2019-07-01

Extensions of the Power Group Enumeration Theorem

Shawn Jeffrey Green
Brigham Young University

Extensions of the Power Group Enumeration Theorem

Shawn Jeffrey Green

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of

Master of Science

Rodney Forcade, Chair
Benjamin Webb
Nathan Priddis

Department of Mathematics

Brigham Young University

ABSTRACT

Extensions of the Power Group Enumeration Theorem

Shawn Jeffrey Green
Department of Mathematics, BYU
Master of Science

The goal of this paper is to develop extensions of Pólya enumeration methods which count orbits of functions. De Bruijn, Harary, and Palmer all worked on this problem and created generalizations which involve permuting the codomain and domain of functions simultaneously. We cover their results and specifically extend them to the case where the group of permutations need not be a direct product of groups. In this situation, we develop a way of breaking the orbits into subclasses based on a characteristic of the functions involved. Additionally, we develop a formula for the number of orbits made up of bijective functions. As a final extension, we also expand the set we are acting on to be the set of all relations between finite sets. Then we show how to count the orbits of relations.

Keywords: Pólya enumeration, De Bruijn, power group, cycle index

Acknowledgments

I am indebted to my advisor Dr. Rodney Forcade for his constant guidance and support. I appreciate him letting me explore a concept in mathematics that I have always been fond of. I am also grateful for the mathematics department at Brigham Young University. Both the faculty and program have been enlightening and I have enjoyed working on this material as part of my program. Overall, I would not have been able to achieve these things without their support.

Contents

List of Tables

List of Figures

Chapter 1. Introductory Material

A interesting combinatorial question to ask is how many different ways can one color the sides of a shape with a fixed number of colors available? For example, how many ways can we color the sides of a square having three colors to use, but not needing to use all three of them. This would be easily answered if we did not consider the symmetry of the shape in question. A square can be rotated multiple ways, each rotation transforming colorings into other equivalent colorings. This leads us to ask how many ways there are to color the shape that do not differ from each other by a mere symmetry. In our example, how many symmetrically distinct ways can one color the sides of a square when three (or more) colors are available? These are the sorts of problems solved by Pólya's Theorem.

In terms of mathematical concepts, we can think of colorings as functions. We are mapping the set of vertices, sides, or faces to a set of colors. We also have symmetries which permute the domain of these functions. The symmetries form a group acting on the set of all such functions. Pólya's Theorem answers the question of how many symmetrically distinct functions there will be. In this chapter we will cover several formulations of Pólya's Theorem to lay the groundwork for generalizing it in Chapter 2. Additionally, we will discuss some extensions of Pólya's Theorem already discovered by De Bruijn, Harary, and Palmer. Finally, we will discuss some limitations of these generalizations to prepare for further extension in Chapters 2 and 3.

1.1 Burnside's Lemma and Basic Pólya Enumeration

We will begin with the foundational result that leads to Pólya's Theorem and most of the results in this paper, Burnside's lemma [1]. While the result is actually attributed to Cauchy and Frobenius [3], we will still use the name Burnside's lemma to fit in with the literature. It is an elementary fact from abstract algebra that a group G acting on a set S via a group action $*$ partitions S into equivalence classes. Here $s_1 \sim s_2$ if and only if $g * s_1 = s_2$ for some

$g \in G$. The equivalence classes of this equivalence relation are usually called *orbits*. We will denote the collection of orbits by S/G and the number of orbits by $|S/G|$. Burnside's lemma provides an effective way of counting the orbits when a finite group is acting on a finite set. To state the result, we will define $S^g = \{s \in S : g * s = s\}$. Thus S^g is the set of elements from S invariant under the action by group element g.

Lemma 1.1. ***(Burnside's)***

Let G be a finite group with a group action defined on a finite set S. Then

$$|S/G| = \frac{1}{|G|} \sum_{g \in G} |S^g|.$$

This lemma essentially says the number of orbits is the average number of elements of S fixed under the action by elements of G. Thus to find the number of orbits induced by a group action we only need to find out how many fixed elements there are for a given group element. While we will not prove Burnside's lemma in this paper, it is not difficult to show and can be found in many books about combinatorics or algebra. Specifically of the cited sources in this paper, [4],[7], and [10] all have proofs of Burnside's lemma.

We now turn our attention to the work of Pólya [8]. Technically, these results were first discovered by Redfield [9], but again the literature generally refers to it as Pólya's Theorem. This result follows from Burnside's lemma in a specific situation of permutations acting on functions. From here on we will take X (the domain of our functions) and Y (the codomain of our functions) to be finite sets with $|X| = n$ and $|Y| = m$. Pólya's Theorem addresses the situation where we take a subgroup G of the group of permutations on X denoted by S_X. We define a group action on the set Y^X of all functions from X to Y. For $\sigma \in G$ we can define the action of σ on $f \in Y^X$ by

$$f \mapsto f\sigma^{-1}. \tag{1.1}$$

Some others would define it as $f \mapsto f\sigma$, but either definition creates the same orbits. We will use the former for reasons that will become clear in Chapter 2. Pólya's Theorem then tells us how many orbits of functions there are under this group action.

To express Pólya's Theorem we will need a tool called the cycle index polynomial designated by Z_G. Using somewhat standard notation we first define $j_k(\sigma)$ for each $k \in \mathbb{N}$ to be the number of k-cycles in the disjoint cycle decomposition of σ. Then for a specific $\sigma \in S_n$ we define the *cycle index monomial* by

$$Z_\sigma(x_1, ..., x_n) = \prod_{i=1}^{n} x_i^{j_i(\sigma)}.$$

This cycle index monomial encodes the entire cycle structure of σ in a multivariate monomial. Then we can define the *cycle index polynomial* for an entire group to be

$$Z_G(x_1, ..., x_n) = \frac{1}{|G|} \sum_{\sigma \in G} Z_\sigma(x_1, ..., x_n).$$

Pólya's Theorem relates the number of orbits to the cycle index polynomial in the following way.

Theorem 1.2. ***(Pólya's)*** *Let G be a subgroup of S_X acting on Y^X as defined in (1.1). Then*

$$\left|Y^X/G\right| = Z_G(m, \ldots, m).$$

Pólya's Theorem is obtained by applying Burnside's lemma to the situation of G acting on Y^X and by counting the number of functions fixed by a $\sigma \in S_X$. The idea behind why the above formula works is that σ fixes a function from X to Y if and only if the function is constant on each collection of elements of X that make up the disjoint cycles of σ. With that idea it then becomes easy to count the number of functions fixed by a group element making Burnside's lemma yield precisely Pólya's Theorem.

A more detailed version of this result can be obtained when we consider that the functions could have weights associated with them. Weights are chosen from S, a commutative ring with identity that contains the rationals. We say a weight function is a map $W : Y^X \to S$ such that W is constant on each orbit in Y^X/G. Then a weighted version of Burnside's lemma [4] can be obtained saying that the sum of the weights of the orbits is actually the average of the weights of the fixed functions.

This weighted version of Burnside's lemma can be used to obtain a weighted version of Pólya Theorem. There are several common ways of defining these sorts of weight functions and we will briefly discuss one of them leading to a weighted version of Pólya's Theorem. Say we have weights associated with the elements of Y. More formally suppose we have a map $w : Y \to S$. Then we can construct a weight function on Y^X by defining for each $f \in Y^X$,

$$W(f) = \prod_{x \in X} w(f(x)). \tag{1.2}$$

Action by σ on f in (1.2) only changes the order of product and thus W is constant within orbits. So W is a weight function. Using W in the weighted form of Burnside's lemma yields the following the weighted version of Polya's Theorem [10].

Theorem 1.3. ***(Pólya's Weighted)*** *Let G be a subgroup of S_X acting on Y^X as defined by (1.1). Then*

$$\sum_{[f] \in (Y^X/G)} W(f) = Z_G(x_1, ...x_n)$$

where we set $x_i = \sum_{y \in Y} w(y)^i$.

As an immediate consequence of Polya's Weighted Theorem, one can obtain descriptions of the types of orbits created by the group action. To do this set $Y = \{y_1', ..., y_m'\}$. Then we can take the ring of weights S to be $\mathbb{Q}[y_1, ..., y_n]$ with $w(y_i') = y_i$ for each i. Thus the weight W of a function using (1.2) encodes the number of times each element of Y is mapped to by an element of X in a polynomial form. The term y_i^k appearing in such a weight means that the element of y_i' has k pre-images in X. Thus given $\epsilon_1 + \cdots + \epsilon_m = n$, we can obtain how many orbits of functions there are where each y'_i has ϵ_i pre-images. The combination of all this information is sometimes called pattern inventory as in [7].

Theorem 1.4. ***(Pólya's Inventory)*** *For each i let y_i be an indeterminate corresponding to the y'_i (the ith element of Y). Let $\epsilon_1 + \cdots + \epsilon_m = n$. The coefficient of $y_1^{\epsilon_1} \cdots y_m^{\epsilon_m}$ in the expanded polynomial*

$$Z_G\left(\sum_{i=1}^{m} y_i, \sum_{i=1}^{m} y_i^2, ..., \sum_{i=1}^{m} y_i^n\right)$$

is the number of distinct orbits of functions where any element of the orbit, $f \in Y^X$, has the stoichiometry defined by $|f^{-1}(y'_i)| = \epsilon_i$ for $i = 1, ..., m$.

Theorem 1.4 makes it not only possible to count how many symmetrically distinct ways one can color a cube for instance, but also allows one to count how many symmetrically distinct ways one could color a cube with two faces red, two blue, and two green. One just needs to substitute $x_i = r^i + b^i + g^i$ for each i between 1 and n into the cycle index polynomial for the group of symmetries on the faces of the cube. Then expand and collect like terms. The coefficient of $r^2b^2g^2$ in the result will be the answer. By performing these computations, we obtain there are actually six symmetrically distinct ways of coloring the faces of the cube with two of each color. Thus Theorem 1.4 gives us more detail on what types of orbits there are.

Together Pólya's results give a method for counting both the number and types of orbits induced by a group acting on the set of functions in the described way. We can apply these results to situations of coloring parts of objects where symmetry is involved. For instance, these results have been applied to counting the number of ways you can color vertices in a graph and how many ways the points in a lattice can be colored. Another application from chemistry is counting how many isomers can be formed from a given collection of atoms in a fixed geometric configuration [6].

There are other equivalent but slightly different looking versions of Pólya's Theorem. One common formulation involves plugging generating functions for the elements of Y by weight into the cycle index polynomial. While we will not cover this in detail, more information can be obtained from Harary and Palmer [4]. In conclusion, there are many ways of expressing Pólya's Theorem but they all have to do with counting orbits of functions (with or without weights) where a group of symmetries acts on the domain. There are two prominent extensions of Pólya's Theorem which we will look at in the next sections. One is offered by De Bruijn and the other was introduced by Harary and Palmer. Both have to do with adding in a group of permutations to act the codomain of the functions simultaneously as we act on the domain.

