

eScholarship

Combinatorial Theory

Title

Permutoric promotion: gliding globs, sliding stones, and colliding coins

Permalink

<https://escholarship.org/uc/item/71r9p20q>

Journal

Combinatorial Theory, 4(2)

ISSN

2766-1334

Authors

Defant, Colin

Madhukara, Rachana

Thomas, Hugh

Publication Date

2024

DOI

10.5070/C64264253

Supplemental Material

<https://escholarship.org/uc/item/71r9p20q#supplemental>

Copyright Information

Copyright 2024 by the author(s). This work is made available under the terms of a Creative Commons Attribution License, available at

<https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

PERMUTORIC PROMOTION: GLIDING GLOBS, SLIDING STONES, AND COLLIDING COINS

Colin Defant^{*1}, Rachana Madhukara², and Hugh Thomas^{†3}

¹*Department of Mathematics, Harvard University, Cambridge, MA 02139, U.S.A.*
colindefant@gmail.com

²*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.*
rachanam@mit.edu

³*Lacim, UQAM, Montréal, QC, H3C 3P8, Canada*
thomas.hugh_r@uqam.ca

Submitted: Jul 13, 2023; Accepted: Jul 3, 2024; Published: Sep 30, 2024

© The authors. Released under the CC BY license (International 4.0).

Abstract. Defant recently introduced *toric promotion*, an operator that acts on the labelings of a graph G and serves as a cyclic analogue of Schützenberger’s promotion operator. Toric promotion is defined as the composition of certain *toggle* operators, listed in a natural cyclic order. We consider more general *permutoric promotion* operators, which are defined as compositions of the same toggles, but in permuted orders. We settle a conjecture of Defant by determining the orders of all permutoric promotion operators when G is a path graph. In fact, we completely characterize the orbit structures of these operators, showing that they satisfy the cyclic sieving phenomenon. The first half of our proof requires us to introduce and analyze new *broken promotion* operators, which can be interpreted via globs of liquid gliding on a path graph. For the latter half of our proof, we reformulate the dynamics of permutoric promotion via stones sliding along a cycle graph and coins colliding with each other on a path graph.

Keywords. Promotion, toric promotion, Coxeter element, cyclic sieving phenomenon

Mathematics Subject Classifications. 05E18

1. Introduction

In his study of the Robinson–Schensted–Knuth correspondence, Schützenberger [Sch63, Sch72, Sch76] introduced a beautiful bijective operator called *promotion*, which acts on the set of linear extensions of a finite poset. Haiman [Hai92] and Malvenuto–Reutenauer [MR94] found that

*Supported by NSF grant No. 2201907.

†Supported by NSERC Discovery Grant RGPIN-2022-03960 and the Canada Research Chairs program, grant number CRC-2021-00120.

promotion could be defined as a composition of local *toggle operators* (also called *Bender–Knuth involutions*). There are now several articles connecting promotion to other areas [AKS14, EG87, HR22, Hua20, PPR09, PS18, Rho10, Sta09, SW12] and generalizing promotion in different directions [AKS14, BSV21, DK23, DPS17, DSV19, Sta09]. Promotion is now one of the most extensively studied operators in the field of dynamical algebraic combinatorics.

Following the approach first considered by Malvenuto and Reutenauer [MR94], we define promotion on labelings of graphs instead of linear extensions of posets. All graphs in this article are assumed to be simple. Let $G = (V, E)$ be a graph with n vertices. A *labeling* of G is a bijection $V \rightarrow \mathbb{Z}/n\mathbb{Z}$. We denote the set of labelings of G by Λ_G . Given distinct $a, b \in \mathbb{Z}/n\mathbb{Z}$, let $(a\ b)$ be the transposition that swaps a and b . For $i \in \mathbb{Z}/n\mathbb{Z}$, the *toggle operator* $\tau_i: \Lambda_G \rightarrow \Lambda_G$ is defined by

$$\tau_i(\sigma) = \begin{cases} (i\ i+1) \circ \sigma & \text{if } \{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \notin E; \\ \sigma & \text{if } \{\sigma^{-1}(i), \sigma^{-1}(i+1)\} \in E. \end{cases}$$

In other words, τ_i swaps the labels i and $i+1$ if those labels are assigned to nonadjacent vertices of G , and it does nothing otherwise. Define *promotion* to be the operator $\text{Pro}: \Lambda_G \rightarrow \Lambda_G$ given by

$$\text{Pro} = \tau_{n-1} \cdots \tau_2 \tau_1.$$

Here and in the sequel, concatenation of operators represents composition.

A recent trend in algebraic combinatorics aims to find cyclic analogues of more traditional “linear” objects (see [ARR20, DMR16] and the references therein). Defant recently defined a cyclic analogue of promotion called *toric promotion* [Def23]; this is the operator $\text{TPro}: \Lambda_G \rightarrow \Lambda_G$ given by

$$\text{TPro} = \tau_n \tau_{n-1} \cdots \tau_2 \tau_1 = \tau_n \text{Pro}.$$

Defant proved the following theorem, which reveals that toric promotion has remarkably nice dynamical properties when G is a forest.

Theorem 1.1 ([Def23]). *Let G be a forest with $n \geq 2$ vertices, and let $\sigma \in \Lambda_G$ be a labeling. The orbit of toric promotion containing σ has size*

$$(n-1) \frac{t}{\gcd(t, n)},$$

where t is the number of vertices in the connected component of G containing $\sigma^{-1}(1)$. In particular, if G is a tree, then every orbit of $\text{TPro}: \Lambda_G \rightarrow \Lambda_G$ has size $n-1$.

Theorem 1.1 stands in stark contrast to the wild dynamics of promotion on most forests. For example, even when G is a path graph with 7 vertices, the order of $\text{Pro}: \Lambda_G \rightarrow \Lambda_G$ is 3224590642072800, whereas all orbits of $\text{TPro}: \Lambda_G \rightarrow \Lambda_G$ have size 6.

In [Def23], Defant (taking a suggestion from Tom Roby) proposed studying a generalization of toric promotion in which the toggle operators τ_1, \dots, τ_n can be composed in any order. In what follows, we let $[n] = \{1, \dots, n\}$.

Definition 1.2. Let G be a graph with n vertices, and let $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ be a bijection. The *permutoric promotion* operator $\text{TPro}_\pi: \Lambda_G \rightarrow \Lambda_G$ is defined by

$$\text{TPro}_\pi = \tau_{\pi(n)} \cdots \tau_{\pi(2)} \tau_{\pi(1)}.$$

One would ideally hope to have an extension of Theorem 1.1 to arbitrary permutoric promotion operators. Unfortunately, trying to completely describe the orbit structure of $\text{TPro}_\pi: \Lambda_G \rightarrow \Lambda_G$ for arbitrary forests G and arbitrary permutations π seems to be very difficult. However, it turns out that we *can* do this when G is a path.¹ To state our main result, we need a bit more terminology.

Let $[k]_q = \frac{1-q^k}{1-q} = 1 + q + \cdots + q^{k-1}$ and $[k]_q! = [k]_q [k-1]_q \cdots [1]_q$. The *q-binomial coefficient* $\begin{bmatrix} k \\ r \end{bmatrix}_q$ is the polynomial $\frac{[k]_q!}{[r]_q! [k-r]_q!} \in \mathbb{C}[q]$.

Let X be a finite set, and let $f: X \rightarrow X$ be an invertible map of order ω (i.e., ω is the smallest positive integer such that $f^\omega(x) = x$ for all $x \in X$). Let $F(q) \in \mathbb{C}[q]$ be a polynomial in the variable q . Following [RSW04], we say the triple $(X, f, F(q))$ *exhibits the cyclic sieving phenomenon* if for every integer k , the number of elements of X fixed by f^k is $F(e^{2\pi ik/\omega})$.

Although we view the set $\mathbb{Z}/n\mathbb{Z}$ as a “cyclic” object, it will often be convenient to identify $\mathbb{Z}/n\mathbb{Z}$ with the “linear” set $[n]$ and consider the total ordering of its elements given by $1 < 2 < \cdots < n$. If $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ is a bijection, then a *cyclic descent* of π^{-1} is an element $i \in \mathbb{Z}/n\mathbb{Z}$ such that $\pi^{-1}(i) > \pi^{-1}(i+1)$ (note that n is permitted to be a cyclic descent).

Let Path_n denote the path graph with n vertices. In [Def23, Conjecture 4.1], Defant conjectured (using different language) that for every bijection $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$, the order of the map $\text{TPro}_\pi: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ is $d(n-d)$, where d is the number of cyclic descents of π^{-1} . Our main theorem not only proves this conjecture, but also determines the entire orbit structure of permutoric promotion in this case.

Theorem 1.3. *Let $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ be a bijection, and let d be the number of cyclic descents of π^{-1} . The order of the permutoric promotion operator $\text{TPro}_\pi: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ is $d(n-d)$. Moreover, the triple*

$$\left(\Lambda_{\text{Path}_n}, \text{TPro}_\pi, n(d-1)!(n-d-1)![n-d]_{q^d} \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q \right)$$

exhibits the cyclic sieving phenomenon.

Note that when $d = 1$, the sieving polynomial in Theorem 1.3 is $n(n-2)![n-1]_q$, which agrees with Theorem 1.1.

Suppose B is a proper subset of $\mathbb{Z}/n\mathbb{Z}$. In order to understand permutoric promotion and prove Theorem 1.3, we define the *broken promotion* operator $\text{Bro}_B: \Lambda_G \rightarrow \Lambda_G$ as follows. Let B_1, \dots, B_k be the vertex sets of the connected components of the subgraph of Cycle_n

¹We will see later that the dynamics of permutoric promotion on the path graph can be explained using coins that move around on the path and collide with each other. The path graph is especially nice because there are always at most two possible directions for a coin to move. One could try to mimic our approach for other trees, but additional complications would likely arise when considering vertices of degree greater than 2.

induced by B . For $1 \leq i \leq k$, let us write $B_i = \{a(i), a(i) + 1, \dots, b(i)\}$, and let $\text{Bro}_{B_i} = \tau_{b(i)} \cdots \tau_{a(i)+1} \tau_{a(i)}$. We then define $\text{Bro}_B = \text{Bro}_{B_1} \cdots \text{Bro}_{B_k}$ (the order does not matter since $\text{Bro}_{B_1}, \dots, \text{Bro}_{B_k}$ commute with each other).

Let $\text{cyc}: \Lambda_G \rightarrow \Lambda_G$ be the *cyclic shift* operator defined by $(\text{cyc}(\sigma))(v) = \sigma(v) + 1$. In Section 3, we give a description of the operator cyc Bro_B in terms of “gliding globs” of liquid. Roughly speaking, some of the labels are immersed in globs of liquid, these globs (and their labels) glide along paths in G in a *jeu de taquin* fashion, and then some of the labels are changed appropriately. We also show that certain indicator functions are *homomesic* for cyc Bro_B (see Proposition 3.6). In Section 4, we specialize to the case when $G = \text{Path}_n$ and establish useful connections between broken promotion and permutoric promotion. The purpose of Section 5 is to prove that all of the sizes of the orbits of TPro_π are divisible by $\text{lcm}(d, n - d)$ (where $G = \text{Path}_n$ and d is the number of cyclic descents of π^{-1}). In Section 6, we use this divisibility result to reformulate the analysis of permutoric promotion in terms of “sliding stones” and “colliding coins.” Roughly speaking, we place some stones on the cycle graph and allow them slide around as we apply toggle operators. At the same time, we place coins on the path graph and allow them to move around and collide with one another. It turns out that the dynamical properties of permutoric promotion are closely related to those of the *stones diagrams* and *coins diagrams*; this allows us to complete the proof of Theorem 1.3. Let $\text{Comp}_d(n)$ be the set of compositions of n into d parts, and define $\text{Rot}_{n,d}: \text{Comp}_d(n) \rightarrow \text{Comp}_d(n)$ by $\text{Rot}_{n,d}(a_1, a_2, \dots, a_d) = (a_2, \dots, a_d, a_1)$. We show how to associate an orbit of $\text{Rot}_{n,d}$ to the dynamics of the coins diagrams by recording how far each coin must travel when passing from one collision to the next. It will turn out that the form of the sieving polynomial in Theorem 1.3 arises from the fact that the triple $(\text{Comp}_d(n), \text{Rot}_{n,d}, \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q)$ exhibits the cyclic sieving phenomenon. In Section 7, we apply Theorem 1.3 to derive the following theorems.

Theorem 1.4. *Let d and n be integers such that $1 \leq d \leq n - 1$. The order of the operator $\text{cyc Bro}_{\{1, \dots, d\}}: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ is $(n - d)n$. Moreover, the triple*

$$\left(\Lambda_{\text{Path}_n}, \text{cyc Bro}_{\{1, \dots, d\}}, (d - 1)!(n - d - 1)! [n]_{q^{n-d}} [n - d]_{q^d} \begin{bmatrix} n - 1 \\ d - 1 \end{bmatrix}_q \right)$$

exhibits the cyclic sieving phenomenon.

For any real number x , let $\llbracket x \rrbracket$ denote the integer closest to x , with the convention that $\llbracket x \rrbracket = x - 1/2$ if $x - 1/2 \in \mathbb{Z}$.

Theorem 1.5. *Let d and n be positive integers such that $1 \leq d \leq \lfloor n/2 \rfloor$. For $i \in \mathbb{Z}$, let $s_i = \llbracket in/d \rrbracket$, and let $\mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus \{s_1 - 1, \dots, s_d - 1\}$. The order of the operator $\text{cyc Bro}_{\mathcal{R}}: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ is dn . Moreover, the triple*

$$\left(\Lambda_{\text{Path}_n}, \text{cyc Bro}_{\mathcal{R}}, (d - 1)!(n - d - 1)! [n]_{q^d} [n - d]_{q^d} \begin{bmatrix} n - 1 \\ d - 1 \end{bmatrix}_q \right)$$

exhibits the cyclic sieving phenomenon.

While noteworthy on its own, the homomesy result from Section 3 also ends up being useful for proving Theorems 1.4 and 1.5.

2. Basics

Let n be a positive integer. Given integers $x \leq y$, we let $[x, y]_n$ denote the tuple of elements in $\mathbb{Z}/n\mathbb{Z}$ obtained by reducing the entries in the tuple $(x, x + 1, \dots, y)$ modulo n . For example, $[3, 7]_3$ is the tuple $(0, 1, 2, 0, 1)$, where the elements are in $\mathbb{Z}/3\mathbb{Z}$. Slightly abusing notation, we will also sometimes view $[x, y]_n$ as a (multi)set. In this context, $[3, 7]_3$ is the multiset $\{0, 0, 1, 1, 2\}$.

Given a finite set X and an invertible map $f : X \rightarrow X$, we let Orb_f denote the set of orbits of f . We will need the following technical lemma concerning the cyclic sieving phenomenon.

Lemma 2.1. *Let $f : X \rightarrow X$ and $g : \tilde{X} \rightarrow \tilde{X}$ be invertible maps, where X and \tilde{X} are finite sets. Let $\{k_i^{m_i} : 1 \leq i \leq \ell\}$ be the multiset of orbit sizes of f , where we use superscripts to denote multiplicities. Let ω be the order of f , and let $F(q) \in \mathbb{C}[q]$ be such that the triple $(X, f, F(q))$ exhibits the cyclic sieving phenomenon. If $N \in \mathbb{Z}_{>0}$ and $\alpha \in \mathbb{Q}_{>0}$ are such that $\{(Nk_i)^{\alpha m_i} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of g , then g has order $N\omega$, and the triple $(\tilde{X}, g, \alpha[N]_{q^\omega} F(q))$ exhibits the cyclic sieving phenomenon.*

Proof. It is clear that g has order $N\omega$. Fix an integer k . When we evaluate the polynomial $\alpha[N]_{q^\omega} F(q)$ at $q = e^{2\pi ik/(N\omega)}$, we obtain $\alpha[N]_{e^{2\pi ik/N}} F(e^{2\pi i(k/N)/\omega})$; we want to show that this is the number of elements of \tilde{X} that are fixed by g^k . If k is not divisible by N , then there are no such elements because all orbits of g have sizes divisible by N ; in this case, we are done because the factor $[N]_{e^{2\pi ik/N}}$ is 0. Now suppose k is divisible by N . Because $(X, f, F(q))$ exhibits the cyclic sieving phenomenon, $F(e^{2\pi i(k/N)/\omega})$ is the number of elements of X fixed by $f^{k/N}$. Therefore, $\alpha N F(e^{2\pi i(k/N)/\omega})$ is the number of elements of \tilde{X} fixed by g^k . This completes the proof because $\alpha[N]_{e^{2\pi ik/N}} F(e^{2\pi i(k/N)/\omega}) = \alpha N F(e^{2\pi i(k/N)/\omega})$. \square

Remark 2.2. The previous lemma also admits another proof, as was kindly pointed out to us by one of the referees. Fix a positive integer ω , and consider the cyclic group of order ω , which we identify with the ω -th roots of unity, denoted μ_ω . Let ρ be a permutation of some finite set, with order dividing ω . The permutation action of μ_ω according to which $e^{2\pi i/\omega}$ acts by ρ has a character, which we denote $\chi_\rho : \mu_\omega \rightarrow \mathbb{C}$. If ρ is a single cycle of length a (which by hypothesis divides ω), then one checks that $\chi_\rho(q) = [a]_{q^{\omega/a}}$. Further, the character of an action consisting of several disjoint cycles is the product of the characters of each of the cycles. To suppose, as in the statement of the lemma, that $(X, f, F(q))$ exhibits the cyclic sieving phenomenon, is equivalent to saying that $F(q)$ is congruent to the character $\chi_f(q)$ modulo $q^\omega - 1$. Using the fact that $[N]_{q^\omega} [a]_{q^{\omega/a}} = [Na]_{q^{\omega/a}}$, we deduce that $\alpha[N]_{q^\omega} F(q)$ is congruent to $\chi_g(q)$ modulo $q^{N\omega} - 1$, and thus that $(\tilde{X}, g, \alpha[N]_{q^\omega} F(q))$ exhibits the cyclic sieving phenomenon.

We write Path_n and Cycle_n for the path with n vertices and the cycle with n vertices, respectively. We embed these graphs in the plane, drawing Path_n horizontally. Identify the vertices of Cycle_n with $\mathbb{Z}/n\mathbb{Z}$ in such a way that they appear in the cyclic order $1, 2, \dots, n$ when read clockwise around the cycle. Let v_1, \dots, v_n be the vertices of Path_n , listed from left to right. Just as the purpose of embedding Path_n in the plane is to be able to talk about left and right,

the purpose of embedding Cycle_n in the plane is to be able to talk about clockwise and counterclockwise. As in the introduction, let us fix a graph G and consider the toggle operators τ_i and the permutoric promotion operators TPro_π on Λ_G .

