UNDERSTANDING GENERALIZED GROUP SIEVING THROUGH λ-RINGS

ETHAN WILLIAMS

Abstract. The cyclic sieving phenomenon occurs when a combinatorial polynomial gives meaningful combinatorial values when evaluated at roots of unity. This paper will examine the phenomenon through a lens of both combinatorics and group representations. Then this paper will generalize the phenomenon to any group and consider the utility of a λ-ring to understand the phenomenon. Finally, this paper will consider which groups succeed and fail at producing interesting occurrences of sieving phenomena.

Contents

1. Introduction: A Curious Polynomial Phenomenon

In this paper, we will examine a phenomenon where polynomials capture information about group actions. We will start with one of the simplest cases: choosing corners of an m-gon and considering how rotations to the m-gon permute these choices. For the case of $m = 6$, Figures 1 and 2 show all possible choices of 2 and 3 corners as segments and triangles respectively.

For the case of 2 corners, the first two rows of Figure 1 show how we can rotate some pairs of points 6 times to get 6 unique pairs. However, the final 3 segments in the bottom row do not form a full set of 6 because after 3 rotations these segments repeat. Of the fifteen elements in the set of pairs of corners of a hexagon, two are fixed by a 180 degree rotation.

Like with the first case, for the case of 3 corners shown in Figure 2, the first three rows show how we can rotate some triangles 6 times to get 6 unique triangles, while the final 2 triangles in the bottom row do not form a full set of 6 because

Date: AUGUST 28, 2021.

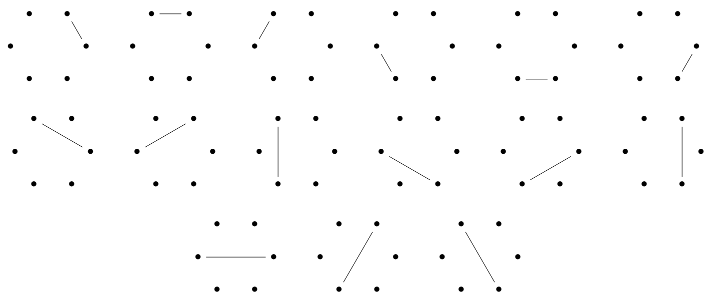

FIGURE 1. The 15 line segments joining corners of a hexagon

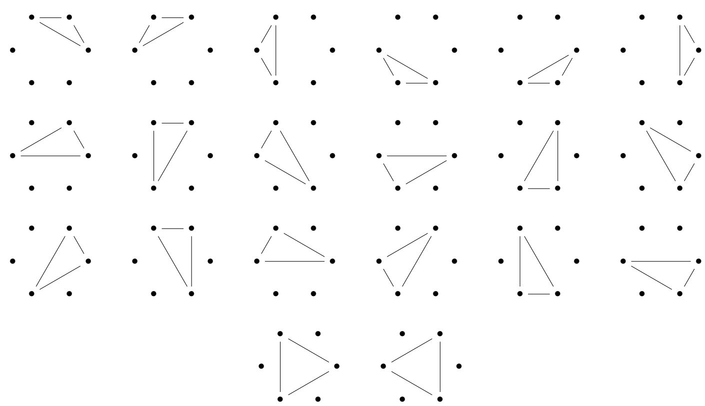

FIGURE 2. The 20 triangles joining corners of a hexagon

after 2 rotations these triangles repeat. Of the twenty elements in the set of sets of three corners of a hexagon, two are fixed by a 120 degree rotation.

For the general case of choosing n corners of an m-gon, there are $\binom{m}{n}$ elements in the set. Most combinations produce m unique choices as we rotate by $1/m$-th of a circle, but some repeat before a full rotation has been made. Calculating this is a combinatorial problem. For a configuration to have symmetry by a $1/k$ rotation, k must divide both m and n. If k does, then $\binom{m/k}{n/k}$ configurations will have symmetry by a $1/k$-th rotation. Reassuringly, all $\binom{m}{n}$ configurations are fixed by a full rotation when $k = 1$.

Let's use this method for the example of a dodecagon. For choosing 0 or 12 corners, there is one configuration which has all symmetries. For choosing 1, 5, 7, or 11 corners, we are able to rotate a set 12 times to get 12 unique configurations because these numbers are coprime to 12. For 2 or 10 corners, most combinations will be unique for each of 12 rotations, but there will be the $\binom{6}{1} = 6$ opposite pairs of points which are not unique by a $1/2$ rotation. For 3 or 9 points, there are $\binom{4}{1} = 4$ equilateral triangles which are not unique by a $1/3$ rotation. For 4 or 8 points, the $\binom{3}{1} = 3$ squares are not unique by a $1/4$ rotation but also any pair of opposite pairs, of which there are $\binom{6}{2} = 15$, is not unique by a $1/2$ rotation. Finally, for 6 points, we have the 2 regular hexagons fixed by a $1/6$ rotation, the $\binom{4}{2} = 6$

Number of Corners	0°	30°	60°	90°	120°	180°
0, 12	1	1	1	1	1	1
1, 11	12	0	0	0	0	0
2, 10	66	0	0	0	0	6
3, 9	220	0	0	0	4	0
4, 8	495	0	0	3	0	15
5, 7	792	0	0	0	0	0
6	924	0	2	0	6	20

TABLE 1. Subsets of n corners of a dodecagon fixed by rotation.

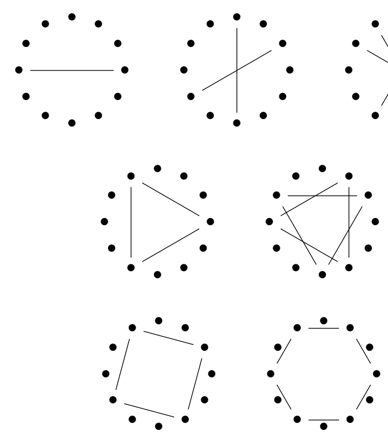

FIGURE 3. Top: Configurations of 1, 2, and 3 segments with 1/2 turn symmetry. Middle: Configurations of 1 and 2 triangles with 1/3 turn symmetry. Bottom: Configurations of a square or a hexagon with 1/4 and 1/6 turn symmetry.

pairs of equilateral triangles fixed by a 1/3 rotation, and the $\binom{6}{3} = 20$ sets of three pairs of opposite points fixed by a 1/2 rotation. These values are summarized in Table 1 above and diagrams of each described configuration are shown in Figure 3 above.

