5-4-2017

Toggling Involutions and Homomesies for Maps on Finite Sets, Noncrossing Partitions, and Independent Sets of Path Graphs

Michael J. Joseph
University of Connecticut, michael.j.joseph@uconn.edu

Toggling Involutions and Homomesies for Maps on Finite Sets, Noncrossing Partitions, and Independent Sets of Path Graphs

Michael James Joseph, Ph.D.

University of Connecticut, 2017

ABSTRACT

This paper explores the orbit structure and homomesy properties of various actions on finite sets. The homomesy phenomenon, meaning constant averages over orbits, was proposed by Propp and Roby in 2011. For many of the known instances of homomesy, Reiner, Stanton, and White's cyclic sieving phenomenon (CSP) is also present. However, we prove homomesy for several maps whose order is large relative to the size of the set, implying that a natural CSP is unlikely. Sometimes we can prove facts about the orbit sizes either as a corollary to the homomesy or by the technique used to prove homomesy.

Many of the actions we describe are products of much simpler ones. Among these, we consider maps defined as products of simple "toggling" involutions. These come from the Striker's theory of generalized toggle groups, an active area of research in dynamical algebraic combinatorics. Several known instances of homomesy have been discovered for elements of toggle groups. While the individual toggles have order two, the order of a composition of several toggles is more difficult to analyze. We also consider an action of "whirling," due to Propp, that can be defined for any family of functions between finite sets. This action is also the composition of simpler ones.

Toggling Involutions and Homomesies for Maps on Finite Sets, Noncrossing Partitions, and Independent Sets of Path Graphs

Michael James Joseph

M.S. University of Connecticut, 2012

M.S. John Carroll University, 2011

B.S. John Carroll University, 2010

A Dissertation

Submitted in Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

at the

University of Connecticut

2017

APPROVAL PAGE

Doctor of Philosophy Dissertation

Toggling Involutions and Homomesies for Maps on Finite Sets, Noncrossing Partitions, and Independent Sets of Path Graphs

Presented by

Michael James Joseph, B.S. Math., M.S. Math.

Major Advisor ________________________
Tom Roby

Associate Advisor ________________________
Ralf Schiffler

Associate Advisor ________________________
David Reed Solomon

University of Connecticut

2017

ACKNOWLEDGMENTS

First, I thank Tom Roby, my advisor, for his guidance on the problems I have worked on and through the entire process. Through the years working with him, I have learned so much about combinatorics, particularly the concepts and problem solving strategies in dynamical algebraic combinatorics, and I have gained a broader understanding of mathematics as a whole. I appreciate him and my other committee members, Ralf Schiffler and David Reed Solomon, for their helpful suggestions and comments on the exposition of the thesis.

I am particularly grateful to James Propp, and his natural skill at discovering conjectures and creating problems to work on. Almost every result I have been a part of began as one of his conjectures or is inspired by one of his conjectures or results. In addition to the many problems he has suggested, he has been very helpful in providing ideas toward solving them.

I thank the American Institute of Mathematics who sponsored a workshop on dynamical algebraic combinatorics that I attended in March 2015. The workshop was organized by James Propp, Tom Roby, Jessica Striker, and Nathan Williams, all of whom I thank. This workshop provided an excellent opportunity for both learning and collaboration. It was there that I began to work with David Einstein, Miriam Farber, Emily Gunawan, Matthew Macauley, James Propp, and Simon Rubinstein-Salzedo, on the problem about toggling noncrossing partitions. The various insights I learned there have not only been useful for that problem but rather across my research as a whole. After the workshop, my understanding of dynamical algebraic combinatorics took a big jump, and I would not have accomplished nearly as much without the workshop or my collaborators. I also appreciate David Einstein

for providing me with Sage code to compute bases for vector spaces of homomesic statistics, which has been invaluable in examining and creating conjectures.

I am grateful for those who I have interacted with and had helpful discussions with at various conferences. There are too many to mention but in addition to those above, I particularly thank Yue Cai, Max Glick, Shahrzad Haddadan, Sam Hopkins, Thomas McConville, Elizabeth Sheridan Rossi, and Richard Stanley.

Lastly, this thesis is the culmination of a learning process that has taken several years. This would not have happened without my parents, Maureen and James Joseph, as well as all of the professors and students I have interacted with at John Carroll University and the University of Connecticut, for their help in leading me down this career path.

Contents

Chapter 1

Introduction

1.1 The homomesy phenomenon

In this paper we investigate several instances of the homomesy phenomenon. The main results are in Chapters 4 through 7. Chapters 1 through 3 present background material and several previous results. The homomesy[1] phenomenon was introduced by Propp and Roby in [32] and is defined as follows.

Definition 1.1.1. Suppose we have a set $\mathcal{S}$, an invertible map $w : \mathcal{S} \to \mathcal{S}$ such that every w-orbit is finite, and a function ("statistic") $f : \mathcal{S} \to \mathbb{K}$, where $\mathbb{K}$ is a field of characteristic 0. Then we say the triple $(\mathcal{S}, w, f)$ exhibits **homomesy** if the average value of f is the same on every orbit. That is, there exists a constant $c \in \mathbb{K}$ such that for every w-orbit $\mathcal{O} \subseteq \mathcal{S}$,

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

[1] Greek for "same middle".

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

We require $\mathbb{K}$ to be a field of characteristic 0 so that we can divide by any positive integer (hence take averages) but we could alternatively replace $\mathbb{K}$ with any $\mathbb{Q}$-vector space. There are extensions of homomesy to non-cyclic group actions, non-invertible monoid actions, and actions that produce infinite length orbits; examples of these can be found in [37]. However, here we will only consider homomesy using the above definition.

Some early isolated examples of homomesy exist in the literature, notably in the conjecture of Panyushev [29], which was proved by Armstrong, Stump, and Thomas [3], and included here as Theorem 2.1.18, but serious investigation of homomesy is quite recent. Although it is a new area of research, the homomesy phenomenon appears to be widespread in numerous combinatorial dynamical systems. Refer to Section 2.1 for several varied examples. Many additional examples are detailed in Roby's survey article [37] as well as in [32, Chapter 2].

In this paper we use the following notation that is common in combinatorics.

- $\mathbb{P} = \{1, 2, 3, 4, \dots\}$. (Think $\mathbb{P}$ for "positive.")
- $\mathbb{N} = \{0, 1, 2, 3, \dots\}$. (Think $\mathbb{N}$ for "nonnegative.")
- $[n] = \{1, 2, \dots, n\}$.
- $\mathfrak{S}_n$ is the symmetric group on $[n]$.
- $\mathfrak{S}_S$ is the symmetric group on the set S.

1.2 Cyclic rotation of binary strings.

An elementary homomesy example described in [32] and [37] is the following. Fix positive integers $k \leq n$, and let $\binom{[n]}{k}$ be the set of length n binary strings with exactly k 1s. Define $C_R : \binom{[n]}{k} \to \binom{[n]}{k}$ to be the rightward cyclic shift map. That is, $C_R(s_1 s_2 \cdots s_n) = s_n s_1 s_2 \cdots s_{n-1}$. Consider the "number of inversions" statistic defined as $\mathrm{inv}(s) := \#\{i < j : s_i > s_j\}$ for $s \in \binom{[n]}{k}$. For example, when $n = 6$ and $k = 2$ there are $\binom{6}{2} = 15$ elements of $\binom{[6]}{2}$, and three orbits under cyclic rotation, as shown in the Figure 1.2.1 (where each column headed "String" represents an orbit).

String	Inv	String	Inv	String	Inv
101000	7	110000	8	100100	6
010100	5	011000	6	010010	4
001010	3	001100	4	001001	2
000101	1	000110	2		
100010	5	000011	0		
010001	3	100001	4		
Average	**4**	**Average**	**4**	**Average**	**4**

FIGURE 1.2.1: Each column is an orbit of $\binom{[6]}{2}$ under cyclic rotation. Notice that the average number of inversions is 4 in each orbit; i.e., the statistic inv is 4-mesic.

Notice that the average number of inversions across every orbit is the same. This is an example of the homomesy phenomenon. The triple $\left(\binom{[6]}{2}, C_R, \mathrm{inv}\right)$ exhibits homomesy with average 4. The general result is that $\left(\binom{[n]}{k}, C_R, \mathrm{inv}\right)$ exhibits homomesy with average $\frac{k(n-k)}{2}$. One of the techniques involved in Propp and Roby's proof of this is to consider **superorbits** of length n. Since C_R is clearly a map of order n, the length of each orbit divides n, and we can repeat each orbit to a superorbit of length n, without changing the average of any statistic across that orbit. In our example, the orbit $(100100, 010010, 001001)$ can be extended to the superorbit $(100100, 010010, 001001, 100100, 010010, 001001)$ and the

average number of inversions is still 4. When considering superorbits that all have the same length, showing that the average number of inversions is the same across every superorbit is equivalent to showing that the total number of inversions is the same across every superorbit.

To conclude the proof, let $s, s' \in \binom{[n]}{k}$ be such that s' is formed from s by changing an occurrence of "10" to "01". Let $\mathcal{O}$ and $\mathcal{O}'$ be the C_R-superorbits of s and s' respectively. Then all but one string in $\mathcal{O}'$ is formed from a string in $\mathcal{O}$ by changing an occurrence of "10" to "01", thus decreasing the number of inversions by 1. The other string in $\mathcal{O}'$ is formed from a string $s'' \in \mathcal{O}$ with a 0 in the first position and a 1 in the final position by swapping the initial 0 and final 1. This swapping adds $n-1$ total inversions. This is because the initial 0 previously was before all k 1s and now is after them (adding k new inversions) and the final 1 previously was before all $n-k$ 0s and now is after them (adding $n-k$ new inversions, one of which we already counted).

There are $n-1$ strings in $\mathcal{O}'$ that each have one less inversion than their corresponding strings in $\mathcal{O}$, and one string in $\mathcal{O}'$ with $n-1$ more inversions than the corresponding string in $\mathcal{O}$. So the total number of inversions is unchanged. Since we can transform any binary string with a fixed length and number of 1s to any other by a series of adjacent swaps, and each swap leaves the total number of inversions unchanged in the corresponding length n superorbits, the average number of inversions is the same for all orbits. See Figure 1.2.2 for an example of how the swapping process affects the number of inversions.

To determine the common average across each orbit, we compute the global average over $\binom{[n]}{k}$. There are $k(n-k)$ pairs consisting of a 0 and 1, each of which will be inverted in half the strings. Thus, the average number of inversions in $\binom{[n]}{k}$ is $\frac{k(n-k)}{2}$. Since every orbit has the same average, it must be the global average $\frac{k(n-k)}{2}$.

Whenever $(\mathcal{S}, w, f)$ exhibits homomesy, the common average of f across every orbit must equal the global average of f across $\mathcal{S}$. Thus $(\mathcal{S}, w, f)$ exhibits homomesy if and only if for

every w-orbit $\mathcal{O}$,

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

This gives an equivalent way to define homomesy. It also means that if c is the average of f on $\mathcal{S}$, then one could prove f to be homomesic under w by showing the average value of f on every w-orbit is always $\leq c$, or by showing the average is always $\geq c$.

String	String	Inversions Change
101000	011000	-1
010100	001100	-1
001010	000110	-1
000101	000011	-1
100010	100001	-1
010001	110000	+5

FIGURE 1.2.2: The first two columns are the orbits under cyclic rotation on $\binom{[6]}{2}$ containing 101000 and 011000, respecively. This demonstrates how changing "10" to "01" in an element of $\binom{[n]}{k}$ decreases 1 from the number of inversions for all but one string in the corresponding superorbits. For the other string, we gain $n-1$ new inversions so the total is unchanged. The changed numbers are displayed in red.

There are other statistics that are homomesic under the action C_R on $\binom{[n]}{k}$. For example, let $I_i(s) := s_i$ be the bit in position i. Then I_i is obviously $\frac{k}{n}$-mesic for any i since each value cycles through position i once during the n iterations that make up a superorbit. The point of this obvious example is to show that homomesy appears all over the place within cyclic group actions on finite sets. Proofs of homomesy range from trivially easy to quite difficult.

1.3 Outline

In Chapter 2, we survey previous work in homomesy demonstrating the significance of breadth of the phenomenon. In Section 2.2, we discuss another recently developed phenomenon called *the cyclic sieving phenomenon* and its connection to homomesy. Chapter 3 provides background on toggle groups, generated by simple involutions called *toggles*.

The new results are in Chapters 4 through 7. The maps we work with in Chapters 4 and 5 are elements in toggle groups, and so we use the material in Chapter 3.

Chapter 4 is about a toggling action on independent sets of a path graph. The main results are Theorem 4.1.13 and its generalization 4.1.33. We apply Coxeter group theory, described in Section 3.3, to explain how we can generalize one result to the other. The main idea in the proof of Theorem 4.1.13 leads to enumerative formulas for the numbers of orbits (Section 4.2) and the sizes of the orbits (Section 4.3). Also, our results can be translated into the language of posets, where we obtain homomesy for well-studied maps *promotion* and *rowmotion* on order ideals of zigzag posets.

Chapter 5 is about toggling noncrossing partitions. Theorem 5.1.5 is the main result. For this, we use a common technique of expressing our statistic as a linear combination of others to prove homomesy. The proof leads to a generalization to graphs; this is the topic of Section 5.3.

Chapters 6 and 7 are about an action called *whirling* on various families of functions between finite sets. In Chapter 6 we study whirling on injections and surjections between finite sets, and generalizations called "m-injections" and "m-surjections". In 7, we examine whirling on the set Park(n) of *parking functions*, which have been well-studied by combinatorialists since their introduction by Konheim and Weiss [26], and on *restricted growth words* which correspond to partitions of sets.

Chapter 2

Survey of previous work in homomesy and cyclic sieving

In this chapter, we detail both the homomesy phenomenon and the *cyclic sieving phenomenon*, another phenomenon introduced fairly recently into dynamical algebraic combinatorics. In Section 2.1, we discuss several previously proven examples of homomesy, along with the phenomenon's significance and breadth. Throughout these examples, we also examine several different common proof techniques. In Section 2.2, we discuss several actions exhibiting the cyclic sieving phenomenon; homomesy has also been discovered in many of these. Due to the tendency for these phenomena to appear in many of the same actions, we have searched for systems with natural homomesic statistics for which a natural cyclic sieving phenomenon seems unlikely. The theme in the main results in the upcoming chapters is that the order of the map is large relative to the size of the ground set.

Readers familiar with homomesy and/or cyclic sieving can skip this chapter (or either section). Some terminology and notation used later is defined in this chapter, such as for rowmotion on posets (Subsection 2.1.4) and noncrossing partitions (Subsection 2.2.1), but

we refer the reader to these in the upcoming chapters.

2.1 Homomesy examples

2.1.1 Suter rotation of Young diagrams.

A **partition** λ of $m \in \mathbb{N}$ is a weakly decreasing sequence $(\lambda_1, \lambda_2, \lambda_3, \cdots)$ of nonnegative integers that add to m. We call λ_i the i^{th} **part** of λ. Since the sum is finite, there exists j such that $\lambda_i = 0$ for all $i > j$. We call the minimum such j the **length** of λ, denoted $\ell(\lambda)$. We usually truncate the zeros and write the partition λ as $(\lambda_1, \lambda_2, \ldots, \lambda_{\ell(\lambda)})$. For example, $(4, 2, 2, 1, 0, 0, \ldots)$ would be written as $(4, 2, 2, 1)$, so the length of this partition is 4. Every partition has an associated **Young diagram** that consists of λ_i boxes in row i, where the rows are counted from the top and left justified. For example, the Young diagram corresponding to $(4, 2, 2, 1)$ is 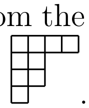. We write $\emptyset$ to denote the empty Young diagram of $(0, 0, \ldots)$, the only partition of 0. The Young diagram of λ is simply another way to display λ, so we consider them to be the same object.

In this homomesy example, we consider the set Y_n of all partitions $\lambda = (\lambda_1, \lambda_2, \lambda_3, \ldots)$ satisfying $\lambda_1 + \ell(\lambda) \leq n$. That is, Y_n consists of all Young diagrams for which the sum of the number of boxes in the first row and left column is at most n (i.e., the box in the top left corner is counted twice in this sum for all nonempty Young diagrams).

In [52], Suter described an action $\text{rot}_n : Y_n \to Y_n$ defined in the following way. Starting with λ, remove the top row (say there are k boxes) and add a column of $n-1-k$ boxes at the left. Then it is elementary to see $\text{rot}_n(\lambda)$ is also in Y_n. For example 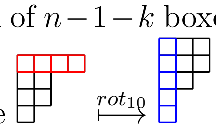

because we removed the top row with 4 boxes and added a left column with $10 - 1 - 4 = 5$ boxes.

Suter showed that $\#Y_n = 2^{n-1}$ for $n \geq 1$ (which is not that difficult to prove) and that rot_n^n is the identity (which is highly nontrivial). Also, he defined a dihedral group action on Y_n generated by rot_n and the **conjugation** operation $\mathrm{conj} : Y_n \to Y_n$ that takes any Young diagram to its transpose, since $\mathrm{rot}_n \circ \mathrm{conj} = \mathrm{conj} \circ \mathrm{rot}_n^{n-1}$.

See Figure 2.1.1 for an illustration of the pentagonal dihedral group action on Y_5, where $\mathrm{rot}_5(\lambda)$ is the Young diagram located by a clockwise rotation of $2\pi/5$ radians about the center  from λ. In particular, rot_5 (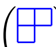) = . The Young diagram found by reflecting across the dashed line from λ is $\mathrm{conj}(\lambda)$. The partition $(2, 1)$ is the only fixed point under rot_5 and it is also self-conjugate. The self-conjugate partitions in Y_5 are those along the dashed line. Lastly, there is an edge in the figure whenever two Young diagrams differ by the addition or removal of a single box; these are the edges in *Young's lattice* [46, p. 254]. Suter showed that rot_n on Y_n exhibits symmetry with respect to this property as well.

David Einstein and James Propp have discovered many homomesic statistics for the action rot_n on Y_n. For example, by looking at the six orbits of rot_6 on Y_6 in Figure 2.1.2, the reader may notice the number of boxes in the top row is $\frac{5}{2}$-mesic, as well as the number of boxes in the left column. In general, these statistics are both $\frac{n-1}{2}$-mesic for rot_n. Note that the number of boxes in the top row and left column are respectively the largest part and length of the corresponding partition.

For the six orbits shown in Figure 2.1.2, the following tuples refer to the numbers of boxes in the top row of each diagram in the orbit.

- Orbit 1: $(0, 1, 2, 3, 4, 5)$
- Orbit 2: $(1, 1, 2, 3, 4, 4)$
- Orbit 3: $(1, 2, 2, 3, 4, 3)$

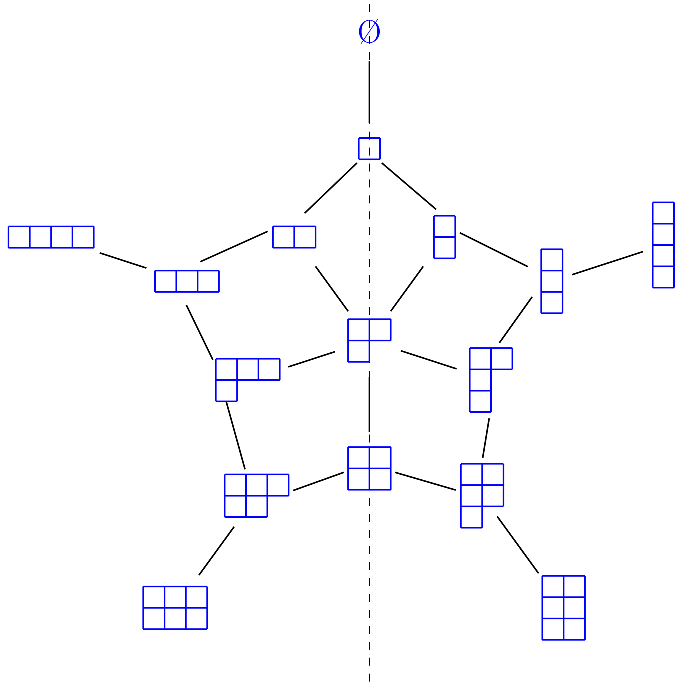

FIGURE 2.1.1: A graph showing the elements of Y_5 as vertices and demonstrating the dihedral symmetry given by rot_5 and conj. For any $\lambda \in Y_5$, $\mathrm{rot}_5(\lambda)$ is the diagram located by a clockwise rotation of $2\pi/5$ radians about the center. Also, $\mathrm{conj}(\lambda)$ is the diagram opposite λ across the dashed line. There are edges between diagrams that differ through addition or removal of one box.

- Orbit 4: $(1, 2, 3, 3, 4, 2)$
- Orbit 5: $(2, 2, 2, 3, 3, 3)$
- Orbit 6: $(2, 3)$

An observant reader may notice that not only is the average $\frac{5}{2}$ for every orbit, but in fact the multiset of numbers of boxes in the top row is symmetric around $\frac{5}{2}$. That is, for any r, any orbit has the same number of partitions with largest part r as it has with largest part $5 - r$. This is in fact true in general.

Theorem 2.1.1. *In any orbit, the multiset of top row (resp. left column) boxes is symmetric around $\frac{n-1}{2}$. This means if there are i diagrams in an orbit with r boxes in the top row, then*

	Average boxes in top row	Average boxes in left column	Average of $I_{1,2} + I_{2,3}$
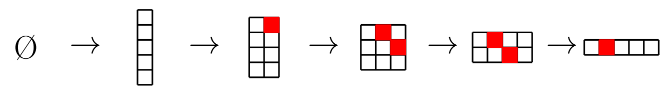	$5/2$	$5/2$	1
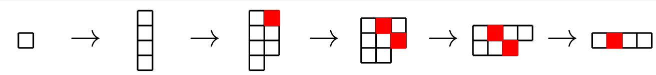	$5/2$	$5/2$	1
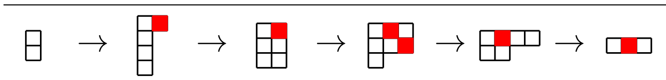	$5/2$	$5/2$	1
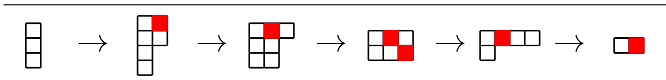	$5/2$	$5/2$	1
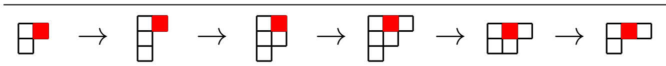	$5/2$	$5/2$	1
	$5/2$	$5/2$	1

FIGURE 2.1.2: The six orbits of rot_6 on Y_6. Note the average number of boxes in the top row and left column is $5/2$ in each orbit. Each occurrence of the box $(1,2)$ or $(2,3)$ is shown in red. Note that $I_{1,2} + I_{2,3}$, the number of red boxes, has average 1 in every orbit; this is an example of Lemma 2.1.4.

there are also i diagrams in that orbit with $n-1-r$ boxes in the top row.

For any $\lambda \in Y_n$ with r boxes in the top row, $\mathrm{rot}_n(\lambda)$ has $n-1-r$ boxes in the left column. Thus we can prove Theorem 2.1.1 for either the top row or left column and it will equivalently hold for the other. We will continue by focusing on this theorem for the top row (largest part).

Theorem 2.1.1 is a stronger result than the aforementioned homomesy. It implies the largest part and length statistics are $\frac{n-1}{2}$-mesic, but having the same average within every orbit does not imply Theorem 2.1.1. However, we can actually restate Theorem 2.1.1 in terms of homomesy! This is useful for the proof.

Definition 2.1.2. For $\lambda \in Y_n$, let $I_{i,j}(\lambda)$ be the indicator function of box (i,j) counted with

matrix coordinates (the box in row i and column j). That is,

$$I_{i,j}(\lambda) = \begin{cases} 1 & \text{if } \lambda_i \geq j, \\ 0 & \text{if } \lambda_i < j. \end{cases}$$

Notice that for $r \geq 1$, λ has largest part r if and only if $I_{1,r}(\lambda) - I_{1,r+1}(\lambda) = 1$. Otherwise $I_{1,r}(\lambda) - I_{1,r+1}(\lambda) = 0$. Thus for given any partition λ and $r \in [n-2]$,

$$I_{1,r}(\lambda) - I_{1,r+1}(\lambda) - I_{1,n-r-1}(\lambda) + I_{1,n-r}(\lambda) = \begin{cases} 1 & \text{if } \lambda \text{ has largest part } r, \\ -1 & \text{if } \lambda \text{ has largest part } n-1-r, \\ 0 & \text{otherwise.} \end{cases} \tag{2.1.1}$$

For the special case $r = 0$ (or equivalently $r = n-1$), Y_n has exactly one partition $\emptyset$ with largest part 0 and one partition $(n-1)$ with largest part $n-1$ because a partition in Y_n with largest part $n-1$ cannot have multiple parts. To compute $\mathrm{rot}_n(n-1)$, we remove the top (and only) row of $(n-1)$, and add a left column with 0 boxes. So $\mathrm{rot}_n(n-1) = \emptyset$. Thus, there is one orbit that contains the partition with largest part 0 and the partition with largest part $n-1$, whereas no other orbit contains any such partitions.

For $r \in [n-2]$, the following theorem is equivalent to Theorem 2.1.1 by Equation (2.1.1).

Theorem 2.1.3. *For any $r \in [n-2]$, $I_{1,r} - I_{1,r+1} - I_{1,n-r-1} + I_{1,n-r}$ is 0-mesic under* rot_n.

To prove Theorem 2.1.3, we employ an often fruitful proof technique of rewriting the statistic as a linear combination of other statistics where homomesy is simpler to show. Given a set S, invertible action τ, and field $\mathbb{K}$, the set of homomesic statistics $f : S \to \mathbb{K}$ forms a vector space [37]. If f is c-mesic and g is d-mesic on τ-orbits, then $f + g$ is $(c+d)$-mesic and kf is kc-mesic for any $k \in \mathbb{K}$. This allows one to form new homomesic statistics from existing ones.

Lemma 2.1.4. *If $i, j, k, \ell \in [n-1]$ satisfy $i + j + k + \ell = n + 2$, then $I_{i,j} + I_{k,\ell}$ is 1-mesic under* rot_n.

Proof. (David Einstein) Without loss of generality assume $i = \ell = 1$. To describe why this can be assumed, first notice that applying rot_n slides every box not in the top row one position up and right. Thus if the block (i, j) is in λ, then $(1, i + j - 1)$ is in $\mathrm{rot}_n^{j-1}(\lambda)$. On the contrary, if the block $(1, i + j - 1)$ is in λ, then (i, j) is in $\mathrm{rot}_n^{1-j}(\lambda)$. So (i, j) appears as a block in an orbit as many times as $(1, i + j - 1)$. Similarly, (k, ℓ) and $(k + \ell - 1, 1)$ appear equally often in any orbit. So we wish to show that $I_{1,j} + I_{n-j,1}$ is 1-mesic. From the definition of rot_n, it is clear that for any λ, exactly one of the following is true:

- $I_{1,j}(\lambda) = 1$, or
- $I_{n-j,1}(\mathrm{rot}_n(\lambda)) = 1$.

Thus, $I_{1,j} + I_{n-j,1}$ is 1-mesic. ■

Let $1 \leq r \leq n - 2$. Lemma 2.1.4 implies that $I_{1,r} + I_{1,n-r}$ and $I_{1,r+1} + I_{1,n-r-1}$ are both homomesic with average 1. To complete the proof of Theorem 2.1.3, note that

$$I_{1,r} - I_{1,r+1} - I_{1,n-r-1} + I_{1,n-r} = (I_{1,r} + I_{1,n-r}) - (I_{1,r+1} + I_{1,n-r-1})$$

is the difference of two 1-mesic statistics and thus 0-mesic.

This example illustrates one of the main reasons to study homomesy. Many results about cyclic actions on a set, like Theorem 2.1.1 can be restated in terms of homomesy. This can allow one to use tools from linear algebra as in the proof above. Furthermore, when working over a finite set $\mathcal{S}$, invertible action w, and set $\{f_1, f_2, \ldots, f_n\}$ of statistics, one can compute the average of each f_i over every w-orbit. Then one can use a computer to determine a basis

for the vector space of all homomesic statistics (under w) that are linear combinations of the f_i statistics. Studying these can lead to insights about the orbits and the action w that would otherwise go unnoticed.

We will use the homomesies in Lemma 2.1.4 as the building blocks to give an alternate proof of homomesy for Suter rotation that first appeared as [32, Proposition 9]. Define the **weight** of the box (i, j) in $\lambda \in Y_n$ to be $n+1-i-j$. That is, the box in the top left corner has weight $n-1$, and the weights decrease as we move right or down. For example, the partition $(5, 4, 1) \in Y_9$ has corresponding weights for each box as shown below:

$$\begin{array}{|c|c|c|c|c|}\hline 8 & 7 & 6 & 5 & 4 \\ \hline 7 & 6 & 5 & 4 \\ \cline{1-4} 6 \\ \cline{1-1}\end{array}\ .$$

In particular, the boxes along any northeast diagonal always have the same weights. Let $f(\lambda)$ be the sum of the weights of the boxes in λ.

The six orbits of rot_6 on Y_6 are shown in Figure 2.1.2. In the order the orbits are displayed, the values of f are as follows: $(0, 15, 24, 27, 24, 15)$, $(5, 14, 23, 26, 23, 14)$, $(9, 18, 21, 24, 21, 12)$, $(12, 21, 24, 21, 18, 9)$, $(13, 16, 19, 22, 19, 16)$, $(19, 16)$. Notice the average within each orbit is $35/2 = (6^3 - 6)/12$. The following is the general result.

Theorem 2.1.5. *Under the action of* rot_n *on* Y_n, f *is homomesic with average* $(n^3 - n)/12$.

Proof. Note that for the box (i, j) to possibly be in $\lambda \in Y_n$, we must have $i, j \in [n-1]$ and $i + j \leq n$. For such i, j, to choose $k, \ell \in [n-1]$ so that $i + j + k + \ell = n + 2$, we first choose k and then ℓ is determined. We have $\ell \in [n-1]$ exactly when $1 \leq k \leq n + 1 - i - j$. Thus the weight of the box (i, j) is the number of possible pairs (k, ℓ). So

$$\sum_{i+j+k+\ell=n+2} (I_{i,j}(\lambda) + I_{k,\ell}(\lambda)) = 2f(\lambda) \tag{2.1.2}$$

because if w is the weight of the box (a, b), then that box is counted on the left side w times as (i, j) and another w times as (k, ℓ).

So f is homomesic with average $\frac{1}{2}\#\{(i,j,k,\ell) \in \mathbb{P}^4 \mid i+j+k+\ell = n+2\}$ by Lemma 2.1.4 and Equation 2.1.2. The number of ordered quadruples (i, j, k, ℓ) that add to $n+2$ (i.e., 4-compositions of $n+2$) is $\binom{n+1}{3} = (n+1)n(n-1)/6$ as these correspond to ways to place three vertical bars in the $n+1$ slots between $n+2$ dots [46, p. 18]. As an example, for $n = 7$,

$$\bullet\bullet\,|\,\bullet\bullet\bullet\bullet\bullet\,|\,\bullet\,|\bullet$$

corresponds to the ordered quadruple $(2, 5, 1, 1)$ that adds to 9. So f is homomesic with average $(n^3 - n)/12$. ■

2.1.2 Homomesy under actions that send one statistic to another.

In combinatorics there are many bijections from a set $\mathcal{S}$ to itself that prove two statistics on $\mathcal{S}$ are equidistributed.

Proposition 2.1.6. *Let $\tau : \mathcal{S} \to \mathcal{S}$ and $f, g : \mathcal{S} \to \mathbb{K}$ be such that $g(\tau(X)) = f(X)$ for all $X \in \mathcal{S}$. If every τ-orbit is finite, then $f - g$ is 0-mesic under the action of τ on $\mathcal{S}$.*

Proof. Let $\mathcal{O} = (X_1, X_2, \ldots, X_\ell)$ be a τ-orbit of $\mathcal{S}$ where $\tau(X_i) = X_{i+1}$ for all $i \in [\ell - 1]$ and

$\tau(X_\ell) = X_1$. Then

$$\begin{aligned}\frac{1}{\ell}\sum_{X\in\mathcal{O}}(f(X)-g(X)) &= \frac{1}{\ell}\left(\sum_{i=1}^{\ell} f(X_i) - \sum_{i=1}^{\ell} g(X_i)\right)\\ &= \frac{1}{\ell}\left(\sum_{i=1}^{\ell} f(X_i) - g(\tau(X_\ell)) - \sum_{i=2}^{\ell} g(\tau(X_{i-1}))\right)\\ &= \frac{1}{\ell}\left(\sum_{i=1}^{\ell} f(X_i) - f(X_\ell) - \sum_{i=1}^{\ell-1} f(X_i)\right)\\ &= 0.\end{aligned}$$

■

Some examples of this occur within the symmetric group $\mathfrak{S}_n$. We have multiple ways to write a permutation $\pi \in \mathfrak{S}_n$. One way is to write π as a product of cycles that includes all of $1, 2, \ldots, n$. The cycle $c = (c_1, c_2, \ldots c_k)$ in π means $c_i \overset{\pi}{\mapsto} c_{i+1}$ for all $i \in [k-1]$ and $c_k \overset{\pi}{\mapsto} c_1$. There are many ways to write a permutation as a disjoint product of cycles. For example, $\pi = (263)(8)(59)(41)(7) \in \mathfrak{S}_9$ could also be written as $\pi = (8)(95)(326)(14)(7)$. We define **canonical cycle decomposition (CCD)** as the unique cycle decomposition that lists the largest element of each cycle first and lists the cycles in increasing order of their largest elements. Thus, the CCD for π is $(41)(632)(7)(8)(95)$. Let $\mathrm{cyc}(\pi)$ denote the number of cycles in a disjoint cycle decomposition of π.

Another common way to represent a permutation $\pi \in \mathfrak{S}_n$ is with its **one-line notation** $\pi(1)\pi(2)\cdots\pi(n)$. For example, the one-line notation for $(2)(3)(541)$ is 52314.

Definition 2.1.7. Let $\pi \in \mathfrak{S}_n$.

- A **left-to-right maximum** (or **record**[1]) of π is a value $\pi(i)$ such that $\pi(i) > \pi(j)$ for

[1]The term *record* comes from an analogy. If we view the one-line notation of π as scores in a game achieved in order, then the left-to-right maxima are precisely the scores that are the record at some point.

all $j < i$.

- A **weak excedance** of π is an $i \in [n]$ for which $\pi(i) \geq i$.

- An **ascent** of π is an $i \in [n-1]$ for which $\pi(i) < \pi(i+1)$.

Let $\underrightarrow{\max}(\pi)$, $\mathrm{wexc}(\pi)$, and $\mathrm{asc}(\pi)$ denote the numbers of left-to-right maxima, weak excedances, and ascents of π, respectively.

Example 2.1.8. The permutation $314526 \in \mathfrak{S}_6$ has four records (3, 4, 5, 6), four weak excedances (1, 3, 4, 6), and three ascents (2, 3, 5).

Let $\widehat{\pi}$ denote the permutation whose one-line notation is formed by dropping the parentheses from the CCD of π. For example, if $\pi = (263)(8)(59)(41)(7)$, then we write π in CCD as $(41)(632)(7)(8)(95)$ so $\widehat{\pi} = 416327895$. The map $\pi \mapsto \widehat{\pi}$, first studied by Alfréd Rényi [35], has several names in the literature; we will call it the Rényi-Foata bijection. It is a bijection as we can recover π from $\widehat{\pi}$ by placing a left parenthesis before each record of $\widehat{\pi}$ and then placing the right parentheses as needed. For example, if $\widehat{\pi} = 314592687$, then $\pi = (31)(4)(5)(92687)$.

# cycles	permutations in $\mathfrak{S}_4$	total
1	$(4123), (4132), (4213), (4231), (4312), (4321)$	6
2	$(21)(43), (31)(42), (32)(41), (1)(432), (1)(423),$ $(2)(413), (2)(431), (3)(412), (3)(421), (321)(4), (312)(4)$	11
3	$(1)(2)(43), (1)(3)(42), (1)(32)(4), (21)(3)(4), (31)(2)(4), (41)(2)(3)$	6
4	$(1)(2)(3)(4)$	1

Figure 2.1.3: The permutations of [4] arranged by number of cycles in CCD. Note that this distribution is the same as that for records, shown in Figure 2.1.4, as a result of the Rényi-Foata bijection.

It is clear by construction that for every $\pi \in \mathfrak{S}_n$, $\underrightarrow{\max}(\widehat{\pi}) = \mathrm{cyc}(\pi)$. Also, for any permutation, weak excedances correspond exactly with occurrences of "ab" for $a < b$ within

# records	permutations in $\mathfrak{S}_4$	total
1	$4123, 4132, 4213, 4231, 4312, 4321$	6
2	$2143, 3142, 3241, 1432, 1423, 2413, 2431, 3412, 3421, 3214, 3124$	11
3	$1243, 1342, 1324, 2134, 3124, 4123$	6
4	1234	1

FIGURE 2.1.4: The permutations of [4] arranged by number of records. Note that this distribution is the same as that for cycles, shown in Figure 2.1.3, as a result of the Rényi-Foata bijection.

# weak excedances	permutations in $\mathfrak{S}_4$	total
1	4123	1
2	$1423, 2143, 2413, 3124, 3142, 3412, 3421, 4132, 4213, 4312, 4321$	11
3	$1243, 1324, 1342, 1432, 2134, 2314, 2341, 2431, 3214, 3241, 4231$	11
4	1234	1

FIGURE 2.1.5: The permutations of [4] arranged by number of weak excedances. Note that this distribution is the same as that for ascents (shifted by 1), shown in Figure 2.1.6, as a result of the Rényi-Foata bijection.

a cycle and the final elements of a cycle in CCD since the largest element is written first. For example, (724351)(8)(96) has weak excedances of 2, 3, 1, 8, 9, 6 (shown in red). When we drop parentheses (e.g., 724351896), there are ascents in the same positions except the final one. It is clear this is true in general; thus $\operatorname{asc}(\widehat{\pi}) + 1 = \operatorname{wexc}(\pi)$.

This is what the Rényi-Foata bijection is used for. It gives a proof that cyc and $\underrightarrow{\max}$

# ascents	permutations in $\mathfrak{S}_4$	total
0	4321	1
1	$1432, 2143, 4312, 3214, 4213, 3142, 4132, 3421, 2431, 4231, 3241$	11
2	$1243, 1324, 1423, 1342, 2134, 3124, 4123, 3412, 2314, 2413, 2341$	11
3	1234	1

FIGURE 2.1.6: The permutations of [4] arranged by number of ascents. Note that this distribution is the same as that for weak excedances (shifted by 1), shown in Figure 2.1.5, as a result of the Rényi-Foata bijection.

are equidistributed statistics on $\mathfrak{S}_n$, and that wexc and asc $+1$ are also equidistributed. See Figures 2.1.3, 2.1.4, 2.1.5, 2.1.6 for $n = 4$. Permutations in $\mathfrak{S}_n$ with a fixed number of cycles (equiv. records) are counted by signless Stirling numbers of the first kind; see [46, p. 26]. Permutations in $\mathfrak{S}_n$ with a fixed number of weak excedances (equiv. ascents) are counted by Eulerian numbers, which are the main focus of a book by Petersen [30].

The following proposition is a direct consequence of Proposition 2.1.6.

Proposition 2.1.9. *Under the Rényi-Foata bijection* $\pi \mapsto \widehat{\pi}$ *on* $\mathfrak{S}_n$*,* $\mathrm{cyc} - \underrightarrow{\max}$ *is 0-mesic and* $\mathrm{wexc} - \mathrm{asc}$ *is 1-mesic.*

Recently, homomesy has been discovered for actions defined by composing the Rényi-Foata bijection with other natural bijections on $\mathfrak{S}_n$ [28].