1.2 De Bruijn's Theorem

The first significant extension of Pólya's Theorem was provided by De Bruijn [2]. In Pólya's original work there are only permutations acting on the domain of the functions. De Bruijn added another group of permutations to simultaneously act on the codomain. He took H and K to be subgroups of S_X and S_Y respectively and defined an action of the group $G = H \times K$ on Y^X. For $f \in Y^X$ and $(\sigma, \tau) \in G$, De Bruijn defined a group action by

$$f \mapsto \tau f \sigma. \tag{1.3}$$

Again one can ask how many orbits of functions there are. However De Bruijn went straight for a weighted generalization so we will have to define weights in this situation.

We will actually use the version of De Bruijn's results that are in Harary and Palmer's paper [5]. So we introduce weights coming from the natural numbers like they did. Suppose that Y can be partitioned into $Y_1, ..., Y_s$ (some of which may be empty) and suppose K is a direct product of groups $K_1 \times \cdots \times K_s$ where each K_i is a subgroup of S_{Y_i}. Set $w(y) = i$ whenever y is in Y_i, and for $f \in Y^X$ set

$$W(f) = \sum_{x \in X} w\left(f(x)\right). \tag{1.4}$$

This will be a weight function since it will be constant on the orbits. Then the following result allows us to count orbits according to weight.

Theorem 1.5. ***(De Bruijn's)*** *Let H and K be subgroups of S_X and S_Y respectively and use the group action from (1.3). The coefficient of x^i after simplification in the below expression is the number of orbits whose functions have weight i as defined by (1.4).*

$$\left[Z_H\left(\frac{\partial}{\partial z_1}, ..., \frac{\partial}{\partial z_n}\right) \prod_{t=0}^{s} Z_{K_t}(p_{t,1}, p_{t,2}, ...) \right]_{z_i=0}$$

where $p_{t,s} = \exp\left(s \sum_{k=1}^{[n/s]} z_{sk}(x^t)^{sk}\right)$.

The above theorem looks quite unpleasant, but it is correct. It is an application of the weighted version of Burnside's lemma. Thus the expression must count the number of functions fixed by a (σ, τ). The idea behind this is that we can decompose σ and τ into a product of disjoint cycles (including one-cycles). Every element of X appears in a cycle of σ and every element of Y appears in a cycle of τ. In order for $f \in Y^X$ to be fixed by (σ, τ), the elements of X that appear in each cycle of σ must be sent to the elements of Y that appear in a cycle of τ where the length of the cycle of τ divides the length of the corresponding cycle of σ. Informally, we will say that functions fixed by a (σ, τ) must send the cycles of σ to cycles of τ where the length of the cycles of τ divide the lengths of the corresponding cycles of σ. We will see why this is the case later in Chapter 2 so if it is not clear now, it will be cleared up later. However, this is the main idea of why De Bruijn's Theorem works.

One big question when attempting to understand De Bruijn's Theorem is why are there partial derivatives and exponential functions? Harary and Palmer made a key observation in describing it [5]. The key observation is:

$$b^j = \left[\left(\frac{\partial}{\partial z}\right)^j \exp(bz) \right]_{z=0}.$$

In other words, the combination of exponential functions and partial derivatives amounts to plugging a quantity (b) into a polynomial. Thus De Bruijn's notation is a creative way of incorporating both cycle index polynomials into the solution. However, the same result can be expressed slightly more intuitively as was done by Harary and Palmer in what is called the Power Group Enumeration Theorem.

1.3 The Power Group Enumeration Theorem

Harary and Palmer also developed a result to count the number of orbits of functions in the same situation as that posed by De Bruijn [5]. They showed their results are just algebraic manipulations of what De Bruijn did. However, their result only makes use of the cycle index polynomial of H and thus does not have to introduce contrivances like partial derivatives and exponential functions to get both cycle index polynomials involved. Again we are taking H and K subgroups of S_X and S_Y respectively with the same group action as De Bruijn (1.3). Then the number of orbits is determined by the following theorem.

Theorem 1.6. ***(Power Group Enumeration Theorem-Constant Form)*** *Let $G = H \times K$ where H and K are subgroups of S_X and S_Y respectively. Let G act on Y^X as defined by (1.3) then*

$$\left|Y^X/G\right| = \frac{1}{|K|}\sum_{\tau\in K} Z_H\big(p_1(\tau), ..., p_n(\tau)\big)$$

where $p_k(\tau) = \sum_{s|k} s j_s(\tau)$.

As before, the reason why this works is that in order for a function to be fixed by this group action, the function must send the elements of each cycle of σ to a given cycle of τ whose length divides the length of the cycle of σ. This idea is encapsulated in the formula below that Harary and Palmer used to prove their result. They establish that

$$\left|(Y^X)^{(\sigma,\tau)}\right| = \prod_{k=1}^{n}\left(\sum_{s|k} s j_s(\tau)\right)^{j_k(\sigma)}. \tag{1.5}$$

The Power Group Enumeration Theorem is obtained by applying Burnside's lemma and then substituting in (1.5).

Harary and Palmer also created a weighted version of their result. We will not state it here since we will not be using it in this paper. They showed this result is equivalent to the result from De Bruijn. While De Bruijn's Theorem appears more similar to Polya's Theorem since the cycle index polynomials of both groups are involved, it also seems much less accessible than the Power Group Enumeration Theorem. Harary and Palmer also applied the Power Group Enumeration Theorem to counting colorings of lines in graphs, finite automata, and self-converse digraphs [4]. While it works in those situations, the Power Group Enumeration Theorem and the work of De Bruijn have a limitation that we will explain in the next section and address in Chapter 2.

1.4 Pathway to Further Generalization

De Bruijn, Harary, and Palmer were all working in the situation where the group was a direct product of subgroups of S_X and S_Y. This is not ideal since many actual situations would not have this property. A group structure of this kind, means there is no connection between the the symmetries on X and the ones on Y. All possible combinations of pairs of symmetries on X and Y must be included for it to be a direct product. However, if there is a connection between the symmetries this situation is no longer achieved.

Consider mapping the vertices of a square to the sides of the same square. When we perform a symmetry on the square it permutes both the vertices and sides in a specific way. The permutations are clearly connected and every possible pairing of them will not be present in the group. Thus the group in this case will not be a direct product. We wonder if we can expand the Power Group Enumeration Theorem to deal with this situation. In Chapter 2 we will address this question and adapt previous results. Furthermore, we will try to get something akin to the pattern inventories which we had for Pólya's Theorem.

Another way we can expand the situation of the Power Group Enumeration Theorem is to change the set of objects we are acting on. A natural set containing Y^X is the set of all relations between X and Y. This larger collection allows the elements of X to be associated with none or multiple elements of Y. In Chapter 3 we will see that we can define a similar group action on the larger set of all relations between finite sets and count the induced orbits.

Chapter 2. Orbits of Functions

In this chapter we will look at a similar situation to the one introduced by De Bruijn except we will no longer require that the group of symmetries be a direct product of groups. To make this work we will have to slightly alter the the group action definition. At the same time we would like our new group action to yield the same orbits as in the case when the group is a direct product and De Bruijn's action is used. Once we have a group action, we will generalize the Power Group Enumeration Theorem to deal with this situation and see if there is anything akin to a pattern inventory (as in Pólya's Theorem) that we can compute.

2.1 Group Action Definition

To replace the direct product structure of the group of symmetries, we are going to take G a subgroup of $S_X \times S_Y$. Thus we still are getting a collection of permutations on X and associated permutations on Y, but now G need not contain every possible combination of pairs of symmetries on X and on Y. However, if we use the same proposed action as De Bruijn, Harary, and Palmer, the operation is no longer a group action since it will not be associative. We can modify the action by defining for each $(\sigma, \tau) \in G$ and $f \in Y^X$

$$f \mapsto \tau f \sigma^{-1}. \tag{2.1}$$

We will denote this action by $*$. It will be a group action since

$$
\begin{aligned}
[(\sigma,\tau)(\sigma',\tau')] * f &= (\sigma\sigma',\tau\tau') * f \\
&= \tau\tau' f(\sigma\sigma')^{-1} \\
&= \tau\tau' f(\sigma')^{-1}(\sigma)^{-1} \\
&= (\sigma,\tau) * \left[\tau' f(\sigma')^{-1}\right] \\
&= (\sigma,\tau) * [(\sigma',\tau') * f]\,.
\end{aligned}
$$

Thus we have a group action in the case where G is not a direct product of subgroups. It is clear in our argument that we need to take the inverse of σ since otherwise the order of σ and σ' would be switched. The idea to take the inverse of σ in the action definition is not a new one. For Pólya's Theorem some mathematicians define the group action as right composition by σ^{-1} which is why we defined it this way in Chapter 1. In that situation we are working with a group so the orbits would be the same as acting by direct composition of σ instead of σ^{-1}. But in our more general situation when we remove the direct product structure in G, we must go back to what will make the proposed action fit the definition.

It actually does not matter that De Bruijn, Harary, and Palmer used a different group action when the group is a direct product. We will see that when G is a direct product of subgroups our action induces exactly the same orbits as the action presented by De Bruijn, Harary, and Palmer. This is because our definition of acting by (σ,τ) is the same as acting by (σ^{-1},τ) in their definition. When G is a direct product it will most certainly contain (σ^{-1},τ). Thus this new definition of the group action will create the same orbits as the ones from both De Bruijn's Theorem and the Power Group Enumeration Theorem in the case G is a direct product of subgroups. Now that we have an action defined, we can find the number of orbits.

2.2 Counting Fixed Functions

The problem of counting orbits in this generalized setting is similar to the situations considered by De Bruijn, Harray, and Palmer. They developed a method of describing the number of functions fixed by a group element (σ, τ) to prove their results. The ideas they developed still work in this more general situation since they did not assume anything about the group structure to prove it. And even though the two actions are slightly different, both actions have the same number of functions fixed after action by (σ, τ) since taking inverses does not change the cycle structure of the permutations. However, since this formula is so essential to the arguments that lie ahead we will provide a brief justification in this section of why those formulas work. We will describe how to count the number of functions fixed by a particular $(\sigma, \tau) \in G$ and from there the results we will prove are not too surprising.

To state the result we will introduce some notation for permutations that we will use in the statement as well as in the rest of the paper. For $\sigma \in S_X$ we will say σ_i is the ith cycle in a fixed disjoint cycle decomposition of the permutation. As a note, we will always include one-cycles when taking our disjoint cycle decompositions. Additionally, there are situations in which we will need to consider cycles as sets instead of permutations. So for $\sigma \in S_X$ we will say $\widehat{\sigma_i}$ is the set of elements of X in the cycle σ_i. Lastly, we will set $\ell(\sigma_i)$ to be the length of the cycle σ_i or equivalently the cardinality of $\widehat{\sigma}_i$.

