Séminaire Lotharingien de Combinatoire **91B** (2024)
Article #2, 12 pp.

Proceedings of the 36th Conference on Formal Power Series and Algebraic Combinatorics (Bochum)

Toric and Permutoric Promotion

Colin Defant*[1], Rachana Madhukara†[2], and Hugh Thomas‡[3]

[1]*Department of Mathematics, Harvard University, Cambridge, MA 02139, USA*
[2]*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[3]*Lacim, UQAM, Montréal, QC, H3C 3P8 Canada*

Abstract. We introduce *toric promotion* as a cyclic analogue of Schützenberger's promotion operator. Toric promotion acts on the set of labelings of a graph G; it is defined as the composition of certain toggle operators, listed in a natural cyclic order. We provide a surprisingly simple description of the orbit structure of toric promotion when G is a forest. We then consider more general permutoric promotion operators, which are defined as compositions of the same toggle operators, but in permuted orders. When G is a path graph, we provide a complete description of the orbit structures of all permutoric promotion operators, showing that they satisfy the cyclic sieving phenomenon.

Keywords: promotion, cyclic analogue, cyclic sieving, toggle operator

This is an extended abstract for the articles [2] and [4]. The first of these articles—written by the first author—focuses on toric promotion, while the second article—written by all three authors—concerns the more general permutoric promotion operators.

1 Introduction

In his study of the Robinson–Schensted–Knuth correspondence, Schützenberger [9, 10] introduced a beautiful bijective operator called *promotion*, which acts on the set of linear extensions of a finite poset. Haiman [6] and Malvenuto–Reutenauer [7] found that promotion could be defined as a composition of local *toggle operators* (also called *Bender–Knuth involutions*). 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 [7], we define promotion on labelings of graphs instead of linear extensions of posets. Let $G = (V, E)$ be a graph with n vertices. A *labeling* of G is a bijection $V \to \mathbb{Z}/n\mathbb{Z}$. We denote the set

*colindefant@gmail.com. Supported by the National Science Foundation under Award No. 2201907 and by a Benjamin Peirce Fellowship at Harvard University.
†rachanam@mit.edu.
‡thomas.hugh_r@uqam.ca.

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 $\tau_i \colon \Lambda_G \to \Lambda_G$ 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 $\mathrm{Pro} \colon \Lambda_G \to \Lambda_G$ given by

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

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 [1, 5] and the references therein). In the same vein, we introduce a cyclic analogue of promotion called *toric promotion*; this is the operator $\mathrm{TPro} \colon \Lambda_G \to \Lambda_G$ given by

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

Our first main result reveals that toric promotion has remarkably nice dynamical properties when G is a forest.

Theorem 1 ([2]). *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 $(n-1)t/\gcd(t,n)$, 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* $\mathrm{TPro} \colon \Lambda_G \to \Lambda_G$ *has size $n-1$.*

Theorem 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 $\mathrm{Pro} \colon \Lambda_G \to \Lambda_G$ is 3224590642072800, whereas all orbits of $\mathrm{TPro} \colon \Lambda_G \to \Lambda_G$ have size 6.

We now consider a generalization of toric promotion in which the toggle operators $\tau_1, \ldots, \tau_n$ can be composed in any order. Let $[n] = \{1, \ldots, n\}$, and let $\pi \colon [n] \to \mathbb{Z}/n\mathbb{Z}$ be a bijection. The *permutoric promotion* operator $\mathrm{TPro}_\pi \colon \Lambda_G \to \Lambda_G$ is defined by

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

One would ideally hope to have an extension of Theorem 1 to arbitrary permutoric promotion operators. Unfortunately, trying to completely describe the orbit structure of $\mathrm{TPro}_\pi \colon \Lambda_G \to \Lambda_G$ 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. 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 $\dfrac{[k]_q!}{[r]_q![k-r]_q!} \in \mathbb{C}[q]$.

Let X be a finite set, and let $f\colon X \to X$ 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 [8], 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 $\pi\colon [n] \to \mathbb{Z}/n\mathbb{Z}$ 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 and Cycle_n be the n-vertex path graph and cycle graph, respectively. In [2], the first author conjectured that for every bijection $\pi\colon [n] \to \mathbb{Z}/n\mathbb{Z}$, the order of $\mathrm{TPro}_\pi\colon \Lambda_{\mathsf{Path}_n} \to \Lambda_{\mathsf{Path}_n}$ is $d(n-d)$, where d is the number of cyclic descents of π^{-1}. Our next main theorem not only proves this conjecture, but also determines the entire orbit structure of TPro_π in this case.