2.1.3 Rotation of permutation matrices.

This example appears as [37, Example 4].

Another way to display a permutation $\pi \in \mathfrak{S}_n$ is by a matrix. The matrix corresponding to $\mathfrak{S}_n$ is the $n \times n$ matrix in which row i contains 1 in column $\pi(i)$ and 0 in all other columns. For example, the matrix representation of 5671423 is

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

which takes the vector $(1, 2, 3, 4, 5, 6, 7)$ to $(5, 6, 7, 1, 4, 2, 3)$.

Permutation matrices are precisely the square binary matrices with exactly one 1 in every

row and every column. Multiplication of permutation matrices corresponds with composition of permutations.

Similarly to how we defined it for strings in Section 1.2, we can define the number of **inversions** statistic on $\mathfrak{S}_n$ by $\operatorname{inv}(\pi) := \#\{i < j : \pi(i) > \pi(j)\}$. Then $\operatorname{inv}(\pi)$ is the number of pairs that are out of order in the one-line notation of π. For example, $\operatorname{inv}(42135) = 4$ because

- 4 is to the left of 1, 2, and 3, which are less than 4, and
- 2 is to the left of 1, which is less than 2.

The minimal length of π as a product of adjacent transpositions $(i \quad i+1)$ is $\operatorname{inv}(\pi)$ [7, Proposition 1.5.2]. For example, $42135 = (34)(12)(23)(12)$ and no smaller length is possible.

Proposition 2.1.10. *Consider the action $\mathcal{Q}$ on $\mathfrak{S}_n$ that rotates the permutation matrix* 90 *degrees clockwise. Then* inv *is homomesic with average $\frac{n(n-1)}{4}$ under this action.*

Example 2.1.11. For $\mathfrak{S}_3$, there are two orbits under rotation, shown below. These have respective inversion numbers $(0, 3)$ and $(1, 2, 1, 2)$.

$$\left(\begin{bmatrix}1&0&0\\0&1&0\\0&0&1\end{bmatrix}, \begin{bmatrix}0&0&1\\0&1&0\\1&0&0\end{bmatrix}\right) \text{ and } \qquad \left(\begin{bmatrix}0&1&0\\1&0&0\\0&0&1\end{bmatrix}, \begin{bmatrix}0&1&0\\0&0&1\\1&0&0\end{bmatrix}, \begin{bmatrix}1&0&0\\0&0&1\\0&1&0\end{bmatrix}, \begin{bmatrix}0&0&1\\1&0&0\\0&1&0\end{bmatrix}\right)$$

Proof. In terms of the matrix representation of π, $\operatorname{inv}(\pi)$ is the number of column pairs $j_1 < j_2$ for which the 1 in column j_1 is below the 1 in column j_2. Also $\operatorname{inv}(\pi)$ is the number of row pairs $i_1 < i_2$ for which the 1 in row i_1 is to the right of the 1 in row i_2.

Let $1 \leq i_1 < i_2 \leq n$. Suppose in π, the 1 in row i_2 is to the right of that of row i_1. Then in $\mathcal{Q}(\pi)$, the 1 in column $n+1-i_1$ is above the 1 in column $n+1-i_2$. On the contrary, suppose that in π, the 1 in row i_2 is to the left of that of row i_1. Then in $\mathcal{Q}(\pi)$, the 1 in column $n+1-i_1$ is below the 1 in column $n+1-i_2$.

The total possible number of pairs i_1, i_2 is $\binom{n}{2} = \frac{n(n-1)}{2}$. If π has r inversions, then $\mathcal{Q}(\pi)$ has $\frac{n(n-1)}{2} - r$ inversions. Thus, across any given orbit, $\mathrm{inv}(\pi)$ alternates between r and $\frac{n(n-1)}{2} - r$, so the average is $\frac{n(n-1)}{4}$. ■

This is one of many instances where the homomesy result follows from an elementary (and stronger) result. This still portrays how widespread the phenomenon is. In some instances, one can discover homomesy by finding a basis for the space of homomesic statistics that are linear combinations of a given set, and then find the more basic fact as a result.

A generalization of permutation matrices is *alternating sign matrices* [8]. Behrend and Roby have generalized this result to alternating sign matrices [4].

Also, in some cases a homomesy result that follows from a more basic fact can generalize to a homomesy result that does not have this property. For example, in Subsection 2.2.1, we describe a map on noncrossing partitions called *Kreweras complementation*. For orbits of this action, the block count alternates between two numbers that total $n+1$, and thus homomesic with average $\frac{n+1}{2}$. In Chapter 5, we generalize this to a large class of actions for which we still have homomesy but not the stronger fact.

2.1.4 Rowmotion on posets

For many posets, homomesic statistics have been discovered under an action called *rowmotion* on order ideals and antichains. Several examples of such homomesy are detailed in [3], [14],[21], [32], and [49]. In Chapter 3 we describe rowmotion as an element of a toggle group.

Definition 2.1.12. A **partially ordered set** (or **poset** for short) is a set P together with a binary relation $\leq$ on P that satisfies the following properties:

- (reflexive) For every $x \in P$, $x \leq x$.
- (antisymmetric) If $x, y \in P$ satisfy $x \leq y$ and $y \leq x$, then $x = y$.
- (transitive) If $x, y, z \in P$ satisfy $x \leq y$ and $y \leq z$, then $x \leq z$.

Similar to the standard notations with the real numbers, we use the notation $x \geq y$ to mean $y \leq x$, $x < y$ to mean "$x \leq y$ and $x \neq y$," and $x > y$ to mean "$x \geq y$ and $x \neq y$."

Definition 2.1.13. For x, y in a poset P, we say that x is **covered** by y (or equivalently y **covers** x), denoted $x \lessdot y$, if $x < y$ and there does not exist z in P with $x < z < y$. The notation $x \gtrdot y$ means that y is covered by x. If neither $x \leq y$ nor $y \leq x$, we say x and y are **incomparable**.

Infinite posets do not necessarily have cover relations, such as $\mathbb{R}$ with the standard $\leq$ relation. However, for finite posets (and some infinite posets as well), all relations can be formed using cover relations and transitivity. Here we will only consider finite posets, which can be depicted using **Hasse diagrams**. In a Hasse diagram, the vertices represent the elements of the poset and the edges represent cover relations. If there is an edge between two vertices x and y in P, then $x \lessdot y$ if y is drawn higher than x and $y \lessdot x$ if x is drawn higher than y.

Example 2.1.14. For the poset whose Hasse diagram is below, we have five elements a, b, c, d, e in the poset, with cover relations $a \lessdot d$, $b \lessdot d$, $c \lessdot e$, and $d \lessdot e$. Only the cover relations are shown in a Hasse diagram, but we also have the relations $a < e$ and $b < e$ by transitivity.

$P =$ 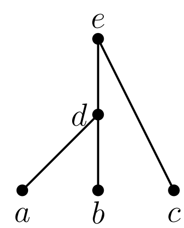

Definition 2.1.15.

- A **subposet**[2] of a poset P is a subset $S \subseteq P$ with the poset structure defined so that if $x, y \in S$ and $x \leq y$ in P, then $x \leq y$ in S also.

- An **order ideal** (resp. **order filter**) of a poset P is a subposet I of P such that if $x \in I$ and $y < x$ (resp. $y > x$) in P, then $y \in I$. The set of order ideals of P is denoted $J(P)$ and the set of order filters of P is denoted $F(P)$.

- An **antichain** of a poset P is a subset $S \subseteq P$ of pairwise incomparable elements. The set of antichains of P is denoted $\mathcal{A}(P)$.

- An element $x \in P$ is a **maximal** (resp. **minimal**) element if P does not contain any $y > x$ (resp. $y < x$).

We have defined the terms used in this paper, but we refer the reader to Stanley's text [46, Ch. 3] for a more thorough introduction to poset theory. Note that for a finite poset P, complementation is a natural bijection between $J(P)$ and $F(P)$. Let $\mathrm{comp}(S)$ denote the complement of a subset $S \subseteq P$. Also, any order ideal (resp. filter) is uniquely determined by its set of maximal (resp. minimal) elements, which is an antichain. Any antichain S of P generates an order ideal $\mathbf{I}(S) := \{x \in P | x \leq y \text{ for some } y \in S\}$ whose set of maximal elements is S and an order filter $\mathbf{F}(S) := \{x \in P | x \geq y \text{ for some } y \in S\}$ whose

[2] Some sources use the term "induced subposet" for this.

set of minimal elements is S. This gives a natural bijection between $J(P)$ and $\mathcal{A}(P)$, and one between $F(P)$ and $\mathcal{A}(P)$.

By composing these bijections, we obtain maps from one of $J(P)$, $\mathcal{A}(P)$, or $F(P)$ into itself.

Definition 2.1.16. For an order ideal $I \in J(P)$, define $\rho_J(I)$ to be the order ideal generated by the minimal elements of the complement of I. For an antichain $A \in \mathcal{A}(P)$, define $\rho_{\mathcal{A}}(A)$ to be the minimal elements of the complement of the order ideal generated by A.

These two maps can each be expressed as the composition of three maps as follows.

$$\begin{aligned} \rho_J \quad &: \quad J(P) \xrightarrow{\text{comp}} F(P) \xrightarrow{\mathbf{F}^{-1}} \mathcal{A}(P) \xrightarrow{\mathbf{I}} J(P) \\ \rho_{\mathcal{A}} \quad &: \quad \mathcal{A}(P) \xrightarrow{\mathbf{I}} J(P) \xrightarrow{\text{comp}} F(P) \xrightarrow{\mathbf{F}^{-1}} \mathcal{A}(P) \end{aligned}$$

Both of these maps are called **rowmotion** and can be written simply as ρ when the context is clear. There is a correspondence between the orbits under these two maps; each $\rho_{\mathcal{A}}$-orbit $\mathcal{O}$ has a corresponding ρ_J-orbit consisting of the order ideals generated by the antichains in $\mathcal{O}$, and vice versa. This relation is depicted by the following commutative diagram.

$$\begin{array}{ccc} J(P) & \xrightarrow{\quad \rho_J \quad} & J(P) \\ \updownarrow & & \updownarrow \\ \mathcal{A}(P) & \xrightarrow[\quad \rho_{\mathcal{A}} \quad]{} & \mathcal{A}(P) \end{array}$$

Rowmotion was introduced as a map on antichains in [9]. It has been studied in various settings by numerous authors and has several names in the literature [51]. The reason for considering the map in both settings is that we can consider statistics for both order ideals

and antichains. We could similarly define $\rho_F : F(P) \to F(P)$ but this is equivalent to ρ_J for the dual poset that swaps the '$\leq$' and '$\geq$' relations, and thus does not give us anything new.

Example 2.1.17. Consider the poset P below. (When we say a poset is a diagram, we mean that diagram is the Hasse diagram of the poset.)

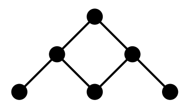

Below we show an example of each of $\rho_{\mathcal{A}}$ acting on an antichain and ρ_J acting on an order ideal as their respective three step processes. In each, hollow circles represent elements of P not in the antichain, order ideal, or order filter.

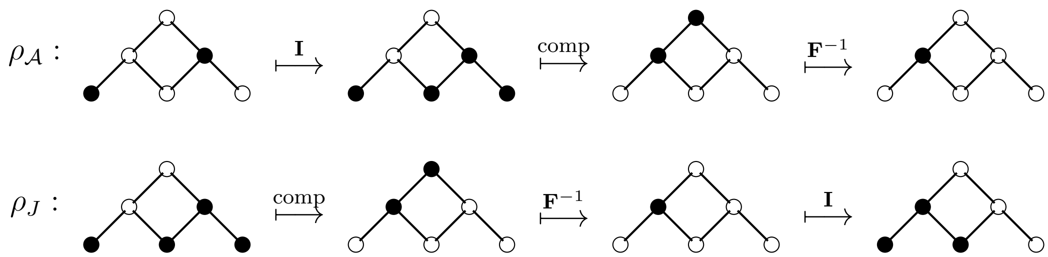

Notice that the chosen order ideal is the one generated by the chosen antichain. After applying ρ to both, we get the order ideal generated by the antichain we obtain.

Homomesy has been found for rowmotion on several families of posets. The following theorem began as a conjecture of Panyushev [29, Conjecture 2.1(iii)] and was later proven by Armstrong, Stump, and Thomas [3, Theorem 1.2]. Panyushev's conjecture is one of the earliest explicit instances of homomesy in the literature, and was proven before the general phenomenon was isolated and named.

Theorem 2.1.18. *Let W be a rank r finite Weyl group and $\Phi^+(W)$ be its positive root poset. Then the cardinality statistic is homomesic with average $r/2$ under the action of $\rho_{\mathcal{A}}$ on $\mathcal{A}(\Phi^+(W))$*

We will neither discuss Weyl groups nor positive root posets, though they are detailed in [7] and arise in representation theory. Up to isomorphism, all but five positive root posets for finite Weyl groups fit into one of three infinite families[3] $\Phi^+(A_n)$, $\Phi^+(B_n)$, and $\Phi^+(D_n)$. The other five are $\Phi^+(W)$ where W is one of E_6, E_7, E_8, F_4, G_2. The elements of these posets are the sets of positive roots listed below. For $\alpha, \beta \in \Phi^+(W)$, the poset relations are given by $\alpha \leq \beta$ if and only if $\beta - \alpha$ is a linear combination of positive roots with nonnegative coefficients. The subscript in each of these Weyl groups is the rank mentioned in Theorem 2.1.18.

- $\Phi^+(A_n) = \{e_i - e_j | 1 \leq i < j \leq n+1\}$.
- $\Phi^+(B_n) = \{e_i \pm e_j | 1 \leq i < j \leq n\} \cup \{e_i | 1 \leq i \leq n\}$.
- $\Phi^+(D_n) = \{e_i \pm e_j | 1 \leq i < j \leq n\}$.

Hasse diagrams for three of these posets are in Figure 2.1.7.

Example 2.1.19. In Figure 2.1.8, we show the three $\rho_{\mathcal{A}}$-orbits of $\mathcal{A}(\Phi^+(A_3))$. Notice that the average cardinality within each orbit is $3/2$, which is consistent with Theorem 2.1.18.

2.2 The cyclic sieving phenomenon

In their 2004 paper [33], Reiner, Stanton, and White introduced another phenomenon that appears in surprisingly many invertible actions on finite sets. Ever since investigation of homomesy began, the cyclic sieving phenomenon has been found in many cyclic actions where homomesy is present, and vice versa.

Before we formally define cyclic sieving, we illustrate it through a simple example.

[3] There is also $\Phi^+(C_n)$ but as posets it is isomorphic to $\Phi^+(B_n)$ so we do not list it.

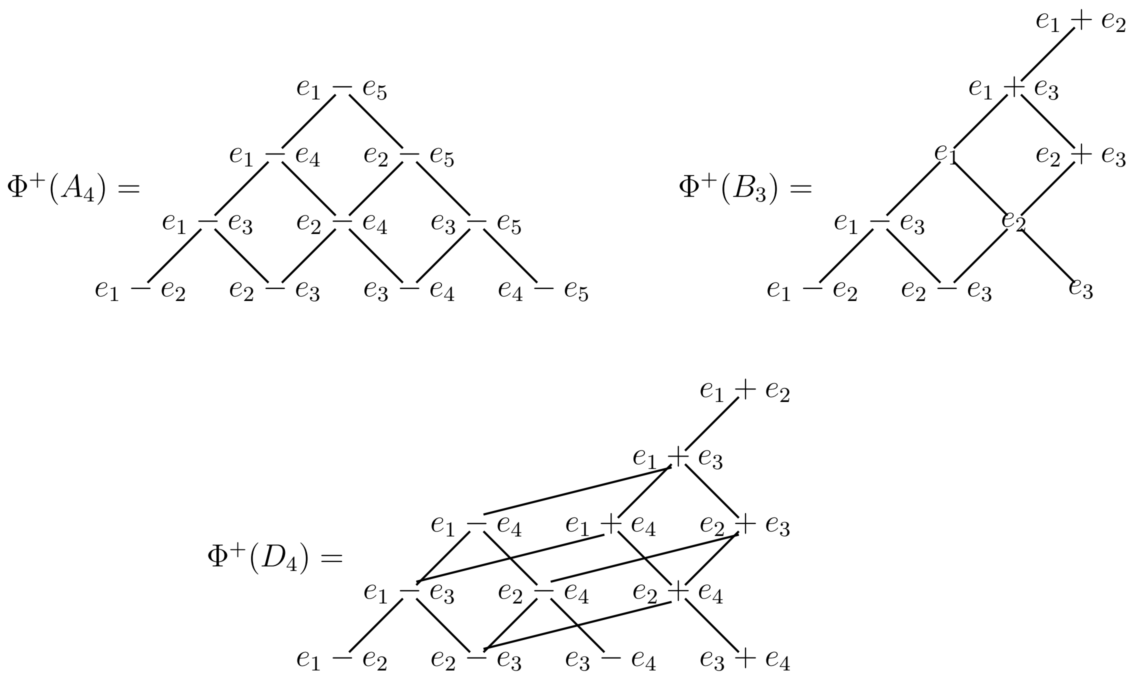

FIGURE 2.1.7: The positive root posets $\Phi^+(A_4)$, $\Phi^+(B_3)$, and $\Phi^+(D_4)$.

Example 2.2.1. Consider the set $\binom{[4]}{2}$ of length four binary strings with exactly two 1s, and the rightward cyclic shift map C_R from Section 1.2. Then there are two orbits under C_R.

$$
\begin{tikzcd}
& 1100 \arrow[dr, bend left, "C_R"] & \\
1001 \arrow[ur, bend left, "C_R"] & & 0110 \arrow[dl, bend left, "C_R"] \\
& 0011 \arrow[ul, bend left, "C_R"] &
\end{tikzcd}
\qquad
\begin{tikzcd}
1010 \arrow[d, bend left=60, "C_R"] \\
0101 \arrow[u, bend left=60, "C_R"]
\end{tikzcd}
$$

There are $\binom{4}{2} = 6$ total elements in $\binom{[4]}{2}$. Of these, none are fixed by C_R or C_R^3, while two are fixed by C_R^2, and all six are fixed by C_R^4 since C_R has order 4.

Now consider the polynomial $X(q) = 1 + q + 2q^2 + q^3 + q^4$. (We will soon explain where this polynomial comes from.) When we plug in powers of $i = \sqrt{-1}$ into the polynomial, we see $X(i) = 0$, $X(i^2) = 2$, $X(i^3) = 0$, and $X(i^4) = 6$. That is, for any $m \in \mathbb{Z}$, the

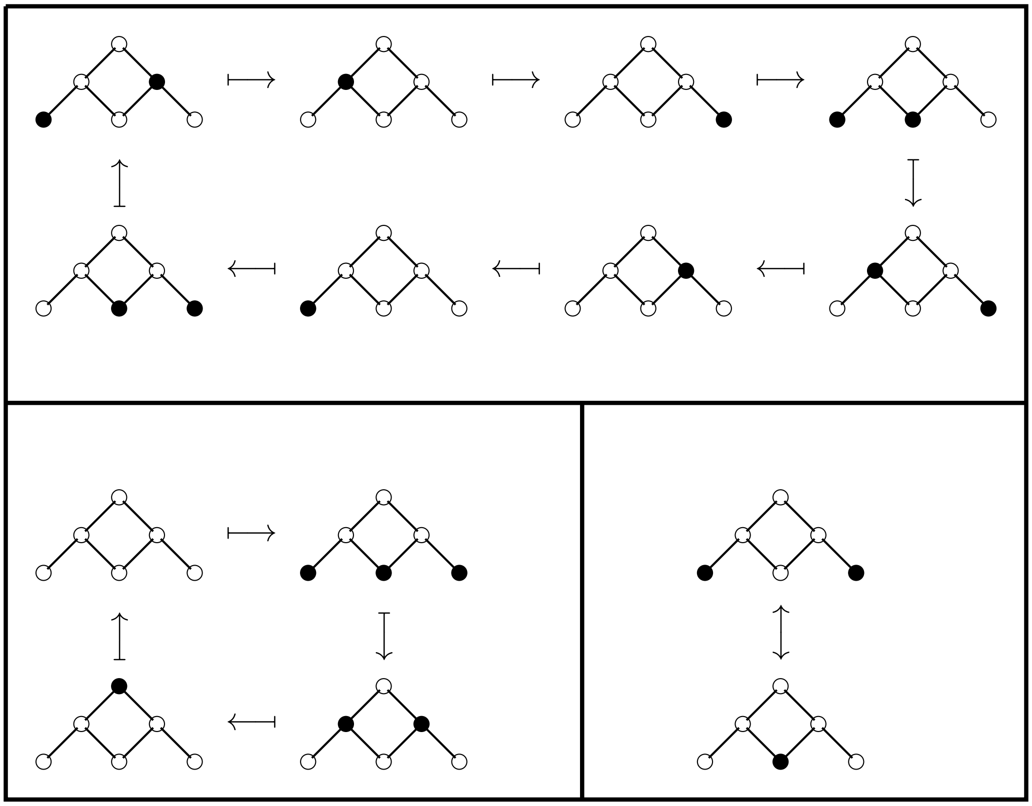

FIGURE 2.1.8: The three orbits of $\rho_{\mathcal{A}}$ on $\mathcal{A}(\Phi^+(A_3))$, each with average cardinality 3/2.

number of elements of $\binom{[4]}{2}$ fixed under C_R^m is $X(i^m)$. For the definition that follows, the triple $\left(\binom{[4]}{2}, X(q), \langle C_R \rangle\right)$ exhibits the cyclic sieving phenomenon.

Definition 2.2.2. Let $\mathcal{C} = \langle c \rangle$ be a cyclic group of order n generated by a map c that acts on a finite set $\mathcal{X}$. Let $\mathcal{X}(q)$ be a polynomial and $\zeta_n = e^{2\pi i/n}$ be a primitive n^{th} root of unity. We say the triple $(\mathcal{X}, X(q), \mathcal{C})$ exhibits the **cyclic sieving phenomenon (CSP)** if for all $m \in \mathbb{Z}$, $X(\zeta_n^m)$ counts the number of elements of $\mathcal{X}$ fixed by c^m.

The CSP is a generalization of Stembridge's $q = -1$ phenomenon [48]. That phenomenon occurs when a set $\mathcal{X}$ has an associated polynomial generating function $X(q)$, for which $X(1)$ counts the cardinality of $\mathcal{X}$, and $X(-1)$ counts the elements of $\mathcal{X}$ fixed under an involution

(map of order 2). The $q = -1$ phenomenon is simply the CSP for an action of order 2.

For more information about the CSP, we direct the reader to the original paper [33], Sagan's survey article [38], and a three page *Notices of the AMS* article [34] meant to be accessible to a general mathematical audience.

The polynomial $X(q)$ from Example 2.2.1 is a ***q*-binomial coefficient** or **Gaussian binomial coefficient** $\binom{4}{2}_q$.

Definition 2.2.3. We define the following q-analogues of common combinatorial numbers.

- $[n]_q := \frac{1-q^n}{1-q} = 1 + q + q^2 + \cdots + q^{n-1}$
- $[n]_q! := [1]_q[2]_q \cdots [n]_q$
- $\binom{n}{k}_q := \frac{[n]_q!}{[k]_q![n-k]_q!}$

Be sure to not confuse $[n]_q$ with $[n] := \{1, 2, \ldots, n\}$. For example,

$$\binom{4}{2}_q = \frac{[4]_q!}{[2]_q![2]_q!} = \frac{1(1+q)(1+q+q^2)(1+q+q^2+q^3)}{1(1+q)1(1+q)} = 1 + q + 2q^2 + q^3 + q^4$$

which is the polynomial from Example 2.2.1.

Notice that when we plug $q = 1$ into $[n]_q$, $[n]_q!$ and $\binom{n}{k}_q$, we get n, $n!$, and $\binom{n}{k}$, respectively. The criterion for a polynomial in q to be a q-analogue of a combinatorial sequence[4] is that we obtain the original sequence by plugging in $q = 1$. There are many q-analogues that have been studied for various combinatorial sequences. They are often used as generating functions that count a given set according to a certain statistic. For a familiar example, 2^n counts the number of subsets of $[n]$. If we replace 2 with $[2]_q = 1 + q$, then from the binomial

[4] The term "sequence" here can also refer to multiple parameter families like the binomial coefficients $\binom{n}{k}$.

theorem

$$[2]_q^n = (1+q)^n = \sum_{k=0}^{n} \binom{n}{k} q^k.$$

Thus, $[2]_q^n$ also counts subsets of $[n]$ but it does so according to a statistic, namely cardinality. The coefficient of q^k is the number of subsets of $[n]$ with cardinality k.

The q-factorial $[n]_q!$ acts in a similar way. We know $n!$ counts the number of permutations in $\mathfrak{S}_n$. There are various statistics f (called *Mahonian* statistics after Percy MacMahon) for which the coefficient of q^k in $[n]_q!$ counts the number of permutations π in $\mathfrak{S}_n$ satisfying $f(\pi) = k$. One Mahonian statistic is the number of inversions described in Subsection 2.1.3. For example

$$[4]_q = 1(1+q)(1 = q+q^2)(1+q+q^2+q^3) = 1+3q+5q^2+6q^3+5q^4+3q^5+q^6$$

and the number of $\pi \in \mathfrak{S}_4$ for which $\mathrm{inv}(\pi)$ is $0,1,2,3,4,5,6$ is $1,3,5,6,5,3,1$, respectively. See Figure 2.2.1.

0 inversions	1234	Total: 1
1 inversion	1243, 1324, 2134	Total: 3
2 inversions	1342, 1423, 2143, 2314, 3124	Total: 5
3 inversions	1432, 2341, 2413, 3142, 3214, 4123	Total: 6
4 inversions	2431, 3241, 3412, 4132, 4213	Total: 5
5 inversions	3421, 4231, 4312	Total: 3
6 inversions	4321	Total: 1

FIGURE 2.2.1: The permutations in $\mathfrak{S}_4$ arranged by inversions. The coefficient of q^k in $[4]_q = 1+3q+5q^2+6q^3+5q^4+3q^5+q^6$ gives the number with k inversions.

Since $\binom{n}{k}_q$ is defined as a rational function, it is not clear a priori that it is always a polynomial. However, it can be shown that it satisfies $\binom{n}{k}_q = \binom{n-1}{k-1}_q + q^k\binom{n-1}{k}_q$ and $\binom{n}{k}_q = q^{n-k}\binom{n-1}{k-1}_q + \binom{n-1}{k}_q$. For $q = 1$, these are both equivalent to the well-known recurrence

$\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k}$ for ordinary binomial coefficients. There is no division in either of the recurrences for $\binom{n}{k}_q$. Using either of them and initial conditions $\binom{n}{0}_q = 1$ and $\binom{0}{k}_q = 0$ for $k \geq 1$, we can construct $\binom{n}{k}_q$ for any n and k. From this, we see that $\binom{n}{k}_q$ is a polynomial.

The q-binomial coefficient also has the purpose of counting objects according to a statistic, but it has another enumerative purpose. For q a power of a prime, let $\mathbb{F}_q$ denote the field of order q. Then $\binom{n}{k}_q$ is the number of k-dimensional subspaces of $\mathbb{F}_q^n$. For example, there are

$$\binom{4}{2}_2 = 1 + 2 + 2(2)^2 + 2^3 + 2^4 = 35$$

subspaces of $\mathbb{F}_2^4$ of dimension 2 and

$$\binom{4}{2}_3 = 1 + 3 + 2(3)^2 + 3^3 + 3^4 = 130$$

subspaces of $\mathbb{F}_3^4$ of dimension 2.

Reiner, Stanton, and White proved the general CSP for C_R on $\binom{[n]}{k}$.

Theorem 2.2.4 ([33, Theorem 1.1(b)])**.** *Let $\binom{[n]}{k}$ be the set of length n binary strings with k 1s and let C_R denote the rightward cyclic shift map. Then $\left(\binom{[n]}{k}, \binom{n}{k}_q, \langle C_R \rangle\right)$ exhibits the CSP.*

There are multiple proofs given in [33] regarding how to prove Theorem 2.2.4. One of them is using an explicit algebraic calculation. If ω is a power of ζ_n, then ω has order d for some $d|n$. The number of strings in $\binom{[n]}{k}$ fixed by $C_R^{n/d}$ is $\binom{n/d}{k/d}$ if $d|k$ and 0 otherwise. This is because such strings consist of concatenating the same length n/d pattern with k/d 1s a total of d times. Then proving Theorem 2.2.4 amounts to showing directly that plugging ω into $\binom{n}{k}_q$ yields $\binom{n/d}{k/d}$ or 0 depending on whether $d|k$ or not. However, this proof technique gives no intuition as to why plugging in roots of unity into a polynomial that is already

well-studied in combinatorics happens to count symmetry classes. For more intuition as to why we have the CSP, another proof uses representation theory, which is a useful technique in proving CSPs. However, it still is not very clear why the CSP has been found for so many natural actions on combinatorial sets.

2.2.1 CSPs for Catalan objects

The sequence $1, 1, 2, 5, 14, 42, 132, 429, 1430, 4862, \ldots$ of Catalan numbers $C_n := \frac{1}{n+1}\binom{2n}{n}$ has at least 214 known combinatorial interpretations. These are published, together with additional exercises, in a recent book by Richard Stanley [47]. The number of antichains (or order ideals or order filters) of $\Phi^+(A_{n-1})$, introduced in Theorem 2.1.18, is C_n. Some notable objects counted by C_n are parenthesizations of multiplication, nonnesting partitions, noncrossing matchings, nonnesting matchings, permutations that avoid a fixed length three pattern, binary trees, and Dyck paths.

In this section, we will consider four sets of Catalan objects that are relevant to our discussion: triangulations of convex polygons, noncrossing partitions, noncrossing $(1, 2)$-configurations, and ballot sequences.

Definition 2.2.5. A **triangulation** of a convex polygon is a division of the polygon into triangles by inserting noncrossing diagonals.

The Catalan number C_n counts the number of triangulations of a convex $(n + 2)$-gon (with labeled vertices). See Figure 2.2.2 for the $C_4 = 14$ triangulations of a regular hexagon.

We defined partitions of integers earlier. Now we define partitions of sets.

Definition 2.2.6. A **partition** of the set $[n]$ is an unordered collection of disjoint subsets of $[n]$, called **blocks**, whose union is $[n]$.

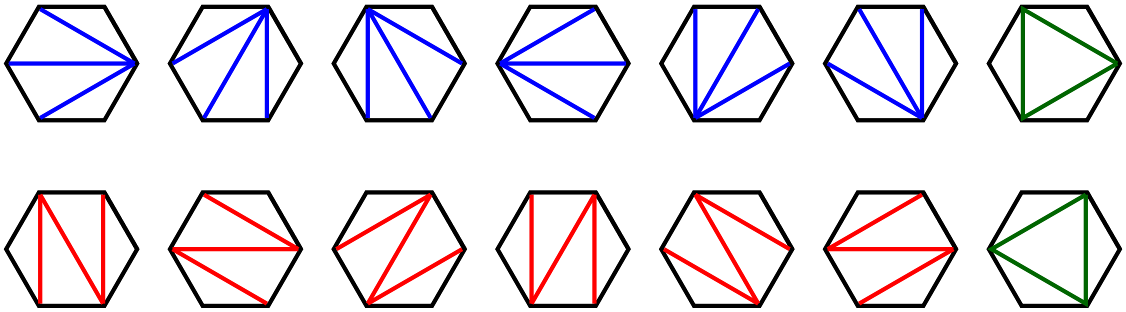

FIGURE 2.2.2: The $C_4 = 14$ triangulations of a regular $(4+2)$-gon. Of these, the six with an "N" shaped edge pattern (shown in red) are fixed under a 180 degree rotation. The two with a triangular edge pattern (shown in green) are fixed under a 120 (or 240) degree rotation. None of them are fixed under a 60 or 300 degree rotation.

For example $\{\{1,2,5\},\{4,7\},\{3,6\}\}$ is a partition of $[7]$. We often write a partition by writing each block without set braces or commas, and using a bar to separate blocks. For example, $\{\{1,2,5\},\{4,7\},\{3,6\}\}$ would be written $125|47|36$. Notice that $125|47|36 = 125|36|47 = 63|152|47$ because the order of the blocks is unimportant as well as the order of the numbers within each block.

Definition 2.2.7. A partition of $[n]$ is said to be **noncrossing** if whenever $1 \leq i < j < k < \ell \leq n$, it is not the case that i and k belong to one block of π with j and ℓ belonging to another block. Let $\mathrm{NC}(n)$ denote the set of noncrossing partitions of $[n]$.

The Catalan number C_n counts the number of noncrossing partitions of $[n]$. See Figure 2.2.3 for the $C_4 = 14$ noncrossing partitions of $[4]$. We show each partition by its *circular representation*, in which n points are drawn around a circle. The motivation behind the term "noncrossing" is that the convex hulls of the blocks do not cross each other. Only one partition of $[4]$ is not noncrossing, namely $13|24$.

The following action on $\mathrm{NC}(n)$, called *Kreweras complementation*, was introduced in [27] and further investigated in [22].

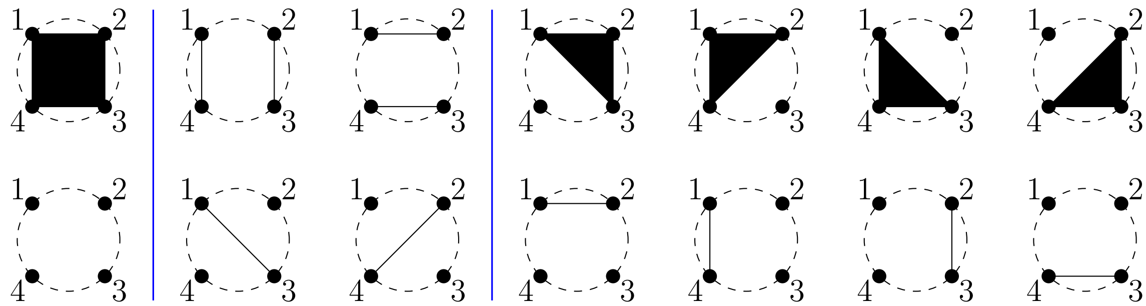

FIGURE 2.2.3: The $C_4 = 14$ noncrossing partitions of [4]. Only the two to the left of the first blue vertical line are fixed under a 90 or 270 degree rotation. All six to the left of the second blue line are fixed under a 180 degree rotation.

Definition 2.2.8. Let $\pi \in \mathrm{NC}(n)$. Draw π on a circle, like those shown in Figure 2.2.3, and insert a new point i' immediately clockwise from i along the circle. The Kreweras complement $\kappa(\pi)$ is the coarsest noncrossing partition of the primed numbers in the complement of π (but considered a partition on $\{1, 2, \ldots, n\}$ not on $\{1', 2', \ldots, n'\}$).

One significance of κ is that it acts as a "square root" of the map r that rotates a noncrossing partition's circular representation counterclockwise by $2\pi/n$ radians.

Proposition 2.2.9. *Let P be a noncrossing partition. Applying the Kreweras complement twice to P rotates P counterclockwise by $2\pi/n$ radians. So the order of κ divides $2n$.*

See Figure 2.2.4 for an example of the Kreweras complement and for an illustration of Proposition 2.2.9.

Definition 2.2.10. Call a subset of $[m]$ a **ball** if it has cardinality 1 and an **arc** if it has cardinality 2. A **(1,2)-configuration** of $[m]$ is a collection of pairwise disjoint balls and arcs chosen from $[m]$. (Some elements of $[m]$ may not be included in a ball or arc at all). A (1,2)-configuration is **noncrossing** if it does not contain a pair of arcs $\{i, k\}$, $\{j, \ell\}$ where $i < j < k < \ell$.

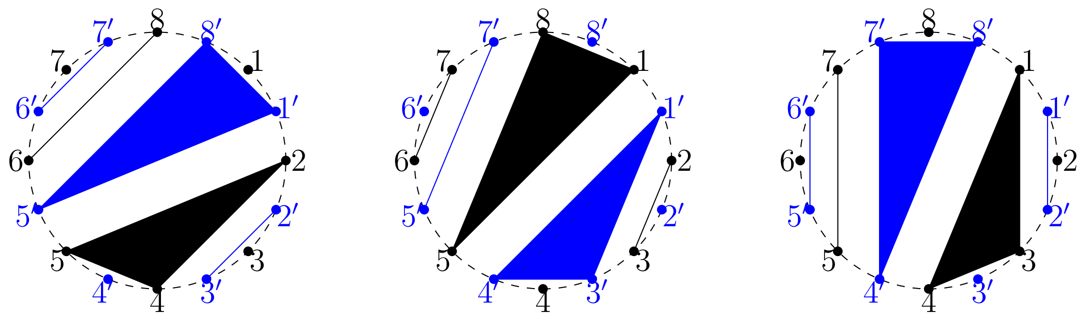

FIGURE 2.2.4: Applying the Kreweras complement κ to the noncrossing partition $P = 1|245|3|68|7$ shown at left yields $\kappa(P) = 158|23|4|67$, which is the blue noncrossing partition on the left, and the black one in the middle. Applying κ twice yields $\kappa^2(P) = 134|2|57|6|8$, shown at right. The convex hulls of $\kappa^2(P)$ and P differ by a counterclockwise rotation of $2\pi/8$ radians.

The Catalan number C_n counts the number of noncrossing (1,2)-configurations of $[n-1]$. We depict these by their *circular representations* in which $1, 2, \ldots, n$ are drawn around the circle. See Figure 2.2.5 for the $C_4 = 14$ noncrossing (1,2)-configurations of $[4-1] = [3]$.

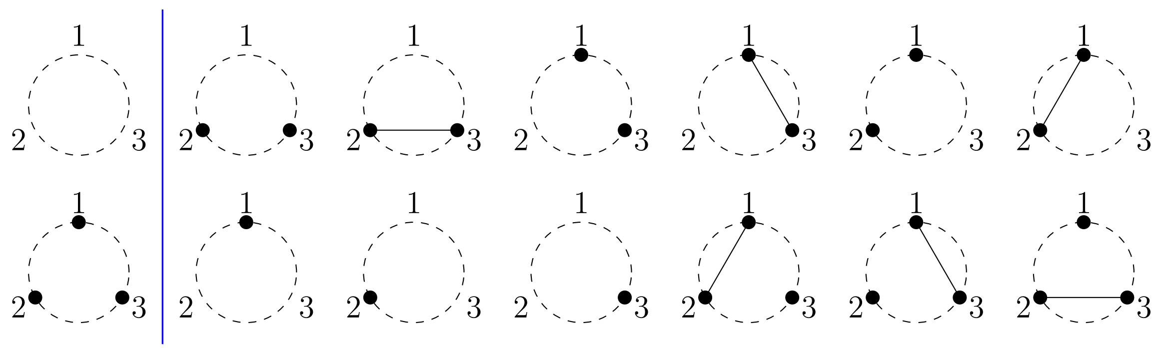

FIGURE 2.2.5: The $C_4 = 14$ noncrossing (1,2)-configurations of [3]. A point by itself denotes a ball, while an arc is depicted by two points connected by a line segment. If there is no point next to a number, then it is neither part of a ball nor an arc. Only the two to the left of the blue vertical line are fixed under a 120 or 240 degree rotation.

Since the Catalan numbers are given by the formula $C_n = \frac{1}{n+1}\binom{2n}{n}$, one natural q-analogue is $C_n = \frac{1}{[n+1]_q}\binom{2n}{n}_q$. This is the cyclic sieving polynomial for many natural actions on Catalan objects, including the ones listed here.

Theorem 2.2.11.