Lemma 2.1. *For $\sigma \in S_X$ and $\tau \in S_Y$ let $\sigma = \sigma_1 \cdots \sigma_u$ and $\tau = \tau_1 \cdots \tau_v$ be the respective disjoint cycle decompositions. Fix $x_1, ..., x_u$ with each $x_i \in \widehat{\sigma}_i$. Given $j_1, ..., j_u$ from $\{1, ..., v\}$ such that $\ell(\tau_{j_i}) | \ell(\sigma_i)$, and $y_i \in \widehat{\tau_{j_i}}$ for each i, there exists a unique function $f : X \to Y$ such that both the following hold.*

i) $f(x_i) = y_i$ for each i.

ii) $(\sigma, \tau) * f = f$.

Proof. We will show the result is true when restricting to a single cycle σ_i. When that is true we can form a union of the functions produced for each cycle and obtain a function described in the statement of the lemma.

Note since $x_i \in \widehat{\sigma_i}$ we can write $\sigma_i = (a_1, a_2, ..., a_s)$ with $x_i = a_1$ and similarly because $y_i \in \widehat{\tau_{j_i}}$ we can write $\tau_{j_i} = (b_1, b_2, ..., b_t)$ with $b_1 = y_i$. We know additionally from assumption that $t|s$.

The result is obvious for $t = 1$ since in that case then we can take the constant function $g(x) = b_1$ which is clearly fixed by (σ_i, τ_j).

When $t > 1$ define $g : \widehat{\sigma_i} \to \widehat{\tau_j}$ as follows.

$$g(a_{tq+r}) = b_{r+1} \text{ where } 0 \leq r < t.$$

Note that

$$\tau_j g(a_{tq+r}) = \begin{cases} b_{r+2} & \text{if } r \leq t-2 \\ b_1 & \text{if } r = t-1 \end{cases}$$

and

$$\begin{aligned} g\sigma_i(a_{tq+r}) &= \begin{cases} g(a_{tq+r+1}) & \text{if } tq+r \neq s \\ g(a_1) & \text{if } tq+r = s \end{cases} \\ &= \begin{cases} b_{r+2} & \text{if } tq+r \neq s \text{ and } r \leq t-2 \\ g(a_{t(q+1)}) & \text{if } tq+r \neq s \text{ and } r = t-1 \\ b_2 & \text{if } tq+r = s \end{cases} \\ &= \begin{cases} b_{r+2} & \text{if } r \leq t-2 \\ b_1 & \text{if } r = t-1. \end{cases} \end{aligned}$$

Thus we get that

$$\tau_j g = g\sigma_i \implies \tau_j g \sigma_i^{-1} = g.$$

Thus g is an example of a function with the desired properties. We now show that is unique.

Suppose that $h : \widehat{\sigma_i} \to \widehat{\tau_{j_i}}$ is another example of such a function. We have that

$$g(a_1) = b_1 = h(a_1).$$

Now to proceed inductively assume $g(a_k) = h(a_k)$. Then we have that

$$g(a_{k+1}) = g(\sigma_i(a_k)) = \tau_j g(a_k) = \tau_j h(a_k) = h(\sigma_i(a_k)) = h(a_{k+1}).$$

Thus $h = g$ and we have uniqueness of the functions for each cycle. $\square$

It is not too difficult to see that the fixed functions described in Lemma 2.1 are actually all the functions fixed by (σ, τ). The functions fixed by (σ, τ) must send cycles of σ to entire cycles of τ where the lengths of cycles from τ divide the corresponding cycle lengths of σ. For those wanting more justification, we will see a formal argument for this in Chapter 3 when we generalize some of these results to relations. Regardless, we have been able to classify the functions fixed by a (σ, τ). We must send each cycle of σ to a cycle of τ where the length of the cycle of τ divides that of σ. Once one point of the function is fixed for each cycle of σ the entire function is determined. Thus for a cycle of τ_j whose length divides the length of σ_i there are $\ell(\tau_j)$ functions fixed by (σ_i, τ_j) going between $\widehat{\sigma_i}$ and $\widehat{\tau_j}$. The following result immediately follows.

Lemma 2.2. *Let $(\sigma, \tau) \in S_X \times S_Y$ with the group action being defined by (2.1). Then*

$$\left| \left(Y^X\right)^{(\sigma,\tau)} \right| = \prod_i \sum_{\ell(\tau_j) | \ell(\sigma_i)} \ell(\tau_j).$$

(Here we let i and j range over how many disjoint cycles there are in σ and τ respectively.)

Proof. By the arguments in the preceding paragraph $\sum_{\ell(\tau_j)|\ell(\sigma_i)} \ell(\tau_j)$ is the number of functions that will be fixed from $\widehat{\sigma_i}$ into Y. This can be done on each cycle of σ giving us how many ways we can make functions that will be fixed from each cycle of σ into Y. To get arbitrary fixed functions we can form a union of the fixed functions on each cycle of σ. Thus there are

$$\prod_i \sum_{\ell(\tau_j)|\ell(\sigma_i)} \ell(\tau_j)$$

ways of doing this. So the result holds. □

Collecting repetitive terms, it is not difficult to see Lemma 2.2 gives the exact same formula as Harary and Palmer (1.5). Now that we have a formula to count the fixed functions, the next step of counting orbits involves applying Burnside's lemma.

2.3 Generalizing the Power Group Enumeration Theorem

We are ready to count the orbits that come from the action of a possibly non-direct product group G acting on Y^X. The next result is how we extend the Power Group Enumeration Theorem to deal with this different group structure. Essentially the main change we need to make is to no longer use the cycle index polynomial of H and instead use only cycle index monomials as needed.

Theorem 2.3. *Let G be a subgroup of $S_X \times S_Y$ that acts on Y^X via $*$. Then*

$$\left|Y^X/G\right| = \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} Z_\sigma\big(p_1(\tau), ..., p_n(\tau)\big)$$

where we set $p_k(\tau) = \sum_{\ell(\tau_j)|k} \ell(\tau_j)$.

Proof. Using Burnside's lemma and (1.5) we get the following chain of logic.

$$
\begin{aligned}
\left|Y^X/G\right| &= \frac{1}{|G|}\sum_{(\sigma,\tau)\in G}\left|(Y^X)^{(\sigma,\tau)}\right| \\
&= \frac{1}{|G|}\sum_{(\sigma,\tau)\in G}\prod_{k=1}^{n}\left(\sum_{s|k} s j_s(\tau)\right)^{j_k(\sigma)} \\
&= \frac{1}{|G|}\sum_{(\sigma,\tau)\in G} Z_\sigma\left(\sum_{s|1} s j_s(\tau), ..., \sum_{s|n} s j_s(\tau)\right).
\end{aligned}
$$

The result as stated follows because for each k

$$
p_k(\tau) = \sum_{\ell(\tau_j)|k} \ell(\tau_j) = \sum_{s|k} s j_s(\tau). \qquad \square
$$

While Theorem 2.3 certainly does count the number of orbits in our more general situation, it is not as beautiful as the Power Group Enumeration Theorem since it is not as directly correlated with the cycle index polynomial of the whole group. However, in the case that $G = H \times K$ as in the Power Group Enumeration Theorem we get

$$
\begin{aligned}
\left|Y^X/G\right| = \frac{1}{|G|}\sum_{(\sigma,\tau)\in G} Z_\sigma\big(p_1(\tau), ..., p_n(\tau)\big) &= \frac{1}{|H||K|}\sum_{\tau\in K}\sum_{\sigma\in H} Z_\sigma\big(p_1(\tau), ..., p_n(\tau)\big) \\
&= \frac{1}{|K|}\sum_{\tau\in K} Z_H\big(p_1(\tau), ..., p_n(\tau)\big).
\end{aligned}
$$

Thus Theorem 2.1 is in fact the Power Group Enumeration Theorem in the situation of the group being a direct product.

We will now see some examples of how these results can be applied in situations where the Power Group Enumeration Theorem will not suffice. The original motivation of this result was to count the number of symmetrically distinct ways to put atoms into a crystal structure where each atom has an associated spin which is permuted by the symmetries of the lattice. This certainly can be achieved with Theorem 2.3, however we will make our

Figure 2.1: Labeled Tetrahedron with Vertices and Faces

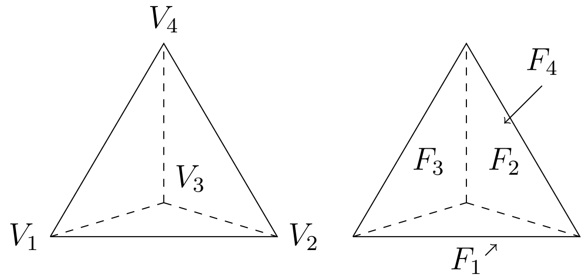

Table 2.1: Fixed Functions from Vertices to Faces of a Tetrahedron

σ	τ	Z_σ	$Z_\sigma\left(p_1(\tau),...,p_4(\tau)\right)$
$(V_1)(V_2)(V_3)(V_4)$	$(F_1)(F_2)(F_3)(F_4)$	x_1^4	4^4
$(V_1\ V_2\ V_3)(V_4)$	$(F_1)(F_2\ F_3\ F_4)$	x_1x_3	$1 \cdot 4$
$(V_1\ V_4\ V_3)(V_2)$	$(F_1\ F_4\ F_2)(F_3)$	x_1x_3	$1 \cdot 4$
$(V_1\ V_2\ V_4)(V_3)$	$(F_1\ F_2\ F_3)(F_4)$	x_1x_3	$1 \cdot 4$
$(V_1\ V_3\ V_2)(V_4)$	$(F_1)(F_2\ F_4\ F_3)$	x_1x_3	$1 \cdot 4$
$(V_1\ V_2)(V_3\ V_4)$	$(F_1\ F_4)(F_2\ F_3)$	x_2^2	4^2
$(V_1\ V_4)(V_2\ V_3)$	$(F_1\ F_2)(F_3\ F_4)$	x_2^2	4^2
$(V_1\ V_3\ V_4)(V_2)$	$(F_1\ F_2\ F_4)(F_3)$	x_1x_3	$1 \cdot 4$
$(V_1)(V_2\ V_4\ V_3)$	$(F_1\ F_4\ F_3)(F_2)$	x_2^2	4^2
$(V_1\ V_3)(V_2\ V_4)$	$(F_1\ F_3)(F_2\ F_4)$	x_2^2	4^2
$(V_1\ V_4\ V_2)(V_3)$	$(F_1\ F_3\ F_2)(F_4)$	x_1x_3	$1 \cdot 4$
$(V_1)(V_2\ V_3\ V_4)$	$(F_1\ F_3\ F_4)(F_2)$	x_1x_3	$1 \cdot 4$

examples more mathematical in nature.