Now that we understand the pattern in this process, we are ready to investigate the polynomials which capture information about this process. We have used the binomial coefficients many times to calculate the fixed configurations, so it may not be too surprising to learn that our polynomials of interest are the q-binomial coefficients.

1.1. **The q-Binomial Coefficients.** The q-binomial coefficients are a series of polynomials with important combinatorial significance. One of their most common applications is counting the number of n-dimensional subspaces of $\mathbb{F}_q^m$ when q is a prime power. We will investigate these polynomials to discover some of the connections ourselves.

Definition 1.1. The *q-bracket* of an integer $n \geq 0$, denoted $[n]_q$, is defined as

$$[n]_q := 1 + q + \cdots + q^{n-1} = \frac{1 - q^n}{1 - q}.$$

The *q-factorial* of an integer $n \geq 0$, denoted $[n]_q!$, is defined as

$$[n]_q! := [n]_q \cdot [n-1]_q \cdots [2]_q \cdot [1]_q.$$

The *q-binomial coefficient* for $0 \leq n \leq m$, denoted $\begin{bmatrix} m \\ n \end{bmatrix}_q$, is defined as

$$\begin{bmatrix} m \\ n \end{bmatrix}_q := \frac{[m]_q!}{[n]_q![m-n]_q!}.$$

At $q = 1$, the q-bracket $[n]_q$ evaluates to n. Thus, at $q = 1$, the q-factorial $[n]_q!$ evaluates to $n!$ and the q-binomial coefficient $\begin{bmatrix} m \\ n \end{bmatrix}_q$ evaluates to the binomial coefficient $\binom{m}{n}$. This shows why these functions are named as such. The q-bracket is a polynomial in q and the q-factorial is a product of q-brackets, so it must be a polynomial as well.

Lemma 1.2. *For all $0 \leq n \leq m$, we have the identity*

$$\begin{bmatrix} m \\ n \end{bmatrix}_q = \begin{bmatrix} m-1 \\ n \end{bmatrix}_q + q^{m-n} \begin{bmatrix} m-1 \\ n-1 \end{bmatrix}_q.$$

Proof. Starting from the right side

$$\begin{aligned}
\begin{bmatrix} m-1 \\ n \end{bmatrix}_q + q^{m-n} \begin{bmatrix} m-1 \\ n-1 \end{bmatrix}_q &= \frac{[m-1]_q!}{[n]_q![m-n-1]_q!} + \frac{q^{m-n}[m-1]_q!}{[n-1]_q![m-n]_q!} \\
&= \frac{[m-1]_q!}{[n-1]_q![m-n-1]_q!} \left(\frac{1}{[n]_q} + \frac{q^{m-n}}{[m-n]_q} \right) \\
&= \frac{[m-1]_q!}{[n-1]_q![m-n-1]_q!} \left(\frac{q^{m-n}[n]_q + [m-n]_q}{[m-n]_q[n]_q} \right).
\end{aligned}$$

If we look at just the second numerator,

$$\begin{aligned}
q^{m-n}[n]_q + [m-n]_q &= q^{m-n}(q^{n-1} + \cdots + q + 1) + (q^{m-n-1} + \cdots + q + 1) \\
&= q^{m-1} + \cdots + q^{m-n} + q^{m-n-1} + \cdots + q + 1 = [m]_q.
\end{aligned}$$

Thus, returning to our first string of equalities we get

$$\begin{aligned}
\begin{bmatrix} m-1 \\ n \end{bmatrix}_q + q^{m-n} \begin{bmatrix} m-1 \\ n-1 \end{bmatrix}_q &= \frac{[m-1]_q!}{[n-1]_q![m-n-1]_q!} \left(\frac{[m]_q}{[m-n]_q[n]_q} \right) \\
&= \frac{[m]_q!}{[n]_q![m-n]_q!} = \begin{bmatrix} m \\ n \end{bmatrix}_q.
\end{aligned}$$

□

Proposition 1.3. *The q-binomial coefficient is always a polynomial with nonnegative integer coefficients.*

Proof. Using the definition of the q-binomial coefficient, $\begin{bmatrix} m \\ 0 \end{bmatrix}_q = \begin{bmatrix} m \\ m \end{bmatrix}_q = 1$ for all m. For an arbitrary coefficient $\begin{bmatrix} m \\ n \end{bmatrix}_q$, we can repeatedly apply Lemma 1.2 until we hit the case of $n = 0$ or $m = n$ for all terms which gives $\begin{bmatrix} m \\ n \end{bmatrix}_q$ as a sum of powers of q. Thus, we see the q-binomial coefficient is always a polynomial with nonnegative integer coefficients. □

Using the definition of the q-binomial coefficient, we can see that any zero of $\begin{bmatrix} m \\ n \end{bmatrix}_q$ must be a root of unity and any primitive root of unity, denoted ω_m, is a zero if $0 < n < m$. The mth roots of unity which are not primitive are not guaranteed to be zeros, so we may wonder what the polynomial evaluates to at these values. Summarized in Table 2 below are the results of evaluating $\begin{bmatrix} 12 \\ n \end{bmatrix}_q$ for different values of n and q. From Proposition 1.3, these are rational polynomials so we need only evaluate at a single primitive root of unity for each divisor of 12.

Unexpectedly, the result of evaluating at each root of unity gives an integer. In fact, these are the same numbers that appeared in the columns of Table 1. We might now recall that the mth roots of unity under multiplication are isomorphic to C_m. If we let the copy of $C_{12} \subset \mathbb{C}$ act on the set of size n subsets of corners of a dodecagon, the number of fixed configurations is equal to $\begin{bmatrix} 12 \\ n \end{bmatrix}_q$ evaluated at the complex number representing each element of C_{12}. The same fact is true for $m = 6$, though not computed here.

Strangely, it seems our combinatorial formula counts something meaningful when we input values beyond the domain of what it intends to represent. While we have not proven that this pattern will hold for all cases, we have managed to stumble upon an instance of the *cyclic sieving phenomenon*. This case is not special, but rather the tip of the iceberg for simpler polynomials counting something when evaluated at roots of unity.

