

CDM

Pólya-Redfield Theory

KLAUS SUTNER

CARNEGIE MELLON UNIVERSITY

SPRING 2024



1 Pólya-Redfield

2 Lamplighters

Let's push the ideas from Burnside's lemma a little bit further to make it easier to deal with questions involving different types of configurations. First, a slight abstraction.

Let

$$V = \{v_1, \dots, v_n\}$$

be a set of **vertices** (arbitrary objects), and

$$C = \{c_1, \dots, c_m\}$$

a **set of colors**. A **coloring** is a map $V \rightarrow C$.

We write $X = V \rightarrow C$ for the set of all colorings.

Let G be a subgroup of \mathfrak{S}_n and define the natural left action of G on X by

$$\begin{aligned}\rho \cdot f &= \rho \circ f \\ [n] &\xrightarrow{\rho} [n] \xrightarrow{f} [m]\end{aligned}$$

Recall that we interpret $g \circ f$ the civilized way: first g , then f .

Thus, $\rho \in G$ permutes the vertices, and then we apply the given coloring map.

Question:

What are the invariant elements $\rho \cdot f = f$ under this action?

Clearly, if ρ is just a single cycle, then f must be constant on the whole cycle.

In the general case, we consider the cycle decomposition of ρ .

Write the cycle decomposition of ρ , including fixed points, as:

$$\rho = (v_{1,1}, \dots, v_{1,q_1}), (v_{2,1}, \dots, v_{2,q_2}), \dots, (v_{p,1}, \dots, v_{p,q_p})$$

Thus, the $v_{i,j}$ are all distinct and $\sum q_i = n$. We write

$$\text{cc}_i(\rho) = \text{number of cycles of length } i \text{ in } \rho$$

$$\#c(\rho) = \sum \text{cc}_i(\rho)$$

The list $\text{cc}_1(\rho), \text{cc}_2(\rho), \dots, \text{cc}_n(\rho)$ is the **cycle shape** of ρ .

Thus, $\#c(\rho)$ is the total number of cycles in ρ , $1 \leq \#c(\rho) \leq n$. Also note that $\sum_i i \text{cc}_i(\rho) = n$.

Lemma

The cardinality of X_ρ is $m^{\#c(\rho)}$.

Proof.

f is in X_ρ iff f is constant on all the cycles of ρ .

But there are exactly m choices for the value of f on any one of the cycles.

□

Note, though, that this requires knowledge of the cycle number for each group element.

How about the cycle structure of all elements of D_n ?

For pure rotations α^k the cycle structure is easy: there are $\gcd(n, k)$ many cycles of length $n / \gcd(n, k)$ each.

But, as we have seen, for reflections things are more complicated; we have to deal with 2-cycles and possibly fixed points.

For example, in an octagon there is a reflection (ignoring convention, we write fixed points for clarity)

$$\rho = (1)(2, 8)(3, 7)(4, 6)(5)$$

From Burnside's lemma we immediately have the

Corollary

The number of distinct orbits is $N = \frac{1}{|G|} \sum_{\rho \in G} m^{\#c(\rho)}$.

For example, when D_4 is acting on the square we have

ρ	$\#c(\rho)$	ρ	$\#c(\rho)$
1	4	β	2
α	1	$\alpha\beta$	3
α^2	2	$\alpha^2\beta$	2
α^3	1	$\alpha^3\beta$	3

Hence

$$N = \frac{1}{8}(m^4 + 2m^3 + 3m^2 + 2m).$$

For D_5 acting on the pentagon we get

ρ	$\#c(\rho)$	ρ	$\#c(\rho)$
1	5	β	3
α	1	$\alpha\beta$	3
α^2	1	$\alpha^2\beta$	3
α^3	1	$\alpha^3\beta$	3
α^4	1	$\alpha^4\beta$	3

Hence

$$N = \frac{1}{10}(m^5 + 5m^3 + 4m).$$

Let us identify two Boolean functions f and g if

$$f(x_1, x_2, \dots, x_k) = g(x_1 \oplus a_1, x_2 \oplus a_2, \dots, x_k \oplus a_k)$$

where the a_i are bits.