If \mathcal{D} is an acyclic directed graph with vertex set \mathcal{V} , then we can define a partial order $\leq_{\mathcal{D}}$ on \mathcal{V} by declaring $v \leq_{\mathcal{D}} v'$ whenever there is a directed path in \mathcal{D} from v to v' ; the resulting poset $(\mathcal{V}, \leq_{\mathcal{D}})$ is called the *transitive closure* of \mathcal{D} .

A *linear extension* of an n -element poset (P, \leq_P) is a word $p_1 \cdots p_n$ whose letters are the elements of P (with each element appearing exactly once) such that $i \leq j$ whenever $p_i \leq_P p_j$. Given a bijection $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$, we obtain an acyclic orientation α_π of Cycle_n by orienting each edge $\{i, i+1\}$ from i to $i+1$ if and only if $\pi^{-1}(i) < \pi^{-1}(i+1)$. If β is any acyclic orientation of Cycle_n , then the linear extensions of its transitive closure $(\mathbb{Z}/n\mathbb{Z}, \leq_\beta)$ are precisely the words $\pi(1) \cdots \pi(n)$ such that $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ is a bijection satisfying $\alpha_\pi = \beta$.

It is well known that any linear extension of a finite poset can be obtained from any other linear extension of the same poset by repeatedly swapping consecutive incomparable elements. If $i, j \in \mathbb{Z}/n\mathbb{Z}$ are incomparable in $(\mathbb{Z}/n\mathbb{Z}, \leq_\beta)$, then they are not adjacent in Cycle_n , so the toggle operators τ_i and τ_j commute. This implies that if $\pi, \pi': [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ are such that $\alpha_\pi = \alpha_{\pi'}$, then the expression for $\text{TPro}_{\pi'}$ as a composition of toggle operators can be obtained from the expression for TPro_π by repeatedly swapping consecutive toggle operators that commute with each other, so $\text{TPro}_\pi = \text{TPro}_{\pi'}$. Therefore, given an acyclic orientation β of Cycle_n , it makes sense to write TPro_β for the permutoric promotion operator TPro_π , where $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ is any bijection such that $\alpha_\pi = \beta$. (This argument showing why TPro_β is well defined is essentially the same as a standard argument showing that Coxeter elements of a Coxeter group correspond bijectively to acyclic orientations of the Coxeter graph.)

A *source* (respectively, *sink*) of an acyclic orientation is a vertex of in-degree (respectively, out-degree) 0. If u is a source (respectively, sink), then we can *flip* u into a sink (respectively, source) by reversing the orientations of all edges incident to u . Two acyclic orientations are *flip equivalent* if one can be obtained from the other by a sequence of flips.

Let us say two maps $f, g: \Lambda_G \rightarrow \Lambda_G$ are *dynamically equivalent* if there is a bijection $\phi: \Lambda_G \rightarrow \Lambda_G$ such that $f \circ \phi = \phi \circ g$. Note that dynamically equivalent invertible maps have the same orbit structure (that is, they have the same number of orbits of each size).

Lemma 2.3. *If β and β' are acyclic orientations of Cycle_n that have the same number of edges oriented counterclockwise, then TPro_β and $\text{TPro}_{\beta'}$ are dynamically equivalent.*

Proof. It is known (see [DMR16]) that two acyclic orientations of Cycle_n have the same number of edges oriented counterclockwise if and only if they are flip equivalent. Therefore, we just need to show that if β and β' are flip equivalent, then TPro_β and $\text{TPro}_{\beta'}$ are dynamically equivalent. It suffices to prove this in the case when β' is obtained from β by flipping a source i into a sink. In this case, one can check that $\text{TPro}_\beta \circ \tau_i = \tau_i \circ \text{TPro}_{\beta'}$. \square

Lemma 2.4. *Let β and β' be acyclic orientations of Cycle_n . Let d and d' be the number of edges oriented counterclockwise in β and β' , respectively. If $d = d'$ or $d = n - d'$, then TPro_β and $\text{TPro}_{\beta'}$ are dynamically equivalent.*

Proof. If $d = d'$, then we are done by Lemma 2.3. Now suppose $d = n - d'$. Define $\phi: \Lambda_G \rightarrow \Lambda_G$ by $(\phi(\sigma))(v) = n + 1 - \sigma(v)$. One can readily check that

$$\phi \circ \tau_i = \tau_{n-i} \circ \phi \tag{2.1}$$

for all $i \in \mathbb{Z}/n\mathbb{Z}$. Let $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ be a bijection such that $\alpha_\pi = \beta$, and define $\pi': [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ by $\pi'(i) = n - \pi(i)$. We have $\text{TPro}_\pi = \text{TPro}_\beta$. It follows from (2.1) that $\phi \circ \text{TPro}_\pi = \text{TPro}_{\pi'} \circ \phi$. This shows that TPro_π and $\text{TPro}_{\pi'}$ are dynamically equivalent. On the other hand, the number of edges oriented counterclockwise in $\alpha_{\pi'}$ is $n - d$, so it follows from Lemma 2.3 that $\text{TPro}_{\pi'}$ is dynamically equivalent to TPro_β . \square

If $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ is a bijection, then the number of cyclic descents of π^{-1} is the same as the number of edges oriented counterclockwise in α_π . This is why cyclic descents appear in Theorem 1.3.

We end this section with a lemma that will allow us to rewrite operators formed as compositions of toggles. We will consider words over the alphabet $\{\tau_1, \dots, \tau_n\}$ both as words and as permutations of Λ_G . Given such a word X , let $X\langle i \rangle$ denote the number of occurrences of τ_i in X .

Lemma 2.5. *Let β be an acyclic orientation of Cycle_n . Let Y be a word over the alphabet $\{\tau_1, \dots, \tau_n\}$ in which each letter appears exactly k times. Suppose that for every suffix X of Y and every arrow $a \rightarrow b$ in β , we have $X\langle a \rangle - X\langle b \rangle \in \{0, 1\}$. When viewed as a bijection from Λ_G to itself, Y is equal to TPro_β^k .*

Proof. For each $i \in \mathbb{Z}/n\mathbb{Z}$, consider k formal symbols $\tau_i^{(1)}, \dots, \tau_i^{(k)}$. Let \mathcal{G} be the group generated by the set $\mathbf{A} = \{\tau_i^{(\ell)} : i \in \mathbb{Z}/n\mathbb{Z}, \ell \in [k]\}$ subject to the relations $\tau_i^{(\ell)} \tau_j^{(m)} = \tau_j^{(m)} \tau_i^{(\ell)}$ whenever $j \notin \{i - 1, i, i + 1\}$. Let \mathcal{D} be the directed graph with vertex set \mathbf{A} and with arrows defined as follows: for each arrow $a \rightarrow b$ in β , the graph \mathcal{D} has arrows $\tau_b^{(\ell)} \rightarrow \tau_a^{(\ell)}$ for all $\ell \in [k]$ and $\tau_a^{(m)} \rightarrow \tau_b^{(m+1)}$ for all $m \in [k - 1]$. Let $(\mathbf{A}, \leq_{\mathcal{D}})$ be the transitive closure of \mathcal{D} .

Let us fix a bijection $\pi: [n] \rightarrow \mathbb{Z}/n\mathbb{Z}$ such that $\text{TPro}_\pi = \text{TPro}_\beta$. Let Y' be the word $(\tau_{\pi(n)} \cdots \tau_{\pi(1)})^k$ over the alphabet $\{\tau_1, \dots, \tau_n\}$. Let Z (respectively, Z') be the word over the alphabet \mathbf{A} obtained from Y (respectively, Y') by replacing the ℓ -th occurrence of the letter τ_i with $\tau_i^{(\ell)}$. The conditions on Y in the hypothesis of the lemma imply that Z is a linear extension of $(\mathbf{A}, \leq_{\mathcal{D}})$; the word Y' satisfies the same conditions, so Z' is also a linear extension of $(\mathbf{A}, \leq_{\mathcal{D}})$. This means that Z' can be obtained from Z by repeatedly swapping consecutive incomparable elements; each such swap corresponds to one of relations defining \mathcal{G} . Thus, Z and Z' represent the same element of \mathcal{G} . There is a natural homomorphism from \mathcal{G} to the group of permutations of Λ_G that sends each generator $\tau_i^{(\ell)}$ to τ_i . This homomorphism sends Z and Z' to the permutations of Λ_G represented by Y and Y' , respectively. Hence, these permutations are the same. This completes the proof because the permutation represented by Y' is TPro_β^k . \square

3. Broken Promotion

In this section, we study the *broken promotion* operators defined in the introduction, describing them in terms of “gliding globs” and relating them to permutoric promotion operators.

3.1. Jeu de Taquin

As before, let us fix an n -vertex graph $G = (V, E)$. Our arguments in this section will require certain *jeu de taquin* operators defined as follows. For $i_1, i_2 \in \mathbb{Z}/n\mathbb{Z}$, define $\text{jdt}_{(i_1, i_2)}: \Lambda_G \rightarrow \Lambda_G$ by

$$\text{jdt}_{(i_1, i_2)}(\sigma) = \begin{cases} (i_1 \ i_2) \circ \sigma & \text{if } \{\sigma^{-1}(i_1), \sigma^{-1}(i_2)\} \in E; \\ \sigma & \text{if } \{\sigma^{-1}(i_1), \sigma^{-1}(i_2)\} \notin E. \end{cases}$$

Thus, $\text{jdt}_{(i_1, i_2)}$ has the effect of trying to “glide” the label i_1 through the label i_2 ; it succeeds in doing so if and only if those labels are on adjacent vertices of G . More generally, if (i_1, \dots, i_r) is a tuple of distinct vertices in V , then we define

$$\text{jdt}_{(i_1, \dots, i_r)} = \text{jdt}_{(i_1, i_r)} \text{jdt}_{(i_1, i_{r-1})} \cdots \text{jdt}_{(i_1, i_2)}.$$

This operator has the effect of trying to glide i_1 through the labels i_2, \dots, i_r in that order. We will primarily be interested in the case when (i_1, \dots, i_r) is such that $i_j = i_1 + j - 1$ for all $1 \leq j \leq r$. In this case, $\{i_1, \dots, i_r\}$ is a cyclic interval $[x, y]_n$, so we simply write $\text{jdt}_{[x, y]_n}$ instead of $\text{jdt}_{(i_1, \dots, i_r)}$.

Example 3.1. If $n = 6$ and σ is the labeling shown on the left in Figure 3.1, then the labeling $\text{jdt}_{[5, 9]_6}(\sigma) = \text{jdt}_{(5, 6, 1, 2, 3)}(\sigma)$ is shown on the right in Figure 3.1.




Figure 3.1: On the left is a labeling of a 6-vertex graph, where the label of each vertex is shown next to it in red. On the right is the labeling $\text{jdt}_{[5, 9]_6}(\sigma) = \text{jdt}_{(5, 6, 1, 2, 3)}(\sigma)$. The labeling $\text{jdt}_{[5, 9]_6}(\sigma)$ was obtained from σ by gliding the label 5 through the labels 6, 1, 3.

3.2. Broken Promotion

Suppose B is a proper subset of $\mathbb{Z}/n\mathbb{Z}$. Recall the definitions of the *broken promotion* operator $\text{Bro}_B: \Lambda_G \rightarrow \Lambda_G$ and the *cyclic shift* operator $\text{cyc}: \Lambda_G \rightarrow \Lambda_G$ from Section 1. We can explicitly describe the action of cyc Bro_B on a labeling $\sigma \in \Lambda_G$ as follows. Let B_1, \dots, B_k be the vertex sets of the connected components of the subgraph of Cycle_n induced by B . For each $1 \leq i \leq k$, let x_i and y_i be such that $B_i = [x_i, y_i - 1]_n$, and imagine immersing the label x_i

in a glob of liquid. The first step is to apply the jeu de taquin operators $\text{jdt}_{[x_i, y_i]_n}$, imagining that the label x_i carries its glob along with it as it glides. For the second step, increase by 1 the label of each vertex in $\sigma^{-1}((\mathbb{Z}/n\mathbb{Z}) \setminus \bigcup_{\ell=1}^k [x_\ell, y_\ell]_n)$. If $x_i - 1 \notin \bigcup_{\ell=1}^k [x_\ell, y_\ell]_n$, this second step will change the label $x_i - 1$ into x_i , so there will be two copies of the label x_i : one in a glob and the other not in a glob. The third and final step is to change each label x_i that is in a glob to the label $y_i + 1$.

It might not be obvious at first that the procedure described in the preceding paragraph does in fact compute $\text{cyc Bro}_B(\sigma)$; however, the verification of this fact is straightforward and can be elucidated through examples.

Example 3.2. Suppose $n = 9$ and $G = \text{Path}_9$. Let $B = \{1, 3, 4, 7, 9\} \subseteq \mathbb{Z}/9\mathbb{Z}$. The connected components of the subgraph of Cycle_9 induced by B have vertex sets

$$B_1 = \{3, 4\} = [3, 4]_9, \quad B_2 = \{7\} = [7, 7]_9, \quad B_3 = \{9, 1\} = [9, 10]_9.$$

Preserving the notation from above, we have $x_1 = 3, y_1 = 5, x_2 = 7, y_2 = 8, x_3 = 9, y_3 = 11$. Recall that the vertices of Path_9 are v_1, \dots, v_9 ; let $\sigma \in \Lambda_{\text{Path}_9}$ be the labeling that sends these vertices to 7, 1, 4, 3, 5, 6, 9, 2, 8, respectively. Figure 3.2 illustrates the three-step procedure for computing $\text{cyc Bro}_B(\sigma)$, showing that $\text{cyc Bro}_B(\sigma)$ sends v_1, \dots, v_9 to 9, 1, 6, 4, 5, 7, 2, 3, 8, respectively.




Figure 3.2: The three steps for applying cyc Bro_B , where $B = \{1, 3, 4, 7, 9\} \subseteq \mathbb{Z}/9\mathbb{Z}$.

3.3. Broken Promotion for the Complement of an Independent Set

Suppose $1 \leq d \leq \lfloor n/2 \rfloor$, and let $\dots < s_{-1} < s_0 < s_1 < s_2 < \dots$ be a bi-infinite sequence of integers such that $s_{i+d} = s_i + n$ and $s_{i+1} \geq s_i + 2$ for all $i \in \mathbb{Z}$. Then $\mathcal{S} = \{s_1, \dots, s_d\}$ is an independent set of size d in Cycle_n . Let $\beta_{\mathcal{S}}$ be the acyclic orientation of Cycle_n in which the elements of \mathcal{S} are sources and all edges not incident to elements of \mathcal{S} are oriented clockwise. The sinks of $\beta_{\mathcal{S}}$ are the elements of $\mathcal{S} - 1 := \{s_1 - 1, \dots, s_d - 1\}$. Let us write $\mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus (\mathcal{S} - 1)$.

In Section 3.2, we gave a three-step description of the action of cyc Bro_B on a labeling σ when B is an arbitrary proper subset of $\mathbb{Z}/n\mathbb{Z}$. This description simplifies when B is \mathcal{R} because in this case, we have $x_i = s_i$ and $y_i = s_{i+1} - 1$, so $\bigcup_{\ell=1}^k [x_\ell, y_\ell]_n$ is all of $\mathbb{Z}/n\mathbb{Z}$ (so the second step in the earlier description has no effect). Hence, we have the following simpler two-step

procedure. Immerse each label s_i in a glob of liquid. The first step is to apply the jeu de taquin operators $\text{jdt}_{[s_i, s_{i+1}-1]_n}$ (for $1 \leq i \leq d$), imagining that the label s_i carries its glob with it as it glides. The second step is to cyclically rotate the labels in the globs, changing each label s_i to s_{i+1} (modulo n).

Example 3.3. Suppose $n = 9$ and $d = 3$. Let $s_1 = 3$, $s_2 = 7$, $s_3 = 9$. Then $\mathcal{S} = \{3, 7, 9\}$, $\mathcal{S} - 1 = \{2, 6, 8\}$, and $\mathcal{R} = (\mathbb{Z}/9\mathbb{Z}) \setminus (\mathcal{S} - 1) = \{1, 3, 4, 5, 7, 9\}$. The first step in the above two-step procedure for applying $\text{cyc Bro}_{\mathcal{R}}$ is to immerse 3, 7, and 9 in globs of liquid and apply $\text{jdt}_{[3,6]_9}$, $\text{jdt}_{[7,8]_9}$, and $\text{jdt}_{[9,11]_9}$. The second step is to cyclically rotate the labels 3, 7, 9. This is illustrated in Figure 3.3.




Figure 3.3: The two steps for applying $\text{cyc Bro}_{\mathcal{R}}$, where $\mathcal{R} = \{1, 3, 4, 5, 7, 9\} \subseteq \mathbb{Z}/9\mathbb{Z}$.

Remark 3.4. Suppose $G = \text{Path}_n$. Neither of the two steps in the above procedure change the relative order in which the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ (i.e., the labels not in the globs) appear from left to right along the path. For example, in Figure 3.3, the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ are 1, 2, 4, 5, 6, 8. In every step of the procedure, these labels appear in the order 1, 4, 5, 6, 2, 8.

3.4. Permutoric Promotion and Broken Promotion

The following lemma shows that the operators $\text{cyc Bro}_{\mathcal{R}}$ and $\text{TPro}_{\beta, \mathcal{S}}$ commute; it will later allow us to use the description of $\text{cyc Bro}_{\mathcal{R}}$ in terms of gliding globs given in Section 3.3 to gain a better understanding of permutoric promotion.