	$q = \omega_{12}^2$	$q = \omega_{12}^3$	$q = \omega_{12}^4$	$q = \omega_{12}^6$	$q = \omega_{12}^{12}$
$n = 0, 12$	1	1	1	1	1
$n = 1, 11$	0	0	0	0	12
$n = 2, 10$	0	0	0	6	66
$n = 3, 9$	0	0	4	0	220
$n = 4, 8$	0	3	0	15	1495
$n = 5, 7$	0	0	0	0	792
$n = 6$	2	0	6	20	924

TABLE 2. Evaluating $\begin{bmatrix} 12 \\ n \end{bmatrix}_q$ at 12th roots of unity for $m = 12$.

2. Cyclic Sieving

Now that we have identified one example, we can define the cyclic sieving phenomenon in full. It occurs when we have a polynomial evaluated at roots of unity giving meaningful information about the action of a cyclic group on some set.

Definition 2.1. Let C_n be a finite cyclic group of order n generated by $c \in C_n$ which acts on a finite set X and let $X(q)$ be a polynomial with integer coefficients. The triple $(X, X(q), C_n)$ exhibits the *cyclic sieving phenomenon* or *CSP* if $X(\omega_n^d)$ is the number of elements of X which are fixed by the action c^d for all $0 \leq d < n$.

It is not clear from this formulation how common a phenomenon like this might be. In fact, we will show at the end of Section 3 that for any C_n-action on some set X, there is some polynomial $X(q)$ which completes the cyclic sieving triple $(X, X(q), C_n)$. However, these polynomials are not guaranteed to be interesting on their own. The phenomenon becomes interesting when the polynomials in question are well known outside of their ability to exhibit CSP. Proofs of the three examples below are given in [9].

Example 2.2. From the previous section, let X be size n subsets of corners of an m-gon, let $C = C_m$ act on this set by rotation, and let $X(q) = \begin{bmatrix} m \\ n \end{bmatrix}_q$. The triple $(X, X(q), C)$ exhibits the cyclic sieving phenomenon.

An equivalent formulation of this example is that X is the size n subsets of $\{0, 1, ..., m-1\}$ and the group action of c^d is adding d to each element and taking the remainder mod m.

Example 2.3. Let X be size n multisets on $\{0, 1, ..., m-1\}$, let $C = C_m$ acting on this set by addition mod m, and let $X(q) = \begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q$. The triple $(X, X(q), C)$ exhibits the cyclic sieving phenomenon. We note that this formula is reminiscent of the formula for calculating the number of size n multisets of $\{0, 1, ..., m-1\}$: $\binom{m+n-1}{n}$.

Example 2.4. Let X be the triangulations of an n-gon with $C = C_n$ acting on this set by rotation. Let $X(q) = \frac{1}{[n-1]_q} \begin{bmatrix} 2n-4 \\ n-2 \end{bmatrix}_q$. The triple $(X, X(q), C)$ exhibits the cyclic sieving phenomenon. Again, this formula bears similarity to the formula for the Catalan numbers which count the number of triangulations of an n-gon: $\frac{1}{n-1}\binom{2n-4}{n-2}$.

Now that we have seen a few examples, we will prove the phenomenon for one of these cases: Example 2.3. While this proof is only for this specific case, it serves as a template for future understanding in Section 5. We will take advantage of symmetric powers of a vector space. Let $c \in C_m$ act on $\mathrm{Sym}^n(\mathbb{C}^m)$ by permuting some basis set $x_1, ..., x_m$ of $\mathbb{C}^m$ cyclically and acting on a symmetric product by $c \cdot x_{i_1} x_{i_2} \cdots x_{i_n} = x_{c \cdot i_1} x_{c \cdot i_2} \cdots x_{c \cdot i_n}$. For example, if $m = 5$ and $n = 3$, the basis vector $x_1 x_3 x_5$ is sent to $x_2 x_4 x_1$ by c. This action is linear.

Proposition 2.5. *Using the construction above, the trace of the linear map induced from $c \in C_m$ is the number of basis vectors of $x_{i_1} \cdots x_{i_n} \in \mathrm{Sym}^n(\mathbb{C}^m)$ which are fixed by this action.*

Proof. The c action sends a basis vector $x_{i_1} \cdots x_{i_n}$ to $x_{i_1+1} \cdots x_{i_n+1}$ where $x_{m+1} = x_1$. After rearranging, we see that this is another basis vector. Thus, the c action is a permutation matrix. Along its diagonal, this matrix has a 1 if a basis vector is sent to itself and a zero otherwise. Thus, we see that the trace of the map induced by this action is number of basis vectors fixed by the c action. □

Proposition 2.6. *Using the construction above, let $c \in C_m$ be a generator. The action induced by c^d is also given by $h_n(1, \omega_m^d, \omega_m^{2d}, ..., \omega_m^{(m-1)d})$ where*

$$h_n(x_1, ..., x_m) = \sum_{1 \le i_1 \le i_2 \le \cdots \le i_n \le m} x_{i_1} x_{i_2} \cdots x_{i_n}$$

Proof. The characteristic polynomial of c acting on $\mathbb{C}^m$ is $x^m - 1$, so there must be an eigenbasis $y_0, ..., y_{m-1}$ with eigenvalues $1, \omega_m, ..., \omega_m^{m-1}$. Because c^d is just repeated application of the c action, the same eigenbasis is still an eigenbasis of c^d with eigenvalues $1, \omega_m^d, ..., \omega_m^{(m-1)d}$. We can form a basis of $\mathrm{Sym}^n(\mathbb{C}^m)$ using the y_i to form products. The c^d action on the basis vector $y_{i_1} \cdots y_{i_n}$ sends it to $\omega_m^{i_1 d} y_{i_1} \cdots \omega_m^{i_n d} y_{i_n} = (\omega_m^{i_1 d} \cdots \omega_m^{i_n d}) y_{i_1} \cdots y_{i_n}$, so $\omega_m^{i_1 d} \cdots \omega_m^{i_n d}$ is an eigenvalue for any combination of $1 \le i_1 \le i_2 \le \cdots \le i_n \le m$. The trace of the c^d action is the sum of all the eigenvalues, which we can see is $h_n(1, \omega_m^d, \omega_m^{2d}, ..., \omega_m^{(m-1)d})$. □