Theorem 2 ([4])**.** *Let $\pi\colon [n] \to \mathbb{Z}/n\mathbb{Z}$ be a bijection, and let d be the number of cyclic descents of π^{-1}. The order of the permutoric promotion operator $\mathrm{TPro}_\pi\colon \Lambda_{\mathsf{Path}_n} \to \Lambda_{\mathsf{Path}_n}$ is $d(n-d)$. Moreover, the following triple exhibits the cyclic sieving phenomenon:*

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

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

Remark 1. Theorem 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 cycles. Theorem 2 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 when π^{-1} has just 2 cyclic descents.

In Section 2, we summarize some of the main ideas that go into the proof of Theorem 2, referring the reader to our full article [4] for the (quite involved) details that we have omitted. We also briefly summarize a proof of Theorem 1 in Section 3, though we refer the reader to [2] for a full proof.

2 Dynamics of Permutoric Promotion

As before, fix a bijection $\pi\colon [n] \to \mathbb{Z}/n\mathbb{Z}$, and let d be the number of cyclic descents of π^{-1}. We assume from now on that G is the path graph Path_n so that TPro_π is an operator on $\Lambda_{\mathsf{Path}_n}$. Given a finite set X and an invertible map $f\colon X \to X$, we write Orb_f for the set of orbits of f.

2.1 A Reduction

Let $\mathsf{Comp}_d(n)$ 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 $\mathrm{Rot}_{n,d}\colon \mathsf{Comp}_d(n) \to \mathsf{Comp}_d(n)$ defined by $\mathrm{Rot}_{n,d}(a_1, a_2, \ldots, a_d) = (a_2, \ldots, a_d, a_1)$. Reiner, Stanton, and White [8] proved that the triple $\left(\mathsf{Comp}_d(n), \mathrm{Rot}_{n,d}, \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q\right)$ exhibits the cyclic sieving phenomenon. As it turns out, this result is responsible for the factor of $\begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q$ in the sieving polynomial in Theorem 2.

Let $\mathrm{cyc}\colon \Lambda_{\mathsf{Path}_n} \to \Lambda_{\mathsf{Path}_n}$ be the *cyclic shift* operator given by $(\mathrm{cyc}(\sigma))(v) = \sigma(v) + 1$. Let $\Phi_{n,d}\colon \Lambda_{\mathsf{Path}_n} \to \Lambda_{\mathsf{Path}_n}$ be the operator

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

Using the identity $\mathrm{cyc}\,\tau_i = \tau_{i+1}\,\mathrm{cyc}$ together with the fact that τ_i and τ_j commute whenever $j \notin \{i-1, i+1\}$, one can show (see [4] for details) that

$$\Phi_{n,d}^{n/\gcd(n,d)} = \mathrm{TPro}_\pi^{\mathrm{lcm}(d,n-d)}. \tag{2.1}$$

Using a result about *friends-and-strangers graphs* from [3], one can prove that every orbit of $\Phi_{n,d}$ has size divisible by $n/\gcd(n,d)$ (see [4, Lemma 6.3]). A substantial portion of our full article is devoted to proving that every orbit of TPro_π has size divisible by $\mathrm{lcm}(d, n-d)$ (see [4, Proposition 5.1]). Together with (2.1), these divisibility results allow us to transfer the problem of describing the orbit structure of TPro_π to that of describing the orbit structure of $\Phi_{n,d}$. Thus, we deduce Theorem 2 from the following proposition and the fact that $\left(\mathsf{Comp}_d(n), \mathrm{Rot}_{n,d}, \begin{bmatrix} n-1 \\ d-1 \end{bmatrix}_q\right)$ exhibits the cyclic sieving phenomenon.

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

2.2 Sliding Stones and Colliding Coins

We now discuss how to construct the map Ω from Proposition 1. 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 $1 \le r \le d$. 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

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

By combining the identity $\mathrm{cyc}\,\tau_i = \tau_{i+1}\,\mathrm{cyc}$ with the fact that cyc^n is the identity map, one can easily verify that $\Phi^m_{n,d} = \theta_{md(n-d)} \cdots \theta_2\theta_1 = \nu_{m(n-d)} \cdots \nu_2\nu_1$ whenever m is a positive multiple of $n/\gcd(n,d)$.

Define a *state* to be a pair $(\sigma, t) \in \Lambda_{\mathsf{Path}_n} \times \mathbb{Z}$; 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 $\sigma \in \Lambda_{\mathsf{Path}_n}$, let $\mathcal{T}_\sigma$ be the unique timeline containing the state $(\sigma, 0)$.