1. *If $\mathcal{X}$ is the set of triangulations of a regular $(n+2)$-gon, and $r : \mathcal{X} \to \mathcal{X}$ represents a counterclockwise rotation of $\frac{2\pi}{n+2}$ radians, then $\left(\mathcal{X}, \frac{1}{[n+1]_q}\binom{2n}{n}_q, \langle r\rangle\right)$ exhibits the CSP.*

2. *If $r : \mathrm{NC}(n) \to \mathrm{NC}(n)$ represents the counterclockwise rotation of $\frac{2\pi}{n}$ radians of the circular representation, then $\left(\mathrm{NC}(n), \frac{1}{[n+1]_q}\binom{2n}{n}_q, \langle r\rangle\right)$ exhibits the CSP.*

3. *For the Kreweras complement $\kappa : \mathrm{NC}(n) \to \mathrm{NC}(n)$, $\left(\mathrm{NC}(n), \frac{1}{[n+1]_q}\binom{2n}{n}_q, \langle \kappa\rangle\right)$ exhibits the CSP.*

4. *The order of the rowmotion action $\rho_{\mathcal{A}}$ on $\mathcal{A}(\Phi^+(A_{n-1}))$ is $2n$. Also $\left(\mathcal{A}(\Phi^+(A_{n-1})), \frac{1}{[n+1]_q}\binom{2n}{n}_q, \langle \rho_{\mathcal{A}}\rangle\right)$ exhibits the CSP.*

5. *If $\mathcal{X}$ is the set of noncrossing (1,2)-configurations of $[n-1]$, and $r : \mathcal{X} \to \mathcal{X}$ represents a counterclockwise rotation of $\frac{2\pi}{n-1}$ radians of the circular representation, then $\left(\mathcal{X}, \frac{1}{[n+1]_q}\binom{2n}{n}_q, \langle r\rangle\right)$ exhibits the CSP.*

While we only need a convex $(n+2)$-gon for the number of triangulations to be counted by C_n, we need the polygon to be regular in order to have the required symmetry in part (1). Note that the same polynomial $\frac{1}{[n+1]_q}\binom{2n}{n}_q$ is a cyclic sieving polynomial for actions of four different orders, respectively $n+2$, n, $2n$, and $n-1$. Part (2) is a direct consequence of (3) by Proposition 2.2.9 but since the rotation action is more natural in its own right, we also mention it separately.

Example 2.2.12. For $n = 4$, let

$$X(q) = \frac{1}{[5]_q}\binom{8}{4}_q = 1 + q^2 + q^3 + 2q^4 + q^5 + 2q^6 + q^7 + 2q^8 + q^9 + q^{10} + q^{12}.$$

1. Let $\mathcal{X}$ be the set of triangulations of a regular hexagon, shown in Figure 2.2.2, and let r denote a rotation of $\frac{2\pi}{6}$ radians counterclockwise. Of the 14 triangulations, none of them are fixed under r or r^5, two of them are fixed under r^2 and r^4, and six of them are fixed under r^3. For $\zeta = \zeta_6 = e^{2\pi i/6} = \frac{1}{2} + \frac{\sqrt{3}}{2}i$, we have
$$X(1) = 14, \qquad X(\zeta) = 0, \qquad X(\zeta^2) = 2, \qquad X(\zeta^3) = 6, \qquad X(\zeta^4) = 2, \qquad X(\zeta^5) = 0.$$

2. Let $\mathcal{X} = \mathrm{NC}(4)$, shown in Figure 2.2.3, and let r denote a rotation of $\frac{2\pi}{4}$ radians counterclockwise. Of the 14 noncrossing partitions, two of them are fixed under r and r^3, and six of them are fixed under r^2. For $i = \zeta_4 = e^{2\pi i/4}$, we have
$$X(1) = 14, \qquad X(i) = 2, \qquad X(i^2) = 6, \qquad X(i^3) = 2.$$

3. Let $\mathcal{X} = \mathrm{NC}(4)$, shown in Figure 2.2.3, and let κ be Kreweras complementation. None of the noncrossing partitions are fixed under κ. When m is odd, κ^m has full order 8, so no $P \in \mathrm{NC}(4)$ is fixed under κ^m either. For $\zeta = \zeta_8 = e^{2\pi i/8} = \frac{\sqrt{2}}{2} + \frac{\sqrt{2}}{2}i$, we have
$$X(\zeta) = 0, \qquad X(\zeta^3) = 0, \qquad X(\zeta^5) = 0, \qquad X(\zeta^7) = 0.$$

4. Let $\mathcal{X} = \mathcal{A}(\Phi^+(A_3))$, shown in Figure 2.1.8. The order of rowmotion ρ is 8. None of the antichains are fixed under ρ^m, for $m = 1, 3, 5, 7$, while two of them are fixed for $m = 2, 6$ and four of them for $m = 4$. This is consistent with the values of $X(q)$ in the above two parts.

5. Let $\mathcal{X}$ be the set of noncrossing (1,2)-configurations of [3], shown in Figure 2.2.5, and let r denote a rotation of $\frac{2\pi}{3}$ radians counterclockwise of the circular representation. Of the 14 configurations, two of them are fixed under r and r^2. For $\zeta = \zeta_3 = e^{2\pi i/3} = -\frac{1}{2} + \frac{\sqrt{3}}{2}i$, we have $X(1) = 14$, $X(\zeta) = 2$, and $X(\zeta^2) = 2$.

Parts (1) and (2) of Theorem 2.2.11 are consequences of more general facts Reiner, Stanton, and White proved in their original paper [33, Theorems 7.1 and 7.2]. Part (3) was

proven by White using direct calculation [5] and later generalized by Armstrong, Stump, and Thomas to "noncrossing partitions" over any finite Coxeter group. They also proved a more general result for (4) describing *any* Weyl group, not just A_n [3, Theorem 1.5]. In fact, they did so by creating an *equivariant bijection* between $\mathrm{NC}(n)$ under κ and $\mathcal{A}(\Phi^+(A_{n-1}))$, i.e., a bijection between $\mathrm{NC}(n)$ and $\mathcal{A}(\Phi^+(A_{n-1}))$ that makes the diagram below commute.

$$\begin{tikzcd}
\mathrm{NC}(n) \arrow[r, "\kappa"] \arrow[d, leftrightarrow] & \mathrm{NC}(n) \arrow[d, leftrightarrow] \\
\mathcal{A}(\Phi^+(A_{n-1})) \arrow[r, "\rho_{\mathcal{A}}"'] & \mathcal{A}(\Phi^+(A_{n-1}))
\end{tikzcd}$$

They defined this bijection uniformly for all Weyl groups (though with a type-by-type proof), though we just consider the type A_n result here. In Figure 2.1.8, we display the three rowmotion orbits on $\mathcal{A}(\Phi^+(A_3))$ with sizes 2, 4, and 8. By this correspondence, we must have three orbits under Kreweras complementation on NC(4) with sizes 2, 4, and 8. In fact, each section of Figure 2.2.3 separated by the vertical lines represents a κ-orbit, though the order within the orbit does not correspond to the listing of the noncrossing partitions in the figure.

Propp and Reiner conjectured $\frac{1}{[n+1]_q}\binom{2n}{n}_q$ to yield a CSP for some set of Catalan objects with order $n-1$. This was based on the fact that plugging in a root of unity of order $n-1$ appeared to always output a nonnegative integer. This led one to search for such a CSP, and Theorem 2.2.11(5) was recently discovered and proven by Thiel [53]. Noncrossing (1,2)-configurations are not among the most well-known and well-studied combinatorial interpretations of C_n. For certain values of n, plugging in a root of unity of an order not dividing $n-1$, $n+2$, or $2n$ into $\frac{1}{[n+1]_q}\binom{2n}{n}_q$ does not necessarily output nonnegative integers, meaning CSPs for maps of those orders are impossible with this polynomial.

As with many polynomials for which the CSP has been discovered, the q-Catalan number $\frac{1}{[n+1]_q}\binom{2n}{n}_q$ has been of interest as a generating function long before it was proven to exhibit the CSP. Fürlinger and Hofbauer introduced the q-Catalan numbers in 1983 in studying ballot sequences [18]. They proved that the coefficient of q^k is the number of length n ballot sequences (which are counted in total by C_n) whose major index is k.

Definition 2.2.13. A **ballot sequence**[5] of length $2n$ is a sequence of n '+' symbols and n '$-$' symbols for which any initial segment contains at least as many '+' as '$-$' symbols. The **major index** of a ballot sequence is the sum of the *positions* containing a '$-$' that is immediately followed by '+'.

For example, $++-+--+-$ is a ballot sequence. Of the four minus signs, only the ones in positions 3 and 6 are immediately followed by a '+'. Thus, this ballot sequence has major index $3+6=9$. See Figure 2.2.6 for all $C_4 = 14$ ballot sequences of length 8, organized according to their major indices.

Ballot Sequence	**Major Index**
$++++----$	0
$+-+++---$	2
$++-++---$	3
$++--++--$	4
$+++-+---$	4
$+++--+--$	5
$+++---+-$	6

Ballot Sequence	**Major Index**
$+-+-++--$	6
$+-++-+--$	7
$+-++--+-$	8
$++-+-+--$	8
$++-+--+-$	9
$++--+-+-$	10
$+-+-+-+-$	12

FIGURE 2.2.6: The $C_4 = 14$ ballot sequences of length 8, arranged by major index. Compare with the q-Catalan number $\frac{1}{[5]_q}\binom{8}{4}_q = 1 + q^2 + q^3 + 2q^4 + q^5 + 2q^6 + q^7 + 2q^8 + q^9 + q^{10} + q^{12}$ for $n = 4$.

[5]The term "ballot sequence" comes in connection with the following voting related problem. Imagine that '+' and '$-$' represent two candidates who both receive n votes in an election. The ballot sequences correspond to the orders of counting the votes for which '+' never trails.

2.2.2 Connections between cyclic sieving and homomesy.

Based on observation, it has long been believed that there are connections between homomesy and the CSP. A lot of actions that have been found to exhibit one of these phenomena also exhibit the other. So far we have discussed both homomesy and cyclic sieving for cyclic rotation of binary strings. Similarly, we have both homomesy and cyclic sieving for rowmotion on antichains and order ideals of $\Phi^+(W)$. We only explicitly stated the CSP for $W = A_{n-1}$ but there is a more general result [3, Theorem 1.5].

It is elementary to verify that for the Kreweras complement $\kappa : \mathrm{NC}(n) \to \mathrm{NC}(n)$, any $P \in \mathrm{NC}(n)$ satisfies $|P| + |\kappa(P)| = n + 1$, where $|P|$ denotes the number of blocks of $|P|$ (Proposition 5.1.3). Thus across any κ-orbit, the block count alternates between two numbers that sum to $n+1$. So as with the homomesy in Subsection 2.1.3, the block count is $(n+1)/2$-mesic under this action. The main motivation for studying the toggle actions on noncrossing partitions in Chapter 5 is to generalize this result to a large family of actions, one of which is κ. These actions w do not, in general, have a predictable order or orbit structure. They also do not generally satisfy $|P| + |w(P)| = n + 1$ but still yield the same homomesy.

A triangulation of a regular $(n+2)$-gon involves $n-1$ diagonals, so $2n-2$ total endpoints of diagonals. Therefore, if we pick a certain vertex location L, then as we rotate the polygon $n+2$ times, every endpoint will pass through location L exactly once. This means that, by considering length $n+2$ superorbits, the statistic e_L of the number of endpoints of diagonals meeting at L is homomesic with average $\frac{2n-2}{n+2}$ under rotation by $\frac{2\pi}{n+2}$ radians.

Perhaps one reason that homomesy appears in most known actions where the CSP has been discovered is that these actions have known orbit structures. The order of the map is known for all the actions with homomesy results discussed in the first two sections except

Proposition 2.1.9 about the Rényi-Foata bijection. However, that homomesy result is not nearly as interesting as the stronger result it was derived from.

There are homomesic statistics that have been proven for maps of finite order on infinite (in fact, uncountable) sets. For example, Lyness 5-cycles and Bloch 5-cycles are maps of order five that act on almost all of $\mathbb{R}^2$. These maps have connections to frieze patterns and cluster algebras. Homomesy under these actions was proven by Andy Hone [32, §2.6]. The fact that these maps have order five is significant for the proofs. In fact, knowing that a map on an infinite set only produces finite orbits (as in our definition of homomesy) usually amounts to knowing the order of the map. On both finite and infinite sets, understanding the order of the map is often important in proving homomesy, and considering superorbits has always been a common proof technique.

The possible relationship between homomesy and cyclic sieving is a bit unclear. One reason is that given any action on a finite set, there are homomesic statistics. Any constant statistic is obviously homomesic. Even considering nonconstant statistics, we can randomly construct homomesic statistics. Consider the two orbits $(1100, 0110, 0011, 1001)$ and $(1010, 0101)$ for cyclic rotation of length 4 binary strings with two 1s. If we create a statistic f by $f(1100) = 9$, $f(0110) = 6$, $f(0011) = 4$, $f(1001) = 3$, $f(1010) = 10$, and $f(0101) = 1$, then f is homomesic with average $11/2$. Clearly this is not an interesting statistic. For potentially more interesting statistics, if there are N orbits under an action τ, and we consider $M \geq N$ total statistics on the ground set, then the rank theorem implies the dimension of the vector space of homomesic statistics that are linear combinations of the considered statistics is at least $M - N + 1$. Homomesy is interesting when we can state a result for a family of sets depending on a parameter.

Just as we always can find homomesy, we can always construct a cyclic sieving polynomial for a given invertible action on a finite set.

Proposition 2.2.14 ([33, Proposition 2.1(ii)])**.** *For a cyclic group $\mathcal{C} = \langle c \rangle$ of order n acting on $\mathcal{X}$ and a polynomial $X(q)$, the triple $(\mathcal{X}, X(q), \mathcal{C})$ exhibits the CSP if and only if*

$$X(q) \equiv \sum_{i=0}^{n-1} a_i q^i \bmod (q^n - 1).$$

where a_i is the number of orbits whose stabilizers have cardinalities dividing i.

As the stabilizer of an orbit of length ℓ is $\{1, q^\ell, q^{2\ell}, \ldots, q^{(n/\ell-1)\ell}\}$, the cardinality of the stabilizer is n/ℓ. In Example 2.2.1, we have two orbits $(1100, 0110, 0011, 1001)$ and $(1010, 0101)$ for C_R on $\binom{[4]}{2}$. The corresponding cyclic sieving polynomial is

$$\binom{4}{2}_q = 1 + q + 2q^2 + q^3 + q^4 \equiv 2 + q + 2q^2 + q^3 \bmod \left(q^4 - 1\right).$$

The two orbits of length 4 and 2 have stabilizer orders 1 and 2, respectively. Both of these orders divide 0 and 2, while only 1 divides 1 and 3. Thus $2 + q + 2q^2 + q^3 \bmod (q^4 - 1)$ is consistent with Proposition 2.2.14. It is easy to see even without Proposition 2.2.14 that two different cyclic sieving polynomials for a cyclic action of order n have to be congruent mod $(q^n - 1)$ since the difference between the two polynomials must be zero upon input of an n^{th} root of unity.

As an example for how to construct a polynomial $X(q)$ to satisfy the cyclic sieving polynomial, we will later see an action[6] that has orbits of sizes 4, 22, 46, and 60. Then the order of this map is $\operatorname{lcm}(4, 22, 46, 60) = 15180$. The orbits have stabilizer cardinalities $\frac{15180}{4} = 3795$, $\frac{15180}{22} = 690$, $\frac{15180}{46} = 330$, and $\frac{15180}{60} = 253$, respectively. Using Proposition 2.2.14, any polynomial

$$\sum_{k=0}^{3} q^{3795k} + \sum_{k=0}^{21} q^{690k} + \sum_{k=0}^{45} q^{330k} + \sum_{k=0}^{59} q^{253k} \bmod \left(q^{15180} - 1\right)$$

[6] It is mentioned shortly after the proof of Corollary 5.1.8.

gives the CSP.

While we get the CSP for any action of a cyclic group on a finite set, it is not interesting if we can only state the polynomial through listing all of the orbit sizes. What makes the cyclic sieving phenomenon phenomenal is that we often get results for infinite families, like in Theorems 2.2.4 (which has parameters n and k) and 2.2.11 (which has parameter n).

For a while, interesting homomesies were conjectured for maps in which there likely is no natural CSP. That is the context in which many problems to be discussed in Chapters 4 through 7 originated. We discuss several instances of homomesy for actions in which the order of the map is less natural and in some cases unknown in general. In situations where the order of the map is quite large in relation to the size of the ground set, we believe a natural CSP seems unlikely, since any polynomial that would give the CSP have large degree.

Chapter 3

Toggle groups

This chapter provides background on toggle groups. Many of our results in the upcoming chapters are for maps expressed as products of toggles.

Definition 3.0.15 ([50])**.** Let E be a set and $\mathcal{L} \subseteq 2^E$ a set of "allowed" subsets of E. Then to every $e \in E$, we define its **toggle** $t_e : \mathcal{L} \to \mathcal{L}$ as

$$t_e(X) = \begin{cases} X \cup \{e\} & \text{if } e \notin X \text{ and } X \cup \{e\} \in \mathcal{L}, \\ X \setminus \{e\} & \text{if } e \in X \text{ and } X \setminus \{e\} \in \mathcal{L}, \\ X & \text{otherwise.} \end{cases}$$

The **toggle group** of $\mathcal{L}$ is the subgroup of the symmetric group $\mathfrak{S}_{\mathcal{L}}$ on $\mathcal{L}$ generated by $\{t_e : e \in E\}$.

Informally, t_e adds or removes e assuming the output is still in $\mathcal{L}$, and otherwise does nothing. While each t_e is an involution (map of order 2). the composition of toggles produces maps whose order is difficult to analyze in general.

3.1 Toggling order ideals of posets

The toggle group was originally introduced by Cameron and Fon-der-Flaass in the setting of order ideals of a poset [12].

Definition 3.1.1. For a poset P, let $\mathrm{Tog}(P)$ denote the toggle group of $J(P)$.

We will only be concerned with finite posets. The rowmotion operation ρ_J on order ideals from Subsection 2.1.4 can be expressed as an element of $\mathrm{Tog}(P)$.

Definition 3.1.2. A sequence $(x_1, x_2, \ldots, x_n)$ containing all of the elements of a finite poset P exactly once is called a **linear extension** of P if it is order-preserving, that is, if $x_i < x_j$ in P then $i < j$.

Proposition 3.1.3 ([12])**.** *Let $(x_1, x_2, \ldots, x_n)$ be any linear extension of a finite poset P. Then $\rho_J = t_{x_1} t_{x_2} \cdots t_{x_n}$.*

We use the convention that a product of toggles is performed right to left. To prove this, we first prove a basic lemma.

Lemma 3.1.4. *Let $I \in J(P)$ and $x \in P$. Then*

$$t_x(I) = \begin{cases} I \cup \{x\} & \text{if } x \text{ is a minimal element of } P \setminus I, \\ I \setminus \{x\} & \text{if } x \text{ is a maximal element of } I, \\ X & \text{otherwise.} \end{cases}$$

$$\begin{array}{ccc} \mathcal{A}(P) & \xrightarrow{\rho_J} & \mathcal{A}(P) \\ \mathbf{I}\big\downarrow & & \big\downarrow\mathbf{I} \\ J(P) & \xrightarrow[\rho_{\mathcal{A}}]{} & J(P) \end{array}$$

Proof. If $x \in I$, then we can remove x from I and still be left with an order ideal if and only if x is a maximal element of I. If $x \notin I$, then we can add x to I and still have an order ideal if and only if everything less than x in P is in I (which is equivalent to x being a minimal element of the complement $P \setminus I$). So this is clearly equivalent to the definition of the toggle t_x. ∎

Proof of Proposition 3.1.3. Let $I \in J(P)$ and $T = t_{x_1} t_{x_2} \cdots t_{x_n}$. Recall that $\rho_J(I)$ is the order ideal generated by the minimal elements not in I. Since $(x_1, x_2, \ldots, x_n)$ is a linear extension, when applying T, we apply t_y before t_x for any pair $x < y$. Thus, from Lemma 3.1.4, the elements inserted by T are precisely the minimal elements not in I. For every $x \in I$, t_x removes x from I unless $x < y$ for some y inserted by T. Thus, $T(I)$ is the order ideal generated by the minimal elements not in I. ∎

Example 3.1.5. Below, we show the effect of applying $t_1 t_2 t_3 t_4 t_5 t_6$ to an order ideal of the root poset $\Phi^+(A_3)$ described in Subsection 2.1.4, where $1, 2, 3, 4, 5, 6$ are the labelings of the poset elements in the diagram. Since $(1, 2, 3, 4, 5, 6)$ is a linear extension of the poset, this composition of toggles is equivalent to rowmotion. Notice that we get the same result that we obtained when applying ρ_J to the same order ideal in Example 2.1.17.

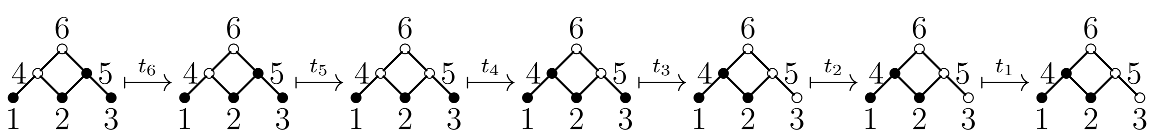

Proposition 3.1.6 ([51])**.** *Two toggles $t_x, t_y \in \operatorname{Tog}(P)$ commute if and only if neither one of x nor y covers the other.*

Proof. **Case 1: x and y are incomparable.** Then whether or not one of x or y can be in an order ideal has no effect on whether the other can so $t_x t_y = t_y t_x$.

Case 2: $x < y$ or $y < x$ but neither one covers the other. Without loss of generality, assume $x < y$. Since y does not cover x, there exists $z \in P$ such that $x < z < y$. Then an order ideal containing y must contain z, and one that does not contain x cannot contain z. Thus, we cannot change whether or not x is in an order ideal and then do the same for y, or vice versa, without changing the status of z. Thus, given any order ideal I, either $t_x(t_y(I)) = t_y(t_x(I)) = I$, or $t_x(t_y(I)) = t_y(t_x(I)) = I\Delta\{x\}$, or $t_x(t_y(I)) = t_y(t_x(I)) = I\Delta\{y\}$ where Δ denotes the symmetric difference operation.

Case 3: either $x \lessdot y$ or $y \lessdot x$. Without loss of generality, assume $x \lessdot y$. Let $I = \{z \in P | z < y\}$ which is an order ideal that has x as a maximal element. So when applying t_x then t_y to I, we first remove x and then cannot add y. On the other hand, when applying t_y then t_x to I, we first add y and then cannot remove x. Since $x \notin (t_y(t_x(I))$ but $x \in (t_x(t_y(I))$, $t_x t_y \neq t_y t_x$. ■

From Proposition 3.1.6, it is clear that when a poset can be placed neatly into "rows" of antichains like $\Phi^+(A_n)$ can, we can apply all toggles within a row simultaneously. In this case, rowmotion applies the toggle maps row by row; hence the name. Graded posets can be placed into rows, given by rank levels.

Definition 3.1.7. We say a poset P is **graded** if it has a well-defined **rank function** $\text{rank} : P \to \mathbb{N}$ satisfying

- $\text{rank}(x) = 0$ for any minimal element x,
- $\text{rank}(y) = \text{rank}(x) + 1$ if $y \gtrdot x$,
- every maximal element x has $\text{rank}(x) = r$, where r is called the **rank** of P.

Let Π be the set of ordered pairs $(m, n) \in \mathbb{Z}^2$ such that m and n have the same parity. We call a finite poset P a **rowed-and-columned poset** (or **rc-poset**) if there exists a

"position" function $\pi : P \to \Pi$ such that whenever $x \lessdot y$ in P and $\pi(x) = (i, j)$ either $\pi(y) = (i-1, j+1)$ or $\pi(y) = (i+1, j+1)$ [51].

Proposition 3.1.8. *For a graded poset of rank r,*

$$\rho_J^i(\emptyset) = \{x \in P |\, \mathrm{rank}(x) \leq i - 1\}$$

for $i \in [r+1]$ and $\rho_J^{r+2}(\emptyset) = \emptyset$.

Proof. Since $\rho_J(I)$ is the order ideal generated by the minimal elements of $P \setminus I$, we have $\rho_J(\emptyset)$ is the set of minimal elements of P, i.e., $\rho_J(\emptyset) = \{x \in P |\, \mathrm{rank}(x) = 0\} = \{x \in P |\, \mathrm{rank}(x) \leq 0\}$. When $i \in [r]$ and

$$I = \{x \in P |\, \mathrm{rank}(x) \leq i - 1\},$$

the minimal elements of $P \notin I$ are the elements of rank i, so

$$\rho_J(I) = \{x \in P |\, \mathrm{rank}(x) \leq i\}.$$

So we inductively have

$$\rho_J^i(\emptyset) = \{x \in P |\, \mathrm{rank}(x) \leq i - 1\}$$

for $i \in [r+1]$. Thus, $\rho_J^{r+1}(\emptyset) = P$. Clearly, there are no minimal elements of $P \setminus P = \emptyset$, so $\rho_J(P) = \emptyset$. Thus, $\rho_J^{r+2}(\emptyset) = \emptyset$. ■

Recall if A is the antichain of maximal elements of an order ideal I, then $\rho_{\mathcal{A}}(A)$ is the antichain of maximal elements of $\rho_J(I)$, so Proposition 3.1.8 gives the following corollary.

Corollary 3.1.9. *For a graded poset of rank r,*

$$\rho_{\mathcal{A}}^{i}(\emptyset) = \{x \in P \mid \operatorname{rank}(x) = i - 1\}$$

for $i \in [r+1]$ and $\rho_{\mathcal{A}}^{r+2}(\emptyset) = \emptyset$.

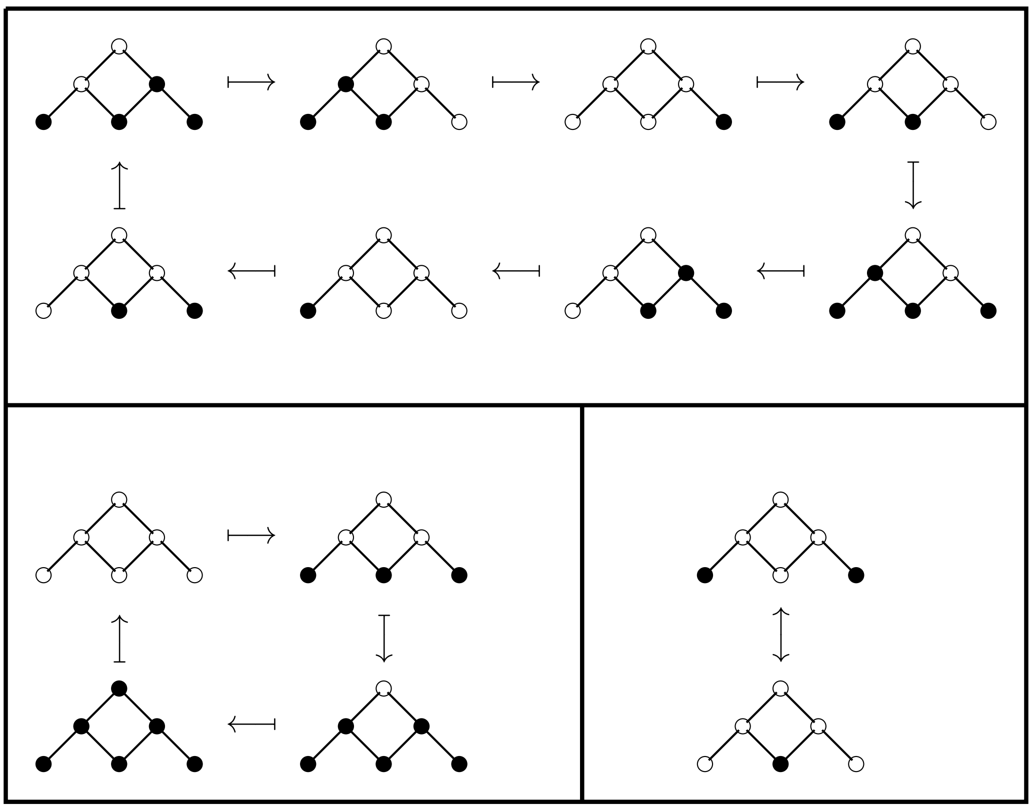

FIGURE 3.1.1: The three orbits of ρ_J on $J(\Phi^+(A_3))$. These are the order ideals generated by the antichains in the orbits of Figure 2.1.8. The average cardinality is $5/2$ for the orbit of size 8, $7/2$ for the orbit of length 4, and $3/2$ for the orbit of length 2. While cardinality is not homomesic, the rank-alternating cardinality is $3/2$-mesic. See Theorem 3.1.10.

Unlike rowmotion for antichains, the cardinality statistic is not homomesic under ρ_J on $J(\Phi^+(A_n))$. However, Haddadan discovered a similar homomesic statistic in terms of the rank of elements.

Theorem 3.1.10 ([21, Corollary 36])**.** *For an order ideal $I \in J(\Phi^+(A_n))$, define the rank-alternating cardinality of I, to be $f(I) = \sum\limits_{x \in I} (-1)^{\text{rank}(x)}$. Then the triple $(J(\Phi^+(A_n)), \rho_J, f)$ exhibits homomesy with average $n/2$.*

See Figure 3.1.1 for the three orbits of ρ_J on $J(\Phi^+(A_3))$. The average cardinalities in the orbits are not equal: $5/2$, $7/2$, and $3/2$, but each orbit has average rank-alternating cardinality $3/2$. The orbit of length four is the one described by Proposition 3.1.8.

Neither of the two posets below are graded. For example, if there is a rank function for the one on the left, then rank(a) would be 0, so $a \lessdot b \lessdot e$ would imply rank$(e) = 2$ but $a \lessdot c \lessdot d \lessdot e$ would imply rank$(e) = 3$. However, the one on the right is an rc-poset, where the coordinates shown describe a position function π.

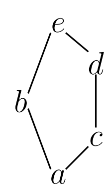

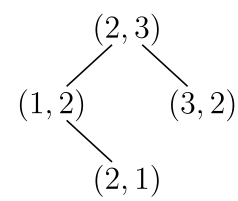

For an rc-poset P with position function π, if we place each $x \in P$ in position $\pi(x)$, then we have a Hasse diagram into antichain rows. So rowmotion can be described as applying all toggles in each row simultaneously starting from the top row down, similar to graded posets. However, rc-posets also allow us to define a similar product of toggles, called promotion, by toggling one column at a time.

Definition 3.1.11. For an rc-poset P with position function π, assume π only outputs positions with first coordinates in the range $[k]$, for some $k \in \mathbb{P}$. This can be assumed without loss of generality. Let c_i be the product of all t_a for which a is in column i (i.e., $\pi(a)$ has first coordinate i). Since all the cover relations go diagonally from one column to

another, c_i is well-defined by Proposition 3.1.6. Then **promotion** Pro : $J(P) \to J(P)$ is defined by $\text{Pro} = c_k \cdots c_2 c_1$.

Note that the definition of Pro not only depends on P but also the position function π. Rc-posets have multiple position functions that meet the criteria, so we always assume a fixed one to define Pro. Like how rowmotion can be considered in multiple settings (such as on antichains), promotion can also. Promotion was originally considered by Schutzënberger as an action on standard Young tableaux [39], which is where the name "promotion" comes from. He then considered it as an action on linear extensions of posets [40]. We will only consider it in the context of order ideals or rc-posets. Generalizations of rowmotion and promotion to allow toggling rows or columns in any specified order is also discussed in [51], but we will not discuss those generalizations.

There have also been numerous natural homomesic statistics discovered for rowmotion and promotion on a product of two chains.

Definition 3.1.12. The poset structure on $[a] \times [b]$ is defined so that $(x_1, y_1) \leq (x_2, y_2)$ if and only if $x_1 \leq x_2$ and $y_1 \leq y_2$.

See Figure 3.1.2 for the Hasse diagram of $[3] \times [5]$. Even though $[a] \times [b]$ is an rc-poset, note that the coordinates do not give a position function as we have to rotate the axes 45 degrees for the Hasse diagram. The rank of (x, y) is $x + y - 2$ and a column consists of all (x, y) for which $y - x$ is the same. Propp and Roby proved the following homomesies for $[a] \times [b]$ in [32].

Theorem 3.1.13. *Let $a \times b \in \mathbb{P}$.*

1. *Under ρ_J on $J([a] \times [b])$, cardinality is homomesic with average $ab/2$.*

2. *Under $\rho_{\mathcal{A}}$ on $\mathcal{A}([a] \times [b])$, cardinality is homomesic with average $ab/(a + b)$.*

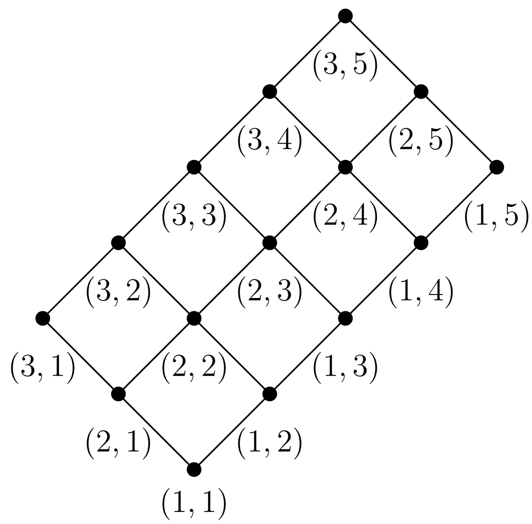

FIGURE 3.1.2: The poset $[3] \times [5]$.

3. *Under* Pro *on* $J([a] \times [b])$*, cardinality is homomesic with average* $ab/2$*.*

The order of any of the maps described in Theorem 3.1.13 is $a + b$ [9], and in fact the set of orbit sizes are precisely the numbers of the form $\frac{a+b}{d}$ for any $d \geq 1$ that divides both a and b [16].

In Section 3.3, we see that we can use the conjugacy of rowmotion and promotion in the toggle group, established in [51, §5] to show that parts 1 and 3 of Theorem 3.1.13 imply each other. More specifically, the theory explains that the set of homomesic statistics which are linear combinations of element indicator functions are the same for ρ_J as for Pro on $J(P)$ for any poset P. So Theorem 3.1.10 implies we also have homomesy for the rank-alternating cardinality under Pro on $J(\Phi^+(A_n))$.

Note that this only applies to a certain family of statistics, as discussed in Remark 3.3.10. Theorem 3.1.13(2) implies that under the action of ρ_J on $J([a] \times [b])$, the number of maximal elements is homomesic with average $ab/(a + b)$; this statistic is *not* homomesic under the action of Pro [32, Example 21].

For certain “nice” posets P, homomesies for toggling $J(P)$ can be extended to generalizations: piecewise-linear and birational toggling. Einstein and Propp [14, Theorems 2 and 4] prove generalizations of the homomesies in Theorem 3.1.13 to these settings in which the order of the map is still finite but the ground set is uncountable.

3.2 Generalized toggle groups

While toggle groups have been studied for order ideals of posets for quite some time, Striker has recently made the case for studying toggle groups over other families of subsets [50]. In this thesis, we analyze actions within toggle groups on the following families of subsets:

- independent sets of graphs (Chapter 4 and Section 5.3),
- noncrossing partitions as sets of “arcs” (Chapter 5),
- subsets within a given cardinality range (Section 6.6).

In Section 4.4, we also analyze rowmotion and promotion on order ideals of zigzag posets, and the homomesies for that map. However, we establish our results by first working in toggle groups for independent sets of path graphs. There is an equivariant bijection between independent sets of a path graph and order ideals of a zigzag poset which allows us to restate our results in the poset setting, but it is much easier and more natural to work on the path graph. This displays the usefulness of Striker’s abstraction to settings beyond posets.

The Kreweras complement κ on noncrossing partitions described in Subsection 2.2.1 can actually be expressed as a product of toggles for linear representations of noncrossing partitions. In Chapter 5, we use Coxeter group theory extend the block count homomesy that is easy to prove for κ, to a large class of actions within the toggle group. Unlike κ, we

have unpredictable orbit sizes in the general setting and it is highly nontrivial to prove the homomesy.

The following describes the order of a product of two toggles within any toggle group.

Proposition 3.2.1. *Let E be a set and $\mathcal{L} \subseteq 2^E$ be the set of allowed subsets. Then for any pair $x, y \in E$, any orbit under the action of $t_x t_y$ has size 1, 2, or 3. So the order of $t_x t_y$ divides 6.*

Proof. Fix $x \neq y$. Since $t_x t_y$ cannot change the status of any elements except x and y, any orbit of $t_x t_y$ can only contain X, $X \cup \{x\}$, $X \cup \{y\}$, and $X \cup \{x, y\}$ for some $X \subseteq 2^E$ that neither contains x nor y. Thus, the size of any orbit cannot be more than 4. If there were an orbit of size 4, then X, $X \cup \{x\}$, $X \cup \{y\}$, and $X \cup \{x, y\}$ must all be in $\mathcal{L}$. However, in this case $X \xmapsto{t_x t_y} X \cup \{x, y\} \xmapsto{t_x t_y} X$ and $X \cup \{x\} \xmapsto{t_x t_y} X \cup \{y\} \xmapsto{t_x t_y} X \cup \{x\}$ so we do not get an orbit of size 4. Thus, no orbit has size greater than 3. ∎

3.3 Coxeter group theory for toggle groups

This section provides a brief summary of material in [13, §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.

Let E be a *finite* set with allowed subsets $\mathcal{L} \subseteq 2^E$. Let $\mathcal{T}$ be the toggle group of $\mathcal{L}$. Since $\mathcal{T}$ is generated by finitely many involutions, it is the quotient of a toggle group. See [7] and [30, Ch. 11–14] to learn about Coxeter groups and their connections to combinatorics.

Definition 3.3.1. A **Coxeter system** is a pair (W, S) consisting of a group W (called a **Coxeter group**) generated by a set $S = \{s_1, \ldots, s_r\}$ of involutions with presentation $W =$

$\langle s_1, \ldots, s_r | s_i^2 = (s_i s_j)^{m_{ij}} = 1 \rangle$, where $m(s_i, s_j) := m_{ij} \geq 2$ for $i \neq j$. It is also permissible for m_{ij} to be infinity, meaning $s_i s_j$ has infinite order.

That is, a Coxeter group is defined by a set S of involution generators together with relations that specify the order of any product of two given generators $s_i s_j$. In the case of toggle groups, the nontrivial[1] toggles are the generators and the product of two distinct generators has order 2, 3, or 6 by Proposition 3.2.1. Toggle groups usually have extra relations not described by the Coxeter relations, such as a restriction on the order of the product of three generators not already implied by the other relations. A toggle group on a finite set clearly must be finite. As all finite Coxeter groups have been classified, we can in many settings determine that the toggle group is not a true Coxeter group but rather a quotient of one.

Knowledge of Coxeter groups is not necessary to understand the rest of the paper, as we explain the theory we use. For a Coxeter system (W, S), a Coxeter element is a product of every generator in S each used exactly once. We borrow this term and extend it to toggle groups.

Definition 3.3.2. An element $w \in \mathcal{T}$ is called a **Coxeter element** if it is a product of every nontrivial toggle map, each used exactly once, in some order.

Many actions within toggle groups for which we have discovered homomesy are Coxeter elements. Most of the toggling actions we will discuss in the upcoming chapters, including rowmotion and promotion, are Coxeter elements.

Definition 3.3.3. For a Coxeter system (W, S), the **Coxeter graph** $\Gamma(W, S)$ is the undirected graph with vertex set S and edges connecting s_i and s_j whenever $m_{ij} \geq 3$.

[1]It is possible to have $x \in E$ for which either all subsets in $\mathcal{L}$ contain x or none contain x. In this case t_x is the trivial map and can be ignored as a generator of the toggle group.

