

Toggling independent sets of a path graph

Michael Joseph

Department of Technology and Mathematics
Dalton State College
Dalton, GA 30720, U.S.A.
mjoseph@daltonstate.edu

Tom Roby

Department of Mathematics
University of Connecticut
Storrs, CT 06269-1009, U.S.A.
tom.robby@uconn.edu

Submitted: Jan 18, 2017; Accepted: Jan 9, 2018; Published: Jan 25, 2018
Mathematics Subject Classifications: 05E18

Abstract

This paper explores the orbit structure and homomesy (constant averages over orbits) properties of certain actions of toggle groups on the collection of independent sets of a path graph. In particular we prove a generalization of a homomesy conjecture of Propp that for the action of a “Coxeter element” of vertex toggles, the difference of indicator functions of symmetrically-located vertices is 0-mesic. Then we use our analysis to show facts about orbit sizes that are easy to conjecture but nontrivial to prove. Besides its intrinsic interest, this particular combinatorial dynamical system is valuable in providing an interesting example of (a) homomesy in a context where large orbit sizes make a cyclic sieving phenomenon unlikely to exist, (b) the use of Coxeter theory to greatly generalize the set of actions for which results hold, and (c) the usefulness of Striker’s notion of generalized toggle groups.

Keywords: Burnside’s Lemma, composition, Coxeter element, homomesy, independent set, involution, orbit, path graph, promotion, rowmotion, toggle group, zigzag poset

1 Introduction

The past several years have seen renewed interest in the area of *dynamical algebraic combinatorics*, where one considers group actions on sets of discrete combinatorial objects and looks for interesting properties of their orbits or orders. These include such actions as *promotion* (on Young tableaux or posets) and *rowmotion* of posets, but there are many others. Many of these can be built up as a sequence of simple involutions, as is the main action we study here on the collection of independent sets of a path graph.

Our goal is to understand the orbit structure and *homomesy* (constant averages over orbits) of this discrete dynamical system. In particular we prove a conjecture of Propp

that, with respect to the action “toggle once at each vertex from left to right”, the difference of indicator functions of symmetrically-located vertices is 0-mesic (Theorem 2.9). By leveraging some basic theory of Coxeter groups, we generalize this result to apply to *any* Coxeter element in the toggle group. Our analysis allows us to deduce facts about orbit sizes that are apparent from the numerical data, but seem difficult to prove directly.

This particular combinatorial dynamical system serves as a valuable case study in several respects. First, it provides an interesting example of the homomesy phenomenon in a context where unwieldy orbit sizes suggest the lack of a natural cyclic sieving phenomenon (CSP) in the sense of Reiner, Stanton, and White [RSW04]. Many combinatorial dynamical systems that support a CSP also have natural homomesic statistics and vice versa, though there appears to be no direct connection between the two (even in specific cases).

Second, by taking a Coxeter theoretic approach, we are able to greatly generalize the set of actions for which our results hold, from the specific action of successively toggling at each vertex to toggling once per vertex in an arbitrary order (Subsection 2.3).

Third, this combinatorial dynamical system displays the usefulness of Striker’s notion of generalized toggle groups [Str16] to settings beyond that of posets. Although there is an equivariant bijection (Proposition 5.6) between the action we study on independent sets and the action of promotion on zigzag posets, it is much easier to establish the homomesy in the former setting first, then translate it to the latter setting.

We now describe the setting and background necessary to understand the problem.


Definition 1.1. Let \mathcal{P}_n denote the **path graph** with vertex set $[n] := \{1, 2, \dots, n\}$ and edge set $\{\{i, i + 1\} : i \in [n - 1]\}$.

Example 1.2. The path graph with seven vertices is



Definition 1.3. An **independent set** of a graph is a subset of the vertices that does not contain a pair of adjacent vertices. Let \mathcal{I}_n denote the set of independent sets of \mathcal{P}_n .

Example 1.4. The set of vertices $\{1, 4, 6\}$ represented



is an independent set of \mathcal{P}_7 , but $\{1, 4, 5, 6\}$ represented



is not. In both of these examples, hollow dots refer to vertices of \mathcal{P}_7 not in the subset.

Although we sometimes write independent sets as subsets of $[n] := \{1, 2, \dots, n\}$ as above, it may not be obvious in that notation what the underlying value of n is. Another notation that is often more convenient for an independent set is its **binary representation**, in which the bit in position i of S is 0 if $i \notin S$ and 1 if $i \in S$. For example 0010010

represents the independent set $\{3, 6\}$ of \mathcal{P}_7 . Thus \mathcal{I}_n can be viewed as the set of length n binary strings that do not contain the subsequence 11 (which would indicate the inclusion of two adjacent vertices). It is well-known and easy to verify that the cardinality of \mathcal{I}_n is a Fibonacci number.

In Section 2, we introduce the *toggle group* \mathcal{T}_n on \mathcal{I}_n , which is generated by basic involutions called *toggles*. The notion of toggle groups goes back to work of Cameron and Fon-der-Flaass, who introduced toggles on *order ideals* of a poset [CF95]. More recently, Striker has studied toggle groups in a more general setting [Str16]. Specifically, given a “ground” set X and a fixed set of “allowed” subsets $\mathcal{L} \subseteq 2^X$, each element $x \in X$ has an associated toggle which removes or inserts x into any given set in \mathcal{L} provided the resulting set is still in \mathcal{L} , and otherwise does nothing. In our situation the ground set is $[n]$ and the set of allowed subsets of $[n]$ is \mathcal{I}_n .

Our main results are a proof of a conjecture of Propp (Theorem 2.9) and a generalization of it (Corollary 2.31). These theorems give examples of the *homomesy* (Greek for “same middle”) *phenomenon*, introduced by Propp and the second author in [PR15] and defined as follows.

Definition 1.5. Suppose we have a set \mathcal{S} , an invertible map $w : \mathcal{S} \rightarrow \mathcal{S}$ such that every w -orbit is finite, and a function (“statistic”) $f : \mathcal{S} \rightarrow \mathbb{K}$, where \mathbb{K} is a field of characteristic 0. Then we say the triple (\mathcal{S}, w, f) exhibits **homomesy** if there exists a constant $c \in \mathbb{K}$ such that for every w -orbit $\mathcal{O} \subseteq \mathcal{S}$,

$$\frac{1}{\#\mathcal{O}} \sum_{x \in \mathcal{O}} f(x) = c.$$

In this case, we say that the function f is **homomesic with average c** , or **c -mesic**, under the action of w on \mathcal{S} .

Some early isolated examples of homomesy exist in the literature, notably a conjecture of Panyushev [Pan09, Conjecture 2.1(iii)] for the rowmotion operator acting on antichains in positive root posets. This was proven by Armstrong, Stump, and Thomas [AST13], but investigation of homomesy as a widespread phenomenon is more recent. Examples now include cyclic actions on partitions, Suter’s action on Young diagrams, rowmotion and promotion of order ideals, Lyness 5-cycles (which has strong connections to cluster algebra theory), and certain toggling actions for noncrossing partitions [EP13, EFGJMPR, Had16, PR15, Rob16, SW12, Str15].

Although it is a new area of research, the homomesy phenomenon has been discovered in a wide variety of combinatorial dynamical systems. Most of the initially proven homomesy results were for systems in which the order of the map was known. In fact, in many cyclic actions where homomesy is present, one also finds the cyclic sieving phenomenon of Reiner, Stanton, and White [RSW04]. (See also Sagan’s more leisurely exposition of the basic ideas and examples [Sag11].) There have long been conjectured homomesies for maps with unpredictable orbit sizes, but the first such proven result came out of a team (including the first author) assembled at an AIM workshop [EFGJMPR]. In this paper,

we eventually discuss how to determine the sizes of the orbits under our maps, but they do not divide a number that is easy to describe without listing all of them.

To prove Propp’s original conjecture (Theorem 2.9), we associate an *orbit board* to each orbit, and partition the ones in the orbit board into *snakes* which begin in the left column and end in the right column. Our technique was inspired by Haddadan’s proof that the “winching” action on k -element subsets of $[n]$ exhibits homomesy (another conjecture of Propp) [Had16]. In addition to proving homomesy, the snake representations lead to many other results, e.g., on the total number of orbits or the existence of orbits of certain sizes, the latter being the focus of Section 4. We first consider a specific action $\varphi \in \mathcal{T}_n$ which toggles left to right, but most of our results can be generalized to certain other actions in \mathcal{T}_n named *Coxeter elements*. This is the focus of Subsection 2.3. We use some theory of Coxeter groups to explain why we can extend proven results for φ to other actions, and this is why we do not start with the more general results.

In Section 5, we explain how our results can be restated in terms of toggling order ideals of zigzag posets via an equivariant bijection $\eta : \mathcal{I}_n \rightarrow J(\mathcal{Z}_n)$, where $J(\mathcal{Z}_n)$ is the set of order ideals of the zigzag poset \mathcal{Z}_n . We describe the connections with the well-studied *promotion* and *rowmotion* operators on the set of order ideals of a poset [SW12]. However, the proofs of our main results are much easier to obtain by working with toggles on \mathcal{I}_n as opposed to those on $J(\mathcal{Z}_n)$, which shows the significance of considering toggle groups in Striker’s generalized setting, instead of the original setting of toggling order ideals from Cameron and Fon-der-Flaass. In fact, the independent sets of \mathcal{P}_n are the antichains of \mathcal{Z}_n in disguise. In [Str16, §3.3, 3.6], Striker discusses toggles on antichains of posets and on independent sets of graphs. In [Jos17], the first author describes an explicit isomorphism between the toggle groups of antichains and of order ideals for a general poset.

2 Toggle Maps on Independent Sets

In this section we state and prove our main homomesy results. Throughout this paper we assume $n \geq 2$. While some of our results also hold for $n = 1$, many do not, and we are not concerned with this trivial case.

2.1 Definitions and main results

We now define the toggles on \mathcal{I}_n .

Definition 2.1. For every $i \in [n]$, define $\tau_i : \mathcal{I}_n \rightarrow \mathcal{I}_n$, the **toggle at vertex i** , in the following way. If $i \in S$, τ_i removes i from S , which still results in an independent set. If $i \notin S$, then $\tau_i(S)$ adds i to S assuming the resulting set is still independent, and otherwise does nothing. Formally,

$$\tau_i(S) = \begin{cases} S \setminus \{i\} & \text{if } i \in S, \\ S \cup \{i\} & \text{if } i \notin S \text{ and } S \cup \{i\} \in \mathcal{I}_n, \\ S & \text{if } i \notin S \text{ and } S \cup \{i\} \notin \mathcal{I}_n. \end{cases}$$

Since the τ_i operate to the left of their arguments, we use the standard convention that a product of toggles is performed from right to left. It is clear that each τ_i is an involution, i.e., τ_i^2 is the identity. We characterize the order of products of two toggles in the following propositions.

Corollary 2.2. *The toggles τ_i and τ_j commute if and only if $|i - j| \neq 1$.*

Proof. If $i = j$, then τ_i and τ_j clearly commute.

Suppose $|i - j| > 1$. Then whether or not i is in an independent set has no effect on whether or not j can be in that set and vice versa. So $\tau_i\tau_j = \tau_j\tau_i$.