Lemma 2.7. *Using the polynomials described in Proposition 2.6, we have an identity*

$$h_n(1, q, ..., q^{m-1}) = \begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q .$$

Proof. We prove this by induction on m and n with our two base cases being $m = 1$ and $n = 0$. First, if $m = 1$, $h_n(1) = 1^n = 1$ and $\begin{bmatrix} n \\ n \end{bmatrix}_q = 1$. If instead $n = 0$, $h_0(1, q, ..., q^{m-1}) = 1$ because it is the single empty product and $\begin{bmatrix} m \\ 0 \end{bmatrix}_q = 1$. Now suppose the identity holds for all combinations of $n \le N$ and $m < M$. The polynomial $h_n(1, q, ..., q^{m-1})$ can be split into two groups, those that do not have a q^{m-1} in their product and those that do. We can write the former as $h_n(1, q, ..., q^{m-2}) = \begin{bmatrix} m+n-2 \\ n \end{bmatrix}_q$ and the latter as $q^{m-1} h_{n-1}(1, q, ..., q^{m-1}) = q^{m-1} \begin{bmatrix} m+n-2 \\ n-1 \end{bmatrix}_q$. This gives

$$\begin{aligned} h_n(1, q, ..., q^{m-1}) &= \begin{bmatrix} m+n-2 \\ n \end{bmatrix}_q + q^{m-1} \begin{bmatrix} m+n-2 \\ n-1 \end{bmatrix}_q \\ &= \begin{bmatrix} (m+n-1)-1 \\ n \end{bmatrix}_q + q^{(m+n-1)-n} \begin{bmatrix} (m+n-1)-1 \\ n-1 \end{bmatrix}_q . \end{aligned}$$

Applying Lemma 1.2, we get

$$h_n(1, q, ..., q^{m-1}) = \begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q .$$

□

Proposition 2.8. *Let X be size n multisets on $\{0, 1, ..., m-1\}$, $C = C_m$ acting on this set by addition mod m, and $X(q) = \begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q$. The triple $(X, X(q), C)$ exhibits the cyclic sieving phenomenon.*

Proof. Multisets index the basis set of $\mathrm{Sym}^n(\mathbb{C}^m)$, so by the construction presented the C_m action on $\mathrm{Sym}^n(\mathbb{C}^m)$ is the same as the C_m action on X. The trace is an invariant, so Propositions 2.5 and 2.6 show that the number of elements in X fixed by $c^d \in C_m$ is given by $h_n(1, \omega_m^d, \omega_m^{2d}, ..., \omega_m^{(m-1)d})$. By Lemma 2.7, this is equal to $\begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q$ evaluated at $q = \omega_m^d$. □

As a final note, the cyclic sieving phenomenon serves to explain a different phenomenon that was identified: the $q = -1$ phenomenon [10]. Stembridge noticed that often evaluating combinatorial formulae at $q = -1$ gives the number of fixed points of many involutions. An involution is a map that is its own inverse, so the process of applying an involution is a C_2 action and -1 is the second root of unity. Thus, we see that the cyclic sieving phenomenon is just a generalization of the $q = -1$ phenomenon.

3. Representation Theory and the Representation Ring

In the proof of Proposition 2.8, we made use of realizing C_m as linear transformations of the vector space $\mathrm{Sym}^n(\mathbb{C}^n)$. Mapping groups into linear transformations is the premise of representation theory. By utilizing some tools from this theory we can better understand what is occurring in the cyclic sieving phenomena.

Definition 3.1. Let G be a group and let V be a vector space over $\mathbb{C}$. A *group representation of* G, shortened to a *G-rep*, is a group homomorphism ρ mapping G to GL(V). When there is no confusion about ρ, the representation is also sometimes referred to by just the vector space V. For the rest of this paper, we will assume V is a finite-dimensional vector space and that G is a finite group.

In our proof of Proposition 2.8, we have made use of two C_m-reps. The first is the representation $\mathbb{C}^m$ where we permuted the basis vectors cyclically and the second is the more complicated representation $\mathrm{Sym}^n(\mathbb{C}^m)$.

Definition 3.2.

(1) Given a representation $\rho : G \to \mathrm{GL(V)}$, a subspace $W \subset V$ is *invariant under* G if for all $w \in W$ and $g \in G$, $\rho(g)w \in W$.
(2) A representation $\rho : G \to \mathrm{GL(V)}$ is *irreducible* if the only subspaces of V invariant under G are V and $\{0\}$.
(3) Two G-reps V and W with maps ρ_V and ρ_W are *isomorphic* if there is a vector space isomorphism $\phi : V \to W$ which also satisfies $\phi(\rho_V(g)v) = \rho_W(g)\phi(v)$ for all $g \in G$ and $v \in V$

Using our above examples, the representation $\mathbb{C}^m$ is not irreducible because the subspace generated by $x_1 + x_2 + \cdots + x_m$ is fixed by all members of C_m. The representation $\rho : C_m \to \mathbb{C}$ given by mapping a generator of C_m to ω_m is irreducible because there are no subspaces besides $\mathbb{C}$ and $\{0\}$.

We will make use of many facts about representations in this paper which will be stated in Propositions 3.3, 3.5, and 3.6. Proofs of these propositions are not particularly illuminating for the purposes of this paper and thus these propositions will be stated without proof. A treatment of group representations which proves all of these propositions can be found in the first two chapters of [3].

Proposition 3.3.

1. *Up to isomorphism, there are only finitely many irreducible representations of a group G.*
2. *Every non-irreducible representation of G is isomorphic to a direct sum of irreducible representations.*
3. *The choice of representations in the direct sum is unique up to rearranging.*

This proposition gives some clarity on why we choose the word *irreducible*: all other representations can be broken down into a direct sum of ineducable representations which can't be broken down any further. Additionally, each irreducible representation in the decomposition of V can be identified with a G-invariant subspace of V for which the restricted linear action on the subspace is isomorphic to the irreducible representation.

Definition 3.4. A *class function* on a group G is a function $f : G \to \mathbb{C}$ which is constant on each conjugacy class of G. If G has n conjugacy classes, then the class functions form an n-dimensional vector space.