Example 2.4. We will find how many symmetrically distinct functions there are from the vertices of a regular tetrahedron to its faces. When we perform a symmetry on the tetrahedron it permutes both vertices and faces and thus fits our situation. In this example we will take the orientation preserving symmetries only. In Table 2.1 we compute the number of functions fixed by each (σ, τ) in the group. Adding up the last column and dividing by the order of the group we get $336/12 = 28$. So we get 28 symmetrically distinct functions from the vertices of the tetrahedron to the faces. By constructing the actual orbits with a computer, we verified this is the correct number.

Figure 2.2: Labeled Square

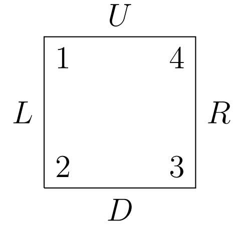

Table 2.2: Fixed Functions from Vertices to Sides of a Square

σ	τ	Z_σ	$Z_\sigma\left(p_1(\tau), ..., p_4(\tau)\right)$
(1)(2)(3)(4)	$(L)(R)(U)(D)$	x_1^4	4^4
(1)(2 4)(3)	$(L\ U)(D\ R)$	$x_1^2x_2$	$0^2 4$
(1 3)(2)(4)	$(L\ D)(U\ R)$	$x_1^2x_2$	$0^2 4$
(1 4)(2 3)	$(L\ R)(U)(D)$	x_2^2	4^2
(1 3)(2 4)	$(U\ D)(L\ R)$	x_2^2	4^2
(1 2)(3 4)	$(U\ D)(L)(R)$	x_2^2	4^2
(1 2 3 4)	$(U\ L\ D\ R)$	x_4	4
(1 4 3 2)	$(U\ R\ D\ L)$	x_4	4

Example 2.5. Suppose we take X to be the set of vertices of a square and Y to be the set of sides of the same square. We will count the number of symmetrically distinct functions from the vertices to the sides of a square. In this example we will consider reflections as symmetries. So the permutations on X alone would be D_4, but each permutation on the vertices has a corresponding permutation on the sides. In the Table 2.2 we compute the number of functions fixed for each (σ, τ) in the group. Here we get the number of orbits to be $312/8 = 39$. Again we verified this count by computer calculation.

Examples 2.4 and 2.5 would not be solvable by the Power Group Enumeration Theorem alone since in both cases the group is not a direct product. In the case of Example 2.5, if we had assumed the group was a direct product and applied the Power Group Enumeration Theorem, we would find that the computed number of orbits is too small. The group of permutations on domain and codomain is D_4 in each case. We know

$$Z_{D_4}(x_1, ..., x_4) = \frac{1}{8}\left(x_1^4 + 2x_2x_1^2 + 3x_2^2 + 2x_4\right).$$

So applying the Power Group Enumeration Theorem we get a total of 13 orbits. This is too few orbits because by making our group a direct product, it added more symmetries and thus reduced the number of symmetrically distinct functions. Assuming the group was a direct product would correspond to the physical situation of mapping the vertices of a square to the sides of a square that is totally separate from the first. When you are trying to map between objects in the same square, the Power Group Enumeration Theorem no longer holds.

Thus Theorem 2.3 is an improvement on the Power Group Enumeration Theorem to allow the group to no longer be a direct product. We could do similar computations to the ones we did with the square and tetrahedron for the cube. However, with the larger number of vertices, faces, sides, and symmetries it makes the computation of the orbits more lengthy so we will not show the details here. Trying to count these orbits for a cube by brute force would be difficult and require a large amount of computational power. Theorem 2.3 provides a much more computationally efficient method of counting such orbits.

2.4 Breaking the Orbits Into Sub-classes

Now that we have a way of counting orbits, we recall that one variation on Pólya's Theorem allowed us to find the pattern inventory. This idea is interesting since it tells us not only how many orbits there are, but also what kinds of orbits there are. In terms of coloring, we would find exactly how many symmetrically distinct colorings there are with a specific number of each color. In the more general group action we defined in Section 2.1, we are simultaneously switching the colors as well as the sites. Thus orbits may not have constant configurations of colorings. We wonder if there is anything similar to the pattern inventory that we can do when orbits are no longer identifiable by which elements of the codomain are mapped to and how many times. Afterwards, we will look specifically at orbits of bijective functions and develop a formula for counting them.

To break the orbits into sub-classes we need an attribute of the functions that is preserved by the action $*$. We will consider how many times each element of Y gets mapped to, but no longer care about which elements in particular are mapped to. To mathematically capture this idea for each $f \in Y^X$ we consider

$$\sum_{y \in Y} \left| f^{-1}(y) \right| = n.$$

Ignoring zeros in the above equation, we have a partition of n. This partition of n shows how many pre-images each element of Y has without specifying which particular elements of Y have those specific sizes of pre-images. Additionally, the partition of n is preserved by the group action. This is because in the below equation the equalities at both (1) and (3) have equal summands and in the equality at (2) the order with which the sum is taken is all that is changing.

$$\sum_{y \in Y} \left| (\tau f \sigma^{-1})^{-1}(y) \right| \overset{(1)}{=} \sum_{y \in Y} \left| \sigma f^{-1} \tau^{-1}(y) \right| \overset{(2)}{=} \sum_{y \in Y} \left| \sigma f^{-1}(y) \right| \overset{(3)}{=} \sum_{y \in Y} \left| f^{-1}(y) \right|.$$

We can then get something akin to the pattern inventories that we got when using Pólya's Theorem. However, it is not quite as elegant since now functions cannot be distinguished by exactly which elements of codomain are mapped to. We can instead ask how many orbits correspond to a given partition of n. As is customary at this point, we will want to plug polynomials into something like the cycle index polynomial to get the answer.

Theorem 2.6. *Let G be a subgroup of $S_X \times S_Y$ and let $*$ be the action of G on Y^X. Define*

$$P(y_1, ..., y_m) = \frac{1}{|G|} \sum_{(\sigma,\tau) \in G} Z_\sigma\big(p_1(\tau), ..., p_n(\tau)\big)$$

where $p_k(\tau)(y_1, ..., y_m) = \sum_{\ell(\tau_j)|k} \ell(\tau_j) \left(\prod_{y'_t \in \widehat{\tau_j}} y_t^{k/\ell(\tau_j)} \right)$. When $P(y_1, ..., y_m)$ is expanded, all the monomials in the polynomial are of degree n. Thus the powers on the y_t's for each monomial are a partition of n. The number of orbits of functions corresponding to a particular *partition of n is the sum of all coefficients of the monomials of $P(y_1, ..., y_m)$ with the same partition of n in the powers of the y_t's.*

Proof. We can restrict Burnside's lemma to the set of functions corresponding to a specific partition of n. Thus to count the number of orbits corresponding to a particular partition of n we merely need to count how many functions corresponding to that partition are fixed by each element of G.

To be fixed by an element (σ, τ) of G, the function must send each cycle of σ to a cycle of τ whose length divides the length of the cycle of σ as in Lemma 2.1. For a given cycle σ_i of length k we can choose any τ_j with $\ell(\tau_j)|k$ and map $\widehat{\sigma_i}$ into $\widehat{\tau_j}$ in a way preserved by the group action. Once a cycle τ_j is chosen, there are $\ell(\tau_j)$ ways of making a fixed function from $\widehat{\sigma_i}$ to $\widehat{\tau_j}$. Furthermore by our argument in Lemma 2.1 each such function has the property that each element in $\widehat{\tau_j}$ has $\ell(\sigma_i)/\ell(\tau_j) = k/\ell(\tau_j)$ pre-images from $\widehat{\sigma_i}$. Thus the monomial

$$\prod_{y_t' \in \widehat{\tau_j}} y_t^{k/\ell(\tau_j)}$$

represents the configuration of outputs for a particular fixed function from $\widehat{\sigma_i}$ to $\widehat{\tau_j}$. As we mentioned, there are $\ell(\tau_j)$ such fixed functions by Lemma 2.1. So the mononmial

$$\ell(\tau_j) \prod_{y_t' \in \widehat{\tau_j}} y_t^{k/\ell(\tau_j)}$$

represents all of the configurations of the fixed functions from $\widehat{\sigma_i}$ to $\widehat{\tau_j}$. Adding up all possible results for different possible $\tau_j's$ where the length divisibility condition is met, gives $p_k(\tau)$ as in the statement of the theorem. Thus $p_k(\tau)$ is essentially a sum of the different possible configurations for fixed functions from $\widehat{\sigma_i}$ to Y.

Now that we have the possible fixed configurations for each σ_i, we multiply them together which is the same as plugging the $p_k(\tau)$'s into the cycle index mononmial. Distribution of this product involves picking one term from each sum to multiply so this is the same as constructing the fixed functions cycle by cycle from σ. Thus we get that the configurations

Table 2.3: Inventory of Fixed Functions from Vertices to Sides of a Square

σ	τ	Z_σ	$Z_\sigma(p_1(\tau), ..., p_4(\tau))$
(1)(2)(3)(4)	(L)(R)(U)(D)	x_1^4	$(L+R+U+D)^4$
(1)(2 4)(3)	(L U)(D R)	$x_1^2x_2$	$0^2(2LU+2DR)$
(1 3)(2)(4)	(L D)(U R)	$x_1^2x_2$	$0^2(2LD+2UR)$
(1 4)(2 3)	(L R)(U)(D)	x_2^2	$(2LR+U^2+D^2)^2$
(1 3)(2 4)	(U D)(L R)	x_2^2	$(2UD+2LR)^2$
(1 2)(3 4)	(U D)(L)(R)	x_2^2	$(2UD+L^2+R^2)^2$
(1 2 3 4)	(U L D R)	x_4	$4ULDR$
(1 4 3 2)	(U R D L)	x_4	$4URDL$

of all functions fixed by (σ, τ) are in that distributed sum. Thus when $P(y_1, ..., y_m)$ is expanded, it is the sum of configurations of functions fixed for each (σ, τ) divided by $|G|$.

Finally, we just need to know how many fixed functions there are corresponding to a specific partition of n and then divide it by $|G|$. So we can look at the sum of all configurations of fixed functions divided by $|G|$, which we have established is $P(y_1, ..., y_m)$. The ones we desire are the ones whose powers on the y_t are the same partition of n as our target partition. So applying Burnside's lemma, we get the desired result. □

While the computation of $P(y_1, ..., y_m)$ can be quite tedious and messy, computers can be used to obtain the desired results. As an example, we will use Theorem 2.6 in the same situation as Example 2.5.