Suppose $|i - j| = 1$. Then $\tau_i(\tau_j(\emptyset)) = \{j\}$ and $\tau_j(\tau_i(\emptyset)) = \{i\}$, so $\tau_i\tau_j \neq \tau_j\tau_i$. ■

Corollary 2.3. *When $n \geq 3$, the order of the map $\tau_i \circ \tau_j$ is*

$$\begin{cases} 1 & \text{if } i = j, \\ 2 & \text{if } |i - j| \geq 2, \\ 6 & \text{if } |i - j| = 1. \end{cases}$$

Proof. The proofs of the first two cases are straightforward since toggles are involutions.

Suppose $|i - j| = 1$. Since τ_i and τ_j are involutions, $(\tau_i\tau_j)^{-1} = \tau_j\tau_i$, so $\tau_i\tau_j$ has the same order as $\tau_j\tau_i$. Thus, we may assume without loss of generality that $i < j$ (so $j = i + 1$). To show that the order of $\tau_i\tau_{i+1}$ is 6, we will show that there is an orbit of size 2, an orbit of size 3, and no orbit with size greater than 3.

Note that the toggles τ_i and τ_{i+1} can only affect whether i and/or $i + 1$ are in a given set, and no independent set can contain both. Thus, every orbit under the action of $\tau_i\tau_{i+1}$ can at the very most contain S , $S \cup \{i\}$, $S \cup \{i + 1\}$, for some $S \in \mathcal{I}_n$. Therefore all orbits have size at most 3.

The orbit $(\emptyset, \{i + 1\}, \{i\})$ has size 3. For $i \geq 2$, the orbit $(\{i - 1\}, \{i - 1, i + 1\})$ has size 2. If $i = 1$, the orbit $(\{3\}, \{1, 3\})$ has size 2. (This is why we needed $n \geq 3$, as the map $\tau_1\tau_2$ on \mathcal{I}_2 has order 3, not 6.) ■

Definition 2.4. Let $\mathfrak{S}_{\mathcal{I}_n}$ denote the symmetric group on \mathcal{I}_n . The **toggle group** of \mathcal{I}_n , denoted \mathcal{T}_n , is the subgroup of $\mathfrak{S}_{\mathcal{I}_n}$ generated by the τ_i toggles.

Definition 2.5. A particular element in \mathcal{T}_n is $\varphi := \tau_n \cdots \tau_2\tau_1$, the map that toggles at each vertex from left to right.

Example 2.6. In \mathcal{I}_5 , $\varphi(10010) = 01001$ by the following steps:

$$10010 \xrightarrow{\tau_1} 00010 \xrightarrow{\tau_2} 01010 \xrightarrow{\tau_3} 01010 \xrightarrow{\tau_4} 01000 \xrightarrow{\tau_5} 01001.$$

Note that $\varphi^{-1} = \tau_1\tau_2 \cdots \tau_n$, which applies the toggles right to left.

Definition 2.7. Given a set $S \in \mathcal{I}_n$ and $j \in [n]$, define $\chi_j(S)$ to be the indicator function of vertex j in S . That is, $\chi_j(S)$ is the j^{th} digit of the binary representation of S .

Example 2.8. $\chi_1(10010) = 1$, $\chi_2(10010) = 0$, $\chi_3(10010) = 0$, $\chi_4(10010) = 1$, $\chi_5(10010) = 0$.

	1	2	3	4	5	6	7
S^0	1	0	1	0	1	0	0
S^1	0	0	0	0	0	1	0
S^2	1	0	1	0	0	0	1
S^3	0	0	0	1	0	0	0
S^4	1	0	0	0	1	0	1
S^5	0	1	0	0	0	0	0
S^6	0	0	1	0	1	0	1
S^7	1	0	0	0	0	0	0
S^8	0	1	0	1	0	1	0
S^9	0	0	0	0	0	0	1
Total	4	2	3	2	3	2	4

Figure 1: The orbit under the action φ on \mathcal{I}_7 containing $S = 1010100$.

One of our main theorems to be proven later is the following conjecture of Propp. We will later extend this result to actions in \mathcal{T}_n that are “Coxeter elements”, i.e., products of every τ_i exactly once in some order (Corollary 2.31).

Theorem 2.9 (Propp’s conjecture). *Under the action of φ on \mathcal{I}_n , $\chi_j - \chi_{n+1-j}$ is 0-mesic for every $1 \leq j \leq n$.*

Definition 2.10. Given an independent set $S \in \mathcal{I}_n$ and $w \in \mathcal{T}_n$, we define the **orbit board** for S and w as follows. Let $S^i = w^i(S)$ for $i \in \mathbb{Z}$ and for any $j \in [n]$, let $S(i, j) = 1$ if $j \in S^i$ and $S(i, j) = 0$ if $j \notin S^i$. (In particular, if $j < 1$ or $j > n$, then $S(i, j) = 0$. These are “out-of-bounds” positions not shown when we display the orbit board.)

Example 2.11. The orbit board for the orbit containing $S = 1010100 \in \mathcal{I}_7$ under the action of φ is shown in Figure 1. This is an orbit of size 10, so $S^{10} = \varphi^{10}(S) = S$. Technically, the orbit board is vertically infinite but periodic, so we only show S^0, S^1, \dots, S^9 and view it as living on a cylinder. The element in row i and column j is $S(i, j)$, with $i \in [0, \ell - 1]$ and $j \in [n]$, where ℓ is the length of S ’s orbit. Notice that the column-sum vector, $(4, 2, 3, 2, 3, 2, 4)$, is palindromic, illustrating Theorem 2.9, since $\chi_j - \chi_{n+1-j}$ has total 0 (and thus average 0) across this orbit for each j .

A homomesy result which is much simpler to prove is the following.

Theorem 2.12. *For $n \geq 2$, under the action of φ on \mathcal{I}_n , the statistics $2\chi_1 + \chi_2$ and $\chi_{n-1} + 2\chi_n$ are both 1-mesic.*

The reader can easily check that this holds for the orbit in Figure 1. This result can be obtained as a corollary of [EFGJMPR, Theorem 7.5], but we include another proof

here. As with Theorem 2.9, we will generalize this theorem to actions by general Coxeter elements of toggles in Subsection 2.3.

Proof. We prove that $2\chi_1 + \chi_2$ is 1-mesic, as the proof for $\chi_{n-1} + 2\chi_n$ is analogous.

The first two bits of any independent set S are either 10, 01, or 00.

If S begins with 10, then when applying φ to S , the first toggle τ_1 removes the first vertex so the first digit is 0. Then τ_2 can sometimes insert the second vertex and sometimes cannot, depending on whether $3 \in S$. Thus, $\varphi(S)$ begins either with 01 or 00.

If S begins with 01, then when applying φ to S , we leave the first vertex out and then remove the second vertex. So $\varphi(S)$ begins with 00.

If S begins with 00, then when applying φ to S , we insert the first vertex and then leave the second vertex out. So $\varphi(S)$ begins with 10.

Thus, when repeatedly applying φ , the first two digits are partitioned into cyclic patterns of $10 \rightarrow 01 \rightarrow 00$ or $10 \rightarrow 00$. (An orbit may contain both types of patterns.) As $2\chi_1 + \chi_2$ has average 1 across both types of patterns, it will across every orbit as well. ■

2.2 Proof of Propp's original conjecture

Our next goal is to prove Theorem 2.9 via a partitioning of the orbit board into “snakes”. We first note what happens in the special case where symmetry of independent sets under reversal makes the result obvious.

Definition 2.13. The **reverse** of a word is that word written in the reverse order. For example, the reverse of 101000010 is 010000101. Denote the reverse of an independent set S as S^{rev} . Writing S as a set,

$$S^{\text{rev}} = \{n + 1 - i \mid i \in S\}.$$

The following is clear because the inverse of φ is the function that composes toggling in the reverse order.

Corollary 2.14. For any S , $(S^{\text{rev}})^{\text{rev}} = S$ and $\varphi(S^{\text{rev}}) = (\varphi^{-1}(S))^{\text{rev}}$.

Definition 2.15. A **symmetrical** independent set is one that is its own reverse. For example, 010010 is symmetrical.

Definition 2.16. A φ -orbit \mathcal{O} is **reversible** if for some $S \in \mathcal{O}$, S^{rev} is also in \mathcal{O} .

For a reversible orbit \mathcal{O} , such as in Example 2.11, there exists one $S \in \mathcal{O}$ whose reverse is also in \mathcal{O} . Then by Proposition 2.14, every set in \mathcal{O} has its reverse in \mathcal{O} . So it is clear that $\chi_j - \chi_{n+1-j}$ has average zero across any reversible orbit. For $n \geq 10$, however, there are φ -orbits on \mathcal{I}_n that are not reversible, so it is surprising a priori that Theorem 2.9 holds in general.

Corollary 2.17. Any φ -orbit containing a symmetrical independent set is reversible. An orbit contains at most two symmetrical independent sets (but even a reversible orbit may not contain any).

Proof. For an orbit \mathcal{O} containing a symmetrical independent set S , $S^{\text{rev}} = S$ is in the orbit, so \mathcal{O} is reversible.

Now assume that an orbit \mathcal{O} contains at least two different symmetrical independent sets S and T . Then there exists $m \geq 1$ such that $\varphi^m(S) = T$ and let m be the least number that satisfies this. Then from Proposition 2.14, we have that

$$\varphi^{-m}(S) = \varphi^{-m}(S^{\text{rev}}) = (\varphi^m(S))^{\text{rev}} = T^{\text{rev}} = T.$$

Thus $S = \varphi^m(T)$, which implies $\varphi^{2m}(S) = S$. Therefore, \mathcal{O} has $2m$ sets, since m was chosen to be minimal. Let $U \neq S, T$ be another set in \mathcal{O} . Then $U = \varphi^k(S)$ for some $k \in [2m - 1]$ with $k \neq m$, and so

$$\varphi^{-k}(S) = \varphi^{-k}(S^{\text{rev}}) = (\varphi^k(S))^{\text{rev}} = U^{\text{rev}}.$$

Since \mathcal{O} has $2m$ sets, we cannot have $\varphi^{-k}(S) = \varphi^k(S)$, so $U \neq U^{\text{rev}}$. Thus, by definition U is not symmetrical, so S and T are the only symmetrical sets in \mathcal{O} . ■

Lemma 2.18. *Consider the action of φ .*

1. *When $S(i, j) = 1$ and $j \neq n$, either $S(i, j + 2) = 1$ or $S(i + 1, j + 1) = 1$, and never both.*
2. *When $S(i, j) = 1$ and $j \neq 1$, either $S(i, j - 2) = 1$ or $S(i - 1, j - 1) = 1$, and never both.*
3. *If $S(i, j) = 1$, then $S(i, j - 1) = S(i, j + 1) = S(i - 1, j) = S(i + 1, j) = 0$.*
4. *If $S(i, j) = 1$, then $S(i + 1, j - 1) = S(i - 1, j + 1) = 0$.*

Proof. (3) is clear because each S^i is an independent set, and if $S(i, j) = 1$, then $j \in S^i$, so $j \notin \varphi(S^i) = S^{i+1}$.

Now we prove (1). If $j \in S^i$, then $j \notin S^{i+1}$, so $j + 1 \in S^{i+1}$ if and only if $j + 2 \notin S^i$.

The proof of (2) is analogous to that of (1) because φ^{-1} applies the toggles in the reverse order.