The *character* of a representation ρ, $\chi_\rho : G \to \mathbb{C}$, is the function $\chi(g) := \mathrm{Tr}(\rho(\mathrm{g}))$. Traces of similar matrices are equal, so the character is a class function.

Proposition 3.5.

1. *The characters of the irreducible representations form a basis for the vector space of class functions on G.*
2. *Two G-representations with the same character are isomorphic.*
3. *Let $V_1, ..., V_n$ be all irreducible representations of some group G. The following equation holds:*

$$|G| = \sum_{i=1}^{n} \deg(V_i)^2.$$

Proposition 3.6.

1. *The character of the direct sum of two representations is the sum of their characters.*
2. *The tensor product of two G-reps V and W with maps ρ_V and ρ_W, written $V \otimes W$, is a G-rep on the tensor product of the vector spaces V and W with $\rho_{V\otimes W}(g)(v \otimes w) = (\rho_V(g)v) \otimes (\rho_W(g)w)$. The character of this representation is the product of the characters of V and W.*
3. *For G-reps U, V, and W, the identity $U \otimes (V \oplus W) \simeq (U \otimes V) \oplus (U \otimes W)$ holds.*
4. *For any group G, the trivial representation V_{triv} is the one-dimensional G-representation where all members of G map to 1. For any G-rep W, $V_{\mathrm{triv}} \otimes W \simeq W$.*

Both direct sums and tensor products are commutative and associative operations on isomorphism classes of representations of some G. From Proposition 3.6, tensor products distribute over direct sums. This fact suggests a possible ring structure. Direct sums do not have additive inverses, so only a semiring can be constructed with the trivial representation as 1. The conditions of a semiring are shown by Proposition 3.6 as well as the commutativity and associativity of the operations.

This semiring can be extended to a ring by letting $\mathbb{N}$ act on members of the semiring by repeated addition, forming an $\mathbb{N}$-semimodule, and taking a tensor product as $\mathbb{N}$-semimodules with the integers. Informally, this process adds an object as the additive inverse for each element in our semiring to make a ring. We call this ring the *representation ring* of G. Because the character commutes with the addition and multiplication operations in the representation ring, we can think of elements in the representation ring as polynomials in the characters of G.

C_m has m conjugacy classes because the group is abelian, so there are m irreducible representations forming a basis due to Proposition 3.5. Let $c \in C_m$ be a generator. We construct representations over the one-dimensional vector space $\mathbb{C}$ by sending c to all m powers of ω_m. The character at c is that power of ω_m, so we have identified all the irreducible representations of C_m. Certainly this set generates the representation ring, but we can do even better. Using the representation that sends c to ω_m, we take the tensor product of this representation with itself to get the other representations. Thus, we see that the representation ring of C_m is generated by a single representation, denoted V_m, that embeds C_m into $\mathbb{C}$.

Definition 3.7. Let G be a group acting on a set X. The permutation G-representation of X is the $|X|$-dimensional vector space with a basis $\{v_x | x \in X\}$ and where each $g \in G$ acts on v_x by sending it to v_{gx}. The representation will be denoted V_X when the group is clear.

The matrix associated to each member $g \in G$ in a permutation representation is a permutation matrix with a single 1 in each row and column. The character of a permutation representation evaluated at each g is the number of elements of X which are fixed by g. The vector that is the sum of all basis vectors is fixed by any action. Thus, if $|X| > 1$, the permutation representation is not irreducible.

Some examples of permutation representations are summarized in Table 3 below for C_6 and generator c. The first representation V_6 is our generating representation of C_6 described above. The next two are the permutation C_6-representations where X is the set of segments in a hexagon from Figure 1 and Y is the set of triangles in a hexagon from Figure 2 with C_6 acting by rotation. As we expect, the character evaluated at an element is the number of elements which are fixed by its action.

The values in this table are also given by the cyclic sieving phenomenon: we can evaluate the sieving polynomials $X(q) = \begin{bmatrix} 6 \\ 2 \end{bmatrix}_q$ and $Y(q) = \begin{bmatrix} 6 \\ 3 \end{bmatrix}_q$ at appropriate sixth roots of unity. Looking at the table, the root of unity we desire is exactly the

Representation V	$\chi_V(1)$	$\chi_V(c)$	$\chi_V(c^2)$	$\chi_V(c^3)$	$\chi_V(c^4)$	$\chi_V(c^5)$
V_6	1	ω_6	ω_6^2	ω_6^3	ω_6^4	ω_6^5
V_X	15	0	0	3	0	0
V_Y	20	0	2	0	2	0

TABLE 3. Characters for three reps of C_6.

character of V_6. Thus, we know that for any $g \in C_6$, $X(\chi_{V_6}(g))$ and $Y(\chi_{V_6}(g))$ give the number of fixed hexagons in the associated group action.

We have seen in Proposition 3.6 that characters distribute over addition and multiplication in the representation ring, so the above equivalences become $\chi_{X(V_6)} = \chi_{V_X}$ and $\chi_{Y(V_6)} = \chi_{V_Y}$. The character uniquely determines a representation, so we must have $X(V_6) = V_X$ and $Y(V_6) = V_Y$ in the representation ring. Our cyclic sieving polynomials not only make sense evaluated at roots of unity, but show a relationship between a generating representation of C_6 and the permutation C_6-representations of two sets.

This process is not unique to C_6. The same process done for any cyclic sieving triple $(X, X(q), C_m)$ yields $X(V_m) = V_X$, an isomorphism in the representation ring. This shows that a cyclic sieving polynomial must exist for any cyclic group action because the polynomial is just a decomposition of the permutation representation into irreducible components. We can also use this isomorphism to motivate what sieving looks like over other types of groups.

4. Sieving Over Any Group

From our construction in the previous paragraph, we have seen that cyclic sieving is equivalent to an isomorphism of representations in the representation ring of a cyclic group. Using this construction, we see how we might generalize the sieving phenomenon to an arbitrary finite group.

Definition 4.1. Let G be a finite group which acts on some finite set X, let $V_1, ..., V_n$ be G-reps which generate the representation ring of G, and let $X(q_1, ..., q_n) \in \mathbb{Z}[q_1, ..., q_n]$ be a polynomial. The quadruple $(X, X(q_1, ..., q_n), \{V_1, ..., V_n\}, G)$ exhibits G-sieving if $X(V_1, ..., V_n)$ is equal to the permutation G-representation of X in the representation ring of G.

