INTRODUCTION TO THE HOMOLOGY GROUPS OF COMPLEXES

RACHEL CARANDANG

Abstract. This paper provides an overview of the homology groups of a 2-dimensional complex. It then demonstrates a proof of the Invariance Theorem and applies the results to the Euler characteristic and map coloring.

Contents

1. Introduction

The algebraic properties of the vertexes, edges, and faces of a 2-dimensional complex are crucial to gaining information about the topological space containing that complex. This paper is an exposition following Michael Henle's *A Combinatorial Introduction to Topology* and covers the most elementary possible starting point of algebraic topology [1]. In this paper, mod 2 homology is used to obtain as much information as possible about triangulated surfaces, using no more sophisticated language or tools.

In Section 2, we introduce the relationships between simplexes, cycles, and boundaries in homology groups. In Section 3, we explore how these relationships hold for a topological space in spite of the many complexes that can be defined on that space. In Section 4, we utilize the information about the invariance of homology groups to relate the number of vertexes, edges, and faces of a surface to its first homology group. Finally, in Section 5, we exploit the fact that the relationship of vertexes, edges, and faces is inherent in map coloring, and apply the Euler characteristic to determine an upper bound on the amount of colors sufficient to color a complex.

2. Homology Groups of a 2-Dimensional Complex

Definition 2.1. A **0-simplex** is a point, a **1-simplex** is an edge, and a **2-simplex** is a triangle.

Figure 2.2. A 0-simplex, 1-simplex, and 2-simplex

Definition 2.3. A **complex** is any topological space constructed from the topological identification of 0, 1, or 2-simplexes.

A special type of complex, consisting of triangles identified in a particular manner, can be used to define the notion of a surface.

Definition 2.4. A **surface** is a topological space that satisfies the following two conditions. First, the topological space must be **triangulable**, meaning it can be obtained from the identification of triangles, where any two triangles are either identified along a single edge, a single vertex, or are disjoint. Second, in addition to being triangulable, each triangle edge must be identified with exactly one other edge. Third, the triangles identified at each vertex must be arranged in a cycle, $T_1, T_2 .... T_k, T_1$ such that adjacent triangles are identified along an edge.

Note that every surface can be made into a complex, since it has a triangulation, but not every complex corresponds to a surface.

Example 2.5. A Mobius strip is not a surface, since it fails to satisfy the property that there exists a triangulation on the Mobius strip such that every triangle edge is identified with exactly one other edge. This is because a Mobius strip has a boundary, depicted as the bolded line in Figure 2.6. Therefore every triangulation will contain triangles that have an edge along this boundary, and this edge will not be identified with some other edge in the triangulation.

Figure 2.6.

Example 2.7. A torus is an example of a surface. The triangulation in Figure 2.8 below shows that a torus is triangulable, and satisfies all the properties of the definition of a surface.

Figure 2.8.

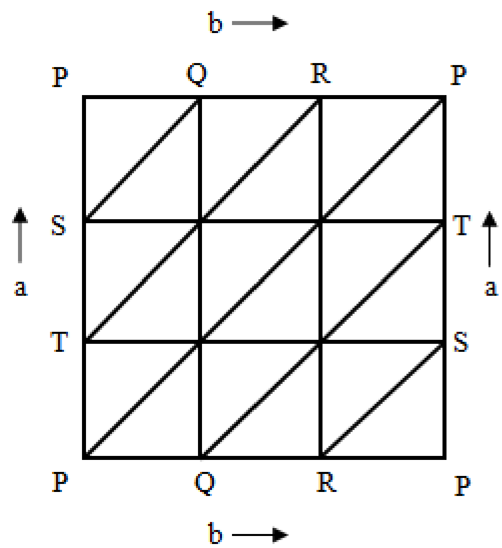

Surfaces can be represented as plane models using the definitions provided below. We will utilize the plane model of a surface throughout this paper to conceptualize the relationship between the k-chains, k-boundaries, and k-cycles of that surface, three concepts which we will define later in this section. Then, we will use k-chains, k-boundaries, and k-cycles to define homology groups.

Definition 2.9. Two edges in a planar diagram are said to be **topologically identified** when 1) each edge is assigned a direction from one endpoint to another and placed in correspondence with the unit interval so that the initial points of the edges correspond to 0 and the end points of the edges correspond to 1, and 2) the points on the edges that correspond to the same value from the unit interval are treated as a single point.