Let $v_1, \ldots, v_n$ be the vertices of Path_n, listed from left to right. For each $\ell \in [n]$, let $\mathbf{v}_\ell$ be a formal symbol associated to v_ℓ; we will call $\mathbf{v}_\ell$ a *replica*. Let $\mathsf{s}_1, \ldots, \mathsf{s}_d$ be stones of different colors. We define the *stones diagram* of a state (σ, t) as follows. Start with a copy of Cycle_n, whose vertices we identify with $\mathbb{Z}/n\mathbb{Z}$. Place $\mathsf{s}_1, \ldots, \mathsf{s}_d$ on the vertices $t+d, \ldots, t+1$, respectively. Then place each replica $\mathbf{v}_\ell$ 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. Now, $(\theta_{d(t-1)+i} \cdots \theta_{d(t-1)+1})(\sigma_{t-1})$ is obtained from $(\theta_{d(t-1)+i-1} \cdots \theta_{d(t-1)+1})(\sigma_{t-1})$ by applying the operator $\theta_{d(t-1)+i} = \tau_{t+d-i}$. If this operator has no effect, then we do not move any of the replicas $\mathbf{v}_1, \ldots, \mathbf{v}_n$ 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 1 illustrates these small steps for a particular example with $n = 8$, $d = 3$, and $t = 1$.

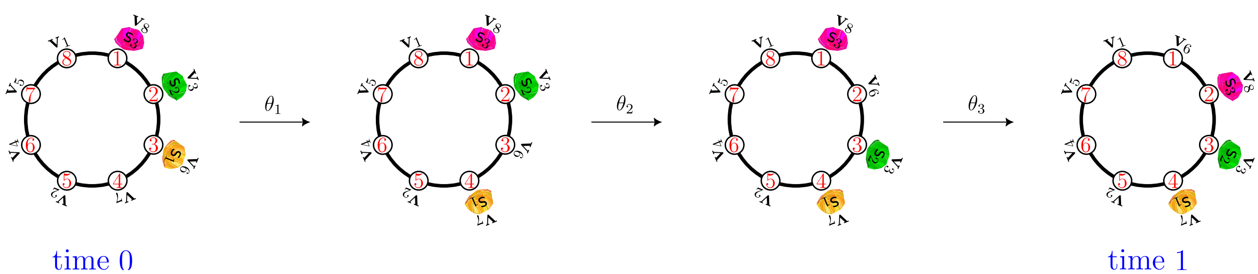

Figure 1: The $d = 3$ small steps transforming the stones diagram of a state at time 0 into the stones diagram of the next state at time 1.

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 $\mathbf{v}_\ell$ sitting on the stone s_i in 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 2 and 3). Note that the set of vertices of Path_n occupied by coins is $\{\sigma^{-1}(t+1), \dots, \sigma^{-1}(t+d)\}$.

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 $\mathbf{v}_\ell$ be the replica sitting on s_i right before the i-th small step, and let $\mathbf{v}_{\ell'}$ 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 $\mathbf{v}_\ell$ along with it or slide from underneath $\mathbf{v}_\ell$ to underneath $\mathbf{v}_{\ell'}$; 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 $\mathsf{c}_1, \dots, \mathsf{c}_d$ 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 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 1 ([4])**.** *When a coin moves, it moves in the direction that it expects to move.*

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 of a collision has to be slightly modified when considering a traffic jam that touches a wall; the reader may refer to [4] for details.

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

Example 1. Suppose $n = 6$ and $d = 3$. Figure 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. $\lozenge$

Example 2. Suppose $n = 6$ and $d = 3$. Figure 3 shows the stones diagrams and coins diagrams of a particular timeline at times $0, 1, \ldots, 17$. For brevity, we have not shown the individual small steps. All of the collisions the 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. $\lozenge$

Let $\mathrm{Coll}_{\mathcal{T}}$ 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 $\mathrm{Coll}_{\mathcal{T}}$ 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 $(\mathrm{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$ be the transitive closure of this directed graph. Let $\mathbf{H}_{\mathcal{T}}$ be the Hasse diagram of $(\mathrm{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$. This Hasse diagram, which is one of our primary tools, has the shape of a bi-infinite chain link fence (see Figure 4). Suppose $\kappa_1 \lessdot_{\mathcal{T}} \kappa_2$ 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 $\mathcal{E}(\kappa_1 \lessdot_{\mathcal{T}} \kappa_2)$, 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 $\kappa_1 \lessdot_{\mathcal{T}} \kappa_2 \lessdot_{\mathcal{T}} \cdots \lessdot_{\mathcal{T}} \kappa_r$ is a saturated chain in $\mathbf{H}_{\mathcal{T}}$, then we write $\mathcal{E}(\kappa_1 \lessdot_{\mathcal{T}} \kappa_2 \lessdot_{\mathcal{T}} \cdots \lessdot_{\mathcal{T}} \kappa_r)$ for the tuple $(\mathcal{E}(\kappa_1 \lessdot_{\mathcal{T}} \kappa_2), \ldots, \mathcal{E}(\kappa_{r-1} \lessdot_{\mathcal{T}} \kappa_r))$ of energies of the edges in the chain.

