Construction of Eilenberg-Maclane Spaces

Ábel Juhász
Supervisors: Lauran Toussaint, Thomas Rot
Vrije Universiteit Amsterdam

June 2023

Contents

1 Introduction

This paper serves as an introduction to Eilenberg-Maclane Spaces – a type of space that is primarily of interest in the field of algebraic topology. While the reader is expected to possess knowledge of the fundamentals of point-set topology (some basic properties of continuous functions, quotient maps, connected spaces, compact spaces, etc.), as well as cursory knowledge of algebraic topology (homotopy classes, homotopy equivalence, the fundamental group, covering spaces), further understanding of algebraic topology is not required to read this paper. Furthermore, in the field of topology, we primarily use functions that are continuous, so any time this paper invokes a generic map between two spaces, the unstated assumption will be that it is continuous, unless explicitly stated otherwise.

To define exactly what an Eilenberg-Maclane space is, we will first have to define the notion of a Homotopy Group, a generalization of the fundamental group for higher dimensions.

Definition 1.1 (Homotopy Group)**.** Let X be a topological space, and x_0 a point in X. Then $\pi_n(X, x_0)$ consists of equivalence classes of maps from I^n to X that send each point in ∂I^n to x_0. Two such functions are equivalent when there is a homotopy from one to the other that is fixed on ∂I^n.

On $\pi_n(X, x_0)$, we define a group structure (with operation $+$) in the following way: if $[f], [g] \in \pi_n(X, x_0)$, then

$$f + g(s) := \begin{cases} f(2s_1, s_2, s_3, \dots, s_n) & \text{if } s_1 \in [0, \frac{1}{2}] \\ g(2s_1 - 1, s_2, s_3, \dots, s_n) & \text{otherwise.} \end{cases}$$

We can define $[f] + [g] := [f + g]$. The nth homotopy group of X with basepoint x_0 is then $\pi_n(X, x_0)$.

For our purposes, there are certain aspects of this definition we will take for granted; namely that homotopy groups are well defined and are groups. We will also take for granted that, if X is path-connected, then $\pi_n(X, x_0)$ is independent (up to isomorphism) of our chosen basepoint, and can thus be written just as $\pi_n(X)$. As most spaces we will work with are path-connected, we will henceforth use the latter notation. [2, pp. 340-341]

Furthermore, there is a one-to-one correspondence between maps $S^n \to X$ that send some basepoint of S^n to x_0, and maps $I^n \to X$ that send ∂I^n to x_0. Explicitly, if $q : I^n \to S^n$ is the quotient map that identifies all points in ∂I^n and sends them to the basepoint of S^n, then a map $S^n \to X$ can be precomposed with q to get a map $I^n \to X$ (the fact that this constitutes a one-to-one correspondance is true by definition). What this entails is that $\pi_n(X, x_0)$ can be equivalently defined as containing equivalence classes of maps with domain S^n instead of I^n. This paper shall therefore use these two definitions interchangeably, depending on context.

A noteworthy property of homotopy groups is that, even for simple spaces like S^n, the associated homotopy groups can be rather chaotic, and their compu-

	π_1	π_2	π_3	π_4	π_5	π_6	π_7	π_8	π_9	π_{10}
S^1	$\mathbb{Z}$	0	0	0	0	0	0	0	0	0
S^2	0	$\mathbb{Z}$	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$
S^3	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{12}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_3$	$\mathbb{Z}_{15}$
S^4	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z} \times \mathbb{Z}_{12}$	$\mathbb{Z}_2^2$	$\mathbb{Z}_2^2$	$\mathbb{Z}_{24} \times \mathbb{Z}_3$
S^5	0	0	0	0	$\mathbb{Z}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$	$\mathbb{Z}_{24}$	$\mathbb{Z}_2$	$\mathbb{Z}_2$

Table 1: The first 10 homotopy groups of spheres of dimension $1, 2, 3, 4$, and 5.

tation is usually extremely difficult. This unpredictable behavior is illustrated in Table 1.[2, p. 339]

Indeed, if one were to intuitively think of the nth homotopy group as a method of classifying n-dimensional holes in a space, this would seem quite strange: it would imply the existence of higher-than-n dimensional holes in S^n, which is visually counter-intuitive. The resolution to this dilemma is that homotopy groups constitute just one possible algebraic invariant used to classify holes in a space. In Section 2.4, we will explore a different, but closely related structure to homotopy groups: homology groups, which give far more intuitive (and easier to compute) results within this domain.

Looking at the chaotic behavior of homotopy groups on simple spaces, one might ask the question: are there any spaces whose homotopy groups are "nice"? An attempt to formally define what we mean by "nice" leads us to the subject of this paper:

Definition 1.2 (Eilenberg-Maclane space)**.** Let G be a group and let n be a positive integer. If X is a path-connected space such that $\pi_k(X) \cong 0$ whenever $k \neq n$, and $\pi_n(X) \cong G$; then X is an Eilenberg-Maclane space. An arbitrary Eilenberg-Maclane space corresponding to G and n is notated $K(G, n)$.

A question raised by this definition is: for which pairs (G, n) of groups and (positive) integers does a $K(G, n)$ space exist? Answering this question will be the main focus of this paper. We can already provide a partial answer using the following theorem:

Theorem 1.1. *Let X be a path-connected topological space, and let $n \geq 2$. Then $\pi_n(X)$ is abelian.*

Proof. Take arbitrary maps $f, g : I^n \to X$ such that all points in ∂I^n map to an arbitrary basepoint x_0. We will construct a homotopy that is fixed on boundary points from $f + g$ to $g + f$. First, for an arbitrary map $h : I^n \to X$ which is x_0 on boundary points, define 2 homotopies U_h (up) and D_h (down) that depend on h, and essentially "squash" h in the respective direction. Formally,

$$U_h(s, t) = \begin{cases} h(s_1, (1+t)s_2 - t, s_3, \ldots, s_n) & \text{if } (1+t)s_2 - t \in [0, 1], \\ x_0 & \text{otherwise.} \end{cases}$$

Confirming that U_h is continuous is just a matter of applying the gluing lemma. The same applies to D_h, which is defined

$$D_h(s,t) = \begin{cases} h(s_1, (1+t)s_2, s_3, \ldots, s_n) & \text{if } (1+t)s_2 \in [0,1], \\ x_0 & \text{otherwise.} \end{cases}$$

With this, one can give a description of the homotopy from $f+g$ to $g+f$ (illustrated in Figure 1): First, take $f+g$ and apply the homotopies D_{f+g} and U_{f+g} restricted to the left ($s_1 \le 0.5$) and right ($s_1 \ge 0.5$) halves of I^n respectively. As $f+g$ is x_0 on all points where $s_1 = 0.5$, the gluing lemma makes this a valid step. The result is a function that is constant on the top-left and bottom-right quadrants of I^n, and squashed versions of f and g on the other 2 quadrants. We can then move f to the right and g to the left. The exact formula for this transformation is

$$H_h(s,t) = \begin{cases} h(s_1 - 0.5t, s_2, \ldots, s_n) & \text{if } s_1 - 0.5t \in [0,1] \wedge s_2 \le 0.5 \\ h(s_1 + 0.5t, s_2, \ldots, s_n) & \text{if } s_1 + 0.5t \in [0,1] \wedge s_2 \ge 0.5 \\ x_0 & \text{otherwise.} \end{cases}$$

Here, we reused the dummy variable h to constitute any map that sends the top-left and bottom-right quadrants (along with boundary points) of I^n to x_0. This homotopy is then continuous by the gluing lemma.

Finally, we apply inverses of the homotopies (i.e. we do the homotopies in the opposite direction) U_{g+f} and D_{g+f} to the left and right parts of the result, and we end up with $g+f$ as desired. This completes the proof. $\square$

This means that for any case where $n > 1$ and G is not abelian, a $K(G,n)$ space cannot exist. What is surprising is that this turns out to be the only criterion one has to impose upon (G,n): in any other case, we can construct a $K(G,n)$ space, regardless of how complicated G is. Constructing these spaces, however, requires some groundwork.

As such, the next section will be focused on building a mathematical framework; on gathering the tools that will be used for the proofs to come later. Specifically, it will introduce cell complexes, the domain in which we will work for the rest of the paper. Section 2.1 will present several important definitions, and discuss the role cell complexes play in homotopy theory. Sections 2.2-2.3 will provide proofs for a multitude of theorems about cell complexes. Of particular note: the proof for Theorem 2.3 – and the accompanying Lemma 2.4 – is fully unique to this paper (barring a coincidence unknown to this paper's author). Section 2 ends with a brief introduction to homology, along with several theorems which will be stated without proof, both for the sake of brevity, and to make the paper more digestible.

