

Posets from Admissible Coxeter Sequences

Matthew Macauley

Department of Mathematical Sciences
Clemson University
Clemson, South Carolina, USA
macaule@clemson.edu

Henning S. Mortveit

Department of Mathematics
Virginia Bioinformatics Institute
Virginia Tech
Blacksburg, Virginia, USA
henning.mortveit@vt.edu

Submitted: Apr 6, 2010; Accepted: Sep 24, 2011; Published: Oct 3, 2011

Mathematics Subject Classification: 20F55;06A06;05C20

Abstract

We study the equivalence relation on the set of acyclic orientations of an undirected graph Γ generated by source-to-sink conversions. These conversions arise in the contexts of admissible sequences in Coxeter theory, quiver representations, and asynchronous graph dynamical systems. To each equivalence class we associate a poset, characterize combinatorial properties of these posets, and in turn, the admissible sequences. This allows us to construct an explicit bijection from the equivalence classes over Γ to those over Γ' and Γ'' , the graphs obtained from Γ by edge deletion and edge contraction of a fixed cycle-edge, respectively. This bijection yields quick and elegant proofs of two non-trivial results: (i) A complete combinatorial invariant of the equivalence classes, and (ii) a solution to the conjugacy problem of Coxeter elements for simply-laced Coxeter groups. The latter was recently proven by H. Eriksson and K. Eriksson using a much different approach.

1 Overview

Let O_Γ be an acyclic orientation of the undirected graph Γ . A cyclic 1-shift (left) of a linear extension π of O_Γ corresponds to converting a source (the element π_1) of O_Γ into a sink, and this gives rise to an equivalence relation on $\text{Acyc}(\Gamma)$ denoted by \sim_κ . We let $\kappa(\Gamma)$ denote the number of equivalence classes in $\text{Acyc}(\Gamma)$ under \sim_κ , and refer to the equivalence classes as κ -classes.

This paper is organized as follows. After terminology and background in Section 2, we show in Section 3 how to associate a poset to each κ -class, and we characterize structural properties of these posets. This helps us better understand admissible sequences as a whole, culminating in a bijection

$$\Theta: \text{Acyc}(\Gamma)/\sim_\kappa \longrightarrow (\text{Acyc}(\Gamma'_e)/\sim_\kappa) \cup (\text{Acyc}(\Gamma''_e)/\sim_\kappa), \quad (1)$$

where Γ'_e and Γ''_e are the graphs formed by deleting and contracting a cycle-edge e of Γ , respectively. From this bijection, the recursion relation for $\kappa(\Gamma)$ in [10] becomes an immediate corollary, enumerating $\kappa(\Gamma)$ through an evaluation of the Tutte polynomial. In Section 4, we use our bijection to construct a complete invariant of $\text{Acyc}(\Gamma)/\sim_\kappa$, the set of κ -classes of Γ . In Section 5, we review a connection to Coxeter theory, and show how the prior results easily solve the conjugacy problem for Coxeter elements in all simply-laced Coxeter groups, and how $\kappa(\Gamma)$ enumerates the conjugacy classes of Coxeter elements. Finally, in the summary, we briefly discuss how the equivalence relation \sim_κ arises in other areas of mathematics such as sequential dynamical systems, the chip-firing game, and the representation theory of quivers. Throughout the paper, we maintain a running example (that we visit five times) using a six-vertex graph Γ that should enhance the paper's readability and motivate the main ideas.

2 Terminology and Background

Let Γ be an undirected, simple and loop-free graph with vertex set $v[\Gamma] = \{1, 2, \dots, n\}$ and edge set $e[\Gamma]$. Let S_Γ denote the set of total orders (i.e., permutations) of $v[\Gamma]$. Define a relation \sim on S_Γ where $\pi \sim \pi'$ if $\pi = \pi_1\pi_2 \cdots \pi_n$ and $\pi' = \pi'_1\pi'_2 \cdots \pi'_n$ differ by a single adjacent transposition $\pi_i\pi_{i+1} \mapsto \pi_{i+1}\pi_i$ where $\{\pi_i, \pi_{i+1}\} \notin e[\Gamma]$. The reflexive transitive closure of \sim is an equivalence relation on S_Γ denoted by \sim_α . We denote the equivalence class containing π by $[\pi]_\Gamma$, and set

$$S_\Gamma/\sim_\alpha = \{[\pi]_\Gamma \mid \pi \in S_\Gamma\}.$$

This corresponds to partially commutative monoids as defined in [4], but restricted to fixed length permutations over $v[\Gamma]$ and with commutation relations encoded by non-adjacency in the graph Γ . Those familiar with Coxeter theory will recognize the similarity of these equivalence classes and the *commutation classes* [20] of reduced expressions of Coxeter elements.