Lemma 3.5. *Let \mathcal{S} be a d -element independent set in Cycle_n , and let $\mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus (\mathcal{S} - 1)$. Then*

$$\text{cyc Bro}_{\mathcal{R}} \text{TPro}_{\beta, \mathcal{S}} = \text{TPro}_{\beta, \mathcal{S}} \text{cyc Bro}_{\mathcal{R}}.$$

Proof. Preserve the notation from Section 3.3. Let $\mathcal{B} = \mathcal{R} \setminus \mathcal{S}$. The vertex sets of the connected components of the subgraph of Cycle_n induced by \mathcal{R} are

$$[s_1, s_2 - 2]_n, [s_2, s_3 - 2]_n, \dots, [s_d, s_{d+1} - 2]_n,$$

SO

$$\begin{aligned} \text{Bro}_{\mathcal{R}} &= \prod_{i=1}^d \text{Bro}_{[s_i, s_{i+1}-2]_n} \\ &= \prod_{i=1}^d \text{Bro}_{[s_i+1, s_{i+1}-2]_n} \tau_{s_i} \\ &= \prod_{i=1}^d \text{Bro}_{[s_i+1, s_{i+1}-2]_n} \prod_{i=1}^d \tau_{s_i} \\ &= \text{Bro}_{\mathcal{B}} \text{Bro}_{\mathcal{S}}. \end{aligned}$$

(We are using the fact that τ_i and τ_j commute whenever $j \notin \{i \pm 1\}$.) A similar computation shows that

$$\text{cyc Bro}_{\mathcal{R}} \text{cyc}^{-1} = \prod_{i=1}^d \text{Bro}_{[s_i+1, s_{i+1}-1]_n} = \text{Bro}_{\mathcal{S}-1} \text{Bro}_{\mathcal{R}}.$$

We also have $\text{cyc Bro}_{\mathcal{S}-1} \text{cyc}^{-1} = \text{Bro}_{\mathcal{S}}$. Therefore,

$$\begin{aligned} \text{cyc Bro}_{\mathcal{R}} \text{Bro}_{\mathcal{S}-1} &= (\text{cyc Bro}_{\mathcal{R}} \text{cyc}^{-1})(\text{cyc Bro}_{\mathcal{S}-1} \text{cyc}^{-1}) \text{cyc} \\ &= \text{Bro}_{\mathcal{S}-1}(\text{Bro}_{\mathcal{R}} \text{Bro}_{\mathcal{S}}) \text{cyc} \\ &= \text{Bro}_{\mathcal{S}-1} \text{Bro}_{\mathcal{R}} \text{cyc}. \end{aligned}$$

This shows that

$$\text{cyc Bro}_{\mathcal{R}} \text{Bro}_{\mathcal{S}-1} \text{Bro}_{\mathcal{R}} = \text{Bro}_{\mathcal{S}-1} \text{Bro}_{\mathcal{R}} \text{cyc Bro}_{\mathcal{R}}.$$

The desired result now follows from the observation that $\text{TPro}_{\beta, \mathcal{S}} = \text{Bro}_{\mathcal{S}-1} \text{Bro}_{\mathcal{R}}$. □

3.5. Homomesy

We end this section with a theorem about broken promotion that will be useful in Section 7 but that we also believe is interesting in its own right. This proposition concerns the notion of *homomesy*, which Propp and Roby introduced in 2015 [PR15]; it is now one of the central focuses in dynamical algebraic combinatorics. Suppose X is a finite set and $f: X \rightarrow X$ is an invertible map. Let Orb_f denote the set of orbits of f . A *statistic* on X is a function $\text{stat}: X \rightarrow \mathbb{R}$. We say the statistic stat is *homomesic* for f with average a if $\frac{1}{|\mathcal{O}|} \sum_{x \in \mathcal{O}} \text{stat}(x) = a$ for every orbit $\mathcal{O} \in \text{Orb}_f$.

Proposition 3.6. *Suppose G is connected. Let v be a vertex of G , and let $i \in \mathbb{Z}/n\mathbb{Z}$. Define $\mathbb{1}_{v,i}: \Lambda_G \rightarrow \mathbb{R}$ by*

$$\mathbb{1}_{v,i}(\sigma) = \begin{cases} 1 & \text{if } \sigma(v) = i; \\ 0 & \text{if } \sigma(v) \neq i. \end{cases}$$

If $B \subseteq \mathbb{Z}/n\mathbb{Z}$ and $i - 1 \notin B$, then $\mathbb{1}_{v,i}$ is homomesic for the map cyc Bro_B with average $1/n$.

Proof. By symmetry, it suffices to prove the result when $i = 1$. We identify $\mathbb{Z}/n\mathbb{Z}$ with $[n]$ and consider the total ordering $1 < 2 < \dots < n$. Given a labeling $\sigma \in \Lambda_G$, we obtain an acyclic orientation η_σ by orienting each edge $\{x, y\}$ from x to y if and only if $\sigma(x) < \sigma(y)$. Observe that $\eta_{\tau_j(\sigma)} = \eta_\sigma$ for every $j \in [n-1]$ and $\sigma \in \Lambda_G$; since $n = i-1 \notin B$, this implies that $\eta_{\text{Bro}_B(\sigma)} = \eta_\sigma$ for every $\sigma \in \Lambda_G$. It is also straightforward to see that $\eta_{\text{cyc}(\sigma)}$ is obtained from η_σ by flipping the vertex $(\text{cyc}(\sigma))^{-1}(1)$ from a sink to a source. Therefore, $\eta_{\text{cyc}(\text{Bro}_B(\sigma))}$ is obtained from η_σ by flipping $(\text{cyc}(\text{Bro}_B(\sigma)))^{-1}(1)$ from a sink to a source.

Let \mathcal{O} be an orbit of cyc Bro_B , and fix $\mu_0 \in \mathcal{O}$. Let $\mu_t = (\text{cyc Bro}_B)^t(\mu_0)$ for all $t \in \mathbb{Z}$. Consider an edge $\{x_0, y_0\}$ in G . Let $\dots < k(0) < k(1) < k(2) < \dots$ be the integers such that $\mu_{k(j)}^{-1}(1) \in \{x_0, y_0\}$. Without loss of generality, assume $x_0 \rightarrow y_0$ is an arrow in $\eta_{\mu_{k(0)}}$. According to the previous paragraph, the orientations of $\{x_0, y_0\}$ are different in $\eta_{\mu_{t-1}}$ and η_{μ_t} if and only if $t \in \{\dots, k(0), k(1), k(2), \dots\}$. Moreover, we have $\mu_{k(j)}(x_0) = 1$ if $x_0 \rightarrow y_0$ is an arrow in $\eta_{\mu_{k(j)}}$, whereas $\mu_{k(j)}(y_0) = 1$ if $y_0 \rightarrow x_0$ is an arrow in $\eta_{\mu_{k(j)}}$. It follows that for $t \in \mathbb{Z}$, we have $\mu_t(x_0) = 1$ if and only if $t = k(j)$ for some even j ; similarly, $\mu_t(y_0) = 1$ if and only if $t = k(j)$ for some odd j . This shows that the number of labelings in \mathcal{O} that send x_0 to 1 is the same as the number of labelings in \mathcal{O} that send y_0 to 1. Because the edge $\{x_0, y_0\}$ was arbitrary and G is connected, it follows that for any two vertices x and y of G , the number of labelings in \mathcal{O} that send x to 1 is the same as the number of labelings in \mathcal{O} that send y to 1. This implies the desired result. \square

Example 3.7. Suppose $n = 5$ and $B = \{1, 3, 4\}$. Figure 3.4 depicts an orbit of cyc Bro_B for a particular choice of a graph G . Select an arbitrary vertex v of G . As predicted by Proposition 3.6, exactly 1 of the 5 labelings in this orbit sends v to 1, and exactly 1 of the 5 labelings in the orbit sends v to 3.




Figure 3.4: An orbit of $\text{cyc Bro}_{\{1,3,4\}}$ for a particular graph G . Note that each of the labels 1 and 3 appears on each vertex exactly once throughout the orbit.

4. Broken Promotion on a Path

Throughout the rest of the article, we will specialize to the case when $G = \text{Path}_n$ is the path with n vertices.

Suppose $1 \leq d \leq \lfloor n/2 \rfloor$. In Section 3, we considered an arbitrary bi-infinite sequence

$$\cdots < s_{-1} < s_0 < s_1 < s_2 < \cdots$$

satisfying $s_{i+d} = s_i + n$ and $s_{i+1} \geq s_i + 2$ for all $i \in \mathbb{Z}$. In this section, we specialize our attention to a particular sequence. We write $\lceil x \rceil$ for the integer closest to a real number x , with the convention that $\lceil x \rceil = x - 1/2$ if $x - 1/2 \in \mathbb{Z}$. For $i \in \mathbb{Z}$, let $s_i = \lceil in/d \rceil$. As in Section 3, we let \mathcal{S} be the independent set $\{s_1, \dots, s_d\}$ of Cycle_n and set $\mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus (\mathcal{S} - 1)$. Let $\beta = \beta_{\mathcal{S}}$ be the acyclic orientation of Cycle_n whose sources are the elements of \mathcal{S} and whose sinks are the elements of $\mathcal{S} - 1$. Then β has exactly d counterclockwise edges.

In what follows, we view tuples $[x, y]_n$ as multisets. When we consider the size of the intersection of a multiset with a set, we count the elements according to their multiplicity in the multiset. For example, in the next proposition, $|[j - q, j - 1]_n \cap (\mathcal{S} - 1)|$ should be interpreted as the number of elements of $[j - q, j - 1]_n$, counted with multiplicity, that are also elements of the set $\mathcal{S} - 1$.

Proposition 4.1. *Suppose γ, q, r are nonnegative integers such that $0 \leq r \leq n - d - 1$ and $\gamma n = q(n - d) + r$ (that is to say, q and r are respectively the quotient and remainder when dividing γn by $n - d$). Let $J = \{j \in [n] : q - \gamma + 1 \leq |[j - q, j - 1]_n \cap (\mathcal{S} - 1)|\}$. Then $|J| = r$, and $\text{TPro}_\beta^\gamma = \text{cyc}^{-q} \text{Bro}_J(\text{cyc Bro}_{\mathcal{R}})^q$.*

Proposition 4.1 will be crucial in the next section when we prove that the sizes of the orbits of TPro_β are all divisible by $\text{lcm}(d, n - d)$. Before proving this proposition, we need a technical lemma.

Lemma 4.2. *Let γ, q, r, J be as in Proposition 4.1. We have $J \cap (\mathcal{S} - 1) = \emptyset$. Also, if $i \in \mathcal{R}$ and $i + 1 \in J$, then $i \in J$.*

Proof. Consider some $s_j - 1 \in \mathcal{S} - 1$. Recall that $s_j = \lceil jn/d \rceil$. It is straightforward to check that

$$|[s_j - 1 - q, s_j - 2]_n \cap (\mathcal{S} - 1)| = |[s_j - 1 - q, s_j - 1]_n \cap (\mathcal{S} - 1)| - 1 \leq \frac{(q + 1)d}{n}.$$

Because $r \leq n - d - 1$, we have

$$|[s_j - 1 - q, s_j - 2]_n \cap (\mathcal{S} - 1)| < \frac{(q + 1)d}{n} + \frac{n - d - r}{n} = q - \frac{q(n - d) + r}{n} + 1 = q - \gamma + 1,$$

so $s_j - 1 \notin J$. This proves that $J \cap (\mathcal{S} - 1) = \emptyset$.

Now suppose $i \in \mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus (\mathcal{S} - 1)$ and $i + 1 \in J$. We have

$$|[i - q, i - 1]_n \cap (\mathcal{S} - 1)| \geq |[i + 1 - q, i + 1 - 1]_n \cap (\mathcal{S} - 1)| \geq q - \gamma + 1,$$

so $i \in J$. □

Proof of Proposition 4.1. As in Lemma 2.5, we consider words over the alphabet $\{\tau_1, \dots, \tau_n\}$ both as words and as permutations of Λ_{Path_n} . Given such a word X , recall that we write $X\langle i \rangle$ for the number of occurrences of τ_i in X .

Let i_1, \dots, i_{n-d} be an ordering of the elements of \mathcal{R} such that $\text{Bro}_{\mathcal{R}} = \tau_{i_{n-d}} \cdots \tau_{i_1}$. Consider the word

$$W = \tau_{i_{n-d}-(q-1)} \cdots \tau_{i_1-(q-1)} \tau_{i_{n-d}-(q-2)} \cdots \tau_{i_1-(q-2)} \cdots \tau_{i_{n-d}-1} \cdots \tau_{i_1-1} \tau_{i_{n-d}} \cdots \tau_{i_1}.$$

For each $i \in \mathbb{Z}/n\mathbb{Z}$, we have $W\langle i \rangle = q - |[i, i+q-1]_n \cap (\mathcal{S} - 1)|$. The size of

$$[i, i+q-1]_n \cap (\mathcal{S} - 1) = [i, i+q-1]_n \cap \{[\lfloor n/d \rfloor] - 1, [\lfloor 2n/d \rfloor] - 1, \dots, [\lfloor dn/d \rfloor] - 1\}$$

must be $\lfloor qd/n \rfloor$ or $\lceil qd/n \rceil$. Using the identity $q(n-d) = \gamma n - r$, we find that $W\langle i \rangle \in \{\gamma - 1, \gamma\}$ for all $i \in \mathbb{Z}/n\mathbb{Z}$. The total number of toggle operators in W is $q(n-d) = \gamma n - r$, so there are exactly r elements $i \in \mathbb{Z}/n\mathbb{Z}$ such that $W\langle i \rangle = \gamma - 1$. Furthermore, we have $W\langle i \rangle = \gamma - 1$ if and only if $i + q \in J$. This proves that $|J| = r$ and that $q - |[j-q, j-1]_n \cap (\mathcal{S} - 1)| = \gamma - 1$ for every $j \in J$.

It follows from Lemma 4.2 that we can choose the ordering i_1, \dots, i_{n-d} so that $\text{Bro}_J = \tau_{i_r} \cdots \tau_{i_1}$. For each $k \in \mathbb{Z}$, we have $\text{cyc}^{-k} \text{Bro}_{\mathcal{R}} \text{cyc}^k = \tau_{i_{n-d-k}} \cdots \tau_{i_{1-k}}$. Thus, when we view W as a permutation of Λ_{Path_n} , it is equal to

$$(\text{cyc}^{-(q-1)} \text{Bro}_{\mathcal{R}} \text{cyc}^{q-1}) \cdots (\text{cyc}^{-1} \text{Bro}_{\mathcal{R}} \text{cyc}) \text{Bro}_{\mathcal{R}} = \text{cyc}^{-q} (\text{cyc Bro}_{\mathcal{R}})^q.$$

When we view the word $W' = \tau_{i_r-q} \cdots \tau_{i_1-q}$ as a permutation, it is equal to $\text{cyc}^{-q} \text{Bro}_J \text{cyc}^q$, so $W'W = \text{cyc}^{-q} \text{Bro}_J (\text{cyc Bro}_{\mathcal{R}})^q$. Every toggle operator τ_i occurs exactly γ times in the word $W'W$. Our goal is to prove that the permutation $W'W$ of Λ_{Path_n} is equal to $\text{TPro}_{\beta}^{\gamma}$. Setting $Y = W'W$ in Lemma 2.5, we find that it suffices to show that if X is a suffix of $W'W$ and $a \rightarrow b$ is an arrow in β , then $X\langle a \rangle - X\langle b \rangle \in \{0, 1\}$.

Given $A \subseteq \mathbb{Z}/n\mathbb{Z}$ and $i \in \mathbb{Z}/n\mathbb{Z}$, let

$$A(i) = \begin{cases} 1 & \text{if } i \in A; \\ 0 & \text{if } i \notin A. \end{cases}$$

Let X be a suffix of $W'W$, and write $|X| = k(n-d) + m$ for some nonnegative integers k and m with $0 \leq m \leq n-d-1$. Then

$$X = \tau_{i_m-k} \cdots \tau_{i_1-k} \tau_{i_{n-d}-(k-1)} \cdots \tau_{i_1-(k-1)} \cdots \tau_{i_{n-d}-1} \cdots \tau_{i_1-1} \tau_{i_{n-d}} \cdots \tau_{i_1}.$$

Let $Q = \{i_1, \dots, i_m\} \subseteq \mathcal{R}$. Let $a \rightarrow b$ be an arrow in β ; we want to show that $X\langle a \rangle - X\langle b \rangle \in \{0, 1\}$. To do this, let us first prove that

$$(\mathcal{S} - 1)(\ell) + Q(\ell) - Q(\ell + 1) \in \{0, 1\} \text{ for all } \ell \in \mathbb{Z}/n\mathbb{Z}. \quad (4.1)$$

Because $(\mathcal{S} - 1) \cap Q = \emptyset$, we must have $(\mathcal{S} - 1)(\ell) + Q(\ell) - Q(\ell + 1) \leq 1$. Suppose by way of contradiction that $(\mathcal{S} - 1)(\ell) + Q(\ell) - Q(\ell + 1) < 0$. Then $(\mathcal{S} - 1)(\ell) = Q(\ell) = 0$ and $Q(\ell + 1) = 1$. This implies that $\ell + 1$ is not in the set \mathcal{S} of sources of β , so there is an arrow $\ell \rightarrow \ell + 1$ in β . Hence, ℓ appears before $\ell + 1$ in the ordering i_1, \dots, i_{n-d} . Since $\ell + 1 \in Q = \{i_1, \dots, i_m\}$, this forces $\ell \in Q$, which is a contradiction.

We can now prove that $X\langle a \rangle - X\langle b \rangle \in \{0, 1\}$; we consider two cases.

Case 1. Suppose $b = a + 1$. In this case, $X\langle a \rangle = k - |[a, a + k - 1]_n \cap (\mathcal{S} - 1)| + Q(a + k)$ and $X\langle b \rangle = k - |[a + 1, a + k]_n \cap (\mathcal{S} - 1)| + Q(a + k + 1)$, so

$$X\langle a \rangle - X\langle b \rangle = -(\mathcal{S} - 1)(a) + (\mathcal{S} - 1)(a + k) + Q(a + k) - Q(a + k + 1).$$

Because $a \rightarrow a + 1$ is an arrow in β , we know that $(\mathcal{S} - 1)(a) = 0$. If we set $\ell = a + k$ in (4.1), we find that $X\langle a \rangle - X\langle b \rangle = (\mathcal{S} - 1)(a + k) + Q(a + k) - Q(a + k + 1) \in \{0, 1\}$.

Case 2. Suppose $b = a - 1$. In this case, $X\langle a \rangle = k - |[a, a + k - 1]_n \cap (\mathcal{S} - 1)| + Q(a + k)$ and $X\langle b \rangle = k - |[a - 1, a + k - 2]_n \cap (\mathcal{S} - 1)| + Q(a + k - 1)$, so

$$X\langle a \rangle - X\langle b \rangle = (\mathcal{S} - 1)(a - 1) - (\mathcal{S} - 1)(a + k - 1) + Q(a + k) - Q(a + k - 1).$$

Because $a \rightarrow a - 1$ is an arrow in β , it follows from the definition of β that $(\mathcal{S} - 1)(a - 1) = 1$. If we set $\ell = a + k - 1$ in (4.1), we find that $(\mathcal{S} - 1)(a + k - 1) + Q(a + k - 1) - Q(a + k) \in \{0, 1\}$. Therefore, $X\langle a \rangle - X\langle b \rangle = 1 - ((\mathcal{S} - 1)(a + k - 1) + Q(a + k - 1) - Q(a + k)) \in \{0, 1\}$. \square

5. Divisibility of Permutoric Promotion Orbit Sizes

Our goal in this section is to prove the following proposition.