Section 3 will use the knowledge attained in the previous section to construct $K(G,n)$ spaces for arbitrary G and n (assuming G is abelian when $n > 1$), thus achieving the paper's main objective. Finally, there will be a short subsection

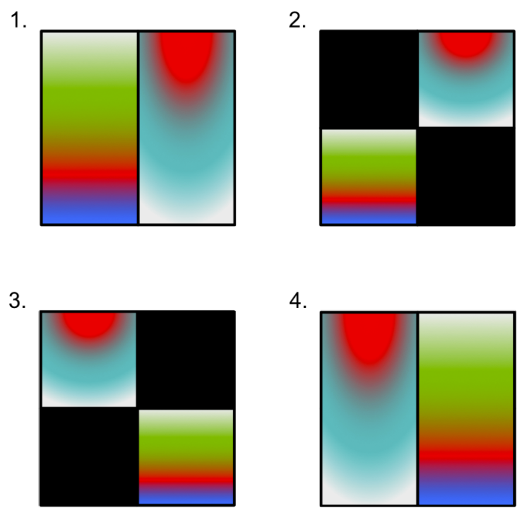

Figure 1: A visual illustration of a homotopy from $f + g$ to $g + f$, for the case of 2 dimensions. The colors indicate where a point on the square maps to, with black points being those that map to the base point.

discussing other miscellaneous properties $K(G, n)$ spaces have, followed by a conclusion.

2 Cell Complexes

2.1 Motivation and Definition

The act of constructing spaces with desirable homotopy groups is an iterative process. We usually start with a simple space we understand, but one which does not have the homotopy groups we are looking for, and then proceed to manipulate it. Step by step, we attach points to, remove points from, or glue points together on our topological space, with each step creating a space that is visually more complicated, but has homotopy groups closer to what we want. Out of the myriad methods there are to manipulate a space, the most useful for our purposes will be the notion of attaching cells:

Definition 2.1 (n-cells)**.** Let X be a topological space. We say that X' is attained by attaching n-cells to X when there exists some indexing set I and an attaching map $\varphi_i : \partial D_i^n \to X$ for each $i \in I$, such that X' is homeomorphic to the quotient space of $X \cup \bigcup_{i \in I} D_i^n$ (a disjoint union of X and indexed n-dimensional disks) induced by the attaching maps. Specifically, the equivalence relation that induces the quotient map makes all $x \in \partial D_i^n$ equivalent to $\varphi_i(x)$. The image of the interior of D_i^n is called an n-cell in X' and is notated e_i^n (where i is just there to index this specific n-cell).

What naturally follows is a type of space that arises by requiring it to consist entirely of cells.

Definition 2.2 (cell complex)**.** A cell complex, or CW-complex [2, p. 5], is a topological space X for which the following holds:

There exists a sequence of spaces

$$X^0 \subset X^1 \subset X^2 \subset \dots$$

such that X^0 is a discrete space (which we say contains 0-cells), and all other X^k is attained by attaching k-cells (or potentially attaching nothing) to X^{k-1}. X^k is called the k-skeleton of X, and X equals the union of all its skeletons. Furthermore, the topology on X is the induced limit topology: a set U is open in X if and only if $U \cap X^k$ is open in X^k for all k.

Cell complexes are of great importance to algebraic topology, for several reasons. The first is practical: they generally behave well with respect to homotopy, as will be seen in Sections 2.2 and 2.3; and their iterative nature allows us to prove many statements in an inductive manner, something that cannot usually be done for general topological spaces. The second reason is more philosophical.

One can define a slightly weaker type of equivalence (than homotopy equivalence) on topological spaces called a weak homotopy equivalence. [1]

Definition 2.3 (Weak equivalence)**.** Let X and Y be spaces, and let $f : X \to Y$ be a map, such that the induced map between path components $f_* : \pi_0(X) \to \pi_0(Y)$ is a bijection, and for all points $x_0 \in X$, $f_* : \pi_n(X, x_0) \to \pi_n(Y, f(x_0))$ is an isomorphism. Then f is called a weak homotopy equivalence (or just weak equivalence).

Surprisingly, any topological space is weakly homotopy equivalent to a cell complex, and any weak homotopy equivalence between two cell complexes is a (non-weak) homotopy equivalence. This latter result in known as Whitehead's theorem.[2, pp. 346-347, pp. 352-353]

For those well-versed in category theory, this means that the category **Hotop** (of topological spaces where homotopic maps are considered the same) can be altered slightly (through a process called localization) to attain a category that is equivalent to the subcategory of **Hotop** only containing cell complexes. For everyone else, this can be translated to "we do not lose much generality by restricting our attention to cell complexes."

2.2 Homotopy Extension Property

Often, when constructing a homotopy of a map $f : X \to Y$, we first construct it on some subspace $A \subset X$, and then extend it to the full space. It would be useful for our purposes if A was a subspace for which this was always possible, so examining when that is the case is a natural starting point.

[1] This technically does not induce an equivalence relation, instead it induces a relation R that is reflexive, transitive, but not symmetric. However, we can just use the smallest equivalence relation that contains R.

Definition 2.4 (Homotopy Extension Property)**.** Let X and $A \subset X$ be topological spaces. The pair (X, A) has the homotopy extension property if and only if for all topological spaces Y, and all maps $f : X \to Y$, any homotopy $A \times [0,1] \to Y$ from $f|_A$ can be extended to a homotopy $X \times [0,1] \to Y$ from f.

Lemma 2.1. *A pair (X, A) has the homotopy extension property if and only if $S = X \times \{0\} \cup A \times [0,1]$ is a retract of $X \times [0,1]$.*

Proof. First, suppose $r : X \times [0,1] \to S$ is a retract. Now suppose $f : X \to Y$ is some map, and let $H : A \times [0,1] \to Y$ be a homotopy of the restriction of f to A. By the gluing lemma one can construct a continuous function $H' : S \to Y$ as being equal to f on $X \times \{0\}$ and being equal to H on $A \times [0,1]$. The map H' can be precomposed with r to get a map $H'' : X \times [0,1] \to Y$ (that is, $H'' = H' \circ r$), which is the homotopy we were looking for.

Now suppose the homotopy extension property holds for (X, A). Take the homotopy $H : A \times [0,1] \to S$ that is just the inclusion map. As the function $H(.,0)$ can be extended to just the inclusion map $X \times \{0\} \to S$, by the homotopy extension property, there exists a homotopy $H' : X \times [0,1] \to S$ that is the identity on S. This makes H' a retract, as desired. $\square$

Using this lemma, we can prove the following important property of cell complexes:

Theorem 2.2. *If X is a cell complex, and X^k is a skeleton of X, then (X, X^k) has the homotopy extension property.*

Proof. We will construct a retract r from $X \times [0,1]$ to $X \times \{0\} \cup X^k \times [0,1]$, which will, by the previous lemma, complete the proof.

We will start by defining this retract on $X^{k+1} \times [0,1]$. For each $(k+1)$-cell e_α^{k+1}, there is a disk D^{k+1} and an associated map $\varphi_\alpha : D^{k+1} \to X_{k+1}$ such that φ_α restricted to the boundary of D^{k+1} is the attaching map of e_α^{k+1}, and restricted to the interior of D^{k+1}, it is just a homeomorphism onto e_α^{k+1}.

Take the space $D^{k+1} \times \{0\} \cup \partial D^{k+1} \times [0,1]$, and let φ'_α map from it to $X^{k+1} \times \{0\} \cup X^k \times [0,1]$ via the rule $(s,t) \to (\varphi_\alpha(s), t)$.

Using a suitable radial projection, we define $\pi : D^{k+1} \times [0,1] \to D^{k+1} \times \{0\} \cup \partial D^{k+1} \times [0,1]$ to be an arbitrary retract. To elaborate slightly: if one imagines this as being embedded in $\mathbb{R}^{k+2}$, then we can take the point $s = (0,0,0,\ldots,2) \in \mathbb{R}^{k+2}$, and let π map via radial projection from the point s (Figure 2 demonstrates this). Precomposing φ'_α with π gives us a map $\varphi''_\alpha : D^{k+1} \times [0,1] \to X^{k+1} \times \{0\} \cup X^k \times [0,1]$. Finally, from this we can define $r_{k+1} : X^{k+1} \times [0,1] \to X^{k+1} \times \{0\} \cup X^k \times [0,1]$ as

$$r_{k+1}(s,t) : \begin{cases} \phi''_\alpha(\phi_\alpha^{-1}(s), t) & \text{if } s \in e_\alpha^{k+1} \\ (s,t) & \text{if } s \in X^k. \end{cases}$$

We know that r_{k+1} is continuous, as it is continuous on $X^k \times [0,1]$, and on each $D_\alpha^{k+1} \times [0,1]$ (where D_α^{k+1} is an attached disk).

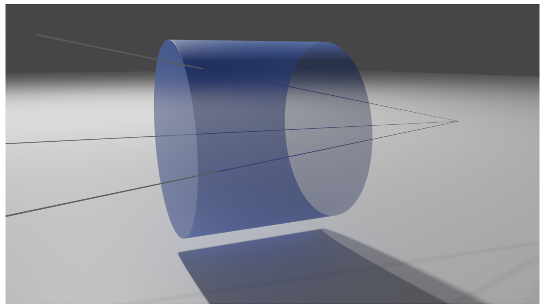

Figure 2: An illustration of radial projection from $D^2 \times [0,1]$ to $D^2 \times \{0\} \cup \partial D^2 \times [0,1]$. The points in the former get mapped to points in the latter (the blue shape) by following radial lines (three of which are represented in the figure).

