic understanding of how surfaces are built.

6.1. **Cell Complexes.** For the purposes of this paper we will define a *2-cell* as an open disk, a *1-cell* as an edge and a *0-cell* as a point. For example, in any triangulation of a surface the vertices are the 0-cells, the edges are the 1-cells and the faces are the 2-cells. The *n-skeleton* of a surface is formed inductively by attaching a collection of n-cells to the $(n-1)$-skeleton, where the 0-skeleton is a discrete collection of 0-cells. We will take for granted the fact that attaching cells of dimension greater than two to a surface does not affect its fundamental group, indicating that these three cell-types are the only ones of interest to us.

Before proving our main theorem, we must first motivate a supporting lemma. Suppose X is a path-connected space. We can attach a collection of 2-cells e_α^2 to X via maps $\phi_\alpha : S^1 \to X$ resulting in a space Y. Choose a basepoint s_0 for S^1 and let ϕ_α denote the loop in Y based at s_0. Now choose a basepoint $x_0 \in X$ and a path φ_α in X from x_0 to s_0 for every α. Then the loop $\varphi_\alpha \phi_\alpha \varphi_\alpha^*$ is a loop based at x_0 for every α. As designed, this loop homotopes trivially in Y after the cell e_α^2 is attached. Therefore, the normal subgroup $N \subset \pi_1(X, x_0)$ generated by the loops $\varphi_\alpha \phi_\alpha \varphi_\alpha^*$ varying over α is contained in the kernel of the map $\pi_1(X, x_0) \to \pi_1(Y, x_0)$ induced by the inclusion $X \hookrightarrow Y$. This results in the following lemma:

Lemma 6.1. *The inclusion $X \hookrightarrow Y$ induces a surjection $\pi_1(X, x_0) \to \pi_1(Y, x_0)$ whose kernel is N. Thus $\pi_1(Y) \approx \pi_1(X)/N$.*

This follows from the van Kampen Theorem.

6.2. **Calculation of π_1.** We begin with a quick definition:

Definition 6.2. Let a, b be generators of a group. The *commutator* of a and b, denoted $[a, b]$, is given by

$$[a, b] = aba^{-1}b^{-1}.$$

Theorem 6.3 (Calculation of π_1)**.** *Let S be a closed, oriented surface of genus g. Then*

$$\pi_1(S) = < a_1, b_1, \cdots, a_g, b_g | [a_1, b_1] \cdots [a_g, b_g] = 1 > .$$

Proof. Consider a wedge sum $\bigvee_{i=1}^{2g} S_i^1$ of $2g$ circles labeled $a_1, a_2, \cdots, a_g, b_1, b_2, \cdots, b_g$. Notice that this is the 1-skeleton of S. By Lemma 5.2 we know that S can be represented by a regular $4g$-gon with boundary $a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} = [a_1, b_1] \cdots [a_g, b_g]$. We can achieve this polygonal representation of S by attaching a 2-cell to our 1-skeleton, resulting in a closed, oriented surface of genus g.

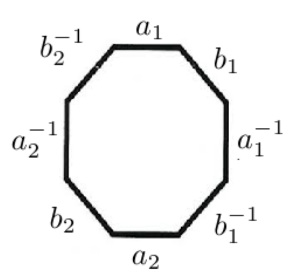

FIGURE 12. 2-skeleton of a closed, oriented surface of genus 2. [j]

We saw in Example 4.2 that π_1 of the wedge sum of $2g$ circles is the free group on $2g$ generators. Therefore by Lemma 6.1 we see that $\pi_1(S) = \pi_1(\bigvee_{i=1}^{2g} S_i^1)/N$, where N is the normal subgroup generated by the loop $a_1 b_1 a_1^{-1} b_1^{-1} \cdots a_g b_g a_g^{-1} b_g^{-1} = [a_1, b_1] \cdots [a_g, b_g]$. Furthermore, we know that this loop is nullhomotopic by the addition of the attached 2-cell. Therefore, we get that

$$\pi_1(S) = < a_1, b_1, \cdots, a_g, b_g | [a_1, b_1] \cdots [a_g, b_g] = 1 >,$$

which is our desired result. ∎

Acknowledgments. I'd like to thank my mentor, Elden Elmanto, for all he has done for me throughout this program. His support was crucial to my understanding of the material and his enthusiasm for the material kept me motivated. I'd also like to thank Peter May for hosting the REU and Alden Walker for answering my questions about surfaces and sparking my interest in this topic.

References

[1] Allen Hatcher Algebraic Topology Cambridge University Press 2001
[2] R. Andrew Kreek Surface Groups VIGRE 2007
[3] Razvan Gelca Lecture Notes Texas Tech University
[4] Anatole Katok Elementary Homotopy Theory Penn State University
[5] Benson Farb Lecture Notes - Math 317 The University of Chicago
[6] Kiyoshi Igusa Lecture Notes - Math 101b Brandeis University
[a]http : //upload.wikimedia.org/wikipedia/commons/1/17/Torus.png
[b]http : //math.uchicago.edu/ chonoles/expository − notes/courses/2012/317/math317notes.pdf
[c]https : //www.simonsfoundation.org/wp − content/uploads/2012/09/Hake$_{s}$ized$_{F}$igure01.jpg
[d]http : //upload.wikimedia.org/wikipedia/commons/thumb/d/d4/4x4$_{g}$rid$_{s}$panning$_{t}$ree.svg/200px − 4x4$_{g}$rid$_{s}$panning$_{t}$ree.svg.p
[e]http : //origin − ars.els − cdn.com/content/image/1 − s2.0 − S1077314206001615 − gr1.jpg
[f]http : //upload.wikimedia.org/wikipedia/commons/8/81/Homotopy$_{c}$urves.png
[g]http : //upload.wikimedia.org/wikipedia/en/3/3d/Fundamental$_{g}$roup$_{o}$f$_{t}$he$_{c}$ircle.svg
[h]http : //upload.wikimedia.org/wikipedia/commons/a/a4/Fundamental$_{g}$roup$_{t}$orus2.png
[i]http : //upload.wikimedia.org/wikipedia/commons/f/fa/7n$_{g}$raph$_{w}$ith$_{m}$inimal$_{s}$panning$_{t}$ree.svg
[j]MatthewActipes