Definition 2.10. The **plane model** of a surface is a representation of that surface in a plane, where symmetrical points on similarly labeled edges indicate that those points are topologically identified on the surface in the direction indicated.

Example 2.11. The plane model of the projective plane. The arrows indicate that diametrically opposed points along a are topologically identified.

Figure 2.12.

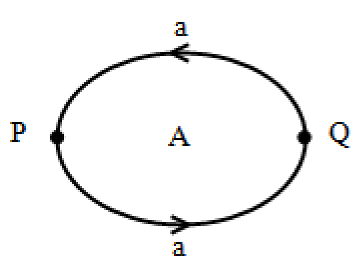

Definition 2.13. The set of **k-chains** of a complex κ, $C_k(\kappa)$, is a vector space over the field $\mathbb{F}_2 = \{0, 1\}$ of two elements. The basis for $C_k(\kappa)$ is the set of k-simplexes. In this paper, we consider k= 0, 1, 2.

Definition 2.14. Let κ be a complex. Let x be a $(k-1)$-simplex and y be a k-simplex for k= 1, 2. The **incident coefficient** of x in y is the number of times x appears in the boundary of y.

Definition 2.15. Let C be a k-chain of κ. The **boundary of C**, $\partial(C)$, is the $(k-1)$-chain consisting of the $(k-1)$-simplexes that have odd incident coefficients on the k-simplexes of C.

The above definition of boundary, which uses mod 2 addition, is motivated by the fact that in a vector space over $\mathbb{F}_2$, each element is its own inverse: $C + C = (1 + 1)C = (0)C = \emptyset$, where the first equality follows from distribution of vector multiplication over scalar addition in $\mathbb{F}_2$.

Proposition 2.16. *The boundary operator $\partial(C_k)$ is a linear transformation from $C_k(\kappa)$ to $C_{k-1}(\kappa)$ for k= 1, 2.*

Proof. Let C_1 and C_2 be k-chains. We will show $\partial(C_1) + \partial(C_2) = \partial(C_1 + C_2)$.

Let S be any $(k-1)$-simplex. Let n_1 and n_2 be the number of times S is incident on C_1 and C_2, respectively. The proof will show that S in $\partial(C_1) + \partial(C_2)$ means $n_1 + n_2$ is odd, and that S in $\partial(C_1 + C_2)$ also means $n_1 + n_2$ is odd. Since this is true for any S, we will have shown that all the simplexes of $\partial(C_1) + \partial(C_2)$ and $\partial(C_1 + C_2)$ must be the same.

Assume S is in $\partial(C_1) + \partial(C_2)$. Then S must be in $\partial(C_1)$ or $\partial(C_2)$ but not both. Since, S can only be in the boundary of a complex if its incident coefficient is odd, it follows that if is S is in $\partial(C_1)$ but not $\partial(C_2)$, n_1 must be odd and n_2 even. If S is in $\partial(C_1)$ but not $\partial(C_2)$, then n_1 must be even and n_2 odd. Thus $n_1 + n_2$ must be odd.

Suppose S is in $\partial(C_1 + C_2)$. This means S is incident an odd number of times on only simplexes in C_1 or C_2 but not both. Let n be the number of times S is incident on the simplexes shared by C_1 and C_2. Thus the incidence of S on all the simplexes of C_1 that are not shared with C_2 is $n_1 - n$. Similarly the incidence of S on all the simplexes of C_2 that are not shared with C_1 is $n_2 - n$.

So the total incidence of S on the k-simplexes of C_1 and C_2 but not both, which is precisely the incidence of S on $C_1 + C_2$, is $(n_1 - n) + (n_2 - n) = n_1 + n_2 - 2n$. If $n_1 + n_2 - 2n$ is odd, $n_1 + n_2$ must be odd. □

Definition 2.17. A **k-boundary**, C, is a k-chain such that $C = \partial(T)$, where T is some $(k + 1)$-chain. By convention, the only 2-boundary is $\emptyset$, since we have not defined k-chains for k greater than 2.

Definition 2.18. A **k-cycle**, C, is a k-chain such that $\partial(C) = \emptyset$. By convention, the boundary of a 0-chain is $\emptyset$, since we have not defined k-chains for k less than 0. So all 0-chains are 0-cycles.