Using the same process, we can create retracts $r_{k+j} : X^{k+j} \times [0,1] \to X^{k+j} \times \{0\} \cup X^{k+j-1} \times [0,1]$. Of note to us is the behavior of the composition of two adjacent retracts of this type: it can be verified that $(r_{k+j} \circ r_{k+j+1})|_{X^{k+j} \times [0,1]} = r_{k+j}$, which of course means $r_{k+j} \circ r_{k+j+1}$ is itself a retract, and it is an extension of r_{k+j}.

This inductively proves that, for each $(k+j)$-skeleton X^{k+j}, there exists a retract $r'_{k+j} : X^{k+j} \times [0,1] \to X^{k+j} \times \{0\} \cup X^k \times [0,1]$, and that these retracts can be assumed to be extensions of each other. Thus, we can define the retract r from $X \times [0,1]$ to $X \times \{0\} \cup X^k \times [0,1]$ as just the unique function such that for all $n > k$, $r|_{X^n \times [0,1]} = r'_n$. As X has the induced limit topology, this automatically implies r is continuous.

This completes the proof.

$\square$

This statement actually holds more generally for any cell complex pair (X, A) where A is a subcomplex of X, but this version of the statement is harder to prove, and will not be needed for our purposes.[2, p. 15]

2.3 Cellular Maps

Another property of maps between cell complexes that would be beneficial if we could just assume to always hold, is that the map in question "respects the cell structure." Formally, this can be defined as

Definition 2.5 (Cellular Map)**.** Let X and Y be cell complexes, and let $f : X \to Y$ be a map. Then f is called cellular when $f(X^n) \subset Y^n$ for all n.

Less formally, this just means f does not increase dimensions. Since we are only really interested in homotopy classes of maps, what we really want to show is that each homotopy class contains a cellular map. The proof of this fact will require considerable setup.

First, we will define the concept of a simplex:

Definition 2.6 (Simplex)**.** Let n be a natural number. The simplex Δ^n is defined by taking a set of $n+1$ points $v_1, v_2, \ldots, v_{n+1}$ (called the vertices of the simplex) and formally making them the basis of an $\mathbb{R}$-vector space (with the usual topology assigned). Then Δ^n is the subspace consisting of points of the form $t_1v_1 + t_2v_2 + \ldots + t_{n+1}v_{n+1}$ such that $t_1 + t_2 + \ldots + t_{n+1} = 1$ and each t_i is in $[0, 1]$.

Simplicies will be useful to us for the following reasons (which we omit proofs of):

1. An n-simplex Δ^n is homeomorphic to an n-disk. Therefore, for a cell complex, an n-cell can be identified with the interior of an n-simplex and the attaching map can be interpreted as having domain $\partial\Delta^n$.

2. If a function f from Δ^n to some convex subset of a euclidean space is already defined on the vertices of Δ^n, it can be linearly extended to the rest of Δ^n via the rule $f(t_1v_1 + t_2v_2 + \ldots t_{n+1}v_{n+1}) := t_1f(v_1) + t_2f(v_2) + \ldots + t_{n+1}f(v_{n+1})$. The image of f has dimension at most n. Furthermore, if two simplicies are attached (in that they share a face in common), then this linear extension will agree on the shared face.

3. Just as cubes can be subdivided into smaller cubes, simplicies can be subdivided into smaller simplicies. The subdivision we will use is called barycentric subdivision. The specifics of it are not relevant for our purposes, other than the fact that each simplex in the subdivision has a diameter of length at most $\frac{n}{n+1}$ times that of the original simplex. [2, pp. 119-120]

All this is sufficient setup to state the following theorem:

Theorem 2.3. *Let $f : \Delta^n \to X$ be a map, where X is a cell complex and Δ^n is an n-simplex. Let $m > n$, and suppose e^m is a cell in X. Furthermore, assume the image of $\partial\Delta^n$ does not contain any element of e^m, and that the image of Δ^n is in X^m. Then f is homotopic to some function g, such that the homotopy is fixed on $\Delta^n - f^{-1}(e^m)$ (in particular, on $\partial\Delta^n$), and the image of g does not contain every point in e^m.*

Proof. The first step is to identify e^m with $(0, 1)^m$.[2] As otherwise we would have nothing to prove, assume the image of f touches every point in $(0, 1)^m$.

Now, we will algorithmically subdivide Δ^n into smaller n-simplices, and then those simplicies into even smaller simplicies, in the following fashion (Figure 3 can act as visual aid for the first half of this proof):

Let $\mathcal{I}_1$ be the set that just contains Δ^n. Each subsequent $\mathcal{I}_k$ contains the simplicies obtained by performing barycentric subdivision on each element of $\mathcal{I}_{k-1}$. For any k, partition $\mathcal{I}_k$ into two collections. The first, $\mathcal{A}_k$, contains each simplex $S \in \mathcal{I}_k$ such that $f(S)$ is entirely contained in $(0, 1)^m$. The second, $\mathcal{B}_k$, is just the complement of the first. Now, let B_k be the union of each element of $\mathcal{B}_k$ (and A_k can be defined the same way but with $\mathcal{A}_k$). If $f(B_k) \cap (0, 1)^m \neq (0, 1)^m$,

[2] Any convex open subset of $\mathbb{R}^m$ which is homeomorphic to e^m would work. I chose $(0, 1)^m$ purely out of personal preference.

then we stop the algorithm and define the sets $\mathcal{A} := \mathcal{A}_k$ and $\mathcal{B} := \mathcal{B}_k$ as the "result" of the algorithm, with sets A and B being the unions of elements of the collections $\mathcal{A}$ and $\mathcal{B}$, respectively.

The algorithm has to conclude after a finite amount of steps. This is for the following reason: if it did not, it would imply for all k that $f(B_k) \cap (0,1)^m$ is equal to $(0,1)^m$. This means, as Δ^n is sequentially compact, we can find a convergent sequence x_k, such that $x_k \in B_k$, $f(x_k) = (0.5, 0.5, \dots, 0.5)$, and the limit of x_k is a point $x \in \bigcap_{k=1}^{\infty} B_k$. By continuity, $f(x) = (0.5, 0.5, \dots, 0.5)$; which is in $(0,1)^m$. This, however, is not possible: as f maps inside X^m, and $(0,1)^m$ is open in X^m, if $f(x) \in (0,1)^m$, then for some $\delta > 0$, an open ball of size δ around x also maps inside $(0,1)^m$. There exists some integer N, such that $\mathcal{I}_N$ only consist of simplicies with diameter less than δ (as the diameter of every simplex decreases by a factor of $\frac{n}{n+1}$ after each subdivision). Any one of those simplicies that contains x is also fully contained in the δ-ball around x, so maps entirely inside $(0,1)^m$. Thus, x is not in B_N, which is a contradiction.

With that out of the way, take the (finite) collection $\mathcal{A}$. Reusing the dummy-variable k, construct a sequence of collections $\mathcal{C}_k$, where for each k, the set $\mathcal{C}_k$ is attained by first taking elements of $\mathcal{A}$, successively performing barycentric subdivision on them k times; and then only taking simplices from that, which do not contain any boundary points of A (or equivalently, which do not contain any points of B). If we define C_k as the union of all elements of $\mathcal{C}_k$, it it follows that $\bigcup_{k=1}^{\infty} C_k = \text{int}(A)$.

From this, we define a function $g : \Delta^n \to X$ iteratively: first, take $g_0 : B \to X$ to just be equal to $f|_B$. The points in $\Delta^n - B$ are precisely the points in $\text{int}(A)$. An important observation to make is that $f(A) \subset (0,1)^m$, and $(0,1)^m$ is convex: we will explicitly use this in the following construction:

For all $k > 0$, we can define a function $g_k : B \cup C_k \to X$. Restricted to B this will always just be defined as f, thus we will only examine the other points. By way of induction, assume g_{k-1} is already defined, and that $g_{k-1}(C_{k-1})$ is in $(0,1)^m$. Then we define, for each point x in C_k that is also a vertex of some simplex in $\mathcal{C}_k$, $g_k(x) := g_{k-1}(x)$ if $x \in C_{k-1}$, and $g_k(x) := f(x)$ otherwise. More generally, the function g_k on arbitrary points of $C \in \mathcal{C}_k$ can be defined as the linear extension of how it is already defined on the verticies of C.

It can be confirmed (basically by definition) that if $j < k$, the function g_k agrees with g_j on C_j. Thus, as each element of Δ^n is either in B or in some set C_k, to define $g : \Delta^n \to X$ it is enough to require that it restricts to g_k for all k.

For a proof confirming g to be continuous, see Lemma 2.4.