So $a_i = 0$ means “leave the bit alone” and $a_i = 1$ means “flip the bit”.

We want to count the Boolean functions modulo this equivalence.

As an example, there are 256 3-bit circuits. Flipping bits we could get the number down to $256/8 = 32$ but that would require all variants to be distinct. So we should expect something like 50 (wild guess).

As we will see shortly, the relevant group here is a Boolean group.

Definition

A group is **Boolean** if every element other than the identity has order 2.

Hence, in a Boolean group, we have

$$x + x = 0$$

for all x .

The additive notation is justified by the following exercise.

Exercise

Show that every Boolean group is commutative.

Example

Let $G = \langle \mathfrak{P}(A), \Delta, \emptyset \rangle$ where Δ denotes symmetric difference:
 $X \Delta Y = (X - Y) \cup (Y - X)$. Then G is a Boolean group.

Example


Any finite Boolean group arises in the following way:

$$\mathbf{2}_k = \langle \mathbf{2}^k, \oplus, \mathbf{0} \rangle$$


where $\mathbf{2}^k$ means binary lists of length k and \oplus is point-wise exclusive or.
Note that $\mathbf{2}_k$ has k generators $e_i = (0, \dots, 0, 1, 0, \dots, 0)$.

By Cayley's theorem we can identify $\mathbf{2}_k$ with a subgroup H of \mathfrak{S}_{2^k} , the full permutation group on 2^k points.


The permutation \hat{a} associated with $a \in \mathbf{2}_k$ is the map $x \mapsto x \oplus a$.




The function diagrams of the permutations \hat{e}_i for $k = 4$.



These are all the 8 permutations of the input values of a 3-bit circuit we can produce by flipping some of the bits. Figure out which bits were flipped.



A list plot of the permutation of 2^8 arising from flipping bits 2 and 3.



Same for flipping all bits. Somewhat surprising.

We can think of a Boolean function $f : \mathbf{2}^k \rightarrow \mathbf{2}$ as a coloring of points $V = \mathbf{2}^k$ by just two colors (so $n = 2^k$ and $m = 2$).

Now the Boolean group $\mathbf{2}_k$ acts on $\mathbf{2}^k$ in the obvious way:

$$\beta \cdot x = \beta \oplus x.$$

But then either β is the identity (in which case $\beta \cdot x = x$) or $\beta^2 \cdot x = x$. So the number of fixed points is either

$$2^{2^k} \quad \text{or} \quad 2^{2^{k-1}}$$

For example, here is the case where $k = 3$ and $\beta = 1$ (flip all bits).

x	$1 \cdot x$	x	$1 \cdot x$
000	111	100	011
001	110	101	010
010	101	110	001
011	100	111	000

For a Boolean function to be invariant, half the values are determined by the other half. The table shows the case $\beta = 1$, where the first half determines the second half, but it is easy to see that this actually works for any β .

Hence the total number of orbits is

$$1/2^k \left(2^{2^k} + (2^k - 1)2^{2^{k-1}} \right) = 2^{2^{k-1}-k} \left(2^{2^{k-1}} + 2^k - 1 \right)$$

For $k = 1, \dots, 5$ we obtain the following values:

$$3, 7, 46, 4336, 134281216$$

Thus, flipping input bits reduces the number of ternary Boolean functions from 256 to 46. Quite remarkable.

The last result deals just with flipping input bits. Of course, there are several other natural operations we could modify a given circuit to obtain others – and even combinations thereof, see below.

- Flip the output bit.
- Reverse inputs.
- Rotate inputs.
- Permute inputs.

Exercise

Count the number of distinct circuits for some of these operations. How hard would it be to deal with combinations thereof?

So far, we have a good tool to count the total number of orbits.

We even get a general formula that depends only on the group G and applies to different kinds of configurations: m is just the number of colors, and the orbit count is a polynomial in m .

But we do not yet know how to determine the size of a specific orbit along the lines of the Tic-Tac-Toe problem: currently, there is somewhat tedious special computation for each given configuration.