Proposition 2.19. *The set of k-boundaries of a complex κ, $B_k(\kappa)$ is a subspace of $C_{k-1}(\kappa)$, and the set of k-cycles, $Z_k(\kappa)$, is a subspace of $C_k(\kappa)$.*

Proof. By definition, the set of k-boundaries $B_k(\kappa)$ is the image of $\partial : C_k \to C_{k-1}$, and the set of k-cycles $Z_k(\kappa)$ is the kernel of $\partial : C_k \to C_{k-1}$. In general, the image and kernel of a linear transformation is a subspace. □

Theorem 2.20. *Every boundary is a cycle.*

Proof. Every 0-boundary is a 0-cycle, since all 0-boundaries are 0-chains, which are 0-cycles.

Every 2-boundary is a 2-cycle, since the only 2-boundary is $\emptyset$.

Let C be a 2-simplex, or polygon. We see that $\partial(\partial(C))$ is the 0-boundary of the set of 1-chains with odd incident coefficents on C. Note that the 1-chains forming $\partial(C)$ are edges, and in a polygon each vertex is connected to two vertexes. When we count the vertexes on $\partial(C)$ edge by edge, we count each vertex twice. Thus the incident coefficent of each vertex on $(\partial(C))$ is 2, and the 0-boundary

$$\partial(\partial(C)) = \emptyset.$$

This is true for the boundary of the boundary of any 2-chain, not just a 2-simplex. This is because any 2-chain is the sum of 2-simplexes. Since the boundary operator is a linear transformation, the boundary of a boundary of any 2-chain remains the null boundary. □

Definition 2.21. Two k-chains C_1 and C_2 are **homologous**, denoted $C_1 \sim C_2$, when $C_1 + C_2$ is a k-boundary.

Proposition 2.22. *Homology is an equivalence relation.*

Proof. $C \sim C$, since $C + C = \emptyset$, which is a k-boundary.

Let C_1 and C_2 be two k-chains such that $C_1 \sim C_2$. Then $C_2 \sim C_1$, because $C_1 + C_2 = C_2 + C_1$ is a k-boundary.

Let C_1, C_2, and C_3 be three k-chains such that $C_1 \sim C_2$ and $C_2 \sim C_3$. Let A and B be the $(k+1)$-chains such that $C_1 + C_2 = \partial(A)$ and $C_2 + C_3 = \partial(B)$. Then $C1 \sim C3$, because

$$\begin{aligned} C_1 + C_3 &= C_1 + \emptyset + C_3 \\ &= C_1 + C_2 + C_2 + C_3 \\ &= \partial(A) + \partial(B) = \partial(A + B) \end{aligned}$$

since the boundary operator is a linear transformation. □

Definition 2.23. The **k-th homology group** of a complex κ, denoted by $H_k(\kappa)$, is the group of equivalence classes over k-cycles by the homology equivalence relation.

Example 2.24. We will now calculate the homology groups of the plane model of a sphere.

Figure 2.25. The plane model of the sphere, where the points and edges to be topologically identified are labeled.

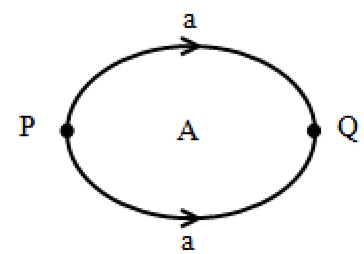

$H_0 \cong \mathbb{Z}_2$. This is because the 0-cycles of a sphere are its 0-chains, which are P, Q, and P+Q. $P \sim Q$, and $P + Q \sim \emptyset$, since $P + Q = \partial(a)$. Therefore H_0 consists of just 2 elements.

$H_1 \cong \emptyset$. The only 1-cycle of the sphere is $\emptyset$, since the only 1-chain of the sphere, a, has boundary $P + Q$.

$H_2 \cong \mathbb{Z}_2$. There is only one 2-chain, A, which is a 2-cycle since $\partial(A) = 0$. Therefore H_2 consists of just $\emptyset$ and A.

Example 2.26. We will calculate the homology groups of the torus.

Figure 2.27. The plane model of the torus.

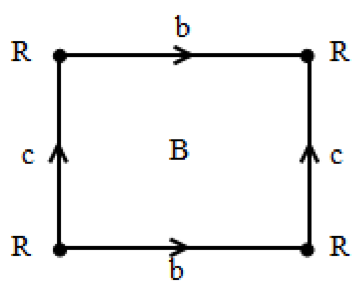

$H_0 \cong \mathbb{Z}_2$, consisting of $\emptyset$ and the only point, R.