To prove (4), assume $j \geq 2$ with $S(i, j) = 1$. Then $j \in S_i$. Then after applying $\tau_{j-1} \cdots \tau_1$ to S_i , we cannot have both $j - 1$ and j in the same independent set. So $S(i + 1, j - 1) = 0$. Thus if $S(i - 1, j + 1) = 1$, then $S(i, j) = 0$ which is a contradiction, so $S(i - 1, j + 1) = 0$. ■

From Lemma 2.18(1), given a 1 in the orbit board (outside of the rightmost column), there is another 1 either in the position two spaces to the right, or the position one space diagonally right and down. From Lemma 2.18(2), for any 1 in the orbit board (outside of the leftmost column), there is another 1 either in the position two spaces to the left, or the position one space diagonally left and up. Therefore, the ones in the orbit board can be partitioned into sequences, called **snakes**, that begin in the left column and end in the right column. For any 1 in the snake, the next 1 is located either two spaces to the right of it, or in the position one space diagonally right and down.

	1	2	3	4	5	6	7
S^0	1	0	1	0	1	0	0
S^1	0	0	0	0	0	1	0
S^2	1	0	1	0	0	0	1
S^3	0	0	0	1	0	0	0
S^4	1	0	0	0	1	0	1
S^5	0	1	0	0	0	0	0
S^6	0	0	1	0	1	0	1
S^7	1	0	0	0	0	0	0
S^8	0	1	0	1	0	1	0
S^9	0	0	0	0	0	0	1
Total	4	2	3	2	3	2	4

Figure 2: The orbit from Figure 1 with colors representing the snakes

Example 2.19. The orbit board from Figure 1, with colors representing the different snakes, is shown in Figure 2.

Therefore, to know where the ones in the orbit board are, it suffices to analyze the snakes. To each φ -orbit on \mathcal{I}_n , we will associate an equivalence class of compositions of $n - 1$ into parts 1 and 2, with each composition representing the snakes.

Definition 2.20. A **composition** of $n \in \mathbb{Z}^+$ is a sequence of positive integers whose sum is n . Two compositions of n are said to be **cyclically equivalent** if one is a cyclic rotation of the other. Otherwise, the compositions are **cyclically inequivalent**.

Example 2.21. 21121, 11212, 12121, 21211, and 12112 are cyclically equivalent compositions of 7.

To associate a composition of $n - 1$ to any given snake in a φ -orbit of \mathcal{I}_n , a step of two positions to the right corresponds to a 2, and a step of one position diagonally right and down corresponds to a 1. Thus, we get a composition of $n - 1$ because we start in the leftmost column and end in the rightmost column.

Definition 2.22. To associate a composition of $n - 1$ to any given snake in a φ -orbit of \mathcal{I}_n , a step of two positions to the right corresponds to a 2, and a step of one position diagonally right and down corresponds to a 1. Thus, we get a composition of $n - 1$ because we start in the leftmost column and end in the rightmost column. This is called the **snake composition** of the snake.

In Example 2.19, the red snake has snake composition 2211, the purple snake has snake composition 2112, the green snake has composition 1122, and the blue snake has composition 1221.

The following lemmata further constrain the possible pattern of ones in an orbit board. They will be used in the proof of Theorem 2.25.

Lemma 2.23. *Under the action of φ , suppose $S(i, j) = 1$ and $S(i + 2, j - 1) = 1$.*

1. *If $S(i, j + 2) = 1$, then $S(i + 2, j + 1) = 1$.*
2. *If $S(i + 1, j + 1) = 1$, then $S(i + 3, j) = 1$.*
3. *If $j = n$, then $S(i + 3, j) = S(i + 3, n) = 1$.*

Proof. Without loss of generality we may assume $i = 0$, as we can start our orbit board anywhere. This means that $j \in S^0$ and $j - 1 \in S^2$.

1. In this scenario, $j + 2 \in S^0$, and we wish to conclude that $j + 1 \in S^2$. By Lemma 2.18(1), $S(0, j) = S(0, j + 2) = 1$ gives $S(1, j + 1) = 0$. Thus $j + 1 \notin S^1$. Also, $j + 2 \in S^0$ implies $j + 2 \notin S^1$. And $j - 1 \in S^2$ implies $j \notin S^2$. Therefore, when applying toggles to S^1 , $j + 1$ gets toggled in, so $j + 1 \in S^2$.
2. In this scenario, $j + 1 \in S^1$, and we wish to conclude that $j \in S^3$. Since $j - 1 \in S^2$, we can use Lemma 2.18(1) to determine that either $j + 1 \in S^2$ or $j \in S^3$. However, $j + 1 \notin S^2$ because $j + 1 \in S^1$. Therefore, $j \in S^3$.
3. Since $S(2, n - 1) = 1$, either $S(2, n + 1) = 1$ or $S(3, n) = 1$ from Lemma 2.18(1). Only the second scenario is possible so $S(3, n) = 1$. ■

Lemma 2.24. *Under the action of φ , suppose $S(i, j) = 1$ and $S(i + 2, j - 2) = 1$.*

1. *If $S(i, j + 2) = 1$, then $S(i + 2, j) = 1$.*
2. *If $S(i + 1, j + 1) = 1$, then $S(i + 3, j - 1) = 1$.*
3. *If $j = n$, then $S(i + 2, j) = S(i + 2, n) = 1$.*

Proof. As in the proof of the previous lemma, assume $i = 0$ without loss of generality. This means that $j \in S^0$ and $j - 2 \in S^2$.

1. In this scenario, $j + 2 \in S^0$, and we wish to conclude that $j \in S^2$. Since $S(0, j) = S(0, j + 2) = 1$, we conclude from Lemma 2.18(1) that $S(1, j + 1) = 0$ and so $j + 1 \notin S^1$. Note that $j - 2 \in S^2$ gives $j - 1 \notin S^2$. Also $j \in S^0$ gives $j \notin S^1$. Thus, when we apply toggles left to right starting with S^1 , we will be able to add vertex j to the set. Thus, $j \in S^2$.
2. In this scenario, $j + 1 \in S^1$, and we wish to conclude that $j - 1 \in S^3$. Since $j + 1 \in S^1$, it follows that $j \notin S^2$ by Lemma 2.18(4). Since $S(2, j - 2) = 1$ and $S(2, j) = 0$, we have $S(3, j - 1) = 1$ by Lemma 2.18(2).
3. Since $n \in S$, we have $n \notin S^1$. Also, $n - 2 \in S^2$ implies $n - 1 \notin S^2$. Thus, when we reach the last vertex when applying φ to S^1 , we insert n . So $n \in S^2$. ■

	1	2	3	4	5	6	7	8	9	10
S^0	1		1		1					
S^1						1				
S^2							1		1	
S^3										1

Figure 3: The ones in the orbit board are an example snake. In Example 2.26, we describe how to generate an entire orbit from one snake.

Theorem 2.25. *In a φ -orbit board, consider a snake starting on the S^i line. Let c be the snake's composition. Consider the least $i' > i$ for which $S(i', 1) = 1$. (This is where the “next” snake begins.)*

1. *If c starts with 1, then $i' = i + 3$.*
2. *If c starts with 2, then $i' = i + 2$.*
3. *The composition for the snake starting on the $S^{i'}$ line is the left cyclic rotation of c .*

Proof. Without loss of generality, assume $i = 0$. Let c' be the composition for the snake that starts on line $S^{i'}$.

If c starts with 1, then $S(0, 1) = S(1, 2) = 1$. Then $1 \notin \tau_1(S^1)$ because $2 \in S^1$. Then applying τ_2 to S_1 removes 2 from the set. Therefore $1, 2 \notin S^2$, and thus we insert 1 when applying τ_1 to S^2 . So $S(3, 1) = 1$, which proves (1). The part of c after the initial 1 describes the sequence of moves for the original snake from $S(1, 2)$ to the rightmost column. Since $S(1, 2) = S(3, 1) = 1$, this same sequence of moves describes the snake that starts on line S^3 from the leftmost column up to column $n - 1$, by Lemma 2.23(1,2). Then the snake with composition c' must finish with a diagonal step, so c' ends with 1. Thus c' is formed from c by moving the initial 1 to the end. This proves (3) for the case where c begins with 1.

Otherwise c starts with 2, so $S(0, 1) = S(0, 3) = 1$. By Lemma 2.18(3), $S(1, 1) = 0$ and by Lemma 2.18(1), $S(1, 2) = 0$. Since $1, 2 \notin S^1$, we insert 1 when applying τ_1 to S^1 , so $S(2, 1) = 1$, which proves (2). The part of c after the initial 2 describes the sequence of moves for the original snake from $S(0, 3)$ to the rightmost column. Since $S(0, 3) = S(2, 1) = 1$, this same sequence of moves describes the snake that starts on line S^2 from the leftmost column up to column $n - 2$, by Lemma 2.24(1,2). Suppose the snake with composition c ends on line S^k , i.e. $S(k, n) = 1$. Then the snake starting on line S^2 contains $S(k + 2, n - 2)$, so by Lemma 2.24(3), $S(k + 2, n) = 1$. Thus c' is formed from c by moving the initial 2 to the end. This proves (3) for the case where c begins with 2. ■

Example 2.26. We show how knowing one snake determines an entire orbit. Suppose we are working in \mathcal{I}_{10} and we have a snake given by the composition 221121. Then we immediately have the part of the orbit board shown in Figure 3.

	1	2	3	4	5	6	7	8	9	10	
S^0	1	0	1	0	1	0	0	1	0	1	Red snake: 221121
S^1	0	0	0	0	0	1	0	0	0	0	
S^2	1	0	1	0	0	0	1	0	1	0	Purple snake: 211212
S^3	0	0	0	1	0	0	0	0	0	1	
S^4	1	0	0	0	1	0	1	0	0	0	Orange snake: 112122
S^5	0	1	0	0	0	0	0	1	0	1	
S^6	0	0	1	0	1	0	0	0	0	0	
S^7	1	0	0	0	0	1	0	1	0	1	Green snake: 121221
S^8	0	1	0	1	0	0	0	0	0	0	
S^9	0	0	0	0	1	0	1	0	1	0	
S^{10}	1	0	1	0	0	0	0	0	0	1	Blue snake: 212211
S^{11}	0	0	0	1	0	1	0	1	0	0	
S^{12}	1	0	0	0	0	0	0	0	1	0	Brown snake: 122112
S^{13}	0	1	0	1	0	1	0	0	0	1	
S^{14}	0	0	0	0	0	0	1	0	0	0	
Total	6	3	4	4	4	4	4	4	3	6	

Figure 4: The unique orbit containing the snake from Figure 3 is the orbit containing $S = 1010100101$ (See Example 2.26).

Using Theorem 2.25, we know that the next snake begins on the S^2 line, and has snake composition 211212. This snake is shown in purple in Figure 4. Also by Theorem 2.25, the next four snakes start on the lines have snake compositions 112122, 121221, 212211, and 122112 respectively and begin on lines S^4 , S^7 , S^{10} and S^{12} . These are shown in orange, green, blue, and brown respectively in Figure 4.

Then the next snake starts on the S^{15} line and has snake composition 221121. However, this is the snake we started with. Therefore $S^0 = S^{15}$. So this orbit has size 15, and the ones in the brown snake on the S^{15} line go on the S^0 line. Every other empty position is a 0 by Lemma 2.18(3). The full orbit board is shown in Figure 4. Notice that this orbit is not reversible, so there is no simple reason for the column-sum vector to be palindromic.

The following should be clear now.