To ameliorate this problem we will need to take a closer look at the cycle shape of the permutation, the number of cycles of each possible length:

$$(\text{cc}_1(\rho), \text{cc}_2(\rho), \dots, \text{cc}_n(\rho))$$

Define a monomial in n variables for each $\rho \in G$, **cycle index monomial** for ρ , by

$$Z_\rho(x_1, \dots, x_n) := x_1^{\text{cc}_1(\rho)} x_2^{\text{cc}_2(\rho)} \dots x_n^{\text{cc}_n(\rho)}$$

and the **cycle index polynomial** to be the sum of all these:

$$Z_G(x) := \frac{1}{|G|} \sum_{\rho \in G} Z_\rho(x)$$

Corollary

The number of distinct orbits is $N = Z_G(m, m, \dots, m)$.

Why should it be any better to write the polynomial

$$x_1^{\text{cc}_1(\rho)} x_2^{\text{cc}_2(\rho)} \dots x_n^{\text{cc}_n(\rho)}$$

rather than the more pedestrian cycle shape vector

$$(\text{cc}_1(\rho), \text{cc}_2(\rho), \dots, \text{cc}_n(\rho))$$

never mind the flaky notation

$$1^{\text{cc}_1(\rho)} + 2^{\text{cc}_2(\rho)} + \dots + n^{\text{cc}_n(\rho)}$$

we briefly used last time? The same information is conveyed in all three cases, we can easily translate back and forth.

True, but polynomials come equipped with a number of operations: addition, multiplication, substitution of integers, substitution of other polynomials. This is the whole idea behind **generating functions**. Ours here are finite (polynomials), but they are just as useful as their infinite cousins.

For example, for $G = D_4$ we have

$$Z_G(\mathbf{x}) = \frac{1}{8}(x_1^4 + 2x_1^2x_2 + 3x_2^2 + 2x_4).$$

Hence

$$N = Z_G(\mathbf{m}) = \frac{1}{8}(m^4 + 2m^3 + 3m^2 + 2m).$$

Addition, multiplication and substitution $x_i \mapsto m$ all conspire to produce the right value.

Not bad, but we still can't quite handle the Tic-Tac-Toe problem in style.

For $f \in X$ let

$$\text{dc}_i(f) := \text{number of vertices } v \text{ such that } f(v) = c_i$$

Note that $\sum_i \text{dc}_i(f) = n$.

Intuitively, the **weight** of a configuration is given by the **color shape** of the configuration, the vector

$$(\text{dc}_1(f), \text{dc}_2(f), \dots, \text{dc}_m(f)) \in \mathbb{N}^m$$

This provides full information about the number of vertices of each color.

Weights are an invariant in the sense that all elements in an orbit of G have the same weight: the permutations can move the colors around, but they cannot change the color counts. So we can talk about the **weight of an orbit**.

Again, it turns out to be more convenient to define the **weight** of a coloring f by the monomial

$$\text{weight}(f) := c_1^{\text{dc}_1(f)} c_2^{\text{dc}_2(f)} \dots c_m^{\text{dc}_m(f)}$$

Note that $\text{weight}(f)$ is a “purely formal expression”, the c_i are just colors.

More precisely, $\text{weight}(f)$ is an element of the polynomial ring $\mathbb{Z}[c_1, \dots, c_m]$.

Generating functions are your friend.

We will see shortly that this trick makes it easy to determine the number of orbits of a given weight.

Suppose we have $n = 8$ and $\rho = (1)(2)(3)(4)(56)(78)$.

Which configurations over three colors r, g, b are invariant under ρ ?

We can think of a coloring as a word

$$u_1 u_2 u_3 u_4 u_5 u_6 u_7 u_8$$

of length 8 over the alphabet $\{r, g, b\}$.

Since our permutation is rather lame, the only constraint for invariance is $u_5 = u_6$ and $u_7 = u_8$, other than that, anything goes. So we have a language $I \subseteq \{r, g, b\}^8$,

We can describe I by a regular expression:

$$(r + g + b)(r + g + b)(r + g + b)(r + g + b)(r^2 + g^2 + b^2)(r^2 + g^2 + b^2) = \\ (r + g + b)^4(r^2 + g^2 + b^2)^2$$

By expanding this regex we get an explicit list of the words in I , a language of cardinality $729 = 3^6$.