Example 2.7. Again we are counting symmetrically distinct functions from the vertices of a square to its sides. Before we had calculated there are 39 distinct orbits, but now we can obtain much more detail on what types of orbits there are. In Table 2.3 we show the computation of $P(U, D, L, R)$. Distributing and adding up the polynomials in the far right column, we obtain the rather impressive polynomial shown on the next page. It is organized according to the partition of 4 in the exponents of the variables.

Table 2.4: Partition Inventory for Functions from Vertices to Sides of a Square

Partition of 4	Number of Orbits
4	1
3+1	6
2+2	7
2+1+1	20
1+1+1+1	5

$$\frac{1}{8}\Big(2D^4 + 2L^4 + 2R^4 + 2U^4 +$$
$$4D^3L+4D^3R+4D^3U+4DL^3+4DR^3+4DU^3+4L^3R+4L^3U+4LR^3+4LU^3+4R^3U+4RU^3+$$
$$6D^2L^2 + 6D^2R^2 + 16D^2U^2 + 16L^2R^2 + 6L^2U^2 + 6R^2U^2+$$
$$16D^2LR + 12D^2LU + 12D^2RU + 12DL^2R + 16DL^2U + 12DLR^2 + 12DLU^2 + 16DR^2U+$$
$$12DRU^2 + 12L^2RU + 12LR^2U + 16LRU^2 + 40DLRU\Big)$$

Table 2.4 shows the number of orbits corresponding to each partition of 4. The $2+1+1$ partition orbits correspond to functions that send two vertices to the same side and the other vertices to different sides. Note that the number of orbits for each partition add up to 39 which is a relief because of our previous calculations. Additionally, we verified by computer calculation that the counts in Table 2.4 are correct. One partition of note is $1+1+1+1$ which corresponds to the bijective functions between the vertices and sides of the square. In Figure 2.3 we show the representatives of the five bijective orbits that we concluded there would be.

Counting the number of bijective orbits is a question of interest and we consider whether it could be computed more easily without that large polynomial. Additionally, we will verify our calculations about the bijective orbits in Example 2.7 in an alternative way. In the case that $n = m$ we have that bijective functions exist. First note that the set of bijective functions

Figure 2.3: Representatives of Bijective Orbits of Functions from Vertices to Sides of a Square

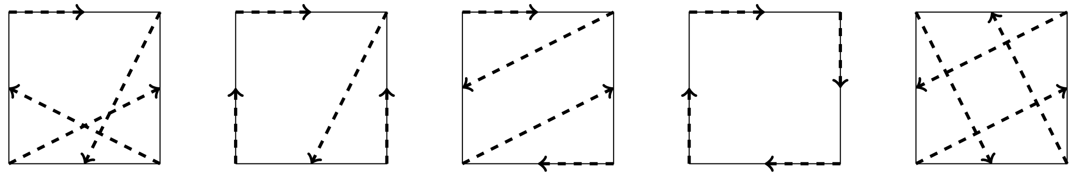

is preserved by the group action since we are just composing bijective functions. Thus some orbits will be made up purely of bijective functions and we can count how many there are. De Bruijn actually dealt with this situation in his paper [2], but he did so assuming the group was a direct product of subgroups. We will do it in the case of an arbitrary subgroup G of $S_X \times S_Y$.

Proposition 2.8. *Let G be a subgroup of $S_X \times S_Y$ with the action $*$ on Y^X. The number of orbits composed of bijective functions is given by*

$$\frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \delta(\sigma,\tau) j_1(\sigma)! \cdots j_n(\sigma)! Z_\sigma(1,2,...,n).$$

Here we are taking the delta function to be defined by the rule

$$\delta(\sigma,\tau) = \begin{cases} 1 & \text{if } j_s(\sigma) = j_s(\tau) \ \forall \ s \in \mathbb{N} \\ 0 & \text{otherwise.} \end{cases}$$

Proof. We can restrict Burnside's lemma to the collection of bijective functions B in Y^X. Thus to solve the problem of how many bijective orbits there are we merely need to count the number of bijective functions fixed by a given $(\sigma,\tau) \in G$. We will show that

$$\left|B^{(\sigma,\tau)}\right| = \delta(\sigma,\tau) j_1(\sigma)! \cdots j_n(\sigma)! Z_\sigma(1,2,...,n) \tag{2.2}$$

which will give us the desired result.

We know for a function to be fixed it must send the cycles of σ to the cycles of τ with the divisibility condition holding. In the case of bijective functions we must bijectively map the cycles of σ to distinct cycles of τ. In other words, in order for it to be possible to make a bijective function fixed by (σ, τ) we must have that σ and τ have exactly the same cycle type. This explains δ in (2.2) because when σ and τ have different cycle types we will get no bijective functions fixed by (σ, τ).

In the case σ and τ have exactly the same cycle structure, fixed bijective functions exist. Each k-cycle of σ needs to be bijectively mapped to a unique k-cycle of τ. When a k-cycle τ_j is chosen in τ there are $\ell(\tau_j) = k$ ways of creating a fixed bijective function into it from a k-cycle of σ. Additionally, there are $j_k(\tau)! = j_k(\sigma)!$ ways of matching the k-cycles of σ with the k-cycles of τ. Thus we get there are $j_k(\sigma)! k^{j_k(\sigma)}$ ways of bijectively assigning the elements of the k-cycles to each other in a way that is preserved under the action by (σ, τ). Multiplying these for each k we get that

$$\begin{aligned}\left|B^{(\sigma,\tau)}\right| &= j_1(\sigma)! \cdots j_n(\sigma)! 1^{j_1(\sigma)} 2^{j_2(\sigma)} \cdots n^{j_n(\sigma)} \\ &= j_1(\sigma)! \cdots j_n(\sigma)! Z_\sigma(1, 2, ..., n).\end{aligned}$$

Thus (2.2) holds. The proposition follows by using (2.2) and Burnside's lemma together. □

We can now apply Proposition 2.8 to Example 2.7 as an alternative and less messy way of getting the number of bijective orbits. The group structure yields the following computation

$$\frac{4! + 2! \cdot 2^2 + 4 + 4}{8} = 5.$$

This is good because it lines up with our previous computations. So Theorem 2.8 allows us to compute the number of symmetrically distinct bijective functions without doing excessive polynomial multiplication.

Thus we have been able to improve on the Power Group Enumeration Theorem allowing us to compute something akin to a pattern inventory of orbits in this more general situation.

We were also able to develop a formula to count the bijective orbits of functions. In the next chapter we will be shifting our focus from orbits of functions to orbits of relations.

Chapter 3. Orbits of Relations

One way we can further extend our results about the Power Group Enumeration Theorem is to make the set of objects we are acting on larger. Previously, our group action had been defined on the set of functions Y^X. In this chapter we will discuss how similar group actions can be defined on the set of all relations between X and Y, $\mathcal{P}(X \times Y)$. Clearly Y^X is contained in $\mathcal{P}(X \times Y)$ so this is similar problem, but we will have to modify our thinking a little. We will start by defining a group action. Then we will describe a method for counting the induced orbits on the set of all relations. Additionally, we will apply this thinking to provide another justification for our results that count orbits of functions.

3.1 Group Action Definition

First we describe a similar action to what has been done before for functions, now on the set of all relations. Let G be a subgroup of $S_X \times S_Y$. We extend the action of G under $*$ to $\mathcal{P}(X \times Y)$ as follows. For $(\sigma, \tau) \in G$ and R a subset of $X \times Y$, define

$$(\sigma, \tau) * R = \Big\{ \big(\sigma(x), \tau(y)\big) : (x, y) \in R \Big\}.$$

That this is a group action is not hard to verify. We will check those details for the sake of completeness. If $(\sigma, \tau), (\sigma', \tau') \in G$ we have the following

$$\begin{aligned}
(\sigma, \tau) * \big[(\sigma', \tau') * R\big] &= (\sigma, \tau) * \Big\{ \big(\sigma'(x), \tau'(y)\big) : (x, y) \in R \Big\} \\
&= \Big\{ \big(\sigma \circ \sigma'(x), \tau \circ \tau'(y)\big) : (x, y) \in R \Big\} \\
&= (\sigma \circ \sigma', \tau \circ \tau') * R \\
&= \big[(\sigma, \tau)(\sigma', \tau')\big] * R.
\end{aligned}$$

That the identity, (id_X, id_Y), fixes R is clear so $*$ defines a group action on $\mathcal{P}(X \times Y)$. Note that in the case where we restrict this action to Y^X we get exactly the same action as was used in Chapter 2. Thus we are justified in using the same symbol $*$. Since we have a group action, the equivalence classes partition the relations between X and Y into orbits. Our goal will be to count how many orbits there are. A similar definition of the group action can be defined on the relations between a finite collection of sets. We will explore this more in Section 3.5. For the time being we will develop some results and machinery to count the number of orbits of relations between two sets.

3.2 Cyclic Fixed Relations Between Two Sets

To count the number of orbits of relations we know from Burnside's lemma that we just need to count the number of relations fixed by a given $(\sigma, \tau) \in G$. To help do that we will establish an object called the cyclic fixed relation. We will find that cyclic fixed relations are the building blocks of every relation fixed by (σ, τ) and will provide a natural way count the total number of fixed relations. Once the fixed relations are counted, we can apply Burnside's lemma to obtain a theorem counting the number of orbits in the next section.

The motivation for defining the cyclic fixed relation is that it should be the smallest relation fixed by (σ, τ) containing a particular point $(x, y) \in X \times Y$. We define the *cyclic fixed relation* as

$$R_{(\sigma,\tau)}(x, y) = \Big\{ \big(\sigma^n(x), \tau^n(y)\big) : n \in \mathbb{Z} \Big\}. \tag{3.1}$$

From this definition it follows that $R_{(\sigma,\tau)}(x, y)$ is a relation invariant under action by (σ, τ). It is also the smallest relation fixed by (σ, τ) that contains the point (x, y). We will now prove a foundational result about the cyclic fixed relations.