The only things left to do are constructing a homotopy from f to g, and proving that the image of g does not contain every point in $(0,1)^m$. The latter can be proven in the following way: for all k, the image of the set C_k under g is an n dimensional subspace of $(0,1)^m$, making it a set with empty interior. Taking the union of all sets C_k (which is just $\text{int}(A)$) shows that, by the Baire category theorem[1, p. 394] – which states that for locally compact Hausdorff spaces, a countable union of sets with empty interior itself has empty interior – the set $g(\text{int}(A))$ has empty interior. By assumption, $g(B)$ does not touch every point in $(0,1)^m$. As B is compact, and X is Hausdorff, the image of B is

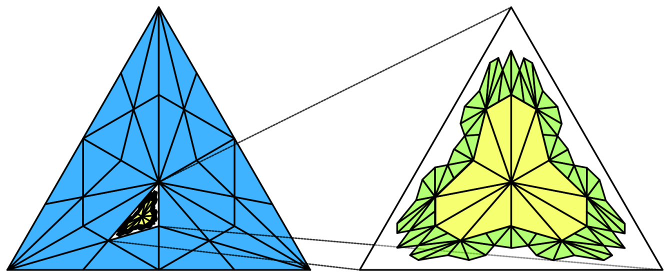

Figure 3: A visual representation of the proof for Theorem 2.3. The blue simplicies in this example are elements of $\mathcal{B}$, and the sole other simplex is the one element of $\mathcal{A}$. Zooming in on A, we can see how the sequence of sets C_k propagates outwards to fill the interior of A. Specifically, C_1 in this case is empty, C_2 is the union of all the yellow triangles, and C_3 is the union of the yellow and green triangles.

closed in X, thus $g(B) \cap (0,1)^m$ is closed in $(0,1)^m$. By definition, this means its (nonempty) complement is open. As $g(A)$ has an empty interior, it cannot touch every point in the aforementioned open set, which means $g(\Delta^n) \cap (0,1)^m = g(A) \cup (g(B) \cap (0,1)^m) \neq (0,1)^m$, proving that g does not touch every point in $(0,1)^m$.

Finally, the homotopy from f to g is

$$H(s,t) = \begin{cases} f(s) & \text{if } s \in B \\ (1-t)f(s) + tg(s) & \text{otherwise.} \end{cases}$$

This is a valid homotopy by the gluing lemma, and is fixed on all points that aren't in $f^{-1}((0,1)^m)$. This completes the proof.

□

To fill in the gap in the proof, we can prove the following lemma:

Lemma 2.4. *The function g constructed in the previous proof is continuous.*

Proof. First, on the set B, the function g is equal to f by definition, so $g|_B$ is continuous. By the gluing lemma, we will only need to show that $g|_A$ is continuous. To show that $g|_{\text{int}(A)}$ is continuous, we use the fact that $\{\text{int}(C_k)\}_{k=1}^{\infty}$ makes up an open cover of $\text{int}(A)$, and restricted to each of these open sets, g is continuous. This means $g|_{\text{int}(A)}$ is continuous.

Consequently, we only need to show continuity on boundary points of A. In order to achieve this, we will first need the following to hold: for any $k > 1$, the boundary of C_k does not intersect the boundary of C_{k-1}. This immediately follows from a previously unmentioned property of barycentric subdivision: if a simplex S has a vertex v, and the face of S that opposes v is labeled F, then no simplex in the subdivision of S which contains v will intersect F (proving from here that the boundary of C_k does not intersect the boundary of C_{k-1} is a matter of induction).

This implies by the definition of g, that for any k, if v is on the boundary of C_k, and v is a vertex of a simplex in $\mathcal{C}_k$; then $g(v) = f(v)$. It follows that if we have some sequence $x_i \in A$ that converges to some x in the boundary of A, then given that each x_i is on the boundary of some C_k, $g(x_i)$ converges to $g(x)$. This is because each of those x_i is inside a face of a simplex in some $\mathcal{C}_k$, and that face can be assumed to be on the boundary of C_k. We will label the n verticies of that face $(v_{i,1}, v_{i,2}, \dots, v_{i,n})$. Now, because the diameters of these faces go to 0 as i goes to infinity, for any fixed j, the sequence $v_{i,j}$ converges to x. Therefore,

$$\lim_{i\to\infty} g(x_i) = \lim_{i\to\infty} \sum_{j=1}^{n} t_{i,j} g(v_{i,j}) = \lim_{i\to\infty} \sum_{j=1}^{n} t_{i,j} f(v_{i,j}) =$$

$$= \lim_{i\to\infty} \sum_{j=1}^{n} t_{i,j} f(x) = f(x) = g(x).$$

We can now extend this proof method for a generic convergent sequence: Reusing the dummy variables, let $x_i \in A$ be a sequence that converges to a point x in the boundary of A. Each of these x_i is in some set $C_k - C_{k-1}$, and if we then choose any arbitrary simplex in $\mathcal{C}_k$ that contains x_i, let $(v_{i,1}, v_{i,2}, \dots, v_{i,n+1})$ be the $n+1$ vertices of the simplex we chose. By the same argument as last time, for each fixed j, the sequence $v_{i,j}$ converges to x. Furthermore, by construction, each one of these verticies is in the boundary of one of the sets C_k. Therefore,

$$\lim_{i\to\infty} g(x_i) = \lim_{i\to\infty} \sum_{j=1}^{n+1} t_{i,j} g(v_{i,j}) = \lim_{i\to\infty} \sum_{j=1}^{n+1} t_{i,j} g(x) = g(x).$$

This means g is continuous, which completes the proof.

□

It is quite remarkable that such a technical proof is required for a statement that is, purely visually, highly intuitive. Fortunately, this is enough foundation to prove the main theorem of this section.

Theorem 2.5. *Let X and Y be cell complexes, and let $f : X \to Y$ be a map. Then f is homotopic to a cellular map. Furthermore, if we assume f restricted to X^k is a cellular map, the homotopy can be assumed to be fixed on X^k.*

Proof. This will be proven inductively on the skeletons of X. First, the fact that there exists a homotopy from f to some function that maps 0-cells to 0-cells is an immediate consequence of the homotopy extension theorem. To elaborate: any path-component of Y must contain at least one 0-cell. Therefore, for any 0-cell x in X, there exists a path from $f(x)$ to a 0-cell in Y. These paths can be combined into a homotopy from $f|_{X^0}$ to a map from X^0 to Y whose image is in Y^0. The homotopy extension property of (X, X^0) then implies f can be homotoped to a function that maps 0-cells to 0-cells.

So, by induction, assume that f is already cellular on X^k. Take each $(k+1)$-cell e_α^{k+1}, and let $\varphi_\alpha : \Delta^{k+1} \to X^{k+1}$, be a function that maps the inside of the simplex homeomorphically to e_α^{k+1}, and suppose $\varphi_\alpha|_{\partial\Delta^{k+1}}$ is just the attaching map.

The first thing to note is that $f \circ \varphi_\alpha(\Delta^{k+1}) \subset Y^m$ for some m (without loss of generality we can then assume m to be the smallest such number). The proof of this is the following: By way of contradiction, assume no such m exists. Take a sequence x_i in Δ^{k+1} such that $y_i := f \circ \varphi_\alpha(x_i)$ is a sequence where each term is in a strictly higher-dimensional cell than the previous one. Also, as Δ^{k+1} is sequentially compact, without loss of generality we can assume x_i to be convergent, thus y_i also converges to some point y. We can also assume that y is not equal to y_i for any i (otherwise we can just remove that point from the sequence). Now, we want to show that $\{y_i\}_{i=1}^{\infty}$ is closed, as that would immediately create a contradiction. As Y has the induced limit topology, this is true whenever $Y^j \cap \{y_i\}_{i=1}^{\infty}$ is closed in Y^j for all j. These sets are all finite, meaning they are closed, which means y is not the limit of y_i, which is a contradiction.

We can employ the exact same proof-method (with minor adjustments) to show that there are only finitely many m-cells that $f \circ \varphi_\alpha$ meets. Now, if $m = k+1$, then we are done with this part of the proof. Otherwise, for each of those m-cells e_i^m, by Theorem 2.3, we can homotope $f \circ \varphi_\alpha$ in such a way that the homotopy is fixed on $\Delta^{k+1} - (f \circ \varphi_\alpha)^{-1}(e_i^m)$, and such that the result of this homotopy does not meet every point of e_i^m. As there are only finitely many m-cells to consider, we can concatenate such homotopies to get one whose end result misses at least one point in all m-cells of Y. Let's call this new function we get h.

To sum up, h is now a map whose image is in Y^m, and only meets finitely many m-cells, in each of which it misses at least one point. We can follow radial lines coming from those missed points to homotope h to a function that maps entirely inside Y^{m-1}. We can inductively keep doing this until we get a function that maps inside Y^{k+1}

All of this implies that there exists a homotopy from $f|_{X^{k+1}}$ to some function g, such that the homotopy is fixed on $X^{k+1} - e_\alpha^{k+1}$, and such that g maps e_α^{k+1} inside Y^{k+1}. Doing these homotopies for each α simultaneously gets us a cellular map from X^{k+1} to Y. By the homotopy extension theorem, this homotopy can be extended to one with domain X instead of X^{k+1}.

The only step left is combining everything together.