So there are 729 invariant configurations. Nice, but we would like to answer questions about configurations with a specific weight (remember tic-tac-toe).

Types, Schmypes: We can also think of this regex as a polynomial in $\mathbb{Z}[r, g, b]$.

Now the variables commute and we can expand the expression, and collect terms with the same weight. We wind up 45 terms, each of degree 8.

$$\begin{aligned}
 & b^8 + 4b^7g + 8b^6g^2 + 12b^5g^3 + 14b^4g^4 + 12b^3g^5 + 8b^2g^6 + 4bg^7 + g^8 + 4b^7r + \\
 & 12b^6gr + 20b^5g^2r + 28b^4g^3r + 28b^3g^4r + 20b^2g^5r + 12bg^6r + 4g^7r + \\
 & 8b^6r^2 + 20b^5gr^2 + 32b^4g^2r^2 + 40b^3g^3r^2 + 32b^2g^4r^2 + 20bg^5r^2 + 8g^6r^2 + \\
 & 12b^5r^3 + 28b^4gr^3 + 40b^3g^2r^3 + 40b^2g^3r^3 + 28bg^4r^3 + 12g^5r^3 + 14b^4r^4 + \\
 & 28b^3gr^4 + 32b^2g^2r^4 + 28bg^3r^4 + 14g^4r^4 + 12b^3r^5 + 20b^2gr^5 + 20bg^2r^5 + \\
 & 12g^3r^5 + 8b^2r^6 + 12bgr^6 + 8g^2r^6 + 4br^7 + 4gr^7 + r^8
 \end{aligned}$$

Consider the term $32 r^2 g^2 b^4$.

The 32 choices for an invariant configuration of weight $r^2 g^2 b^4$ are

rrggbbbb, rrbbggbb, rrbbbb, rrggbbbb, rgrgbbbb, rgrgbbbb, rbrbggbb, rbrbbb, rbbgggbb, rbbrrbbg, grrgbbbb, grgrbbbb, ggrrbbbb, ggbbrrbb, ggbbrrrr, gbgrrrrb, gbgbbrrr, gbbgrrbb, gbbgrrrr, brrbggbb, brrbbb, brbrggbb, brbrbbgg, bggbrrbb, bggbrrrr, bbggrrbb, bbggrrrr, bbrrggbb, bbrrrbbg, bbggrrbb, bbggrrrr, bbbrrrgg, bbbggrrr

These are precisely the words over alphabet $\{r, g, b\}$ with weight $(2, 2, 4)$ subject to the constraint $u_5 = u_6$ and $u_7 = u_8$.

In our example, we have $Z_\rho(x) = x_1^4 x_2^2$.

Our magic polynomial is just $Z_\rho(r + g + b, r^2 + g^2 + b^2)$.

This all hinges on the fact that we can express logic by algebra: we really need to consider choices of colors on the cycles of a given permutation.

But we can translate this counting problem into polynomial arithmetic by introducing formal variables for the colors and then working in the polynomial ring $\mathbb{Z}[r, g, b]$. It is quite surprising how much mileage one can get out of generating functions.

Since we can perform the arithmetic operations easily, this translation allows for easy computation – sort of, a computer algebra system might come in handy.

Here is a more general description of this method.

Let

$$N(a_1, a_2, \dots, a_m) := \# \text{ orbits with weight } c_1^{a_1} c_2^{a_2} \dots c_m^{a_m}.$$

and define the **pattern inventory** of G on X to be

$$\sum_a N(a_1, a_2, \dots, a_m) c_1^{a_1} c_2^{a_2} \dots c_m^{a_m}$$

This is a (typically huge) polynomial in variables c_1, \dots, c_m whose coefficients contain exactly the counting information we are after.

Nice, but utterly useless unless we can somehow compute the pattern inventory (without resorting to brute force, of course).