H_1 consists of four elements. This is because every 1-chain of the torus is a 1-cycle, and no two 1-cycles are homologous. Thus H_1 consists of $\emptyset, b, c$, and $b + c$. In fact, $H_1 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$. Also, H_1 can be generated by b and c, so the dimension $h_1 = 2$.

$H_2 \cong \mathbb{Z}_2$ since there are only two 2-cycles: $\emptyset$ and B.

3. The Invariance Theorem

In this section, we will show that homology groups tell us something fundamental about a surface, and do not depend on the variety of triangulations that the surface can be made into.

At this point, it is not guaranteed that a surface has the same homology groups regardless of the choice of triangulation. To see this, consider the case when the surface can be made into two different triangulations. Since the homology groups of a complex depend on that complex, we cannot be sure that the homology groups of one triangulation are the same as those of another triangulation. The following lemma and theorems will help us resolve this issue.

Lemma 3.1. *Let κ be any complex. Let κ^+ be the complex obtained from κ by drawing a single new edge dividing a single polygon of κ into two polygons. Then the homology groups of κ and κ^+ are the same.*

Figure 3.2. The complex κ and the complex κ^+, formed by adding the edge g.

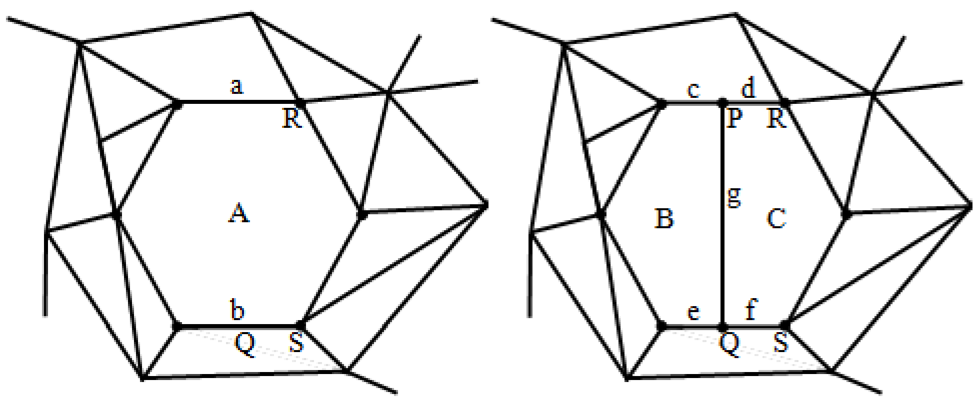

Proof. By convention, any 0-chain on a complex is a 0-cycle. In $H_0(\kappa^+)$, $P \sim R$, since P and R form a 0-boundary of d. Also $Q \sim S$, since Q and S form a 0-boundary of f. Thus any 0-cycle on κ^+ containing P or Q is homologous to a 0-cycle on κ^+ that does not contain P or Q, by the transitive property of homologous k-chains. Clearly, any 0-cycle that does not contain P or Q on κ^+ corresponds to a cycle on κ. Then any 0-cycle on κ^+ can be found to correspond to a 0-cycle on κ. Therefore there exists an isomorphism between the group of 0-cycles of κ and κ^+. $H_0(\kappa) \cong H_0(\kappa^+)$.

Let λ be a 1-cycle of κ^+. If λ contains both c and d or both e and f, and does not contain g, then λ corresponds to a 1-cycle on κ^+. This is because $c + d$ and $e + f$ correspond to a and b on κ, respectively. If λ contains c or d but not both, or contains e or f but not both, or contains g, then this is not necessarily the case.

Assume λ contains c or d but not both. Then one of the 0-chains on the boundary of λ is P. Also, all 0-chains on the boundary of λ must have even coefficents, since $\partial(\lambda) = 0$. Therefore λ must also contain g, since the incident coefficent of P must be even. If λ contains g, then one of the 0-chains on the boundary of λ is Q. Q must also have an even incident coefficent on the boundary of λ, so λ must contain either e or f but not both. Therefore if λ contains c or d but not both, it also contains e or f but not both, and g.