Let H_0 be a homotopy from f to a function that is cellular on X^0. Then we can inductively define H_{k+1} as the homotopy that takes the result of H_k, and moves it to a function that is cellular on X^{k+1}, without moving any points in X^k. The homotopy that combines all of these together will be called H, and it essentially just does H_k in the interval $[1 - 2^{-k}, 1 - 2^{-k-1}]$. Of course, the number 1 is not in any of these intervals, but as for all $x \in X$, there exists a natural number N such that for each $n \geq N$, H_n is fixed on x; $H(x, 1)$ can just be defined as $H_N(x,1)$. That H is continuous is clear on the domain $X \times [0,1)$. To prove continuity on other points, note that for any sequence (x_i, t_i) in $X \times [0,1]$ that converges to some point $(x,1)$, there exists some N such that $x_i \in X^N$ for all i. This means that for $t_i > 1 - 2^{-N-1}$, $H(x_i, t_i) = H(x_i, 1) = H_N(x_i, 1)$, which converges to $H_N(x,1) = H(x,1)$ by the continuity of H_N.

This means H is continuous, and thus completes the proof. $\square$

2.4 Brief Notes on Homology

As alluded to in Section 1, homology is an algebraic invariant assigned to topological spaces. Giving a comprehensive overview of homology is beyond the scope of this paper; rather, this section aims to provide at least a minor amount of intuition for how they work, and state several theorems about them that will be used in Section 3.

Broadly, the purpose of homology is classifying holes, so a reasonable first point of investigation is asking what it means for a space to contain a hole. Informally, we may come to the following pseudo-definition: If an n-dimensional object within our space surrounds something, but in way that cannot be filled in, then there is an n-dimensional hole in our space. Now, this does not really work as a definition, as its informality invites several questions:

- What do we mean by an n-dimensional object?
- What does it mean to surround something?
- What does filling in an n-dimensional object mean?

The first of these questions has a complicated answer for general topological spaces, but for our purposes, we only need to answer this question for cell complexes. So, if X is a cell complex, we can think of n-dimensional objects as just collections of n-cells. As homology is algebraic in nature, we will turn this into a group: The group $C_n(X)$ is just the free abelian group generated by each n-cell in X, and is called the group of n-chains of X.

We can try to use this to answer our other two questions: By "surround something", we mean that these n-cells are patched together in a way that leaves no gaps; or, in other words, our object has no boundary. By "cannot be filled in", we mean that our object is not the boundary of an $(n+1)$-dimensional object. Figure 4 illustrates this idea.

More formally, if we have some boundary map ∂_n that, for each n, sends elements of $C_n(X)$ to elements of $C_{n-1}(X)$, the kernel of $\partial_n : C_n(X) \to C_{n-1}(X)$ constitutes elements in $C_n(X)$ that "surround something" (these are called cycles), while the image of $\partial_{n+1} : C_{n+1}(X) \to C_n(X)$ constitutes elements in $C_n(X)$ that can be "filled in" (these are called boundaries). Furthermore, continuing with our analogy, if an n-dimensional object is a boundary, then it has to surround something; namely, the $(n+1)-$dimensional object it is the boundary of – so each boundary should also be a cycle. This invites the following purely algebraic definition:

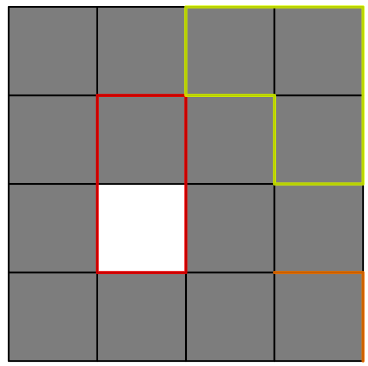

Figure 4: An illustration of how homology captures holes in a space. In our case, the red, yellow, and orange lines illustrate three different 1-chains. The orange one does not surround anything, as it has a boundary. The yellow one does surround something, as it has no boundary, but it does not surround a hole, as it is itself a boundary of a 2-dimensional object in our space. The red one is the only one to surround a hole, as it neither has a boundary, nor is a boundary of a 2-dimensional object in our space.

Definition 2.7 (Chain Complex)**.** For each natural number n, let C_n be an abelian group, and let for each n exist a homomorphism $\partial_n : C_n \to C_{n-1}$ (for $n = 0$, this is $\partial_0 : C_0 \to 0$). This system is a chain complex if for all n, $\partial_{n-1}\partial_n = 0$.

On this chain complex, we define homology groups as

Definition 2.8 (Homology of a Chain Complex)**.** Given a chain complex C, for each n we define two subsets of C_n: the first is B_n, which contains boundaries, and the second is Z_n which contains cycles. The set B_n is the image of $\partial_{n+1} : C_{n+1} \to C_n$, and Z_n is the kernel of $\partial_n : C_n \to C_{n-1}$. The n-th homology group H_n of this chain complex is the quotient group Z_n/B_n.

The intuition behind this definition is that cycles represent objects which might surround a hole, and, out of those objects, boundaries represent the ones which we know can be filled in – thus do not surround a hole. So, if we want to capture the behavior of holes in our space, reducing the former by the latter makes sense.

To define the homology groups of our cell complex X, we just need to define the boundary maps $\partial_n : C_n(X) \to C_{n-1}(X)$ in a way that creates a chain complex. After that, the homology groups of X can be defined to be the homology groups of that chain complex.

In order to construct these boundary maps, we will need the following Theorem:

Theorem 2.6. *For all n, $\pi_n(S^n) \cong \mathbb{Z}$.*

For the sake of brevity, we omit the proof.[5, p. 51] We can use this to define the following notion:

Definition 2.9 (Degree)**.** Let $n \geq 1$, and let $f : S^n \to S^n$ be a map. If $\pi_n(S^n)$ is identified with $\mathbb{Z}$,[3] then the degree of f is the number where $f_* : \pi_n(S^n) \to \pi_n(S^n)$ sends 1. Furthermore, the degree of a map $f : S^0 \to S^0$ is equal to 0 if f is constant, equal to 1 if f is the identity, and is equal to -1 if f switches the two points of S^0.

Useful for our purposes will be the subsequent lemma:

Lemma 2.7. *For each integer k and natural number $n \geq 1$, there exists a function $f : S^n \to S^n$ with degree k. Furthermore, homotopic maps have the same degree, and the constant map has degree 0 (even in the case of $n = 0$).*

Proof. Without loss of generality, assume $[\mathrm{id}] \cong 1$ in our identification $\pi_n(S^n) \cong \mathbb{Z}$ (where id is the identity on S^n). Then any representative $f : S^n \to S^n$ of an element of $\pi_n(S^n)$, such that $[f] \cong k$, has degree k. This is because $[\mathrm{id}] \cong 1$, and f_* sends this element to $[f] \cong k$.

Homotopic maps act the same way on $\pi_n(S^n)$, and the constant map acts by sending everything to 0; which proves the rest of the statement. □

Using the concept of a degree, we can finally define homology on cell complexes.

Definition 2.10 (Cellular Homology)**.** Let X be a cell complex. For each n, $C_n(X)$ is the free abelian group generated by the set of all n-cells of X. For each $n-$cell e^n_α, the associated attaching map is $\varphi_\alpha : S^{n-1} \cong \partial D^n_\alpha \to X^{n-1}$. For each $(n-1)$-cell e^{n-1}_β, we construct a quotient map φ_β from X^{n-1} to S^{n-1} by making all points outside of e^{n-1}_β equivalent under the relation $\sim$, and then applying the standard homeomorphism from $X^{n-1}/\sim$ to S^{n-1}. Composing φ_α and φ_β gets us a map $\varphi_{\alpha\beta} : S^{n-1} \to S^{n-1}$.

From this we define the boundary map $C_n(X) \to C_{n-1}(X)$ as sending each e^n_α to $\sum_\beta d(\varphi_{\alpha\beta})e^{n-1}_\beta$, where $d(\varphi_{\alpha\beta})$ stands for the degree of the map $\varphi_{\alpha\beta}$. The homology groups of the resulting chain-complex are the homology groups of X.

To be able to make swift progress, we will take for granted that this is well defined, that homology is independent of cell structure, and that homotopy equivalent cell complexes have the same homology groups.[2, pp. 137-140] Purely to gain intuition about cellular homology, we will prove the following lemma:

Lemma 2.8. *The n-th homology group of S^n is isomorphic to $\mathbb{Z}$.*

[3] There is a slight complication which arises from how one might choose the basepoint. Generally, if $g, g' : S^n \to S^n$ are representatives of elements in $\pi_n(S^n, x_0)$ and $\pi_n(S^n, y_0)$ respectively, if g and g' are homotopic, then $[g]$ and $[g']$ should be identified with the same number.