Corollary 2.27. *For any φ -orbit, the set of snake compositions is invariant under cyclic rotation. Thus, there is a bijection between φ -orbits of \mathcal{I}_n and cyclically inequivalent compositions of $n - 1$ into parts 1 and 2. Also, an orbit \mathcal{O} is reversible if and only if for each snake in the orbit with snake composition c , there is also a snake in the orbit whose snake composition is c in the reverse order.*

We are now ready to prove Propp's original conjecture.

Proof of Theorem 2.9. We wish to prove that for any $j \in [n]$, $\chi_j - \chi_{n+1-j}$ is 0-mesic. It suffices to show that for any \mathcal{O} ,

$$\sum_{S \in \mathcal{O}} \chi_j(S) = \sum_{S \in \mathcal{O}} \chi_{n+1-j}(S).$$

Since every snake in \mathcal{O} starts in the leftmost column and ends in the rightmost column, the orbit has the same number of ones in the leftmost column as in the rightmost column of the (finite version of the) orbit board. Thus,

$$\sum_{S \in \mathcal{O}} \chi_1(S) = \sum_{S \in \mathcal{O}} \chi_n(S).$$

Each entry in the snake composition refers to how many columns we move right to get to the next 1 in the snake. Thus for $j > 1$, there is a 1 in column j of the orbit board for every snake composition that has an initial segment adding to $j - 1$. Similarly, there is a 1 in column $n + 1 - j$ of the orbit board for every snake composition that has a final segment adding to $j - 1$. By cyclic rotation of snake compositions, we get that there are the same number of snake compositions in \mathcal{O} with an initial segment that adds to $j - 1$ as there are with a final segment that adds to $j - 1$. ■

Example 2.28. For the orbit board in Example 2.26, there is a 1 in column 4 whenever an initial segment of a snake's composition adds to 3. There are two snake compositions associated with this orbit that begin with 12. They are 121221 (green) and 122112 (brown). For each of these, there is a cyclic rotation of the snake composition that ends with 12. These are 122112 (brown) and 211212 (purple). These give ones in column 7 (fourth column from the right).

Also there are two snake compositions associated with this orbit that begin with 21. They are 211212 (purple) and 212211 (blue). For each of these, there is a cyclic rotation of the snake composition that ends with 21. These are 121221 (green) and 221121 (red). These give ones in column 7 (fourth column from the right).

If there were snake compositions for this orbit that began with 111 (the other way to have an initial segment adding to 3), then by cyclic rotation, there would be just as many that end in 111.

2.3 Coxeter groups and extending Propp's conjecture

Theorems 2.9 and 2.12 are for orbits of the specific action $\varphi = \tau_n \cdots \tau_2 \tau_1$. In this subsection, we describe how these homomesy results can be generalized to some other actions in \mathcal{T}_n , namely *Coxeter elements*.

Definition 2.29. An element $w \in \mathcal{T}_n$ is called a **Coxeter element** if it is a product of $\tau_1, \tau_2, \dots, \tau_n$, each used exactly once, in some order.

Theorem 2.30. Let $w, w' \in \mathcal{T}_n$ be two Coxeter elements. Any statistic which is a linear combination of the indicator functions χ_j is c -mesic under the action of w if and only if it is c -mesic under the action of w' .

Since φ is a Coxeter element in \mathcal{T}_n , we use Theorem 2.30 to generalize Theorems 2.9 and 2.12 to the action of *any* Coxeter element. This is the following corollary.

Corollary 2.31. *Let $w \in \mathcal{T}_n$ be a Coxeter element, and consider the action of w on \mathcal{I}_n .*

1. *The statistic $\chi_j - \chi_{n+1-j}$ is 0-mesic for every $1 \leq j \leq n$.*
2. *The statistics $2\chi_1 + \chi_2$ and $\chi_{n-1} + 2\chi_n$ are both 1-mesic.*

The rest of this section leads up to the proof of Theorem 2.30, which is near the end.

As \mathcal{T}_n is generated by finitely many involutions, it is the quotient of a Coxeter group. See [BB05] and [Pet15, Ch. 11–14] to learn about Coxeter groups and their connections to combinatorics. The Coxeter group with presentation


$$\langle \tau_1, \tau_2, \dots, \tau_n \mid \tau_i^2 = 1, (\tau_i \tau_{i+1})^6 = 1, (\tau_i \tau_j)^2 = 1 \text{ for } j \geq i + 2 \rangle,$$

for $n \geq 3$, is well known to be infinite. However, the toggle group \mathcal{T}_n is a subgroup of $\mathfrak{S}_{\mathcal{I}_n}$ and thus finite. Therefore \mathcal{T}_n has extra relations in addition to the Coxeter relations $(\tau_i^2 = 1, (\tau_i \tau_i + 1)^6 = 1, (\tau_i \tau_j)^2 = 1 \text{ for } j \geq i + 2)$, so it is not a true Coxeter group. Knowledge of Coxeter groups is not necessary to understand the rest of the paper, as we explain the theory necessary to describe why we can extend the homomesy results for φ to orbits of any Coxeter element $w \in \mathcal{T}_n$. Even though \mathcal{T}_n is not a true Coxeter group, we borrow the term “Coxeter element” from Coxeter group theory.

The following is a brief summary of material in [EFGJMPR, §3,6], which describes in detail how Coxeter group theory can be applied to toggle groups, and the connections between conjugation of elements and homomesy. We discuss it in the context of \mathcal{T}_n .

To each Coxeter element $w \in \mathcal{T}_n$, we define an orientation of \mathcal{P}_n . That is, we direct each edge of the path graph in such a way that corresponds to our Coxeter element. To do this, we direct the edge connecting i and $i + 1$ in the direction of i if τ_i appears to the right of τ_{i+1} in an expression of w , and in the direction of $i + 1$ if τ_{i+1} appears to the right of τ_i in an expression of w .

Example 2.32. The Coxeter element $w = \tau_3 \tau_4 \tau_2 \tau_6 \tau_7 \tau_5 \tau_1 \in \mathcal{T}_7$ corresponds to the following orientation of \mathcal{P}_7 .



While w has other representations as a product of toggles, each used exactly once, they are all found by taking the expression $\tau_3 \tau_4 \tau_2 \tau_6 \tau_7 \tau_5 \tau_1$ and applying the commutativity relations $\tau_i \tau_j = \tau_j \tau_i$ if $|i - j| \geq 2$. For example, $w = \tau_3 \tau_2 \tau_6 \tau_4 \tau_5 \tau_7 \tau_1$, found by moving τ_4 to the right of $\tau_2 \tau_6$ and swapping the order of τ_5 and τ_7 . Since applying the commutativity relations does not change the relative order of any pair τ_i, τ_{i+1} , this does not change the orientation of any edge. Thus, the orientation of \mathcal{P}_n corresponds uniquely to the Coxeter element, and this orientation is well-defined.

Shi showed that Coxeter elements correspond uniquely to acyclic orientations of the graph associated with a Coxeter group, where the vertices of the graph are the involution generators, and the edges of the graph are between pairs of noncommuting generators [Shi97, Proposition 1.3]. Since \mathcal{P}_n already has no cycles, any orientation of it corresponds to a unique Coxeter element in \mathcal{T}_n , even though \mathcal{T}_n is only the quotient of a Coxeter group. Toggle maps corresponding to source vertices are called **initial** in w because they can be brought to the left by the commutativity relations. Toggle maps corresponding to sink vertices are called **final** in w because they can be brought to the right by the commutativity relations. Vertices that are neither sources nor sinks correspond to toggles that can neither be brought to the left nor right in an expression of w (while using every toggle exactly once).

Since the toggles are involutions, conjugating by an final (resp. initial) toggle τ_i of w corresponds to moving it to the left (resp. right) of an expression for w . For example, τ_5 is final in $w = \tau_3\tau_2\tau_6\tau_4\tau_5\tau_7\tau_1$, so conjugating by τ_5 gives

$$\begin{aligned}\tau_5 w \tau_5 &= \tau_5(\tau_3\tau_2\tau_6\tau_4\tau_5\tau_7\tau_1)\tau_5 \\ &= \tau_5(\tau_3\tau_2\tau_6\tau_4\tau_7\tau_1\tau_5)\tau_5 \\ &= \tau_5\tau_3\tau_2\tau_6\tau_4\tau_7\tau_1.\end{aligned}$$

This conjugation by a final element of w corresponds to changing a sink into a source in the corresponding orientation of \mathcal{P}_n . Conjugation by an initial element would correspond to changing a source into a sink. If we were to conjugate w by τ_4 , which is neither initial nor final, we would get $\tau_4\tau_3\tau_2\tau_6\tau_4\tau_5\tau_7\tau_1\tau_4$ which is not a Coxeter element.

H. and K. Eriksson showed that two Coxeter elements are conjugate if and only if we can transform the orientation for one of them into the other by a sequence that changes sinks into sources or vice versa [EE09]. This corresponds to a sequence of conjugations by generators that are either initial or final; we call such conjugations **admissible**.

In the case of \mathcal{T}_n , any two Coxeter elements are conjugate. This is because \mathcal{T}_n has generators $\tau_1, \tau_2, \dots, \tau_n$ satisfying $\tau_i^2 = 1$ and $(\tau_i\tau_j)^2 = 1$ when $|i - j| > 1$. Thus, any Coxeter elements in \mathcal{T}_n are conjugate [SW12, Lemma 5.1]. Not only that, but the proof given by Striker and Williams shows that we can get from any Coxeter element to any other one by a sequence of *admissible* conjugations. We describe their method for conjugating by toggles to transform any Coxeter element w into $\varphi = \tau_n \cdots \tau_2\tau_1$. The method is slightly modified here because we wish to transform w into $\tau_n \cdots \tau_2\tau_1$ not $\tau_1\tau_2 \cdots \tau_n$. Starting with w , find the largest number k such that τ_k is final in w , then push τ_k to the right and conjugate by τ_k . Repeat this until we arrive at $\varphi = \tau_n \cdots \tau_2\tau_1$.

Example 2.33. Let $w = \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1$ as in Example 2.32. Refer to Figure 5 for a sequence of conjugations to transform w into φ and the corresponding orientations of \mathcal{P}_n at each step. By the process described above, $\varphi = u^{-1}wu$ where $u = \tau_7\tau_5\tau_6\tau_7\tau_4\tau_5\tau_6\tau_7$.

We are now ready for the proof of Theorem 2.30.