Let $\mu = \lambda + \partial(B)$. We will show that μ is a 1-cycle homologous to λ, which corresponds to a 1-cycle in κ. First, μ is a k-cycle because $\partial(\mu) = \partial(\lambda) + \partial(\partial(B)) = 0$, since the boundary of a k-cycle is 0 and the boundary of a boundary of a polygon is 0. Second, $\mu \sim \lambda$ because μ and λ differ by a 1-boundary. Finally, μ corresponds to a 1-cycle in λ because μ contains both c and d, or both e and f, and not g. Therefore any 1-cycle on κ^+ corresponds to a 1-cycle on κ. $H_1(\kappa) \cong H_1(\kappa^+)$.

Any 2-cycle in $H_2(\kappa^+)$ must contain both B and C or neither, since g is a 1-chain incident once on the boundary of B and once on the boundary of C, and all 1-chains on the boundary of any 2-cycle must have even incident coefficents. Thus any 2-cycle in κ^+ containing B and C corresponds to a cycle in κ containing A. $H_2(\kappa) \cong H_2(\kappa^+)$. □

Theorem 3.3. *(Invariance Theorem) Let φ be a compact, connected surface. Then the homology groups of φ are independent of the choice of triangulation.*

Proof. Take any triangulation κ of a compact, connected surface φ. Any triangulation of a surface can be reduced to the plane model of the surface, by a series of cutting and pasting operations that preserve the information about the overall topological identifications of the edges of the surface. Cutting and pasting along an edge of a polygon can be seen as drawing a new edge dividing a polygon or erasing an edge dividing a polygon. By Lemma 3.1 none of these cutting or pasting operations change the homology groups of the surface. Thus, the homology groups of a complex are the same as those of its plane model.

Furthermore, since the homology groups associated with the triangulation κ are the same as those of the plane model of φ, it follows that any other triangulation has the same homology groups as the plane model of φ. Thus any two triangulations of φ have the same homology groups. □

4. The Euler Characteristic

In this section, we will prove that if a complex κ represents a surface φ, and if V, E, and F are the number of vertexes, edges, and faces in κ, then $V - E + F$ is a constant independent of the manner in which φ is divided up to form the complex κ. This constant is called the Euler characteristic of φ, $\chi(\varphi)$.

We now recall a few definitions related to vector spaces. These definitions, along with further material regarding vector spaces, can be found in the chapter "Vector Spaces" in *A Combinatorial Introduction to Topology* [1]. Below all vector spaces are assumed to be finite dimensional.

Definition 4.1. Let A be a linearly independent subset of the vector space V over a field F that spans V. Then A is a **basis** for V.

Definition 4.2. The **dimension** of a vector space V over a field F is the number of elements in a basis for V.

Definition 4.3. Let H be a subspace of the vector space G over a field F. Then the **quotient space of** G **by** H, denoted G/H, is the set of left cosets of H.

Clearly, $H_k(\kappa)$ is a quotient space over $\mathbb{F}_2$. We see that $H_k(\kappa) = Z_k(\kappa)/B_k(\kappa)$, because the left cosets of $B_k(\kappa)$ in $Z_k(\kappa)$ are exactly the equivalence classes over k-cycles by the homology equivalence relation.

Theorem 4.4. *Let H be a vector subspace of the vector space G. Then the dimension of G/H equals the dimension of G minus the dimension of H.*

Theorem 4.5. *(Rank-nullity) Let ϕ be a linear transformation between the vector spaces G_1 and G_2. Then dim(G_1) = dim(kerϕ) + dim(imϕ)*

Example 4.6. Let the dimension of $C_k(\kappa)$ be denoted by c_k, the dimension of $Z_k(\kappa)$ be denoted by z_k, and the dimension of $B_k(\kappa)$ be denoted by b_k.

Since the boundary operator ∂_k is a linear transformation between vector spaces C_k and C_{k-1}, the dimension of $C_k(\kappa)$ can be determined for k= 0, 1, 2 using the rank-nullity theorem. We begin by calculating the dimension of both the kernel and the image of ∂_k for k= 0, 1, 2.