Proof. The standard cellular structure on S^n consists of one n-cell, and one 0-cell. If $n > 1$, this already is enough to imply that $H_n(S^n) \cong \mathbb{Z}$, as the associated chain complex is

$$\ldots \to 0 \to \mathbb{Z} \to 0 \to \ldots \to 0 \to \mathbb{Z} \to 0,$$

the two non-trivial groups being $C_n(S^n)$ and $C_0(S^n)$ respectively. As the boundary maps here are all forced to be zero, the nth homology group is trivially isomorphic to $\mathbb{Z}$. For $n = 1$, the chain complex looks like

$$\ldots \to 0 \to \mathbb{Z} \xrightarrow{\partial_1} \mathbb{Z} \to 0,$$

which means we need to take a closer look at the boundary map $\partial_1 : C_1 \to C_0$. The attaching map of the 1-cell is a constant, which means by Lemma 2.7, the associated degree is 0, making the boundary map trivial. This makes $H_1(S^1) \cong \mathbb{Z}$. $\square$

Going by intuition, n-dimensional spheres contain a single n-dimensional hole, and no other holes. The homology groups of S^n reflect this: for $k \geq 1$, $H_k(S^n) \cong 0$ whenever $k \neq n$, indicating that there is no k-dimensional hole in S^n; while $H_n(S^n) \cong \mathbb{Z}$, implying that there is an n-dimensional hole in S^n. Comparing this to Table 1, we notice an interesting contrast in complexity between higher homotopy groups and homology groups of spheres.

In spite of this contrast, these concepts are not unrelated. The relationship between homotopy groups and homology groups is described by the Hurewicz theorem[2, pp. 366-367], which (partly) states the following:

Theorem 2.9. *Let $n \geq 2$ and let X be an $(n-1)-$connected space (that is, a space whose first $n-1$ homotopy groups are trivial). Then $H_n(X) \cong \pi_n(X)$.*

This theorem gives us the final tool we need for constructing $K(G, n)$ spaces, so we shall begin construction now.

3 Construction of Eilenberg-Maclane Spaces

The previous section has been rather technical, and involved several lengthy proofs. In this section, we get to utilise the fruits of our labor, and prove the statement we posited in the beginning of this paper: a $K(G, n)$ space exists if and only if G is abelian or $n = 1$. As higher homotopy groups behave differently than the fundamental group, it makes sense to treat $n = 1$ as a special case.

3.1 Constructing $K(G, 1)$ Spaces

Theorem 3.1. *Let G be a group. Then a $K(G, 1)$ space exists.*

Proof. The general proof method will be the following: if we know that a space X has a contractible covering space, then we know that all higher homotopy groups of X are trivial. In short, this is because any map $f : S^n \to X$ can be lifted to that covering space[2, p. 61], and then applying a null-homotopy on the lift gives us a null-homotopy on the original map, f. Furthermore, we also know that the group of deck-transformations on that covering space is isomorphic to $\pi_1(X)$.

Therefore, our approach will be to construct a contractible space E on which the group G will act, and from that construct a $K(G,1)$ space X via the covering map $\pi : E \to X$, where the aforementioned group action on E will be just the group of deck-transformations. For this purpose it will be sufficient to require G to act freely of E: that is, aside from the identity element, no element of G will fix a single point in E.

The way we will construct E is by first constructing a suitable 0-skeleton of E that G acts freely on, and then inductively constructing (and extending the group action G to) all n-skeletons of E. Our starting space E^0 will just be the group G equipped with the discrete topology, where the group action is obvious. From this space we construct E^1 by, for any pair of distinct elements $a, b \in G$, attaching a 1-cell (or path) going from a to b, and a 1-cell going from b to a. The way an element g of G acts of this new space E^1 is by sending the 1-cell connecting a to b (in that order) to the 1-cell connecting ga to gb. This group action is free on E^1.

All higher skeletons we construct by induction: assume that E^k has already been constructed with the free group action G extended (freely) to it. We then take the set of homotopy classes of maps $[S^k, E^k]$. If this only consists of one element – the class of null homotopic maps – then $E^{k+1} := E^k$. Otherwise, we attain E^{k+1} by, for each homotopy class that is not null homotopic, taking some representative $f : S^k \to E^k$, and for each element of $g \in G$, attaching an $(n+1)$-cell to E^k via the attaching map gf. That is, the map f composed with the group action. The method of extending the group action from E^k to E^{k+1} is thus obvious: $g \in G$ sends an $(n+1)-$cell with attaching map f to one with attaching map gf.

Note that this may be excessive: that is, it is possible that we attach more cells than necessary at each step. For practical purposes, as long as $\pi_k(E^{k+1})$ becomes trivial, and our group action can be extended to E^{k+1}, it does not matter how many cells are attached (the argument detailing why this is the case is laid out in the next paragraph). In this proof we intentionally attach more cells than necessary, so it works on any choice of G.

We define E as the union of the sets E^k, with the standard induced limit topology (as is the case for all infinite dimensional cell complexes). To see that E is contractible, we will first show that all homotopy groups of E are zero: let $f : S^n \to E$ be a map, and without loss of generality, assume the basepoint of S^n maps to a 0-cell. By cellular approximation, f is homotopic to a cellular map, and we can assume said homotopy does not move the basepoint, so again without loss of generality, assume f is a cellular map. That means its image is entirely contained in E^n, or, more importantly for our purposes, E^{n+1}. By construction f is null-homotopic in E^{n+1} , so $\pi_n(E) \cong 0$.

This makes E weakly homotopy equivalent to a point. As E is a cell complex, Whitehead's Theorem implies E is therefore contractible.

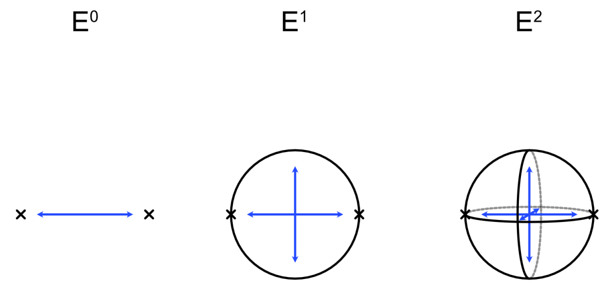

Figure 5: The first three skeletons of the covering space E in the construction of a $K(\mathbb{Z}_2, 1)$ space. The blue arrows indicate how $\mathbb{Z}_2$ acts on these skeletons by swapping antipodal points.

With all this setup, we can finally construct X: Let the map $\pi : E \to X$ be the quotient map that makes two points x and y in E equivalent if there exists some $g \in G$ such that $gx = y$. The function π is a covering map (the short reason why: around any point in E, one can take a sufficiently small neighborhood U, such that for any non-identity element g, gU is disjoint from U. This implies that $\pi^{-1}(\pi(U)) \cong G \times U$ where G is equipped with the discrete topology, which proves π is a covering map), and the group of deck-transformations on E is isomorphic to G, so X is a $K(G, 1)$ space.

□

For most groups, this construction cannot really be practically followed in any way that produces a visually interpretable result. As an example, following step one of this proof (i.e. the construction of E^1) for a group like $\mathbb{Q}$ immediately yields a non-drawable confluence of lines. However, for a finite group, one may actually be able to manually follow the proof and give a tangible cell decomposition of a $K(G, 1)$ space.

As an example, for $\mathbb{Z}_2$, our proof gives us the following construction: E^0 is just a two-point discrete space on which $\mathbb{Z}_2$ acts by switching the points. By attaching two 1-cells to these two points, we get E^1, which is a circle; the group action on it sending each point x to $-x$. Then, we get E^2 by attaching two disks to E^1, thus getting the space S^2; with the group action again sending x to $-x$.

In general, we see that $E^k = S^k$, meaning $E = S^\infty$, and the group action on E swaps antipodal points. The space we get by making antipodal points on S^k equivalent is called a real projective space (and is notated $\mathbb{RP}^k$), so $\mathbb{RP}^\infty$ is a $K(\mathbb{Z}_2, 1)$ space. Figure 5 visually illustrates this construction.

The fact that S^∞ is contractible is an interesting result in and of itself, and a whole class of $K(G, 1)$ spaces can be constructed by letting G act on S^∞ by "rotating" it, and then reducing the space by the group action (these are called infinite-dimensional lens spaces). To describe the process in more detail:

The $(2n-1)$-dimensional sphere is usually defined as a subspace of $\mathbb{R}^{2n}$. However, $\mathbb{R}^{2n}$ can be identified with $\mathbb{C}^n$, and if S^{2n-1} is treated as a subspace of $\mathbb{C}^n$, then

$$S^{2n-1} = \{(z_1, z_2, \dots, z_n) \in \mathbb{C}^n : \sum_{k=1}^{n} |z_k|^2 = 1\}.$$

This means that, since

$$S^1 \subset S^3 \subset S^5 \subset \dots \subset S^\infty,$$

and because S^∞ has the induced limit topology, S^∞ can be written as a subspace of $\mathbb{C}^\infty$. Namely, it is the space of all sequences that are eventually 0 (as each point in S^∞ is also in some S^{2n-1}), and for which $\sum_{k=1}^{\infty} |z_k|^2 = 1$.

Given a positive integer p, and a sequence of integers $(q_1, q_2, \dots)$ which are all coprime to p, we can define a map $S^\infty \to S^\infty$ with the rule

$$(z_1, z_2, z_3, \dots) \to (e^{\frac{2q_1\pi i}{p}} z_1, e^{\frac{2q_2\pi i}{p}} z_2, e^{\frac{2q_3\pi i}{p}} z_3, \dots).$$

This map is continuous (as it is just a set-theoretical product of continuous maps), and it generates a free group action on S^∞. These are essentially rotations of S^∞, and reducing by this group action gives us a $K(G,1)$ space. By way of example, setting p to any number greater than 1 and setting each q_k equal to 1 gives us a $K(\mathbb{Z}_p, 1)$ space.