Proposition 5.1. *If β is an acyclic orientation of Cycle_n with d counterclockwise edges, then every orbit of TPro_β has size divisible by $\text{lcm}(d, n - d)$.*

Lemma 2.4 tells us that it suffices to prove Proposition 5.1 when $1 \leq d \leq \lfloor n/2 \rfloor$. Furthermore, if $d = 1$, then TPro_β is dynamically equivalent to the toric promotion operator TPro , so it follows from Theorem 1.1 (specialized to the case when $G = \text{Path}_n$) that all orbits of TPro_β have size $n - 1$. Thus, we may assume in what follows that $2 \leq d \leq \lfloor n/2 \rfloor$. By Lemma 2.4, we only need to prove Proposition 5.1 for one specific choice of an acyclic orientation β with d counterclockwise edges. As in Section 4, let $s_i = \lfloor \lfloor in/d \rfloor \rfloor$, and let \mathcal{S} be the independent set $\{s_1, \dots, s_d\}$ of Cycle_n . Let $\mathcal{R} = (\mathbb{Z}/n\mathbb{Z}) \setminus (\mathcal{S} - 1)$. Let $\beta = \beta_{\mathcal{S}}$ be the acyclic orientation of Cycle_n whose sources are the elements of \mathcal{S} and whose sinks are the elements of $\mathcal{S} - 1$. We will prove that every orbit of TPro_β has size divisible by $\text{lcm}(d, n - d)$.

Fix a labeling $\lambda \in \Lambda_{\text{Path}_n}$. Let γ be the size of the orbit of TPro_β containing λ . Using the division algorithm, we can write $\gamma n = q(n - d) + r$, where q and r are nonnegative integers and $0 \leq r \leq n - d - 1$. As in Proposition 4.1, let

$$J = \{j \in [n] : q - \gamma + 1 \leq |[j - q, j - 1]_n \cap (\mathcal{S} - 1)|\}.$$

Since Proposition 4.1 allows us to rewrite TPro_β^γ in terms of the operator $\text{cyc Bro}_{\mathcal{R}}$, we will want to consider the orbit of λ under $\text{cyc Bro}_{\mathcal{R}}$. Thus, we let

$$\mathcal{M} = \{(\text{cyc Bro}_{\mathcal{R}})^t(\lambda) : t \in \mathbb{Z}\}.$$

In Section 3.3, we described how to compute the action of $\text{cyc Bro}_{\mathcal{R}}$ on a labeling via a two-step procedure involving gliding globs. As mentioned in Remark 3.4, neither of the two steps in this procedure change the relative order in which the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ appear along the path. Thus, we have the following lemma.

Lemma 5.2. *For every $\mu \in \mathcal{M}$, the order in which the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ appear along the path in μ is the same as the order in which they appear along the path in λ .*

We are now in a position to prove that γ is divisible by $n - d$.

Lemma 5.3. *If $\lambda \in \Lambda_{\text{Path}_n}$ belongs to an orbit of TPro_β of size γ , then γ is divisible by $n - d$.*

Proof. Recall that we write $\gamma n = q(n - d) + r$ using the division algorithm. The map TPro_β^γ fixes λ . Lemma 3.5 tells us that TPro_β commutes with $\text{cyc Bro}_\mathcal{R}$, so TPro_β^γ acts as the identity on \mathcal{M} and thus trivially restricts to a bijection from \mathcal{M} to itself. Since

$$\text{TPro}_\beta^\gamma = \text{cyc}^{-q} \text{Bro}_J (\text{cyc Bro}_\mathcal{R})^q$$

by Proposition 4.1, the map $\text{cyc}^{-q} \text{Bro}_J$ also restricts to a bijection from \mathcal{M} to itself. Let u_1, \dots, u_{n-d} be the elements of $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$, listed in the order in which they appear from left to right along the path in λ . It follows from Lemma 4.2 that there exist integers $\dots, y_0, y_1, y_2, \dots$ satisfying $y_{i+d} = y_i + n$ and $s_i \leq y_i \leq s_{i+1} - 1$ for all i such that $J = \bigcup_{i=1}^d [s_i, y_i - 1]_n$ (viewing J as a subset of $\mathbb{Z}/n\mathbb{Z}$). For each $1 \leq i \leq d$, we have that $y_i \notin J$ and $y_i - 1 \in J$, so it follows from the definition of J that $|[y_i - q, y_i - 1] \cap (\mathcal{S} - 1)| < |[y_i - 1 - q, y_i - 2] \cap (\mathcal{S} - 1)|$. We deduce that $y_i - 1 - q \in \mathcal{S} - 1$ for all $1 \leq i \leq d$. Therefore,

$$\mathcal{S} = \{y_1 - q, \dots, y_d - q\}. \quad (5.1)$$

Let

$$\zeta = \begin{cases} 0 & \text{if } u_1 \in \bigcup_{\ell=1}^d [s_\ell, y_\ell]_n; \\ 1 & \text{otherwise.} \end{cases}$$

Note that, regardless of the value of ζ , the element $u_1 + \zeta - 1$ cannot be of the form y_i for any integer i . Therefore, it follows from (5.1) that

$$u_1 + \zeta - q - 1 \notin \mathcal{S}. \quad (5.2)$$

As mentioned above, $\text{cyc}^{-q} \text{Bro}_J$ restricts to a bijection from \mathcal{M} to itself; thus, it follows from Lemma 5.2 that the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ appear in the order u_1, \dots, u_{n-d} from left to right along the path in the labeling $\text{cyc}^{-q} \text{Bro}_J(\lambda)$.

Consider applying cyc Bro_J to λ using the three-step gliding-globs procedure described in Section 3. We immerse the labels s_1, \dots, s_d and then apply the jeu de taquin operators $\text{jdt}_{[s_i, y_i]_n}$, imagining that the label s_i carries its glob along with it as it glides. After this initial step, the label u_1 will be on some vertex z ; at this point in time, all of the vertices to the left of z have globs of liquid on them, while z does not. We claim that $\lambda(z) \in \bigcup_{\ell=1}^d [s_\ell, y_\ell]_n$ if and only if $u_1 \in \bigcup_{\ell=1}^d [s_\ell, y_\ell]_n$. This is obvious if $\lambda(z) = u_1$. On the other hand, if $\lambda(z) \neq u_1$, then it follows from the definition of the jeu de taquin operators that $\lambda(z)$ and u_1 must both be in $\bigcup_{\ell=1}^d [s_\ell, y_\ell]_n$. This proves the claim, which is equivalent to the statement that $\zeta = 1$ if and only if $z \in \lambda^{-1}((\mathbb{Z}/n\mathbb{Z}) \setminus \bigcup_{\ell=1}^d [s_\ell, y_\ell]_n)$. The second step in the gliding-globs procedure increases by 1 the label of each vertex in $\lambda^{-1}((\mathbb{Z}/n\mathbb{Z}) \setminus \bigcup_{\ell=1}^d [s_\ell, y_\ell]_n)$; therefore, the label of z is $u_1 + \zeta$ after the second step. Note that the second step does not move any of the globs of

liquid. The third step of the procedure changes the label in each glob of liquid to a label of the form $y_i + 1$. It follows that in the labeling $\text{cyc Bro}_J(\lambda)$, the labels of the vertices to the left of z are all of the form $y_i + 1$, and the label of z is $u_1 + \zeta$. This means that in the labeling $\text{cyc}^{-q} \text{Bro}_J(\lambda) = \text{cyc}^{-q-1}(\text{cyc Bro}_J(\lambda))$, the labels of the vertices to the left of z are all of the form $y_i - q$ (i.e., they are in \mathcal{S} by (5.1)), and the label of z is $u_1 + \zeta - q - 1$. Combining this with (5.2), we find that $u_1 + \zeta - q - 1$ is the label in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ that appears farthest to the left in the labeling $\text{cyc}^{-q} \text{Bro}_J(\lambda)$. As mentioned above, $\text{cyc}^{-q} \text{Bro}_J$ sends \mathcal{M} to itself, so it follows from Lemma 5.2 that the labels in $(\mathbb{Z}/n\mathbb{Z}) \setminus \mathcal{S}$ appear in the order u_1, \dots, u_{n-d} in $\text{cyc}^{-q} \text{Bro}_J(\lambda)$. Consequently, $u_1 + \zeta - q - 1 = u_1$. This proves that q is congruent to 0 or -1 modulo n .

We defined q and r so that $\gamma n = q(n - d) + r$ and $0 \leq r \leq n - d - 1$. This implies that $r \not\equiv -d \pmod{n}$. Reading the first equation modulo n yields $r \equiv qd \pmod{n}$, so $q \not\equiv -1 \pmod{n}$. Therefore, we must have $q \equiv 0 \pmod{n}$ and $r = 0$. Writing $q = mn$, we find that $\gamma = m(n - d)$, which completes the proof. \square

Finally, we can complete the proof of the main result of this section.

Proof of Proposition 5.1. As discussed at the beginning of this section, it suffices to prove that every orbit of TPro_β is divisible by $\text{lcm}(d, n - d)$, where $\beta = \beta_{\mathcal{S}}$ is the acyclic orientation of Cycle_n coming from the independent set \mathcal{S} defined above. As before, let $\lambda \in \Lambda_{\text{Path}_n}$, and let γ be the size of the orbit of TPro_β containing λ . Lemma 5.3 tells us that γ is divisible by $n - d$, so we just need to show that γ is also divisible by d . Using the division algorithm, we can write $\gamma n = q(n - d) + r$. Since $n - d$ divides γ , we find that $r = 0$ and that q is divisible by n . Thus, it follows from Proposition 4.1 that the set J is empty and that we can write $\text{TPro}_\beta^\gamma = \text{cyc}^{-q}(\text{cyc Bro}_{\mathcal{S}})^q = (\text{cyc Bro}_{\mathcal{S}})^q$.

Given a labeling $\sigma \in \Lambda_{\text{Path}_n}$, let $\psi(\sigma)$ be the sequence obtained by reading the labels in \mathcal{S} in the order in which they appear from left to right along the path in σ . Recall from Section 3.3 the two-step gliding-globs procedure for computing the action of $\text{cyc Bro}_{\mathcal{S}}$. In the first step of this procedure, none of the globs of liquid can glide through each other. In the second step, we simply cyclically permute the d labels in the globs of liquid. This shows that $\psi(\text{cyc Bro}_{\mathcal{S}}(\sigma))$ is obtained from $\psi(\sigma)$ by cyclically permuting the labels in \mathcal{S} in the cyclic order s_1, \dots, s_d . It follows that every orbit of $\text{cyc Bro}_{\mathcal{S}}$ has size divisible by d . Since $\lambda = \text{TPro}_\beta^\gamma(\lambda) = (\text{cyc Bro}_{\mathcal{S}})^q(\lambda)$, we find that d divides q . The equation $\gamma n = q(n - d)$ then forces $d(n - d)$ to divide γn . Since $\text{gcd}(n, d)$ divides $n - d$, this implies that d divides $\gamma(n / \text{gcd}(n, d))$. But d and $n / \text{gcd}(n, d)$ are coprime, so d divides γ . \square

6. Orbit Structure of Permutoric Promotion

Throughout this section, we continue to fix G to be the path graph Path_n . Our primary goal is to prove Theorem 1.3.

6.1. A Reformulation

One of the advantages of Lemma 2.4 is that it allows us to work with whichever acyclic orientation β is most convenient for our purposes. In Section 5, we chose to work with the acyclic orientation $\beta_{\mathcal{S}}$ whose sources were the elements of an independent set \mathcal{S} and whose sinks were the elements of $\mathcal{S} - 1$. However, in this section, we will fix β to be the acyclic orientation of Cycle_n whose unique source is d and whose unique sink is n .

The purpose of Section 5 was to prove Proposition 5.1, which tells us that the sizes of the orbits of TPro_{β} are all divisible by $\text{lcm}(d, n-d)$. The reason this is necessary is that it allows us to reduce the problem of determining the orbit sizes of TPro_{β} to the problem of determining the orbit sizes of TPro_{β}^d . The following proposition allows us to rewrite TPro_{β}^d in a more convenient form.

Proposition 6.1. *We have*

$$\text{TPro}_{\beta}^d = \prod_{i=n}^1 (\tau_i \tau_{i+1} \cdots \tau_{i+d-1}) = (\tau_n \tau_{n+1} \cdots \tau_{d+n-1}) \cdots (\tau_2 \tau_3 \cdots \tau_{d+1}) (\tau_1 \tau_2 \cdots \tau_d).$$

Proof. Think of $(\tau_n \tau_{n+1} \cdots \tau_{d+n-1}) \cdots (\tau_2 \tau_3 \cdots \tau_{d+1}) (\tau_1 \tau_2 \cdots \tau_d)$ as a word Y over the alphabet $\{\tau_1, \dots, \tau_n\}$. Note that every letter in this alphabet appears exactly d times in Y . By Lemma 2.5, we just need to show that if X is a suffix of Y and $a \rightarrow b$ is an arrow in β , then $X\langle a \rangle - X\langle b \rangle \in \{0, 1\}$; this is straightforward to check directly. \square

Remark 6.2. By combining Proposition 6.1 with the identity $\text{cyc}^{-1} \tau_{i+1} = \tau_i \text{cyc}^{-1}$ and the fact that cyc^n is the identity map, one can readily check that $\text{TPro}_{\beta}^d = \left(\text{cyc}^{-1} \text{Bro}_{\{1, \dots, d\}}^{-1} \right)^n$.

Define a map $\Phi_{n,d}: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ by

$$\Phi_{n,d} = \text{cyc}^d \prod_{i=n-d}^1 (\tau_i \tau_{i+1} \cdots \tau_{i+d-1}) = \text{cyc}^d (\tau_{n-d} \tau_{n-d+1} \cdots \tau_{n-1}) \cdots (\tau_2 \tau_3 \cdots \tau_{d+1}) (\tau_1 \tau_2 \cdots \tau_d).$$

Using the identity $\text{cyc} \tau_i = \tau_{i+1} \text{cyc}$ together with Proposition 6.1, one can check that

$$\Phi_{n,d}^{n/\text{gcd}(n,d)} = \text{TPro}_{\beta}^{\text{lcm}(d,n-d)}. \quad (6.1)$$

Lemma 6.3. *Every orbit of $\Phi_{n,d}: \Lambda_{\text{Path}_n} \rightarrow \Lambda_{\text{Path}_n}$ has size divisible by $n/\text{gcd}(n,d)$.*

Proof. Let $\text{FS}(\overline{\text{Path}}_n, \text{Cycle}_n)$ be the graph with vertex set Λ_{Path_n} in which two distinct labelings σ, σ' are adjacent if and only if there exists $i \in \mathbb{Z}/n\mathbb{Z}$ such that $\sigma' = \tau_i(\sigma)$. In the language of the article [DK21], this is the *friends-and-strangers graph* of $\overline{\text{Path}}_n$ and Cycle_n ,

where $\overline{\text{Path}}_n$ is the complement of Path_n . For $\sigma \in \Lambda_{\text{Path}_n}$, let H_σ be the connected component of $\text{FS}(\overline{\text{Path}}_n, \text{Cycle}_n)$ containing σ . It follows from Theorem 4.1 and Proposition 4.4 in [DK21] that there is a well-defined action of the group $\langle \text{cyc} \rangle \cong \mathbb{Z}/n\mathbb{Z}$ on the set of connected components of $\text{FS}(\overline{\text{Path}}_n, \text{Cycle}_n)$ given by $\text{cyc} \cdot H_\sigma = H_{\text{cyc}(\sigma)}$; moreover, these results from [DK21] imply that all orbits of this action have size n . Note that $H_{\Phi_{n,d}(\sigma)} = \text{cyc}^d \cdot H_\sigma$. If $\Phi_{n,d}^k(\sigma) = \sigma$, then $H_\sigma = H_{\Phi_{n,d}^k(\sigma)} = \text{cyc}^{dk} \cdot H_\sigma$, so k is divisible by $n/\text{gcd}(n, d)$. \square

Let $\text{Comp}_d(n)$ denote the set of compositions of n with d parts (i.e., d -tuples of positive integers that sum to n). There is a natural *rotation* operator $\text{Rot}_{n,d}: \text{Comp}_d(n) \rightarrow \text{Comp}_d(n)$ defined by $\text{Rot}_{n,d}(a_1, a_2, \dots, a_d) = (a_2, \dots, a_d, a_1)$. Our goal in the next subsection will be to relate $\Phi_{n,d}$ and $\text{Rot}_{n,d}$ via the following proposition. Recall that we write Orb_f for the set of orbits of an invertible map f .

Proposition 6.4. *There is a map $\Omega: \text{Orb}_{\Phi_{n,d}} \rightarrow \text{Orb}_{\text{Rot}_{n,d}}$ such that $|\Omega(\mathcal{O})| = \frac{d}{n}|\mathcal{O}|$ for every $\mathcal{O} \in \text{Orb}_{\Phi_{n,d}}$ and $|\Omega^{-1}(\widehat{\mathcal{O}})| = d!(n-d)!$ for every $\widehat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$.*

Before proceeding to the proof of Proposition 6.4, let us see why it implies Theorem 1.3.