Orientations of Γ are represented as maps $O_\Gamma: e[\Gamma] \rightarrow v[\Gamma] \times v[\Gamma]$, which may also be viewed as directed graphs. The set of acyclic orientations of Γ is denoted by $\text{Acyc}(\Gamma)$, and we set $\alpha(\Gamma) = |\text{Acyc}(\Gamma)|$. Each acyclic orientation defines a partial ordering on $v[\Gamma]$ where the covering relations are $i \leq_{O_\Gamma} j$ if $\{i, j\} \in e[\Gamma]$ and $O_\Gamma(\{i, j\}) = (i, j)$. The set of linear extensions of O_Γ contains precisely the permutations $\pi \in S_\Gamma$ such that if $i \leq_{O_\Gamma} j$, then i precedes j in π . Through the ordering of $v[\Gamma]$, every permutation $\pi \in S_\Gamma$ induces a canonical linear order on $v[\Gamma]$. Moreover, each permutation $\pi \in S_\Gamma$ induces an acyclic orientation $O_\Gamma^\pi \in \text{Acyc}(\Gamma)$ defined by $O_\Gamma^\pi(\{i, j\}) = (i, j)$ if i precedes j in π and $O_\Gamma^\pi(\{i, j\}) = (j, i)$ otherwise. The canonical bijection

$$f_\Gamma: S_\Gamma/\sim_\alpha \rightarrow \text{Acyc}(\Gamma), \quad f_\Gamma([\pi]_\Gamma) = O_\Gamma^\pi, \quad (2)$$

identifies equivalence classes and acyclic orientations, and thus the number of equivalence classes under \sim_α is $\alpha(\Gamma)$.

For $O_\Gamma \in \text{Acyc}(\Gamma)$ and $e = \{v, w\} \in e[\Gamma]$, let $O_\Gamma^{\rho(e)}$ be the orientation of Γ obtained from O_Γ by reversing the orientation of the edge e . Let Γ'_e and Γ''_e denote the graphs obtained from Γ by deletion and contraction (see, e.g., [13, p. 415]) of e , respectively, and let $O_{\Gamma'}$ and $O_{\Gamma''}$ denote the orientations of O_Γ inherited under these operations. (Since our graphs are assumed to be loop-free, when we contract an edge $\{v, w\}$, we remove the resulting loop.) The bijection

$$\beta_e: \text{Acyc}(\Gamma) \longrightarrow \text{Acyc}(\Gamma'_e) \cup \text{Acyc}(\Gamma''_e) \tag{3}$$

defined by

$$O_\Gamma \xrightarrow{\beta_e} \begin{cases} O_{\Gamma'} , & O_\Gamma^{\rho(e)} \notin \text{Acyc}(\Gamma) , \\ O_{\Gamma'} , & O_\Gamma^{\rho(e)} \in \text{Acyc}(\Gamma) \text{ and } O_\Gamma(e) = (v, w) , \\ O_{\Gamma''} , & O_\Gamma^{\rho(e)} \in \text{Acyc}(\Gamma) \text{ and } O_\Gamma(e) = (w, v) , \end{cases} \tag{4}$$

is well-known, and shows that one may compute $\alpha(\Gamma)$ through the recursion relation

$$\alpha(\Gamma) = \alpha(\Gamma'_e) + \alpha(\Gamma''_e) ,$$

valid for any $e \in e[\Gamma]$. It basically removes the edge $e = \{v, w\}$ if it cannot be contracted, and otherwise, it either contracts or removes it depending on its orientation. We illustrate this with the following example, which we will revisit four more times throughout this article.

Example 2.1. *For an explicit example of β_e , see Figure 1, which shows three acyclic orientations of the same graph Γ , and a fixed edge $e = \{v, w\}$. (The vertices v' , w' , and z , which will be referred to later, are only labeled once for clarity.) Call these orientations O_Γ^a , O_Γ^b , and O_Γ^c , respectively. The map β_e removes edge e from O_Γ^a because contracting it would result in a directed cycle. Neither O_Γ^b nor O_Γ^c have a directed path from v to w other than the edge (v, w) , so in both cases, contracting e would give an acyclic orientation of Γ''_e . By the definition of β_e in (4), β_e removes e in O_Γ^b , and contracts it in O_Γ^c .*

Via the bijection in (2), it is clear that mapping $\pi = \pi_1\pi_2 \cdots \pi_n \in [\pi]_\Gamma$ to $\pi_2 \cdots \pi_n\pi_1$ corresponds precisely to converting the source vertex π_1 in O_Γ^π into a sink. We call such a conversion a *source-to-sink operation*, or a *click*. Two orientations $O_\Gamma, O'_\Gamma \in \text{Acyc}(\Gamma)$ where O_Γ can be transformed into O'_Γ by a sequence of clicks are said to be click-related. We write this as $\mathbf{c}(O_\Gamma) = O'_\Gamma$ where $\mathbf{c} = c_1c_2 \cdots c_k$ with $c_i \in v[\Gamma]$. To clarify notation, we mean

$$\mathbf{c}(O_\Gamma) = c_k(c_{k-1}(\cdots c_2(c_1(O_\Gamma)))) .$$

Such as sequence \mathbf{c} is called an *admissible sequence*, or a *click-sequence*. The former term comes from the representation theory of quivers [1, 8], but we will usually stick to the latter due to brevity, the overuse of the term “admissible sequence” throughout mathematics, and the convenience of “click” doubling as a verb. It is straightforward to verify that this click-relation is an equivalence relation on $\text{Acyc}(\Gamma)$, and we also refer to click-related