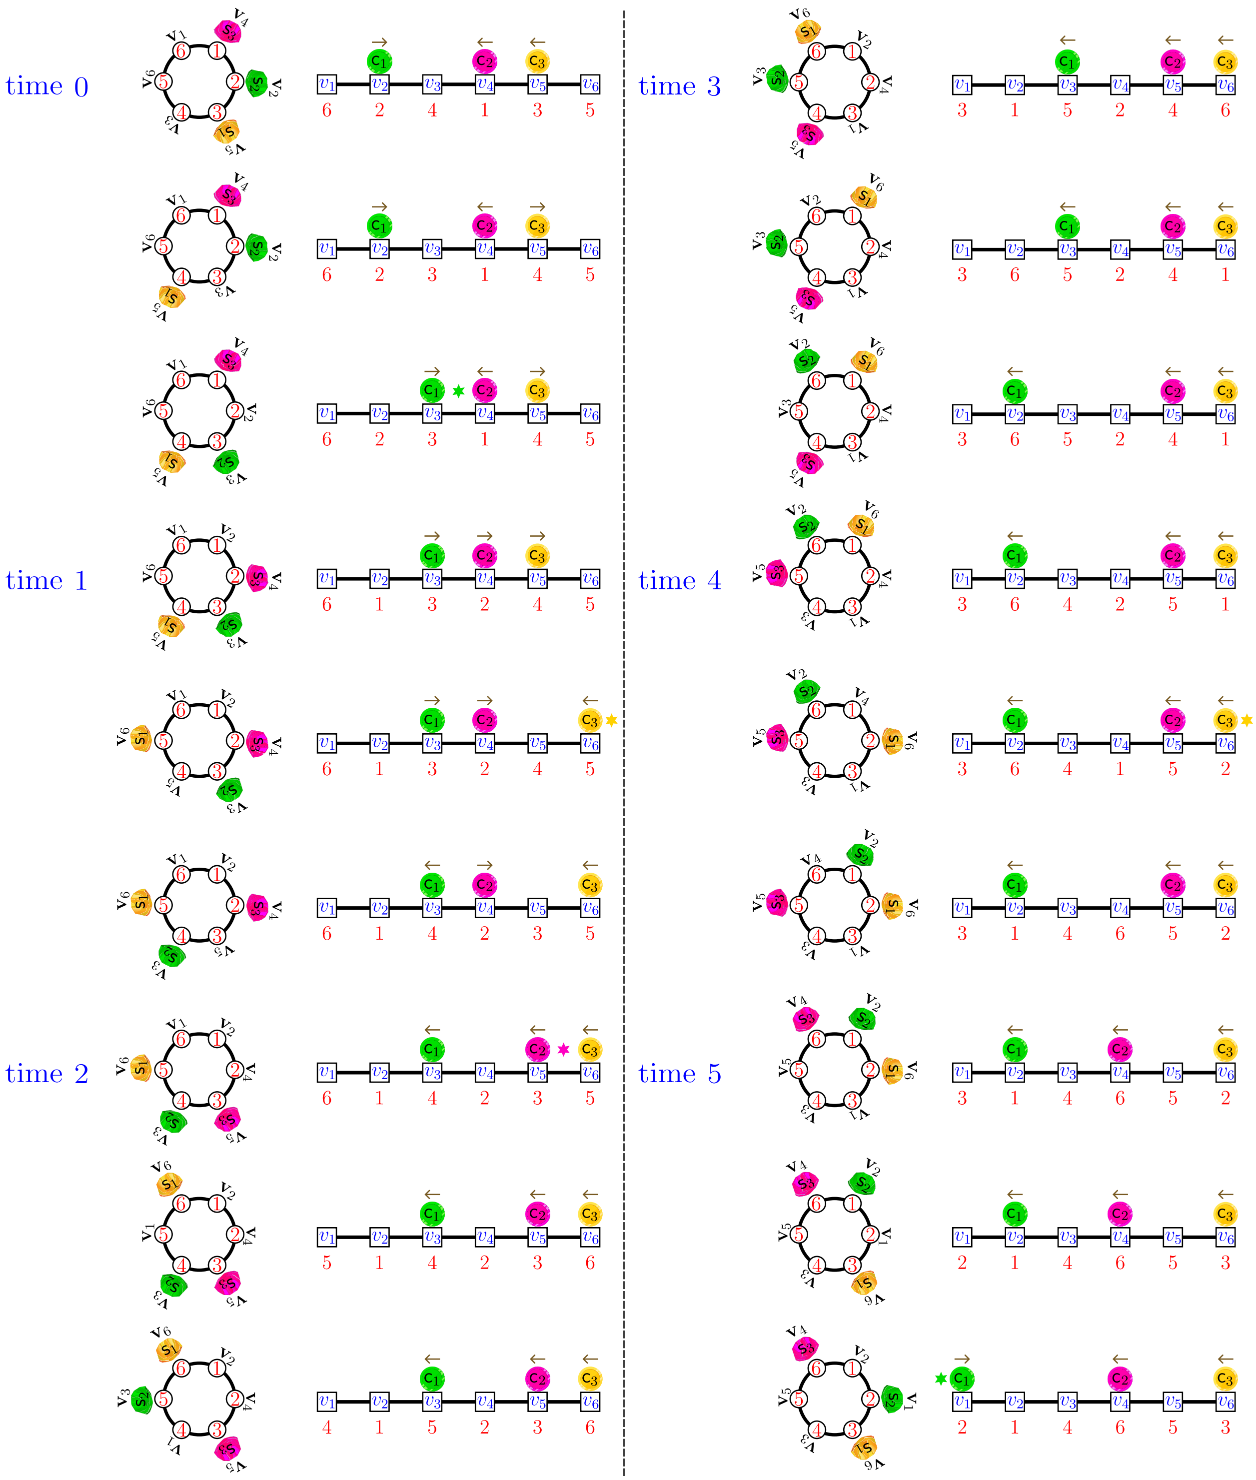

Figure 2: The evolution of stones diagrams and coins diagrams over time, with each individual small step illustrated. At each moment, we have drawn an arrow over each coin to indicate which direction it expects to move. Each collision is indicated by a star whose color is the same as that of the stone that moved to cause the collision. Each labeling is depicted in red numbers below the path.