3.2 Moore Spaces

The construction for $n > 1$ is more involved than the one for $n = 1$. We begin by constructing the homology analogue of a $K(G,n)$ space: a Moore space.

Definition 3.1 (Moore Space)**.** Let G be an abelian group, and $n \geq 1$. A Moore space $M(G,n)$ is a path-connected space whose n-th homology group is isomorphic to G, such that for any other $k \geq 1$, $H_k(M(G,n))$ is trivial.

Before we begin creating such a space, the following short algebraic preliminary is required:

Theorem 3.2. *Every abelian group is isomorphic to a quotient group of a free abelian group.*

Proof. We take an abelian group G, we take the abelian free group $\mathbb{Z}^{(G)}$, and map each generator e_g to g via a surjective homomorphism h. The group G is then isomorphic to $\mathbb{Z}^{(G)}/\ker h$ by the first isomorphism theorem. $\square$

This gives us enough ammunition to prove the existence of Moore spaces.

Theorem 3.3. *Let G be an abelian group and $n \geq 2$. Then $M(G,2)$ exists.*

Proof. Take a single-point space, and attach to it one n-cell e_g^n for each $g \in G$. Call this space X^n. Let N be a subgroup of $\mathbb{Z}^{(G)}$ such that $G \cong \mathbb{Z}^{(G)}/N$ (this can be done by the previous theorem). As N is a subgroup of an abelian free group, it is itself is an abelian free group.[3, p. 9-16] Let $S \subset \mathbb{Z}^{(G)}$ be a set of independent generators for N. Our next step is to attach one $(n+1)$-cell e_s^{n+1} to our space for each element s of S such that the attaching map $f_s : S^n \to X^n$ has the following property:

For each e_g^n, the degree of the map $q_g \circ f_s : S^n \to X^n/(X^n - e_g^n) \cong S^n$, where $q_g : X^n \to X^n/(X^n - e_g^n)$ is a quotient map; is equal to the coefficient of g in s. That is, if $s = \sum_{h \in G} c_h h$, where $c_h \in \mathbb{Z}$ (and only finitely many are non-zero), then the degree of $q_g \circ f_s$ should be c_g.

Thus, we will construct these functions: For each integer d, we define $h_d : S^n \to S^n$ to be some fixed map of degree d. By Lemma 2.7, such a map exists for all d. Now, for some $s \in S$, let $s = c_1 g_1 + c_2 g_2 + \ldots + c_k g_k$, where c_i are integers and g_i are elements of G. Let x_0 stand for some basepoint of S^n, let $U \in S^n$ be some open set that does not contain x_0, and let $f : S^n \to S^n$ be an arbitrary map. Then f can be homotoped into a function f' that is supported on U: that is, $f'|_{S^n - U}$ is just the constant function x_0.

The easiest way to show this is by assuming f already maps x_0 to x_0, in which case f can be pulled back to a function from I^n to S^n where the the boundary points of I^n all map to x_0. Let us call this new function p. Take a small closed n-cube V in I^n that is contained in U (or, more accurately, maps inside U via the quotient map identifying all boundary points of I^n). Suppose the image of the projection of V onto the ith coordinate is the interval $[x_i - \epsilon, x_i + \epsilon]$. From this we can write up the homotopy

$$H_i : I^n \times [0,1] \to S^n$$

$$H_i(y,t) : \begin{cases} p(y_1, y_2, \ldots, x_i + (1 + t(\frac{1}{\epsilon} - 1))(y_i - x_i), \ldots, y_n) & \text{if } x_i + (1 + t(\frac{1}{\epsilon} - 1)) \\ & (y_i - x_i) \in [0,1], \\ x_0 & \text{otherwise.} \end{cases}$$

Informally, this makes p supported on $I^{i-1} \times [x_i - \epsilon, x_i + \epsilon] \times I^{n-i}$. A simple application of the gluing lemma tells us H_i is continuous. Doing these homotopies for each i consecutively (though, obviously in such a way that the p within the definition of H_i changes to whichever function we are applying the homotopy to) gives us a function supported on V. Applying the naturally induced homotopy to f gives us a function supported on U, as desired.

Now, consider that for each n-cell e_g^n in X, there is a natural embedding $\varphi_g : S^n \to X$ that maps x_0 to the one 0-cell, and the rest of S^n to e_g^n. We are ready to construct our attaching map f_s: Take k disjoint open subsets $U_1, U_2, \ldots, U_k$ of S^n, such that none of them contain x_0. For each function h_{c_i}, we can take a homotopic function that is supported on U_i, and by Lemma 2.7, homotopic functions have the same degree, so without loss of generality we can just assume h_{c_i} is supported on U_i. This means we can construct $f_s : S^n \to X^n$ as:

$$f_s(x) : \begin{cases} \varphi_{g_i} \circ h_{c_i}(x) & \text{if } x \in U_i \\ \text{The 0-cell} & \text{if } x \notin U_i \text{ for all } i \end{cases}$$

We can confirm f_s to be continuous by the gluing lemma. Attaching each e_s^{n+1} via f_s gives us the space $X = X^{n+1}$. The chain complex attached to X is

$$\ldots \to 0 \to \mathbb{Z}^{(S)} \xrightarrow{\partial_{n+1}} \mathbb{Z}^{(G)} \to 0 \to \ldots \to 0 \to \mathbb{Z} \to 0.$$

The potentially non-trivial homology groups (other than H_0) are H_n and H_{n+1}. We know that H_{n+1} is trivial, as by construction ∂_{n+1} is just the inclusion map $N \to \mathbb{Z}^{(G)}$, meaning it is injective, and thus $H_{n+1}(X) \cong \ker \partial = 0$.

The nth homology group on the other hand is $H_n(X) \cong \mathbb{Z}^{(G)}/N \cong G$. This makes X an $M(G, n)$ space.

□

Important to us is the following corollary:

Corollary 3.3.1. *For all abelian groups G and integers $n \geq 2$, an $(n-1)$-connected space whose n-th homotopy group is G exists.*

Proof. Take the $M(G, n)$ space constructed in the previous proof. As for all $k < n$, the k-skeleton of that space is is just a point, by cellular approximation, any map $S^k \to X$ is null-homotopic. Thus $M(G, n)$ (or at least the version of it in the previous proof) is $(n-1)$-connected. The Hurewicz Theorem then implies that $G \cong H_n(X) \cong \pi_n(X)$. □

3.3 Killing Higher Homotopy Groups

By the previous corollary, the only homotopy groups we still need to worry about are the ones higher than n. Hence, what we need is some way to kill higher homotopy groups of the construction we already have, without affecting lower ones. This will be done by attaching higher and higher dimensional cells until each homotopy group above n is trivial. The next theorem ensures that this process will not affect lower homotopy groups. Note that we have used a similar statement for the construction of $K(G, 1)$ spaces: there, in different words, we used the fact that the inclusion map $\iota : X^{k+1} \to X$ induces a surjection $\iota_* : \pi_k(X^{k+1}) \to \pi_k(X)$. What we have not yet shown is that this is also an injection.

Theorem 3.4. *If $\iota : X^{n+1} \to X$ is the inclusion map, then $\iota_* : \pi_k(X^{n+1}) \to \pi_k(X)$ is an isomorphism for all $k \leq n$, and a surjection for $k = n+1$.*

Proof. Let $k \leq n+1$, and let $\iota_* : \pi_k(X^{n+1}) \to \pi_k(X)$ be the homomorphism induced by the inclusion map. By cellular approximation, if we assume $f : S^k \to X$ maps the basepoint of f to a 0-cell, the map f can be homotoped to one whose image is contained in X^{n+1} without moving the basepoint. Therefore, ι_* is surjective.

To prove it is also injective when $k \leq n$, suppose $f : S^k \to X$ and $g : S^k \to X$ have a basepoint-preserving homotopy H between them. Also, without loss of generality assume f and g to be cellular.

The domain of H, $S^k \times I$, can be given a cell structure as follows: take two points as the 0-cells, and attach them together with a 1-cell. Attach a k-cell to both 0-cells to get the spheres at each end of $S^k \times I$. Finally, to attach a single $k+1$-cell: Imagine our cell e^{k+1} as the interior of I^{k+1}. Let $\pi : I^{k+1} \to I$ be the projection onto the first coordinate. We define the attaching map $\varphi : \partial I^{k+1} \to (S^k \times I)^k$ (the final k denotes that it is the k-skeleton) in the following way: restricted to $\{1\} \times I^k$ and $\{0\} \times I^k$, the maps are just the natural quotient maps to the spheres on the two sides of the cylinder we are constructing (that is, in our case, it maps the boundary of $\{0\} \times I^k$ to the first 0-cell and the interior of $\{0\} \times I^k$ to the attached k-cell, while the set $\{1\} \times I^k$ maps similarly to the other 0-cell–k-cell pair). For all other points in ∂I^{k+1}, the attaching map is just the map π followed by the natural homeomorphism from there to the 1-cell.