Theorem (Pólya-Redfield)

Set $y_i = c_1^i + c_2^i + \dots + c_m^i$. Then the pattern inventory is $Z_G(y_1, y_2, \dots, y_n)$.

Likewise, $Z_\rho(y_1, y_2, \dots, y_n)$ is the generating function for the number of patterns of the indicated weight, invariant under ρ .

It may still seem rather difficult to compute the pattern inventory, but with a little computer algebra it is not too hard: we need to deal with permutation groups and some polynomial algebra.

First we compute all the 8 monomials, based on cycle shapes. Here are some examples.

ρ	Z_ρ	ρ	Z_ρ
1	x_1^4	β	x_2^2
α	x_4	$\alpha\beta$	$x_1^2x_2$
α^2	x_2^2	$\alpha^2\beta$	x_2^2
α^3	x_4	$\alpha^3\beta$	$x_1^2x_2$

Then we take the average of all the monomials and we get the CIP for D_4 :

$$\frac{1}{8}(x_1^4 + 2x_1^2x_2 + 3x_2^2 + 2x_4)$$

Assume $m = 2$.

$$Z_{\alpha}(\mathbf{y}) = c_1^4 + c_2^4$$


$$Z_{\alpha^2}(\mathbf{y}) = c_1^4 + 2 c_1^2 c_2^2 + c_2^4$$

$$Z_{\alpha^3\beta}(\mathbf{y}) = c_1^4 + 2 c_1^3 c_2 + 2 c_1^2 c_2^2 + 2 c_1 c_2^3 + c_2^4$$

Lastly, here is the pattern inventory:

$$Z_G(\mathbf{y}) = c_1^4 + c_1^3 c_2 + 2 c_1^2 c_2^2 + c_1 c_2^3 + c_2^4$$

There are 6 orbits, two of weight $c_1^2 c_2^2$ (columns 3 and 4 below).



Pattern inventory for three colors for $n = 4$.

$$c_1^4 + c_1^3 c_2 + 2c_1^2 c_2^2 + c_1 c_2^3 + c_2^4 + c_1^3 c_3 + 2c_1^2 c_2 c_3 + 2c_1 c_2^2 c_3 + \\ c_2^3 c_3 + 2c_1^2 c_3^2 + 2c_1 c_2 c_3^2 + 2c_2^2 c_3^2 + c_1 c_3^3 + c_2 c_3^3 + c_3^4$$

Lastly, two colors and $n = 10$.

$$\frac{1}{20} (x_1^{10} + 5x_1^2 x_2^4 + 6x_2^5 + 4x_5^2 + 4x_{10})$$

$$c_1^{10} + c_1^9 c_2 + 5c_1^8 c_2^2 + 8c_1^7 c_2^3 + 16c_1^6 c_2^4 + 16c_1^5 c_2^5 + \\ 16c_1^4 c_2^6 + 8c_1^3 c_2^7 + 5c_1^2 c_2^8 + c_1 c_2^9 + c_2^{10}$$









For the Tic-Tac-Toe problem we have $|V| = 9$ and $|C| = 3$.

The group G acting on the board is isomorphic to D_4 but has degree 9. The cycle decompositions G are different and the CIP is:

$$\frac{1}{8}(x_1^9 + 2x_1x_4^2 + x_1x_2^4 + 4x_1^3x_2^3)$$

Of the 1680 333-patterns, 36 have orbits of length 4 and 192 have orbits of length 8.

If we only consider patterns of size 4 but additionally allow for permutation of colors, there are only 8 super-patterns:



CIP

$$\frac{1}{8}(c_1+c_2+c_3)\left((c_1+c_2+c_3)^8+4(c_1+c_2+c_3)^2(c_1^2+c_2^2+c_3^2)^3+(c_1^2+c_2^2+c_3^2)^4+2(c_1^4+c_2^4+c_3^4)^2\right)$$

Pattern inventory, note the $228c_1^3c_2^3c_3^3$ term.