Proof of Theorem 2.30. As w and w' differ by a sequence of admissible conjugations, it suffices to consider the case where w' differs from w by a single conjugation. Let $w \in \mathcal{T}_n$ be a Coxeter element, τ_k initial or final in w , and $w' = \tau_k w \tau_k$.

$$\begin{array}{lcl}
& \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \rightarrow \bullet \leftarrow \bullet \rightarrow \bullet \\
= & \tau_3\tau_4\tau_2\tau_6\tau_5\tau_1\tau_7 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_7 & \\
& \tau_7\tau_3\tau_4\tau_2\tau_6\tau_5\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \rightarrow \bullet \leftarrow \bullet \leftarrow \bullet \\
= & \tau_7\tau_3\tau_4\tau_2\tau_6\tau_1\tau_5 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_5 & \\
& \tau_5\tau_7\tau_3\tau_4\tau_2\tau_6\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \\
= & \tau_5\tau_7\tau_3\tau_4\tau_2\tau_1\tau_6 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_6 & \\
& \tau_6\tau_5\tau_7\tau_3\tau_4\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \\
= & \tau_6\tau_5\tau_3\tau_4\tau_2\tau_1\tau_7 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_7 & \\
& \tau_7\tau_6\tau_5\tau_3\tau_4\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \\
= & \tau_7\tau_6\tau_5\tau_3\tau_2\tau_1\tau_4 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_4 & \\
& \tau_4\tau_7\tau_6\tau_5\tau_3\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \leftarrow \bullet \\
= & \tau_4\tau_7\tau_6\tau_3\tau_2\tau_1\tau_5 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_5 & \\
& \tau_5\tau_4\tau_7\tau_6\tau_3\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \leftarrow \bullet \\
= & \tau_5\tau_4\tau_7\tau_3\tau_2\tau_1\tau_6 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_6 & \\
& \tau_6\tau_5\tau_4\tau_7\tau_3\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \rightarrow \bullet \\
= & \tau_6\tau_5\tau_4\tau_3\tau_2\tau_1\tau_7 & 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \\
& \text{conj} \downarrow \tau_7 & \\
& \tau_7\tau_6\tau_5\tau_4\tau_3\tau_2\tau_1 & \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet \leftarrow \bullet
\end{array}$$

Figure 5: A demonstration of how to write $\tau_7\tau_6\tau_5\tau_4\tau_3\tau_2\tau_1$ as a conjugation of $\tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1$, with the corresponding orientations of \mathcal{P}_7 at every step.

We have the following commutative diagram.

$$\begin{array}{ccc} S & \xrightarrow{w} & w(S) \\ \tau_k \downarrow & & \downarrow \tau_k \\ \tau_k(S) & \xrightarrow{w'} & \tau_k(w(S)) \end{array}$$

Write a w -orbit $\mathcal{O} = (S^0, S^1, \dots, S^{m-1})$, where as with the orbit board, $S^i = w(S^{i-1})$. Consider the superscripts to be mod m , the size of the orbit, so $S^0 = S^m = w(S^{m-1})$. Then, via the commutative diagram above, $\mathcal{O}' = (\tau_k(S^0), \tau_k(S^1), \dots, \tau_k(S^{m-1}))$ is a w' -orbit. This creates a bijection between w -orbits and w' -orbits that preserves the orbit sizes.

When $j \neq k$, $\chi_j(S^i) = \chi_j(\tau_k(S^i))$ as τ_k can only change the value of χ_k , so

$$\sum_{S \in \mathcal{O}} \chi_j(S) = \sum_{S' \in \mathcal{O}'} \chi_j(S').$$

If τ_k is final in w , then it is the first toggle applied when performing w . Also τ_k only appears once when applying w , since w is a Coxeter element. Thus, in this scenario $\chi_k(\tau_k(S^i)) = \chi_k(S^{i+1})$. So

$$\sum_{S' \in \mathcal{O}'} \chi_k(S') = \chi_k(S^1) + \chi_k(S^2) + \dots + \chi_k(S^{m-1}) + \chi_k(S^0) = \sum_{S \in \mathcal{O}} \chi_k(S).$$

Analogously, if τ_k is initial in w , then $\chi_k(\tau_k(S^j)) = \chi_k(S^{j-1})$. Thus,

$$\sum_{S' \in \mathcal{O}'} \chi_k(S') = \chi_k(S^{m-1}) + \chi_k(S^0) + \chi_k(S^1) + \dots + \chi_k(S^{m-2}) = \sum_{S \in \mathcal{O}} \chi_k(S).$$

Therefore, the sum of any χ_j is the same in \mathcal{O} and \mathcal{O}' . As these orbits have the same size, the average of any χ_j is also the same across them. See Example 2.35. ■

Remark 2.34. In general, conjugation does not preserve homomesy of statistics under the respective maps. Theorem 2.30 describes only a specific scenario in which homomesy is preserved, namely when the conjugation is admissible (as is the case for any two Coxeter elements) and the statistic is a linear combination of the indicator functions χ_j . If we were to find other types of homomesic statistics for orbits under one Coxeter element, they would not necessarily remain homomesic under a different Coxeter element.

Example 2.35. Figure 6 contains an orbit under the action $w = \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1$ starting with 1010010 and an orbit under the action of φ starting with 1010100. Notice that in the sequence of admissible conjugations to go from w to φ , we conjugate by τ_4 once, τ_5 and τ_6 each twice, and τ_7 thrice, and each conjugation is by a final toggle, as described in Example 2.33. As in the proof of Theorem 2.30, in the w -orbit board, if we slide columns 4, 5, 6, and 7 up by 1 row, 2 rows, 2 rows, and 3 rows respectively, we get the φ -orbit board.

	1	2	3	4	5	6	7
S	1	0	1	0	0	1	0
$w(S)$	0	0	0	0	0	0	0
$w^2(S)$	1	0	1	0	1	0	1
$w^3(S)$	0	0	0	0	0	1	0
$w^4(S)$	1	0	0	1	0	0	0
$w^5(S)$	0	1	0	0	0	0	1
$w^6(S)$	0	0	1	0	1	0	0
$w^7(S)$	1	0	0	0	0	0	1
$w^8(S)$	0	1	0	0	1	0	0
$w^9(S)$	0	0	0	1	0	0	1

	1	2	3	4	5	6	7
S'	1	0	1	0	1	0	0
$\varphi(S')$	0	0	0	0	0	1	0
$\varphi^2(S')$	1	0	1	0	0	0	1
$\varphi^3(S')$	0	0	0	1	0	0	0
$\varphi^4(S')$	1	0	0	0	1	0	1
$\varphi^5(S')$	0	1	0	0	0	0	0
$\varphi^6(S')$	0	0	1	0	1	0	1
$\varphi^7(S')$	1	0	0	0	0	0	0
$\varphi^8(S')$	0	1	0	1	0	1	0
$\varphi^9(S')$	0	0	0	0	0	0	1

Figure 6: Left: the orbit under the action of $w = \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1$ containing $S = 1010010$. Right: the orbit under the action of φ containing $S' = 1010100$. See Example 2.35.

Remark 2.36. Notice that while conjugation in the toggle group preserves the corresponding orbit structures (total number of orbits and multiset of orbit sizes) and that a sequence of admissible conjugations preserves the homomesic property of any statistic that is a linear combination of the χ_j statistics, many other propositions we have made along the way do not hold for generic Coxeter elements. In particular, any statement about which independent sets are in a given orbit does not extend to general Coxeter elements. In the orbit on the left in Figure 6, notice that parts 1, 2, and 4 of Lemma 2.18 are violated, though part 3 holds for orbits under any Coxeter element by the same proof. Also notice that the orbit contains four symmetrical independent sets, but contains independent sets whose reverses are not also in the orbit. This shows that Proposition 2.17 does not hold for arbitrary Coxeter elements. In fact, when n is odd, it can be shown that any given orbit under the map $\tau_2\tau_4 \cdots \tau_{n-1}\tau_1\tau_3\tau_5 \cdots \tau_n$ always either consists entirely of symmetrical independent sets or contains no symmetrical independent sets. This is because $\tau_1\tau_3\tau_5 \cdots \tau_n$ consists of entirely commuting toggles, as does $\tau_2\tau_4 \cdots \tau_{n-1}$. So by symmetry τ_1 acts as τ_n , and τ_3 acts as τ_{n-2} , and so on.

3 Enumerating Independent Sets and φ -Orbits

In this section we present enumerative formulas for the numbers of φ -orbits and reversible φ -orbits of \mathcal{I}_n . Numerical data and the Online Encyclopedia of Integer Sequences led us to a conjectured formula for the number of φ -orbits and connected them with binary necklaces and bracelets. This also helped point us towards the snake-partition of orbit boards used in the proof of Propp's original conjecture in Section 2. Our main tools are Burnside's Lemma and some bijections.

The first proposition is well-known and easy to see directly from the defining recursion.

Proposition 3.1. *The number of independent sets in \mathcal{I}_n is F_{n+2} , the $(n+2)^{\text{nd}}$ Fibonacci number, with $F_0 = 0$, $F_1 = 1$ and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$.*

Proposition 3.2. *Let $k \in \mathbb{Z}^+$.*

- (a) *The number of symmetrical independent sets in \mathcal{I}_{2k} is F_{k+1} .*
- (b) *The number of symmetrical independent sets in \mathcal{I}_{2k-1} is F_{k+2} .*

Proof. (a) By symmetry, $k \in S$ if and only if $k+1 \in S$. This means neither k nor $k+1$ can be in S , since S is independent. Therefore, the first $k-1$ vertices can be any independent set, and the last $k-1$ vertices is its reverse. Thus, the symmetrical sets in \mathcal{I}_{2k} are in bijection with \mathcal{I}_{k-1} , so the number of them is F_{k+1} .

(b) If the middle vertex k is not in the set, then the first $k-1$ vertices can be any independent set, and the last $k-1$ vertices is its reverse. Therefore, there are $\#\mathcal{I}_{k-1} = F_{k+1}$ symmetrical independent sets of this type. If vertex k is in the set, then neither $k-1$ nor $k+1$ can be. Then the first $k-2$ vertices can be any independent set, and the last $k-2$ vertices is its reverse. There are $\#\mathcal{I}_{k-2} = F_k$ symmetrical independent sets of this type. Thus, the total number of symmetrical sets in \mathcal{I}_{2k-1} is $F_k + F_{k+1} = F_{k+2}$. ■

To count the number of φ -orbits, there is a connection with binary necklaces, which we now introduce.

Definition 3.3. As with compositions, two binary strings of length n are said to be **cyclically equivalent** if one is a cyclic rotation of the other. Otherwise, the strings are **cyclically inequivalent**. A **binary necklace** is an equivalence class of binary strings under cyclic equivalence (i.e., under the action of the cyclic group C_n). A **binary bracelet** of length n is an equivalence class of length n binary strings under the action of the dihedral group D_n generated by cyclic rotation and reversal. The **length** of a binary necklace or bracelet is the length of the any string in the equivalence class.

Example 3.4. There are six binary necklaces of length 4:

$\{0000\}$,
 $\{1000, 0100, 0010, 0001\}$,
 $\{1100, 1001, 0011, 0110\}$,
 $\{1010, 0101\}$,
 $\{1110, 1101, 1011, 0111\}$, and
 $\{1111\}$,

all of which are also binary bracelets. For $n \leq 5$, we get the same equivalence classes under C_n and D_n . For $n = 6$, there are 14 binary necklaces, but only 13 binary bracelets, with the (distinct) C_6 -classes of 101100 and of 001101, which are reversals of one another, combining into a single class under the D_6 -action [Slo16, Seq. A000029, A000031].

Now we use Burnside's Lemma [Sta99, Lemma 7.24.5], to count length n binary necklaces and bracelets with no subsequence 11. For this we first need to count binary strings that will contain no subsequence 11 even when the group action makes the first and last elements adjacent.