Proof of Theorem 1.3 assuming Proposition 6.4. Let k_1, \dots, k_ℓ be the different sizes of the orbits of $\text{Rot}_{n,d}$, and let m_i be the number of orbits of $\text{Rot}_{n,d}$ of size k_i . Then $\{k_i^{m_i} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of $\text{Rot}_{n,d}$, where we use superscripts to denote multiplicities. If we assume Proposition 6.4, then we find that the multiset of orbit sizes of $\Phi_{n,d}$ is

$$\left\{ \left(\frac{n}{d} k_i \right)^{d!(n-d)!m_i} : 1 \leq i \leq \ell \right\}.$$

It then follows from (6.1) and Lemma 6.3 that the multiset of orbit sizes of $\text{TPro}_\beta^{\text{lcm}(d,n-d)}$ is

$$\left\{ \left(\frac{\text{gcd}(n, d)}{n} \frac{n}{d} k_i \right)^{(n/\text{gcd}(n,d))d!(n-d)!m_i} : 1 \leq i \leq \ell \right\},$$

and we can then invoke Proposition 5.1 to see that the multiset of orbit sizes of TPro_β is

$$\begin{aligned} & \left\{ \left(\text{lcm}(d, n-d) \frac{\text{gcd}(n, d)}{n} \frac{n}{d} k_i \right)^{(1/\text{lcm}(d,n-d))(n/\text{gcd}(n,d))d!(n-d)!m_i} : 1 \leq i \leq \ell \right\} \\ & = \left\{ ((n-d)k_i)^{n(d-1)!(n-d-1)!m_i} : 1 \leq i \leq \ell \right\}. \end{aligned}$$

Since $\text{Rot}_{n,d}$ has order d , this implies that TPro_β has order $d(n-d)$. It is well known [RSW04] that the triple

$$\left(\text{Comp}_d(n), \text{Rot}_{n,d}, \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q \right)$$

exhibits the cyclic sieving phenomenon. Hence, Lemma 2.1 (with $f = \text{Rot}_{n,d}$ and $g = \text{TPro}_\beta$) implies that

$$\left(\Lambda_{\text{Path}_n}, \text{TPro}_\beta, n(d-1)!(n-d-1)![n-d]_{q^d} \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q \right)$$

exhibits the cyclic sieving phenomenon. \square

6.2. Sliding Stones and Colliding Coins

Our aim is now to prove Proposition 6.4, which, as we have just seen, implies our main theorem about the orbit structure of permutoric promotion. Code implementing several of the combinatorial constructions described in this section can be found at <https://cocalc.com/hrthomas/permutoric-promotion/implementation>.

For each integer k , let $\theta_k = \tau_{q+d+1-r}$, where q and r are the unique integers satisfying $k = qd + r$ and $1 \leq r \leq d$. Let

$$\nu_\ell = \theta_{d\ell} \theta_{d\ell-1} \cdots \theta_{d(\ell-1)+2} \theta_{d(\ell-1)+1}.$$

Observe that $\theta_{k+dn} = \theta_k$ for all integers k . We have

$$\Phi_{n,d} = \text{cyc}^d \theta_{d(n-d)} \cdots \theta_2 \theta_1 = \text{cyc}^d \nu_{n-d} \cdots \nu_2 \nu_1.$$

By combining the identity $\text{cyc} \tau_i = \tau_{i+1} \text{cyc}$ with the fact that cyc^n is the identity map, one can easily verify that

$$\Phi_{n,d}^m = \theta_{md(n-d)} \cdots \theta_2 \theta_1 = \nu_{m(n-d)} \cdots \nu_2 \nu_1 \quad (6.2)$$

whenever m is a positive multiple of $n/\text{gcd}(n, d)$.

Define a *state* to be a pair $(\sigma, t) \in \Lambda_{\text{Path}_n} \times \mathbb{Z}$; we call σ the *labeling* of the state and say that the state is at *time* t . A *timeline* is a bi-infinite sequence $\mathcal{T} = (\sigma_t, t)_{t \in \mathbb{Z}}$ of states such that $\sigma_t = \nu_t(\sigma_{t-1})$ for all $t \in \mathbb{Z}$. Note that every state belongs to a unique timeline. For $\sigma \in \Lambda_{\text{Path}_n}$, let \mathcal{T}_σ be the unique timeline containing the state $(\sigma, 0)$.

Let v_1, \dots, v_n be the vertices of Path_n , listed from left to right. For each $\ell \in [n]$, let \mathbf{v}_ℓ be a formal symbol associated to v_ℓ ; we will call \mathbf{v}_ℓ a *replica*. Let s_1, \dots, s_d be stones of different colors. We define the *stones diagram* of a state (σ, t) as follows. Start with a copy of Cycle_n . Place s_1, \dots, s_d on the vertices $t+d, \dots, t+1$, respectively. Then place each replica \mathbf{v}_ℓ on the vertex $\sigma(v_\ell)$ of Cycle_n ; if this vertex already has a stone sitting on it, then we place the replica on top of the stone.

Suppose we have a timeline $\mathcal{T} = (\sigma_t, t)_{t \in \mathbb{Z}}$. We want to describe how the stones diagrams of the states evolve as we move through the timeline. We will imagine transforming the stones diagram of $(\sigma_{t-1}, t-1)$ into that of (σ_t, t) via a sequence of d *small steps*. The i -th small step moves s_i one space clockwise. The labeling $(\theta_{d(t-1)+i} \cdots \theta_{d(t-1)+1})(\sigma_{t-1})$ is obtained from $(\theta_{d(t-1)+i-1} \cdots \theta_{d(t-1)+1})(\sigma_{t-1})$ by applying the toggle operator $\theta_{d(t-1)+i} = \tau_{t+d-i}$. If this operator has no effect (i.e., $(\theta_{d(t-1)+i} \cdots \theta_{d(t-1)+1})(\sigma_{t-1}) = (\theta_{d(t-1)+i-1} \cdots \theta_{d(t-1)+1})(\sigma_{t-1})$), then we do not move any of the replicas $\mathbf{v}_1, \dots, \mathbf{v}_n$ during the i -th small step (in this case, the stone s_i slides from underneath one replica to underneath a different replica). Otherwise, $\theta_{d(t-1)+i}$ has the effect of swapping the labels $t+d-i$ and $t+d-i+1$, so we swap the replicas that were sitting on the vertices $t+d-i$ and $t+d-i+1$ (in this case, the stone s_i carries the replica sitting on it along with it as it slides). Figure 6.1 illustrates these small steps for a particular example with $n = 8$, $d = 3$, and $t = 1$.

Now consider d coins of different colors such that the set of colors of the coins is the same as the set of colors of the stones. We define the *coins diagram* of a state (σ, t) as follows. Start with a copy of Path_n . For each $i \in [d]$, there is a replica \mathbf{v}_ℓ sitting on the stone s_i in




Figure 6.1: The $d = 3$ small steps transforming the stones diagram of a state at time 0 into the stones diagram of the next state at time 1.

the stones diagram of (σ, t) ; place the coin with the same color as the stone s_i on the vertex v_ℓ (see Figures 6.2 and 6.3). Note that the set of vertices of Path_n occupied by coins is $\{\sigma^{-1}(t + 1), \dots, \sigma^{-1}(t + d)\}$. (While states are uniquely determined by their stones diagrams, the coins diagrams will be crucial for analyzing the dynamics.)

Consider how the coins diagrams evolve as we move through a timeline $\mathcal{T} = (\sigma_t, t)_{t \in \mathbb{Z}}$. Let us transform the stones diagram of $(\sigma_{t-1}, t-1)$ into that of (σ_t, t) via the d small steps described above. Let v_ℓ be the replica sitting on s_i right before the i -th small step, and let $v_{\ell'}$ be the replica sitting on the vertex one step clockwise from s_i right before the i -th small step. When s_i moves in the i -th small step, it will either carry its replica v_ℓ along with it or slide from underneath v_ℓ to underneath $v_{\ell'}$; the latter occurs if and only if $\ell' = \ell \pm 1$. In the former case, no coins move during the i -th small step; in the latter case, a coin moves from v_ℓ to the adjacent vertex $v_{\ell'}$ (which did not have a coin on it right before this small step).

If we watch the coins diagrams evolve as we move through the timeline, then by the previous paragraph, the coins will move around on Path_n , but they will never move through each other. Therefore, it makes sense to name the coins c_1, \dots, c_d in the order they appear along the path from left to right, and this naming only depends on the timeline (not the specific state in the timeline). Define a *traffic jam* to be a maximal nonempty collection of coins that occupy a contiguous block of vertices (so the vertices occupied by the coins in a particular traffic jam induce a connected subgraph of Path_n). Note that a traffic jam could have just a single coin. We say a traffic jam *touches a wall* if it contains a coin that occupies v_1 or v_n .

At any time, a coin has an idea of the direction in which it expects to move next (our coins are conscious now). Note that this is not necessarily the direction in which it will move next because it may change its mind before it moves. The way that a coin c decides which direction it expects to move is as follows. Suppose c currently occupies vertex v_j , and suppose the coins in the traffic jam containing c occupy the vertices v_r, v_{r+1}, \dots, v_s . The coin c looks at the stones diagram and reads ahead in the clockwise direction, starting from the stone of its color, and it determines whether it first sees \mathbf{v}_{r-1} or \mathbf{v}_{s+1} . If it first sees \mathbf{v}_{r-1} , it expects to move left; if it first sees \mathbf{v}_{s+1} , it expects to move right. If $r - 1$ is not the index of a replica (because $r = 1$), the first replica that c sees will be \mathbf{v}_{s+1} ; similarly, if $s + 1$ is not the index of a replica (because $s = n$), the first replica c sees will be \mathbf{v}_{r-1} .

Figure 6.2 shows several stones diagrams and coins diagrams. In each coins diagram, an arrow has been placed over each coin to indicate which direction it expects to move.

Lemma 6.5. *When a coin moves, it moves in the direction that it expects to move.*

Proof. Suppose c occupies v_j and is about to move left. The stone of the same color as c is under the replica \mathbf{v}_j , and the next replica clockwise is \mathbf{v}_{j-1} , which has no stone under it. It follows that c is the leftmost coin in a traffic jam. When c reads through the stones diagram looking for one of two replicas, it first sees \mathbf{v}_{j-1} , so it indeed expects to move left. The analysis for coins moving to the right is the same. \square

The following lemma tells us under what circumstances a coin can change its mind about which way it is going to move.

Lemma 6.6. *Let c be a coin, and let s be the stone with the same color as c . Consider a small step, and let v_j be the vertex occupied by c right before the small step. Let v_r, v_{r+1}, \dots, v_s be the vertices occupied by the coins in the traffic jam that contains c right before the small step. During this small step, c changes its mind about which direction it expects to move if and only if one of the following situations occurs:*

- *The stone s slides through \mathbf{v}_{r-1} or \mathbf{v}_{s+1} and carries \mathbf{v}_j along with it as it slides (so c does not move in the coins diagram), and the traffic jam containing c does not touch a wall (so $1 < r \leq s < n$).*
- *The coin c moves, and the traffic jam that contains c after the small step touches a wall.*

Proof. First of all, note that c will not change its mind about which way it is going to move except during a small step when s moves. Indeed, even though the traffic jam containing c may change during other small steps, it is straightforward to check that these small steps will not change the direction that c expects to move.

While c is in a traffic jam that touches a wall, there is only one way that it can expect to move: away from that wall. Thus, it does not change its mind about which way it is moving before it actually moves, but it does change its mind the moment it arrives in the traffic jam (i.e., when the second bulleted item in the statement of the lemma is satisfied).

Now consider a small step during which c moves, and suppose the traffic jam that contains c after the small step does not touch a wall. For simplicity, let us assume that c expects to move

left before this small step. Then during the small step, c does in fact move left (by Lemma 6.5). Let us say c moves from v_j to v_{j-1} . Then during this small step, s slides from underneath the replica v_j to underneath the replica v_{j-1} . After the small step, the vertices occupied by the coins in the traffic jam containing c are $v_r, v_{r+1}, \dots, v_{j-1}$ for some $r \in \{2, \dots, j-1\}$, so c looks in the stones diagram for either v_{r-1} or v_j . It will certainly see v_{r-1} first since v_j is one step behind s (in the clockwise order) at this time. Thus, c still expects to move left after the small step.

Finally, consider the situation from the first bulleted item in the statement of the lemma. Let us again assume for simplicity that c expects to move left before the small step. Then s slides through v_{r-1} . After the small step, when c reads through the stones diagram to determine which direction it expects to move, it again searches for v_{r-1} and v_{s+1} (because no coins moved during the small step). It will see v_{s+1} before v_{r-1} because v_{r-1} is now right behind s in the clockwise order. So c expects to move right after the small step. \square

The importance of understanding the direction in which a coin expects to move is that it will enable us to understand *collisions*. There are *two-coins collisions*, which involve two coins that occupy adjacent vertices of Path_n ; there are *left-wall collisions*, which can occur when c_1 occupies v_1 ; and there are *right-wall collisions*, which can occur when c_d occupies v_n . The prototypical examples of collisions are when two non-adjacent coins move to become adjacent or when a coin moves to become adjacent to a wall, but other examples are possible when traffic jams of size greater than 1 are involved.

The precise definition of a two-coins collision that occurs in a traffic jam that does not touch a wall is as follows. We say coins c_i and c_{i+1} are *butting heads* if they occupy adjacent vertices and c_i expects to move right while c_{i+1} expects to move left. We say c_i and c_{i+1} are involved in a two-coins collision at a small step if they are not butting heads immediately before the small step and they are butting heads immediately after the small step. This can happen either because the two coins were not adjacent prior to the small step, but it can also happen because the two coins were adjacent but one of them changed its mind about the direction it expected to move.

The definition has to be slightly modified in a traffic jam that touches a wall. Consider first the case when a small step occurs during which a coin c moves so as to join a traffic jam that touches the wall. At the same time, c changes its mind so that it now expects to move away from the wall that the traffic jam touches. Nonetheless, if there is a coin c' adjacent to c after the small step, we still count this as a two-coins collision between c and c' . (We can imagine that there was a brief instant of time right after c moved to join the traffic jam but right before it changed its mind about which way it expected to move, thus resulting in c butting heads with c' very briefly.) Similarly, if c moved onto v_1 (respectively, v_n) during this small step (so it is in a traffic jam of size 1 that touches a wall), then we count this as a left-wall (respectively, right-wall) collision.

We now discuss how to define a collision that occurs in the “interior” of a traffic jam of size at least 2 that touches a wall. In such a traffic jam, all the coins always want to move away from the wall, so by the above definition, there would be no collisions within the traffic jam. However, this is not what we want. Instead, suppose we are considering a coin c_i that occupies v_j . Assume the coins in the traffic jam containing c_i occupy vertices v_1, v_2, \dots, v_k , where $j < k$. Thus, we are assuming the traffic jam touches the left wall, but the symmetrical considerations apply if the traffic jam touches the right wall. The stone with the same color as c_i carries the replica v_j . Suppose there is a small step during which the stone with the same color as c_i slides through v_{k+1} ,

carrying v_j along with it as it slides. Note that c_i does not move during this small step. In this case, we say c_i collides with c_{i-1} (or is involved in a left-wall collision if $i = 1$). To explain heuristically why this collision occurs, we can imagine that c_i has a “flicker of confusion” when it sees the stone with its same color slide through v_{k+1} . When it sees this, c_i “thinks” it should change its mind and expect to move left. But then it realizes that it cannot expect to move left because it is in a traffic jam that touches the left wall, so it quickly goes back to expecting to move right. During this brief instant, the collision occurs because c_i “thinks” it should be butting heads with c_{i-1} (or with the left wall if $i = 1$).

We say a collision occurs at time t if it occurs during a small step between times $t - 1$ and t .

Example 6.7. Suppose $n = 6$ and $d = 3$. Figure 6.2 shows some stones diagrams and coins diagrams evolving over time. At each stage, the arrow over a coin points in the direction that the coin expects to move. Collisions are indicated in the coins diagrams by stars, and each star is colored to indicate which stone moves in the small step during which the collision occurs. Note that the right-wall collision at time 5 (marked with a gold star in the first small step after time 4) occurs because c_3 has a “flicker of confusion” when the gold stone s_1 slides through v_4 (carrying v_6 along with it as it slides).

Example 6.8. Suppose $n = 6$ and $d = 3$. Figure 6.3 shows the stones diagrams and coins diagrams of a particular timeline at times $0, 1, \dots, 17$. For brevity, we have not shown the individual small steps. All of the collisions that occur at time t (i.e., during the small steps between time $t - 1$ and time t) are indicated in the coins diagram at time t . The color of the star can be used to determine the small step during which the collision occurs. One can check that the states in this timeline are periodic with period 18.

Let $\text{Coll}_{\mathcal{T}}$ be the set of all collisions that take place in the coins diagrams of the states of the timeline \mathcal{T} . We define a directed graph with vertex set $\text{Coll}_{\mathcal{T}}$ by drawing an arrow from a collision κ to a collision κ' whenever there is a coin involved in both κ and κ' and the collision κ occurs before κ' . Let $(\text{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$ be the transitive closure of this directed graph. Let $\mathbf{H}_{\mathcal{T}}$ be the Hasse diagram of $(\text{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$. This Hasse diagram, which will be one of our primary tools, has the shape of a bi-infinite chain link fence (see Figure 6.4). Suppose $\kappa_1 \leq_{\mathcal{T}} \kappa_2$ is an edge in $\mathbf{H}_{\mathcal{T}}$. Then κ_1 and κ_2 are collisions that both use some coin c ; we define the *energy* of this edge, denoted $\mathcal{E}(\kappa_1 \leq_{\mathcal{T}} \kappa_2)$, to be the number of different vertices that c occupies between these two collisions, including the vertices occupied by c when the collisions occur. More generally, if $\kappa_1 \leq_{\mathcal{T}} \kappa_2 \leq_{\mathcal{T}} \dots \leq_{\mathcal{T}} \kappa_r$ is a saturated chain in $\mathbf{H}_{\mathcal{T}}$, then we write $\mathcal{E}(\kappa_1 \leq_{\mathcal{T}} \kappa_2 \leq_{\mathcal{T}} \dots \leq_{\mathcal{T}} \kappa_r)$ for the tuple $(\mathcal{E}(\kappa_1 \leq_{\mathcal{T}} \kappa_2), \dots, \mathcal{E}(\kappa_{r-1} \leq_{\mathcal{T}} \kappa_r))$ of energies of the edges in the chain.

Example 6.9. If \mathcal{T} is the timeline containing the states whose stones diagrams and coins diagrams are shown in Figure 6.3, then (a finite part of) $\mathbf{H}_{\mathcal{T}}$ is shown in Figure 6.4. Each collision is represented by a color-coded star, and the blue number inside the star is the time when the collision occurs. Each edge is labeled by its energy.

A *diamond* in $\mathbf{H}_{\mathcal{T}}$ consists of collisions $\kappa_1, \kappa_2, \kappa_3, \kappa_4$ together with four edges given by cover relations $\kappa_1 \leq_{\mathcal{T}} \kappa_2, \kappa_1 \leq_{\mathcal{T}} \kappa_3, \kappa_2 \leq_{\mathcal{T}} \kappa_4, \kappa_3 \leq_{\mathcal{T}} \kappa_4$. A *half-diamond* in $\mathbf{H}_{\mathcal{T}}$ consists of collisions $\kappa'_1, \kappa'_2, \kappa'_3$, where κ'_1 and κ'_3 are either both left-wall collisions or both right-wall collisions,




Figure 6.2: The evolution of stones diagrams and coins diagrams over time, with each individual small step illustrated. At each moment, we have drawn an arrow over each coin to indicate which direction it expects to move. Each collision is indicated by a star whose color is the same as that of the stone that moved to cause the collision. Each labeling is depicted in red numbers below the path.