Similarly, the **Coxeter graph** $\Gamma(\mathcal{T})$ of a toggle group $\mathcal{T}$ is the undirected graph with vertices x for every $x \in E$ whose toggle t_x is nontrivial, and edges connecting x and y whenever t_x and t_y do not commute.

For $\mathrm{Tog}(P)$, the toggle group for order ideals of a poset P, the Hasse diagram of P is the Coxeter graph by Proposition 3.1.6.

For a toggle group $\mathcal{T}$ and Coxeter element $w \in \mathcal{T}$, we define an acyclic orientation of $\Gamma(\mathcal{T})$ by orienting any edge between x and y in the direction of x if t_x appears to the right of t_y in an expression for w, and in the direction of y if t_y appears to the right of t_x in an expression for w.

Example 3.3.4. For the poset $\Phi^+(A_3)$ shown at left, rowmotion $\rho_J = t_1t_2t_3t_4t_5t_6$ is described by the graph orientation in the middle. Promotion $\mathrm{Pro} = t_3t_5t_2t_6t_4t_1$ is described by the graph orientation on the right.

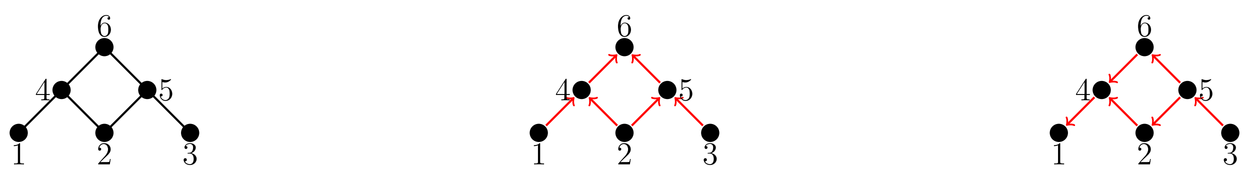

$t_1t_2t_3t_4t_5t_6$ $t_3t_5t_2t_6t_4t_1$

While there are multiple equivalent expressions for these maps as a product of every toggle each used exactly once, they are all formed by applying the commutativity relations. For example, $\rho_J = t_2t_3t_5t_1t_4t_6$, formed from $t_1t_2t_3t_4t_5t_6$ by commuting t_5 and t_4 and then moving t_1 past t_2, t_3, and t_5, all of which it commutes with. Since this commutation does not swap any pair connected by an edge in the Coxeter graph, the graph's orientation is unchanged. Thus, any acyclic orientation of $\Gamma(\mathcal{T})$ corresponds uniquely to the Coxeter element, and this orientation is well-defined. In fact, Shi showed that Coxeter elements correspond uniquely to acyclic orientations of the graph associated with a Coxeter group [41, Proposition 1.3].

For a Coxeter element $w \in \mathcal{T}$, toggles 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 t_x of w corresponds to moving it to the left (resp. right) of an expression for w. For example, t_3 is initial in $w = t_1t_2t_3t_4t_5t_6$ so conjugating by t_3 gives

$$\begin{aligned} t_3wt_3 &= t_3(t_1t_2t_3t_4t_5t_6)t_3 \\ &= t_3(t_3t_1t_2t_4t_5t_6)t_3 \\ &= t_1t_2t_4t_5t_6t_3. \end{aligned}$$

Conjugation by a final element of w corresponds to changing a sink into a source in the corresponding orientation of $\Gamma(\mathcal{T})$. Conjugation by an initial element corresponds to changing a source into a sink. If we were to conjugate w by t_4, which is neither initial nor final, we would get $t_4t_3t_1t_2t_4t_5t_6t_4$ which is likely not a Coxeter element as probably the relations do not allow us to rewrite this with t_4 only appearing once. Since we are only working over the *quotient* of a Coxeter group, we do not know all the relations so we cannot be entirely sure.

H. and K. Eriksson showed that two Coxeter elements in a Coxeter group 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 [15]. This corresponds to a sequence of conjugations by generators that are either initial or final; we call such conjugations **admissible**. In the

quotient of a Coxeter group, it is still the case that if the Coxeter graphs for two elements differ by a sequence that changes sinks into sources or vice versa, then they are conjugate, but the converse is not necessarily true. It may be possible for two elements to be conjugate for non-Coxeter-theoretic reasons.

Striker and Williams proved that for a poset P, ρ_J and Pro are conjugate elements in $\mathrm{Tog}(P)$ (and that this extends to their generalizations of rowmotion and promotion) [51, §5]. Their proof in fact shows that they are conjugate via a sequence of *admissible* conjugations.

Example 3.3.5. In $\mathrm{Tog}(\Phi^+(A_3))$, we can get from $\rho = t_1t_2t_3t_4t_5t_6$ to $\mathrm{Pro} = t_3t_5t_2t_6t_4t_1$ by conjugating by t_1 then t_2 then t_4 then t_1. All of these are conjugations by initial elements, and we have $\mathrm{Pro} = (t_1t_4t_2t_1)\rho(t_1t_4t_2t_1)^{-1}$.

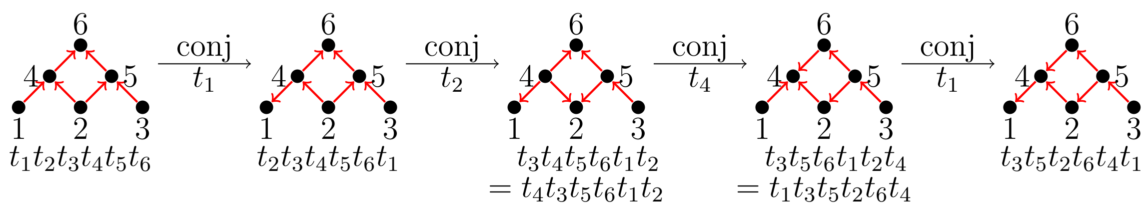

We also define a class of toggle group elements that we will use in Chapter 5.

Definition 3.3.6. An element $w \in \mathcal{T}$ is called a **partial Coxeter element** if it is a product of toggle maps, each used at most once.

Obviously a Coxeter element is also a partial Coxeter element but not vice versa. Any partial Coxeter element containing the toggles for elements $e_1, e_2, \ldots, e_n \in E$ corresponds to an acyclic orientation of the subgraph of $\Gamma(\mathcal{T})$ containing $e_1, e_2, \ldots, e_n$ in a manner analogous to that of Coxeter elements. If there is an edge between e_i and e_j in this subgraph, then we orient it in the direction of whichever of e_i or e_j occurs later in an expression of w.

Example 3.3.7. On $\mathrm{Tog}(\Phi^+(A_3))$ with the elements labeled above, a partial Coxeter element is $w = t_1t_6t_3t_4$. So w corresponds to the orientation

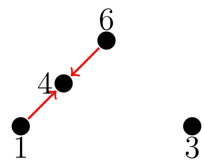

of the subgraph of the Hasse diagram of this poset containing only $\{1, 3, 4, 6\}$. Note that since 3 is an isolated vertex, t_3 commutes with all of the other toggles in w.

The following definitions are the same as for Coxeter elements. We call t_x **initial** (resp. **final**) in a partial Coxeter element w if it can be brought to the left (resp. right) by the commutativity relations. Initial and final toggles correspond to source and sink vertices, respectively, in the corresponding acyclic subgraph orientations. We call a conjugation by an initial or final toggle **admissible** and this corresponds to changing a source into a sink or vice versa in the subgraph orientation.

Definition 3.3.8. For a set E and element $e \in E$, the **indicator function** $I_e : 2^E \to \{0, 1\}$ of e is given by

$$I_e(X) = \begin{cases} 1 & \text{if } e \in X, \\ 0 & \text{if } e \notin X. \end{cases}$$

For a subset $\mathcal{L} = 2^E$ over which a toggle group is defined, we often consider the restriction of I_e to $\mathcal{L}$. The following explains how we can often extend homomesies known for one (partial) Coxeter element to a broader class.

Theorem 3.3.9. *Let $\mathcal{T}$ be the toggle group for some $\mathcal{L} \subseteq 2^E$. Let $w, w' \in \mathcal{T}$ be partial Coxeter elements for which we can transform w into w' via a sequence of admissible conjugations. Let $f : \mathcal{L} \to \mathbb{K}$ be a statistic which is a $\mathbb{K}$-linear combination of indicator functions. Then $(\mathcal{L}, w, f)$ exhibits homomesy with average c if and only if $(\mathcal{L}, w', f)$ does.*

Proof. It suffices to show that these homomesies are preserved under an admissible conjugation by a single toggle, since the case for a sequence of toggles will follow inductively. That is, without loss of generality assume $w' = t_x w t_x$ where t_x is initial or final in w. It suffices to show that any statistic which is a linear combination of the indicator functions I_e is c-mesic under the action of w if and only if it is c-mesic under the action of w'.

We have the following commutative diagram.

$$\begin{tikzcd} X \arrow[r, "w"] \arrow[d, "t_x"'] & w(X) \arrow[d, "t_x"] \\ t_x(X) \arrow[r, "w'"'] & t_x(w(X)) \end{tikzcd}$$

Write a w-orbit $\mathcal{O} = (X_1, X_2, \ldots, X_\ell)$ where for every $i \in [\ell - 1]$, $w(X_i) = X_{i+1}$ and $w(X_\ell) = X_1$. From the commutative diagram above, $\mathcal{O}' = (t_x(X_1), t_x(X_2), \ldots, t_x(X_\ell))$ is a w'-orbit. This creates a bijection between w-orbits and w'-orbits that preserves the orbit sizes. To avoid special cases throughout this proof, we consider the subscripts mod ℓ, so in particular $X_0 = X_\ell$ and $X_{\ell+1} = X_1$.

When $e \neq x$, $I_e(X_i) = I_e(t_x(X_i))$ as I_x is the only indicator function that t_x can change. So

$$\sum_{X \in \mathcal{O}} I_e(X) = \sum_{X' \in \mathcal{O}'} I_e(X').$$

If t_x is final in w, then it is the first toggle applied when performing w. Also t_x only appears once when applying w, since w is a partial Coxeter element. Thus, in this scenario

$I_x(t_x(X_i)) = I_x(X_{i+1})$. So

$$\begin{aligned}\sum_{X' \in \mathcal{O}'} I_x(X') &= I_x(X_2) + I_x(X_3) + \cdots + I_x(X_\ell) + I_x(X_{\ell+1}) \\ &= I_x(X_2) + I_x(X_3) + \cdots + I_x(X_\ell) + I_x(X_1) \\ &= \sum_{X \in \mathcal{O}} I_x(X).\end{aligned}$$

Analogously, if t_x is initial in w, then $I_x(t_x(X_i)) = I_x(X_{i-1})$. In this scenario,

$$\begin{aligned}\sum_{X' \in \mathcal{O}'} I_x(X') &= I_x(X_0) + I_x(X_1) + I_x(X_2) + \cdots + I_x(X_{\ell-1}) \\ &= I_x(X_\ell) + I_x(X_1) + I_x(X_2) + \cdots + I_x(X_{\ell-1}) \\ &= \sum_{X \in \mathcal{O}} I_x(X).\end{aligned}$$

Thus the sum of any I_e is the same in $\mathcal{O}$ and $\mathcal{O}'$. As these orbits have the same size, the average of any I_e is also the same across them. This of course extends to linear combinations.

■

Theorem 3.3.9 explains that it suffices to prove either part 1 or 3 of Theorem 3.1.13 to obtain the other, since cardinality is the sum of all of the indicator functions. Also, the rank-alternating cardinality is a linear combination of indicator functions, so the rowmotion homomesy of Theorem 3.1.10 also extends to promotion. See Figure 3.3.1 for the promotion orbits on $J(\Phi^+(A_3))$. By the conjugation described in Example 3.3.5, for any ρ_J orbit $\mathcal{O}$ on this set, we have a corresponding promotion orbit $\{t_1t_4t_2t_1(I) | I \in \mathcal{O}\}$.

Remark 3.3.10. In general, conjugation does not preserve homomesy of statistics under the respective maps. Theorem 3.3.9 describes only a specific scenario in which homomesy is preserved, namely when the conjugation is admissible and the statistic is a linear combination

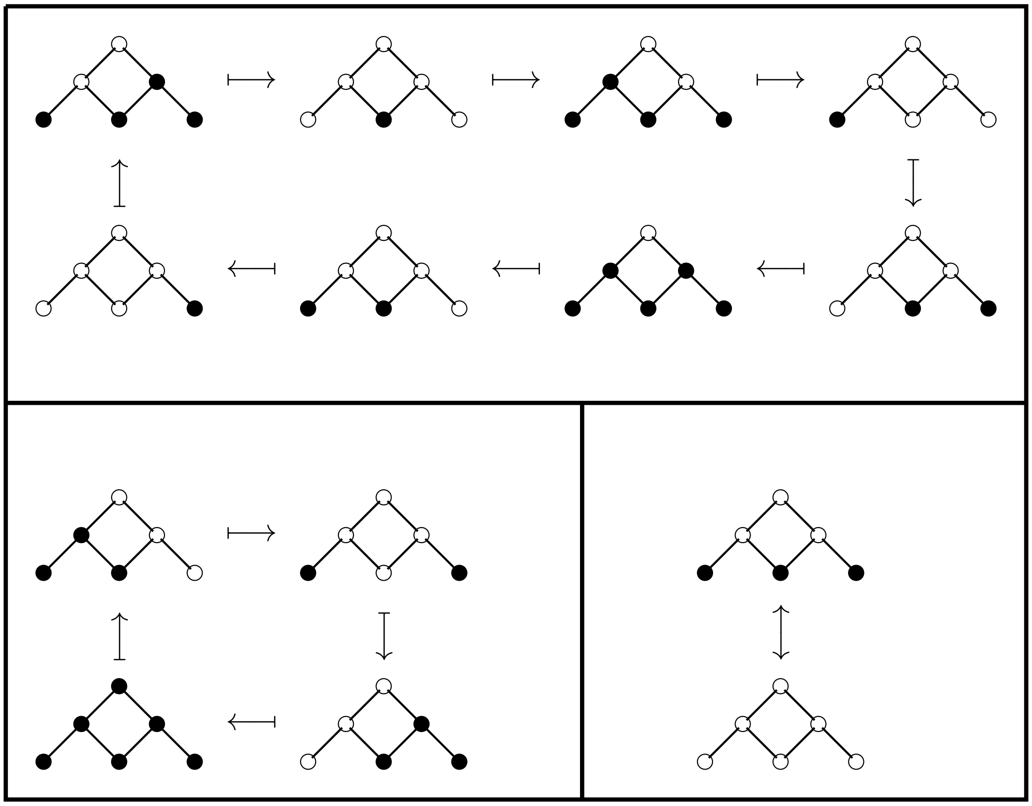

FIGURE 3.3.1: The three orbits of Pro on $J(\Phi^+(A_3))$. Via the conjugation in Example 3.3.5, each of these orbits is $(t_1t_4t_2t_1)$ applied to the order ideals in the corresponding rowmotion orbits, shown in Figure 3.1.1. Again the rank-alternating cardinality is 3/2-mesic, as a consequence of Theorems 3.1.10 and 3.3.9. On the other hand, for ρ_J orbits, the number of maximal elements is 3/2-mesic, a consequence of Theorem 2.1.18. This is not the case for these orbits as the average number of maximal elements across the orbits are 13/8, 5/4, 3/2. See Remark 3.3.10.

of the indicator functions. Theorems 2.1.18 and 3.1.13(2) imply two scenarios for which the number of maximal elements is homomesic under ρ_J.[2] However these do not extend to Pro. Notice in Figure 3.1.1 that within every ρ_J-orbit, the average number of maximal elements is 3/2, but in the Pro-orbits shown in Figure 3.3.1, the average differs from orbit to orbit.

[2]Recall that the $\rho_{\mathcal{A}}$-orbits consist of the sets of maximal elements of the order ideals in the ρ_J-orbits.

Chapter 4

Toggling independent sets of a path graph

This chapter is joint work with Tom Roby. Most of this exposition can be found in [25] with small modifications and additions here.

4.1 Toggles and homomesy

We explore the orbit structure and homomesy properties of certain actions of toggle groups on the collection of independent sets of a path graph. In particular we prove a conjecture of Propp that with respect to the action of a particular Coxeter element of vertex toggles, the difference of indicator functions of symmetrically-located vertices is 0-mesic (Theorem 4.1.13). Using the theory of Section 3.3, we generalize this result to apply to any Coxeter element in the toggle group. We then use our analysis to show facts about orbit sizes that are easy to conjecture but nontrivial to prove.

This provides an interesting example of homomesy in a context where unwieldy orbit

sizes make a natural CSP unlikely to exist. We determine the orbit sizes in Section 4.3, but the order of the map is not easy to state without listing all the orbit sizes.

This problem also displays the usefulness of Striker's notion of generalized toggle groups [50] to settings beyond that of posets. Although there is an equivariant bijection (Proposition 4.4.4) 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.

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

Definition 4.1.1. Let $\mathcal{P}_n$ denote the **path graph** with vertex set $[n]$ and edge set $\{\{i, i+1\} : i \in [n-1]\}$. Call two vertices **adjacent** if they are connected by an edge.

Example 4.1.2. The path graph with seven vertices is

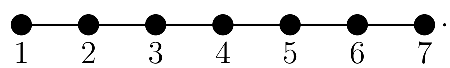

Definition 4.1.3. An **independent set** of a graph is a subset of the vertices that does not contain an adjacent pair. Let $\mathcal{I}_n$ denote the set of independent sets of $\mathcal{P}_n$.

Example 4.1.4. The set of vertices $\{1, 4, 6\}$ represented as

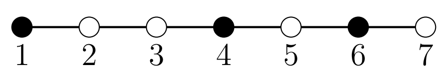

is an independent set of $\mathcal{P}_7$, but $\{1, 4, 5, 6\}$ represented as

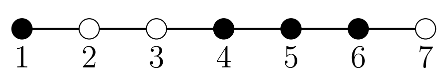

is not. In both of these examples, hollow dots refer to vertices of $\mathcal{P}_7$ not in the subset.

Although we sometimes write independent sets as subsets of $[n]$ 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). The following enumerative result is well-known.

Proposition 4.1.5. *The number of independent sets in $\mathcal{I}_n$ is F_{n+2}, the $(n+2)^{nd}$ Fibonacci number, with $F_0 = 0$, $F_1 = 1$ and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$.*

Proof. When $n = 1$, there are two independent sets: 0 and 1. When $n = 2$, there are three independent sets: 00, 01, and 10. So the formula holds for $n = 1$ and $n = 2$. Let $n \geq 3$ and let $S \in \mathcal{I}_n$. If $n \notin S$, then there are no restrictions for the other $n-1$ vertices, so we can choose any set in $\mathcal{I}_{n-1}$ and attach 0 on at the end. If $n \in S$, then $n - 1 \notin S$, so we can choose any set in $\mathcal{I}_{n-2}$ and attach 01 on at the end. By induction $\#\mathcal{I}_n = \#\mathcal{I}_{n-1} + \#\mathcal{I}_{n-2} = F_{n+1} + F_n = F_{n+2}$. ■

In this section we state and prove our main homomesy results. Throughout this chapter 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.

4.1.1 Definitions and main results

We consider toggles on $\mathcal{P}_n$, with allowed subsets given by $\mathcal{I}_n$, as in Chapter 3. As we will later refer to toggles on zigzag posets and relate the two, we use τ_i instead of t_i to denote the toggle at vertex i.

Definition 4.1.6. For every $i \in [n]$, $\tau_i : \mathcal{I}_n \to \mathcal{I}_n$ is the toggle at vertex i. 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,

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

Let $\mathcal{T}_n$ denote the toggle group on $\mathcal{I}_n$.

Proposition 4.1.7. *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 $\tau_i \tau_j \neq \tau_j \tau_i$. ■

Proposition 4.1.8. *When $n \geq 3$, the order of the map $\tau_i \tau_j$ is*

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

Proof. The first two cases are straightforward from Proposition 4.1.7 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 $j = i + 1$). From Proposition 3.2.1, we know that no orbit has size greater than 3. Thus to show that the order of $\tau_i \tau_{i+1}$ is 6, it suffices to show that there exists both an orbit of size 2 and orbit of size 3.

The orbit $(\varnothing, \{i+1\}, \{i\})$ 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 4.1.9. Let $\varphi := \tau_n \cdots \tau_2\tau_1$ be the particular Coxeter element in $\mathcal{T}_n$ that toggles at each vertex from left to right.

Example 4.1.10. In $\mathcal{I}_5$, $\varphi(10010) = 01001$ by the following steps:

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

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

Definition 4.1.11. 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.[1] That is, $\chi_j(S)$ is the j^{th} digit of the binary representation of S.

Example 4.1.12. $\chi_1(1010) = 1, \chi_2(1010) = 0, \chi_3(1010) = 1, \chi_4(1010) = 0.$

One of our main theorems to be proven later is the following, which began as a conjecture of Propp. We will later extend this result to the actions of Coxeter elements in $\mathcal{T}_n$ (Theorem 4.1.33). We will do this by explaining why any two Coxeter elements in $\mathcal{T}_n$ are conjugate via a sequence of admissible conjugations, so the result for φ can be extended to any Coxeter element by Theorem 3.3.9.

Theorem 4.1.13 (Propp's conjecture)**.** *Under the action of φ on $\mathcal{I}_n$, $\chi_j - \chi_{n+1-j}$ is 0-mesic for every $1 \leq j \leq n$.*

[1] We use χ_j here instead of I_j for indicator functions like in Definition 3.3.8 to avoid similarity with the set $\mathcal{I}_n$ of independent sets.

Definition 4.1.14. 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 4.1.15. The orbit board for the orbit containing $S = 1010100 \in \mathcal{I}_7$ under the action of φ is shown in Figure 4.1.1. 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 $S^0, S^1, \ldots, S^9$ 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. This illustrates Theorem 4.1.13 since $\chi_j - \chi_{n+1-j}$ has total 0 (and thus average 0) across this orbit for each j.

	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 4.1.1: The orbit under the action φ containing $S = 1010100$.

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

Theorem 4.1.16. *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 4.1.1. This result follows directly from Theorem 5.3.3, a result about more general graphs, but we include another proof here. As with Theorem 4.1.13, we will later generalize this theorem to the action of any Coxeter element (Theorem 4.1.33).

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 $10 \to 01 \to 00$ or $10 \to 00$. (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. ■

4.1.2 Proof of Propp's original conjecture

Our next goal is to prove Theorem 4.1.13 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 4.1.17. 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, $S^{\text{rev}} = \{n+1-i | i \in S\}$.

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

Proposition 4.1.18. *For any S, $(S^{\text{rev}})^{\text{rev}} = S$ and $\varphi(S^{\text{rev}}) = (\varphi^{-1}(S))^{\text{rev}}$.*

Definition 4.1.19. A **symmetrical** independent set is one that is its own reverse. For example, 010010 is symmetrical.

Definition 4.1.20. A φ-orbit $\mathcal{O}$ is **reversible** if for a given (equivalently for all by Proposition 4.1.18) $S \in \mathcal{O}$, S^{rev} is also in $\mathcal{O}$.

Across any reversible orbit, like the one in Example 4.1.15, it is clear that $\chi_j - \chi_{n+1-j}$ has average zero across the 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 4.1.13 holds in general.

Proposition 4.1.21. *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 $\varphi^m(S) = T$; let m be the least such number. Then from Proposition 4.1.18, we have that

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

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

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

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 4.1.22. *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 $S(i, j-1) = S(i, j+1) = S(i-1, j) = S(i+1, j) = 0$.*

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. ■

From Lemma 4.1.22(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 4.1.22(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 1s 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.

Example 4.1.23. The orbit board from Figure 4.1.1, with colors representing the different snakes, is shown in Figure 4.1.2.

	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 4.1.2: The orbit from Figure 4.1.1 with colors representing the snakes.

Therefore, to know where the 1s 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 4.1.24. A **composition** of $n \in \mathbb{P}$ 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 4.1.25. 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 4.1.26. The **snake composition** for a snake is the composition that corresponds to the snake in the way just described.

In Example 4.1.23, 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 lemmas further constrain the possible pattern of 1s in an orbit board.

Lemma 4.1.27. *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 4.1.22(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 4.1.22(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 4.1.22(1). Only the second scenario is possible so $S(3, n) = 1$.

■

Lemma 4.1.28. *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 4.1.22 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 $j \notin S^2$. This is because we cannot have both j and $j+1$ in an independent set after applying $\tau_j \cdots \tau_1$ to S^1. Since $S(2, j-2) = 1$ and $S(2, j) = 0$, we have $S(3, j-1) = 1$ by Lemma 4.1.22.

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 $n \in S^2$.

■

Theorem 4.1.29. *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 by one position.*

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 $2 \in S^1$. So $1, 2 \notin S^2$, 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 4.1.27(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 4.1.22, $S(1, 1) = S(1, 2) = 0$. 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 4.1.28(1,2). Suppose the snake with composition c ends on line S^k, i.e., $S(k, n) =$

	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 4.1.3: The 1s in the orbit board are an example snake. In Example 4.1.30, we describe how to generate an entire orbit from one snake.

1. Then the snake starting on line S^2 contains $S(k+2, n-2)$, so by Lemma 4.1.28(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 4.1.30. 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 4.1.3.

Using Theorem 4.1.29, 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 4.1.4. Also by Theorem 4.1.29, 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 4.1.4.

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 1s in the brown snake on the S^{15} line go on the S^0 line. Every other empty position is a 0 by Lemma 4.1.22(3). The full orbit board is shown in Figure 4.1.4. Notice that this orbit is not reversible, so there is no simple reason for the column-sum vector to be palindromic.

The following should now be clear.

	1	2	3	4	5	6	7	8	9	10
S^0	1	0	1	0	1	0	0	1	0	1
S^1	0	0	0	0	0	1	0	0	0	0
S^2	1	0	1	0	0	0	1	0	1	0
S^3	0	0	0	1	0	0	0	0	0	1
S^4	1	0	0	0	1	0	1	0	0	0
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
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
S^{11}	0	0	0	1	0	1	0	1	0	0
S^{12}	1	0	0	0	0	0	0	0	1	0
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

Red snake: 221121

Purple snake: 211212

Orange snake: 112122

Green snake: 121221

Blue snake: 212211

Brown snake: 122112

FIGURE 4.1.4: The unique orbit containing the snake from Figure 4.1.3 is the orbit containing $S = 1010100101$ (See Example 4.1.30).

Proposition 4.1.31. *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 4.1.13. 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 1s in the leftmost column as in the rightmost column of the (finite version of the) orbit board. Thus,

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

Now 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 4.1.32. For the orbit board in Example 4.1.30, 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 1s 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 1s 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.

4.1.3 Generalizing to Coxeter elements

Theorems 4.1.13 and 4.1.16 are for orbits of the specific action $\varphi = \tau_n \cdots \tau_2\tau_1$. Now we utilize the theory introduced in Section 3.3 to describe how we can generalize the homomesies to the action of any Coxeter element $w \in \mathcal{T}_n$. Recall that a Coxeter element is a product of the toggles $\tau_1, \tau_2, \ldots, \tau_n$, each used exactly once, in some order.

The following generalizes Theorems 4.1.13 and 4.1.16 to the action of *any* Coxeter element. The rest of this section leads up to the proof of Theorem 4.1.33, which is near the end.

Theorem 4.1.33. *Let $w \in \mathcal{T}_n$ be a Coxeter element. Under 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.*

From Proposition 4.1.7, the noncommuting pairs of toggles are exactly the pairs τ_i, τ_{i+1}, so the Coxeter graph $\Gamma(\mathcal{T}_n)$ is the path graph $\mathcal{P}_n$. Recall from Section 3.3 that Coxeter elements correspond to acyclic orientations of the Coxeter graph. For $\mathcal{P}_n$, there are no cycles so all orientations are acyclic.

Example 4.1.34. 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 $\mathcal{P}_7$.

$$\begin{array}{ccccccccccccc} \bullet & \longleftarrow & \bullet & \longleftarrow & \bullet & \longrightarrow & \bullet & \longrightarrow & \bullet & \longleftarrow & \bullet & \longrightarrow & \bullet \\ 1 & & 2 & & 3 & & 4 & & 5 & & 6 & & 7 \end{array}$$

Since $\mathcal{T}_n$ has generators $\tau_1, \tau_2, \ldots, \tau_n$ satisfying $\tau_i^2 = 1$ and $(\tau_i\tau_j)^2 = 1$ when $|i - j| > 1$, any two Coxeter elements in $\mathcal{T}_n$ are conjugate [51, Lemma 5.1]. This is not the case for every toggle group, but rather the case for this special case of toggling independent sets of path graphs. 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 4.1.35. Let $w = \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1$ as in Example 4.1.34. Refer to Figure 4.1.5 for a sequence of conjugations to transform w into φ and the corresponding orientations of $\mathcal{P}_n$ at each step. Each of these conjugations is by a final element and thus corresponds to changing a sink to a source in the corresponding orientation. 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 now complete the proof of Theorem 4.1.33.

Proof of Theorem 4.1.33. Let $w \in \mathcal{T}_n$ be a Coxeter element. Then since we can transform w into φ via a sequence of admissible conjugations, Theorem 3.3.9 says that any statistic that is a linear combination of $\chi_1, \chi_2, \ldots, \chi_n$ statistics is c-mesic under the action of w if and only if it is c-mesic under φ. Therefore, Theorem 4.1.13 implies (1) and Theorem 4.1.16 implies (2). ■

Example 4.1.36. Figure 4.1.6 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

$$
\begin{array}{rll}
& \tau_3\tau_4\tau_2\tau_6\tau_7\tau_5\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longrightarrow 4 \longrightarrow 5 \longleftarrow 6 \longrightarrow 7 \\
= & \tau_3\tau_4\tau_2\tau_6\tau_5\tau_1\tau_7 & \\
& \text{conj}\big\downarrow \tau_7 & \\
& \tau_7\tau_3\tau_4\tau_2\tau_6\tau_5\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longrightarrow 4 \longrightarrow 5 \longleftarrow 6 \longleftarrow 7 \\
= & \tau_7\tau_3\tau_4\tau_2\tau_6\tau_1\tau_5 & \\
& \text{conj}\big\downarrow \tau_5 & \\
& \tau_5\tau_7\tau_3\tau_4\tau_2\tau_6\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longrightarrow 4 \longleftarrow 5 \longrightarrow 6 \longleftarrow 7 \\
= & \tau_5\tau_7\tau_3\tau_4\tau_2\tau_1\tau_6 & \\
& \text{conj}\big\downarrow \tau_6 & \\
& \tau_6\tau_5\tau_7\tau_3\tau_4\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longrightarrow 4 \longleftarrow 5 \longleftarrow 6 \longrightarrow 7 \\
= & \tau_6\tau_5\tau_3\tau_4\tau_2\tau_1\tau_7 & \\
& \text{conj}\big\downarrow \tau_7 & \\
& \tau_7\tau_6\tau_5\tau_3\tau_4\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longrightarrow 4 \longleftarrow 5 \longleftarrow 6 \longleftarrow 7 \\
= & \tau_7\tau_6\tau_5\tau_3\tau_2\tau_1\tau_4 & \\
& \text{conj}\big\downarrow \tau_4 & \\
& \tau_4\tau_7\tau_6\tau_5\tau_3\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longleftarrow 4 \longrightarrow 5 \longleftarrow 6 \longleftarrow 7 \\
= & \tau_4\tau_7\tau_6\tau_3\tau_2\tau_1\tau_5 & \\
& \text{conj}\big\downarrow \tau_5 & \\
& \tau_5\tau_4\tau_7\tau_6\tau_3\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longleftarrow 4 \longleftarrow 5 \longrightarrow 6 \longleftarrow 7 \\
= & \tau_5\tau_4\tau_7\tau_3\tau_2\tau_1\tau_6 & \\
& \text{conj}\big\downarrow \tau_6 & \\
& \tau_6\tau_5\tau_4\tau_7\tau_3\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longleftarrow 4 \longleftarrow 5 \longleftarrow 6 \longrightarrow 7 \\
= & \tau_6\tau_5\tau_4\tau_3\tau_2\tau_1\tau_7 & \\
& \text{conj}\big\downarrow \tau_7 & \\
& \tau_7\tau_6\tau_5\tau_4\tau_3\tau_2\tau_1 & 1 \longleftarrow 2 \longleftarrow 3 \longleftarrow 4 \longleftarrow 5 \longleftarrow 6 \longleftarrow 7
\end{array}
$$

FIGURE 4.1.5: A demonstration showing 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.

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 4.1.35. As in the proof of Theorem 4.1.33, 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 4.1.6: 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 4.1.36.

Remark 4.1.37. 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 what independent sets are in a given orbit does not extend to general Coxeter elements. In the orbit on the left in Figure 4.1.6, notice that parts 1 and 2 of Lemma 4.1.22 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 4.1.21 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.

4.2 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 [43] led Tom Roby and the author 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 4.1. Our main tools are Burnside's Lemma and some bijections.

Recall Proposition 4.1.5 that $\#\mathcal{I}_n = F_{n+2}$ where $F_0 = 0$, $F_1 = 1$ and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$.

Proposition 4.2.1. *Let $k \in \mathbb{P}$.*

1. *The number of symmetrical independent sets in $\mathcal{I}_{2k}$ is F_{k+1}.*

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

Proof. (1) 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}.

(2) 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 (or w-orbits for a given Coxeter element w), there is a connection with binary necklaces, which we now introduce.

Definition 4.2.2. 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 n-gonal dihedral group D_n generated by cyclic rotation and reversal. The **length** of a binary necklace or bracelet is the length of any string in the equivalence class.

Example 4.2.3. There are six binary necklaces of length 4: $\{0000\}$, $\{1000, 0100, 0010, 0001\}$, $\{1010, 0101\}$, $\{1100, 1001, 0011, 0110\}$, $\{1110, 1101, 1011, 0111\}$, $\{1111\}$, all of which are also binary bracelets. For $n \leq 5$, we get the same equivalence classes under C_n as for 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 [43, Seq. A000029, A000031].

Now we use Burnside's Lemma [45, 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.

Lemma 4.2.4. *The number of binary strings of length n with no subsequence* 11 *that do not both start and end with 1 is $F_{n-1} + F_{n+1}$, for $n \geq 1$.*

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 1: If s begins with 1, then it has the form 10_________0, where the blank space represents an independent set of $\mathcal{P}_{n-3}$. So there are F_{n-1} strings of this form.

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.

Altogether there are $F_{n-1} + F_{n+1}$ such binary strings. ∎

Lemma 4.2.5 (Burnside's Lemma)**.** *Let Y be a finite set and G a subgroup of $\mathfrak{S}_Y$. For each $w \in G$, let $\mathrm{Fix}(w) = \{y \in Y : w(y) = y\}$ be the set of elements of Y fixed by w. Let Y/G be the set of orbits of G. Then*

$$\#(Y/G) = \frac{1}{\#G} \sum_{w \in G} \# \, \mathrm{Fix}(w).$$

Proposition 4.2.6. *The number of binary necklaces of length n with no subsequence 11 is given by [43, Seq. A000358]:*

$$\frac{1}{n} \sum_{d|n} \phi(n/d)(F_{d-1} + F_{d+1}), \tag{4.2.1}$$

where ϕ represents Euler's totient function.

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 4.2.4. By elementary group theory, there are $\phi(n/d)$ elements of order n/d, and the result follows. ∎

Proposition 4.2.7. *The number of binary bracelets of length n with no subsequence 11 is given by*

$$\frac{1}{2}\left(F_{\lfloor n/2\rfloor+2} + \frac{1}{n}\sum_{d|n}\phi(n/d)(F_{d-1}+F_{d+1})\right). \tag{4.2.2}$$

This sequence appears as [43, Seq. A129526].

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

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

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

Case 1: n odd. Here all lines of reflection of the n-gon leave exactly one point fixed, as shown in Figure 4.2.1. To determine which strings are fixed by a reflection, we assume without loss of generality that we are working with the reflection that reverses a string. The strings fixed by reversal are the palindromic strings. Since our binary strings cannot both start and end with 1, any palindromic string both begins and ends with 0.

We count the palindromic strings with 0 in the center.

These are of the form $0\underline{\hspace{2em}}0\underline{\hspace{2em}}0$ where the blue string is the reversal of the red and both the red and blue strings do not contain the subsequence 11.

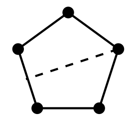

FIGURE 4.2.1: Showing an example reflection of the pentagon.

For example, if the red string is 10010, then the blue string is 01001 and we have the palindromic string 0100100010010 of length 13.

The red string determines the entire string, and the red string can be any independent set of $\mathcal{P}_{(n-3)/2}$, so there are $F_{(n+1)/2}$ strings of this form.

We count the palindromic strings with 1 in the center.

The 1 in the center is surrounded by zeros, so these are of the form 0_______010_______0 where the blue string is the reversal of the red and both the red and blue strings do not contain the subsequence 11.

For example, if the red string is 1010, then the blue string is 0101 and we have the palindromic string 0101001001010 of length 13.

The red string determines the entire string, and the red string can be any independent set of $\mathcal{P}_{(n-5)/2}$, so there are $F_{(n-1)/2}$ strings of this form.

In total there are $F_{(n+1)/2} + F_{(n-1)/2} = F_{(n+3)/2}$ strings fixed by a reflection.

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

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

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. These can be seen in Figure 4.2.2.

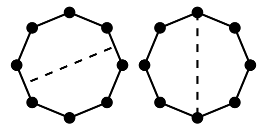

FIGURE 4.2.2: Showing an example of each type of reflection of the octagon.

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. The strings fixed by reversal are the palindromic strings. Since our binary strings cannot both start and end with 1, any palindromic string both begins and ends with 0. The center digits must also be 0. These palindromic strings are of the form 0______00______0 where the blue string is the reversal of the red and both the red and blue strings do not contain the subsequence 11.

For example, if the red string is 10010, then the blue string is 01001 and we have the palindromic string 01001000010010 of length 14.

The red string determines the entire string, and the red string can be any independent set of $\mathcal{P}_{(n-4)/2}$. So there are $F_{n/2}$ strings of this form.

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 $n/2+1$. Depending on the digits in positions 1 and $n/2+1$, there are four possibilities. In each possibility below, the blue string is the reversal of the red string, so the red string determines the entire string.

1st possibility. 10______0______0

The red string can be any independent set of $\mathcal{P}_{(n-4)/2}$, so there are $F_{n/2}$ such strings.

2nd possibility. 0______0______

The red string can be any independent set of $\mathcal{P}_{(n-2)/2}$, so there are $F_{n/2+1}$ such strings.

3rd possibility. 0______010______

The red string can be any independent set of $\mathcal{P}_{(n-4)/2}$, so there are $F_{n/2}$ such strings.

4th possibility. 10______010______0

The red string can be any independent set of $\mathcal{P}_{(n-6)/2}$, so there are $F_{n/2-1}$ such strings.

Adding together the four possibilities, the number of strings is

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

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

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

Our formulas for both the odd and even cases are equivalent to Equation 4.2.2. ■

Corollary 4.2.8. *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 $S + D =$ (4.2.1) and $S + \frac{1}{2}D =$ (4.2.2) (equations above). Eliminating D gives the result. ■

Example 4.2.9. 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 4.2.8. There are $8 = F_6$ other necklaces.

Theorem 4.2.10. *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 4.2.6, the number of cyclically inequivalent compositions of n with each part equal to 1 or 2 is [43, Seq. A000358] given by*

$$\frac{1}{n}\sum_{d|n}\varphi\left(\frac{n}{d}\right)(F_{d-1}+F_{d+1}).$$

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

$$010010001 \quad \longleftrightarrow \quad 212112.$$

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.