Proof. This formula can be easily confirmed for $n = 1$ and $n = 2$, so assume $n \geq 3$. There are two types of binary strings s :

Case 2: If s begins with 0, then it has the form 0____, where the blank space represents an independent set of \mathcal{P}_{n-1} . So there are F_{n+1} strings of this form. ■

$$\#(Y/G) = \frac{1}{\#G} \sum_{w \in G} \# \text{Fix}(w).$$
$$\frac{1}{n} \sum_{d|n} \phi(n/d) (F_{d-1} + F_{d+1}), \quad (1)$$

Proof. In the context of Burnside's Lemma, G is the cyclic group C_n of order n and Y is the set of binary strings of length n with no subsequence 11 that also do not both start and end with 11. For any $d|n$, the number of elements of Y fixed by a group element of order n/d in C_n is $F_{d-1} + F_{d+1}$ by Lemma 3.5. By elementary group theory, there are $\phi(n/d)$ elements of order n/d , and the result follows. \blacksquare

$$\frac{1}{2} \left(F_{\lfloor n/2 \rfloor + 2} + \frac{1}{n} \sum_{d|n} \phi(n/d) (F_{d-1} + F_{d+1}) \right). \quad (2)$$

Proof. The cases of $n \leq 4$ can be computed case by case, so we assume $n \geq 5$. As above, Y is the set of binary strings of length n with no subsequence 11 that also do not both start and end with 1, but now $G = D_n$. By Burnside's Lemma, the number of binary bracelets with no subsequence 11 is

$$\begin{aligned} \frac{1}{2n} \sum_{w \in D_n} \# \text{Fix}(w) &= \frac{1}{2n} \left(\sum_{w \text{ reflection in } D_n} \# \text{Fix}(w) + \sum_{w \text{ rotation in } D_n} \# \text{Fix}(w) \right) \\ &= \frac{1}{2n} \left(\sum_{w \text{ reflection in } D_n} \# \text{Fix}(w) + \sum_{d|n} \phi(n/d) (F_{d-1} + F_{d+1}) \right). \end{aligned}$$

The question now is what is fixed by a reflection? We split this into two cases.

Case 1: n odd. Here each reflection leaves exactly one point fixed. Without loss of generality, assume that we are working with the reflection that reverses a string, which fixes exactly the palindromic strings. Since our binary strings cannot both start and end with 1, any palindromic string both begins and ends with 0.

Similar to the proof of Lemma 3.5, it is easy to show that there are $F_{(n+1)/2}$ palindromic strings with 0 in the center, and $F_{(n-1)/2}$ with 1 in the center. Thus, there are $F_{(n+1)/2} + F_{(n-1)/2} = F_{(n+3)/2}$ strings fixed by a reflection.

Thus, in the odd n case, the number of binary bracelets with no subsequence 11 is

$$\frac{1}{2n} \left(nF_{(n+3)/2} + \sum_{d|n} \phi(n/d)(F_{d-1} + F_{d+1}) \right) = \frac{1}{2} \left(F_{(n+3)/2} + \frac{1}{n} \sum_{d|n} \phi(n/d)(F_{d-1} + F_{d+1}) \right).$$

Case 2: n even. Here there are two types of reflections of a regular n -gon, those that fix no vertices and those that fix two vertices.

For the $n/2$ reflections that fix no vertices, we assume without loss of generality that we are working with the reflection that reverses a string, which fixes exactly the palindromic strings. It is straightforward to show there are $F_{n/2}$ such strings.

For the $n/2$ reflections that fix two vertices, we assume without loss of generality that we are working with the reflection that fixes the digits in positions 1 and $n/2 + 1$. There are four possibilities, because the digits in positions 1 and $n/2 + 1$ can each be 0 or 1. It is easy to show that there are $F_{n/2+1}$ such strings with 0 in both fixed positions, $F_{n/2-1}$ with 1 in both fixed positions, and $2F_{n/2}$ with different values in the fixed positions.

Adding together the four possibilities, there are

$$\begin{aligned} F_{n/2-1} + F_{n/2} + F_{n/2} + F_{n/2+1} &= F_{n/2+1} + F_{n/2} + F_{n/2+1} \\ &= F_{n/2+1} + F_{n/2+2} \end{aligned}$$

such strings.

Thus, in the even n case, the number of binary bracelets with no subsequence 11 is

$$\begin{aligned} &\frac{1}{2n} \left(\frac{n}{2} (F_{n/2} + F_{n/2+1} + F_{n/2+2}) + \sum_{d|n} \phi(n/d)(F_{d-1} + F_{d+1}) \right) \\ &= \frac{1}{2} \left(F_{n/2+2} + \frac{1}{n} \sum_{d|n} \phi(n/d)(F_{d-1} + F_{d+1}) \right). \end{aligned}$$

In both the odd and even cases, the formula holds. ■

Corollary 3.9. *The number of length n binary necklaces with no subsequence 11 which are the same under reversal is $F_{\lfloor n/2 \rfloor + 2}$.*

Proof. Let S be this number, and D the number of such necklaces which are not equal to their reversals. Then $S + D = (1)$ and $S + \frac{1}{2}D = (2)$ (equations above). Eliminating D gives the result. ■

Example 3.10. There are 10 binary necklaces of length 9 without the subsequence 11. One representative from each of them is as follows: 000000000, 100000000, 101000000, 100100000, 100010000, 101010000, 100100100, 101010100, 101001000, 000100101. For each of the last two, the reverse is not a cyclic rotation of itself, and thus not the same necklace. These are the necklaces that do not count in Corollary 3.9. There are $8 = F_6$ other necklaces.

Theorem 3.11. *There is a bijection between the set of binary necklaces of length n with no subsequence 11, and cyclically inequivalent compositions of n with each part equal to 1 or 2. Therefore, by Proposition 3.7, the number of cyclically inequivalent compositions of n with each part equal to 1 or 2 is $\frac{1}{n} \sum_{d|n} \varphi(n/d)(F_{d-1} + F_{d+1})$ [Slo16, Seq. A000358].*

Proof. Take a binary necklace of length n with no subsequence 11, and write it in the form of a string s so the first character is 0, which exists because there is no subsequence 11. Then create a composition of n into parts 1 and 2, by replacing each occurrence of 01 in s with 2, and each occurrence of 0 in s not followed by 1 with a 1. For example

$$010010001 \longleftrightarrow 212112.$$

To show that this bijection is well-defined, if we cyclically rotate the composition to the left, we get a cyclically equivalent composition, but this corresponds to rotating the string s to the left once or twice depending on if the first part of the composition is 1 or 2. Therefore, the new string is part of the same necklace. An example of this is shown below. The strings in the necklace that begin with 1 have no corresponding composition.

$$\begin{array}{ll} 010010001 & \longleftrightarrow 212112 \\ 100100010 & \longleftrightarrow \\ 001000101 & \longleftrightarrow 121122 \\ 010001010 & \longleftrightarrow 211221 \\ 100010100 & \longleftrightarrow \\ 000101001 & \longleftrightarrow 112212 \\ 001010010 & \longleftrightarrow 122121 \\ 010100100 & \longleftrightarrow 221211 \\ 101001000 & \longleftrightarrow \end{array}$$

■

Corollary 3.12. *The number of cyclically inequivalent compositions c of n with each part equal to 1 or 2, for which the reverse order of c is cyclically equivalent to c is $F_{\lfloor n/2 \rfloor + 2}$.*

Proof. This follows easily from Corollary 3.9 and the proof of Theorem 3.11. ■

Theorem 3.13. *The total number of φ -orbits of \mathcal{I}_n is $\frac{1}{n-1} \sum_{d|(n-1)} \phi((n-1)/d)(F_{d-1} + F_{d+1})$.*

The total number of reversible orbits is $F_{\lfloor n/2 \rfloor + 1}$.

Proof. Combine the bijection of Proposition 2.27 with Theorem 3.11 for orbits and with Corollary 3.12, for reversible orbits. Note that $\lfloor (n-1)/2 \rfloor + 2 = \lfloor n/2 \rfloor + 1$. ■

A more user-friendly version of the latter formula is:

	1	2	3	4	5	6	7
S^0	1	0	1	0	0	0	0
S^1	0	0	0	1	0	1	0
S^2	1	0	0	0	0	0	1
S^3	0	1	0	1	0	0	0
S^4	0	0	0	0	1	0	1
S^5	1	0	1	0	0	0	0
S^6	0	0	0	1	0	1	0
S^7	1	0	0	0	0	0	1
S^8	0	1	0	1	0	0	0
S^9	0	0	0	0	1	0	1

Figure 7: The φ -orbit with snake composition 2121 has size 5 not 10 because of the periodicity of the composition.

- The number of reversible orbits in \mathcal{I}_{2k-1} is F_{k+1} .
- The number of reversible orbits in \mathcal{I}_{2k} is F_{k+1} .

4 Sizes of Orbits

In the initial data gathering, the authors observed several mysterious patterns in the sizes of φ -orbits. For example, almost all of the orbits had size congruent to $1 - n \pmod 4$ and certain orbit sizes like 4 and 6 never appeared at all (for any n). These mysteries were also cleared up via the snake compositions of Section 2, which leads to a simple characterization (Proposition 4.3) of the orbit size corresponding to a given composition. From this we derive nontrivial consequences for the existence of orbits of a fixed sizes as n varies, summarized for small sizes in Figure 8.

The basic idea uses Theorem 2.25: whenever the snake composition c contains a 2, the next snake starts two positions down, and when a snake composition contains a 1, the next snake starts three positions down. So the size of the orbit should be $|c'|$, where c' is obtained from c by replacing each 1 with a 3. For example, if $c = 221121$, then $c' = 223323$ is a composition of 15, which is the size of the orbit in Example 2.26.

This naive approach can fail, however, in the case where periodicity of c leads one to create a “super-orbit” rather than an orbit, e.g., the orbit of \mathcal{I}_7 in Figure 7 given by snake composition $c = 2121$. Here $|c'| = 2 + 3 + 2 + 3 = 10$, but $S^5 = S^0$ so the orbit actually has size 5, and the board repeats itself. Therefore, given a snake composition such as 2121 made up entirely of a repeated segment (here 21), we must divide by the number of times the smallest repeated segment repeats itself (here 2) to get the correct orbit size.

Definition 4.1. Call a composition c **periodic** if it consists entirely of adjacent copies of

the same repeated substring. Given a composition c , let $\psi(c)$ be the number of times the smallest repeated segment repeats itself to make up c . If $\psi(c) = 1$, we call c **aperiodic**.

Example 4.2. The composition 21221 is aperiodic, but $\psi(22122212) = 2$ and $\psi(222) = 3$.

Corollary 4.3. Given a φ -orbit \mathcal{O} containing snake composition c , let $N_1(c)$ be the number of occurrences of 1 in c and $N_2(c)$ be the number of occurrences of 2 in c . Then the size of the orbit \mathcal{O} is $\frac{3N_1(c)+2N_2(c)}{\psi(c)}$.

Therefore, given any orbit size, we can characterize exactly for which n there is an orbit of \mathcal{I}_n with that size, and how many such orbits.

Corollary 4.4. There is an orbit of \mathcal{I}_n of size 2 exactly when n is odd. In this case, the orbit is unique.

Proof. The only composition of 2 into parts 2 and 3 is the composition 2. By Proposition 2.27, an orbit has size 2 if and only if a snake composition corresponding to the orbit is of the form $222 \cdots 2$, with k 2s repeated in \mathcal{I}_{2k+1} . ■

Example 4.5. For n odd, the two independent sets in a φ -orbit of \mathcal{I}_n of size 2 are the empty set and the set $\{1, 3, 5, \dots, n\}$, as shown in the following orbit board in \mathcal{I}_9 .

	1	2	3	4	5	6	7	8	9
S^0	0	0	0	0	0	0	0	0	0
S^1	1	0	1	0	1	0	1	0	1

Corollary 4.6. For any $n \geq 2$, there is a unique orbit of \mathcal{I}_n of size 3.