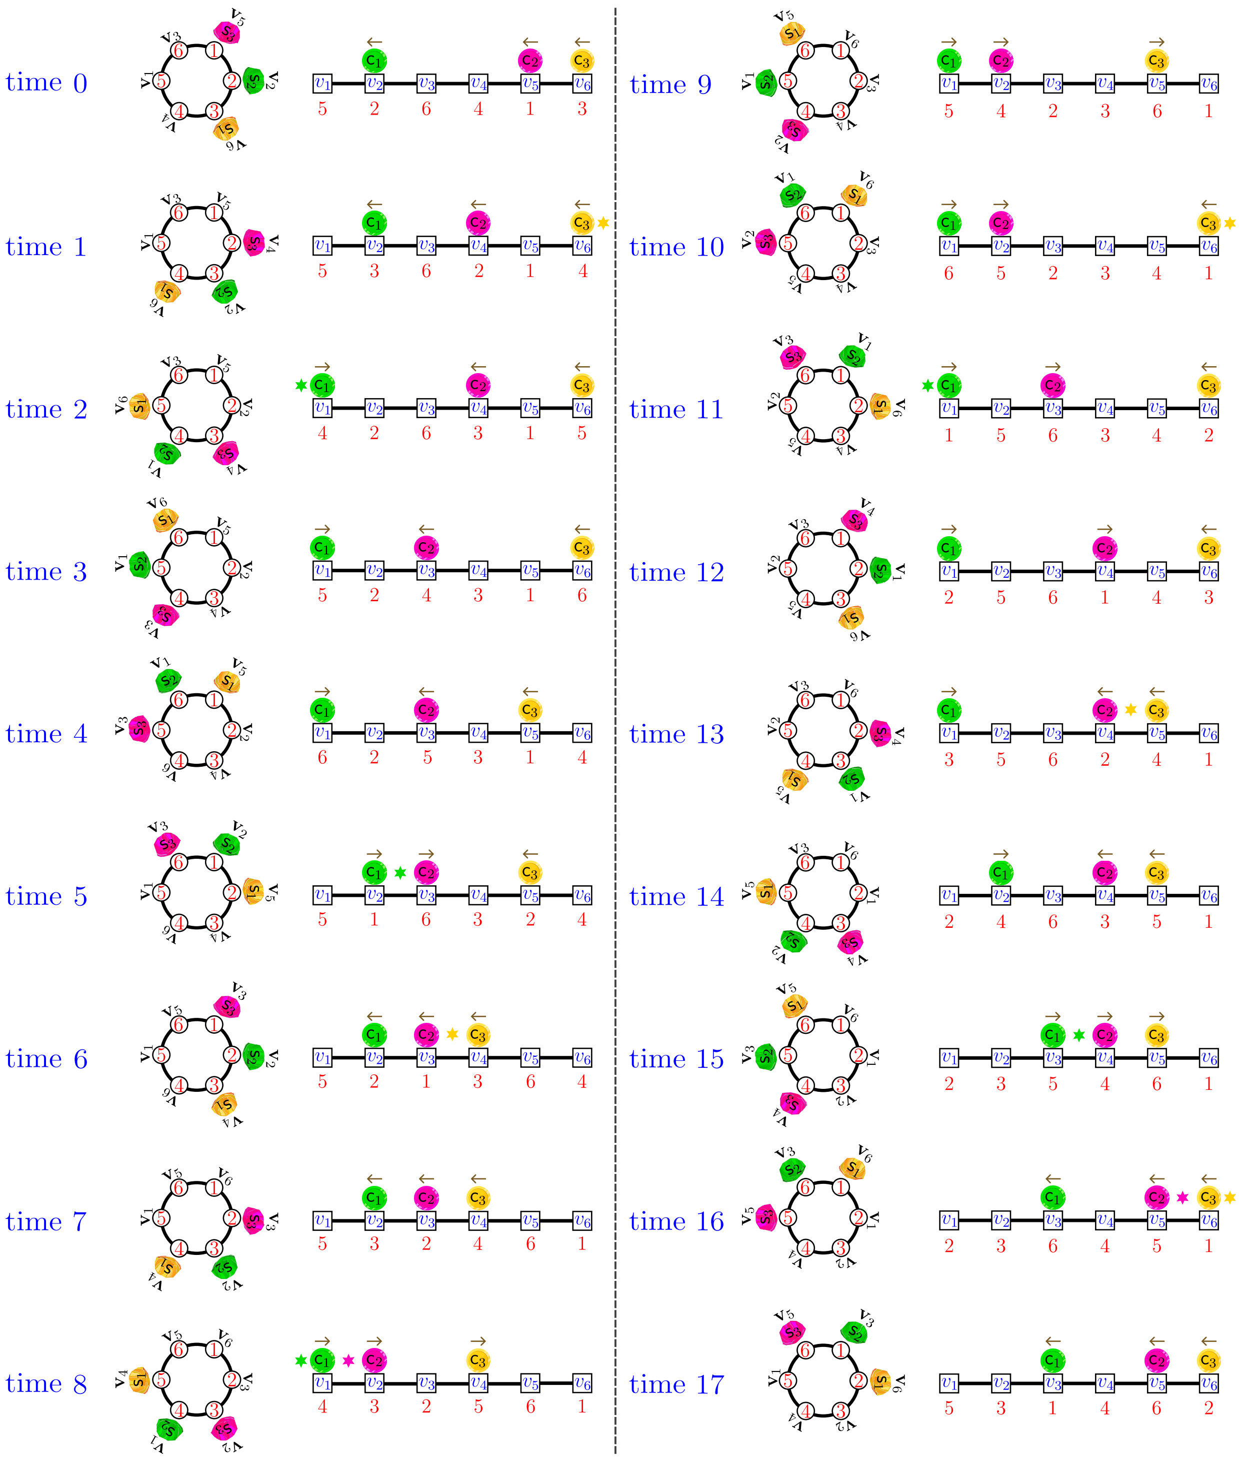

Figure 3: The stones diagrams and coins diagrams of the states in a timeline at times $0, 1, \ldots, 17$. Here, $n = 6$ and $d = 3$. The collisions that occur during the small steps between times $t - 1$ and t are represented by color-coded stars in the coins diagram at time t. Each labeling is depicted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

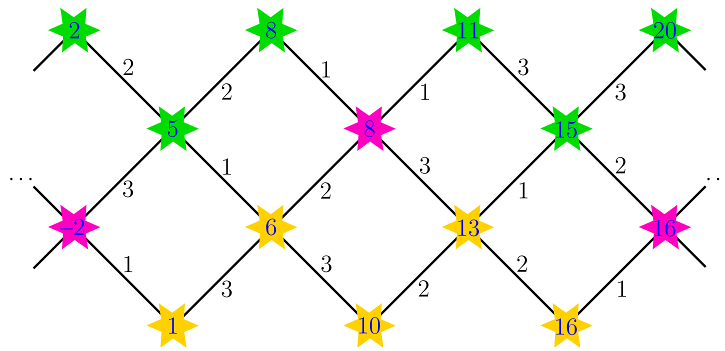

Figure 4: The Hasse diagram $\mathbf{H}_\mathcal{T}$, where $\mathcal{T}$ is the timeline containing the states whose stones diagrams and coins diagrams are shown in Figure 3. We have drawn the Hasse diagram sideways (to save vertical space), so each cover relation $\kappa \lessdot_\mathcal{T} \kappa'$ is drawn with κ to the left of κ'. 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.