Equivalently, the quadruple $(X, X(q_1, ..., q_n), \{V_1, ..., V_n\}, G)$ exhibits G-sieving if $X(\chi_{V_1}, ..., \chi_{V_n}) : G \to \mathbb{C}$ evaluated at any element $g \in G$ is equal to the number of elements of X which are fixed by the action by g.

We may now wonder what these polynomials look like for groups beyond that of a cyclic group. In several of our examples for CSP, we use the cyclic action on polygons. This suggests that a natural next option might be the dihedral group which also acts on polygons by rotations and reflections. As we will see, the dihedral groups have an easy to understand representation ring.

Definition 4.2. A dihedral group D_n is the group generated by two elements r and s with the additional relations $r^n = s^2 = 1$ and $rs = sr^{-1}$. This group has $2n$ elements and can be realized as the rotation and reflection symmetries of an n-gon.

Proposition 4.3. *Let n be odd. The irreducible representations of D_n are generated by*

$$V_{triv} : r \mapsto 1 \text{ and } s \mapsto 1,$$

$$V_{sign} : r \mapsto 1 \text{ and } s \mapsto -1,$$

$$\text{and } V_i : r \mapsto \begin{bmatrix} \omega_n^i & 0 \\ 0 & \omega_n^{-i} \end{bmatrix} \text{ and } s \mapsto \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \text{ for } 1 \le i \le \frac{n-1}{2}.$$

Proof. The character of each representation is distinct, so each representation is unique. The eigenvectors of the two matrices described for each V_1 are distinct, so we must have that all representations are irreducible. Using the sum in Proposition 3.5, we see this must be all representations. □

Using the notation above, we notice that $V_{sign} \otimes V_{sign} = V_{triv}$ and $V_i = V_1^{\otimes i}$, so we generate the representation ring using just V_{sign} and V_1. Conveniently, this means that we construct polynomials in two variables which exhibit the dihedral sieving phenomenon. For the case of even n, we run into two new issues. First, there are two additional one-dimensional representations which make the representation ring harder to generate. Second, the two ways we might choose to have s act on a polygon – a reflection that fixes opposite corners or through opposite edges fixing no corners – give different results for numbers of fixed points so we have two different polynomials depending on this choice. With odd n, the natural action on the polygon results in the same polynomials no matter which line of reflection we choose.

Definition 4.4. The *generalized Fibonacci polynomials* are a sequence of polynomials in s and t defined as

$$\{0\}_{s,t} = 0,$$

$$\{1\}_{s,t} = 1,$$

$$\{n+2\}_{s,t} = s\{n+1\}_{s,t} + t\{n\}_{s,t}.$$

Like with the q-brackets, we define a factorial operation as

$$\{n\}_{s,t}! = \{n-1\}_{s,t} \cdots \{2\}_{s,t}\{1\}_{s,t}.$$

Finally, the *Fibonacci coefficients* are a set of polynomials defined as

$$\left\{ \begin{matrix} m \\ n \end{matrix} \right\}_{s,t} = \frac{\{m\}_{s,t}!}{\{n\}_{s,t}!\{m-n\}_{s,t}!}.$$

Using these polynomials, we list a few examples of dihedral sieving which are proven in [6].

Example 4.5. Let X be the set of size n subsets of corners of an m-gon with m odd. Let $G = D_m$ act on this set naturally. Let $X(s,t) = \left\{ \begin{matrix} m \\ n \end{matrix} \right\}_{s,t}$. The quadruple $(X, X(s,t), \{V_{sign}, V_1\}, G)$ exhibits the dihedral sieving phenomenon.

Example 4.6. Let X be size n multisets on $\{1, ..., m\}$ with m odd, $G = G_m$ acting on this set by natural dihedral permutation, and $X(s,t) = \left\{ \begin{matrix} m+n-1 \\ n \end{matrix} \right\}_{s,t}$. The quadruple $(X, X(s,t), \{V_{sign}, V_1\}, G)$ exhibits the dihedral sieving phenomenon.

We might notice that these two formulae are extremely similar to the cyclic sieving phenomena on the same set X with the q-binomial coefficients swapped for Fibonacci coefficients. This similarity suggests that there is some underlying connection between these groups and their respective polynomial structures. As we will soon find out, this is a side effect of a *λ-ring* structure on both representation rings.

5. λ-Rings

Before we define what a λ-ring is, we must define a few polynomials that are used in the construction.

Definition 5.1. The jth elementary symmetric polynomial in the k variables $x_1, ..., x_k$, denoted by $e_{j,k}$, is given by

$$e_{j,k} = \sum_{1 \leq i_1 < i_2 < \cdots < i_j \leq k} x_{i_1} x_{i_2} \cdots x_{i_j}.$$

We now define for all natural n and m the polynomial $P_{n,m}$ to be the unique polynomial in nm variables such that $P_{n,m}(e_{1,nm}, ..., e_{nm,nm})$ is equal to the coefficient on the t^n term of the expression

$$\prod_{1 \leq i_1 < i_2 < \cdots < i_m \leq nm} (1 + x_{i_1} x_{i_2} \cdots x_{i_k} t).$$

Let $e_{j,k}$ be as constructed above and $f_{j,k}$ be the same polynomials but over different variables $y_1, ..., y_k$. We now define for all natural n the polynomial P_n to be the unique polynomial in $2n$ variables such that $P_n(e_{1,n}, ..., e_{n,n}, f_{1,n}, ..., f_{n,n})$ is equal to the coefficient on the t^n term of the expression

$$\prod_{i,j=1}^{n} (1 + x_i y_j t).$$

The existence and uniqueness of these polynomials is guaranteed by the fundamental theorem of symmetric polynomials. It states that every symmetric polynomial can be uniquely written as a polynomial in the elementary polynomials and is Corollary 31 of Section 14.6 of [1]. Using these polynomials, we define the λ-ring structure.