This construction gives an $(n+1)$-dimensional cell structure on $S^k \times I$. If we assume H maps the two 0-cells (and thus the 1-cell as well) to the basepoint, then H restricted to the k-skeleton of $S^k \times I$ is cellular. By the cellular map theorem, this means that H can be homotoped into a basepoint-preserving homotopy from f to g contained entirely in X^{k+1}, and as $k+1 \leq n+1$, this makes ι_* injective. Therefore, ι_* is an isomorphism on the first n homotopy groups, completing the proof. □

We can use this to construct the $K(G, n)$ spaces we are looking for:

Theorem 3.5. *If $n \geq 2$ and G is an abelian group, a $K(G, n)$ space exists.*

Proof. Take the space $X^{n+1} = M(G, n)$ attained by our method of construction. As seen in Corollary 3.3.1, the first n homotopy groups are already as they should be. So, we will take the group $\pi_{n+1}(X^{n+1})$ and kill it. Specifically, by way of induction, suppose that for some $k > 0$, all homotopy groups below $n+k$ of the space X^{n+k} are as they should be: trivial, except for $\pi_n(X^{n+k}) \cong G$. Then construct X^{n+k+1} by attaching $(n+k+1)$-cells whose attaching maps are representatives of non-zero elements of $\pi_{n+k}(X^{n+k})$ (if there are any). The group $\pi_{n+k}(X^{n+k+1})$ will then be trivial, and by the previous theorem, the homotopy groups below $n+k$ are unaffected.

Repeating this process forever results in a (likely infinite-dimensional) space X, which is a $K(G, n)$ space. □

The main question posited in the introduction has thus been answered in the affirmative. That is to say, we have proven the following:

Theorem 3.6. *A $K(G, n)$ space exists if and only if $n = 1$ or G is abelian.*

3.4 Uniqueness and Other Miscellaneous Properties

While the primary focus of this paper is proving the existence of $K(G, n)$ spaces, it would be amiss not to at least mention other notable characteristics they possess. A decision made in writing this paper was to – up until this point – never refer to any space as being "the $K(G, n)$ space", electing to instead refer to them as "a $K(G, n)$ space"; because the former would imply some level of uniqueness, which we had no reason to assume. Therefore, it may be surprising that the following holds:

Theorem 3.7. *Let X and Y both be $K(G, n)$ spaces. Then X and Y are weakly homotopy equivalent.*

Proof. It is sufficient to show that any arbitrary $K(G, n)$ space Y is weakly equivalent to the $K(G, n)$ space we explicitly constructed, which we will label X.

We will construct a weak equivalence $f : X \to Y$. First, as $f_* : \pi_k(X) \to \pi_k(Y)$ is trivially an isomorphism for $k \neq n$ for any map, the only thing we need to ensure is that f acts as an isomorphism between the nth homotopy groups of X and Y.

Since our construction of X was different for $n = 1$ than for $n > 1$, one might expect the proof to be different for those two cases. However, this is not the case: regardless of the value of n, X^n consists of only one 0-cell, and no k-cells with $0 < k < n$, so it is just a wedge-sum of n-spheres. For each of those spheres, there is a homeomorphism c_i from S^n to that sphere sending the basepoint to the 0-cell. By cellular approximation, the collection $[c_i]$ generates $\pi_n(X)$.

This is enough information to create a map $f_n : X^n \to Y$. Specifically, take an arbitrary isomorphism $\varrho : \pi_n(X) \to \pi_n(Y)$. Suppose for each i, that g_i is a representative of $\varrho([c_i])$. Then let f_n be the unique map such that $f_n \circ c_i = g_i$ for all i.

Now, by construction, if f_n can be extended to X, the resulting map f will be a weak equivalence; so the only step left of this proof is the construction of such an extension. This will be done inductively: for $k \geq n$, assume $f_k : X^k \to Y$ is already defined, and is an extension of f_n. For any $(k+1)$-cell e_i^{k+1} in X^{k+1}, let $\varphi_i : S^k \to X^k$ be the attaching map. If $k = n$, then $f_k \circ \varphi_i$ is null-homotopic by construction; otherwise, it is null-homotopic because $\pi_k(Y)$ is trivial.

Let $H : S^k \times [0, 1] \to Y$ be a homotopy from $f_k \circ \varphi_i$ to some constant function. This induces a map H' from the quotient space $S^k \times [0, 1]/S^k \times \{1\}$ to Y such that if $q : S^k \times [0, 1] \to S^k \times [0, 1]/S^k \times \{1\}$ is the quotient map, then $H' \circ q = H$. Furthermore, D^{k+1} is homeomorphic to $S^k \times [0, 1]/S^k \times \{1\}$, meaning there exists an induced map $D^{k+1} \to Y$ that is equal to $f_k \circ \varphi_i$ on the boundary. Identifying the interior of D^{k+1} with e_i^{k+1} then induces a map from $X^k \cup e_i^{k+1}$ to Y. Doing this process for all $(k + 1)$-cells of X^{k+1} gives us the map $f_{k+1} : X^{k+1} \to Y$, which is an extension of f_k by construction.

The only thing left is to define $f : X \to Y$ as the unique map that restricts to f_k for each $k \geq n$. The map f is continuous, as X is equipped with the induced limit topology; and f is a weak homotopy equivalence, completing the proof. $\square$

In cell complexes weak equivalences and homotopy equivalences are one and the same, so within the domain of cell complexes, $K(G,n)$ spaces are unique up to homotopy equivalence.

A second major result is that $K(G,n)$ spaces can act as "building blocks" for all connected cell complexes (up to homotopy equivalence). The statement, in its raw form, is:

Theorem 3.8. *For any connected cell complex X, there exists a sequence of spaces (called a Postnikov tower)*

$$\{*\} = X_0 \leftarrow X_1 \leftarrow X_2 \leftarrow X_3 \leftarrow \ldots,$$

where each X_k has only trivial homotopy groups above k, and each map $X_k \to X_{k-1}$ is a fibration (which can be thought of as a generalisation of a projection map), whose fiber is $K(\pi_k(X), k)$. Furthermore, there exists maps $f_k : X \to X_k$ which induce isomorphisms for all homotopy groups below $k+1$, and which commute with the other maps in the tower.

There is a simple way to interpret this statement. If one thinks of this as "building X from the ground up", then $\{*\}$ is the ground, and a fibration $X_k \to X_{k-1}$ is a valid way to build on top of that. Going from X_k to X_{k+1} is just going from a space that agrees with X on the first k homotopy groups (and is trivial on all others) to one that agrees with X on the first $k+1$ homotopy groups (and is trivial on all others). In a way X then becomes the limit of the sequence $(X_0, X_1, X_2, \ldots)$. In fact, the space X is the category theoretical limit of the Postnikov tower, so the word "limit" is less of an analogy than the previous sentence may have implied. The ability to decompose a space in this manner has many applications. As an example, it can be used to compute the homotopy group $\pi_{n+1}(S^n)$ for any value of n.[4, p. 63]

4 Conclusion

In conclusion, we have shown the existence of $K(G,n)$ spaces given certain regularity conditions of G and n; namely, that G has to be abelian whenever $n > 1$. We have also shown that Eilenberg-Maclane spaces are unique up to weak equivalence.

Furthermore, as incremental steps to achieving the aforementioned main results, this paper also contains proofs for several important theorems relating to cell-complexes and homotopy groups. These are, in order: the nonexistence of higher non-abelian groups, the homotopy extension property of cell-complexes, and the cellular map theorem.

There are several questions one could ask about $K(G,n)$ spaces which were beyond the scope of this paper, and are subject to future research. First, for which (G,n) does a finite-dimensional $K(G,n)$ space exist? The method of construction included in this paper makes no attempt to keep the dimension of the resulting $K(G,n)$ space finite, so searching for an alternate construction that – perhaps for a more restricted class of (G,n)-pairs – produces finite dimensional $K(G,n)$ spaces is a suitable area to investigate next. In a similar vein, another question one might ask is: which $K(G,n)$ spaces have a cell decomposition that contains finitely many k-cells for each k? Or, a more uniform version of this question: which $K(G,n)$ spaces have a cell decomposition, such that for some natural number m, it holds that for any k, the amount of k-cells in the cell decomposition is at most m?

The previous three questions were about limitations imposed upon the size of $K(G,n)$ spaces. Questions that are more algebraic in nature are also worthy of investigation. These include questions about homology groups (or any other algebraic invariant) of $K(G,n)$ spaces; or questions about how group-theoretic properties of G correspond to topological properties of a $K(G,n)$ space. Further investigation into the role of Eilenberg-Maclane spaces within topology and algebra is therefore warranted.

References

[1] Beckenstein E. and L. Narici. *Topological Vector Spaces.* Chapman and Hall/CRC, 2 edition, 2010.

[2] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2 edition, 2002.

[3] D. Johnson. *Topics in the Theory of Group Presentations.* Cambridge University Press, 1980.

[4] L. Maxim. *Lecture Notes on Homotopy Theory and Applications.* University of Wisconsin-Madison.

[5] J. Milnor. *Topology from the Differentiable Viewpoint.* Princeton University Press, 1997.