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 $\kappa_1 \lessdot_\mathcal{T} \kappa_2$, $\kappa_1 \lessdot_\mathcal{T} \kappa_3$, $\kappa_2 \lessdot_\mathcal{T} \kappa_4$, $\kappa_3 \lessdot_\mathcal{T} \kappa_4$. 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, together with two edges given by cover relations $\kappa'_1 \lessdot_\mathcal{T} \kappa'_2$ and $\kappa'_2 \lessdot_\mathcal{T} \kappa'_3$. Our arguments rely crucially on the following lemma.

Lemma 2 ([4])**.** *In any half-diamond in the Hasse diagram $\mathbf{H}_\mathcal{T}$, the two edges have the same energy. In any diamond in the Hasse diagram $\mathbf{H}_\mathcal{T}$, opposite edges have the same energy.*

For each collision $\kappa \in \mathrm{Coll}_\mathcal{T}$, 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 $\varphi(\kappa_1 \lessdot_\mathcal{T} \kappa_2 \lessdot_\mathcal{T} \cdots \lessdot_\mathcal{T} \kappa_m) = \varphi(\kappa_1) \lessdot_\mathcal{T} \varphi(\kappa_2) \lessdot_\mathcal{T} \cdots \lessdot_\mathcal{T} \varphi(\kappa_m)$. 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 $\mathscr{T} = (\kappa_0 \lessdot_\mathcal{T} \kappa_1 \lessdot_\mathcal{T} \cdots \lessdot_\mathcal{T} \kappa_d)$ 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 $\mathscr{T}$ to be the tuple $\mathcal{E}(\mathscr{T}) = (\varepsilon_1, \ldots, \varepsilon_d)$, where ε_i is the energy of the edge $\kappa_{i-1} \lessdot_\mathcal{T} \kappa_i$; note that $\mathcal{E}(\mathscr{T}) \in \mathsf{Comp}_d(n)$.

Lemma 3 ([4])**.** *Let $\mathcal{T}$ be a timeline, and let $\mathscr{T}$ be a transversal of $\mathbf{H}_\mathcal{T}$. Then $\mathcal{E}(\varphi(\mathscr{T})) = \mathrm{Rot}_{n,d}(\mathcal{E}(\mathscr{T}))$. The period of $\mathbf{H}_\mathcal{T}$ is equal to the size of the orbit of $\mathrm{Rot}_{n,d}$ containing $\mathcal{E}(\mathscr{T})$.*

Proof. The second statement follows from the first because, by Lemma 2, 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 Lemma 2. □

Example 3. Suppose $n = 6$ and $d = 3$. Let $\mathbf{H}_{\mathcal{T}}$ be the Hasse diagram from Figure 4, and let $\mathscr{T} = (\kappa_0 \lessdot_{\mathcal{T}} \kappa_1 \lessdot_{\mathcal{T}} \kappa_2 \lessdot_{\mathcal{T}} \kappa_3)$ be the transversal consisting of the collisions that occur at times $2, 5, 6, 10$. Then $\mathcal{E}(\mathscr{T}) = (2, 1, 3) \in \mathsf{Comp}_3(6)$. The period of $\mathbf{H}_{\mathcal{T}}$ is 3, which is the size of the $\mathrm{Rot}_{6,3}$-orbit containing $(2, 1, 3)$. The transversal $\varphi(\mathscr{T})$ consists of both the collisions that occur at time 8 along with the collisions at times 13 and 16. We have $\mathcal{E}(\varphi(\mathscr{T})) = (1, 3, 2) = \mathrm{Rot}_{6,3}(\mathcal{E}(\mathscr{T}))$. Similarly, $\mathcal{E}(\varphi^2(\mathscr{T})) = (3, 2, 1) = \mathrm{Rot}^2_{6,3}(\mathcal{E}(\mathscr{T}))$. ♢