Proof. The only composition of 3 into parts 2 and 3 is the composition 3. Therefore, a φ -orbit has size 3 if and only if the corresponding snake composition has the form $111 \cdots 1$, with k ones repeated in \mathcal{I}_{k+1} . ■

Example 4.7. The three independent sets S^0, S^1, S^2 in the orbit of size 3 are the sets of all elements of $[n]$ congruent to 0, 1, 2 mod 3, respectively. An example of the orbit board for this orbit of \mathcal{I}_7 is shown below.

	1	2	3	4	5	6	7
S^0	0	0	1	0	0	1	0
S^1	1	0	0	1	0	0	1
S^2	0	1	0	0	1	0	0

Corollary 4.8. For every n , there are no orbits of \mathcal{I}_n of size 4 or 6.

Proof. The only composition of 4 into parts 2 and 3 is the composition $2 + 2$. However, this composition is periodic and therefore gives an orbit of size 2. Similarly, the only compositions of 6 into parts 2 and 3 are $2 + 2 + 2$ and $3 + 3$, both periodic. These give the orbits of size 2 and 3 respectively. ■

Orbit size m	Aperiodic cyclically inequivalent compositions of m into parts 2 or 3	Corresponding snake composition type	$n \geq 2$ for which \mathcal{I}_n has an orbit of size m
2	2	$222 \cdots 2$	$n \equiv 1 \pmod{2}$
3	3	$111 \cdots 1$	all n
4	none	none	none
5	$3 + 2$	$1212 \cdots 12$	$n \equiv 1 \pmod{3}$
6	none	none	none
7	$3 + 2 + 2$	$122122 \cdots 122$	$n \equiv 1 \pmod{5}$
8	$3 + 3 + 2$	$112112 \cdots 112$	$n \equiv 1 \pmod{4}$
9	$3 + 2 + 2 + 2$	$12221222 \cdots 1222$	$n \equiv 1 \pmod{7}$
10	$3 + 3 + 2 + 2$	$11221122 \cdots 1122$	$n \equiv 1 \pmod{6}$
11	$3 + 3 + 3 + 2$	$11121112 \cdots 1112$	$n \equiv 1 \pmod{5}$
11	$3 + 2 + 2 + 2 + 2$	$1222212222 \cdots 12222$	$n \equiv 1 \pmod{9}$
12	$3 + 2 + 3 + 2 + 2$	$1212212122 \cdots 12122$	$n \equiv 1 \pmod{8}$
12	$3 + 3 + 2 + 2 + 2$	$1122211222 \cdots 11222$	$n \equiv 1 \pmod{8}$

Figure 8: Classification of path graphs which give certain φ -orbit sizes

Corollary 4.9. *For $n \geq 2$, there is an orbit of \mathcal{I}_n of size 5 exactly when $n \equiv 1 \pmod{3}$. In this case, the orbit is unique.*

Proof. The only compositions of 5 into parts 2 and 3 are $2 + 3$ and $3 + 2$, which are cyclically equivalent. Therefore, an orbit has size 5 if and only if a snake composition corresponding to the orbit is of the form $1212 \cdots 12$, with k total patterns of 12 repeated (thus a composition of $3k$). This snake composition is in an orbit of \mathcal{I}_{3k+1} . ■

Example 4.10. The orbit board for the φ -orbit of \mathcal{I}_7 of size 5 is below.

	1	2	3	4	5	6	7
S^0	1	0	1	0	0	0	0
S^1	0	0	0	1	0	1	0
S^2	1	0	0	0	0	0	1
S^3	0	1	0	1	0	0	0
S^4	0	0	0	0	1	0	1

We can continue this classification for $n \geq 2$ as shown in Figure 8. For example \mathcal{I}_n has an orbit of size 7 if and only if $n \equiv 1 \pmod{5}$, in which case the orbit is unique. An orbit of size 11 exists if and only if $n \equiv 1 \pmod{5}$ or $n \equiv 1 \pmod{9}$, and this orbit is unique except when n is both $1 \pmod{5}$ and $1 \pmod{9}$ (i.e., $n \equiv 1 \pmod{45}$) in which case there exist two such orbits. Also, \mathcal{I}_n has an orbit of size 12 if and only if $n \equiv 1 \pmod{8}$, in which case there are always exactly two such orbits.

Corollary 4.11. *For even n , the φ -orbit of \mathcal{I}_n containing the empty set has size $n + 1$.*

	1	2	3	4	5	6
S^0	0	0	0	0	0	0
S^1	1	0	1	0	1	0
S^2	0	0	0	0	0	1
S^3	1	0	1	0	0	0
S^4	0	0	0	1	0	1
S^5	1	0	0	0	0	0
S^6	0	1	0	1	0	1

Figure 9: The φ -orbit of \mathcal{I}_6 containing \emptyset .

Proof. It is easy to see that for even n , $\varphi(0000 \cdots 00) = 1010 \cdots 10$ and $\varphi^2(0000 \cdots 00) = 0000 \cdots 01$. This gives the first three rows of an orbit board, corresponding to aperiodic snake composition $\underbrace{22 \cdots 2}_{(n-2)/2} 1$, whose orbit size is $2^{\frac{n-2}{2}} + 3 = n + 1$. See Figure 9 for an example of this orbit when $n = 6$. ■

Theorem 4.12. *Let \mathcal{O} be an φ -orbit of \mathcal{I}_n and c be a snake composition that appears in \mathcal{O} . If c is aperiodic, then the size of \mathcal{O} is congruent to $1 - n \pmod{4}$. Furthermore, even when c is periodic, the size of \mathcal{O} divides an integer $m \equiv 1 - n \pmod{4}$ for $m \leq 3(n - 1)$ (where m depends on the orbit \mathcal{O}).*

Proof. Using the notation of Proposition 4.3, the size of \mathcal{O} is $\frac{3N_1(c) + 2N_2(c)}{\psi(c)}$. If $c = \underbrace{111 \cdots 1}_{n-1}$, then

$$3N_1(c) + 2N_2(c) = 3(n - 1) \equiv 1 - n \pmod{4}.$$

Any other composition of $n - 1$ into parts 1 or 2 can be formed starting with $111 \cdots 1$ and replacing strings of 11 with 2. Each time a 11 in a snake composition is changed to a 2, the sum $3N_1(c) + 2N_2(c)$ is decreased by 4. Thus, $3N_1(c) + 2N_2(c) \equiv 1 - n \pmod{4}$ and is at most $3(n - 1)$. Thus, the size of \mathcal{O} divides $3N_1(c) + 2N_2(c)$ and when c is aperiodic, the orbit size is given by $3N_1(c) + 2N_2(c) \equiv 1 - n \pmod{4}$. ■

Corollary 4.13. *For even n , every φ -orbit of \mathcal{I}_n has odd size. Furthermore, when $n \equiv 3 \pmod{4}$, there exist no orbits with size divisible by 4.*

Proof. Theorem 4.12 tells us that the size of any orbit \mathcal{O} divides an integer $m \equiv 1 - n \pmod{4}$, which is odd for n even. Meanwhile, $n \equiv 3 \pmod{4}$ forces $m \equiv 2 \pmod{4}$, which $\#\mathcal{O}$ must divide. ■

Corollary 4.14. *When n is even, any reversible φ -orbit of \mathcal{I}_n contains exactly one symmetrical independent set.*

Proof. For a reversible orbit \mathcal{O} and any $S \in \mathcal{O}$, S^{rev} is also in \mathcal{O} . This partitions the independent sets in \mathcal{O} into pairs, with the only unpaired ones being the sets that satisfy

$S = S^{\text{rev}}$, i.e., the symmetrical independent sets. By Corollary 4.13, \mathcal{O} has odd size so there must be an odd number of symmetrical independent sets in \mathcal{O} . By Proposition 2.17, an orbit can contain at most two symmetrical independent sets, so \mathcal{O} contains exactly one symmetrical independent set. ■

5 Connections with Order Ideals in Zigzag Posets


The original problem about independent sets is connected with other well-studied maps, called *promotion* and *rowmotion*, on zigzag posets. Rowmotion was introduced as a map on antichains of a poset in [BS74], while promotion was originally formulated by Schützenberger in the context of standard Young tableaux [Sch63]. Promotion and rowmotion have both been studied in various settings by numerous authors. Here, we consider them as maps on order ideals, as discussed by Striker and Williams [SW12].

Our main result is an equivariant bijection between the φ -action on \mathcal{I}_n and the rowmotion action on order ideals of a zigzag poset. This allows us to carry over results from the previous sections to the new context, which we now describe.

A **partially ordered set** (or **poset** for short) is a set P together with a binary relation \leq that is reflexive, antisymmetric, and transitive. (A reader unfamiliar with posets can find the necessary background in Stanley’s text [Sta11, Ch. 3].) An **order ideal** of a poset P is a subset I of P such that if $x \in I$ and $y \leq x$ in P , then $y \in I$. The set of order ideals of P is denoted $J(P)$.

Definition 5.1. The **zigzag poset** with n elements, denoted \mathcal{Z}_n , is the poset consisting of elements a_1, \dots, a_n and relations $a_{2i-1} < a_{2i}$ and $a_{2i+1} < a_{2i}$. (Such posets are also called **fence posets** and are discussed in [Sta11, p. 367].)

The zigzag posets have Hasse diagrams that can be drawn in a zigzag formation, hence the name. For example



The toggle group on $J(\mathcal{Z}_n)$ is defined analogously to that of \mathcal{I}_n . The main conceptual difference is that while an element can always be toggled out of an independent set, the same is not true for all elements in an order ideal of P .

Definition 5.2. The **toggle** $t_i : J(\mathcal{Z}_n) \rightarrow J(\mathcal{Z}_n)$ is defined as

$$t_i(I) := \begin{cases} I \cup \{a_i\} & \text{if } a_i \notin I \text{ and } I \cup \{a_i\} \in J(\mathcal{Z}_n) \\ I \setminus \{a_i\} & \text{if } a_i \in I \text{ and } I \setminus \{a_i\} \in J(\mathcal{Z}_n) \\ I & \text{otherwise} \end{cases}$$

The **toggle group** of $J(\mathcal{Z}_n)$, denoted $\text{Tog}(\mathcal{Z}_n)$, is generated by the toggles t_i for $i \in [n]$.

As with toggles on \mathcal{I}_n , it is clear that $t_i^2 = 1$ for any $i \in [n]$. The following is an analogue of Proposition 2.2.

Corollary 5.3. *Two toggles $t_i, t_j \in \text{Tog}(\mathcal{Z}_n)$ commute if and only if $|i - j| \neq 1$.*

Proof. When $|i - j| > 1$, a_i and a_j are incomparable (i.e., neither $a_i \leq a_j$ nor $a_j \leq a_i$) so inclusion of one of a_i or a_j cannot affect whether the other can be included. ■

Two special elements of $\text{Tog}(\mathcal{Z}_n)$ are **promotion** $\text{Pro} = t_n \cdots t_2 t_1$ and **rowmotion**

$$\text{Row} = \begin{cases} t_{n-1} t_{n-3} \cdots t_3 t_1 t_n t_{n-2} \cdots t_2 & \text{if } n \text{ is even} \\ t_n t_{n-2} \cdots t_3 t_1 t_{n-1} t_{n-3} \cdots t_2 & \text{if } n \text{ is odd} \end{cases}.$$

These maps have been studied on general posets by many authors. Rowmotion can be defined for any poset, and promotion can be defined for any poset which can be embedded as a rowed-and-columned poset (such as \mathcal{Z}_n) [SW12].