Definition 5.2. A *λ-ring* is a commutative ring X with 1 equipped with a sequence of unary functions $\lambda^n : X \to X$ which satisfy all of the following for all $x, y \in X$ and $n, m \in \mathbb{N}$:

- $\lambda^0(x) = 1$
- $\lambda^1(x) = x$
- $\lambda^n(1) = 0$ if $n \geq 2$
- $\lambda^n(x + y) = \sum_{i=0}^{n} \lambda^i(x) \lambda^{x-i}(x)$
- $\lambda^n(xy) = P_n(\lambda^1(x), \lambda^2(x), ..., \lambda^n(x), \lambda^1(y), ..., \lambda^n(y))$
- $\lambda^n(\lambda^m(x)) = P_{n,m}(\lambda^1(x), ..., \lambda^{mn}(x))$

Example 5.3. Let $X = \mathbb{Z}$ be the ring integers and let

$$\lambda^n(x) = \binom{x}{n} = \frac{x(x-1)\cdots(x-n+1)}{n!}$$

be the binomial coefficient defined on all integers x. Then X is a λ-ring.

Example 5.4. Let X be the representation ring for any finite group G and let $\lambda^n(V)$ be the nth exterior power of the representation V with g acting on each component in the product. Then X is a λ-ring [11].

Example 5.5. Let X be the ring $\mathbb{Z}[q]$ and define $\lambda^n : \mathbb{Z}[q] \to \mathbb{Z}[q]$ by its action on additive generators $\lambda^n(1 + q + q^2 + \cdots + q^m) = \begin{bmatrix} m+n-1 \\ n \end{bmatrix}_q$. Equipped with these λ^n, X is a λ-ring.

We might notice that the second and third examples are two parts of our story of cyclic sieving. It turns out that taking the character is a λ-ring homomorphism: the character commutes with the λ-operations in each λ-ring. Because we are able to describe the permutation ring of subsets using λ-operations and can easily map from a single generator of one λ-ring to the other, it is easy to construct the polynomials by using the known λ-operations in the polynomial ring.

For the other examples, we are often able to describe other operations, such as exterior powers, using these λ-operations. This is why we end up seeing the q-binomial coefficients frequently in formulae for cyclic sieving which make use of lots of symmetry: the underlying λ-operation of q-binomial coefficients is the same.

The λ-ring structure is also why we see similarities in several formulae in both cyclic and dihedral sieving. While the polynomial ring is more complex, the Fibonacci polynomials are serving as a λ-operation on the associated ring. Going forward, we will use this strategy of identifying the λ-ring structures in each ring and mapping between the simplest cases in order to build up our sieving polynomials.

6. Representations of the Symmetric Group

Now that we have seen how sieving works with two simple cases of cyclic and dihedral groups, another type of group to consider is the symmetric group. To do this, we need to first investigate Young tableaux and their connection to representations of the symmetric group.

Definition 6.1. Let $\pi = (n_1, ..., n_k)$ be a partition of n (where $n_1 \geq n_2 \geq \cdots \geq n_k$ and $n_1 + \cdots + n_k = n$). A *Young diagram* is an arrangement of n squares into k rows where the ith row has n_i squares and the leftmost square in each row are aligned vertically. See Figure 4 for a visual example of these Young diagrams. A *Young tableau* is a Young diagram with the numbers $1, ..., n$ each placed exactly one to a box. A *Young tabloid* for a given partition is an equivalence class of all Young tableaux which have the same numbers in each row. A Young tabloid is drawn without vertical bars. An example of both Young tableaux and Young tabloids are shown in Figure 5.

For a given tableau t, we use $\{t\}$ to denote the coresponding tabloid. The symmetric group S_n acts on the set $1, ..., n$ by permutation. This action extends to the set of Young tableaux or tabloids for a given partition by permuting the

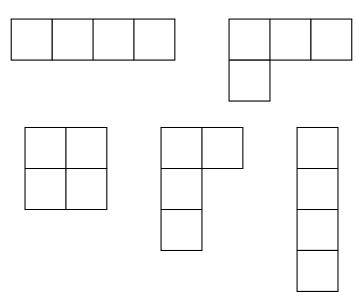

Figure 4. The 5 Young diagrams corresponding to partitions of 4. Top: (4) and $(3, 1)$. Bottom: $(2, 2)$, $(2, 1, 1)$, and $(1, 1, 1, 1)$.

3	5	1
4	2	
6		

1	5	3
2	4	
6		

1	3	5
2	4	
6		

FIGURE 5. Two Young tableau (left) of the partition (3,2,1) which are in the same equivalence class of the Young tabloid (right)

numbers in each box. With this action, a member of S_n which only permutes numbers within the row they start in fixes the given tabloid. If $\{t\}$ is a tabloid and $\sigma \in S_n$ is a group element, we use $\{\sigma t\}$ to denote this action. We also use $x_{\{t\}}$ to refer to the basis vector corresponding to $\{t\}$ in the permutation representation of tabloids.

Definition 6.2. The *column group* of a tableau t, denoted C_t, is the subgroup of S_n which fixes which column each number is in. As tabloids are tableaux which differ by row permutations, each member of the column group of t sends $\{t\}$ to a unique tabloid.

For the first tableau in Figure 5, the column group is $\langle (3,4),(3,6),(2,5)\rangle \subset S_6$.

Definition 6.3. For a given tableau t, its *polytabloid* e_t is the vector in the permutation representation of tabloids given by

$$\sum_{\sigma \in C_n} \operatorname{sign}(\sigma)\{\sigma t\}.$$

For each partition π of n, the subspace of the permutation representation of tabloids corresponding to that partition spanned by the polytabloids of all Young tableaux with shape π is an S_n representation because S_n acts on this vector space. However, being just a representation is not very interesting.

Proposition 6.4. *For each partition π of n, the subspace of the permutation representation of tabloids corresponding to that partition spanned by the polytabloids of all Young tableaux with shape π is a unique irreducible S_n representation.*

A proof of Proposition 6.4 is given in Chapter 4 of [3].

Corollary 6.5. *Each partition of n corresponds to a conjugacy class in S_n, so we have a unique irreducible representation for each conjugacy class of S_n. By Proposition 3.5, these must be all irreducible representations of S_n.*

With this catalog of irreducible representations of S_n, we better understand the representation ring and can investigate the possibility of S_n-sieving.

7. Understanding (Symmetric) Sieving with λ-Rings

From Section 5 of this paper, we saw that understanding the λ-ring structure of the representation ring gives insight into constructing sieving polynomials. In this section, we work through a similar process for S_n and see where the process breaks down.