Proposition 3.1.

i) Two cyclic fixed relations of (σ, τ) are either disjoint or equal.

ii) If $x_1, x_2 \in X$ are in different cycles of σ or $y_1, y_2 \in Y$ are in different cycles of τ then

$$R_{(\sigma,\tau)}(x_1, y_1) \neq R_{(\sigma,\tau)}(x_2, y_2).$$

Proof.

i) Suppose that (x', y') is in both $R_{(\sigma,\tau)}(x_1, y_1)$ and $R_{(\sigma,\tau)}(x_2, y_2)$. Then we have for some $k, j \in \mathbb{Z}$,

$$\left(\sigma^k(x_1), \tau^k(y_1)\right) = (x', y') = \left(\sigma^j(x_2), \tau^j(y_2)\right).$$

Thus we obtain $\left(\sigma^{k-j}(x_1), \tau^{k-j}(y_1)\right) = (x_2, y_2)$ and $\left(\sigma^{j-k}(x_2), \tau^{j-k}(y_2)\right) = (x_1, y_1)$. Therefore we have that $(x_1, y_1) \in R_{(\sigma,\tau)}(x_2, y_2)$. However, $R_{(\sigma,\tau)}(x_1, y_1)$ is the smallest relation fixed by (σ, τ) containing (x_1, y_1). Thus

$$R_{(\sigma,\tau)}(x_1, y_1) \subset R_{(\sigma,\tau)}(x_2, y_2).$$

The other inclusion follows similarly.

ii) Assume x_1 and x_2 are in different cycles of σ. It is clear that $\sigma^k(x_1) \neq x_2$ for all k. So it follows that $(x_2, y_2) \notin R_{(\sigma,\tau)}(x_1, y_1)$. A similar argument holds when looking at the case where y_1 and y_2 are in different cycles of τ. □

We now get to establish why the cyclic fixed relation is so foundational to the collection of all fixed relations. For any relation R fixed by (σ, τ) the following is true:

$$R = \cup_{(x,y)\in R} \left[R_{(\sigma,\tau)}(x, y)\right]. \tag{3.2}$$

We can eliminate repetition in the union of (3.2), because cyclic fixed relations are either disjoint or equal. Thus we get that every relation fixed by (σ, τ) is a unique union of cyclic fixed relations. So cyclic fixed relations are the building blocks of every fixed relation.

We will consider whether every union of cyclic fixed relations is also fixed by (σ, τ). If so, we will have a convenient way of constructing all relations fixed by (σ, τ). It is not too difficult to see that every union of cyclic fixed relations is also fixed by (σ, τ), but we will show a slightly stronger result to justify it.

Proposition 3.2. *The finite union of relations fixed by (σ, τ) is also fixed by (σ, τ).*

Proof. Let R_1 and R_2 be relations between X and Y that are fixed by (σ, τ). Then observe that

$$\begin{aligned}(\sigma, \tau) * (R_1 \cup R_2) &= \Big\{ \big(\sigma(x), \tau(y)\big) : (x, y) \in R_1 \cup R_2 \Big\} \\ &= [(\sigma, \tau) * R_1] \cup [(\sigma, \tau) * R_2] \\ &= R_1 \cup R_2.\end{aligned}$$

The general result then follows by induction. □

Now with that minor detail out of the way, we know that any union of a collection of cyclic fixed relations is also fixed by (σ, τ). Thus to find the total number of relations fixed by (σ, τ) we now only need to find how many different cyclic fixed relations there are. Once we obtain that quantity, we can count how many ways there are to form a union out of different combinations of them. Specifically we obtain the number of fixed relations is two raised to the number of different cyclic fixed relations. Thus for a given $(\sigma, \tau) \in G$, our goal will be to count the number of different cyclic fixed relations as we vary the point (x, y). As a step in this direction, we will find the size of the cyclic fixed relations themselves.

Proposition 3.3. *Let $\sigma \in S_X$, $\tau \in S_Y$, $x \in \widehat{\sigma_i}$, and $y \in \widehat{\tau_j}$. Then we have that*

$$\big| R_{(\sigma,\tau)}(x, y) \big| = \operatorname{lcm}\big(\ell(\sigma_i), \ell(\tau_j)\big).$$

Proof. Let $a = \operatorname{lcm}\big(\ell(\sigma_i), \ell(\tau_j)\big)$. It is clear from the definition in equation (3.1) that $R_{(\sigma,\tau)}(x, y)$ is the same as $R_{(\sigma_i,\tau_j)}(x, y)$. We claim that

$$R_{(\sigma_i,\tau_j)}(x, y) = \Big\{ (x, y), \big(\sigma_i(x), \tau_j(y)\big), ..., \big(\sigma_i^{a-1}(x), \tau_j^{a-1}(y)\big) \Big\}. \tag{3.3}$$

First we must show that all of the elements in the set on the right hand side of (3.3) are distinct. Suppose that

$$\sigma_i^b(x) = \sigma_i^c(x) \text{ and } \tau_j^b(y) = \tau_j^c(y).$$

where b and c are nonnegative and less than a. Then without loss of generality we may assume $b \geq c$. Then we obtain

$${\sigma_i}^{b-c}(x) = x \text{ and } {\tau_j}^{b-c}(y) = y.$$

Thus σ_i^{b-c} and τ_j^{b-c} must both be the identity. So $b - c$ is a common multiple of both cycle lengths since the respective cycle lengths are the order of the cycles. Since a is the least common multiple of the cycle lengths, $b-c$ cannot be smaller than a or else it must be zero. But $b - c$ must be smaller than a because of the size of b and c thus we must have that $b - c = 0$. So all elements in the right hand side of (3.3) are distinct.

Now we show equality of the sets in (3.3). One inclusion is immediate from definition (3.1). For the other, let k be a nonzero integer. Then using the division algorithm we can write $k = aq + r$ for some $0 \leq r \leq a - 1$. Observe that

$$\begin{aligned}\left(\sigma_i^k(x), \tau_j^k(y)\right) &= \left(\sigma_i^{aq+r}(x), \tau_j^{aq+r}(y)\right)\\ &= \left(\sigma_i^r(x), \tau_j^r(y)\right).\end{aligned}$$

Thus (3.3) holds, immediately giving us the desired result. $\square$

Proposition 3.3 is not too surprising since the order of the group element $\sigma_i\tau_j$ in $S_{X\cup Y}$ is also the least common multiple of the lengths of σ_i and τ_j. One consequence of Proposition 3.3 is that if $x_1, x_2 \in \widehat{\sigma_i}$ and $y_1, y_2 \in \widehat{\tau_j}$, then $\left|R_{(\sigma,\tau)}(x_1, y_1)\right| = \left|R_{(\sigma,\tau)}(x_2, y_2)\right|$. This fact will be very useful moving forward. We now have everything we need to answer the question of how many different cyclic fixed relations there are.

Proposition 3.4. *Let σ_i and τ_j be cycles of σ and τ, respectively. The number of different cyclic fixed relations $R_{(\sigma,\tau)}(x,y)$ with x in $\widehat{\sigma_i}$ and y in $\widehat{\tau_j}$ is* $\gcd\big(\ell(\sigma_i), \ell(\tau_j)\big)$.

Proof. From Proposition 3.1 i) we know the cyclic fixed relations from points in $\widehat{\sigma_i} \times \widehat{\tau_j}$ partition $\widehat{\sigma_i} \times \widehat{\tau_j}$. By Proposition 3.3 we also know each cyclic fixed relation is the same size namely $\operatorname{lcm}\big(\ell(\sigma_i), \ell(\tau_j)\big)$. Thus it follows that the number of cyclic fixed relations with the property described above is

$$\frac{|\widehat{\sigma_i} \times \widehat{\tau_j}|}{\operatorname{lcm}\big(\ell(\sigma_i), \ell(\tau_j)\big)} = \frac{\ell(\sigma_i)\ell(\tau_j)}{\operatorname{lcm}\big(\ell(\sigma_i), \ell(\tau_j)\big)} = \gcd\big(\ell(\sigma_i), \ell(\tau_j)\big). \qquad \square$$

With these tools in hand we finally obtain an equation for the number of relations fixed by a $(\sigma, \tau) \in G$ which was our goal.

Lemma 3.5. *Let $(\sigma, \tau) \in S_X \times S_Y$ and let $*$ be the group action on $\mathcal{P}(X \times Y)$. Then the following holds:*

$$\left|\mathcal{P}(X \times Y)^{(\sigma,\tau)}\right| = \exp\left(\ln(2) \sum_{i,j} \gcd\big(\ell(\sigma_i), \ell(\tau_j)\big)\right).$$

Proof. We know from Proposition 3.1 ii) that cyclic fixed relations generated from points in different cycles of domain or codomain create distinct cyclic fixed relations. By Proposition 3.4 we know the number of cyclic fixed relations when points are chosen from the direct product of the elements in a cycle in the domain and a cycle in the codomain is the greatest common divisor of their lengths. Thus adding up all possible greatest common divisors of pairs of lengths of cycles from the domain and codomain, yeilds the total number of distinct cyclic fixed relations. We know all fixed relations by (σ, τ) are unions of these $\sum_{i,j} \gcd\big(\ell(\sigma_i), \ell(\tau_j)\big)$ cyclic fixed relations. The number of unions that can be formed is two raised to the number of cyclic fixed relations. Thus

$$\left|\mathcal{P}(X \times Y)^{(\sigma,\tau)}\right| = \exp\left(\ln(2) \sum_{i,j} \gcd\big(\ell(\sigma_i), \ell(\tau_j)\big)\right). \qquad \square$$

Having a count of the number of relations fixed by a group element and recalling Burnside's lemma, we are ready to count the orbits of relations.

3.3 Counting the Orbits of Relations on Two Sets

We know from Burnside's lemma that we only need to count the number of relations fixed by a particular (σ, τ) in order to get the number of orbits. In the last section we counted the number relations fixed by a particular (σ, τ). Our method for getting this was to notice every fixed relation can be written as a union of cyclic fixed relations. Since these are disjoint or equal we just needed to count the number of different cyclic fixed relations as we varied the point (x, y). Then it was easy to count how many different unions we could form. This thinking leads to the following result.

Theorem 3.6. *Let G be a subgroup of $S_X \times S_Y$. Then*

$$|\mathcal{P}(X \times Y)/G| = \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} Z_\sigma\big(p_1(\tau), ..., p_n(\tau)\big)$$

where $p_k(\tau) = Z_\tau\left(2^{\gcd(k,1)}, ..., 2^{\gcd(k,m)}\right)$.

Proof. This result is obtained applying Burnside's lemma and Lemma 3.5 with some algebraic manipulations as shown on the next page.

$$
\begin{aligned}
|\mathcal{P}(X \times Y)/G| &= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \left|\mathcal{P}(X \times Y)^{(\sigma,\tau)}\right| \\
&= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \exp\left(\ln(2) \sum_{i,j} \gcd\big(\ell(\sigma_i), \ell(\tau_j)\big)\right) \\
&= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \exp\left(\sum_{k=1}^{n} \ln(2) \sum_{j} j_k(\sigma) \gcd\big(k, \ell(\tau_j)\big)\right) \\
&= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \prod_{k=1}^{n} \exp\left(\ln(2) \sum_{j} j_k(\sigma) \gcd\big(k, \ell(\tau_j)\big)\right) \\
&= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} \prod_{k=1}^{n} \left(\exp\left(\ln(2) \sum_{j} \gcd\big(k, \ell(\tau_j)\big)\right)\right)^{j_k(\sigma)} \\
&= \frac{1}{|G|} \sum_{(\sigma,\tau)\in G} Z_\sigma\left(2^{\sum_j \gcd(1,\ell(\tau_j))}, ..., 2^{\sum_j \gcd(n,\ell(\tau_j))}\right).
\end{aligned}
$$

To get the result as stated we just need to note that for each k,

$$
\begin{aligned}
2^{\sum_j \gcd\left(k,\ell(\tau_j)\right)} &= \exp\left(\ln(2) \sum_{j} \gcd\big(k, \ell(\tau_j)\big)\right) \\
&= \exp\left(\ln(2) \sum_{t=1}^{m} j_t(\tau) \gcd(k,t)\right) \\
&= \prod_{t=1}^{m} \exp\big(\ln(2) j_t(\tau) \gcd(k,t)\big) \\
&= \prod_{t=1}^{m} \Big(\exp\big(\ln(2) \gcd(k,t)\big)\Big)^{j_t(\tau)} \\
&= Z_\tau\left(2^{\gcd(k,1)}, ..., 2^{\gcd(k,m)}\right). \qquad \square
\end{aligned}
$$

Thus we have a result that counts the number of orbits of relations. The helpfulness of this result can be seen in the following example. Many of the other examples we have seen in this paper have been easily verifiable by computer programming. However, the set of all relations gets large quickly as X and Y do. This is the first example where we can really appreciate how much more powerful these results are than performing a brute force method.