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

■

Corollary 4.2.11. *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 4.2.8 and the proof of Theorem 4.2.10. ∎

Theorem 4.2.12. *The total number of φ-orbits of $\mathcal{I}_n$ is*

$$\frac{1}{n-1} \sum_{d|(n-1)} \phi\left(\frac{n-1}{d}\right)(F_{d-1} + F_{d+1}).$$

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

Proof. Combine the bijection of Proposition 4.1.31 with Theorem 4.2.10 for orbits and with Corollary 4.2.11, for reversible orbits. Note that $\lfloor (n-1)/2 \rfloor + 2 = \lceil n/2 \rceil + 1$. ∎

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

- 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}.

Data for $n \leq 26$ is in Figure 4.2.3. Note that while the number of reversible orbits follows the Fibonacci sequence (with each term repeated twice), non-reversible orbits begin at $n \geq 10$ and eventually account for the majority of orbits.

4.3 Sizes of orbits

In the initial data gathering, Tom Roby and the author observed several mysterious patterns in the sizes of φ-orbits. For example, almost all of the orbits had size congruent to $1-n$ mod 4 (see Figure 4.2.3) and certain orbit sizes like 4 and 6 never appeared at all (for any n).

Vertices	Total Orbits	Reversible Orbits	Size 0 mod 4	Size 1 mod 4	Size 2 mod 4	Size 3 mod 4
2	1	1	0	0	0	1
3	2	2	0	0	1	1
4	2	2	0	1	0	1
5	3	3	1	0	1	1
6	3	3	0	0	0	3
7	5	5	0	1	3	1
8	5	5	0	4	0	1
9	8	8	6	0	1	1
10	10	8	0	1	0	9
11	15	13	0	0	12	3
12	19	13	0	18	0	1
13	31	21	26	1	3	1
14	41	21	0	0	0	41
15	64	34	0	4	59	1
16	94	34	0	91	0	3
17	143	55	141	0	1	1
18	211	55	0	0	0	211
19	329	89	0	1	319	9
20	493	89	0	492	0	1
21	766	144	751	0	12	3
22	1170	144	0	5	0	1165
23	1811	233	0	18	1792	1
24	2787	233	0	2786	0	1
25	4341	377	4336	1	3	1
26	6713	377	0	0	0	6713

FIGURE 4.2.3: The total numbers of orbits, reversible φ-orbits, and the breakdown of orbits by size mod 4. The general formulas for orbits and reversible orbits are given in Theorem 4.2.12. Note that the sizes of the orbits are $1 - n \bmod 4$ the vast majority of the time for $n \geq 6$.

These mysteries were also cleared up via the snake compositions of Section 4.1, which leads to a simple characterization (Theorem 4.3.3) 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 in Figure 4.3.2.

The basic idea uses Theorem 4.1.29: 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 4.1.30.

This naive approach can fail, however, in the case where periodicity of c leads one to create a "superorbit" rather than an orbit, e.g., the orbit of $\mathcal{I}_7$ in Figure 4.3.1 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. Unlike the superorbits we considered in various examples in Chapters 1 and **??**, the ones we consider here are not all the same size.

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

Definition 4.3.1. 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 4.3.2. The composition 21221 is aperiodic, but $\psi(22122212) = 2$ and $\psi(222) = 3$.

Theorem 4.3.3. *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.

Proposition 4.3.4. *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. Therefore, 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 4.3.5. 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 $\{1, 3, 5, \ldots, n\}$, 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

Proposition 4.3.6. *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 1s repeated in $\mathcal{I}_{k+1}$. ■

Example 4.3.7. 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 $0, 1, 2 \bmod 3$, 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

Proposition 4.3.8. *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 22. 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 222 and 33, both periodic. These give orbits of size 2 and 3 respectively. ∎

Proposition 4.3.9. *For $n \geq 2$, there is an orbit of $\mathcal{I}_n$ of size 5 exactly when $n \equiv 1 \bmod 3$. In this case, the orbit is unique.*

Proof. The only compositions of 5 into parts 2 and 3 are 23 and 32, 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 $1212\cdots12$, 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 4.3.10. 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

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	$222\cdots 2$	$n \equiv 1 \bmod 2$
3	3	$111\cdots 1$	all n
4	none	none	none
5	$3+2$	$1212\cdots 12$	$n \equiv 1 \bmod 3$
6	none	none	none
7	$3+2+2$	$122122\cdots 122$	$n \equiv 1 \bmod 5$
8	$3+3+2$	$112112\cdots 112$	$n \equiv 1 \bmod 4$
9	$3+2+2+2$	$12221222\cdots 1222$	$n \equiv 1 \bmod 7$
10	$3+3+2+2$	$11221122\cdots 1122$	$n \equiv 1 \bmod 6$
11	$3+3+3+2$	$11121112\cdots 1112$	$n \equiv 1 \bmod 5$
11	$3+2+2+2+2$	$1222212222\cdots 12222$	$n \equiv 1 \bmod 9$
12	$3+2+3+2+2$	$1212212122\cdots 12122$	$n \equiv 1 \bmod 8$
12	$3+3+2+2+2$	$1122211222\cdots 11222$	$n \equiv 1 \bmod 8$

FIGURE 4.3.2: Classification of path graphs which give certain φ-orbit sizes.

We can continue this classification for $n \geq 2$ as shown in Figure 4.3.2. For example $\mathcal{I}_n$ has an orbit of size 7 if and only if $n \equiv 1 \bmod 5$, in which case the orbit is unique. An orbit of size 11 exists if and only if $n \equiv 1 \bmod 5$ or $n \equiv 1 \bmod 9$, and this orbit is unique except when n is both 1 mod 5 and 1 mod 9 (i.e., $n \equiv 1 \bmod 45$) in which case there exist two such orbits. Also, $\mathcal{I}_n$ has an orbit of size 12 if and only if $n \equiv 1 \bmod 8$, in which case there are always exactly two such orbits.

Proposition 4.3.11. *For even n, the φ-orbit of $\mathcal{I}_n$ containing the empty set has size $n+1$.*

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 4.3.3 for an example of this orbit when $n = 6$. ■

Theorem 4.3.12. *Let $\mathcal{O}$ be an φ-orbit of $\mathcal{I}_n$ and c be a snake composition that appears in*

	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 4.3.3: The φ-orbit of $\mathcal{I}_6$ contaning $\emptyset$.

$\mathcal{O}$. If c is aperiodic, then the size of $\mathcal{O}$ is congruent to $1-n \bmod 4$. Furthermore, even when c is periodic, the size of $\mathcal{O}$ divides an integer $m \equiv 1-n \bmod 4$ for $m \leq 3(n-1)$ (where m depends on the orbit $\mathcal{O}$).

Proof. Using the notation of Theorem 4.3.3, the size of $\mathcal{O}$ is $\frac{3N_1(c)+2N_2(c)}{\psi(c)}$. If $c = \underbrace{111\cdots 1}_{n-1}$, then

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

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, $3N_1(c)+2N_2(c) \equiv 1-n \bmod 4$ 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 $3N_1(c)+2N_2(c) \equiv 1-n \bmod 4$. ∎

Corollary 4.3.13. *For even n, every φ-orbit of $\mathcal{I}_n$ has odd size. Furthermore, when $n \equiv 3 \bmod 4$, there exist no orbits with size divisible by 4.*

Proof. Theorem 4.3.12 tells us that the size of any orbit $\mathcal{O}$ divides an integer $m \equiv 1-n \bmod 4$, which is odd for n even. Meanwhile, $n \equiv 3 \bmod 4$ forces $m \equiv 2 \bmod 4$, which $\#\mathcal{O}$ must divide. ∎

Corollary 4.3.14. *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 4.3.13, $\mathcal{O}$ has odd size so there must be an odd number of symmetrical independent sets in $\mathcal{O}$. By Proposition 4.1.21, an orbit can contain at most two symmetrical independent sets, so $\mathcal{O}$ contains exactly one symmetrical independent set. ■

4.4 Homomesy for toggling order ideals of zigzag posets

In Chapter 3 we described toggle groups for order ideals of posets, particularly the well-studied maps promotion and rowmotion. The original problem about independent sets is connected with this for a particular type of poset, called zigzag posets.

Definition 4.4.1. The **zigzag poset** with n elements, denoted $\mathcal{Z}_n$, is the poset consisting of elements $a_1, ..., a_n$ and relations $a_{2i-1} < a_{2i}$ and $a_{2i+1} < a_{2i}$. (Such posets are also called **fence posets** and are discussed in [46, p. 367].)

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

$\mathcal{Z}_6 =$ 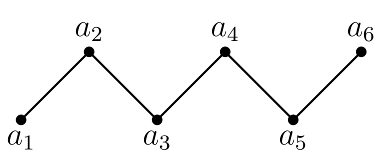

and $\mathcal{Z}_7 =$ 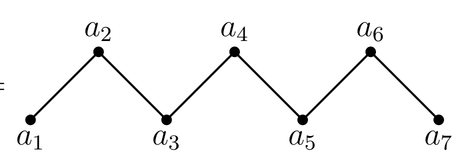

.

To avoid unnecessary double subscripts like t_{a_i} using the notation of Chapter 3, we instead

use t_i for the toggle of a_i. So for any order ideal $I \in J(\mathcal{Z}_n)$,

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

where $I\Delta\{a_i\} := (I \setminus \{a_i\}) \cup (\{a_i\} \setminus I)$ is the symmetric difference of I and $\{a_i\}$. The toggle group generated by $\{t_i | i \in [n]\}$ is denoted $\mathrm{Tog}(\mathcal{Z}_n)$.

The only cover relations are between elements a_i, a_{i+1} for $i \in [n-1]$, so by Propositon 3.1.6, two toggles $t_i, t_j \in \mathrm{Tog}(\mathcal{Z}_n)$ commute if and only if $|i - j| \neq 1$.

Zigzag posets are rc-posets; consider the position function $\pi(a_i) = (i, 1)$ if i is odd and $\pi(a_i) = (i, 2)$ if i is even. Using this position function, $\mathrm{Pro} = t_n \cdots t_2 t_1$.

Also using Proposition 3.1.3,

$$\rho_J = \begin{cases} t_{n-1}t_{n-3}\cdots t_3t_1t_nt_{n-2}\cdots t_2 & \text{if } n \text{ is even} \\ t_nt_{n-2}\cdots t_3t_1t_{n-1}t_{n-3}\cdots t_2 & \text{if } n \text{ is odd} \end{cases}$$

gives rowmotion on $J(\mathcal{Z}_n)$.

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

Proposition 4.4.2. *The map $\eta : \mathcal{I}_n \to J(\mathcal{Z}_n)$ 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 4.4.3. 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 $\eta(S)$.

$1001010 \quad \overset{\eta}{\longmapsto}$

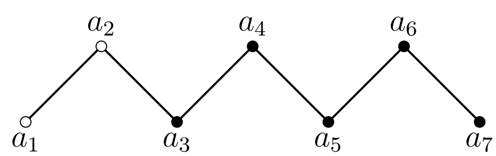

Proposition 4.4.4. *For every $i \in [n]$, $\eta \circ \tau_i = t_i \circ \eta$. Thus, $\eta \circ \varphi = \mathrm{Pro} \circ \eta$, making η an **equivariant** bijection, as shown in the following commutative diagrams.*

$$\begin{tikzcd} \mathcal{I}_n \arrow[r, "\eta"] \arrow[d, "\tau_i"'] & J(\mathcal{Z}_n) \arrow[d, "t_i"] \\ \mathcal{I}_n \arrow[r, "\eta"'] & J(\mathcal{Z}_n) \end{tikzcd} \qquad \begin{tikzcd} \mathcal{I}_n \arrow[r, "\eta"] \arrow[d, "\varphi"'] & J(\mathcal{Z}_n) \arrow[d, "\mathrm{Pro}"] \\ \mathcal{I}_n \arrow[r, "\eta"'] & J(\mathcal{Z}_n) \end{tikzcd}$$

While promotion and rowmotion are conjugate for any poset, we get that any two Coxeter elements in $\mathrm{Tog}(\mathcal{Z}_n)$ are also, by the conjugacy of Coxeter elements in $\mathcal{T}_n$ and Proposition 4.4.4. Thus, the orbit structure of $J(\mathcal{Z}_n)$ under ρ, or any other Coxeter element in $\mathrm{Tog}(\mathcal{Z}_n)$, is the same as that of Pro, which by Proposition 4.4.4, is the same as the orbit structure of φ on $\mathcal{I}_n$.

Using Proposition 4.4.4, we can restate Theorem 4.1.33 for toggling in $J(\mathcal{Z}_n)$, as follows.

Theorem 4.4.5. *Let w be a Coxeter element in $\mathrm{Tog}(\mathcal{Z}_n)$. Let $\chi_{a_j} : J(Z_n) \to \{0,1\}$ be the indicator function of a_j. Then on w-orbits in $J(Z_n)$, 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 Theorem 4.1.33 into the language of $J(\mathcal{Z}_n)$ via Proposition 4.4.4. ■

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.

Zigzag posets are graded with

$$\mathrm{rank}(a_i) = \begin{cases} 0 & \text{if } i \text{ is even,} \\ 1 & \text{if } i \text{ is odd.} \end{cases}$$

So by Proposition 3.1.8, $J(\mathcal{Z}_n)$ has a rowmotion orbit of size 3 consisting of

- the empty order ideal,
- the order ideal consisting of the bottom row elements,
- $\mathcal{Z}_n$ entirely.

It is not directly obvious that this is the only orbit of this size. But since the orbit structure of ρ is the same as that of φ on $\mathcal{I}_n$, uniqueness follows from Proposition 4.3.6.

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.

Chapter 5

Toggling noncrossing partitions

This chapter is joint work with David Einstein, Miriam Farber, Emily Gunawan, Matthew Macauley, James Propp, and Simon Rubinstein-Salzedo that began at a workshop on Dynamical Algebraic Combinatorics hosted by American Institute of Mathematics in March 2015. This exposition is published in [13] with some modifications and reorganization here.

5.1 Linear representations and toggles for noncrossing partitions

Recall the definition, basic facts, and related terminology of noncrossing partitions from Section 2.2. To employ Striker's notion of toggles, as detailed in Chapter 3, to $\mathrm{NC}(n)$, we make use of the **linear representation** of noncrossing partitions, as shown in Figure 5.1.1. This representation of $P \in \mathrm{NC}(n)$ depicts the numbers $1, \ldots, n$ as equally-spaced points on a horizontal line and consists of arcs above the line joining points i and j whenever i and j are *successive* elements of the same block. Formally, the linear representation of P consists

of those pairs (i, j) with $1 \leq i < j \leq n$ with the property that i and j are in the same block of P but none of $i+1, i+2, \ldots, j-1$ (the "interior" of the arc (i, j)) are also in that block. The noncrossing property of the partition guarantees that if two arcs belong to P, then their interiors are disjoint, their left endpoints are distinct, and their right endpoints are distinct. (That is, we never see two arcs exhibiting the three forbidden configurations depicted in Figure 5.1.2, called respectively *crossing*, *left-nesting*, and *right-nesting*. Note however that nested arcs are allowed.) Conversely, any collection of arcs satisfying these conditions determines a unique noncrossing partition P.

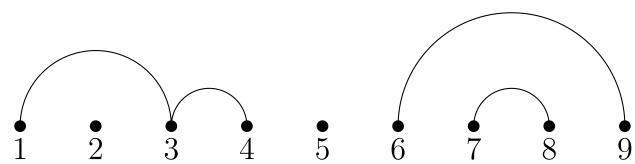

FIGURE 5.1.1: The linear representation $\{(1,3),(3,4),(6,9),(7,8)\}$ of the partition 134|2|5|69|78.

We consider the linear representation to be another way of writing a noncrossing partition. Thus, for the partition of [9] in Figure 5.1.1, $\{\{1,3,4\},\{2\},\{5\},\{6,9\},\{7,8\}\}$, 134|2|5|69|78, and $\{(1,3),(3,4),(6,9),(7,8)\}$ are all considered to be equal. Note that the last one is the linear representation, which only corresponds to a unique noncrossing partition if we know that $n = 9$. We assume we know what n is when we write a noncrossing partition as a set of arcs.

We will make use of both the circular and linear representations of noncrossing parti-

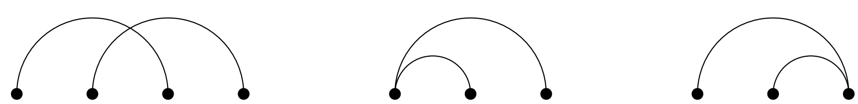

FIGURE 5.1.2: Disallowed pairs of arcs in a noncrossing partition: crossing, left-nesting, and right-nesting, respectively.

tions. Recall that the circular representation of $P \in \mathrm{NC}(n)$ depicts the numbers $1, \ldots, n$ as equally-spaced points on a circle (by convention arranged clockwise) and the blocks as convex hulls. Figure 5.1.3 shows the linear and circular representations of the noncrossing partition $1|245|3|68|7 = \{(2,4),(4,5),(6,8)\}$. The noncrossing property ensures that the convex hulls are pairwise disjoint, i.e., the blocks are "noncrossing."

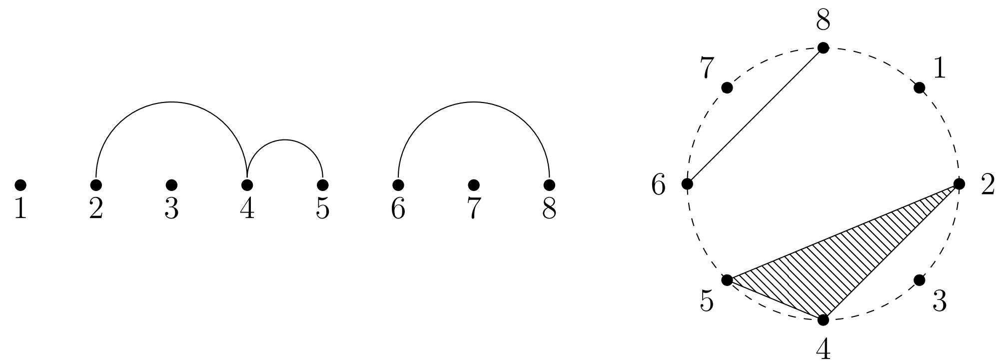

FIGURE 5.1.3: The linear and circular representations of the noncrossing partition $1|245|3|68|7 = \{(2,4),(4,5),(6,8)\}$.

The linear representation allows us to write $\mathrm{NC}(n)$ as a set of subsets of $\{(i,j)|1 \leq i < j \leq n\}$. Viewing $\mathrm{NC}(n)$ in this way, we can define a *toggle* $\tau_{i,j}$ for each arc (i,j), as we discussed in Chapter 3. Unlike toggles for many families of subsets (but like toggles for independent sets of graphs), we can always remove an arc from an noncrossing partition and be left with a noncrossing partition, though we cannot always add an arc. This gives the following slightly simpler expression for $\tau_{i,j}$:

$$\tau_{i,j}(P) = \begin{cases} P \cup \{(i,j)\} & \text{if } (i,j) \notin P \text{ and } P \cup \{(i,j)\} \in \mathrm{NC}(n,) \\ P \setminus \{(i,j)\} & \text{if } (i,j) \in P, \\ P & \text{otherwise.} \end{cases}$$

Definition 5.1.1. Let W_n denote the toggle group on $\mathrm{NC}(n)$ generated by the $\binom{n}{2}$ toggle operations.

We defined the Kreweras complement $\kappa : \mathrm{NC}(n) \to \mathrm{NC}(n)$ in Definition 2.2.8, with an example in Figure 2.2.4.

This gives some natural statistics to consider on $\mathrm{NC}(n)$.

Definition 5.1.2. The **arc count statistic** $\alpha(P)$ of a noncrossing partition is the number of pairs (i, j) with $1 \leq i < j \leq n$ appearing in P. The **block count statistic** $\beta(P) = |P|$ is the number of blocks of P.

Since a block with h elements contains $h-1$ arcs, it follows that $\alpha(P) + \beta(P) = n$. This relation shows that α is homomesic under any action if and only if β also is.

The convex hulls of the blocks divide the circular representation of $P \in \mathrm{NC}(n)$ into regions. Replacing a set of singletons (within a region) with a block with h elements adds $h-1$ new regions. So if P has k total blocks, then it divides the circle into $n+1-k$ regions. From the way κ is defined, this implies $\kappa(P)$ has $n+1-k$ blocks, so we have the following.

Proposition 5.1.3. *For any $P \in \mathrm{NC}(n)$, $|P| + |\kappa(P)| = n + 1$.*

Across any κ-orbit, the block count alternates between two (possibly the same) numbers whose average is $\frac{n+1}{2}$. Thus $(\mathrm{NC}(n), \kappa, \beta)$ exhibits homomesy with average $\frac{n+1}{2}$ and $(\mathrm{NC}(n), \kappa, \alpha)$ exhibits homomesy with average $\frac{n-1}{2}$. Recall Proposition 2.2.9 that κ^2 rotates the circular representation counterclockwise $2\pi/n$ radians. So the order of κ divides $2n$.

Also, Theorem 2.2.11(3) describes a CSP for the map κ. The original motivation behind studying toggles on $\mathrm{NC}(n)$ is that we can extend the same homomesy result to a large class of actions w (one of which is κ). In general, for such w, the order of w and general orbit structure is unknown, so a natural CSP seems highly unlikely. Also, we do *not* in general

have the $|P| + |w(P)| = n = 1$, so the more specific result for κ does not extend. Only the homomesy does and the proof is highly nontrivial.

It turns out that κ is an element in the toggle group W_n. The proof of the following result is long, so we omit it and refer the reader to the main paper about this problem.

Theorem 5.1.4 ([13, Theorem 5.2])**.** *Let $P \in \mathrm{NC}(n)$. Then*

$$\kappa(P) = \tau_{1,2}\tau_{1,3}\tau_{1,4}\cdots\tau_{1,n}\tau_{2,3}\tau_{2,4}\cdots\tau_{2,n}\cdots\tau_{n-2,n-1}\tau_{n-2,n}\tau_{n-1,n}(P).$$

In the above expression, $\tau_{i,j}$ appears to the left of $\tau_{k,\ell}$ whenever one of the following conditions hold:

- *$i < k$, or*
- *$i = k$ and $j < \ell$.*

So κ is a Coxeter element in W_n. The generalized result is the following which states we have the same homomesy for a large class of partial Coxeter elements. The following is the main theorem of this chapter.

Theorem 5.1.5. *Let $w \in W_n$ be any partial Coxeter element that contains every toggle of the form $\tau_{i,i+1}$. Then the triple $(\mathrm{NC}(n), \alpha, w)$ exhibits homomesy with average $\frac{n-1}{2}$. This implies also that $(\mathrm{NC}(n), \beta, w)$ exhibits homomesy with average $\frac{n+1}{2}$.*

Example 5.1.6. Figure 5.1.4 shows the five orbits of $w = \tau_{3,4}\tau_{1,2}\tau_{2,3}\tau_{1,4}$ on $\mathrm{NC}(4)$. Note that w satisfies the necessary conditions in Theorem 5.1.5 but is not a Coxeter element, since it does not contain $\tau_{1,3}$ or $\tau_{2,4}$. The figure shows that arc count is $3/2$-mesic under the action of w.

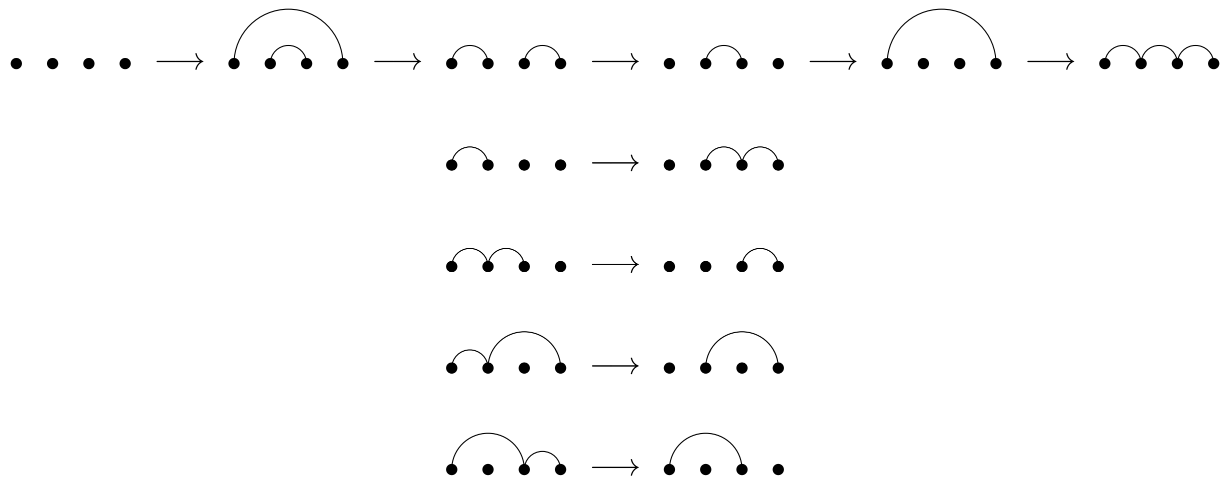

FIGURE 5.1.4: The five orbits of $w = \tau_{3,4}\tau_{1,2}\tau_{2,3}\tau_{1,4}$ on NC(4). Notice that in any orbit, the average of the arc count is $\frac{3}{2}$. Also notice that in general, $\alpha(P) + \alpha(w(P)) \neq 3$, as is the case for Kreweras complementation.

We will prove Theorem 5.1.5 in Section 5.2. As a special case, this theorem has the following corollary.

Corollary 5.1.7. *Let $w \in W_n$ be any Coxeter element. Then the triple $(\mathrm{NC}(n), \alpha, w)$ exhibits homomesy with average $\frac{n-1}{2}$. This implies also that $(\mathrm{NC}(n), \beta, w)$ exhibits homomesy with average $\frac{n+1}{2}$.*

Hence the arc count statistic is *simultaneously* homomesic for all partial Coxeter elements that contain every $\tau_{i,i+1}$ (which includes all Coxeter elements). We will also show in Section 5.2 that there are more refined statistics that are homomesic for certain partial Coxeter elements.

Another consequence of Theorem 5.1.5 is the following.

Corollary 5.1.8. *Let n be even and $w \in W_n$ be any partial Coxeter element that contains every toggle of the form $\tau_{i,i+1}$. Then each w-orbit of $\mathrm{NC}(n)$ contains an even number of noncrossing partitions.*

Proof. The arc count of any noncrossing partition is an integer. Therefore, the only way for the average arc count across an orbit to be $\frac{n-1}{2}$, which is not an integer for even n, is if the orbit contains an even number of noncrossing partitions. ■

This gives an example as to how homomesy can be used to prove statements that neither mention homomesy nor the statistic that is homomesic. There is no other known way to prove Corollary 5.1.8, as there does not appear to be a way to characterize the orbit sizes in general. For example, in NC(6), the sizes of the orbits of the Coxeter element

$$w = \tau_{4,6}\tau_{3,6}\tau_{2,4}\tau_{1,5}\tau_{2,5}\tau_{1,3}\tau_{3,4}\tau_{1,2}\tau_{1,6}\tau_{2,6}\tau_{3,5}\tau_{2,3}\tau_{1,4}\tau_{5,6}\tau_{4,5}$$

are 4, 22, 46, and 60. There is no noticeable pattern aside from the fact that they are all even.

Figure 5.1.4 displays an example of Corollary 5.1.8, as each orbit in the example contains either two or six noncrossing partitions.

Unlike the toggling problem on independent sets of path graphs detailed in Chapter 4, not all Coxeter elements are conjugate. In particular, the orbit structures are generally different. So we cannot use only Coxeter group theory to extend the result from κ to any Coxeter element.

Simion and Ullman defined an bijection $\lambda : \mathrm{NC}(n) \to \mathrm{NC}(n)$ as $\lambda = \eta \circ \kappa$, where η is the relabeling map that replaces i by $n - i$ for $1 \leq i < n$ and leaves n fixed [42]. This satisfies $|P| + |\lambda(P)| = n + 1$, since relabeling does not affect the number of blocks. So α and β are homomesic under λ as well. Unlike κ, λ is an involution. In general, λ cannot be expressed as one of the maps in Theorem 5.1.5. One way to determine this is by checking all the possible maps for $n = 3$, or by noting that a statistic we show in Section 5.2 to be 1-mesic is not under λ.

It will be helpful for what follows to classify which pairs of toggles do and do not commute. Any pair of distinct arcs (i, j) and (k, ℓ) can be classified into one of six types (possibly after swapping (i, j) with (k, ℓ)):

1. $i < j < k < \ell$ (disjoint),
2. $i < k < \ell < j$ (nesting),
3. $i < j = k < \ell$ (m-shaped),
4. $i = k < j < \ell$ (left-nesting),
5. $i < k < j = \ell$ (right-nesting),
6. $i < k < j < \ell$ (crossing).

The type is sufficient to determine whether or not the pair of toggles commutes.

Proposition 5.1.9. *Let $\tau_{i,j}$ and $\tau_{k,\ell}$ be distinct toggles. Then $\tau_{i,j}$ and $\tau_{k,\ell}$ commute if and only if the arcs (i, j) and (k, ℓ) are disjoint, nesting, or m-shaped.*

Proof. Suppose first that (i, j) and (k, ℓ) are left-nesting, right-nesting, or crossing. So (i, j) and (k, ℓ) are not allowed to be in the same noncrossing partition. Let $P = \{\}$ be the empty noncrossing partition of $[n]$. Then $\tau_{i,j}\tau_{k,\ell}(P) = \{(k, \ell)\}$, whereas $\tau_{k,\ell}\tau_{i,j}(P) = \{(i, j)\}$. Thus $\tau_{i,j}$ and $\tau_{k,\ell}$ do not commute.

On the other hand, suppose that (i, j) and (k, ℓ) are disjoint, nested, or m-shaped. Then (i, j) and (k, ℓ) are allowed to be in the same noncrossing partition, so adding or removing (i, j) does not interfere with adding or removing (k, ℓ). Therefore $\tau_{i,j}$ and $\tau_{k,\ell}$ commute. ■

The commuting pairs are illustrated in Figure 5.1.5, and the non-commuting pairs were shown back in Figure 5.1.2. Note that the non-commuting toggles correspond exactly with

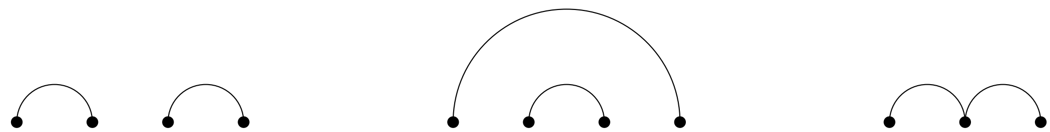

FIGURE 5.1.5: Commuting pairs of toggles: disjoint, nesting, and m-shaped, respectively.

the disallowed pairs of arcs, and the commuting toggles with the allowed pairs of arcs, which was the key idea of the proof above.

Corollary 5.1.10. *Given $i < j$ with $j - i = h$, there are $h(n + 1 - h) - 2$ toggles $\tau_{k,\ell}$ that do not commute with $\tau_{i,j}$.*

Proof. Let $i < j$ with $j - i = h$. We will classify the toggles that do not commute with $\tau_{i,j}$.

If $\ell > j$, then $\tau_{i,\ell}$ does not commute with $\tau_{i,j}$, and if $k < i$, then $\tau_{k,j}$ does not commute with $\tau_{i,j}$. This type of toggle is in one-to-one correspondence with the numbers in $[n]$ that are less than i or greater than j, and there are $n - h - 1$ such numbers.

The other way that $\tau_{k,\ell}$ will not commute with $\tau_{i,j}$ is if one of k or ℓ is strictly between i and j, and the other is not strictly between i and j. There are $h - 1$ numbers in $[n]$ that are strictly between i and j, and the other $n + 1 - h$ numbers in $[n]$ are not strictly between i and j, so there are $(h - 1)(n + 1 - h)$ such pairs k, ℓ.

Adding up the two cases, there are $n - h - 1 + (h - 1)(n + 1 - h) = h(n + 1 - h) - 2$ toggles $\tau_{k,\ell}$ that do not commute with $\tau_{i,j}$. ∎

The following describes the order of a product of two toggles.

Proposition 5.1.11. *For any pair of toggles $\tau_{i,j}$ and $\tau_{k,\ell}$, let $m(\tau_{i,j}\tau_{k,\ell})$ denote the order of*

the element $\tau_{i,j}\tau_{k,\ell}$ *in* W_n*. Then*

$$m(\tau_{i,j},\tau_{k,\ell}) = \begin{cases} 1 & \text{if } (i,j) = (k,\ell), \\ 2 & \text{if } \tau_{i,j}, \tau_{k,\ell} \text{ commute and } (i,j) \neq (k,\ell), \\ 6 & \text{if } \tau_{i,j}, \tau_{k,\ell} \text{ do not commute.} \end{cases}$$

Proof. The first two cases are straightforward since toggles are involutions.

Suppose $\tau_{i,j}$ and $\tau_{k,\ell}$ do not commute. By Proposition 5.1.9, the arcs are either left-nesting, right-nesting, or m-shaped, and so no noncrossing partition can contain both (i,j) and (k,ℓ). From Proposition 3.2.1, no orbit has size greater than 3. We proceed by showing the existence of orbits of sizes 3 and 2.

Applying $\tau_{i,j}\tau_{k,\ell}$ to $\{\}$ gives $\{(k,\ell)\}$. Applying $\tau_{i,j}\tau_{k,\ell}$ to $\{(k,\ell)\}$ gives $\{(i,j)\}$, and then applying $\tau_{i,j}\tau_{k,\ell}$ again gives $\{\}$, so this is an orbit of size 3.

Let A be any arc that can be in the same noncrossing partition as one of (i,j) and (k,ℓ), but not the other. Figure 5.1.2 shows that we can always find such an arc. Specifically, if these arcs are crossing with $i < k < j < \ell$, or left-nesting with $i = k < j < \ell$, then $A = (j,\ell)$ will work. The right-nesting case is analogous. Without loss of generality, assume A can be in the same noncrossing partition as (i,j) but not (k,ℓ). Then $(\{A\}, \{A,(i,j)\})$ is an orbit of size 2. ■

The commutativity relations between the toggles can be described by a undirected graph, called the *base graph* Γ_n. Note that Γ_n is the Coxeter graph $\Gamma(W_n)$ for the toggle group, hence the choice of gamma. However, since the noncommuting pairs of toggles refer to precisely the disallowed pairs of arcs, $\mathrm{NC}(n)$ can be described as the set of independent sets of Γ_n. This gives the extension to Section 5.3 is why we do not only think of it as the Coxeter graph.

Definition 5.1.12. The base graph Γ_n of W_n is the graph $(V, \mathcal{E})$, where $V = \{\tau_{i,j} \mid i < j\}$, and $\mathcal{E}$ consists of edges of the form $\{\tau_{i,j}, \tau_{k,\ell}\}$, where $\tau_{i,j}$ and $\tau_{k,\ell}$ are non-commuting toggles.

It is easiest to arrange the vertex set $\{\tau_{i,j} \mid 1 \leq i < j \leq n\}$ in an upper-triangular grid, as shown in Figure 5.1.6. Each row is a clique (complete subgraph), as is each column; these correspond to the half-nesting pairs. Finally, there are some "diagonal edges," which correspond to crossing pairs: (i, j) and (k, ℓ) where $i < k < j < \ell$. All of these are negatively sloped.

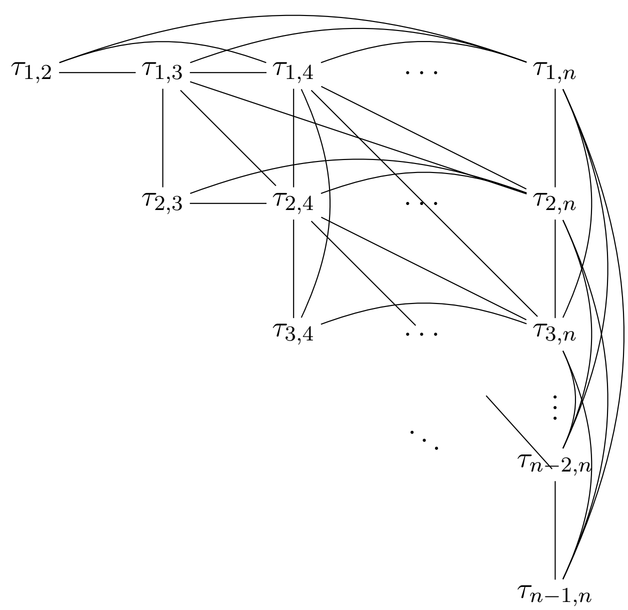

Figure 5.1.6: The base graph Γ_n of the toggle group W_n.

In summary, the base graph Γ_n, when drawn as in Figure 5.1.6, has three types of edges: horizontal, vertical, and diagonal. Thus, it is possible to describe certain acyclic orientations of Γ_n (describing some Coxeter elements) as e.g. "orienting all edges east, south, and southeast" (which describes κ) or "orienting all edges east, north, and southeast."

5.2 Proof of Theorem 5.1.5

When searching for homomesies, it often helps to consider simple indicator function statistics, and then determine which linear combinations of these are homomesic.

Definition 5.2.1. We denote the indicator function of the arc (i,j) by $\chi_{i,j} : \mathrm{NC}(n) \to \{0,1\}$ defined as

$$\chi_{i,j}(P) = \begin{cases} 1 & \text{if } (i,j) \in P, \\ 0 & \text{if } (i,j) \notin P. \end{cases}$$

We utilize the following linear combinations of the indicator functions.

Definition 5.2.2. Given $k \in [n-1]$ and $P \in \mathrm{NC}(n)$, define the statistic $\psi_k : \mathrm{NC}(n) \to \mathbb{Z}$ in the following way:

$$\begin{aligned} \psi_k(P) &= 2\chi_{k,k+1}(P) + \sum_{1 \le i \le k-1} \chi_{i,k+1}(P) + \sum_{k+2 \le j \le n} \chi_{k,j}(P) \\ &= \sum_{1 \le i \le k} \chi_{i,k+1}(P) + \sum_{k+1 \le j \le n} \chi_{k,j}(P) \end{aligned}$$

where $\chi_{i,j}$ is the indicator function of the arc (i,j).

Due to the restrictions on arcs with a common left or right endpoint, and arcs that cross, any noncrossing partition can only contain at most one arc that is of the form $(i,k+1)$ or (k,j). Thus, for any $P \in \mathrm{NC}(n)$, $\psi_k(P) \in \{0,1,2\}$. Also, $\psi_k(P)$ is fully determined by the following three cases.

- $\psi_k(P) = 0$ if and only if P contains no arcs of the form $(i,k+1)$ or (k,j).
- $\psi_k(P) = 1$ if and only if P contains an arc of the form $(i,k+1)$ or (k,j) that is not the arc $(k,k+1)$.

- $\psi_k(P) = 2$ if and only if P contains the arc $(k, k+1)$.

Theorem 5.2.3. *Given $k \in [n-1]$, let T either be a Coxeter element, or a partial Coxeter element that contains $\tau_{k,k+1}$. Then ψ_k is 1-mesic on orbits of T.*

Proof. Without loss of generality, we assume that in T, the toggle $\tau_{k,k+1}$ is the toggle that is applied last. If this is not the case, then we may conjugate T by the toggles that are performed after $\tau_{k,k+1}$, and then the homomesy and orbit sizes will be unchanged by Theorem 3.3.9.

To prove that ψ_k is 1-mesic, it is equivalent to prove that in any orbit $\mathcal{O}$,

$$\#\{P \in \mathcal{O} : \psi_k(P) = 0\} = \#\{P \in \mathcal{O} : \psi_k(P) = 2\}.$$

The general strategy is to show that when $\psi_k(P) = 0$ and r is the smallest positive value such that $\psi_k\left(T^r(P)\right) \neq 1$, then $\psi_k\left(T^r(P)\right) = 2$ and vice versa.