The kernel of ∂_k is every k-chain which has null boundary, which is the set of k-cycles of C_k. So for k= 0, 1, 2,

$$dim(\ker \partial_k) = z_k \tag{4.7}$$

The image of ∂_k is every $(k-1)$-chain that is a boundary of a k-chain, which is the set of $(k-1)$-boundaries of C_{k-1}. So for $k = 0, 1, 2$,

$$dim(im\partial_k) = b_k \tag{4.8}$$

Thus by Theorem 4.5, (4.7), and (4.8),

$$\begin{aligned} c_0 &= dim(ker\partial_0) + dim(im\partial_0) \\ &= z_0 \\ c_1 &= dim(ker\partial_1) + dim(im\partial_1) \\ &= z_1 + b_0 \\ c_2 &= dim(ker\partial_2) + dim(im\partial_2) \\ &= z_2 + b_1 \end{aligned}$$

Theorem 4.9. *Let φ be a surface. Then the Euler characteristic $\chi(\varphi)$ is independent of the choice of complex κ to divide up φ.*

Proof. Let the dimension of $H_k(\kappa)$ be denoted by h_k. By Thm 4.4 above,

$$h_0 = z_0 - b_0 \tag{4.10}$$

$$h_1 = z_1 - b_1 \tag{4.11}$$

$$h_2 = z_2 - b_2 \tag{4.12}$$

Also, $c_0 = V$, the number of vertexes. This is because a basis for the group of 0-chains is the set of 0-simplexes in κ. The dimension of C_0 is therefore the number of 0-simplexes, or the number of vertices.

By similar argument, $c_1 = E$, the number of edges, and $c_2 = F$, the number of faces.

Then

$$\begin{aligned} V - E + F &= c_0 - c_1 + c_2 \\ &= z_0 - (z_1 + b_0) + (z_2 + b_1)(\text{See Example 4.6}) \\ &= z_0 - b_0 - (z_1 - b_1) + z_2 - b_2 \\ &= h_0 - h_1 + h_2 \end{aligned}$$

Because the invariance theorem states that the homology groups of a complex are independent of the triangulation of that complex, $h_0 - h_1 + h_2$ is independent of κ. ☐

Proposition 4.13. *The value h_1, the number of elements in the equivalence classes of 1-chains of a surface, is the twice the genus of a surface, where the genus indicates how many holes the surface has.*

Example 4.14. In Example 2.26 we saw that $h_1 = 2$ for the torus, which is indeed twice the genus of the torus. The genus of the torus is 1, since the torus has one hole.

Example 4.15. In this example, we will demonstrate how the dimension of H_1, denoted h_1, relates to the number of holes in an infinite plane with a grating. In this paper, we define a grating as a decomposition of the plane into rectangles, whose sides are determined by an infinite number of horizontal and vertical lines. A grating is a complex, and we can regard the vertexes as 0-simplexes, the lines connecting the vertexes as 1-simplexes, and the rectangular faces as 2-simplexes.

If a plane has no holes, every 1-cycle divides the plane into two faces, an inside and an outside. Every 1-cycle is thus a 1-boundary of two 2-chains. In Figure 4.16, A bounds G_1 and G_2.

Figure 4.16.

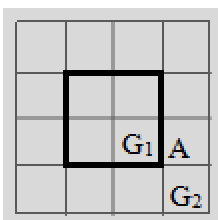

Note that G_2 is indeed a 2-chain, since it the set of an infinite number of 2-simplexes. Because every 1-cycle is a boundary, H_1 consists of $\emptyset$. Thus $h_1 = 0$ for a plane with no holes.

If a plane has one hole, then every 1-cycle in G is a 1-boundary of one 2-chain, and h_1 remains 0. In Figure 4.17, A bounds G_1 and B bounds G_2. By the same reasoning as above, h_1=0 for a plane with one hole.

Figure 4.17. Examples of 1-cycles in a plane with one hole

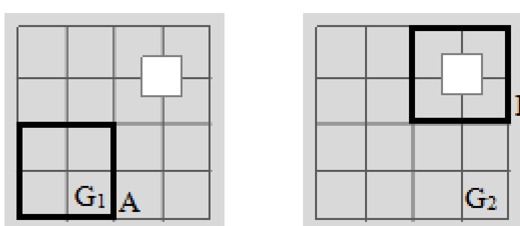

If a plane has two holes, then given a grating, some 1-cycles in G are not boundaries in G. In Figure 4.18, A does not bound anything, because both its inside and outside have holes. However, Figure 4.18 below shows that the union of B_1 and B_2 form a 1-boundary of G.