Table 3.1: Fixed Relations from Vertices to Sides of a Square

σ	τ	Z_σ	Z_τ	$Z_\sigma\left(p_1(\tau), ..., p_4(\tau)\right)$
$(1)(2)(3)(4)$	$(L)(R)(U)(D)$	x_1^4	x_1^4	$(2^4)^4$
$(1)(2\ 4)(3)$	$(L\ U)(D\ R)$	$x_1^2x_2$	x_2^2	$[(2^1)^2]^2 \cdot (2^2)^2$
$(1\ 3)(2)(4)$	$(L\ D)(U\ R)$	$x_1^2x_2$	x_2^2	$[(2^1)^2]^2 \cdot (2^2)^2$
$(1\ 4)(2\ 3)$	$(L\ R)(U)(D)$	x_2^2	$x_1^2x_2$	$[(2^1)^2 \cdot 2^2]^2$
$(1\ 3)(2\ 4)$	$(U\ D)(L\ R)$	x_2^2	x_2^2	$[(2^2)^2]^2$
$(1\ 2)(3\ 4)$	$(U\ D)(L)(R)$	x_2^2	$x_1^2x_2$	$[(2^1)^2 \cdot 2^2]^2$
$(1\ 2\ 3\ 4)$	$(U\ L\ D\ R)$	x_4	x_4	2^4
$(1\ 4\ 3\ 2)$	$(U\ R\ D\ L)$	x_4	x_4	2^4

Example 3.7. We will consider the familiar example of X being the set of vertices of a square and Y being the set of sides of that same square. Previously we had counted the number of symmetrically distinct functions from the vertices to the sides. This meant each vertex was sent to exactly one side. Relations for this example correspond to sending a vertex to possibly multiple sides or no sides at all. Before starting, we note that we have four vertices and four sides. Thus we have a cardinality of sixteen for $X \times Y$. This leads to a total of $2^{16} = 65,536$ possible relations that can be formed. This certainly would not be trivial to check computationally. However, with the aid of Theorem 3.6 the procedure becomes much simpler. In Table 3.1 we compute the number of fixed relations for each group element. This gives us the number of orbits to be $66848/8 = 8356$. Thus there are 8,356 symmetrically distinct ways to relate the vertices of a square with the sides. This is quite a bit smaller than the 65,536 relations we had to start with.

We have now solved the question of how many orbits of relations between two sets there are. In Chapter 2 we focused exclusively on orbits of functions. Since functions are also relations we wonder if we can apply our thinking about counting orbits of relations to count orbits of functions. This will be covered in the next section, and then we will go on to generalizing these ideas for relations among a finite collection of sets.

3.4 Alternate Proof for Function Enumeration Theorems Using Relations

We know Y^X is contained in $\mathcal{P}(X \times Y)$ so we ask ourselves if we can see our results about orbits of functions inside of our results about orbits of relations. This line of thinking will lead to a alternative proof of the results we covered in Chapter 2, which were adapted from Harary and Palmer's work. We recall our method of counting orbits of relations was to look at the cyclic fixed relations. This leads us to ask about classifying when a cyclic fixed relation will be a function.

Proposition 3.8. *Let $\sigma \in S_X$ and $\tau \in S_Y$ with $x \in \widehat{\sigma_i}$ and $y \in \widehat{\tau_j}$. Then $R_{(\sigma,\tau)}(x, y)$ is a function defined from $\widehat{\sigma_i}$ to Y if and only if $\ell(\tau_j)|\ell(\sigma_i)$.*

Proof.
$(\implies)$: If the cyclic fixed relation $R_{(\sigma,\tau)}(x, y)$ is a function defined on $\widehat{\sigma_i}$, then

$$\left|R_{(\sigma,\tau)}(x, y)\right| = \ell(\sigma_i).$$

But we know by Proposition 3.3 that $|R_{(\sigma,\tau)}(x, y)|$ is actually the least common multiple of $\ell(\sigma_i)$ and $\ell(\tau_j)$. Thus $\ell(\sigma_i)$ must be a multiple of $\ell(\tau_j)$. So $\ell(\tau_j)|\ell(\sigma_i)$.
$(\impliedby)$: If $\ell(\tau_j)|\ell(\sigma_i)$ we get again that $|R_{(\sigma,\tau)}(x, y)| = \ell(\sigma_i)$. It is clear that the first coordinate in the cyclic fixed relation ranges between all elements in $\widehat{\sigma_i}$. So each x in $\widehat{\sigma_i}$ must be used exactly once as the first coordinate in a point in $R_{(\sigma,\tau)}(x, y)$. Thus $R_{(\sigma,\tau)}(x, y)$ is a function defined on the elements in the cycle σ_i. □

This is the connection we need between relations and functions that make it possible to prove Lemma 2.2 in an alternate way. That result was essential to our extension of the Power Group Enumeration Theorem from Chapter 2. We do know a function fixed by (σ, τ) is also a relation and is thus a union of cyclic fixed relations (which will have to also be functions). Thus we need to count the number of fixed functions from each cycle of σ to the elements of Y. Let σ_i be a cycle of σ. Using Proposition 3.4 we know the number of fixed functions from $\widehat{\sigma_i}$ to Y is

$$\sum_{\ell(\tau_j)|\ell(\sigma_i)} \gcd\left(\ell(\sigma_i), \ell(\tau_j)\right) = \sum_{\ell(\tau_j)|\ell(\sigma_i)} \ell(\tau_j).$$

Then the fixed functions for each of the cycles of σ can be glued together to make a function from all of X to Y that is fixed by (σ, τ). All fixed functions can be constructed this way. So it follows that in total the number of fixed functions will be

$$\left|(Y^X)^{(\sigma,\tau)}\right| = \prod_i \sum_{\ell(\tau_j)|\ell(\sigma_i)} \ell(\tau_j).$$

Note the above formula is the same we obtained in Lemma 2.1. Thus we can connect cyclic fixed relations back to solving the problem of finding orbits of functions as well.

3.5 Counting the Orbits of Relations on More Sets

As one final extension of what we have been doing, we apply the concepts previously discussed to a collection of more than two sets. Let $X_1, ..., X_q$ all be finite sets and G a subgroup of $S_{X_1} \times \cdots \times S_{X_q}$. Here we similarly define the group action $*$. With $\left(\sigma^{(1)}, ..., \sigma^{(q)}\right) \in G$, we define the action on a relation $R \in \mathcal{P}\left(X_1 \times \cdots \times X_q\right)$ by

$$\left(\sigma^{(1)}, ..., \sigma^{(q)}\right) * R = \left\{\left(\sigma^{(1)}(x_1), ..., \sigma^{(q)}(x_q)\right) : (x_1, ..., x_q) \in R\right\}. \tag{3.4}$$

Note that we are using parentheses in exponents to denote components in the direct product since subscripts have already been used in cycle decompositions. A similar argument to the case when $q = 2$, which we showed in Section 3.1, can be made showing that this is indeed a group action. We can again ask how many orbits of relations there are. Similar results to ones in Section 3.2 work by the same arguments just applied to longer tuples. We again can define the cyclic fixed relations the same way as before and they will be either disjoint or equal. We again have that

$$\left|R_{\left(\sigma^{(1)},...,\sigma^{(q)}\right)}(x_1,...x_q)\right| = \operatorname{lcm}\Bigg(\ell\left(\sigma^{(1)}_{k_1}\right),...,\ell\left(\sigma^{(q)}_{k_q}\right)\Bigg),$$

where each x_i is in the cycle $\sigma^{(i)}_{k_i}$ of $\sigma^{(i)}$.

An analog of Proposition 3.4 can be obtained in this more general circumstance. The number of cyclic fixed relations by $\left(\sigma^{(1)},...,\sigma^{(q)}\right)$ coming from points in $\widehat{\sigma^{(1)}_{k_1}} \times \cdots \times \widehat{\sigma^{(q)}_{k_q}}$ is

$$\frac{\ell\left(\sigma^{(1)}_{k_1}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(\ell\left(\sigma^{(1)}_{k_1}\right),...,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}. \tag{3.5}$$

Unfortunately, we cannot transform this to a greatest common divisor as we did before. This is because for more than two numbers that result does not hold. For instance, consider the set of numbers $\{2,2,4\}$. Taking the least common multiple we get four. However, the greatest common factor (two) is not the product of the numbers divided by the least common multiple. Thus we can tell our notation is going to get messier than in the case where we only had two sets.