In other words, we will prove that in any orbit, the number of partitions that do not contain any arcs of the form $(i, k+1)$ or (k, j) is equal to the number of partitions that contain the arc $(k, k+1)$.

Let $\{A_1, \ldots, A_m\}$ be the (possibly empty) set of arcs with left endpoint k or right endpoint $k+1$ whose toggles are contained in T, excluding $(k, k+1)$. We will index the A_i's in the order that they are being toggled in T. Note that T may contain other toggles in addition to $\tau_{A_1}, \ldots, \tau_{A_m}$ and $\tau_{k,k+1}$. However, these other arcs do not affect whether or not $(k, k+1)$ can be inserted into a partition, although they may affect whether or not the $A_1, \ldots, A_m$ arcs can be inserted. The toggles that are significant in this proof are $\tau_{A_1}, \ldots, \tau_{A_m}$ and $\tau_{k,k+1}$.

Let P be such that $(k, k+1) \in P$, so $\psi_k(P) = 2$. Then when computing $T(P)$, every toggle τ_{A_i} will not be able to add the arc A_i because $(k, k+1)$ is in the partition, and then the final toggle $\tau_{k,k+1}$ will remove $(k, k+1)$ from the partition. Thus, $T(P)$ contains no arcs of the form $(i, k+1)$ or (k, j), so $\psi_k(T(P)) = 0$.

Let P be such that $\psi_k(P) = 0$. So P contains no arcs of the form $(i, k+1)$ or (k, j). When computing $T(P)$, the toggle τ_{A_1} (if A_1 exists) will attempt to add the arc A_1 to the partition. It may or not be possible to add that arc, depending on other arcs in the partition. If A_1 cannot be added, then τ_{A_2} (if A_2 exists) will attempt to add the arc A_2 to the partition. Again, that may or may not be possible, and the process continues. There are two cases that can happen.

Case 1: An arc A_i is added to the partition. Then the toggles $\tau_{A_{i+1}}, \ldots, \tau_{A_m}, \tau_{k,k+1}$ will do nothing. So $T(P)$ contains the arc A_i and thus $\psi_k(T(P)) = 1$.

Case 2: None of the arcs $A_1, \ldots, A_m$ can be added to the partition when the toggles $\tau_{A_1}, \ldots, \tau_{A_m}$ are applied. Then there are no arcs that interfere with the ability to add the arc $(k, k+1)$, so the final toggle $\tau_{k,k+1}$ adds this arc. Therefore $\psi_k(T(P)) = 2$.

Note that if the word T contains no arcs with left endpoint k or right endpoint $k+1$ other than $(k, k+1)$, then $\{A_1, \ldots, A_m\} = \emptyset$, so we go to Case 2 automatically.

Now let P be a partition that contains A_i for some i. When computing $T(P)$, the toggles τ_{A_j} for $j < i$ do nothing, then τ_{A_i} removes A_i from the partition. Then there are two cases for what happens when the toggles τ_{A_j} for $j > i$ and $\tau_{k,k+1}$ are applied.

Case 1: An arc A_j for some $j > i$ is added to the partition. Then $\tau_{A_{j+1}}, \ldots, \tau_{A_m}, \tau_{k,k+1}$ will do nothing. So $T(P)$ contains the arc A_j and thus $\psi_k(T(P)) = 1$.

Case 2: None of the arcs A_j for $j > i$ can be added when those respective toggles are applied. Then there are no arcs that interfere with the ability to add the arc $(k, k+1)$, so the final toggle $\tau_{k,k+1}$ adds this arc. Therefore $\psi_k(T(P)) = 2$.

From this, it is clear that when applying T repeatedly to P, the next partition $T^r(P)$ for which $\psi_k\left(T^r(P)\right) \neq 1$ satisfies $\psi_k\left(T^r(P)\right) = 2$. Thus, in any orbit $\mathcal{O}$,

$$\#\{P \in \mathcal{O} : \psi_k(P) = 0\} = \#\{P \in \mathcal{O} : \psi_k(P) = 2\},$$

so ψ_k is 1-mesic on orbits of T. ■

The arc count statistic is $\sum_{1\leq i<j\leq n}\chi_{i,j}$. Given any $i<j$ with $j-i\geq 2$, the coefficient of $\chi_{i,j}$ in ψ_i and ψ_{j-1} is 1, and the coefficient of $\chi_{i,j}$ in all other ψ_k is 0. For any i, the coefficient of $\chi_{i,i+1}$ in ψ_i is 2, and the coefficient of $\chi_{i,i+1}$ in all other ψ_k is 0. Therefore, the arc count statistic is equal to $\frac{1}{2}\sum_{k=1}^{n-1}\psi_k$. Theorem 5.1.5 now follows.

Proof of Theorem 5.1.5. By Theorem 5.2.3, ψ_k is 1-mesic on orbits of T, for every $k\in[n-1]$. So the arc count statistic $\alpha=\frac{1}{2}\sum_{k=1}^{n-1}\psi_k$ is $\frac{n-1}{2}$-mesic. ■

5.3 Extension to independent sets of graphs

In this section we generalize our main results to the toggles on independent sets of general simple graphs. We first define some notation.

Definition 5.3.1. Let $G=(V,\mathcal{E})$ be a simple graph. For $v\in V$, we denote by $N(v)$ the set of neighbors of v, called the **neighborhood** of v. A set $W\subset V$ of vertices is called **independent** if no two vertices in W are adjacent. We denote by $\mathrm{card}(W)$ the cardinality of W and $\mathrm{Ind}(G)$ the set of all independent sets of G.

We let τ_v be the toggle on $\mathrm{Ind}(G)$ associated with the vertex $v\in V$. Removing a vertex from an independent set does not change the independence property, so τ_v can be characterized as

$$\tau_v(W)=\begin{cases} W\cup\{v\} & \text{if } v\notin W \text{ and } W\cup\{v\}\in\mathrm{Ind}(G),\\ W\setminus\{v\} & \text{if } v\in W,\\ W & \text{otherwise.}\end{cases}$$

We let $\mathcal{T}_G$ denote the toggle group on $\mathrm{Ind}(G)$ generated by $\{t_v | v \in V\}$.

Since noncrossing partitions $P \in \mathrm{NC}(n)$ are equivalently independent sets of the base graph Γ_n, the toggles on $\mathrm{NC}(n)$ are a special case of toggles on $\mathrm{Ind}(G)$. An example of this is in Figure 5.3.1 where an element of NC(5) is displayed as an independent set of Γ_5. Thus, the action of the group W_n on $\mathrm{NC}(n)$ is isomorphic to the action of the group $\mathcal{T}_{\Gamma_n}$ on Γ_n.

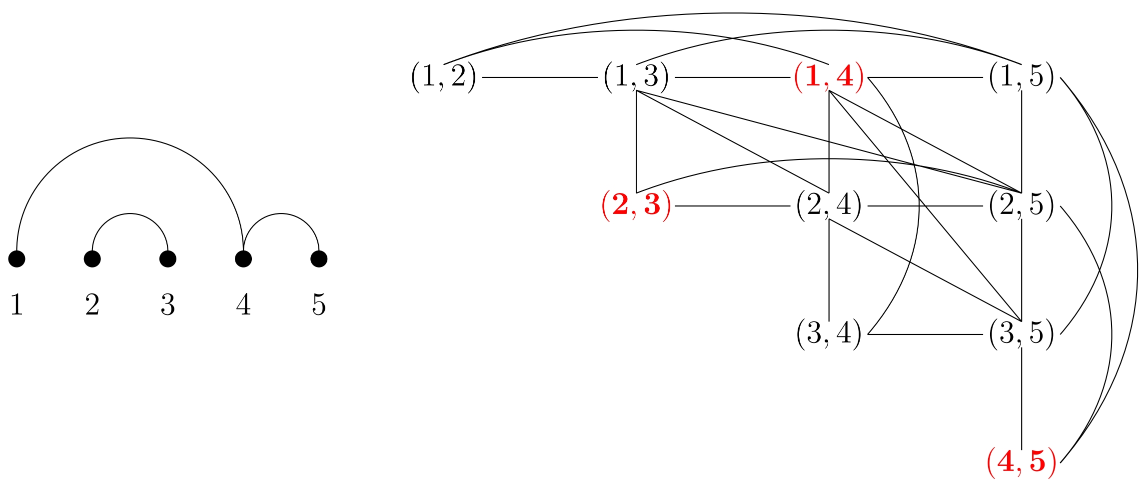

FIGURE 5.3.1: The noncrossing partition $\{(1,4),(2,3),(4,5)\}$ of [5] shown at left corresponds to the independent set of Γ_5 displayed in red and bold on the right.

Next, let us introduce an analogue to the ψ_k statistics on $\mathrm{NC}(n)$.

Definition 5.3.2. Given $G = (V, \mathcal{E})$ and $v \in V$, define the statistic $\psi_v : \mathrm{Ind}(G) \to \mathbb{Z}$ in the following way:

$$\psi_v(W) = 2\chi_v(W) + \sum_{u \in N(v)} \chi_u(W)$$

where χ_v is the indicator function of the vertex v.

For $G = \Gamma_n$ and $v = (k, k+1)$ this definition coincides with Definition 5.2.2. The following result is a generalization of Theorem 5.2.3.

Theorem 5.3.3. *Let $G = (V, \mathcal{E})$ be a simple graph. Given $v \in V$, let T be a partial Coxeter element that contains τ_v. If $N(v)$ forms a clique (complete subgraph) in G, then the statistic ψ_v is 1-mesic on orbits of T.*

Proof. Let us examine the proof of Theorem 5.2.3. The only property of the graph Γ_n that is used in the proof is that for any k, the set of neighbors of the vertex $v = (k, k+1)$ is a clique. This implies that at most one of the vertices in $\{v\} \cup N(v)$ can be contained in an independent set in Γ_n. Thus, similarly to the proof of Theorem 5.2.3, we conclude that ψ_v is 1-mesic on orbits of T. ■

For the setting of the path graph $\mathcal{P}_n$, Theorem 5.3.3 leads directly to Theorem 4.1.16 and the more general Theorem 4.1.33(2) from the previous chapter; for each of 1 and n has only one neighbor, so the neighborhood forms a clique. This cannot be used for other vertices.

We are now ready to present a generalization of Theorem 5.1.5.

Theorem 5.3.4. *Let $G = (V, \mathcal{E})$ be a simple graph with maximal independent set U of vertices that satisfies the following two properties:*

1. *For any $u \in U$, the set of vertices $N(u)$ forms a clique in G.*

2. *Any vertex in $V \setminus U$ has exactly two neighbors in U.*

Let T be a partial Coxeter element containing all toggles τ_u for $u \in U$. Then card *is $|U|/2$-mesic under the action of T.*

Proof. By Theorem 5.3.3 $\sum_{u \in U} \psi_u$ is $|U|$-mesic on orbits of T. On the other hand, property (2) implies that $\sum_{u \in U} \psi_u = 2 \sum_{v \in V} \chi_v = 2\,\mathrm{card}$. Therefore card is $|U|/2$-mesic. ■

Definition 5.3.5. Let $G = (V, \mathcal{E})$ be a graph. Then we say that G is **2-cliquish** if for *some* maximal independent set U, the conditions of Theorem 5.3.4 are satisfied.

It may seem that the definition of 2-cliquish is very restrictive and would not apply to many interesting graphs other than the base graph Γ_n for noncrossing partitions. However, it is actually easy to construct various classes of 2-cliquish graphs. The rest of this section is devoted to describing some such classes.

Example 5.3.6. Some examples of 2-cliquish graphs are as follows.

- A complete graph with a single edge removed is 2-cliquish. The two vertices without an edge connecting them form the maximal independent set. An example of this type of graph is in Figure 5.3.2.

- Given any graph G, define a graph G' in the following way. Start with G. For every vertex v in the graph, add two new vertices and connect v to the two new vertices. Then G' is 2-cliquish and the maximal independent set is the added vertices. An example of this type of graph is in Figure 5.3.3.

- Let C_n denote the cycle graph with n vertices. For every edge e in C_n add a vertex v_e to the graph and add edges from v_e to each endpoint of e. This graph is 2-cliquish. Its maximal independent set is the set of n added vertices $\{v_e\}$. An example of this type of graph is in Figure 5.3.4.

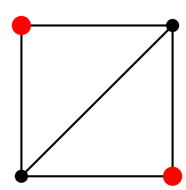

FIGURE 5.3.2: The complete graph K_4 with an edge removed. The maximal independent set is shown with the large red vertices.

The following theorem describes ways to form 2-cliquish graphs from other 2-cliquish graphs.

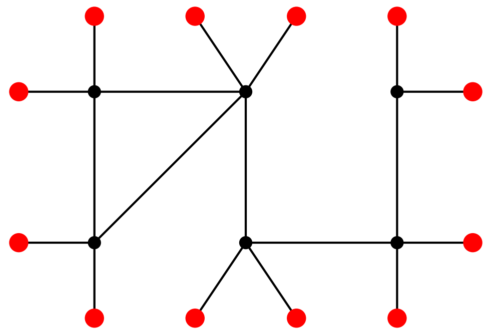

FIGURE 5.3.3: A 2-cliquish graph formed starting with a graph (the subgraph of black vertices) and attaching two new vertices to each vertex. The large red vertices form the maximal independent set.

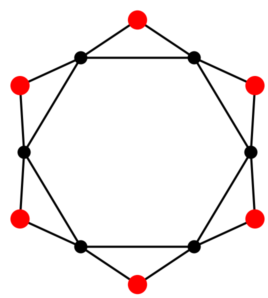

FIGURE 5.3.4: The graph formed from C_6 with an extra vertex added for each edge adjacent to the endpoints of the corresponding edge. This graph is 2-cliquish and the large red vertices form the maximal independent set.

Theorem 5.3.7. *Let $G = (V, \mathcal{E})$ and $G' = (V', \mathcal{E}')$ be 2-cliquish graphs with maximal independent sets U and U' respectively.*

1. *The disjoint union of G and G' is 2-cliquish with maximal independent set $U \cup U'$.*
2. *Let e be an edge not in $\mathcal{E}$ with endpoints in $V \setminus U$. Then $(V, \mathcal{E} \cup \{e\})$ is 2-cliquish with maximal independent set U.*
3. *Let e be an edge in $\mathcal{E}$ with endpoints $v, w \in V \setminus U$ such that v and w do not have a common neighbor in U. Then $(V, \mathcal{E} \setminus \{e\})$ is 2-cliquish with maximal independent set U.*

Proof. Part (1) is clear from the definition. For part (2), consider two vertices $v, w \in V \setminus U$ that are not adjacent in G. Since v and w are not adjacent, they cannot have a common neighbor in U, since $N(u)$ is a clique for all $u \in U$. Thus, adding an edge e between v and w does not change the fact that $N(u)$ is a clique for all $u \in U$. Since the endpoints of e are in $V \setminus U$, the graph $(V, \mathcal{E} \cup \{e\})$ also satisfies the condition that every vertex in $V \setminus U$ has exactly two neighbors in U. It also does not change the fact that U is a maximal independent set. Therefore, $(V, \mathcal{E} \cup \{e\})$ is 2-cliquish.

To prove (3), let e be an edge in $\mathcal{E}$ with endpoints $v, w \in V \setminus U$ such that v and w do not have a common neighbor in U. Since e has no endpoints in U, $(V, \mathcal{E} \setminus \{e\})$ satisfies the condition that every vertex in $V \setminus U$ has exactly two neighbors in U. Also, $N(u)$ is a clique for all $u \in U$ because this is true for G so it is true when removing an edge between vertices without a common neighbor in U. Thus, $(V, \mathcal{E} \cup \{e\})$ is 2-cliquish. ■

We now discuss how to generate all 2-cliquish graphs with a given number of vertices.

Definition 5.3.8. Let $G = (V, \mathcal{E})$ be 2-cliquish.

- We say G is **skeletal** if no edges can be removed from it as in part (3) of Theorem 5.3.7.

- The graph formed from removing edges from G when possible in accordance with part (3) of Theorem 5.3.7 is said to be the **skeletalization** of G.

In order to generate 2-cliquish graphs, it suffices to begin with the skeletal graphs, and add edges when possible as in part (2) of Theorem 5.3.7. Figure 5.3.5 shows an example of this. A skeletal graph G is on the left. There are two pairs of elements that can be connected by edges as in part (2) of Theorem 5.3.7. This leads to the four 2-cliquish graphs whose skeletalization is G.

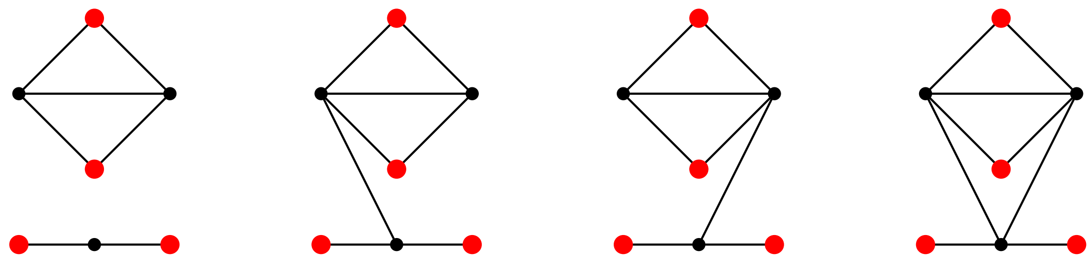

FIGURE 5.3.5: The four 2-cliquish graphs whose skeletalization is the graph on the left. The two in the middle are isomorphic, so they are considered the same unlabeled graph.

A *multigraph* is a graph that may contain multiple edges with the same pair of endpoints.

Theorem 5.3.9. *There is a bijection between pairs* (Γ, U)*, where* Γ *is a skeletal 2-cliquish graph with* n *vertices and* U *is a maximal independent set of* Γ*, and loopless multigraphs* $G = (V, \mathcal{E})$ *that satisfy* $|V| + |\mathcal{E}| = n$*.*

Proof. Let $\Gamma = (V, \mathcal{E})$ be a skeletal 2-cliquish graph with maximal independent set U, and consider the following construct of a loopless multigraph $(V', \mathcal{E}')$ from Γ: Let $V' = U$, and for each vertex $v \in V \setminus U$, construct an edge $e' \in \mathcal{E}'$, as follows: v has exactly two neighbors in U, say u_1 and u_2; let e' be an edge connecting u_1 and u_2. Thus $|V'|+|\mathcal{E}'| = |U|+|V \setminus U| = |V| = n$, and it is clear that $(V', \mathcal{E}')$ is loopless.

For the reverse direction, let $(V', \mathcal{E}')$ be a loopless multigraph with $|V'| + |\mathcal{E}'| = n$. Construct a graph $\Gamma = (V, \mathcal{E})$ whose vertices are in bijection with those of V' and $\mathcal{E}'$; let $U = V'$ and $V \setminus U = \mathcal{E}'$. Connect $v \in V \setminus U$ and $u \in U$ with an edge if u is an endpoint of v in V', and connect $v_1, v_2 \in V \setminus U$ with an edge if they share a common endpoint in V'. (We never connect two elements of U with an edge.) It is clear that $(V', \mathcal{E}')$ is 2-cliquish with maximal independent set U, and that this map is the inverse of the one in the other direction. ∎

Example 5.3.10. An example of this bijection can be seen in Figure 5.3.6. We start with

the multigraph M on the left and construct the skeletal 2-cliquish graph on the right. The vertices A, B, C, D, E of the multigraph correspond to the vertices a, b, c, d, e in the skeletal graph. The set $\{a, b, c, d, e\}$ is the maximal independent set of our skeletal graph. The edges e_1, e_2, e_3, e_4 correspond to the vertices v_1, v_2, v_3, v_4 of the skeletal graph. We use the multigraph to determine which two vertices in $\{a, b, c, d, e\}$ the other vertices are adjacent to. For example, the edge e_1 has endpoints A and B, so we add edges from v_1 to a and b. Lastly, whenever two vertices have a common neighbor in the independent set $\{a, b, c, d, e\}$, we must add an edge connecting them. Therefore, we place an edge between v_1 and v_2, another between v_1 and v_3, another between v_2 and v_3, and another between v_3 and v_4.

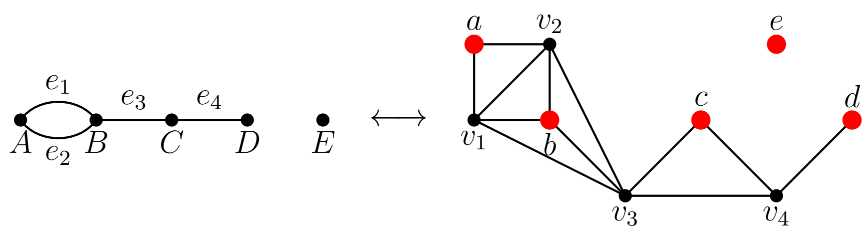

FIGURE 5.3.6: An example showing the bijection in Theorem 5.3.9.

The following corollary is clear using the bijection constructed in the proof of Theorem 5.3.9.

Corollary 5.3.11. *A pair (Γ, U) refers to a skeletal 2-cliquish graph Γ with maximal independent set U as in Theorem 5.3.9.*

1. *There is a bijection between pairs (Γ, U) such that $|\Gamma| = n$ and Γ has no isolated vertices, and loopless multigraphs $G = (V, \mathcal{E})$ with no isolated vertices that satisfy $|V| + |\mathcal{E}| = n$.*

2. *There is a bijection between pairs (Γ, U) with Γ connected and $|\Gamma| = n$, and connected loopless multigraphs $G = (V, \mathcal{E})$ that satisfy $|V| + |\mathcal{E}| = n$.*

3. There is a bijection between pairs (Γ, U) *satisfying* $|\Gamma| = n$ *and* $|U| = A$*, and loopless multigraphs* $G = (V, \mathcal{E})$ *that satisfy* $|V| = A$ *and* $|\mathcal{E}| = n - A$*.*

Note that if one is interested in generating 2-cliquish graphs without isolated vertices, it is enough to start with skeletal 2-cliquish graphs without isolated vertices. However, if one is interested in generating connected 2-cliquish graphs, it is not enough to begin with connected skeletal 2-cliquish graphs, as the skeletalization of a connected graph can be disconnected, as can be seen in Figure 5.3.5.

Chapter 6

Whirling injections, surjections, and generalizations

The results of Chapters 6 and 7 are joint work with James Propp and Tom Roby and will appear in a forthcoming paper.

6.1 m-injections and m-surjections

The main result of this chapter (Theorem 6.2.9) started as a conjecture from Propp about injective and surjective functions between finite sets. In addition to proving the conjecture, we have generalized it to m-injective functions and conjecture it for m-surjective functions, as defined in this section.

Definition 6.1.1. A function $f : S \to T$ is m**-injective** if every element $t \in T$ appears as an output of f at most m times, i.e., $\#f^{-1}(t) \leq m$ for every $t \in T$. An m-injective function is also called an m**-injection**.

Definition 6.1.2. A function $f : S \to T$ is m**-surjective** if every element $t \in T$ appears as an output of f at least m times, i.e., $\#f^{-1}(t) \geq m$ for every $t \in T$. An m-surjective function is also called an m**-surjection**.

We denote the set of all m-injective functions from S to T by $\mathrm{Inj}_m(S,T)$ and the set of all m-surjective functions from S to T by $\mathrm{Sur}_m(S,T)$. In our case, where the domain of functions is $[n]$ and the codomain is $[k]$, we write $\mathrm{Inj}_m(n,k)$ and $\mathrm{Sur}_m(n,k)$ to mean $\mathrm{Inj}_m([n],[k])$ and $\mathrm{Sur}_m([n],[k])$ respectively. We write functions $f : [n] \to [k]$ in one-line notation as $f(1)f(2)\cdots f(n)$. For example, the function $f(1) = 5$, $f(2) = 1$, $f(3) = 2$, $f(4) = 1$ is written 5121. (If we allow outputs with multiple digits, we can insert commas, e.g. $4, 6, 11, 10, 5, 6$. However, in the examples we will write out, the codomain is $[k]$ with $k \leq 9$.)

Despite the name, it should be noted that m-injections are not injective in general for $m \geq 2$. Such functions are called "width-m restricted functions" in [54], which discusses their enumeration. Injections are the same as 1-injections and surjections are the same as 1-surjections. Clearly, if $m_1 \leq m_2$, then $\mathrm{Inj}_{m_1}(S,T) \subseteq \mathrm{Inj}_{m_2}(S,T)$ and $\mathrm{Sur}_{m_2}(S,T) \subseteq \mathrm{Sur}_{m_1}(S,T)$.

Using well-known techniques of generating functions, such as those in [55] or [45, Ch. 5] it is straightforward to show that we can count m-injections via the generating function

$$\sum_{n=0}^{\infty} \#\,\mathrm{Inj}_m(n,k)\frac{x^n}{n!} = \left(\sum_{i=0}^{m} \frac{x^i}{i!}\right)^k.$$

The cardinality of $\mathrm{Sur}_m(n,k)$ is $k!\left\{ {n \atop k} \right\}_{\geq m}$ where $\left\{ {n \atop k} \right\}_{\geq m}$ represents number of partitions of $[n]$ into k blocks, where each block has cardinality at least m. These numbers are called the *m-associated Stirling numbers of the second kind* [24]. It is elementary to show that they

satisfy the recurrence

$$\left\{ {n \atop k} \right\}_{\geq m} = k \left\{ {n-1 \atop k} \right\}_{\geq m} + \binom{n-1}{m-1} \left\{ {n-m \atop k-1} \right\}_{\geq m}$$

and generating function

$$\sum_{n=0}^{\infty} \left\{ {n \atop k} \right\}_{\geq m} \frac{x^n}{n!} = \frac{1}{k!} \left(\sum_{i=m}^{\infty} \frac{x^i}{i!} \right)^k .$$

6.2 The whirling action

Definition 6.2.1. Let $\mathcal{F}$ denote a family of functions $f : [n] \to [k]$. Define a map $w_i : \mathcal{F} \to \mathcal{F}$, called **whirling at index** i, in the following way. Given $f \in \mathcal{F}$, repeatedly add 1 (mod k) to the value of $f(i)$ until we get a function in $\mathcal{F}$. The new function is $w_i(f)$.

Example 6.2.2. If $\mathcal{F} = \mathrm{Inj}_2(6,4)$ and $f = 422343$. Then to find $w_3(f)$, we first add 1 to $f(3) = 2$ and get the function 423343, which is not 2-injective since 3 appears as an output three times. So, we add 1 again to the third position and get 424343, which again is not 2-injective. Therefore, we add 1 (mod 4) again and get the 2-injection 421343, and thus $w_3(f) = 421343$.

Remark 6.2.3. The map w_i depends on which family of functions $\mathcal{F}$ we are using. For example,

- if $\mathcal{F} = \mathrm{Inj}_3(7,3)$, then $w_1(1221332) = 3221332$,
- if $\mathcal{F} = \mathrm{Sur}_1(7,3)$, then $w_1(1221332) = 2221332$,
- if $\mathcal{F} = \mathrm{Sur}_2(7,3)$, then $w_1(1221332) = 1221332$.

Whenever we talk about whirling, we must first make it clear what $\mathcal{F}$ is. While in Chapter 7 we will consider whirling on other families of functions, in this chapter $\mathcal{F}$ either refers to $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_m(n,k)$ for some $n,k,m \in \mathbb{P}$.

Remark 6.2.4. In the case that $\mathcal{F} = \mathrm{Sur}_m(n,k)$, the only possible reason we could not add 1 mod k to $f(i)$ is because the value $f(i)$ only appears as an output of the function m times. Therefore, in this case, w_i either adds 1 mod k to $f(i)$, or it leaves f alone. This is not the case for m-injective functions.

Definition 6.2.5. Let $\mathcal{F}$ denote either $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_m(n,k)$ for a given $n,k,m \in \mathbb{P}$. We define the **whirling** map, denoted $\mathbf{w} : \mathcal{F} \to \mathcal{F}$, as whirling at indices $1, 2, \ldots, n$ in that order, i.e., $\mathbf{w} = w_n \circ \cdots \circ w_2 \circ w_1$.

Example 6.2.6. Let $\mathcal{F} = \mathrm{Inj}_1(4,7)$. Then

$$2753 \xmapsto{w_1} 4753 \xmapsto{w_2} 4153 \xmapsto{w_3} 4163 \xmapsto{w_4} 4165$$

so $\mathbf{w}(2753) = 4165$.

Proposition 6.2.7. *On any family $\mathcal{F}$, the map w_i is invertible. Given $f \in \mathcal{F}$, repeatedly subtract 1 (mod k) from the value of $f(i)$ until we get a function in $\mathcal{F}$, and this new function is $w_i^{-1}(f)$. Therefore*

$$\mathbf{w}^{-1} = w_1^{-1} \circ w_2^{-1} \circ \cdots \circ w_n^{-1}.$$

So we can discuss orbits and homomesy for $\mathbf{w}$. Theorem 6.2.9 is the main result of this chapter.

Definition 6.2.8. For $j \in [k]$, define $\eta_j(f) = \#f^{-1}(\{j\})$ to be the number of times j appears as an output of the function f.

Theorem 6.2.9. *Fix $\mathcal{F}$ to be either $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_1(n,k)$ for a given $n,k,m \in \mathbb{P}$. Then under the action of $\mathbf{w}$ on $\mathcal{F}$, η_j is $\frac{n}{k}$-mesic for any $j \in [k]$; equivalently, $\eta_i - \eta_j$ is 0-mesic for all $i,j \in [k]$.*

The equivalence is clear since both statements mean that i and j appear as outputs of functions the same number of times across any orbit. We also conjecture this result for m-surjections in general.

Conjecture 6.2.10. *Let $\mathcal{F} = \mathrm{Sur}_m(n,k)$ for $m,n,k \in \mathbb{P}$. Then under the action of $\mathbf{w}$ on $\mathcal{F}$, η_j is $\frac{n}{k}$-mesic for any $j \in [k]$; equivalently, $\eta_i - \eta_j$ is 0-mesic for all $i,j \in [k]$.*

6.3 The proof of Theorem 6.2.9 for injections

In this section, we prove Theorem 6.2.9 when $\mathcal{F} = \mathrm{Inj}_1(n,k)$. Even though this is a special case of m-injections, it is easier to understand the proof in this simpler situation. We will utilize much of the same notation and terminology for the other cases.

Let $\mathcal{O}$ be an orbit under the action of $\mathbf{w}$ on $\mathcal{F}$. We draw a **board** for the orbit $\mathcal{O}$ by placing some $f \in \mathcal{O}$ on the top line. The function in row $i+1$ is $\mathbf{w}^i(f)$ for $i \in [0, \ell(\mathcal{O})-1]$, where $\ell(\mathcal{O})$ is the length of $\mathcal{O}$. For example, in Figure 6.3.1, we show a board for the orbit containing $f = 621$ in $\mathcal{F} = \mathrm{Inj}_1(3,6)$, so $\mathbf{w}(f) = 342$, $\mathbf{w}^2(f) = 563$, etc., and $\mathbf{w}^{10}(f) = f$.

Notice that within an orbit board, if f is a given line, then the "partially whirled element" $(w_i \circ \cdots \circ w_1)(f)$ is given by the first i numbers on the line below f and the last $n-i$ numbers of f. (In using the word "below", the function below the bottom line is the top line again, as we consider the orbit board to be cylindrical.) For instance, in the first two lines of Figure 6.3.1, $f = 621$, $w_1(f) = 321$, and $(w_2 \circ w_1)(f) = 341$.

We use the term **reading** the orbit board to refer to this action where we start at a

6	2	1
3	4	2
5	6	3
1	2	4
3	5	6
4	1	2
5	3	4
6	5	1
2	6	3
4	1	5

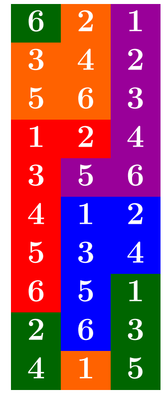

FIGURE 6.3.1: *Left:* The orbit under the action on $\mathbf{w}$ on $\mathrm{Inj}_1(3,6)$ containing $f = 621$.
Right: This same orbit but with colors representing the [6]-chunks.

certain position P of the orbit, and continue to the right until we reach the end of the line, and then go to the leftmost position of the line below and continue. This is because it is exactly like reading a book (except that continuing past the bottom line means returning to the top line). When we refer to a certain position being x positions "before" or "after" the position P, or say the "previous" or "next" position, we always mean in the reading order.

Definition 6.3.1. For a given position P in an orbit board, let (P, h) denote the position h places after P in the reading order (or $-h$ places before P if $h < 0$). Let $(P, [a, b]) = \{(P, h) | a \le h \le b\}$.

Example 6.3.2. Consider the following orbit on $\mathrm{Inj}_1(4, 5)$. Let P be the position in the second row and second column, shown below surrounded by a black rectangle. Then $(P, [1, 4])$ consists of the four positions circled in red. Also $(P, [0, 4])$ is $(P, [1, 4])$ together with P, while $(P, [-1, 2])$ is the second row. As the orbit is cylindrical, the bottom right corner is both $(P, 14)$ and $(P, -6)$. Note that P refers only to the *position*, not the value in that position. So $P \ne (P, 5)$ since they are different positions both containing the value 3. Similarly, we will never write, e.g., $P = 3$.

$$\begin{array}{cccc}
3 & 2 & 1 & 5 \\
4 & \boxed{\mathbf{3}} & 2 & 1 \\
5 & 4 & 3 & 2 \\
1 & 5 & 4 & 3 \\
2 & 1 & 5 & 4
\end{array}$$

Note that if P is in row i and column j, then $(P, [1, n])$ always consists of all positions to the right of P in row i, together with the leftmost j positions of row $i + 1$. Also note that $P([-n, -1])$ consists of all points left of P in row i together with the rightmost $n - j + 1$ positions of row $i - 1$.

Lemma 6.3.3. *Suppose i is in position P of a board for the $\mathbf{w}$-orbit $\mathcal{O}$ on $\mathrm{Inj}_1(n, k)$.*

1. *There is exactly one occurrence of $i + 1 \bmod k$ within $(P, [1, n])$.*

2. *There is exactly one occurrence of $i - 1 \bmod k$ within $(P, [-n, -1])$.*

Proof. To prove (1), suppose position P is in column j. So $f(j) = i$ for some $f \in \mathcal{O}$. Then $(P, [1, n])$ contains the multiset of outputs of $(w_j \circ \cdots \circ w_1)(f)$. If $(w_{j-1} \circ \cdots \circ w_1)(f)$ does not have $i + 1 \bmod k$ already as an output, then by definition w_j changes the output corresponding to the input j from i to $i + 1 \bmod k$. So there is an occurrence of $i + 1 \bmod k$ within $(P, [1, n])$. Since $(w_j \circ \cdots \circ w_1)(f)$ cannot have any output more than once, this occurrence is unique.

The proof of (2) is analogous to (1) using the inverse of whirling instead. ∎

Proof of Theorem 6.2.9 for $\mathcal{F} = \mathrm{Inj}_1(n, k)$. The idea of the proof is to partition any given orbit into $[k]$-**chunks** that contain every number $1, 2, \ldots, k$ exactly once. Within any orbit $\mathcal{O}$, we will assume without loss of generality that 1 appears as an output at least as often as any other number $2, \ldots, k$. This is because if $i > 1$ appears as an output more times than 1,

6	2	1	6	2	1	6	2	1	6	2	1	6	2	1	6	2	1
3	4	2	3	4	2	3	4	2	3	4	2	3	4	2	3	4	2
5	6	3	5	6	3	5	6	3	5	6	3	5	6	3	5	6	3
1	2	4	1	2	4	1	2	4	1	2	4	1	2	4	1	2	4
3	5	6	3	5	6	3	5	6	3	5	6	3	5	6	3	5	6
4	1	2	4	1	2	4	1	2	4	1	2	4	1	2	4	1	2
5	3	4	5	3	4	5	3	4	5	3	4	5	3	4	5	3	4
6	5	1	6	5	1	6	5	1	6	5	1	6	5	1	6	5	1
2	6	3	2	6	3	2	6	3	2	6	3	2	6	3	2	6	3
4	1	5	4	1	5	4	1	5	4	1	5	4	1	5	4	1	5

FIGURE 6.3.2: We demonstrate how to place a [6]-chunk on this $\mathbf{w}$-orbit on $\mathcal{F} = \mathrm{Inj}_1(3,6)$. In compartment i from left to right, the positions highlighted in green are the first i positions in the chunk. The entries in red are the next $n = 3$ positions after the one containing the i. We choose the unique $i+1$ among the red entries to be in the chunk.

then we can renumber $i, i+1, \ldots, k, 1, \ldots, i-1$ as $1, 2, \ldots, k$, and the outputs remain the same relative to each other (mod k).

Choose a 1 in the orbit board and call this position P_1. Then by Lemma 6.3.3(1), there exists a unique occurrence of 2 within $(P_1, [1, n])$; place this 2 in the same chunk as 1. For every i in a chunk in position P_i, choose the $i+1$ within $(P_i, [1, n])$. Continue this until the chunk contains $1, 2, \ldots, k$. Refer to Figure 6.3.2 for an example of this process. In each step shown, the red and bold entries are the next $n = 3$ positions after the position placed in the chunk.

To start a new chunk, we choose a 1 entry that is not already part of a chunk, and continue the same process. If Q is the position of the $i+1$ in a chunk, the i in the same chunk clearly must be in $Q([-n, -1])$. By Lemma 6.3.3(2), there is only one i entry that can be in the same chunk as the given $i+1$ entry. Thus, our $[k]$-chunks are disjoint.

Once every 1 in the orbit board is part of a completed $[k]$-chunk, there are no more entries not already in a chunk, since we assumed 1 appears as an output at least as often as

any other number. The chunking process shows that $1, 2, \ldots, k$ appear as outputs the same number of times in any orbit. ■

6.4 The proof of Theorem 6.2.9 for m-injections

Now we prove Theorem 6.2.9 for the case $\mathcal{F} = \mathrm{Inj}_m(n, k)$. We use a similar technique as for the $\mathrm{Inj}_1(n, k)$ case. We will again partition orbit boards into $[k]$-chunks, where each chunk contains $1, 2, \ldots, k$ and the instance of $i + 1$ within a chunk is at most n positions after the instance of i (in the reading order). Unlike in the $\mathrm{Inj}_1(n, k)$ case, there is no longer a *unique* way of partitioning the orbit into chunks; see Figure 6.4.1 for two different ways to partition the orbit of $\mathrm{Inj}_2(4, 4)$ containing 1441 into [4]-chunks. However, all that matters to prove Theorem 6.2.9 is the existence of a partitioning into $[k]$-chunks. Nonetheless, the proof becomes more complicated due to lack of uniqueness.

We will again use the notations (P, h) and $(P, [a, b])$ as we did for the injections proof.