$$\begin{aligned} & c_1^9+3c_1^8c_2+8c_1^7c_2^2+16c_1^6c_2^3+23c_1^5c_2^4+23c_1^4c_2^5+16c_1^3c_2^6+8c_1^2c_2^7+3c_1c_2^8+c_2^9+3c_1^8c_3+ \\ & 12c_1^7c_2c_3+38c_1^6c_2^2c_3+72c_1^5c_2^3c_3+89c_1^4c_2^4c_3+72c_1^3c_2^5c_3+38c_1^2c_2^6c_3+12c_1c_2^7c_3+ \\ & 3c_2^8c_3+8c_1^7c_3^2+38c_1^6c_2c_3^2+108c_1^5c_2^2c_3^2+174c_1^4c_2^3c_3^2+174c_1^3c_2^4c_3^2+108c_1^2c_2^5c_3^2+ \\ & 38c_1c_2^6c_3^2+8c_2^7c_3^2+16c_1^6c_3^3+72c_1^5c_2c_3^3+174c_1^4c_2^2c_3^3+228c_1^3c_2^3c_3^3+174c_1^2c_2^4c_3^3+ \\ & 72c_1c_2^5c_3^3+16c_2^6c_3^3+23c_1^5c_3^4+89c_1^4c_2c_3^4+174c_1^3c_2^2c_3^4+174c_1^2c_2^3c_3^4+89c_1c_2^4c_3^4+ \\ & 23c_2^5c_3^4+23c_1^4c_3^5+72c_1^3c_2c_3^5+108c_1^2c_2^2c_3^5+72c_1c_2^3c_3^5+23c_2^4c_3^5+16c_1^3c_3^6+38c_1^2c_2c_3^6+ \\ & 38c_1c_2^2c_3^6+16c_2^3c_3^6+8c_1^2c_3^7+12c_1c_2c_3^7+8c_2^2c_3^7+3c_1c_3^8+3c_2c_3^8+c_3^9 \end{aligned}$$

1 Pólya-Redfield

2 **Lamplighters**

Most of the groups we have used in Pólya counting so far are fairly straightforward: they correspond directly to geometric symmetries. Of course, things get messier in higher dimensions, but still.

This might create the impression that the groups associated with actions on combinatorial objects are always easy to understand. Not so. Sometimes it requires quite a bit of effort to come up with the right group that describes some (intuitively clear) action.

Here is a notorious example.

Suppose you have a ring of n lamps; each lamp is either on or off.



There is an eponymous lamplighter, some dude who walks around and turns lights on and off.

The lamplighter can perform two atomic actions:

α move to the next lamp, or

τ toggle the state of the current lamp.

The actions are clearly reversible, so there must be a group plus action hiding somewhere.



It is obvious that $\alpha^n = 1$ and $\tau^2 = 1$.

The group does not commute, $\alpha\tau \neq \tau\alpha$ (assuming $n > 1$).

But what exactly is the group, and how does it act?

Clearly we can describe the space of configurations as

$$X = 2^n \times \mathbb{Z}_n$$



So we are dealing with bitvectors and modular numbers.

The picture shows the configurations (000000, 0) and (101100, 2).

We need the group G generated by α and τ , something like

$$\langle \alpha, \tau \mid \alpha^n = 1, \tau^2 = 1, \dots \rangle$$

but we don't have all the necessary identities for this kind of description. For example, it is true that

$$(\tau \alpha \tau \alpha^{n-1})^2 = 1$$

It seem like a good idea to try to write the group elements in two parts:

- part 1 describes the changes in lights (toggle pattern), and
- part 2 describes the movement of the lamplighter.

So G should be a product, something like $G = A \times B$.

Traditionally, one considers a bi-infinite row of lamps rather than a finite circle. So the configurations are

$$X = \prod_{\mathbb{Z}} \mathbf{2} \times \mathbb{Z}$$

If only finitely many lamps are ever lit (which is automatically the case if we start with all lights off), we get

$$X = \coprod_{\mathbb{Z}} \mathbf{2} \times \mathbb{Z}$$

So the space of configurations is fairly straightforward in all three cases, but we need to figure out what exactly the group is. We'll ignore these infinite groups.