First, we identify three important representations of the symmetric group. The first is the permutation representation that comes from S_n acting on n letters. As we have seen, identifying how the permutation representation breaks down into irreducible representations lets us use λ-operations in a λ-ring isomorphism when we construct our polynomials. The second representation is the representation coresponding to the permutation (n) for which there is only one tabloid and all permutations act trivially on it, so this representation, called the trivial representation, sends all group members to 1. We have seen that there is always a subspace of the permutation representation for which S_n acts trivially, so it appears at least once in the the decomposition into irreducible representations. If we remove one copy of the trivial representation from the permutation representation's decomposition, we are left with some $(n-1)$-dimensional representation that is the direct sum of all other irreducible representations which we call the standard representation. A fact not proven here is that the standard representation is irreducible and corresponds to the partition $(n-1,1)$ [3]. Thus, we have $V_{S_n \text{perm}} = V_{(n)} \oplus V_{(n-1,1)}$.

Now, like with the proof of cyclic sieving, we use this to more easily describe our permutation representations of more interesting sets. We know that $\bigwedge^n(V_{S_m \text{perm}}) = \bigwedge^n(V_{(n)} \oplus V_{(n-1,1)})$ should be isomorphic to the permutation representation of size n subsets of $\{1,...,m\}$ and $\text{Sym}^n(V_{S_m \text{perm}}) = \text{Sym}^n(V_{(n)} \oplus V_{(n-1,1)})$ should be isomorphic to the permutation representation of size n multisets of $\{1,...,m\}$. If we could expand these powers into a polynomial in the representations, we could generate our sieving polynomial.

In fact, this is almost possible. While we can't expand directly into representations, we can write a formula for the character of these representations. For a G-rep V, the simplest case of $n = m$ gives

$$\chi_{\bigwedge^2 V}(g) = \frac{\chi_v(g)^2 - \chi_V(g^2)}{2},$$

$$\chi_{\text{Sym}^2 V}(g) = \frac{\chi_v(g)^2 + \chi_V(g^2)}{2}.$$

The higher powers are written as a determinant of a matrix with simple entries as shown in [4]. The first problem we face is the terms of the form $\chi_V(g^k)$ which show up in all of these formulae. Because S_n has terms of all orders less than or equal to n, it is challenging to determine how this is a sum of irreducible representations. Even just tensor powers in S_n become quite complicated due to the varied nature of the irreducible representations. Finally, after deducing which irreducible representations the formula decomposes into, there is not a uniform way to generate the representation ring, so the number of variables in a sieving polynomial must grow with n.

With these challenges, finding sieving polynomials of S_n is more of an exercise in simplifying a simple yet tedious expression rather than utilizing already existing formulae like the case of cyclic and dihedral sieving. The first two relations for multisets of $\{1,2,3,4\}$ are listed:

$$\text{Sym}^2 V_{S_4 \text{ perm}} = V_{(4)} \oplus V_{(3,1)} \oplus V_{(2,2)},$$

$$\text{Sym}^3 V_{S_4 \text{ perm}} = 3V_{(4)} \oplus 4V_{(3,1)} \oplus V_{(2,1,1)} \oplus V_{(2,2)}.$$

These could be turned into polynomials in 5 variables for each of the irreducible representations of S_4 and could be further reduced into a polynomial in just $V_{(1,1,1,1)}$ and $V_{3,1}$, but the relations that make these formulae neater do not illuminate much about the structure and are more an indication that S_4 is small enough to be generated with 2 generators. This also shows why the existence of a polynomial is uninteresting; any permutation representation can be written as a sum of irreducible representations. In the cases of cyclic sieving – where all representations are powers of one generator – or dihedral sieving – where the small degrees of irreducible representations limit what results might be – we make use of well known polynomials in our λ-ring homomorphism to give nice answers to these sieving questions but this method does not work in general.

Acknowledgments

I would like to thank mentors Trevor Hyde and Peng Hui How for providing resources and explanations to help me understand these topics. I would also like to thank them for proofreading and providing suggestions to help this paper reach its full potential. I would like to thank Peter May along with everyone else who helped The University of Chicago Mathematics REU 2021 run as smoothly as possible.

References

[1] David Steven Dummit and Richard M. Foote. *Abstract algebra*. John Wiley & Sons, 2004.

[2] Jesse Elliott. "Binomial rings, integer-valued polynomials, and λ-rings". In: *Journal of Pure and Applied Algebra* 207.1 (2006), pp. 165–185. ISSN: 0022-4049. DOI: https://doi.org/10.1016/j.jpaa.2005.09.003.

[3] William Fulton and Joe Harris. *Representation theory: A first course*. Springer, 2004.

[4] Darij Grinberg and Victor Reiner. *Hopf Algebras in Combinatorics*. 2020. arXiv: 1409.8356 [math.CO].

[5] D. Knutson. *Lambda-rings and the representation theory of the symmetric group*. Springer, 1973.

[6] Sujit Rao and Joe Suk. *Dihedral Sieving Phenomena*. 2019. arXiv: 1710.06517 [math.CO].

[7] V. Reiner, D. Stanton, and D. White. "The cyclic sieving phenomenon". In: *Journal of Combinatorial Theory, Series A* 108.1 (2004), pp. 17–50. ISSN: 0097-3165. DOI: 10.1016/j.jcta.2004.04.009.

[8] Victor Reiner, Dennis Stanton, and Dennis White. "WHAT IS...Cyclic Sieving?" In: *Notices of the American Mathematical Society* 61 (Feb. 2014), pp. 169–171. DOI: 10.1090/noti1084.

[9] Bruce E. Sagan. *The cyclic sieving phenomenon: a survey*. 2011. arXiv: 1008.0790 [math.CO].

[10] John R Stembridge. "Some hidden relations involving the ten symmetry classes of plane partitions". In: *Journal of Combinatorial Theory, Series A* 68.2 (1994), pp. 372–409. ISSN: 0097-3165. DOI: https://doi.org/10.1016/0097-3165(94)90112-0.

[11] Marcus Zibrowius. *Symmetric Representation Rings are λ-Rings*. 2015. arXiv: 1308.0796 [math.KT].