For $k, t \in \mathbb{Z}$, let $\sigma_t^{(k)} = \mathrm{cyc}^{-k}(\sigma_{t+k})$. 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 $\mathrm{Coll}_{\mathcal{T}^{(k)}}$ that occurs at time t, then there is a collision $\psi_k(\kappa) \in \mathrm{Coll}_{\mathcal{T}}$ that occurs at time $t + k$. The resulting map $\psi_k \colon \mathrm{Coll}_{\mathcal{T}^{(k)}} \to \mathrm{Coll}_{\mathcal{T}}$ is an isomorphism from $(\mathrm{Coll}_{\mathcal{T}^{(k)}}, \leq_{\mathcal{T}^{(k)}})$ to $(\mathrm{Coll}_{\mathcal{T}}, \leq_{\mathcal{T}})$; 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 $\mathcal{T}_\sigma$ for the unique timeline containing the state $(\sigma, 0)$. It follows from Lemma 3 that the energy compositions of the transversals of $\mathbf{H}_{\mathcal{T}_\sigma}$ form a single orbit $\widetilde{\Omega}(\sigma)$ of $\mathrm{Rot}_{n,d}$. 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 $\widetilde{\Omega}(\sigma_0) = \widetilde{\Omega}(\Phi_{n,d}(\sigma_0))$. Thus, we obtain a map

$$\Omega \colon \mathrm{Orb}_{\Phi_{n,d}} \to \mathrm{Orb}_{\mathrm{Rot}_{n,d}}$$

that sends the $\Phi_{n,d}$-orbit containing a labeling μ to $\widetilde{\Omega}(\mu)$. In [4], we prove that this map satisfies the conditions in Proposition 1; we omit the proof here.

3 Toric Promotion on a Forest

Let us briefly mention how the perspective of stones and coins diagrams can be used to prove Theorem 1. Let $G = (V, E)$ be a forest. Let $v_1, \ldots, v_n$ be the vertices of G, and let $\mathbf{v}_1, \ldots, \mathbf{v}_n$ be their replicas. We can represent a labeling $\sigma \in \Lambda_G$ by placing each replica $\mathbf{v}_k$ on the vertex $\sigma(v_k)$ of Cycle_n. We again place a stone on a vertex of Cycle_n to indicate which toggle we are about to apply, and we put a coin on the vertex of G whose replica sits on the stone.

Let T be the connected component of G containing $\sigma^{-1}(1)$; this is the connected component on which the coin always sits. Let t be the number of vertices of T. As we apply the sequence of toggles $\tau_1, \tau_2, \ldots$ (repeating cyclically), the coin will move around to all of the vertices in T. One can show that for all vertices $v_j, v_{j'} \in V$ such that v_j is in T and $j \neq j'$, there is a unique time in the interval $[1, t(n-1)]$ during which $\mathbf{v}_j$ sits on the stone and $\mathbf{v}_{j'}$ sits one space clockwise of the stone. This implies that if v_k is a vertex of degree δ in T, then there are $n - \delta - 1$ times in the interval $[1, t(n-1)]$ when $\mathbf{v}_k$ rides clockwise one space on the stone, and there are $t - \delta - 1$ times in the interval $[1, t(n-1)]$ when $\mathbf{v}_k$ moves counterclockwise one space because the stone slides through it. On the other hand, if v_k is a vertex that is not in T, then $\mathbf{v}_k$ never rides on the stone, and there are t times in the interval $[1, t(n-1)]$ when $\mathbf{v}_k$ moves counterclockwise one space because the stone slides through it. It follows that applying $t(n-1)$ toggles has the effect of rotating the stone and all of the replicas counterclockwise by t. This implies the desired result.

References

[1] R. M. Adin, V. Reiner, and Y. Roichman. "On cyclic descents for tableaux". *Int. Math. Res. Not. IMRN* **2020** (2020), pp. 10231–10276.

[2] C. Defant. "Toric promotion". *Proc. Amer. Math. Soc.* **150** (2022).

[3] C. Defant and N. Kravitz. "Friends and strangers walking on graphs". *Comb. Theory* **1** (2021).

[4] C. Defant, R. Madhukara, and H. Thomas. "Permutoric promotion: gliding globs, sliding stones, and colliding coins". 2023. arXiv:2305.19961.

[5] M. Develin, M. Macauley, and V. Reiner. "Toric partial orders". *Trans. Amer. Math. Soc.* **368** (2016), pp. 2263–2287.

[6] M. D. Haiman. "Dual equivalence with applications, including a conjecture of Proctor". *Discrete Math.* **99** (1992), pp. 79–113.

[7] C. Malvenuto and C. Reutenauer. "Evacuation of labelled graphs". *Discrete Math.* **132** (1994), pp. 137–143.

[8] V. Reiner, D. Stanton, and D. White. "The cyclic sieving phenomenon". *J. Combin. Theory Ser. A* **108** (2004), pp. 17–50.

[9] M.-P. Schützenberger. "Quelques remarques sur une construction de Schensted". *Canad. J. Math.* **13** (1961), pp. 117–128.

[10] M.-P. Schützenberger. "Promotion des morphismes d'ensembles ordonnés". *Discrete Math.* **2** (1972), pp. 73–94.