Next we define the equivariant bijection η that takes us between independent sets of \mathcal{P}_n and order ideals of \mathcal{Z}_n .

Corollary 5.4. *The map $\eta : \mathcal{I}_n \rightarrow J(\mathcal{Z}_n)$ defined below is a bijection:*

$$\eta(S) := \{a_i \mid i \in [n], i \text{ odd}, i \notin S\} \cup \{a_i \mid i \in [n], i \text{ even}, i \in S\}.$$


Proof. Note that for $I \subseteq \mathcal{Z}_n$ to be an order ideal, it must satisfy the property that whenever a_{2j} is in I , so are a_{2j-1} and a_{2j+1} (or just a_{2j-1} if $2j = n$).

Suppose $a_{2j} \in \eta(S)$. Then $2j \in S$. Since S is independent, $2j - 1$ and $2j + 1$ are not in S . So $a_{2j-1} \in \eta(S)$ and (when $2j \neq n$) $a_{2j+1} \in \eta(S)$. Hence $\eta(S) \in J(\mathcal{Z}_n)$. The inverse of η is given by

$$\eta^{-1}(I) = \{i \mid i \in [n], i \text{ odd}, a_i \notin I\} \cup \{i \mid i \in [n], i \text{ even}, a_i \in I\}$$

which always produces an independent set for $I \in J(\mathcal{Z}_n)$ by analogous reasoning. Thus η is a bijection. ■

Example 5.5. Let $n = 7$ and $S = 1001010 = \{1, 4, 6\}$. Then $a_1 \notin \eta(S)$ and $a_3, a_5, a_7 \in \eta(S)$ because $1 \in S$ and $3, 5, 7 \notin S$. Also, $a_2 \notin \eta(S)$ and $a_4, a_6 \in \eta(S)$, since $2 \notin S$ and $4, 6 \in S$. This correspondence is shown below, where the hollow circles are included in \mathcal{Z}_7 but not the order ideal $\eta(S)$.



Corollary 5.6. *For every $i \in [n]$, $\eta \circ \tau_i = t_i \circ \eta$. Thus, $\eta \circ \varphi = \text{Pro} \circ \eta$, making η an equivariant bijection, as shown in the following commutative diagrams.*

$$\begin{array}{ccc}
\mathcal{I}_n & \xrightarrow{\eta} & J(\mathcal{Z}_n) \\
\tau_i \downarrow & & \downarrow t_i \\
\mathcal{I}_n & \xrightarrow{\eta} & J(\mathcal{Z}_n)
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{I}_n & \xrightarrow{\eta} & J(\mathcal{Z}_n) \\
\varphi \downarrow & & \downarrow \text{Pro} \\
\mathcal{I}_n & \xrightarrow{\eta} & J(\mathcal{Z}_n)
\end{array}$$

As with \mathcal{T}_n , $\text{Tog}(\mathcal{Z}_n)$ is the quotient of a Coxeter group, and it is clearly isomorphic to \mathcal{T}_n via Proposition 5.6. We define Coxeter elements analogously to \mathcal{T}_n , and Pro and Row are two examples of Coxeter elements in $\text{Tog}(\mathcal{Z}_n)$.

Definition 5.7. A **Coxeter element** in $\text{Tog}(\mathcal{Z}_n)$ is a product of each of the n toggles t_1, t_2, \dots, t_n each used exactly once, in some order.

Theorem 5.8. Any two Coxeter elements in $\text{Tog}(\mathcal{Z}_n)$ are conjugate. In particular, promotion and rowmotion are conjugate.

In [SW12, §5], Striker and Williams prove that Pro and Row are conjugate for any rowed-and-columned poset. However, in the case of $\text{Tog}(\mathcal{Z}_n)$, the conjugacy of any two Coxeter elements follows from the conjugacy of Coxeter elements in \mathcal{T}_n and Proposition 5.6. Thus, the orbit structure of $J(\mathcal{Z}_n)$ under Row, or any other Coxeter element in $\text{Tog}(\mathcal{Z}_n)$, is the same as that of Pro, which by Proposition 5.6, is the same as the orbit structure of φ on \mathcal{I}_n .

Using Proposition 5.6, we can restate Corollary 2.31 for toggling in $J(\mathcal{Z}_n)$, as follows.

Corollary 5.9. Let w be a Coxeter element in $\text{Tog}(\mathcal{Z}_n)$. Let $\chi_{a_j} : J(\mathcal{Z}_n) \rightarrow \{0, 1\}$ be the indicator function of a_j . Then on w -orbits in $J(\mathcal{Z}_n)$, the following statistics are homomesic.

- If n is odd, then $\chi_{a_j} - \chi_{a_{n+1-j}}$ is 0-mesic for every $j \in [n]$. Also $2\chi_{a_1} - \chi_{a_2}$ and $2\chi_{a_n} - \chi_{a_{n-1}}$ are both 1-mesic.
- If n is even, then $\chi_{a_j} + \chi_{a_{n+1-j}}$ is 1-mesic for every $j \in [n]$. Also $2\chi_{a_1} - \chi_{a_2}$ is 1-mesic and $2\chi_{a_n} - \chi_{a_{n-1}}$ is 0-mesic.

Proof. From the definition of η , it is clear that for any $S \in \mathcal{I}_n$,

$$\chi_{a_j}(\eta(S)) = \begin{cases} \chi_j(S) & \text{if } j \text{ is even} \\ 1 - \chi_j(S) & \text{if } j \text{ is odd} \end{cases}.$$

The rest of the proof follows by restating Corollary 2.31 into the language of $J(\mathcal{Z}_n)$ via Proposition 5.6. ■

Notice that our statements above are significantly more complicated to state, forcing us to divide into odd and even cases. This would also make direct proofs of them in the $J(\mathcal{Z}_n)$ -setting more unwieldy. It is much easier to handle them via translation to the \mathcal{I}_n context. This shows the efficacy of Striker's notion of generalized toggling.

It is well-known and not hard to see that for any graded poset P of rank r , there is a rowmotion orbit on $J(P)$ of size $r + 1$ generated by the empty ideal, where $\text{Row}^i(\emptyset)$ consists of all elements of rank $\leq i - 1$. In particular, $J(\mathcal{Z}_n)$ has a rowmotion orbit of size 3. It is not directly obvious that this is the only orbit of this size. But since the orbit structure of Row is the same as that of φ on \mathcal{I}_n , uniqueness follows from Proposition 4.6.

In other proven examples of homomesy for rowmotion on posets, the map generally has a small order and a cyclic sieving phenomenon has been found. However, the rowmotion map on $J(\mathcal{Z}_n)$ has a large order, and thus a natural cyclic sieving result is unlikely, which makes the homomesy for this poset particularly interesting.

Acknowledgements

The authors are grateful to James Propp for suggesting this problem initially and helping us understand its broader context. We have benefitted from useful discussions with David Einstein, Max Glick, Darij Grinberg, Shahrzad Haddadan, Matthew Macauley, Vic Reiner, Elizabeth Sheridan Rossi, Richard Stanley, Jessica Striker, and Nathan Williams. Computations leading to the initial conjectures were done in Sage [S⁺16]. The Online Encyclopedia of Integer Sequences [Slo16] was invaluable for connecting our data with previously known sequences. We also thank an anonymous referee for comments that helped us improve the exposition of this paper.

References

- [AST13] D. Armstrong, C. Stump, and H. Thomas. A uniform bijection between nonnesting and noncrossing partitions. *Transactions of the American Mathematical Society*, 365(8):4121–4151, 2013. Also available as [arXiv:1101.1277v2](#).
- [BB05] A. Björner and F. Brenti. *Combinatorics of Coxeter groups*. Springer Verlag, New York, 2005.
- [BS74] A. Brouwer and L. Schrijver. On the period of an operator, defined on antichains. *Stichting Mathematisch Centrum. Zuivere Wiskunde*, (ZW 24/74):1–13, 1974.
- [CF95] P. Cameron and D. Fon-der-Flaass. Orbits of antichains revisited. *European J. Combin.*, 16(6):545–554, 1995.
- [EE09] H. Eriksson and K. Eriksson. Conjugacy of Coxeter elements. *Electron. J. Combin.*, 16(2):#R4, 2009.
- [EFGJMPR] D. Einstein, M. Farber, E. Gunawan, M. Joseph, M. Macauley, J. Propp, and S. Rubinstein-Salzedo. Noncrossing partitions, toggles, and homomesies. *Electron. J. Combin.*, 23(3), 2016. Also available at [arXiv:1510.06362v2](#).
- [EP13] D. Einstein and J. Propp. Combinatorial, piecewise-linear, and birational homomesy for products of two chains. [arXiv:1310.5294](#), 2013.

- [Had16] S. Haddadan. Some instances of homomesy among ideals of posets. [arXiv:1410.4819v3](#), 2016.
- [Jos17] M. Joseph. Antichain toggling and rowmotion. [arXiv:1709.09331](#), 2017.
- [Pan09] D. Panyushev. On orbits of antichains of positive roots. *European J. Combin.*, 30(2):586–594, 2009. Also available at [arXiv:0711.3353v2](#).
- [Pet15] T. K. Petersen. *Eulerian Numbers*. Springer, New York, 2015.
- [PR15] J. Propp and T. Roby. Homomesy in products of two chains. *Electron. J. Combin.*, 22(3), 2015. Also available at [arXiv:1310.5201v5](#).
- [Rob16] T. Roby. Dynamical algebraic combinatorics and the homomesy phenomenon. In *Recent Trends in Combinatorics*, pages 619–652. Springer, 2016. Also available at <http://www.math.uconn.edu/~troby/homomesyIMA2015Revised.pdf>.
- [RSW04] V. Reiner, D. Stanton, and D. White. The cyclic sieving phenomenon. *Journal of Combinatorial Theory, Series A*, 108(1):17–50, 2004.
- [S⁺16] W. A. Stein et al. *Sage Mathematics Software (Version 7.3)*. The Sage Development Team, 2016. <http://www.sagemath.org>.
- [Sag11] B. Sagan. The cyclic sieving phenomenon: A survey. *London Math. Soc. Lecture Note Ser.*, 392:183–233, 2011. Also available at [arXiv:1008.0790v3](#).
- [Sch63] M. P. Schützenberger. Quelques remarques sur une construction de Schensted. *Math. Scand.*, 12:117–128, 1963.
- [Shi97] J. Shi. The enumeration of Coxeter elements. *J. Algebraic Combin.*, 6(2):161–171, 1997.
- [Slo16] N. J. A. Sloane. *OEIS, Online Encyclopedia of Integer Sequences*, 2016. <https://oeis.org/>.
- [Sta99] R. Stanley. *Enumerative combinatorics, volume 2*. Cambridge University Press, 1999.
- [Sta11] R. Stanley. *Enumerative combinatorics, volume 1, 2nd edition*. Cambridge University Press, 2011. Also available at <http://math.mit.edu/~rstan/ec/ec1/>.
- [Str15] J. Striker. The toggle group, homomesy, and the Razumov-Stroganov correspondence. *Electron. J. Combin.*, 22(2):P2–57, 2015. Also available at [arXiv:1503.08898v1](#).
- [Str16] J. Striker. Rowmotion and generalized toggle groups. [arXiv:1601.03710](#), 2016.
- [SW12] J. Striker and N. Williams. Promotion and rowmotion. *European J. Combin.*, 33:1919–1942, 2012. Also available at [arXiv:1108.1172v3](#).