Recall our general strategy: we want to combine the atomic operations and their inverses in arbitrary ways. For example, there has to be a group element for “toggle 0 and 1” (coordinates are relative to the lamplighter's position). and so forth. And the group operation applied to “toggle 0 and 1” and “toggle 1 and 2” should be “toggle 0 and 2”.

So one would expect something like

$$G = 2^n \times \mathbb{Z}_n$$

The first component keeps track of toggles, the second of the lamplighter's position.

(a, s) should act on (x, p) as follows: flip the lights in x according to a , then change p according to s .

Let $e_i \in \mathbf{2}^n$ be the unit vector with the 1 in position $i \in (n)$, so e.g. $e_3 = 00010000$. Then we would like to have the following interpretation:

$\tau \rightsquigarrow (e_0, 0)$: toggle the lamp at the current position.

$\alpha \rightsquigarrow (\mathbf{0}, 1)$: the lamplighter moves forward by one.

And $(e_0 \oplus e_1, 3)$ means: toggle the lights in relative positions 0 and 1, then move forward by 3 places, corresponding to $\tau\alpha\tau\alpha^2$.

Does $G = \mathbf{2}^n \times \mathbb{Z}_n$ properly reflect this intuition? Recall that in a product group, the two group operations are applied separately in each of the two components.

$$(a, r)(b, s) = (a \oplus b, r + s)$$

Consider the group element $\tau\alpha\tau\alpha$. According to the definition of our group, we have

$$(e_0, 1)(e_0, 1) = (e_0 \oplus e_0, 2) = (\mathbf{0}, 2)$$

where \oplus stands for bit-wise xor. Our model thinks that $\tau\alpha\tau\alpha = \alpha^2$, which is clearly false. The position is right, but not the toggle pattern: simply adding the vectors does not reflect the change in position.

What we need instead as the result of the multiplication is

$$(e_0, 1)(e_0, 1) = (e_0 \oplus e_1, 2)$$

We need to take into account the position of the lamplighter: when the lamplighter moves, the next switch-vector has to be adjusted accordingly.

More technically, a modular number s acts on \mathbf{a} by rotating the sequence by s places.


We write σ for the cyclic shift. To fix our group, we keep the same carrier set, but we adjust the group operation to

$$(\mathbf{a}, r)(\mathbf{b}, s) = (\mathbf{a} \oplus \sigma^r(\mathbf{b}), r + s)$$

One might suspect that group theorists are well aware of this type of modified product: they are called **semidirect product** or **wreath product**. We are not going to explain these notions here, take a look at a standard group theory text.

Recall that $\mathbf{2}_n$ is the Boolean group of order 2^n , in the special case $n = 1$ we get the additive group \mathbb{Z}_2 , which also happens to be the symmetric group on two letters. At any rate, the lamplighter group on a ring of size n is given by the wreath product

$$G = \mathbf{2} \wr \mathbb{Z}_n = \mathbf{2}^n \rtimes \mathbb{Z}_n$$



Exercise

Verify $(\alpha\tau^i\alpha\tau^{-i})^2 = 1$.

Exercise

Interpret $(\alpha\tau^i\alpha\tau^{-i})^2 = 1$ geometrically.


Exercise

What would happen if we were to use the “move first, then toggle” interpretation?

The real, infinite Lamplighter group has lots of interesting properties and has been studied extensively; Google Scholar shows 15,900 hits.

One fun fact: the lamplighter group is closely related to finite state machines.

Consider the following 2-state alphabetic transducer:



Note that state 0 copies the next bit while state 1 flips it. That's in, just copy or flip.

It is easy to see that picking one of the states k as initial state defines a permutation \underline{k} of all binary words. Moreover, the permutation is length and prefix preserving; it can be pictured as an automorphism of the infinite binary tree.

The automaton is just an elegant description of a recursive system of equations:

$$\underline{0}(ax) = a \underline{a}(x)$$

$$\underline{1}(ax) = \bar{a} \underline{a}(x)$$

But then we can consider the group G generated by the permutations $\underline{0}$ and $\underline{1}$, a subgroup of the group of all automorphism of the binary tree.

Big Surprise: G is isomorphic to the lamplighter group.