Lemma 6.4.1. *Suppose i is in position P of a board for the orbit $\mathcal{O}$ of $\mathbf{w}$ on $\mathrm{Inj}_m(n, k)$.*

1. *There are at most m occurrences of $i + 1 \bmod k$ within $(P, [1, n])$.*

2. *If the position (P, n) directly below P does not contain $i + 1 \bmod k$, then there are exactly m occurrences of $i + 1 \bmod k$ within $(P, [1, n - 1])$.*

3. *There are at most m occurrences of $i - 1 \bmod k$ within $(P, [-n, -1])$.*

4. *If the position $(P, -n)$ directly above P does not contain $i - 1 \bmod k$, then there are exactly m occurrences of $i - 1 \bmod k$ within $(P, [-(n - 1), -1])$.*

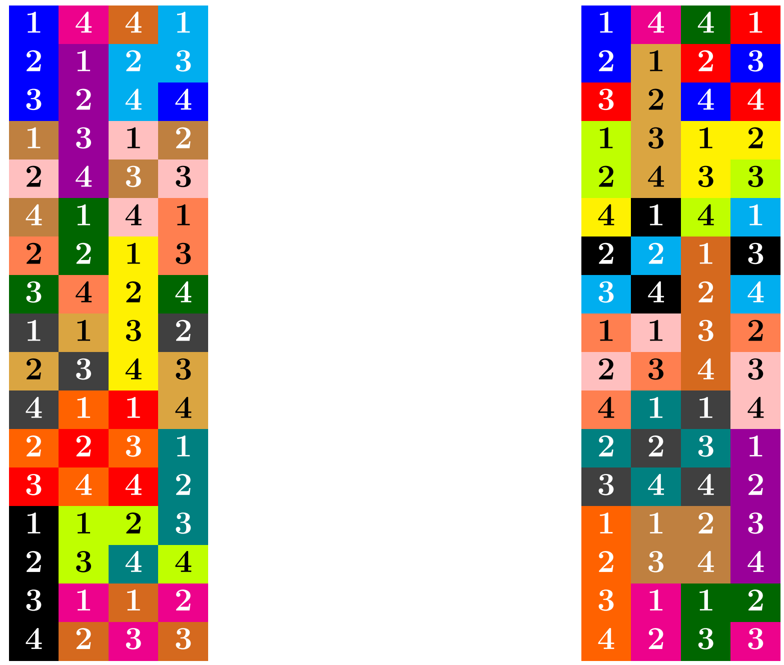

FIGURE 6.4.1: Two different ways to partition the same $\mathbf{w}$-orbit of $\mathrm{Inj}_2(4,4)$ into $[4]$-chunks.

Proof. Suppose position P is in column j. So $f(j) = i$ for some $f \in \mathcal{O}$, and $(P, [1, n])$ contains the multiset of outputs of $(w_j \circ \cdots \circ w_1)(f)$. By m-injectivity, there cannot be more than m occurrences of the $i+1 \bmod k$, proving (1).

If $(w_{j-1} \circ \cdots \circ w_1)(f)$ does not have $i+1 \bmod k$ already m times as an output, then by definition w_j changes the output corresponding to the input j from i to $i+1 \bmod k$. This proves (2).

The proofs of (3) and (4) are analogous to (1) and (2), using the inverse of whirling instead. ∎

Remark 6.4.2. For any position P in column j, $(P, [1, n])$ contains the multiset of outputs of $(w_j \circ \cdots \circ w_1)(f)$. Thus, any set of n consecutive positions (in the reading order) cannot contain more than m equal values.

Proof of Theorem 6.2.9 for $\mathcal{F} = \mathrm{Inj}_m(n, k)$. By the same relabeling argument from the in-

jections proof, we will assume without loss of generality that within any given orbit $\mathcal{O}$, 1 appears as an output at least as often as any other number $2, \ldots, k$.

Choose a 1 in the orbit board and call this position P_1. Then by Lemma 6.4.1(2), there exists at least one occurrence of 2 in $(P_1, [1, n])$; pick such a 2 and place it in the same chunk as 1. For every i in a chunk, call the position P_i and select an $i+1$ within $(P_i, [1, n])$ to be in the same chunk. Continue this until the chunk contains $1, 2, \ldots, k$.

To start a new chunk, we choose a 1 entry that is not already part of a chunk, and wish to continue the same process. However, unlike for 1-injections, there is not necessarily a unique occurrence of $i+1$ within the next n positions after an i, nor a unique occurrence of $i-1$ within the previous n positions before an i. When placing an $i+1$ entry in the same chunk as i in position P, then we want to choose $i+1$ that is not already part of a chunk. When $(P, [1, n])$ contains such an $i+1$ entry, we choose one of them. However, such an $i+1$ entry may not exist depending on how we chose earlier chunks.

See the left side of Figure 6.4.2 for an example of this problem. Let P be the position of the 2 in the purple chunk. Then the only 3 in $(P, [1, n])$ is already part of the brown chunk. In this case, we reassign the part of the brown chunk starting with 3 to be in the purple chunk, and then continue where the purple chunk is complete and we now attempt to complete the brown chunk.

In general, suppose for a given i in position P, all instances of $i+1$ within $(P, [1, n])$ are already part of chunks. Then we claim that at least one of these $i+1$ entries is in the same chunk as an i entry in $(P, [1, n-1])$. To explain this we consider two cases.

Case 1: The entry directly below position P is $i+1$. Let $Q = (P, n)$ be this position below P. Then the $i+1$ in position Q is already in a chunk with an i entry in $(Q, [-n, -1]) = (P, [0, n-1])$. However, we know the i in position P is not already in a chunk with an $i+1$. So the $i+1$ in position Q must be in the same chunk as an i entry in

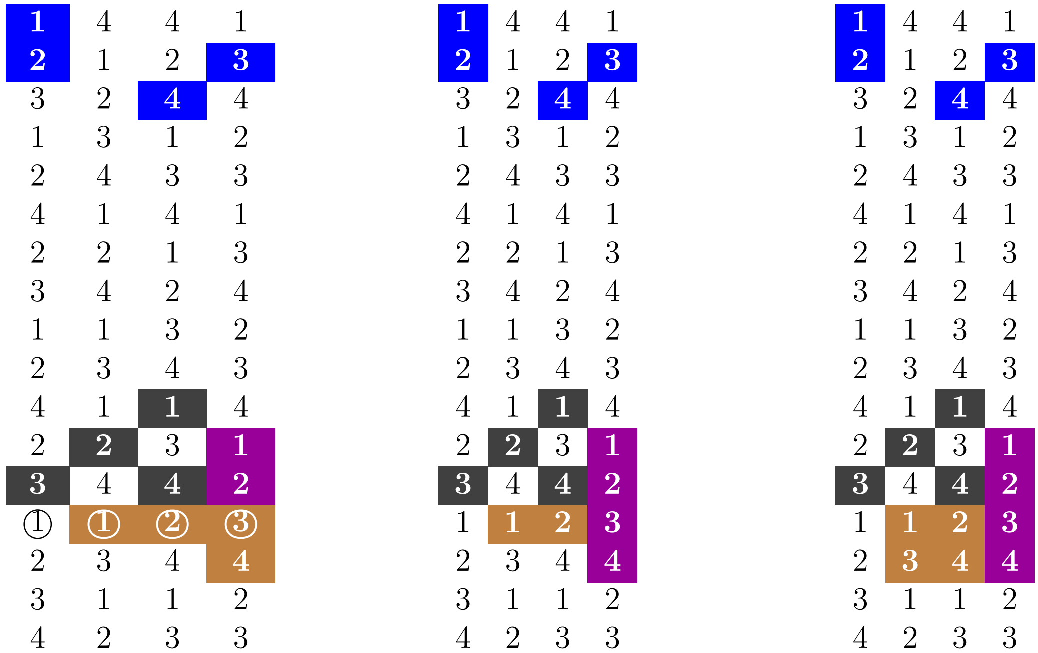

FIGURE 6.4.2: In the left diagram, we cannot complete the purple chunk because the only position containing 3 within the next $n = 4$ positions after the purple chunk's 2 is in the brown chunk already. Thus, we have to take the end of the brown chunk starting from the 3, place it in the purple chunk (middle), and then complete the brown chunk (right).

$(P, [1, n-1])$.

Case 2: The entry directly below position P is not $i+1$. Then Lemma 6.4.1(2) implies that there are m occurrences of $i+1$ within $(P, [1, n-1])$. Each of them is in a chunk with an i that is at most n positions before, and thus must be in $(P, [-(n-1), n-2])$. If all of these are in $(P, [-(n-1), -1])$, then there would be m different i entries in $(P, [-(n-1), -1])$. Adding in the i in position P, this implies there are $m+1$ different i entries in $(P, [-(n-1), 0])$, contradicting Remark 6.4.2. So at least one of the occurrences of $i+1$ within $(P, [1, n-1])$ is in a chunk with an i in $(P, [1, n-1])$.

Suppose that we get stuck in partitioning the orbit board into $[k]$-chunks. In this case, when creating a chunk C through $\{1, \ldots, i\}$, we have i in position P and all $i+1$ entries in

$(P, [1, n])$ are already part of chunks. Then we will choose a chunk C' for which both the i and $i+1$ entries are within $(P, [1, n])$, which exists by the above argument. We will take the positions of the $i+1, \ldots, n$ entries and make them part of the chunk C instead of C'. Then it is just as if C was created previously, and we now must complete C'. This can cause a chain reaction of chunk reassignment.

At any point, let Q be the position of $i+1$ for which we reassign $i+1, \ldots, k$ to be in a different chunk. Then we are always choosing to place it with an i that is earlier (in the reading order) within $(Q, [-n, -1])$. Since there are only finitely many i entries that can be in the same chunk as the given $i+1$, this process will terminate at some point, showing that we will eventually partition the orbit board into chunks as required. ■

Remark 6.4.3. While partitioning the orbit board into chunks, suppose whenever we have an i in position P, we always choose the *last* occurrence of $i+1$ within $(P, [1, n])$ not already within a chunk. Then if we were to have to reassign, we would be unable to do so while choosing a chunk for which both the i and $i+1$ entries are within $(P, ([1, n]))$. Since we know we can always reassign and the process ends, always choosing the last possible $i+1$ in $(P, [1, n])$ means we can partition the entire orbit without reassigning. However, the only known way to prove this is by allowing reassignments and then seeing this scenario describes a special case.

6.5 The proof of Theorem 6.2.9 for surjections

In this section, we prove Theorem 6.2.9 in the case where $\mathcal{F} = \mathrm{Sur}_1(n, k)$.

Recalling Remark 6.2.4, w_i either adds 1 mod k to the value of $f(i)$ or it leaves f alone. Thus, given any $i \in [n]$ and $f \in \mathrm{Sur}_1(n, k)$, $\mathbf{w}$ either adds 1 mod k to $f(i)$ or does not change

the value $f(i)$. Therefore, in an orbit board, such as the one in Figure 6.5.1, each entry in a column is either the same or 1 greater (mod k) as the one directly above it.

Definition 6.5.1. Call a position P in an orbit board a **red light** if it has the same value as the position below it.

This term comes from the analogy with traffic lights. As we look down a column, we always add 1 to the value, except we have to stop when a position is a red light. The red lights are surrounded by red circles in Figure 6.5.1.

3	1	1	1	4	4	2	4
3	2	2	1	1	1	3	4
4	3	2	2	2	1	4	1
1	3	3	3	2	2	4	2
1	4	4	3	3	3	1	2
2	1	4	4	4	3	2	3

FIGURE 6.5.1: An example $\mathbf{w}$-orbit of $\mathrm{Sur}_1(8,4)$ containing $f = 31114424$.

Lemma 6.5.2. *Let P be a position in an orbit board of $\mathrm{Sur}_1(n,k)$, and let i be the value in that position.*

1. *If $(P,[1,n-1])$ does not contain i in any position, then the position (P,n) directly below P contains the value i.*

2. *If $(P,[1,n-1])$ contains the value i in some position, then the position (P,n) directly below P contains the value $i+1 \bmod k$.*

Proof. Suppose P is in column j and let f be the function on the row with P. So $f(j)=i$. Then $(P,[0,n-1])$ consists of the multiset of entries of $(w_{j-1}\circ\cdots\circ w_1)(f)$.

If $(P, [1, n-1])$ does not contain i in any position, then $(w_{j-1} \circ \cdots \circ w_1)(f)$ only contains i as an output once. So to maintain surjectivity, $((w_j \circ \cdots \circ w_1)(f))(j) = i$, and hence (P, n) contains the value i.

On the other hand, if $(P, [1, n-1])$ contains i in any position, then $(w_{j-1} \circ \cdots \circ w_1)(f)$ contains i as an output more than once. Thus applying w_j changes the output corresponding to j from i to $i+1 \bmod k$, since we still get a surjective function. In this case, (P, n) contains the value $i + 1 \bmod k$. ■

Note that Lemma 6.5.2 implies that a position is a red light if and only if it does not contain the same value as any of the next $n - 1$ positions.

Lemma 6.5.3. *Let P be a red light position that contains i. Let Q be the last position (in the reading order) in $(P, [1, n-1])$ that contains $i + 1 \bmod k$. Then Q is a red light. In this case, P is the earliest position containing i within $(Q, [-(n-1), 1])$.*

Proof. We will refer to $i + 1 \bmod k$ as $i + 1$. Via Lemma 6.5.2, it suffices to prove that $(Q, [1, n-1])$ does not contain $i+1$ in any position. Every position in $(Q, [1, n-1])$ is either

- in $(P, [1, n-1])$,
- the position below P, or
- the position directly below R for some R between P and Q.

Since Q is the last position in $(P, [1, n-1])$ that contains $i+1$, any position simultaneously in both $(Q, [1, n-1])$ and $(P, [1, n-1])$ cannot contain $i+1$. The position below P contains i since P is a red light. Let R be a position between P and Q. If the position below R contains $i + 1$, then R contains either i or $i + 1$. Since P is a red light, no position in $(P, [1, n-1])$ contains i, so R cannot contain i. If R contains $i+1$, then R is not a red light

because $Q \in (R, [1, n-1])$. So the position below R cannot contain $i+1$. Thus, there are no positions in $(Q, [1, n-1])$ containing $i+1$.

By using the inverse of $\mathbf{w}$, an analogous proof explains that P is the earliest position containing i within $(Q, [-(n-1), 1])$. ■

Proof of Theorem 6.2.9 for $\mathcal{F} = \mathrm{Sur}_1(n, k)$. Within any column in an orbit board, if we skip over the red light positions, then every entry is 1 mod k more than the entry above it. Thus by the cyclic nature of the columns, the entries in the non-red light positions of any column are equidistributed between $1, 2, \ldots, k$. Therefore, it suffices to show that the entries in the red light positions in the orbit are equidistributed between $1, 2, \ldots, k$.

If there are no red light positions in the orbit, then we are done. If there is a red light position, pick one and call in P. Let i be the entry in P and circle it. Then by Lemma 6.5.3, we can pick the last $i+1$ mod k entry in $(P, [1, n-1])$ and call the position Q. Then we circle Q, which is a red light. Then using Lemma 6.5.3, the last $i+2$ mod k entry in $(Q, [1, n-1])$ is a red light, that we can circle. We can continue this and by the finiteness of positions in the orbit, we will eventually return to P in which case we stop. This chain of red lights clearly has the entries $1, 2, \ldots, k$ equidistributed by construction.

Now we look at the orbit again. If there are no red lights not already circled, we are done. Otherwise, pick a red light, circle it, and begin the same process again. By Lemma 6.5.3, given any red light Q, there is a unique position P containing i for which Q is the last position containing $i+1$ mod k in $(P, [1, n-1])$. Therefore, we will not circle a position circled in the previous chain.

We will continue this process until there are no more red lights to circle. Thus the circled red light positions in the orbit will be equidistributed between $1, 2, \ldots, k$. ■

6.6 Consequences of the homomesy

Let $\mathcal{F}$ be either $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_1(n,k)$. Then given any $j \in [k]$, $\eta_j(f)$ is always an integer. Thus, Theorem 6.2.9 leads to a corollary about orbit sizes for the case when the average value of η_j across every orbit is not an integer.

Corollary 6.6.1. *Suppose $\mathcal{F}$ is either $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_1(n,k)$. The size of any $\mathbf{w}$-orbit of $\mathcal{F}$ is a multiple of $k/\gcd(n,k)$.*

This is similar to Corollary 5.1.8 for toggling noncrossing partitions in that we can use homomesy to prove a property of the orbit sizes for which no other proof is known.

Proof. Theorem 6.2.9 says that the average value of η_j across any $\mathbf{w}$-orbit is n/k. When reduced to lowest terms, the denominator of n/k is $k/\gcd(n,k)$. Thus, the size of any orbit must be a multiple of $k/\gcd(n,k)$ in order for the average value of η_j to be n/k. ■

Theorem 6.2.9 (and therefore Corollary 6.6.1 also) extend if we replace $\mathbf{w}$ by any product of the w_i maps, each used exactly once, in some order. (These are analogous to extending to Coxeter elements in earlier chapters, but we do not use that term here since we are not in the quotient of a Coxeter group.)

Theorem 6.6.2. *Let π be a permutation on $[n]$ and $\mathbf{w}_\pi := w_{\pi(n)} \circ \cdots \circ w_{\pi(2)} \circ w_{\pi(1)}$. Fix $\mathcal{F}$ to be either $\mathrm{Inj}_m(n,k)$ or $\mathrm{Sur}_1(n,k)$ for a given $n,k,m \in \mathbb{P}$. Then under the action of $\mathbf{w}_\pi$ on $\mathcal{F}$, η_j is $\frac{n}{k}$-mesic for any $j \in [k]$.*

Proof. Notice that for $\mathcal{F} = \mathrm{Inj}_m(n,k)$ or $\mathcal{F} = \mathrm{Sur}_m(n,k)$, we have

$$f \in \mathcal{F} \iff f \circ \pi \in \mathcal{F}.$$

Therefore whirling f at index $\pi(i)$ is like whirling $f \circ \pi$ at index i, i.e., $w_i(f \circ \pi) = w_{\pi_i}(f) \circ \pi$. So $\mathbf{w}(f \circ \pi) = \mathbf{w}_\pi(f) \circ \pi$. Thus for any $\mathbf{w}_\pi$-orbit $\mathcal{O} = (f_1, f_2, \ldots, f_\ell)$, there exists a corresponding $\mathbf{w}$-orbit $\mathcal{O}' = (f'_1, f'_2, \ldots, f'_\ell)$ such that $f'_i = f_i \circ \pi$ for all i. Since $\eta_j(f \circ \pi) = \eta_j(f)$, Theorem 6.2.9 implies η_j is also $\frac{n}{k}$-mesic on orbits of $\mathbf{w}_\pi$. ■

Refer to Figure 6.6.1 for an illustration of the idea in the above proof.

f	142
$\mathbf{w}_\pi(f)$	365
$\mathbf{w}_\pi^2(f)$	421
$\mathbf{w}_\pi^3(f)$	543
$\mathbf{w}_\pi^4(f)$	615
$\mathbf{w}_\pi^5(f)$	236
$\mathbf{w}_\pi^6(f)$	451
$\mathbf{w}_\pi^7(f)$	612
$\mathbf{w}_\pi^8(f)$	324
$\mathbf{w}_\pi^9(f)$	536

g	124
$\mathbf{w}(g)$	356
$\mathbf{w}^2(g)$	412
$\mathbf{w}^3(g)$	534
$\mathbf{w}^4(g)$	651
$\mathbf{w}^5(g)$	263
$\mathbf{w}^6(g)$	415
$\mathbf{w}^7(g)$	621
$\mathbf{w}^8(g)$	342
$\mathbf{w}^9(g)$	563

FIGURE 6.6.1: Left: The orbit of $\mathbf{w}_\pi$ on $\mathrm{Inj}_1(3,6)$ containing $f = 142$ for the permutation $\pi = 132$. Right: The orbit of $\mathbf{w}$ on $\mathrm{Inj}_1(3,6)$ containing $g = 124 = f \circ \pi$. Notice that, as in the proof of Theorem 6.6.2, $\mathbf{w}^i(g) = \mathbf{w}^i(f) \circ \pi$. Thus, we can transform the $\mathbf{w}_\pi$-orbit on the left to the $\mathbf{w}$-orbit on the right by swapping the second and third columns, so it is clear that since $1, 2, \ldots, 6$ appear as outputs of functions equally often in the $\mathbf{w}$-orbit, they do so for the $\mathbf{w}_\pi$-orbit too.

For the specific case $\mathcal{F} = \mathrm{Inj}_m(n, 2)$, we can restate our homomesy result in terms of toggle groups. Like many results in Chapters 3, 4, and 5, Theorem 6.6.2 leads to a homomesy result for Coxeter elements in a toggle group over subsets of $[n]$ whose cardinality is restricted between r and $n - r$.

Let $n \in \mathbb{P}$ and $r \in \mathbb{N}$ with $r \le n/2$. We consider toggles where the ground set $E = [n]$ and $\mathcal{L}_r(n) := \{X \subseteq [n] \mid r \le \#X \le n - r\}$ is the set of allowed subsets. Then for each

$e \in [n]$, the toggle $t_e : \mathcal{L}_r(n) \to \mathcal{L}_r(n)$ is

$$t_e(X) = \begin{cases} X \cup \{e\} & \text{if } e \notin X \text{ and } \#X \leq n-r-1, \\ X \setminus \{e\} & \text{if } e \in X \text{ and } \#X \geq r+1, \\ X & \text{otherwise.} \end{cases}$$

Theorem 6.6.3. *Let π be a permutation on $[n]$ and $T_\pi := t_{\pi(n)} \circ \cdots \circ t_{\pi(2)} \circ t_{\pi(1)}$. Then under the action of T_π on $\mathcal{L}_r(n)$, the cardinality statistic is $n/2$-mesic.*

Proof. Let $\mathcal{F} = \mathrm{Inj}_{n-r}(n, 2)$. Any function $f \in \mathcal{F}$ is clearly associated with a subset $S(f) \subseteq [n]$ with cardinality between r and $n-r$, given by $S(f) := \{i \in [n] \mid f(i) = 1\}$. This relation goes both ways so S is a bijection. It is straightforward to see that $t_i \circ S = S \circ w_i$ shown in the commutative diagram below.

$$\begin{tikzcd} \mathrm{Inj}_{n-r}(n,2) \arrow[r, "S"] \arrow[d, "w_i"'] & \mathcal{L}_r(n) \arrow[d, "t_i"] \\ \mathrm{Inj}_{n-r}(n,2) \arrow[r, "S"'] & \mathcal{L}_r(n) \end{tikzcd}$$

So $T_\pi \circ S = S \circ \mathbf{w}_\pi$, where $\mathbf{w}_\pi := w_{\pi(n)} \circ \cdots \circ w_{\pi(2)} \circ w_{\pi(1)}$ as in Theorem 6.6.2. Thus any T_π-orbit on $\mathcal{L}_r(n)$ can be written as $(S(f_1), S(f_2), \ldots, S(f_\ell))$ where $(f_1, f_2, \ldots, f_\ell)$ is a $\mathbf{w}_\pi$-orbit of $\mathcal{F}$. Via Theorem 6.6.2, η_1 has average $n/2$ on $(f_1, f_2, \ldots, f_\ell)$. Since $\eta_1(f) = \#S(f)$, the average cardinality in the T_π-orbit $(S(f_1), S(f_2), \ldots, S(f_\ell))$ is $n/2$. ∎

Chapter 7

Whirling other families of functions

The definition of the whirling maps w_i and $\mathbf{w}$ from Section 6.2 make sense for any family of functions $\mathcal{F}$ between finite domain $[n]$ and codomain $[k]$. Thus it is natural to consider these for other families $\mathcal{F}$ of combinatorial interest. In Section 7.1, we consider $\mathbf{w}$ on parking functions, leading to a previously unpublished homomesy result of Nathan Williams. In Section 7.2, we consider $\mathbf{w}$ on restricted growth words.

7.1 Whirling parking functions

Parking functions, originally defined by Konheim and Weiss [26] in 1966, have been studied extensively in the years since. The main result of this section, Theorem 7.1.5, is due to Nathan Williams.

Definition 7.1.1. A **parking function** is an ordered tuple $(a_1, a_2, \ldots, a_n)$ of integers in $[n]$ such that if $A_1 \leq A_2 \leq \cdots \leq A_n$ is the weakly increasing rearrangement of $a_1, a_2, \ldots, a_n$, then $A_i \leq i$ for all $i \in [n]$.

Considering parking functions as functions $f : [n] \to [n]$ where $f(i) = a_i$ in the above definition, then the criterion to be a parking function is equivalent to $\#f^{-1}([i]) \geq i$ for all $i \in [n]$. We will write parking functions in one-line notation $f(1)f(2)\ldots f(n) = a_1a_2\cdots a_n$, assuming $n \leq 9$, which is the case in our examples. The set of parking functions $f : [n] \to [n]$ is denoted $\mathrm{Park}(n)$.

Example 7.1.2. For $n = 6$, 355121 is a parking function, but 145641 is not since there are only two entries ≤ 3.

Whirling is defined the same way on $\mathrm{Park}(n)$ as on any family of functions, but now our domain and codomain are the same. Thus, w_i adds 1 mod n repeatedly to the value of $f(i)$ until we obtain a function in $\mathrm{Park}(n)$. By the above criterion, if we cannot change a value from j to $j+1$, then it is because $\#f^{-1}([j]) = j$. Therefore, that value cannot be changed to anything greater than j. So it will become 1 in the definition of w_i. This leads to the following proposition whose proof is straightforward.

Proposition 7.1.3. *Let $f \in \mathrm{Park}(n)$. Then*

$$(w_i(f))(i) = \begin{cases} f(i)+1 & \text{if } \#f^{-1}([f(i)]) > f(i) \\ 1 & \text{if } \#f^{-1}([f(i)]) = f(i) \end{cases}.$$

Example 7.1.4. In Park(4),

$$1332 \xmapsto{w_1} 1332 \xmapsto{w_2} 1432 \xmapsto{w_3} 1412 \xmapsto{w_4} 1413$$

so $\mathbf{w}(1332) = 1413$.

Theorem 7.1.5 (Nathan Williams)**.** *Let $n \geq 2$ and consider orbits of $\mathbf{w}$ on* $\mathrm{Park}(n)$.

1. *Every orbit has size $n + 1$.*

2. *Given $i \in [n]$, every orbit has exactly two functions $f \in \mathrm{Park}(n)$ for which $f(i) = 1$.*

We will prove Theorem 7.1.5 later in this section after we discuss the necessary theory. It is well-known that $\#\,\mathrm{Park}(n) = (n+1)^{n-1}$. Konheim and Weiss proved this originally, but a more straightforward proof (included in many combinatorics texts) is due to Pollak [36]. Thus, Theorem 7.1.5 implies that there are $(n+1)^{n-2}$ orbits under $\mathbf{w}$.

Notice that if we let

$$I_{i\mapsto 1}(f) = \begin{cases} 1 & \text{if } f(i) = 1 \\ 0 & \text{if } f(i) \neq 1 \end{cases}$$

then from Theorem 7.1.5, $I_{i\mapsto 1}$ is $2/(n+1)$-mesic on $\mathbf{w}$-orbits for all $i \in [n]$.

Williams's proof of Theorem 7.1.5 involves Stanley's bijection between parking functions on $[n]$ and factorizations of the cycle $(n+1, n, \ldots, 2, 1) \in \mathfrak{S}_{n+1}$ as the product of n transpositions [44], [6][1]. Stanley showed that there is a bijection between parking functions $(a_1, a_2, \ldots, a_n)$ and factorizations $(a_1b_1)(a_2b_2)\cdots(a_nb_n)$ of $(n+1, n, \ldots, 2, 1)$, in which a_i is the lesser number in the i^{th} cycle.

Example 7.1.6. Since 1332 is a parking function, there exist unique $b_1 > 1, b_2 > 3, b_3 > 3, b_4 > 2$ such that $(1b_1)(3b_2)(3b_3)(2b_4) = (54321)$. We can see by trial and error that this factorization is $(15)(34)(35)(23)$. However, there are no $b_1 > 3, b_2 > 4, b_3 > 3, b_4 > 1$ such that $(3b_1)(4b_2)(3b_3)(1b_4) = (54321)$ because $3431 \notin \mathrm{Park}(4)$.

See Figure 7.1.1 for the four orbits on Park(3), as well as the corresponding factorization of (4321) for each parking function. The reader can easily verify Theorem 7.1.5 holds for

[1] The norm in the literature for this bijection is to use $(1, 2, \ldots, n, n+1)$. However, we use $(n+1, n, \ldots, 2, 1)$ because it makes our proof simpler while using the convention that a product of cycles is performed right-to-left. If the reader would alternatively prefer to use $(1, 2, \ldots, n, n+1)$, then consider the products to be left-to-right instead.

$111 \leftrightarrow (12)(13)(14)$	$112 \leftrightarrow (13)(14)(23)$	$211 \leftrightarrow (23)(12)(14)$	$131 \leftrightarrow (12)(34)(13)$
$221 \leftrightarrow (23)(24)(12)$	$213 \leftrightarrow (24)(12)(34)$	$321 \leftrightarrow (34)(23)(12)$	$212 \leftrightarrow (23)(14)(24)$
$312 \leftrightarrow (34)(13)(23)$	$121 \leftrightarrow (13)(23)(14)$	$132 \leftrightarrow (14)(34)(23)$	$311 \leftrightarrow (34)(12)(13)$
$123 \leftrightarrow (14)(24)(34)$	$231 \leftrightarrow (24)(34)(12)$	$113 \leftrightarrow (12)(14)(34)$	$122 \leftrightarrow (14)(23)(24)$

FIGURE 7.1.1: Each of the four columns represents a $\mathbf{w}$-orbit on Park(3). Within any column, applying $\mathbf{w}$ to any parking functions gives the one below it, and applying $\mathbf{w}$ to the parking function on the bottom gives the top. For each parking function, we also show the corresponding factorization of (4321) into transpositions.

these orbits.

Lemma 7.1.7. *Define* $\overline{w} : \mathrm{Park}(n) \to \mathrm{Park}(n)$ *to be the function that applies* w_1 *and then moves the first digit of the parking function to the end. Then* $\overline{w}^n = \mathbf{w}$.

Proof. Notice from the definition of parking function that $(a_1, a_2, \ldots, a_n)$ is a parking function if and only if any rearrangement is as well. When we apply $\overline{w}$ for the i^{th} time, the digit originally i places from the left is now the first digit. So applying $\overline{w}^i$ produces the same result as applying $w_i \circ \cdots \circ w_2 \circ w_1$ except that the initial substring of i digits has been moved to the end. Compare Examples 7.1.4 and 7.1.8. ∎

Example 7.1.8. In Park(4),

$$1332 \overset{\overline{w}}{\longmapsto} 3321 \overset{\overline{w}}{\longmapsto} 3214 \overset{\overline{w}}{\longmapsto} 2141 \overset{\overline{w}}{\longmapsto} 1413$$

so $\mathbf{w}(1332) = 1413$.

Instead of considering $\mathbf{w}$ to be the composition of n different maps, we will now study the single map $\overline{w}$, which can be considered an "n^{th} root" of $\mathbf{w}$. The following lemma is key to the proof of Theorem 7.1.5.

Lemma 7.1.9. *Define* $\pi^c := (n+1, n, \ldots, 2, 1)^{-1}\pi(n+1, n, \ldots, 2, 1)$ *to be the conjugation of* $\pi \in \mathfrak{S}_{n+1}$ *by* $(n+1, n, \ldots, 2, 1)$. *Suppose the factorization of* $(n+1, n, \ldots, 2, 1)$ *corresponding to* $(a_1, a_2, \ldots, a_n) \in \mathrm{Park}(n)$ *is* $(a_1b_1)(a_2b_2)\cdots(a_nb_n)$. *Then the factorization corresponding to* $\overline{w}(a_1, a_2, \ldots, a_n)$ *is* $(a_2b_2)\cdots(a_nb_n)(a_1b_1)^c$.

Proof. (Nathan Williams) First it is clear that applying the conjugation to a 2-cycle (a_1b_1) adds 1 mod $(n+1)$ to both a_1 and b_1. Also notice that if $(a_1b_1)(a_2b_2)\cdots(a_nb_n) = (n+1, n, \ldots, 2, 1)$, then

$$\begin{aligned}(a_2b_2)\cdots(a_nb_n)(a_1b_1)^c &= (a_1b_1)(n+1, n, \ldots, 2, 1)(n+1, n, \ldots, 2, 1)^{-1}(a_1b_1)(n+1, n, \ldots, 2, 1) \\ &= (n+1, n, \ldots, 2, 1).\end{aligned}$$

Since $(a_2b_2)\cdots(a_nb_n)(a_1b_1)^c$ is a factorization of $(n+1, n, \ldots, 2, 1)$ into n transpositions, the smaller entries of each cycle form a parking function on $[n]$.

Case 1: $b_1 \neq n+1$. Then $(a_2, a_3, \ldots, a_n, a_1+1)$ is the parking function corresponding to $(a_2b_2)\cdots(a_nb_n)(a_1b_1)^c$. Since $(a_2, a_3, \ldots, a_n, a_1+1)$ is a parking function, it must be $\overline{w}(a_1, a_2, \ldots, a_n)$.

Case 2: $b_1 = n+1$. Then $(a_1b_1)^c = (1, a_1+1)$. Suppose $(a_2, a_3, \ldots, a_n, a_1+1)$ is a parking function. Then it corresponds to $(a_2b_2')(a_3b_3')\cdots(a_nb_n')(a_1+1, b_1'+1)$ where $b_1' \leq n$. Since $(a_1b_1')^c = (a_1+1, b_1'+1)$, $(a_1b_1')(a_2b_2')(a_3b_3')\cdots(a_nb_n') = (n+1, n, \ldots, 2, 1)$. Since $b_1' \leq n$, this gives two different factorizations for which the smaller elements in each cycles are $a_1, a_2, \ldots, a_n$ in order. Since each $(a_1, a_2, \ldots, a_n) \in \mathrm{Park}(n)$ corresponds to a unique factorization of $(n+1, n, \ldots, 2, 1)$, this is a contradiction. So $(a_2, a_3, \ldots, a_n, a_1+1)$ is not a parking function which means $\overline{w}(a_1, a_2, \ldots, a_n) = (a_2, \ldots, a_n, 1)$. This corresponds to $(a_2b_2)\cdots(a_nb_n)(1, a_1+1) = (a_2b_2)\cdots(a_nb_n)(a_1b_1)^c$. ■

We are now ready to prove Theorem 7.1.5.

Proof of Theorem 7.1.5. (Nathan Williams) From Lemmas 7.1.7 and 7.1.9, if $f \in \mathrm{Park}(n)$ corresponds to $(a_1b_1)(a_2b_2)\dots(a_nb_n)$, then $\mathbf{w}(f)$ corresponds to $(a_1b_1)^c(a_2b_2)^c\dots(a_nb_n)^c$. Since each conjugation $(ab) \mapsto (ab)^c$ adds 1 mod $(n+1)$ to a and b, it takes $n+1$ conjugations to return to (ab), except when $|a-b| = (n+1)/2$ in which case it takes $(n+1)/2$ conjugations. However, if $b_i = a_i + (n+1)/2$ for all i, then $(n+1, n, \dots, 2, 1)$ would be a product of disjoint transpositions. This would imply $(n+1, n, \dots, 2, 1)$ has order 1 or 2, which is not the case since $n \geq 2$. Therefore, at least one cycle (a_ib_i) requires the full $n+1$ conjugations to return to itself. So f requires $n+1$ iterations of $\mathbf{w}$ to return to f.

Given any cycle (a_ib_i) and $j \in [n+1]$, as we repeatedly add 1 mod $(n+1)$ to each entry, there will be exactly two times in the orbit for which this cycle contains j. In the case $j = 1$, 1 must be the lesser entry in the cycle. So there is a 1 in position i for exactly two parking functions in the $\mathbf{w}$-orbit. ■

Using an analogous proof, we can generalize Theorem 7.1.5 to maps that whirl at every index, but in an arbitrary order.

Theorem 7.1.10. *Let π be a permutation on $[n]$, where $n \geq 2$, and let $\mathbf{w}_\pi := w_{\pi(n)} \circ \dots \circ w_{\pi(2)} \circ w_{\pi(1)}$.*

1. *Every $\mathbf{w}_\pi$-orbit has size $n+1$.*
2. *Given $i \in [n]$, every $\mathbf{w}_\pi$-orbit contains exactly two functions $f \in \mathrm{Park}(n)$ for which $f(i) = 1$.*

Proof. Note that any $f : [n] \to [n]$ satisfies $f \in \mathrm{Park}_n$ if and only if $f \circ \pi \in \mathrm{Park}_n$. Therefore whirling f at index $\pi(i)$ is the same as whirling $f \circ \pi$ at index i, i.e., $w_i(f \circ \pi) = w_{\pi_i}(f) \circ \pi$. So $\mathbf{w}(f \circ \pi) = \mathbf{w}_\pi(f) \circ \pi$. Thus for any $\mathbf{w}_\pi$-orbit $\mathcal{O} = (f_1, f_2, \dots, f_\ell)$, there exists a corresponding $\mathbf{w}$-orbit $\mathcal{O}' = (f_1', f_2', \dots, f_\ell')$ such that $f_j' = f_j \circ \pi$ for all i. By Theorem 7.1.5, the $\mathbf{w}$-orbit

has length $\ell = n+1$, so the $\mathbf{w}_\pi$-orbit does too. For $i \in [n]$ Theorem 7.1.5 tells us that there are exactly two functions $f' \in \mathcal{O}'$ satisfying $f'(\pi^{-1}(i)) = 1$. Thus, there are exactly two functions $f \in \mathcal{O}$ satisfying $f(i) = 1$. ■

An interesting application of the whirling map is to better understand how to construct the factorization of $(n+1, n, \ldots, 2, 1)$ corresponding to a given $f \in \mathrm{Park}(n)$. The proof given in [44] and [6] proves existence and uniqueness but does not describe an easy way of constructing the factorization.

Example 7.1.11. There is a factorization $(4b_1)(3b_2)(1b_3)(4b_4)(1b_5)(6b_6) = (7654321)$ since $431416 \in \mathrm{Park}(6)$. We know that $b_6 = 7$ since $b_6 > 6$ and all $b_i \in [7]$, but *a priori* b_1 could be 5, 6, or 7, and b_3 could be 2, 3, 4, 5, 6, or 7. Using trial and error, there would be 1296 possibilities to check to find which one matches (7654321).

However, if we consider the $\mathbf{w}$-orbit containing 431416, then we can analyze the corresponding factorizations for all seven parking functions in the orbit. If we know the i^{th} cycle for one factorization in the orbit, then we can easily determine it for all of them. In particular, let g be such that the i^{th} cycle in the corresponding factorization contains 1. Then the i^{th} cycle in the factorization that corresponds to $\mathbf{w}^{-1}(g)$ contains $n+1$. Refer to Figure 7.1.2, in which we find a 1 in a cycle (shown in red), then place $n+1 = 7$ in the cycle one row above (also in red) and then determine that (45)(37)(12)(47)(13)(67) is the one corresponding to 431416.

Consider points $1, 2, 3, \ldots, n+1$ placed clockwise around a circle. There is a bijection, described in [19], between factorizations of $(n+1, n, \ldots, 2, 1)$ as a product of n transpositions and trees with labeled edges on these points that satisfy:

- the edges meet only at endpoints (i.e., do not cross each other),

431416	(45)(37)(12)(47)(13)(67)
512121	(56)(14)(23)(15)(24)(17)
623231	(67)(25)(34)(26)(35)(12)
134342	(17)(36)(45)(37)(46)(23)
145153	(12)(47)(56)(14)(57)(34)
216214	(23)(15)(67)(25)(16)(45)
321325	(34)(26)(17)(36)(27)(56)

FIGURE 7.1.2: An example showing how to construct the factorization of (7654321) that corresponds to the parking function 431416 ∈ Park(6) under Stanley's bijection. See Example 7.1.11.

- the edges are labeled with the numbers from 1 to n such that the labels for the edges meeting at any given vertex v increase when moving clockwise across the circle's interior around v.

In this bijection, if $(a_i b_i)$ is cycle i in the factorization, then the edge labeled i connects a_i to b_i. By this construction, adding 1 mod $(n+1)$ to each number in the factorization corresponds to rotating the edges of the tree $\frac{2\pi}{n+1}$ radians clockwise. See Figure 7.1.3 for an example **w**-orbit on Park(4) together with the corresponding factorizations of (54321) and trees.