Figure 6.3: The stones diagrams and coins diagrams of the states in a timeline at times $0, 1, \dots, 17$. Here, $n = 6$ and $d = 3$. The collisions that occur during the small steps between times $t - 1$ and t are represented by color-coded stars in the coins diagram at time t . Each labeling is depicted by the red numbers below the path.

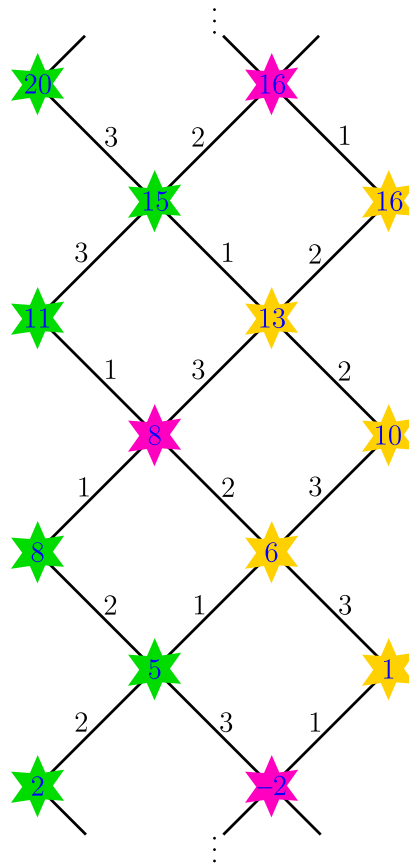


Figure 6.4: A Hasse diagram H_T . Each collision is represented by a star whose color is the same as that of the stone that moved to cause the collision. Blue numbers indicate the times when the collisions occur. Edges are labeled by their energies.

together with two edges given by cover relations $\kappa'_1 \leq_T \kappa'_2$ and $\kappa'_2 \leq_T \kappa'_3$. Our arguments in the next subsection rest on the following three lemmas.

Lemma 6.10. *In any half-diamond in the Hasse diagram H_T , the two edges have the same energy.*

Proof. Fix a half-diamond in H_T with edges given by the cover relations $\kappa'_1 \leq_T \kappa'_2 \leq_T \kappa'_3$. By symmetry, we may assume κ'_1 and κ'_3 are both left-wall collisions. Then κ'_1 and κ'_3 occur when c_1 occupies v_1 . Say κ'_2 occurs when c_1 occupies v_m . Both edges of the half-diamond have energy m . □

Lemma 6.11. *In any diamond in the Hasse diagram H_T , opposite edges have the same energy.*

Proof. Fix a diamond in H_T with edges given by cover relations $\kappa_1 \leq_T \kappa_2, \kappa_1 \leq_T \kappa_3, \kappa_2 \leq_T \kappa_4, \kappa_3 \leq_T \kappa_4$. Say the collision κ_1 involves coins c_i and c_{i+1} and takes place when c_i occupies vertex v_m and c_{i+1} occupies vertex v_{m+1} . Without loss of generality, suppose κ_2 involves c_i

and κ_3 involves c_{i+1} . Then κ_2 occurs when c_i occupies some vertex v_k with $k \leq m$, and κ_3 occurs when c_{i+1} occupies some vertex v_ℓ with $\ell \geq m+1$. The collision κ_4 involves the coins c_i and c_{i+1} . We claim that κ_4 occurs when c_i occupies $v_{k+\ell-m-1}$ and c_{i+1} occupies $v_{k+\ell-m}$; this will imply that the edges $\kappa_1 \leq_{\mathcal{T}} \kappa_2$ and $\kappa_3 \leq_{\mathcal{T}} \kappa_4$ both have energy $m - k + 1$ and that the edges $\kappa_1 \leq_{\mathcal{T}} \kappa_3$ and $\kappa_2 \leq_{\mathcal{T}} \kappa_4$ both have energy $\ell - m$.

By symmetry, we may assume that $m - k \leq \ell - m - 1$. Let $x, y \in [d]$ be such that s_x and s_y are the stones with the same colors as c_i and c_{i+1} , respectively. Consider starting at the time when κ_1 occurs and watching the stones diagrams and coins diagrams evolve as we move forward in time. We will assume that $k < m$, that $x > y$, and that c_i moves from v_m to v_{m-1} before c_{i+1} moves from v_{m+1} to v_{m+2} ; the other cases are similar. Let t_0 be the first time after the collision κ_1 when c_i moves from v_m to v_{m-1} . The coin c_i will move from v_m to v_{m-1} and then to v_{m-2} and so on until reaching v_k ; it will then turn around and move back across v_{k+1}, \dots, v_m and then continue on toward $v_{k+\ell-m-1}$. For $j \in \{k, \dots, k + \ell - m - 1\}$, let ζ_j be the amount of time that c_i spends on v_j during this trip. The coin c_{i+1} stays on v_{m+1} for some time after t_0 ; it then moves to the right until reaching v_ℓ , where it turns around and heads back to $v_{k+\ell-m+1}$. For $j' \in \{m+1, \dots, \ell\}$, let $\xi_{j'}$ be the amount of time after t_0 that c_{i+1} spends on $v_{j'}$ during this trip.

By analyzing the stones diagrams, one can show that $\zeta_j = \xi_{j'} = n - d$ for all $j \in \{k, \dots, m-1\}$ and all $j' \in \{k + \ell - m + 1, \dots, \ell\}$. Similarly, $\zeta_j = \xi_{j+1}$ for all $j \in \{m, \dots, k + \ell - m - 1\}$. Let $N = (m - k)(n - d) + \sum_{j=m}^{k+\ell-m-1} \zeta_j$. At time $t_0 + N$, either the coin c_i moves from $v_{k+\ell-m-1}$ to $v_{k+\ell-m}$, or the coin c_{i+1} moves from $v_{k+\ell-m+1}$ to $v_{k+\ell-m}$. It follows from the assumption that $x > y$ that, in fact, c_{i+1} moves from $v_{k+\ell-m+1}$ to $v_{k+\ell-m}$ at time $t_0 + N$. This proves the claim. \square

Example 6.12. Suppose $n = 6$ and $d = 3$, and let \mathcal{T} be the timeline from Examples 6.8 and 6.9. Let $\kappa_1, \kappa_3, \kappa_4$ be the collisions that occur at times 6, 10, 13, respectively, and let κ_2 be the two-coins collision at time 8. In the notation of the proof of Lemma 6.11, we have $i = 2$, $m = 3$, $k = 2$, $\ell = 6$, and $t_0 = 8$. We have $\zeta_2 = \xi_6 = 3 = n - d$, $\zeta_3 = \xi_4 = 1$, and $\zeta_4 = \xi_5 = 1$. Thus, $N = (m - k)(n - d) + \sum_{j=m}^{k+\ell-m-1} \zeta_j = 5$. As explained in the proof of Lemma 6.11, the coin c_3 moves from v_6 to v_5 at time $t_0 + N = 13$.

Lemma 6.13. *If the two edges of a half-diamond in the Hasse diagram $\mathbf{H}_{\mathcal{T}}$ have energy m , then the amount of time between the collisions at the bottom and the top of the half-diamond is $m(n - d)$.*

Proof. Without loss of generality, assume the bottom and top collisions in the half-diamond are left-wall collisions.

If $m > 1$, the same style of argument as in the proof of Lemma 6.11 proves that the total amount of time that c_1 spends on each of the vertices v_1, \dots, v_m is exactly $n - d$, which proves the claim.

If $m = 1$, a similar argument also applies. Let s be the stone with the same color as c_1 . A left-wall collision can only occur during a small step in which s moves. Moreover, such a small step results in a left-wall collision if and only if, right before the small step occurs, the replica one space clockwise of s is v_ℓ , where ℓ is the smallest of all the indices of replicas that do not

sit on stones at that time (equivalently, the traffic jam containing c_1 has size $\ell - 1$). The relative cyclic order of the indices of the replicas that do not sit on stones remains constant over time, so the time between these left-wall collisions is exactly $n - d$. \square

6.3. The Map Ω

Equipped with Lemmas 6.10, 6.11 and 6.13, we now turn to constructing and analyzing the map Ω from Proposition 6.4.

For each collision $\kappa \in \text{Coll}_{\mathcal{T}}$, let $\varphi(\kappa)$ be the collision involving the same set of coins as κ that occurs next after κ . In other words, if κ is the bottom element of a diamond (respectively, half-diamond), then $\varphi(\kappa)$ is the top element of that same diamond (respectively, half-diamond). We extend this notation to saturated chains in $\mathbf{H}_{\mathcal{T}}$ (including edges) by letting

$$\varphi(\kappa_1 \leq_{\mathcal{T}} \kappa_2 \leq_{\mathcal{T}} \cdots \leq_{\mathcal{T}} \kappa_m) = \varphi(\kappa_1) \leq_{\mathcal{T}} \varphi(\kappa_2) \leq_{\mathcal{T}} \cdots \leq_{\mathcal{T}} \varphi(\kappa_m).$$

We define the *period* of $\mathbf{H}_{\mathcal{T}}$ to be the smallest positive integer p such that e and $\varphi^p(e)$ have the same energy for every edge e of $\mathbf{H}_{\mathcal{T}}$. A *transversal* of $\mathbf{H}_{\mathcal{T}}$ is a saturated chain $\mathcal{T} = (\kappa_0 \leq_{\mathcal{T}} \kappa_1 \leq_{\mathcal{T}} \cdots \leq_{\mathcal{T}} \kappa_d)$ such that κ_0 is a left-wall collision, κ_d is a right-wall collision, and κ_i involves the stones c_i and c_{i+1} for every $i \in [d - 1]$. In other words, a transversal is a saturated chain that moves from left to right across $\mathbf{H}_{\mathcal{T}}$. We define the *energy composition* of \mathcal{T} to be the tuple $\mathcal{E}(\mathcal{T}) = (\varepsilon_1, \dots, \varepsilon_d)$, where ε_i is the energy of the edge $\kappa_{i-1} \leq_{\mathcal{T}} \kappa_i$; note that $\mathcal{E}(\mathcal{T}) \in \text{Comp}_d(n)$.

Lemma 6.14. *Let \mathcal{T} be a timeline, and let \mathcal{T} be a transversal of $\mathbf{H}_{\mathcal{T}}$. Then*

$$\mathcal{E}(\varphi(\mathcal{T})) = \text{Rot}_{n,d}(\mathcal{E}(\mathcal{T})).$$

Moreover, the period of $\mathbf{H}_{\mathcal{T}}$ is equal to the size of the orbit of $\text{Rot}_{n,d}$ containing $\mathcal{E}(\mathcal{T})$.

Proof. The second statement follows from the first because, by Lemmas 6.10 and 6.11, the energies of all edges in $\mathbf{H}_{\mathcal{T}}$ are determined by the energy composition of a single transversal of $\mathbf{H}_{\mathcal{T}}$. The first statement is also immediate from Lemmas 6.10 and 6.11. \square

Example 6.15. Suppose $n = 6$ and $d = 3$. Let $\mathbf{H}_{\mathcal{T}}$ be the Hasse diagram from Figure 6.4, and let $\mathcal{T} = (\kappa_0 \leq_{\mathcal{T}} \kappa_1 \leq_{\mathcal{T}} \kappa_2 \leq_{\mathcal{T}} \kappa_3)$ be the transversal consisting of the collisions that occur at times 2, 5, 6, 10. Then $\mathcal{E}(\mathcal{T}) = (2, 1, 3) \in \text{Comp}_3(6)$. The period of $\mathbf{H}_{\mathcal{T}}$ is 3, which is the size of the $\text{Rot}_{6,3}$ -orbit containing $(2, 1, 3)$. The transversal $\varphi(\mathcal{T})$ consists of both the collisions that occur at time 8 along with the collisions at times 13 and 16. We have $\mathcal{E}(\varphi(\mathcal{T})) = (1, 3, 2) = \text{Rot}_{6,3}(\mathcal{E}(\mathcal{T}))$. Similarly, $\mathcal{E}(\varphi^2(\mathcal{T})) = (3, 2, 1) = \text{Rot}_{6,3}^2(\mathcal{E}(\mathcal{T}))$.

Let S_r be the symmetric group consisting of all permutations of $[r]$. Suppose v_{i_1}, \dots, v_{i_r} is a sequence of distinct vertices of Path_n . We define the *standardization* of this sequence to be the unique permutation in S_r that has the same relative order as i_1, \dots, i_r when written in one-line notation. For example, the standardization of v_3, v_5, v_1, v_6 is 2314. Let $\mathcal{T} = (\sigma_t, t)_{t \in \mathbb{Z}}$ be a timeline. Recall that the stones s_1, \dots, s_d sit on the vertices $t+d, \dots, t+1$, respectively, in the stones diagram of (σ_t, t) . Let $\text{stand}_t(\mathcal{T})$ be the standardization of the

sequence $\sigma_t^{-1}(t+d), \dots, \sigma_t^{-1}(t+1)$. Alternatively, $\text{stand}_t(\mathcal{T})$ is the permutation $\rho: [d] \rightarrow [d]$ such that the stone s_i and the coin $c_{\rho(i)}$ have the same color for every $i \in [d]$. Let us also define $\overline{\text{stand}}_t(\mathcal{T})$ to be the standardization of $\sigma_t^{-1}(1), \dots, \sigma_t^{-1}(t), \sigma_t^{-1}(t+d+1), \dots, \sigma_t^{-1}(n)$ (i.e., the standardization of the sequence obtained from $\sigma_t^{-1}(1), \sigma_t^{-1}(2), \dots, \sigma_t^{-1}(n)$ by deleting $\sigma_t^{-1}(i)$ for all $t+1 \leq i \leq t+d$). It follows from the analysis of how stones diagrams evolve through a timeline that $\text{stand}_t(\mathcal{T}) = \text{stand}_{t+1}(\mathcal{T})$ and $\overline{\text{stand}}_t(\mathcal{T}) = \overline{\text{stand}}_{t+1}(\mathcal{T})$. In other words, $\text{stand}_t(\mathcal{T})$ and $\overline{\text{stand}}_t(\mathcal{T})$ only depend on the timeline \mathcal{T} and not on the time t . Thus, it makes sense to drop the subscripts and just write $\text{stand}(\mathcal{T})$ and $\overline{\text{stand}}(\mathcal{T})$. Note that there are $d!$ possibilities for $\text{stand}(\mathcal{T})$ and $(n-d)!$ possibilities for $\overline{\text{stand}}(\mathcal{T})$; this will end up being responsible for the appearance of $d!(n-d)!$ in Proposition 6.4.

For $k, t \in \mathbb{Z}$, let $\sigma_t^{(k)} = \text{cyc}^{-k}(\sigma_{t+k})$. It follows immediately from the definition of a timeline that the sequence $\mathcal{T}^{(k)} = (\sigma_t^{(k)}, t)_{t \in \mathbb{Z}}$ is also a timeline; that is, $\nu_t(\sigma_{t-1}^{(k)}) = \sigma_t^{(k)}$ for all $t \in \mathbb{Z}$. Furthermore, the stones diagram of $(\sigma_t^{(k)}, t)$ is obtained from that of $(\sigma_{t+k}, t+k)$ by moving all stones and replicas k positions counterclockwise. It follows that the coins diagrams of $(\sigma_t^{(k)}, t)$ and $(\sigma_{t+k}, t+k)$ are identical. Therefore, if κ is a collision in $\text{Coll}_{\mathcal{T}^{(k)}}$ that occurs at time t , then there is a collision $\psi_k(\kappa) \in \text{Coll}_{\mathcal{T}}$ that occurs at time $t+k$. The resulting map $\psi_k: \text{Coll}_{\mathcal{T}^{(k)}} \rightarrow \text{Coll}_{\mathcal{T}}$ is an isomorphism from $(\text{Coll}_{\mathcal{T}^{(k)}}, \leq_{\mathcal{T}^{(k)}})$ to $(\text{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$; furthermore, under this isomorphism, corresponding edges of the Hasse diagrams $\mathbf{H}_{\mathcal{T}^{(k)}}$ and $\mathbf{H}_{\mathcal{T}}$ have the same energy.

Recall that we write \mathcal{T}_σ for the unique timeline containing the state $(\sigma, 0)$. It follows from Lemma 6.14 that the energy compositions of the transversals of $\mathbf{H}_{\mathcal{T}_\sigma}$ form a single orbit $\tilde{\Omega}(\sigma)$ of $\text{Rot}_{n,d}$. If $\mathcal{T}_\sigma = (\sigma_t, t)_{t \in \mathbb{Z}}$ (so $\sigma_0 = \sigma$), then $\Phi_{n,d}(\sigma_0) = \sigma_0^{(n-d)}$, so $\mathcal{T}_{\Phi_{n,d}(\sigma_0)} = \mathcal{T}_\sigma^{(n-d)}$. Using the isomorphism ψ_{n-d} , we find that $\tilde{\Omega}(\sigma_0) = \tilde{\Omega}(\Phi_{n,d}(\sigma_0))$. Thus, we obtain a map

$$\Omega = \Omega_{n,d}: \text{Orb}_{\Phi_{n,d}} \rightarrow \text{Orb}_{\text{Rot}_{n,d}}$$

that sends the $\Phi_{n,d}$ -orbit containing a labeling μ to $\tilde{\Omega}(\mu)$. We will prove that this map satisfies the conditions in Proposition 6.4.

Lemma 6.16. *For any labeling $\sigma \in \Lambda_{\text{Path}_n}$, we have*

$$\text{stand}(\mathcal{T}_\sigma) = \text{stand}(\mathcal{T}_{\Phi_{n,d}(\sigma)}) \quad \text{and} \quad \overline{\text{stand}}(\mathcal{T}_\sigma) = \overline{\text{stand}}(\mathcal{T}_{\Phi_{n,d}(\sigma)}).$$

Hence, $\text{stand}(\mathcal{T}_\sigma)$ and $\overline{\text{stand}}(\mathcal{T}_\sigma)$ only depend on the orbit of $\Phi_{n,d}$ containing σ .