Let us revisit the problem of counting Boolean functions modulo equivalence. We know how to handle two cases:

- Permutation of variables: \mathfrak{S}_n acting on 2^n
- Negation of variables: 2_n acting on 2^n

It is tempting to ask how to deal with permutation and negation together.

For example

$$x \wedge (y \vee z) \text{ and } z \wedge (\bar{x} \vee y)$$

are in a sense the same Boolean function: naming of variables does not really matter and exchanging x and \bar{x} does not matter (much) either.

We refer to two Boolean functions that can be obtained from each other by permuting and/or negating variables as **PN-equivalent**.

Counting the number of PN-equivalence classes should be a problem that falls well within the reach of the counting methods we have just seen.

It is even clear that $V = 2^n$ and $C = 2$ is the right framework.

Question: What is the group G acting on V in this case?

No doubt both \mathfrak{S}_n and 2_n are involved, but exactly how?

Note that we can figure out the cardinality of G : $2^n \cdot n!$

An entirely reasonable first guess would be that

$$G = \mathbf{2}_n \times \mathfrak{S}_n$$

is the right group; after all, the pieces make sense and we get the right size.

Alas, that's plain wrong. To see why, consider $g = (e_1, (12)) \in G$. Informally, g means “negate the first variable, then swap it with the second.” The effect of g^2 ought to be something like

$$x y z \dots \mapsto y \bar{x} z \mapsto \bar{x} \bar{y} z \dots$$

Disaster strikes: in G we have $g^2 = 1$.

Both in the full and restricted wreath product we considered so far, the group B acts on sequences, $\prod_B A$ or $\coprod_B A$.

But we also could have B acting on some set X and then consider sequences in $\prod_X A$ or $\coprod_X A$.

The definitions don't change at all: we can still “multiply” a sequence index with a group element in B .

That's all we need!

One writes $A \wr_X B$ for this version of the wreath product.

The problem is again that the ordinary Cartesian product simply fails here, we need something slightly more complicated: the (generalized) wreath product $G = \mathfrak{S}_2 \wr_{[n]} \mathfrak{S}_n$. Writing the operations additively and multiplicatively, we have in G :

$$(x, f)(y, g) = (x + f * y, fg)$$

where

$$f * y = (y_{f(1)}, \dots, y_{f(n)}).$$

is the vanilla left action.

If you find this description of G a bit obscure, think of permutation matrices instead: a permutation matrix of order n is 0/1-matrix of size n by n such that every column and every row contains exactly a single 1.

It is not hard to see that these matrices correspond precisely to permutations of $[n]$. Under ordinary matrix multiplication, they produce a group that is isomorphic to \mathfrak{S}_n .

To represent $G = \mathfrak{S}_2 \wr_{[n]} \mathfrak{S}_n$ we need to use **signed permutation matrices**: the entries now can be $+1$ or -1 .

Again, with ordinary matrix multiplication, we obtain a group isomorphic to G . Note that the cardinalities match up.

Exercise

Verify the last claim.

Still, for very small n we can compute G directly and determine its cycle index polynomial. For example, for $n = 5$ we get

$$\frac{1}{3840} \left(x_1^{32} + 20x_1^{16}x_2^8 + 60x_1^8x_2^{12} + 231x_2^{16} + 80x_1^8x_3^8 + 240x_1^4x_2^2x_4^6 + 240x_2^4x_4^6 + 520x_4^8 + 384x_1^2x_5^6 + 160x_1^4x_2^2x_3^4x_6^2 + 720x_2^4x_6^4 + 480x_8^4 + 384x_2x_{10}^3 + 320x_4^2x_{12}^2 \right)$$

The number of PN-equivalent Boolean functions of 5 variables is therefore

$$1,228,158$$

as opposed to 4,294,967,296 functions in the set-theoretic sense.

A brute force computation of CIPs as on the last slide is not really the right answer: we should try to compute the CIP for G given the CIPs for the Boolean group $\mathbf{2}_n$ and the full symmetric group \mathfrak{S}_n .

The good news is: this can be done and produces a computationally satisfactory answer.

The bad new is: the calculation is quite complicated and would lead us too far astray.