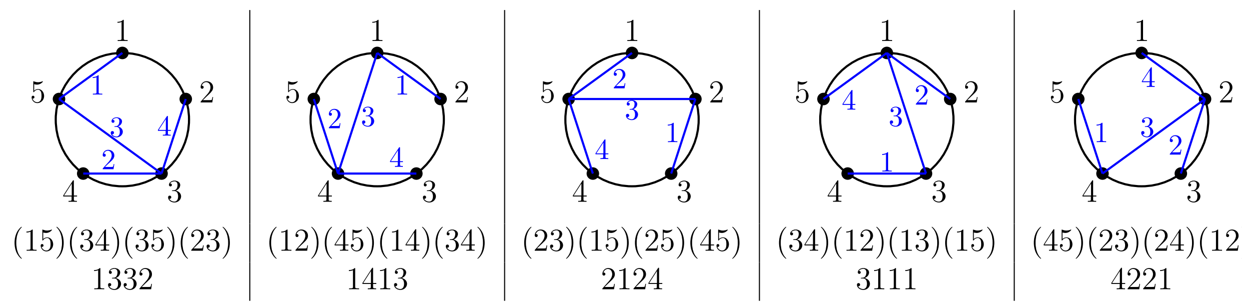

FIGURE 7.1.3: The orbit $(1332, 1413, 2124, 3111, 4221)$ of **w** on Park(5), with each parking function shown under the factorization of (54321) and tree it yields. Notice that **w** rotates the corresponding tree clockwise $2\pi/5$ radians.

7.2 Whirling restricted growth words

Another way of encoding a set partition is using its restricted growth word.

Definition 7.2.1. Let π be a partition of $[n]$ with k blocks. First order the blocks according to their least elements. The **restricted growth word** (or **RG-word**) according to π is the function $f : [n] \to [k]$ where $f(i) = j$ if i is in the j^{th} block of π.

Example 7.2.2. For the partition $125|46|37$ of $[7]$, we reorder the blocks as $125|37|46$. Then its RG-word is 1123132.

Let $\mathrm{RG}(n, k)$ denote the set of RG-words corresponding to partitions of $[n]$ with exactly k blocks and $\mathrm{RG}(n)$ the set of RG-words corresponding to all partitions of $[n]$ (without specifying the number of blocks). That is

$$\mathrm{RG}(n) = \bigcup_{k=1}^{n} \mathrm{RG}(n, k).$$

It is clear that any RG-word corresponds to a unique set partition, and that an RG-word of length n is in $\mathrm{RG}(n, k)$ where k is the maximum entry.

Let $\left\{ {n \atop k} \right\}$ denote the number of partitions of $[n]$ with exactly k blocks and B_n denote the total number of partitions of $[n]$. These are called the Stirling numbers of the second kind and Bell numbers respectively, e.g., see [46]. RG-words were first introduced by Hutchinson [23] and have been studied recently by Cai and Readdy [11] for their connections to the q and "q-$(1 + q)$" analogues for Stirling numbers of the second kind.

The following proposition is clear from the definition of RG-words, since the blocks are ordered via least elements.

Proposition 7.2.3.

- *A function* $f : [n] \to [k]$ *is in* $\mathrm{RG}(n,k)$ *if and only if it is surjective and for all* $1 \leq j \leq k-1$, $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$.

- *A function* $f : [n] \to [n]$ *is in* $\mathrm{RG}(n)$ *if and only if for any* $1 \leq j \leq n-1$,

  - *if* $j+1$ *is in the range of* f*, then so is* j*, and in this case,*

  - $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$.

The condition $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$ means that in the one-line notation of f, the first occurrence of j occurs before the first $j+1$. This is where the term "restricted growth" comes from. In particular, any RG-word f satisfies $f(1) = 1$.

Note that we consider the codomain of functions in $\mathrm{RG}(n,k)$ to be $[k]$ and of $\mathrm{RG}(n)$ to be $[n]$. Therefore for $\mathcal{F} = \mathrm{RG}(n,k)$, w_i adds 1 mod k repeatedly to the value $f(i)$ until we get a function in $\mathcal{F}$, but for $\mathcal{F} = \mathrm{RG}(n)$, w_i adds 1 mod n repeatedly instead.

Proposition 7.2.4. *Let* $\mathcal{F} = \mathrm{RG}(n,k)$*,* $f \in \mathcal{F}$*, and* $i \in [n]$*. Let* $f(i) = j$*. Then* w_i *changes the output at* i *in the following way.*

- *If* $j \neq k$*, then*

$$(w_i(f))(i) = \begin{cases} j+1 & \text{if there exists } i' < i \text{ such that } f(i') = j, \\ j & \text{if the only } i' < \min\{h|f(h) = j+1\} \text{ s.t. } f(i') = j \text{ is } i' = i, \\ 1 & \text{otherwise.} \end{cases}$$

- *If* $j = k$*, then*

$$(w_i(f))(i) = \begin{cases} k & \text{if the only value } i' \text{ for which } f(i') = k \text{ is } i' = i, \\ 1 & \text{otherwise.} \end{cases}$$

Refer to Example 7.2.6 for an example of Proposition 7.2.4. This says that in the one-line notation of f, if the value in position i is $j \neq k$, then w_i adds 1 to it if it is not the first occurrence of j. If it is the first occurrence of j, then w_i leaves it alone if there is not another j to the left of the first $j+1$, and otherwise changes the value to 1. In the case where the value in position i is k, w_i leaves it alone if it is the only k, and otherwise changes it to 1.

Proof. **Case 1: $f(i) \neq k$.** We have two subcases to consider.

Case 1a: There exists $i' < i$ such that $f(i') = j$. Then changing the value of $f(i)$ will not change the RG-word criterion $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$ from Proposition 7.2.3. So $(w_i(f))(i) = j + 1$.

Case 1b: There does not exist $i' < i$ such that $f(i') = j$. Then we cannot change the value of $f(i)$ to anything larger than j and still have an RG-word. If the only $i' < \min\{h|f(h) = j+1\}$ such that $f(i') = j$ is $i' = i$, then changing the value of $f(i)$ to something other than j will violate $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$ and not be an RG-word. Otherwise, changing the value of $f(i)$ to 1 results in an RG-word.

Case 2: $f(i) = k$. Since an RG-word must be surjective, if there does not exist $i' \neq i$ satisfying $f(i') = k$, then we cannot change the value of $f(i)$ to something other than k. So $(w_i(f))(i) = k$. Otherwise, changing the value of $f(i)$ to 1 still results in an RG-word, and so $(w_i(f))(i) = 1$. ■

For the $\mathcal{F} = \mathrm{RG}(n)$ case, we have the following slightly different characterization of w_i.

Proposition 7.2.5. *Let $\mathcal{F} = \mathrm{RG}(n)$, $f \in \mathcal{F}$, and $i \in [n]$. Let $f(i) = j$. Then*

$$(w_i(f))(i) = \begin{cases} j+1 & \text{if there exists } i' < i \text{ such that } f(i') = j, \\ j & \text{if the only } i' < \min\{h|f(h) = j+1\} \text{ s.t. } f(i') = j \text{ is } i' = i, \\ 1 & \text{otherwise.} \end{cases}$$

If there is no h satisfying $f(h) = j+1$, then we are not in the second case.

Proof. If there exists $i' < i$ such that $f(i') = j$, then changing the value of $f(i)$ will not change the RG-word criterion $\min\{i|f(i) = j\} \leq \min\{i|f(i) = j+1\}$ from Proposition 7.2.3. So $(w_i(f))(i) = j+1$.

Otherwise we cannot change the value of $f(i)$ to anything larger than j, so applying w_i will cycle the value of $f(i)$ around to 1 before possibly landing on a function in $\mathcal{F}$. If $j+1$ is in the range of f and $i' = i$ is the only $i' < \min\{h|f(h) = j+1\}$ such that $f(i') = j$, then we cannot change the value of $f(i)$ from j and still satisfy the necessary criterion. Otherwise, we can change the value of $f(i)$ to 1. ∎

Example 7.2.6. For $\mathcal{F} = \text{RG}(7,4)$,

$$1213341 \overset{w_1}{\mapsto} 1213341 \overset{w_2}{\mapsto} 1213341 \overset{w_3}{\mapsto} 1223341 \overset{w_4}{\mapsto} 1221341 \overset{w_5}{\mapsto} 1221341 \overset{w_6}{\mapsto} 1221341 \overset{w_7}{\mapsto} 1221342$$

so $\mathbf{w}(1213341) = 1221342$.

On the other hand, for $\mathcal{F} = \text{RG}(7)$,

$$1213341 \overset{w_1}{\mapsto} 1213341 \overset{w_2}{\mapsto} 1213341 \overset{w_3}{\mapsto} 1223341 \overset{w_4}{\mapsto} 1221341 \overset{w_5}{\mapsto} 1221341 \overset{w_6}{\mapsto} 1221311 \overset{w_7}{\mapsto} 1221312$$

so $\mathbf{w}(1213341) = 1221312$.

Propositions 7.2.4 and 7.2.5 show that for $\mathcal{F} = \text{RG}(n,k)$ or $\mathcal{F} = \text{RG}(n)$, whirling at index i either adds 1 to the value $f(i)$ or leaves $f(i)$ alone or changes the value $f(i)$ to 1.

For $\mathcal{F} = \text{RG}(n,k)$ or $\mathcal{F} = \text{RG}(n)$, w_1 acts trivially since every RG-word f satisfies $f(1) = 1$. Thus, $\mathbf{w} = w_n \circ \cdots w_3 \circ w_2$ on these families of functions. So when we consider generalized toggle orders we define $\mathbf{w}_\pi = w_{\pi(n)} \circ \cdots \circ w_{\pi(3)} \circ w_{\pi(2)}$ where π is a permutation of $\{2, 3, \ldots, n\}$.

Definition 7.2.7. Let
$$I_{i\mapsto 1}(f) = \begin{cases} 1 & \text{if } f(i) = 1, \\ 0 & \text{if } f(i) \neq 1. \end{cases}$$

The main homomesy result is the following.

Theorem 7.2.8. *Let $n \geq 2$. Fix $\mathcal{F}$ to be either* $\mathrm{RG}(n)$ *or* $\mathrm{RG}(n,k)$ *for some $k \in [n]$. Let π be a permutation of $\{2,3,\ldots,n\}$. Under the action of $\mathbf{w}_\pi$ on $\mathcal{F}$, $I_{i\mapsto 1} - I_{j\mapsto 1}$ is 0-mesic for any $i,j \in \{2,3,\ldots,n\}$.*

See Figure 7.2.1 for an illustration of Theorem 7.2.8 for the orbits under the action of $\mathbf{w}$ (i.e., π is the identity) over $\mathcal{F} = \mathrm{RG}(5,3)$. This theorem is another instance where there is homomesy under an action that produces unpredictable orbit sizes, and for which the order of the map is unknown in general.

For RG-words, results for $\mathbf{w}$ need not necessarily extend to other $\mathbf{w}_\pi$ products like they did for $\mathrm{Inj}_m(n,k)$, $\mathrm{Sur}_m(n,k)$, and $\mathrm{Park}(n)$, since a rearrangement of an RG-word is not always an RG-word. In fact, these other whirling orders do not always yield the same orbit structure as $\mathbf{w}$. Therefore, we prove this result keeping an arbitrary π in mind.

We will use several lemmas in the proof of Theorem 7.2.8. For the rest of this section, assume $n \geq 2$.

Lemma 7.2.9. *Let $\mathcal{F}$ be either* $\mathrm{RG}(n)$ *or* $\mathrm{RG}(n,k)$ *for some $k \in [n]$. Let π, σ be permutations on $\{2,3,\ldots,n\}$ where $\sigma(i) = \pi(i+1)$ for $2 \leq i \leq n-1$ and $\sigma(n) = \pi(2)$. Then for any $\mathbf{w}_\pi$-orbit $\mathcal{O} = (f_1, f_2, \ldots, f_\ell)$, there is a $\mathbf{w}_\sigma$-orbit $\mathcal{O}' = (f_1', f_2', \ldots, f_\ell')$ for which $f_j' = w_{\pi(2)}(f_j)$ for all $j \in [\ell]$.*

Also, given $h \in [n]$, the multiset of values $f(h)$ as f ranges over $\mathcal{O}$ is the same as that for f ranging over $\mathcal{O}'$.

f	12123
$\mathbf{w}(f)$	11231
$\mathbf{w}^2(f)$	12312
$\mathbf{w}^3(f)$	12323
$\mathbf{w}^4(f)$	12131
$\mathbf{w}^5(f)$	12232
$\mathbf{w}^6(f)$	11233
$\mathbf{w}^7(f)$	12311
$\mathbf{w}^8(f)$	12322
$\mathbf{w}^9(f)$	12333
$\mathbf{w}^{10}(f)$	12113
$\mathbf{w}^{11}(f)$	12223
$\mathbf{w}^{12}(f)$	11123
$\mathbf{w}^{13}(f)$	12231
$\mathbf{w}^{14}(f)$	11232
$\mathbf{w}^{15}(f)$	12313

g	12213
$\mathbf{w}(g)$	11223
$\mathbf{w}^2(g)$	12331
$\mathbf{w}^3(g)$	12132
$\mathbf{w}^4(g)$	12233
$\mathbf{w}^5(g)$	11213
$\mathbf{w}^6(g)$	12321
$\mathbf{w}^7(g)$	12332
$\mathbf{w}^8(g)$	12133

FIGURE 7.2.1: The two $\mathbf{w}$-orbits for $\mathcal{F} = \mathrm{RG}(5,3)$. The left orbit has length 16 and the right orbit has length 9. Notice that in the left orbit, there are the same numbers (four each) of functions h satisfying each of $h(2) = 1$, $h(3) = 1$, $h(4) = 1$, and $h(5) = 1$. In the right orbit, there are also the same numbers (two each) of functions h satisfying each of $h(2) = 1$, $h(3) = 1$, $h(4) = 1$, and $h(5) = 1$. This demonstrates Theorem 7.2.8 for the case $n = 5$, $k = 3$, and π is the identity.

Proof. Let $\mathcal{O} = (f_1, f_2, \ldots, f_\ell)$ be a $\mathbf{w}_\pi$-orbit satisfying $\mathbf{w}_\pi(f_j) = f_{j+1}$ for all j. Consider the subscripts to be mod ℓ, the length of the orbit, so $f_{\ell+1} = f_1$ for instance. Note that $\mathbf{w}_\pi = w_{\pi(n)} \circ \cdots \circ w_{\pi(3)} \circ w_{\pi(2)}$ and $\mathbf{w}_\sigma = w_{\pi(2)} \circ w_{\pi(n)} \circ \cdots \circ w_{\pi(3)}$. Therefore, if we let $f'_j = w_{\pi(2)}(f_j)$ for all j, then $\mathbf{w}_\sigma(f'_j) = f'_{j+1}$ for all j. So $\mathcal{O}' = (f'_1, f'_2, \ldots, f'_\ell)$ is a $\mathbf{w}_\sigma$-orbit.

For $h \neq \pi(2)$, $f_i(h) = f'_i(h)$. Also, $f'_i(\pi(2)) = (w_{\pi(2)}(f))(\pi(2)) = \mathbf{w}_\pi(2) = f_{i+1}(2)$. Thus for any $h \in [n]$, the multiset of values as $f(h)$ is the same for f ranging over $\mathcal{O}$ is the same as that for $\mathcal{O}'$. ∎

Remark 7.2.10. Suppose we have $\mathbf{w}_\pi = w_{\pi(n)} \circ \cdots \circ w_{\pi(3)} \circ w_{\pi(2)}$ and $\mathbf{w}_\sigma$ is some cyclic shift

of this order of composition. Then by applying Lemma 7.2.9 repeatedly, we get that every $\mathbf{w}_\pi$-orbit $\mathcal{O}$ corresponds uniquely with a $\mathbf{w}_\sigma$-orbit $\mathcal{O}'$ where the multisets of $f(h)$ values are the same for f ranging over $\mathcal{O}$ as for $\mathcal{O}'$.

For example, on $\mathcal{F} = \mathrm{RG}(6)$, $w_3 \circ w_5 \circ w_4 \circ w_6 \circ w_2$ and $w_4 \circ w_6 \circ w_2 \circ w_3 \circ w_5$ satisfy the above orbit correspondence. However in general two orders for applying the whirling maps do not yield the same orbit structure. For example, $w_3 \circ w_5 \circ w_4 \circ w_6 \circ w_2$ does not have the same orbit structure as $w_5 \circ w_6 \circ w_3 \circ w_2 \circ w_4$ because these are not cyclic rotations of each other.

Lemma 7.2.11. *Let $\mathcal{F}$ be either $\mathrm{RG}(n)$ or $\mathrm{RG}(n,k)$ for some $k \in [n]$ and π be a permutation of $\{2,3,\dots,n\}$. Let $i = \pi(2)$ and suppose $2 \le i \le n-1$. Then for $f \in \mathcal{F}$, if $f(i) = 1$, then $\left(\mathbf{w}_\pi^{-1}(f)\right)(i) \ge \left(\mathbf{w}_\pi^{-1}(f)\right)(i+1)$.*

Proof. Let $f \in \mathcal{F}$ such that $f(i) = 1$. Let $g = \mathbf{w}_\pi^{-1}(f)$. Since w_i is the first map we apply when we apply $\mathbf{w}$ to g, w_i changes the value of $g(i)$ to 1. By Propositions 7.2.4 and 7.2.5, either i is the least i' satisfying $g(i') = g(i)$, or $g(i) = k$ (the latter option only being possible or making sense if $\mathcal{F} = \mathrm{RG}(n,k)$).

Assume by way of contradiction that $g(i) < g(i+1)$. This could only be possible in the case where i is the least i' satisfying $g(i') = g(i)$. From the restricted growth condition, $g(i+1) = g(i)+1$. Then Propositions 7.2.4 and 7.2.5 say that w_i would leave the value $g(i)$ as is, not change it to 1, which is a contradiction. Thus, $g(i) \ge g(i+1)$. ■

Lemma 7.2.12. *Let $i \in \{2,3,\dots,n\}$.*

1. *The number of $f \in \mathrm{RG}(n,k)$ satisfying $f(i) = 1$ is $\left\{ {n-1 \atop k} \right\}$.*

2. *The number of $f \in \mathrm{RG}(n)$ satisfying $f(i) = 1$ is B_{n-1}.*

Proof. The condition $f(i) = 1$ means 1 and i are in the same block of the partition of $[n]$ corresponding to f. Such a set partition can be formed by first choosing a partition of $\{2, 3, \ldots, n-1\}$ and then placing 1 in the same block as i. ■

We are now ready for the proof of Theorem 7.2.8.

Proof of Theorem 7.2.8. Let $\mathcal{O}$ be a $\mathbf{w}_\pi$-orbit. The goal is to prove that in $\mathcal{O}$, there are the same number of functions $f \in \mathcal{O}$ satisfying $f(i) = 1$ as there are satisfying $f(j) = 1$, when $2 \leq i, j \leq n$. We will approach this by showing that $\mathcal{O}$ has the same number of functions f satisfying $f(i) = 1$ as $f(i+1) = 1$, for $2 \leq i \leq n-1$.

Let $i \in \{2, 3, \ldots, n-1\}$. Without loss of generality, assume $i = \pi(2)$. This is because by Remark 7.2.10, we could cyclically change π to get this condition, without changing the property we wish to prove for orbits of $\mathbf{w}_\pi$.

Let $f \in \mathcal{O}$ be a function satisfying $f(i) = 1$ and $f(i+1) \neq 1$. By Lemma 7.2.11, $\left(\mathbf{w}_\pi^{-1}(f)\right)(i) \geq \left(\mathbf{w}_\pi^{-1}(f)\right)(i{+}1)$. Now let r be the least positive integer such that $\left(\mathbf{w}_\pi^{r}(f)\right)(i) \geq \left(\mathbf{w}_\pi^{r}(f)\right)(i+1)$. (Such an r must exist because when repeatedly applying $\mathbf{w}_\pi$ to f, we must eventually cycle back around to $\mathbf{w}_\pi^{-1}(f)$.)

Let $g = \mathbf{w}_\pi^{r-1}$. Then

$$g(i) < g(i+1) \text{ and } (\mathbf{w}_\pi(g))(i) \geq (\mathbf{w}_\pi(g))(i+1). \tag{7.2.1}$$

Due to the way whirling at an index works, either $(\mathbf{w}_\pi(g))(i) = g(i){+}1$ or $(\mathbf{w}_\pi(g))(i{+}1) = 1$.

Suppose $(\mathbf{w}_\pi(g))(i{+}1) \neq 1$. Then $(\mathbf{w}_\pi(g))(i{+}1) = g(i{+}1)$ or $(\mathbf{w}_\pi(g))(i{+}1) = g(i{+}1){+}1$, the latter of which violates 7.2.1 because $(\mathbf{w}_\pi(g))(i) = g(i)+1$. So $(\mathbf{w}_\pi(g))(i+1) = g(i+1)$. Then from 7.2.1, we have $(\mathbf{w}_\pi(g))(i) = g(i) + 1 = g(i+1)$. However, since we apply w_{i+1} after w_i, w_{i+1} would have to add 1 to the value $g(i{+}1)$ by Propositions 7.2.4 and 7.2.5. This

contradicts $(\mathbf{w}_\pi(g))(i+1) = g(i+1)$. Thus,

$$(\mathbf{w}_\pi^r(f))(i+1) = (\mathbf{w}_\pi(g))(i+1) = 1.$$

By Lemma 7.2.11, $(\mathbf{w}_\pi^m(f))\,(i) \neq 1$ for all $m \in [r]$. Therefore, given any function $f \in \mathcal{O}$ satisfying $f(i) = 1$ and $f(i+1) \neq 1$, as we repeatedly apply $\mathbf{w}_\pi$ to f, we obtain a function that sends $i+1$ to 1 before we get another function that sends i to 1. This implies $\#\{f \in \mathcal{O} | f(i) = 1\} \leq \#\{f \in \mathcal{O} | f(i+1) = 1\}$.

By Lemma 7.2.12, $\#\{f \in \mathcal{F} | f(i) = 1\} = \#\{f \in \mathcal{F} | f(i+1) = 1\}$. Therefore, since $\#\{f \in \mathcal{O} | f(i) = 1\} \leq \#\{f \in \mathcal{O} | f(i+1) = 1\}$ for *every* $\mathbf{w}_\pi$-orbit $\mathcal{O}$ over $\mathcal{F}$, we must have $\#\{f \in \mathcal{O} | f(i) = 1\} = \#\{f \in \mathcal{O} | f(i+1) = 1\}$ for every orbit $\mathcal{O}$. ■

Theorem 7.2.8 implies another homomesic statistic only for the action of $\mathbf{w}_\pi$ only on $\mathrm{RG}(n,k)$.

Corollary 7.2.13. *Let $n \geq 2$ and π a permutation of $\{2, 3, \ldots, n\}$. Under the action of $\mathbf{w}_\pi$ on $\mathrm{RG}(n,k)$, the statistic*

$$f \mapsto \binom{k}{2} f(2) - f(n)$$

is homomesic with average $k(k-2)$.

Proof. For $f \in \mathrm{RG}(n,k)$, $f(2)$ is always 1 or 2. Consider a $\mathbf{w}_\pi$-orbit $\mathcal{O}$ with length $\ell(\mathcal{O})$. If $f(2) = 2$ and $f(n) = k$ for every $f \in \mathcal{O}$, then the total value of $f \mapsto \binom{k}{2} f(2) - f(n)$ across $\mathcal{O}$ is $\frac{k(k-1)}{2} 2\ell(\mathcal{O}) - k\ell(\mathcal{O}) = k(k-2)\ell(\mathcal{O})$.

Suppose instead that there are c functions $f \in \mathcal{O}$ that satisfy $f(2) = 1$. Then this decreases the total value of $\binom{k}{2} f(2)$ across $\mathcal{O}$ by $\binom{k}{2} c$.

From Theorem 7.2.8, there are also c functions $f \in \mathcal{O}$ satisfying $f(n) = 1$. If $f(n) \neq k$, then n cannot be the least i for which $f(i) = f(n)$ according to Proposition 7.2.3. Thus,

when $f(n) \neq k$, $\mathbf{w}_\pi$ adds 1 to the value of $f(n)$. When $f(n) = k$, either $(\mathbf{w}_\pi(f))(n) = 1$ or $(\mathbf{w}_\pi(f))(n) = k$ via Proposition 7.2.4. So there are also c functions satisfying $f(n) = j$ for any $j \in [k-1]$. In relation to the original case where $f(n) = k$ for all f, this decreases the total value of $f(n)$ across $\mathcal{O}$ by $c(1+2+3+\cdots+(k-1)) = \binom{k}{2}c$.

Therefore, the total value of $f \mapsto \binom{k}{2}f(2) - f(n)$ across $\mathcal{O}$ is

$$k(k-2)\ell(\mathcal{O}) - \binom{k}{2}c + \binom{k}{2}c = k(k-2)\ell(\mathcal{O}).$$

Hence this statistic has average $k(k-2)$ on $\mathcal{O}$. ■

Definition 7.2.14. A **noncrossing RG-word** is an RG-word whose corresponding partition is noncrossing. Let $\mathrm{RG}_{\mathrm{nc}}(n,k)$ and $\mathrm{RG}_{\mathrm{nc}}(n)$ denote the sets of all noncrossing RG-words in $\mathrm{RG}(n,k)$ and $\mathrm{RG}(n)$ respectively.

Proposition 7.2.15. *An RG-word f is noncrossing if and only if in the one-line notation, we do not have abab in order (not necessarily consecutively).*

Example 7.2.16. Let $f = 122133143$. Then $f \in \mathrm{RG}(9,4)$ but $f \notin \mathrm{RG}_{\mathrm{nc}}(9,4)$ because the red values are 1313 in relative order.

Due to the noncrossing condition, there is not such a simple description (like Proposition 7.2.4 or 7.2.5) of how w_i acts for $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n,k)$ or $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n)$. In the following example, we see w_i can decrease $f(i)$ to something other than 1, or increase $f(i)$ by more than one.

Example 7.2.17. For $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(6,4)$,

$$123442 \xmapsto{w_1} 123442 \xmapsto{w_2} 123442 \xmapsto{w_3} 123442 \xmapsto{w_4} 123242 \xmapsto{w_5} 123242 \xmapsto{w_6} 123244$$

so $\mathbf{w}(123442) = 123244$.

For $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(6)$,

$$123442 \xmapsto{w_1} 123442 \xmapsto{w_2} 123442 \xmapsto{w_3} 123442 \xmapsto{w_4} 123242 \xmapsto{w_5} 123222 \xmapsto{w_6} 123224$$

so $\mathbf{w}(123442) = 123224$.

We conjecture homomesy for whirling noncrossing RG-words. Again w_1 acts trivially for $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n)$ or $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n,k)$, so for simplicity we ignore w_1.

Conjecture 7.2.18. *Let $n \geq 2$. Fix $\mathcal{F}$ to be either $\mathrm{RG}_{\mathrm{nc}}(n)$ or $\mathrm{RG}_{\mathrm{nc}}(n,k)$ for some $k \in [n]$. Let π be a permutation of $\{2, 3, \ldots, n\}$. Under the action of $\mathbf{w}_\pi$ on $\mathcal{F}$, $I_{2\mapsto 1} - I_{n\mapsto 1}$ is 0-mesic.*

This conjecture says that for noncrossing RG-words, we have the specific case $i = 2$, $j = n$ of the homomesy in Theorem 7.2.8. The proof of Theorem 7.2.8 relied on the fact from Lemma 7.2.12 that for $\mathcal{F} = \mathrm{RG}(n,k)$ or $\mathcal{F} = \mathrm{RG}(n)$, $\#\{f \in \mathcal{F} | f(i) = 1\} = \#\{f \in \mathcal{F} | f(j) = 1\}$ when $2 \leq i, j \leq n$. For $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n,k)$ or $\mathcal{F} = \mathrm{RG}_{\mathrm{nc}}(n)$, we only have this when $i = 2$, $j = n$. For example, over the set of noncrossing partitions on [4], there are five partitions containing 1 and 2 in the same block, four partitions containing 1 and 3 in the same block, and five partitions containing 1 and 4 in the same block.

Under the whirling left-to-right order (π is the identity) Conjecture 7.2.18 has been confirmed for all $n \leq 9$ and relevant k values ($k \in [n]$). It has also been tested and confirmed for many random whirling orders (with $n \leq 9$).

Chapter 8

Future directions

There are many systems consisting of a set of combinatorial interest together with an action for which one can search for homomesic statistics. For most of the ones studied so far, we have discovered homomesy and we gain more insight into this phenomenon by continuing to explore more sets and actions.

This thesis considers dynamical actions on some of the most basic objections in combinatorics: independent sets of a path graph, various classes of set partitions, and various classes of functions between finite sets. Many of these can be generalized in several directions. For example, as mentioned in Chapter 3, homomesies for toggling $J(P)$ have been found to generalize to a piecewise-linear analogue for some posets (like $P = [a] \times [b]$) and not for others (like our zigzag poset results in Section 4.4). More insight on our results and their significance can be gained by studying various generalizations.

Parking functions can be generalized in several ways. Given any finite Weyl group W, Armstrong, Reiner, and Rhodes describe a set of "parking functions" corresponding to W [2]. In this analogue, parking functions in the usual sense correspond to parking functions on

symmetric groups. Postnikov and Shapiro describe the set of G-parking functions for a graph G; in this analogue parking functions in the usual sense are those on complete graphs [31]. Bruner and Panholzer describe generalizations to rooted labeled trees and to directed graphs of mappings [10]. For the rooted labeled trees case, parking functions in the usual sense correspond to those on the path graph $\mathcal{P}_n$ with n as the root. These give us several ways to try to extend the results on whirling parking functions.

In the case $n \leq \frac{k}{2}$, the whirling action on $\mathrm{Inj}_m(n,k)$ can be restated in terms of periodic box-ball systems, as described in [17] and [56]. For this, we obtain a sequence of actions $T_1, T_2, T_3, \ldots$ that eventually stabilize to a map that is equivalent (under an equivariant bijection) to $\mathbf{w}$. This gives a natural place to study to see if the homomesy can be extended to these other actions.

To continue the work on noncrossing partitions, there are several directions to explore generalizations for future work. Noncrossing partitions can be defined for any finite Coxeter group W. They are defined as an interval in a poset called the "absolute poset" corresponding to W. Details can be found in [1]. Constructing a generalization of a map (or set of objects) to other Coxeter groups is not easy. Right now, we do not know a way to extend our linear representations to general W, since they are not stated only in terms of the generators s_i, but if we can find a generalization, then it is a good place to look to see if our theory and homomesy extends. Generalizations of combinatorial sets to other Coxeter types is presently an active research area, as it is mostly mysterious why many results that hold for type A (the symmetric group $\mathfrak{S}_n$) also hold for all finite Coxeter groups.

Another natural question to ask is what happens for piecewise-linear and birational analogues of various toggle actions studied in this thesis. These analogues are defined for toggling over order ideals of posets; see in [14] and [20]. Similarly, there is at least a piecewise-linear analogue for toggling independent sets of graphs, where each vertex is assigned a value be-

tween 0 and 1 and "independent" means for any pair of adjacent vertices, the vertex labels sum to no more than 1. While our homomesy results for independent sets of path graphs and for noncrossing partitions (as independent sets of the base graph) here have been shown *not* to hold for the piecewise-linear setting, maybe some other properties discussed here to extend.

As there are over 200 known sets of objects counted by the sequence of Catalan numbers, we also have many known bijections between them. It is likely that studying an action that is naturally defined on one set of objects and translating it into the language of another leads to interesting results. In particular, we may be able to restate our results for noncrossing partitions (or the conjectured ones for noncrossing RG-words) in a natural way for a map on another set of Catalan objects.

Bibliography

[1] D. Armstrong, *Generalized noncrossing partitions and combinatorics of Coxeter groups*, vol. 202, American Mathematical Society, 2009.

[2] D. Armstrong, V. Reiner, and B. Rhoades, *Parking spaces*, Advances in Mathematics **269** (2015), 647–706.

[3] D. Armstrong, C. Stump, and H. Thomas, *A uniform bijection between nonnesting and noncrossing partitions*, Transactions of the American Mathematical Society **365** (2013), no. 8, 4121–4151, also available as `arXiv:1101.1277v2`.

[4] R. Behrend and T. Roby, *Dihedral homomesy for alternating sign matrices*, in preparation.

[5] D. Bessis and V. Reiner, *Cyclic sieving of noncrossing partitions for complex reflection groups*, Annals of Combinatorics **15** (2011), no. 2, 197–222.

[6] P. Biane, *Parking functions of types A and B*, Electron. J. Combin **9** (2002), no. 1, N7.

[7] A. Björner and F. Brenti, *Combinatorics of Coxeter groups*, Springer Verlag, New York, 2005.

[8] D. Bressoud, *Proofs and confirmations: The story of the alternating-sign matrix conjecture*, Cambridge University Press, 1999.

[9] A. Brouwer and L. Schrijver, *On the period of an operator, defined on antichains*, Stichting Mathematisch Centrum. Zuivere Wiskunde (1974), no. ZW 24/74, 1–13.

[10] M. Bruner and A. Panholzer, *Parking functions for trees and mappings*, arXiv preprint arXiv:1504.04972 (2015).

[11] Y. Cai and M. Readdy, *q-Stirling numbers: A new view*, Advances in Applied Mathematics **86** (2017), 50–80.

[12] P. Cameron and D. Fon-der-Flaass, *Orbits of antichains revisited*, European J. Combin. **16** (1995), no. 6, 545–554.

[13] 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** (2016), no. 3, also available at `arXiv:1510.06362v2`.

[14] D. Einstein and J. Propp, *Combinatorial, piecewise-linear, and birational homomesy for products of two chains*, `arXiv:1310.5294` (2013).

[15] H. Eriksson and K. Eriksson, *Conjugacy of Coxeter elements*, Electron. J. Combin. **16** (2009), no. 2, #R4.

[16] D. Fon-der-Flaass, *Orbits of antichains in ranked posets*, European journal of combinatorics **14** (1993), no. 1, 17–22.

[17] Kaori Fukuda, *Box-ball systems and Robinson-Schensted-Knuth*, Journal of Algebraic Combinatorics **19** (2004), 67–89.

[18] J. Fürlinger and J. Hofbauer, *q-Catalan numbers*, Journal of Combinatorial Theory, Series A **40** (1985), no. 2, 248–264.

[19] I. Goulden and A. Yong, *Tree-like properties of cycle factorizations*, Journal of Combinatorial Theory, Series A **98** (2002), no. 1, 106–117.

[20] D. Grinberg and T. Roby, *Iterative properties of birational rowmotion*, `arXiv:1402.6178` (2014).

[21] S. Haddadan, *Some instances of homomesy among ideals of posets*, `arXiv:1410.4819v3` (2016).

[22] C. Heitsch, *Counting orbits under Kreweras complementation*, In preparation.

[23] G. Hutchinson, *Partioning algorithms for finite sets*, Communications of the ACM **6** (1963), no. 10, 613–614.

[24] S. Ilie, D. Jeffrey, R. Corless, and X. Zhang, *Computation of Stirling numbers and generalizations*, Symbolic and Numeric Algorithms for Scientific Computing, IEEE, 2015, pp. 57–60.

[25] M. Joseph and T. Roby, *Toggling independent sets of a path graph*, `arXiv:1701.04956` (2017).

[26] A. Konheim and B. Weiss, *An occupancy discipline and applications*, SIAM Journal on Applied Mathematics **14** (1966), no. 6, 1266–1274.

[27] G. Kreweras, *Sur les partitions non croisées d'un cycle*, Discrete Math. **1** (1972), no. 4, 333–350. MR 0309747 (46 #8852)

[28] M. LaCroix and T. Roby, *Foatic actions of the symmetric group and fixed-point homomesy*, In preparation.

[29] D. Panyushev, *On orbits of antichains of positive roots*, European J. Combin. **30** (2009), no. 2, 586–594, also available at `arXiv:0711.3353v2`.

[30] T. K. Petersen, *Eulerian numbers*, Springer, New York, 2015.

[31] A. Postnikov and B. Shapiro, *Trees, parking functions, syzygies, and deformations of monomial ideals*, Transactions of the American Mathematical Society **356** (2004), no. 8, 3109–3142.

[32] J. Propp and T. Roby, *Homomesy in products of two chains*, Electron. J. Combin. **22** (2015), no. 3, also available at `arXiv:1310.5201v5`.

[33] V. Reiner, D. Stanton, and D. White, *The cyclic sieving phenomenon*, Journal of Combinatorial Theory, Series A **108** (2004), no. 1, 17–50.

[34] V. Reiner, D. Stanton, and D. White, *What is... cyclic sieving*, Notices Amer. Math. Soc **61** (2014), no. 2, 169–171.

[35] A. Rényi, *Théorie des éléments saillants d'une suite d'observations*, Annales scientifiques de l'Université de Clermont. Mathématiques **8** (1962), no. 2, 7–13.

[36] J. Riordan, *Ballots and trees*, Journal of Combinatorial Theory **6** (1969), no. 4, 408–411.

[37] T. Roby, *Dynamical algebraic combinatorics and the homomesy phenomenon*, Recent Trends in Combinatorics, Springer, 2016, pp. 619-652, also available at `http://www.math.uconn.edu/~troby/homomesyIMA2015Revised.pdf`.

[38] B. Sagan, *The cyclic sieving phenomenon: A survey*, London Math. Soc. Lecture Note Ser **392** (2011), 183–233, also available at `arXiv:1008.0790v3`.

[39] M. Schützenberger, *Quelques remarques sur une construction de Schensted*, Math. Scand. **12** (1963), 117–128. MR 0190017.

[40] M. Schützenberger, *Promotion des morphismes d'ensembles ordonnés*, Discrete Math. **2** (1972), 73–94. MR 0299539

[41] J. Shi, *The enumeration of Coxeter elements*, J. Algebraic Combin. **6** (1997), no. 2, 161–171. MR 1436533 (98d:20048).

[42] R. Simion and D. Ullman, *On the structure of the lattice of noncrossing partitions*, Discrete Math. **98** (1991), 193–206.

[43] N. Sloane, *OEIS, online encyclopedia of integer sequences*, 2016, `https://oeis.org/`.

[44] R. Stanley, *Parking functions and noncrossing partitions*, Electron. J. Combin **4** (1997), no. 2, R20.

[45] R. Stanley, *Enumerative combinatorics, volume 2*, Cambridge University Press, 1999.

[46] R. Stanley, *Enumerative combinatorics, volume 1, 2nd edition*, Cambridge University Press, 2011, also available at `http://math.mit.edu/~rstan/ec/ec1/`.

[47] R. Stanley, *Catalan numbers*, Cambridge University Press, 2015.

[48] J. Stembridge, *Some hidden relations involving the ten symmetry classes of plane partitions*, Journal of Combinatorial Theory, Series A **68** (1994), no. 2, 372–409.

[49] J. Striker, *The toggle group, homomesy, and the Razumov-Stroganov correspondence*, Electron. J. Combin. **22** (2015), no. 2, P2–57, also available at `arXiv:1503.08898v1`.

[50] J. Striker, *Rowmotion and generalized toggle groups*, arXiv:1601.03710 (2016).

[51] J. Striker and N. Williams, *Promotion and rowmotion*, European J. Combin. **33** (2012), 1919–1942, also available at `arXiv:1108.1172v3`.

[52] R. Suter, *Young's lattice and dihedral symmetries*, European J. Combin. **23** (2002), no. 2, 233–238.

[53] M. Thiel, *A new cyclic sieving phenomenon for catalan objects*, Discrete Mathematics **340** (2017), no. 3, 426–429.

[54] D. Walsh, *Width-restricted finite functions*, `http://capone.mtsu.edu/dwalsh/width_2.pdf`.

[55] H. Wilf, *Generatingfunctionology, 3rd edition*, CRC Press.

[56] D. Yoshihara, F. Yura, and T. Tokihiro, *Fundamental cycle of a periodic box–ball system*, Journal of Physics A: Mathematical and General **36** (2002), no. 1, 99.