Proof. Let $\mathcal{T}_\sigma = (\sigma_t, t)_{t \in \mathbb{Z}}$ (so $\sigma_0 = \sigma$). Let $\mu = \Phi_{n,d}(\sigma)$. The stones diagram of the state $(\sigma_0^{(n-d)}, 0) = (\mu, 0)$ is obtained from that of $(\sigma_{n-d}, n-d)$ by moving all stones and replicas $n-d$ positions counterclockwise, so $\text{stand}(\mathcal{T}_\sigma) = \text{stand}(\mathcal{T}_{\Phi_{n,d}(\sigma)})$. Since $(\sigma_{n-d}, n-d)$ is in the timeline \mathcal{T}_σ and $(\mu, 0)$ is in the timeline \mathcal{T}_μ , the permutations $\text{stand}(\mathcal{T}_\sigma)$ and $\text{stand}(\mathcal{T}_\mu)$ are the standardizations of the sequences $\sigma_{n-d}^{-1}(1), \sigma_{n-d}^{-1}(2), \dots, \sigma_{n-d}^{-1}(n-d)$ and $\mu^{-1}(d+1), \mu^{-1}(d+2), \dots, \mu^{-1}(n)$, respectively. But $\mu = \text{cyc}^d(\sigma_{n-d})$, so these sequences are equal. \square

For $\rho \in S_r$, let $\text{rev}(\rho)$ be the permutation whose one-line notation is obtained by reversing that of ρ . Let $\delta: \mathbb{Z}/n\mathbb{Z} \rightarrow \mathbb{Z}/n\mathbb{Z}$ be the automorphism of Cycle_n defined by $\delta(i) = d+1-i$.

Given an orbit $\widehat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$, let $\text{rev}(\widehat{\mathcal{O}}) \in \text{Orb}_{\text{Rot}_{n,d}}$ be the orbit obtained by reversing all the compositions in $\widehat{\mathcal{O}}$.

Lemma 6.17. *For every $\sigma \in \Lambda_{\text{Path}_n}$, we have*

$$\text{stand}(\mathcal{T}_{\delta \circ \sigma}) = \text{rev}(\text{stand}(\mathcal{T}_\sigma)) \quad \text{and} \quad \overline{\text{stand}}(\mathcal{T}_{\delta \circ \sigma}) = \text{rev}(\overline{\text{stand}}(\mathcal{T}_\sigma)).$$

Furthermore, $\widetilde{\Omega}(\delta \circ \sigma) = \text{rev}(\widetilde{\Omega}(\sigma))$.

Proof. The first statement is immediate from the definitions. To see why the second statement is true, note that we can obtain the coins diagrams of the states in $\mathcal{T}_{\delta \circ \sigma}$ by “going backward in time” through the coins diagrams of the states in \mathcal{T}_σ and permuting the colors of the coins. To be more precise, let us write $\mathcal{T}_\sigma = (\sigma_t, t)_{t \in \mathbb{Z}}$ and $\mathcal{T}_{\delta \circ \sigma} = (\sigma'_t, t)_{t \in \mathbb{Z}}$ (so $\sigma_0 = \sigma$ and $\sigma'_0 = \delta \circ \sigma$). Then for every $t \in \mathbb{Z}$, the coins diagram of (σ'_t, t) is obtained from that of $(\sigma_{-t}, -t)$ by permuting the colors of the coins. Let $\mathcal{S} = (\kappa_0 \prec_{\mathcal{T}_\sigma} \dots \prec_{\mathcal{T}_\sigma} \kappa_d)$ be a transversal of $\mathbf{H}_{\mathcal{T}_\sigma}$ with energy composition $\mathcal{E}(\mathcal{S}) = (\varepsilon_1, \dots, \varepsilon_d)$. Then $\widetilde{\Omega}(\sigma)$ is the orbit of $\text{Rot}_{n,d}$ containing $(\varepsilon_1, \dots, \varepsilon_d)$. If κ_j occurs at time t_j and involves c_i , then there is a collision κ'_j in the timeline $\mathcal{T}_{\delta \circ \sigma}$ that occurs at time $-t_j$ and involves c_i (though c_i may have a different color in the coins diagrams of this timeline). In particular, κ'_d is a right-wall collision, κ'_0 is a left-wall collision, and $\kappa'_d \prec_{\mathcal{T}_{\delta \circ \sigma}} \dots \prec_{\mathcal{T}_{\delta \circ \sigma}} \kappa'_0$ is a saturated chain in $\mathbf{H}_{\mathcal{T}_{\delta \circ \sigma}}$. We have $\mathcal{E}(\kappa'_d \prec_{\mathcal{T}_{\delta \circ \sigma}} \dots \prec_{\mathcal{T}_{\delta \circ \sigma}} \kappa'_0) = (\varepsilon_d, \dots, \varepsilon_1)$. Starting with this saturated chain, one can straightforwardly apply Lemmas 6.10 and 6.11 to find that there is a transversal \mathcal{S}' in $\mathbf{H}_{\mathcal{T}_{\delta \circ \sigma}}$ with energy composition $\mathcal{E}(\mathcal{S}') = (\varepsilon_d, \dots, \varepsilon_1)$. Thus, $\widetilde{\Omega}(\delta \circ \sigma)$ is the orbit of $\text{Rot}_{n,d}$ containing $(\varepsilon_d, \dots, \varepsilon_1)$, which is $\text{rev}(\widetilde{\Omega}(\sigma))$. \square

Lemma 6.18. *For every $\mathcal{O} \in \text{Orb}_{\Phi_{n,d}}$, we have $|\Omega(\mathcal{O})| = \frac{d}{n} |\mathcal{O}|$.*

Proof. Fix $\mathcal{O} \in \text{Orb}_{\Phi_{n,d}}$, and let $\mathcal{T} = (\sigma_t, t)_{t \in \mathbb{Z}}$ be a timeline such that $\sigma_0 \in \mathcal{O}$. Consider a transversal $\mathcal{S} = (\kappa_0 \prec_{\mathcal{T}} \kappa_1 \prec_{\mathcal{T}} \dots \prec_{\mathcal{T}} \kappa_d)$ of $\mathbf{H}_{\mathcal{T}}$, and let $\mathcal{E}(\mathcal{S}) = (\varepsilon_1, \dots, \varepsilon_d)$. Then $\mathcal{E}(\mathcal{S})$ is in the orbit $\Omega(\mathcal{O})$. Let us define ε_k for all $k \in \mathbb{Z}$ by declaring $\varepsilon_{i+d} = \varepsilon_i$. Let t_j be the time when the collision κ_j occurs.

Consider the stones diagrams. Between times t_0 and t_1 , the stone with the same color as c_1 slides along the cycle carrying v_1 until sliding from underneath v_1 to underneath v_2 , which it carries until sliding underneath v_3 , and so on until it finally slides underneath v_{ε_1} . The positions of $v_1, \dots, v_{\varepsilon_1}$ throughout this interval of time are completely determined by the value of ε_1 , the permutations $\text{stand}(\mathcal{T})$ and $\overline{\text{stand}}(\mathcal{T})$, and the residue of t_0 modulo n . It follows that $t_1 - t_0$ is determined by ε_1 , $\text{stand}(\mathcal{T})$, $\overline{\text{stand}}(\mathcal{T})$, and the residue of t_0 modulo n . Between times t_1 and t_2 , the stone with the same color as c_2 slides along the cycle carrying v_{ε_1+1} until sliding from underneath v_{ε_1+1} to underneath v_{ε_1+2} , which it carries until sliding underneath v_{ε_1+3} , and so on until it finally slides underneath $v_{\varepsilon_1+\varepsilon_2}$. The positions of $v_{\varepsilon_1+1}, \dots, v_{\varepsilon_1+\varepsilon_2}$ throughout this interval of time are determined by the pair $(\varepsilon_1, \varepsilon_2)$, the permutations $\text{stand}(\mathcal{T})$ and $\overline{\text{stand}}(\mathcal{T})$, and the residue of t_1 modulo n . Thus, $t_2 - t_1$ is determined by $(\varepsilon_1, \varepsilon_2)$, $\text{stand}(\mathcal{T})$, $\overline{\text{stand}}(\mathcal{T})$, and the residue of t_0 modulo n . In general, the values of $t_1 - t_0, \dots, t_d - t_{d-1}$ are determined by the energy composition $(\varepsilon_1, \dots, \varepsilon_d)$, the permutations $\text{stand}(\mathcal{T})$ and $\overline{\text{stand}}(\mathcal{T})$, and the residue of t_0 modulo n .

Let p be the period of $\mathbf{H}_{\mathcal{T}}$. By Lemma 6.14, p is equal to $|\Omega(\mathcal{O})|$, the size of the orbit of $\text{Rot}_{n,d}$ containing the composition $\mathcal{E}(\mathcal{T}) = (\varepsilon_1, \dots, \varepsilon_d)$. Hence, $\varepsilon_1 + \dots + \varepsilon_p = \frac{p}{d}(\varepsilon_1 + \dots + \varepsilon_d) = pn/d$. Let t_j^* be the time when the collision $\varphi^p(\kappa_j)$ occurs. Using Lemmas 6.10 and 6.11, we find that the edges in the half-diamond of $\mathbf{H}_{\mathcal{T}}$ between $\varphi^{i-1}(\kappa_0)$ and $\varphi^i(\kappa_0)$ both have energy ε_i . Therefore, by Lemma 6.13, the time between the collisions $\varphi^{i-1}(\kappa_0)$ and $\varphi^i(\kappa_0)$ is $(n-d)\varepsilon_i$. This shows that $\varphi^p(\kappa_0)$ occurs at time $t_0^* = t_0 + (n-d)(\varepsilon_1 + \dots + \varepsilon_p) = t_0 + pn(n-d)/d$. Because p is the size of an orbit of $\text{Rot}_{n,d}$, it is divisible by $d/\gcd(n, d)$; this implies that $t_0^* \equiv t_0 \pmod{n}$. By the definition of p , the transversal $\overline{\varphi^p(\mathcal{T})}$ has the same energy composition $(\varepsilon_1, \dots, \varepsilon_d)$ as \mathcal{T} . Since the permutations $\text{stand}(\mathcal{T})$ and $\overline{\text{stand}(\mathcal{T})}$ only depend on \mathcal{T} , it follows from the preceding paragraph that $t_j^* - t_{j-1}^* = t_j - t_{j-1}$ for all $1 \leq j \leq d$; consequently, $t_j^* = t_j + pn(n-d)/d$ for all $0 \leq j \leq d$. From this, we deduce that $\sigma_t = \sigma_{t+pn(n-d)/d}$ for all $t \in \mathbb{Z}$. In fact, pn/d is the smallest positive integer ℓ such that $\sigma_t = \sigma_{t+\ell(n-d)}$ for all $t \in \mathbb{Z}$ (otherwise, we could reverse this argument to find that the period of $\mathbf{H}_{\mathcal{T}}$ is smaller than p).

According to (6.2), we have $\Phi_{n,d}^{pn/d} = \nu_{pn(n-d)/d} \cdots \nu_2 \nu_1$, so $\Phi_{n,d}^{pn/d}(\sigma_0) = \sigma_{pn(n-d)/d} = \sigma_0$. Hence, $|\mathcal{O}|$ divides pn/d . On the other hand, since Lemma 6.3 tells us that $|\mathcal{O}|$ is divisible by $n/\gcd(n, d)$, we can use (6.2) to find that

$$\sigma_0 = \Phi_{n,d}^{|\mathcal{O}|}(\sigma_0) = (\nu_{|\mathcal{O}|(n-d)} \cdots \nu_2 \nu_1)(\sigma_0) = \sigma_{|\mathcal{O}|(n-d)}.$$

Since $|\mathcal{O}|(n-d)$ is divisible by n (by Lemma 6.3), we have $\nu_{t+|\mathcal{O}|(n-d)} = \nu_t$ for all integers t . Consequently, $\sigma_t = \sigma_{t+|\mathcal{O}|(n-d)}$ for all integers t . Appealing to the last sentence in the previous paragraph, we deduce that $|\mathcal{O}| \geq pn/d = \frac{n}{d}|\Omega(\mathcal{O})|$. As $|\mathcal{O}|$ divides pn/d , the proof is complete. \square

Recall that Lemma 6.16 tells us that $\text{stand}(\mathcal{T}_\sigma)$ and $\overline{\text{stand}(\mathcal{T}_\sigma)}$ only depend on the orbit of $\Phi_{n,d}$ containing σ . In order to complete the proof of Proposition 6.4, we just need to show that $|\Omega^{-1}(\widehat{\mathcal{O}})| = d!(n-d)!$ for every $\widehat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$. We will do this by showing that for each pair of permutations $(\rho, \bar{\rho}) \in S_d \times S_{n-d}$, there exists a unique orbit $\mathcal{O} \in \Omega^{-1}(\widehat{\mathcal{O}})$ such that $\text{stand}(\mathcal{T}_\sigma) = \rho$ and $\overline{\text{stand}(\mathcal{T}_\sigma)} = \bar{\rho}$ for every $\sigma \in \mathcal{O}$. We start by proving existence; uniqueness will then follow from a simple counting argument. We implore the reader to consult Example 6.20 while reading the proof of the next lemma.

Lemma 6.19. *Suppose $\widehat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$ and $(\rho, \bar{\rho}) \in S_d \times S_{n-d}$. There exists an orbit $\mathcal{O} \in \Omega^{-1}(\widehat{\mathcal{O}})$ such that $\text{stand}(\mathcal{T}_\sigma) = \rho$ and $\overline{\text{stand}(\mathcal{T}_\sigma)} = \bar{\rho}$ for every $\sigma \in \mathcal{O}$.*

Proof. If $d = 1$, then the result is obvious because $\widehat{\mathcal{O}} = \text{Comp}_1(n) = \{(n)\}$. Therefore, we may assume $d \geq 2$ and proceed by induction on d . It follows from Lemma 6.17 that the following two statements are equivalent:

1. There exists an orbit $\mathcal{O} \in \Omega^{-1}(\widehat{\mathcal{O}})$ such that $\text{stand}(\mathcal{T}_\sigma) = \rho$ and $\overline{\text{stand}(\mathcal{T}_\sigma)} = \bar{\rho}$ for every $\sigma \in \mathcal{O}$.
2. There exists an orbit $\mathcal{O} \in \Omega^{-1}(\text{rev}(\widehat{\mathcal{O}}))$ such that $\text{stand}(\mathcal{T}_\sigma) = \text{rev}(\rho)$ and $\overline{\text{stand}(\mathcal{T}_\sigma)} = \text{rev}(\bar{\rho})$ for every $\sigma \in \mathcal{O}$.

Therefore, we may assume² without loss of generality that the number 1 appears to the left of the number 2 in the one-line notation of ρ (otherwise, replace ρ , $\bar{\rho}$, and $\hat{\mathcal{O}}$ by $\text{rev}(\rho)$, $\text{rev}(\bar{\rho})$, and $\text{rev}(\hat{\mathcal{O}})$, respectively).

Since $d < n$, every composition in $\hat{\mathcal{O}}$ has a part that is strictly greater than 1. Thus, we may choose a composition $(\varepsilon_1, \dots, \varepsilon_d) \in \hat{\mathcal{O}}$ such that $\varepsilon_2 \geq 2$. We will also assume for simplicity that $\varepsilon_1 \geq 2$; the case when $\varepsilon_1 = 1$ is similar. Let ρ' be the permutation in S_{d-1} obtained from ρ by deleting the entry 1 and decreasing the remaining entries by 1. Let $\bar{\rho}'$ be the permutation in $S_{d-\varepsilon_1+1}$ obtained from $\bar{\rho}$ by deleting the entries $1, \dots, \varepsilon_1 - 1$ and decreasing the remaining entries by $\varepsilon_1 - 1$. Let $\hat{\mathcal{O}}'$ be the orbit of $\text{Rot}_{n-\varepsilon_1, d-1}$ containing $(\varepsilon_2, \dots, \varepsilon_d)$. By induction, there exists an orbit $\mathcal{O}' \in \Omega_{n-\varepsilon_1, d-1}^{-1}(\hat{\mathcal{O}}')$ such that $\text{stand}(\mathcal{T}_{\sigma'}) = \rho'$ and $\text{stand}(\mathcal{T}_{\sigma'}) = \bar{\rho}'$ for every $\sigma' \in \mathcal{O}'$ (the timeline $\mathcal{T}_{\sigma'}$ is defined with the parameters $n - \varepsilon_1$ and $d - 1$ replacing n and d).

Fix $\sigma'_0 \in \mathcal{O}'$, and consider the timeline $\mathcal{T}_{\sigma'_0} = (\sigma'_t, t)_{t \in \mathbb{Z}}$ (defined with the parameters $n - \varepsilon_1$ and $d - 1$). Let us identify $\text{Path}_{n-\varepsilon_1}$ with the subgraph of Path_n obtained by deleting the vertices $v_1, \dots, v_{\varepsilon_1}$. Thus, the leftmost vertex in $\text{Path}_{n-\varepsilon_1}$ is v_{ε_1+1} , and the replicas appearing in the stones diagrams of states in $\mathcal{T}_{\sigma'_0}$ are $\mathbf{v}_{\varepsilon_1+1}, \dots, \mathbf{v}_n$. Let $\kappa'_1 \prec_{\mathcal{T}_{\sigma'_0}} \dots \prec_{\mathcal{T}_{\sigma'_0}} \kappa'_d$ be a transversal of $\mathbf{H}_{\mathcal{T}_{\sigma'_0}}$ with energy composition $(\varepsilon_2, \dots, \varepsilon_d)$. Let $k + 1$ be the first time after the collision κ'_1 when the leftmost coin moves. Because $\varepsilon_2 \geq 2$, the leftmost coin occupies v_{ε_1+1} in the coins diagram of (σ'_k, k) and occupies v_{ε_1+2} in the coins diagram of $(\sigma'_{k+1}, k + 1)$. Let v_η be the vertex of $\text{Path}_{n-\varepsilon_1}$ such that $\sigma'_k(v_\eta) = d + k \in \mathbb{Z}/(n - \varepsilon_1)\mathbb{Z}$. In the stones diagram of (σ'_k, k) , the replica \mathbf{v}_η sits one space clockwise from the consecutive block of stones.

We can construct the stones diagram of a state (μ, m) with $\mu \in \Lambda_{\text{Path}_n}$ from the stones diagram of (σ'_k, k) by inserting vertices to replace $\text{Cycle}_{n-\varepsilon_1}$ by Cycle_n , adding one stone, and adding the replicas $\mathbf{v}_1, \dots, \mathbf{v}_{\varepsilon_1}$. When we do this, we make sure to keep the stones on the consecutive block of vertices $m + d, \dots, m + 1$, and we make sure that the replicas that were sitting on stones remain on stones. We place the replica $\mathbf{v}_{\varepsilon_1}$ on the newly inserted stone. We can also ensure that $\text{stand}(\mu, m) = \rho$ and $\text{stand}(\mu, m) = \bar{\rho}$, and we can choose m so that \mathbf{v}_η is on the vertex $m + d + 1$. (In fact, these conditions uniquely determine μ and uniquely determine m modulo n .) Let σ_0 be the unique labeling in Λ_{Path_n} such that the timeline \mathcal{T}_{σ_0} contains the state (μ, m) .