Figure 4.18. Examples of 1-cycles in a plane with two holes

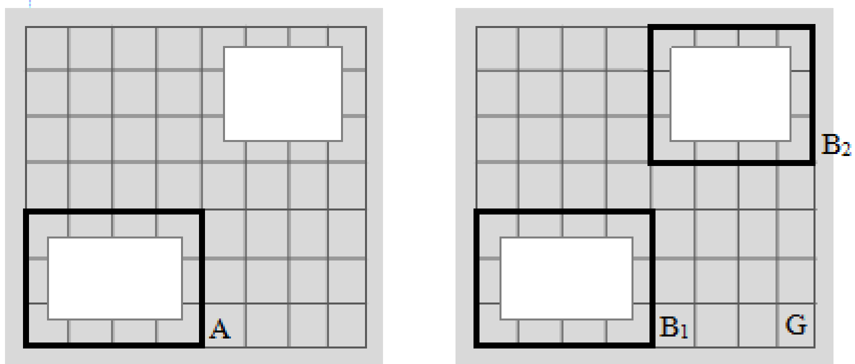

Thus in a plane with two holes, any two 1-cycles that are not boundaries are homologous to each other, since their union forms a 1-boundary. All 1-cycles that are boundaries are clearly homologous to $\emptyset$. Therefore the homology group has two elements. For a plane with two holes, $h_1 = 1$.

Theorem 4.19. *(The Euler-Poincare formula)* $\chi(\varphi) = 2 - h_1$ *for all compact, connected surfaces.*

Proof. For all connected surfaces, $h_0 = 1$. This is because all 0-cycles are vertexes, and all vertexes are homologous to each other because any two vertexes in a connected surface form the boundary of the of the 1-chain between them. So a basis for H_0 consists of just one element.

For all compact, connected surfaces without boundary, $h_2 = 1$. The only nonempty 2-cycle is the sum of all the 2-simplexes of the surface. This is the only 2-chain that has no boundary. This is because the surface is compact and connected, so each edge of any 2-simplex on the surface touches the edge of another 2-simplex. Thus every edge is counted twice for each face it appears on the boundary of, and each edge has an incident coefficent of 2. There are no edges with odd incident coefficients, and the boundary of the sum of all the 2-simplexes is $\emptyset$. So a basis for H_2 consists of just one element.

Therefore for all compact, connected surfaces

$$\begin{aligned}\chi(\varphi) &= V - E + F\\ &= h_0 - h_1 + h_2\\ &= 1 - h_1 + 1\\ &= 2 - h_1\end{aligned}$$

□

5. Map Coloring

In order to extract patterns from coloring the faces of a map, we must necessarily involve the relationship between the vertexes, edges, and faces of a map. Thus in this section, we will see an application of the Euler characteristic in finding an upper bound of the number of colors needed to color a map.

Definition 5.1. A map is a 2-dimensional complex. A map is **N-colorable** if N colors are sufficient to color its faces so that any two faces sharing an edge have different colors.

Lemma 5.2. *Let φ be a compact, connected surface. Let κ be a complex on φ with F faces, E edges, and V vertexes. Let $a = \frac{2E}{F}$. If N is a positive integer such that $a < N$ for every complex κ on φ , then any complex κ on φ is N-colorable.*

Proof. We will prove this lemma by induction on the number of faces. Let a be the average number of edges per face on a surface κ, and let $a < N$. Let k be the number of faces of κ. If $k < N$, then κ is clearly N-colorable, since every face can be painted a different color. This handles the base case of the induction.

Now, assume κ has k faces and is N-colorable. We will show that this implies that a complex κ_1 with $k+1$ faces is also N-colorable.

Note that the ratio $a = \frac{2E}{F}$ in Lemma 5.2 represents the average number of edges per face. This is because $2E$ represents the total number of edges counted face by face, since each edge connects two faces and is counted twice, and F represents the total number of faces.

We are given that the average number of edges per face, a, is less than N, so there must be at least one face with less than N edges.

Take this face and distribute it among its adjacent faces, as demonstrated in Figure 5.3. The total number of faces on κ decreases by 1, so κ has k faces. The total number of edges on each adjacent face increases by 1, since it gains 2 edges but loses 1 edge.

Figure 5.3. Removing a face by distributing it among adjacent faces.