Theorem 3.9. *Let G be a subgroup of $S_{X_1}\times\cdots\times S_{X_q}$ acting on $\mathcal{P}\left(X_1\times\cdots\times X_q\right)$ as defined in (3.4). Then*

$$\left|\mathcal{P}\left(X_1\times\cdots\times X_q\right)/G\right| = \frac{1}{|G|}\sum_{\left(\sigma^{(1)},...,\sigma^{(q)}\right)\in G} Z_{\sigma^{(1)}}\left(2^{p_1},...,2^{p_n}\right).$$

Where

$$p_k = p_k\left(\sigma^{(2)},...,\sigma^{(q)}\right) = \sum_{k_2,...,k_q}\frac{k\cdot\ell\left(\sigma^{(2)}_{k_2}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(k,\ell\left(\sigma^{(2)}_{k_2}\right),...,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}.$$

Proof. For less messy notation we set $P=\mathcal{P}(X_1\times\cdots\times X_q)$. We know (3.5) expresses the number of cyclic fixed relations using points from $\widehat{\sigma^{(1)}_{k_1}}\times\cdots\times\widehat{\sigma^{(q)}_{k_q}}$. Thus adding the possibilities over all possible combinations of cycles yields the total number of cyclic fixed relations. The number of ways to form a union from these cyclic fixed relations will be two raised to that power. This yields the equation

$$\begin{aligned}
\left|P^{\left(\sigma^{(1)},\ldots,\sigma^{(q)}\right)}\right| &= \exp\left(\ln(2)\sum_{k_1,\ldots,k_q}\frac{\ell\left(\sigma^{(1)}_{k_1}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(\ell\left(\sigma^{(1)}_{k_1}\right),\ldots,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}\right)\\
&= \exp\left(\ln(2)\sum_{k=1}^{n}\left[j_k\left(\sigma^{(1)}\right)\sum_{k_2,\ldots,k_q}\frac{k\cdot\ell\left(\sigma^{(2)}_{k_2}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(k,\ell\left(\sigma^{(2)}_{k_2}\right),\ldots,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}\right]\right)\\
&= \prod_{k=1}^{n}\exp\left(\ln(2)j_k\left(\sigma^{(1)}\right)\sum_{k_2,\ldots,k_q}\frac{k\cdot\ell\left(\sigma^{(2)}_{k_2}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(k,\ell\left(\sigma^{(2)}_{k_2}\right),\ldots,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}\right)\\
&= \prod_{k=1}^{n}\left(\exp\left(\ln(2)\sum_{k_2,\ldots,k_q}\frac{k\cdot\ell\left(\sigma^{(2)}_{k_2}\right)\cdots\ell\left(\sigma^{(q)}_{k_q}\right)}{\operatorname{lcm}\left(k,\ell\left(\sigma^{(2)}_{k_2}\right),\ldots,\ell\left(\sigma^{(q)}_{k_q}\right)\right)}\right)\right)^{j_k\left(\sigma^{(1)}\right)}\\
&= Z_{\sigma^{(1)}}\left(2^{p_1},\ldots,2^{p_n}\right).
\end{aligned}$$

Then the desired result follows by applying Burnside's lemma. □

This formula that we developed is rather lengthy but it does accomplish its goal effectively. As a demonstration, we will look back at our example of the regular tetrahedron and apply Theorem 3.9. As it was in the case with only two sets, using this theorem is much more efficient than a brute force method.

Example 3.10. Observe the regular tetrahedron has vertices, faces, and sides which are all simultaneously permuted when performing symmetries on the structure. We can now ask and answer the question of how many symmetrically distinct relations are there between the vertices, faces, and sides. This is a massive problem computationally since the total number of relations will be 2^{96} which is approximately 7.9228163×10^{28}. Unphazed, in Table 3.2 we compute the number of relations fixed by each symmetry. Applying Theorem 3.9 we get the number of orbits to be about 6.6023468×10^{27}.

Thus Theorem 3.9 makes something that would be impossible by brute force actually quite routine. The solution to counting orbits of relations between two sets is clearly more

Figure 3.1: Labeled Tetrahedron with Vertices, Faces, and Sides

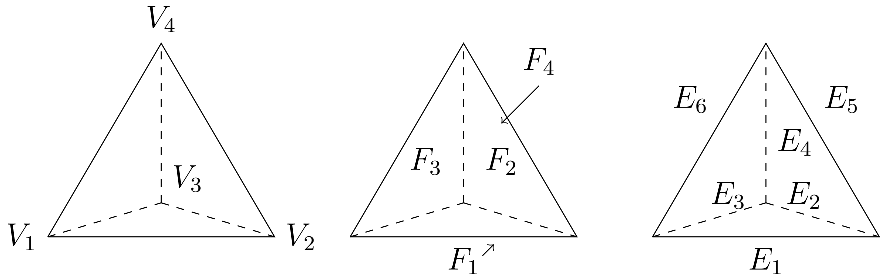

Table 3.2: Fixed Relations between Vertices, Faces, and Sides of a Tetrahedron

$\sigma^{(1)}$	$\sigma^{(2)}$	$\sigma^{(3)}$	$Z_{\sigma^{(1)}}$	$Z_{\sigma^{(1)}}(a_1, ..., a_4)$ where $a_k = 2^{p_k}$
$(V_1)(V_2)(V_3)(V_4)$	$(F_1)(F_2)(F_3)(F_4)$	$(E_1)(E_2)(E_3)(E_4)(E_5)(E_6)$	x_1^4	$(2^{24})^4$
$(V_1\ V_2\ V_3)(V_4)$	$(F_1)(F_2\ F_3\ F_4)$	$(E_1\ E_2\ E_3)(E_4\ E_6\ E_5)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1\ V_4\ V_3)(V_2)$	$(F_1\ F_4\ F_2)(F_3)$	$(E_1\ E_5\ E_2)(E_3\ E_6\ E_4)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1\ V_2\ V_4)(V_3)$	$(F_1\ F_2\ F_3)(F_4)$	$(E_1\ E_5\ E_6)(E_2\ E_4\ E_3)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1\ V_3\ V_2)(V_4)$	$(F_1)(F_2\ F_4\ F_3)$	$(E_1\ E_3\ E_2)(E_4\ E_5\ E_6)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1\ V_2)(V_3\ V_4)$	$(F_1\ F_4)(F_2\ F_3)$	$(E_1)(E_2\ E_6)(E_3\ E_5)(E_4)$	x_2^2	$(2^{24})^2$
$(V_1\ V_4)(V_2\ V_3)$	$(F_1\ F_2)(F_3\ F_4)$	$(E_1\ E_4)(E_2)(E_3\ E_5)(E_6)$	x_2^2	$(2^{24})^2$
$(V_1\ V_3\ V_4)(V_2)$	$(F_1\ F_2\ F_4)(F_3)$	$(E_1\ E_2\ E_5)(E_3\ E_4\ E_6)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1)(V_2\ V_4\ V_3)$	$(F_1\ F_4\ F_3)(F_2)$	$(E_1\ E_6\ E_3)(E_2\ E_5\ E_4)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1\ V_3)(V_2\ V_4)$	$(F_1\ F_3)(F_2\ F_4)$	$(E_1\ E_4)(E_2\ E_6)(E_3)(E_5)$	x_2^2	$(2^{24})^2$
$(V_1\ V_4\ V_2)(V_3)$	$(F_1\ F_3\ F_2)(F_4)$	$(E_1\ E_6\ E_5)(E_2\ E_3\ E_4)$	x_1x_3	$2^8 \cdot 2^{24}$
$(V_1)(V_2\ V_3\ V_4)$	$(F_1\ F_3\ F_4)(F_2)$	$(E_1\ E_3\ E_6)(E_2\ E_4\ E_5)$	x_1x_3	$2^8 \cdot 2^{24}$

elegant than this expanded situation of more sets. Either way, we can now count the number of orbits of relations as well as functions.

Chapter 4. Conclusion

Overall Pólya's Theorem allows one to determine how many symmetrically distinct ways of coloring something there are. In mathematical language, it allows us to count the orbits of a group action defined on the set of functions. De Bruijn, Harary, and Palmer generalized those results to allow simultaneous permutations of the domain and codomain of the functions. However, their results are somewhat limited in their applicability because they assumed the structure of the group must be a direct product of subgroups. We used the ideas in their publications to extend their results to the situation of the group not being a direct product. We also created a way of breaking the orbits down further by creating a correspondence between them and the partitions of n. Additionally we extended the ideas further and applied them to the set of all relations and not just functions. The results we provided can be used in several situations. However, there are a few more questions that we can ask beyond what we have answered in this paper.

In regard to orbits of functions there are several avenues we could explore further, but did not because of time.

Q1: We did not create weighted versions of the results, but this could be done.

Q2: We found a particular way of breaking the orbits into subclasses related to partitions of n. We could ask if there are other properties of the functions that would be preserved by the group action and attempt to count orbits with respect to those properties.

Q3: In the actual computations of the number of orbits there was a decent amount of repetition coming from group elements where the cycle structure is identical to other group elements. Is there a better way of expressing the formula to eliminate this redundancy?

Q4: Additionally, in our results we assumed nothing about the group structure itself besides it being a subgroup of $S_X \times S_Y$. If the set of permutations on the codomain had properties where it did not switch certain elements, could we use that to further break the orbits into subclasses more similar to standard Pólya inventories. More formally, if Y could be partitioned into sets Y_i such that the permutations on the codomain preserve the $Y_i's$, could we find the number of orbits of functions mapping ϵ_i elements from the domain into each Y_i? De Bruijn did this in the case of the groups being a direct product so we just wonder how to express the solution when the group is not necessarily a direct product.

Q5: The group action preserves both injective and surjective functions. In Chapter 2 we found a formula for the number of bijective orbits, but are there nice formula's for the injective and surjective orbits individually?

The situation of relations is slightly different than that of functions so we could consider several more things about counting orbits of relations.

Q6: The number of points in a relation is preserved by the group action. We wonder if there is a nice way of expressing how many orbits of relations there are with a given number of points in it? This problem seems like it may not have a nice solution based on our approach of using cyclic fixed relations, but perhaps there is a way of solving this problem in elegant manner that would yield a nice formula.

Q7: Again we wonder if our formulas can be improved to reduce repetition of calculations. It is clear that we could improve the statement of Theorem 3.9, but what we have works and changes to it may complicate notation.

Thus there are more avenues to explore in many of the things we have covered in this paper.

Bibliography

[1] Burnside, W. Theory of Groups of Finite Order. *Cambridge University Press* (1897).

[2] De Bruijn, N. G. Generalization of Polya's Fundamental Theorem in Enumerative Combinatorial Analysis. *Nederl. Akad Wetensch. Proc. Ser. A* **62**: 59-69 (1959).

[3] Frobenius, F. G. Ueber die Congruenz nach einem aus zwei endlichen Gruppen gebildeten Doppelmodul. *Crelle's Journal* **101** (4): 273–299 (1887).

[4] Harary, F. and Palmer, E. M. Graphical Enumeration. *Academic Press, Inc.* 32-50, 135-155 (1973).

[5] Harary, F. and Palmer, E. M. The Power Group Enumeration Theorem. *Journal of Combinatorial Theory* **1**: 157-173 (1966).

[6] McDaniel, D. Restatement of Pólya's Theorem. *Inorganic Chemistry* **11** (11): 2678-2682 (1972) DOI: 10.1021/ic50117a021.

[7] Merris, R. Combinatorics. *Wiley-Interscience Series in Discrete Mathematics and Optimization.* 175-249 (2003).

[8] Pólya, G. Kombinatorische Anzahlbestimmugen fűr Gruppen, Graphen, und chemische Verbindungen, *Acta Math.* **68**: 145-254 (1937).

[9] Redfield, J. H. The Theory of Group-Reduced Distributions. *American Journal of Mathematics*, **49** (3): 433-455 (1927).

[10] van Lint, J. H. and Wilson, R. M. A Course in Combinatorics. *Cambridge University Press*. 461-474 (1992).