Consider watching the coins diagrams of the states in \mathcal{T}_{σ_0} evolve over time. At time m , the coins c_1 and c_2 occupy v_{ε_1} and v_{ε_1+1} , respectively. At time $m + 1$, the coin c_2 moves to v_{ε_1+2} , and c_1 stays on v_{ε_1} (we are using the fact that 1 appears to the left of 2 in ρ). This implies (since $\varepsilon_1 \geq 2$) that the last collision involving c_1 that occurred at or before time m must have been a two-coins collision involving c_1 and c_2 ; let us call this collision κ_1 and say that it occurred at time m' . After time $m + 1$, c_1 will move leftward until reaching v_1 , where it will take part in a left-wall collision κ_0^* . Meanwhile, c_2 will travel rightward until reaching $v_{\varepsilon_1+\varepsilon_2}$, where it will collide with c_3 in a two-coins collision κ_2 . Then c_3 will move rightward until reaching $v_{\varepsilon_1+\varepsilon_2+\varepsilon_3}$, where it will collide with c_4 in a two-coins collision κ_3 , and so on. Eventually, c_d moves rightward and takes part in a right-wall collision κ_d . The key ob-

²This assumption might seem innocuous, but it is actually imperative for our argument. Thus, Lemma 6.17 really is quite crucial.

servation here is that throughout this process, in the stones diagrams, any stone carrying a replica \mathbf{v}_ℓ with $\ell \geq \varepsilon_1 + 2$ will just slide through any replica $\mathbf{v}_{\ell'}$ with $\ell' \leq \varepsilon_1$. This means that the replicas $\mathbf{v}_1, \dots, \mathbf{v}_{\varepsilon_1}$ that we inserted when passing from $\text{Cycle}_{n-\varepsilon_1}$ to Cycle_n will not affect where the collisions $\kappa_2, \kappa_3, \dots, \kappa_d$ occur. This is why $\kappa_2, \dots, \kappa_d$ occur at the same places (though possibly at different times) as $\kappa'_2, \dots, \kappa'_d$, respectively. We find that $\mathcal{E}(\kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_0^*) = \varepsilon_1$ and $\mathcal{E}(\kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_2 \triangleleft_{\mathcal{T}_{\sigma_0}} \dots \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_d) = (\varepsilon_2, \dots, \varepsilon_d)$. The edge $\kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_0^*$ is the top edge in a half-diamond; let $\kappa_0 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_1$ be the bottom edge of the same half-diamond. Then Lemma 6.10 tells us that $\mathcal{E}(\kappa_0 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_1) = \varepsilon_1$, so the transversal $\kappa_0 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \dots \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_d$ of $\mathbf{H}_{\mathcal{T}_{\sigma_0}}$ has energy composition $(\varepsilon_1, \dots, \varepsilon_d)$. Thus, $\tilde{\Omega}(\sigma_0)$ is the orbit of $\text{Rot}_{n,d}$ containing $(\varepsilon_1, \dots, \varepsilon_d)$. If \mathcal{O} is the orbit of $\Phi_{n,d}$ containing σ_0 , then $\Omega(\mathcal{O}) = \hat{\mathcal{O}}$. Furthermore, $\text{stand}(\mathcal{T}_{\sigma_0}) = \rho$ and $\overline{\text{stand}}(\mathcal{T}_{\sigma_0}) = \bar{\rho}$. According to Lemma 6.16, we have $\text{stand}(\mathcal{T}_\sigma) = \rho$ and $\overline{\text{stand}}(\mathcal{T}_\sigma) = \bar{\rho}$ for all $\sigma \in \mathcal{O}$. \square




Figure 6.5: The stones diagrams and coins diagrams of the states at times $0, 1, \dots, 7$ in the timeline $\mathcal{T}_{\sigma'_0}$ from Example 6.20.

Example 6.20. Let us illustrate the proof of Lemma 6.19. Suppose $n = 8$, $d = 3$, $\rho = 132$, and $\bar{\rho} = 52413$. Let $\hat{\mathcal{O}}$ be the orbit of $\text{Rot}_{8,3}$ containing the composition $(\varepsilon_1, \varepsilon_2, \varepsilon_3) = (4, 3, 1)$. Note that 1 appears before 2 in ρ and that $\varepsilon_2 \geq 2$. We have $\rho' = 21$ and $\bar{\rho}' = 21$.

We can choose σ'_0 to be the labeling such that the stones diagrams and coins diagrams of the states of $\mathcal{T}_{\sigma'_0}$ at times $0, 1, \dots, 7$ are shown in Figure 6.5. One can check that the states in this timeline are periodic with period 8. We can choose the transversal $\kappa'_1 \triangleleft_{\mathcal{T}_{\sigma'_0}} \kappa'_2 \triangleleft_{\mathcal{T}_{\sigma'_0}} \kappa'_3$ so that κ'_1 is the left-wall collision at time 0, κ'_2 is the two-coins collision at time 3, and κ'_3 is the right-wall collision at time 5. We have $k = 1$ and $\eta = 7$.

Figure 6.6 illustrates how we construct the stones diagram of (μ, m) from that of $(\sigma'_1, 1)$. In this example, $m = 2$. Four vertices were inserted to transform Cycle_4 into Cycle_8 , and the

vertices were then renamed. Since $\eta = 7$, we have placed v_7 on the vertex $m + d + 1 = 6$. Note that the standardization of v_4, v_6, v_5 is $132 = \rho$ and that the standardization of v_8, v_2, v_7, v_1, v_3 is $52413 = \bar{\rho}$.

Figure 6.7 shows the stones diagrams and coins diagrams of the states in \mathcal{T}_{σ_0} at the times $0, \dots, 11$. (The labelings of the states in this timeline are actually periodic with period 40, but we chose not to draw the diagrams of 40 states.) The collision κ_1 involves c_1 and c_2 and occurs at time 0. Then κ_0^* is the left-wall collision at time 9. The collision κ_2 involves c_2 and c_3 and occurs at time 6, while κ_3 is the right-wall collision at time 11. Observe that $\mathcal{E}(\kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_0^*) = 4 = \varepsilon_1$ and $\mathcal{E}(\kappa_1 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_2 \triangleleft_{\mathcal{T}_{\sigma_0}} \kappa_3) = (3, 1) = (\varepsilon_2, \varepsilon_3)$.




Figure 6.6: On the left is the stones diagram of $(\sigma'_1, 1)$ from Example 6.20. On the right is the stones diagram of $(\mu, 2)$, which is constructed from that of $(\sigma'_1, 1)$ by inserting four new vertices (shaded), a new stone (gold), and the new replicas v_1, v_2, v_3, v_4 .

Proof of Proposition 6.4. We know by Lemma 6.18 that $|\Omega(\mathcal{O})| = \frac{d}{n}|\mathcal{O}|$ for every $\mathcal{O} \in \text{Orb}_{\Phi_{n,d}}$. It follows from Lemma 6.19 that $|\Omega^{-1}(\hat{\mathcal{O}})| \geq d!(n-d)!$ for every $\hat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$. Therefore,

$$\begin{aligned} n! &= |\Lambda_{\text{Path}_n}| = \sum_{\mathcal{O} \in \text{Orb}_{\Phi_{n,d}}} |\mathcal{O}| = \sum_{\hat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}} \sum_{\mathcal{O} \in \Omega^{-1}(\hat{\mathcal{O}})} |\mathcal{O}| = \sum_{\hat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}} |\Omega^{-1}(\hat{\mathcal{O}})| \cdot \frac{n}{d} |\hat{\mathcal{O}}| \\ &\geq d!(n-d)! \frac{n}{d} \sum_{\hat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}} |\hat{\mathcal{O}}| = n(d-1)!(n-d)! |\text{Comp}_d(n)| = n(d-1)!(n-d)! \binom{n-1}{d-1} = n!. \end{aligned}$$

This inequality must actually be an equality, so we must have $|\Omega^{-1}(\hat{\mathcal{O}})| = d!(n-d)!$ for every $\hat{\mathcal{O}} \in \text{Orb}_{\text{Rot}_{n,d}}$. □

As discussed at the end of Section 6.1, Proposition 6.4 implies Theorem 1.3.




Figure 6.7: The stones diagrams and coins diagrams of the states at times $0, 1, \dots, 11$ in the timeline \mathcal{T}_{σ_0} from Example 6.20. Blue numbers indicate times. Collisions are represented by color-coded stars.

7. Orbit Structure of Broken Promotion

In this final section, we prove Theorems 1.4 and 1.5, which describe the orbit structure of cyc Bro_B for particular choices of the subset $B \subseteq \mathbb{Z}/n\mathbb{Z}$.

Proof of Theorem 1.4. Let β be the acyclic orientation of Cycle_n whose unique source is d and whose unique sink is n . To ease notation, let $F(q) = n(d-1)!(n-d-1)![n-d]_{q^d} \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q$. Theorem 1.3 tells us that TPro_β has order $d(n-d)$ and that the triple $(\Lambda_{\text{Path}_n}, \text{TPro}_\beta, F(q))$ exhibits the cyclic sieving phenomenon. Since the sizes of the orbits of TPro_β are all divisible by d (by Proposition 5.1), it follows that TPro_β^d has order $n-d$ and that the triple $(\Lambda_{\text{Path}_n}, \text{TPro}_\beta^d, F(q))$ also exhibits the cyclic sieving phenomenon. By Remark 6.2, we have

$$\text{TPro}_\beta^d = \left(\text{cyc}^{-1} \text{Bro}_{\{1, \dots, d\}}^{-1} \right)^n = (\text{Bro}_{\{1, \dots, d\}} \text{cyc})^{-n} = \text{cyc}^{-1} (\text{cyc Bro}_{\{1, \dots, d\}})^{-n} \text{cyc},$$

so TPro_β^d and $(\text{cyc Bro}_{\{1, \dots, d\}})^n$ have the same orbit structure. Consequently, $(\text{cyc Bro}_{\{1, \dots, d\}})^n$ has order $n-d$, and the triple $(\Lambda_{\text{Path}_n}, (\text{cyc Bro}_{\{1, \dots, d\}})^n, F(q))$ satisfies the cyclic sieving phenomenon. It follows immediately from Proposition 3.6 that the orbit sizes of $\text{cyc Bro}_{\{1, \dots, d\}}$ are all divisible by n . Therefore, $\text{cyc Bro}_{\{1, \dots, d\}}$ has order $(n-d)n$, and if $\{k_i^{m_i} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of $(\text{cyc Bro}_{\{1, \dots, d\}})^n$, then $\{(nk_i)^{m_i/n} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of $\text{cyc Bro}_{\{1, \dots, d\}}$. According to Lemma 2.1, the triple $(\Lambda_{\text{Path}_n}, \text{cyc Bro}_{\{1, \dots, d\}}, \frac{1}{n}[n]_{q^{n-d}} F(q))$ exhibits the cyclic sieving phenomenon. This completes the proof. \square

Proof of Theorem 1.5. Let d, n, s_i , and \mathcal{R} be as in the statement of the theorem. Let β be the acyclic orientation of Cycle_n whose sources are the elements of the set $\mathcal{S} = \{s_1, \dots, s_d\}$ and whose sinks are the elements of $\mathcal{S} - 1$. Let $F(q) = n(d-1)!(n-d-1)![n-d]_{q^d} \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q$. Theorem 1.3 tells us that TPro_β has order $d(n-d)$ and that the triple $(\Lambda_{\text{Path}_n}, \text{TPro}_\beta, F(q))$ exhibits the cyclic sieving phenomenon. Since the sizes of the orbits of TPro_β are all divisible by $n-d$ (by Proposition 5.1), it follows that TPro_β^{n-d} has order d and that the triple $(\Lambda_{\text{Path}_n}, \text{TPro}_\beta^{n-d}, F(q))$ also exhibits the cyclic sieving phenomenon. If we set $\gamma = n-d$, $q = n$, and $r = 0$ in Proposition 4.1, we find that $\text{TPro}_\beta^{n-d} = (\text{cyc Bro}_\mathcal{R})^n$. It follows from Proposition 3.6 that the sizes of the orbits of $\text{cyc Bro}_\mathcal{R}$ are all divisible by n . Therefore, $\text{cyc Bro}_\mathcal{R}$ has order dn , and if $\{k_i^{m_i} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of $(\text{cyc Bro}_\mathcal{R})^n$, then $\{(nk_i)^{m_i/n} : 1 \leq i \leq \ell\}$ is the multiset of orbit sizes of $\text{cyc Bro}_\mathcal{R}$. According to Lemma 2.1, the triple $(\Lambda_{\text{Path}_n}, \text{cyc Bro}_\mathcal{R}, \frac{1}{n}[n]_{q^d} F(q))$ exhibits the cyclic sieving phenomenon. \square

8. Future Directions

Theorem 1.1 determines the orbit structure of toric promotion when G is a forest. It is still open to understand the dynamics of toric promotion for other graphs, including cycle graphs.

Theorem 1.3 determines the orbit structure of any permutoric promotion operator when G is a path. It would be interesting to gain a better understanding of TPro_π when G is another type of tree, even in the special case when π^{-1} has 2 cyclic descents. A natural place to start could

be the case when G is obtained from Path_{n-1} by adding a new vertex that is adjacent to v_{n-2} (i.e., G is the Dynkin diagram of type D_n).

Acknowledgements

Colin Defant thanks Tom Roby for initially suggesting the generalization from toric promotion to permutoric promotion. We thank Nathan Williams for inspiring conversations. We thank the anonymous referees for their helpful feedback.

References

- [AKS14] Arvind Ayyer, Steven Klee, and Anne Schilling. Combinatorial Markov chains on linear extensions. *J. Algebraic Combin.*, 39(4):853–881, 2014. doi:10.1007/s10801-013-0470-9.
- [ARR20] Ron M. Adin, Victor Reiner, and Yuval Roichman. On cyclic descents for tableaux. *Int. Math. Res. Not. IMRN*, 2020(24):10231–10276, 2020. doi:10.1093/imrn/rny280.
- [BSV21] Joseph Bernstein, Jessica Striker, and Corey Vorland. P -strict promotion and B -bounded rowmotion, with applications to tableaux of many flavors. *Comb. Theory*, 1:Paper No. 8, 46, 2021. doi:10.5070/C61055369.
- [Def23] Colin Defant. Toric promotion. *Proc. Amer. Math. Soc.*, 151(1):45–57, 2023. doi:10.1090/proc/16079.
- [DK21] Colin Defant and Noah Kravitz. Friends and strangers walking on graphs. *Comb. Theory*, 1:Paper No. 6, 34, 2021. doi:10.5070/C61055363.
- [DK23] Colin Defant and Noah Kravitz. Promotion sorting. *Order*, 40(1):199–216, 2023. doi:10.1007/s11083-022-09603-9.
- [DMR16] Mike Develin, Matthew Macauley, and Victor Reiner. Toric partial orders. *Trans. Amer. Math. Soc.*, 368(4):2263–2287, 2016. doi:10.1090/tran/6356.
- [DPS17] Kevin Dilks, Oliver Pechenik, and Jessica Striker. Resonance in orbits of plane partitions and increasing tableaux. *J. Combin. Theory Ser. A*, 148:244–274, 2017. doi:10.1016/j.jcta.2016.12.007.
- [DSV19] Kevin Dilks, Jessica Striker, and Corey Vorland. Rowmotion and increasing labeling promotion. *J. Combin. Theory Ser. A*, 164:72–108, 2019. doi:10.1016/j.jcta.2018.12.007.
- [EG87] Paul Edelman and Curtis Greene. Balanced tableaux. *Adv. in Math.*, 63(1):42–99, 1987. doi:10.1016/0001-8708(87)90063-6.
- [Hai92] Mark D. Haiman. Dual equivalence with applications, including a conjecture of Proctor. *Discrete Math.*, 99(1-3):79–113, 1992. doi:10.1016/0012-365X(92)90368-P.

- [HR22] Sam Hopkins and Martin Rubey. Promotion of Kreweras words. *Selecta Math. (N.S.)*, 28(1):Paper No. 10, 38, 2022. doi:10.1007/s00029-021-00714-6.
- [Hua20] Brice Huang. Cyclic descents for general skew tableaux. *J. Combin. Theory Ser. A*, 169:105120, 45, 2020. doi:10.1016/j.jcta.2019.105120.
- [MR94] Claudia Malvenuto and Christophe Reutenauer. Evacuation of labelled graphs. *Discrete Math.*, 132(1-3):137–143, 1994. doi:10.1016/0012-365X(92)00569-D.
- [PPR09] T. Kyle Petersen, Pavlo Pylyavskyy, and Brendon Rhoades. Promotion and cyclic sieving via webs. *J. Algebraic Combin.*, 30(1):19–41, 2009. doi:10.1007/s10801-008-0150-3.
- [PR15] James Propp and Tom Roby. Homomesy in products of two chains. *Electron. J. Combin.*, 22(3):Paper 3.4, 29, 2015. doi:10.37236/3579.
- [PS18] Svetlana Poznanović and Kara Stasikelis. Properties of the promotion Markov chain on linear extensions. *J. Algebraic Combin.*, 47(3):505–528, 2018. doi:10.1007/s10801-017-0783-1.
- [Rho10] Brendon Rhoades. Cyclic sieving, promotion, and representation theory. *J. Combin. Theory Ser. A*, 117(1):38–76, 2010. doi:10.1016/j.jcta.2009.03.017.
- [RSW04] V. Reiner, D. Stanton, and D. White. The cyclic sieving phenomenon. *J. Combin. Theory Ser. A*, 108(1):17–50, 2004. doi:10.1016/j.jcta.2004.04.009.
- [Sch63] M. P. Schützenberger. Quelques remarques sur une construction de Schensted. *Math. Scand.*, 12:117–128, 1963. doi:10.7146/math.scand.a-10676.
- [Sch72] M. P. Schützenberger. Promotion des morphismes d’ensembles ordonnés. *Discrete Math.*, 2(1):73–94, 1972. doi:10.1016/0012-365X(72)90062-3.
- [Sch76] M. P. Schützenberger. Evacuations. In *Colloquio Internazionale sulle Teorie Combinatorie (Roma, 1973), Tomo I*, pages 257–264. Accad. Naz. Lincei, Rome, 1976.
- [Sta09] Richard P. Stanley. Promotion and evacuation. *Electron. J. Combin.*, 16(2):Research Paper 9, 24, 2009. doi:10.37236/75.
- [SW12] Jessica Striker and Nathan Williams. Promotion and rowmotion. *European J. Combin.*, 33(8):1919–1942, 2012. doi:10.1016/j.ejc.2012.05.003.