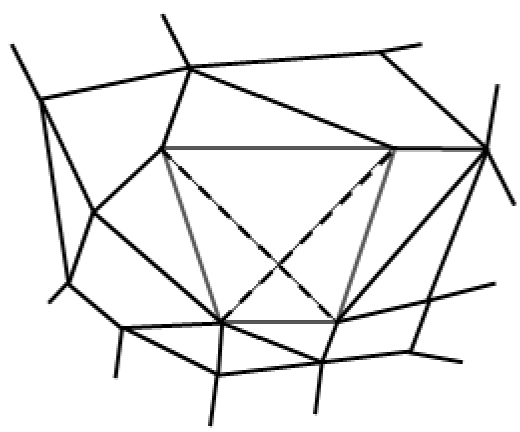

Since by the induction hypothesis we assumed any complex with k faces was N-colorable, κ^- has k faces and is N-colorable.

After all the faces of κ^- have been colored, recover the face with less than N edges by undoing the step shown Figure 5.3. Since this face has less than N adjoining faces, color this face with a color that is not present in its adjoining faces. Therefore κ^- is N-colorable. $\square$

Theorem 5.4. *Six colors are sufficient to color any map on the sphere or projective plane.*

Proof. Let κ be a complex on φ with F faces, E edges, and V vertexes.

Find all vertexes of κ where only two edges meet and combine the edges into one edge. Now every vertex lies on each 3 edges of κ. This changes the number of edges, but not the number of faces.

Note the total number of edges counted vertex by vertex is $2E$, since each edge connects two vertexes and is therefore counted twice. Thus the average number of edges per vertex, $\frac{2E}{V}$, must be at least three.

$$\frac{2E}{V} \geq 3$$
$$V \leq \frac{2E}{3}$$

Also, Theorem 4.19 tells us that $V = \chi - F + E$

$$\text{Substituting for } V\text{, we get } \chi - F + E \leq \frac{2E}{3}$$
$$\text{Solving for } E\text{, we get } E \leq -3\chi + 3f$$
$$\text{Substituting into } a\text{, we get } a = \frac{2E}{F} = \frac{2(-3\chi + F)}{F}$$

So we get

$$a = 6\left(1 - \frac{\chi}{F}\right) \tag{5.5}$$

We will only consider $F > \chi$. This is because $\chi < 6$, so if $F \leq \chi < 6$, then the complex is clearly 6-colorable.

We know the Euler characteristic of a sphere is 2 and that of a projective plane is 1. Then for the sphere and projective plane,

$$a \leq 6\left(1 - \frac{\chi}{F}\right) < 6$$

Therefore by Lemma 5.2, 6 colors are sufficient to color any complex on a sphere or projective plane. ☐

Theorem 5.6. *Let φ be any surface of characteristic $\chi \leq 0$. Then any map on φ can be colored by N_χ colors, where*

$$N_\chi = \left\lceil \frac{7 + \sqrt{49 - 24\chi}}{2} \right\rceil$$

Proof. Let $\chi \leq 0$. As in the previous proof, $a \leq 6\left(1 - \frac{\chi}{F}\right)$. We want to find N so that

$$6\left(1 - \frac{\chi}{F}\right) < N$$

If the above equation holds, then by Lemma 5.2, we will know how many colors are sufficient to color the complex. For every complex, note that if $F < N$, the complex is clearly N-colorable. Now assume $F \geq N + 1$.

Then $6\left(1-\dfrac{\chi}{F}\right) < N$

Solving for N, we get $N^2 - 5N - 6 + 6\chi > 0$

which yields $N > \dfrac{5 \pm \sqrt{49-24\chi}}{2}$

We will take

$$N > \frac{5 + \sqrt{49-24\chi}}{2}$$

since this is the stronger of the two requirements.

Adding 1 to both sides gives

$$N + 1 > \frac{7 + \sqrt{49-24\chi}}{2}$$

The smallest integer satisfying this equation is also the largest integer satisfying

$$N \leq \frac{7 + \sqrt{49-24\chi}}{2}$$

Thus

$$N_\chi = \lceil \frac{7 + \sqrt{49-24\chi}}{2} \rceil$$

N_χ is the number of colors sufficient to color a map on a surface with a Euler characteristic less than or equal to zero. □

Acknowledgments. It is a pleasure to thank my mentor, Bena Tshishiku, and Dr. Peter May for their guidance on this paper.

References

[1] Michael Henle. A Combinatorial Introduction to Topology. General Publishing Company, Ltd. 1979.