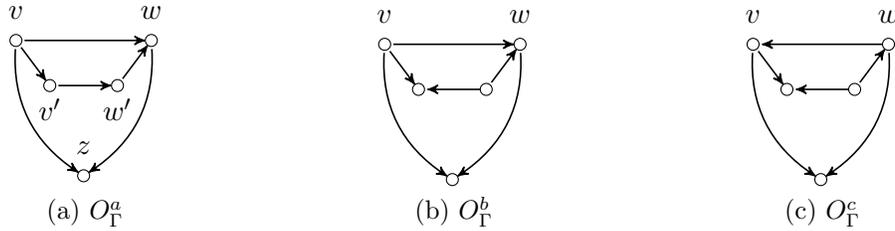


Figure 1: An example of the map β_e applied to three acyclic orientations of a graph Γ . If contracting the edge $e = \{v, w\}$ would introduce a directed cycle (as in O_Γ^a), then we must delete it. Otherwise, we can either delete or contract it, so we pick the convention that we delete it if it is oriented (v, w) (as in O_Γ^b), and contract it if it is oriented (w, v) (as in O_Γ^c).

acyclic orientations as κ -equivalent. Clearly, and as pointed out by V. Reiner [15, p. 309], one may also approach this in the setting of total orders on $v[\Gamma]$ by identifying elements that differ by (i) flips of adjacent elements not connected in Γ and (ii) cyclic shifts. However, for our purposes, approaching this at the level of acyclic orientations seems more natural in light of the bijection (2).

3 Constructing the Bijection Θ

The bijection $\beta_e: \text{Acyc}(\Gamma) \rightarrow \text{Acyc}(\Gamma'_e) \cup \text{Acyc}(\Gamma''_e)$ in (3) does not extend to a well-defined map on κ -classes, i.e., $\text{Acyc}(\Gamma)/\sim_\kappa \rightarrow (\text{Acyc}(\Gamma'_e)/\sim_\kappa) \cup (\text{Acyc}(\Gamma''_e)/\sim_\kappa)$. Thus, we need to take a different approach to construct our bijection. An edge e of an undirected graph Γ is a *bridge* if removing e increases the number of connected components of Γ . An edge that is not a bridge is a *cycle-edge*, or equivalently, an edge e is a cycle-edge if it is contained in a cycle traversing e precisely once.

Throughout, we will let $e = \{v, w\}$ denote a fixed cycle-edge of the connected graph Γ , and, for ease of notation, we set $\Gamma' = \Gamma'_e$ and $\Gamma'' = \Gamma''_e$. Recall that for $O_\Gamma \in \text{Acyc}(\Gamma)$ we let $O_{\Gamma'}$ and $O_{\Gamma''}$ denote the inherited orientations of Γ' and Γ'' . Notice that $O_{\Gamma'}$ is always acyclic, while $O_{\Gamma''}$ is acyclic if and only if there is no directed path with endpoints v and w in $O_{\Gamma''}$. Finally, we let $[O_\Gamma]$ denote the κ -class containing O_Γ .

The interval $[a, b]$ of a poset \mathcal{P} (where $a \leq b$) is the subposet consisting of all $c \in \mathcal{P}$ such that $a \leq c \leq b$. Viewing a finite poset \mathcal{P} as a directed graph $D_{\mathcal{P}}$, the interval $[a, b]$ contains precisely the vertices that lie on a directed path from a to b , and thus is a vertex-induced subgraph of $D_{\mathcal{P}}$. By assumption, $e[\Gamma]$ contains $\{v, w\}$, so for all $O_\Gamma \in \text{Acyc}(\Gamma)$ either $v \leq_{O_\Gamma} w$ or $w \leq_{O_\Gamma} v$. In this section, we will study the interval $[v, w]$ in the poset O_Γ (when $v \leq_{O_\Gamma} w$) and its behavior under clicks.

Definition 3.1. Let $\text{Acyc}_{\leq}(\Gamma)$ be the set of acyclic orientations of vertex-induced subgraphs of Γ . Define the map

$$\mathcal{I}: \text{Acyc}(\Gamma) \rightarrow \text{Acyc}_{\leq}(\Gamma)$$

by $\mathcal{I}(O_\Gamma) = [v, w]$ if $v \leq_{O_\Gamma} w$, and by $\mathcal{I}(O_\Gamma) = \emptyset$ (the null graph) otherwise. We will refer to $\mathcal{I}(O_\Gamma)$ as the vw -interval of O_Γ .

Elements of $\text{Acyc}(\Gamma)$ can be thought of as posets over $v[\Gamma]$, and elements of $\text{Acyc}_{\leq}(\Gamma)$ can be thought of as certain subposets of these, though they need not be induced (because two vertices on a directed path in Γ need not be on a directed path in an induced subgraph of Γ). Through a slight abuse of notation, we will at times refer to $\mathcal{I}(O_\Gamma)$ as a poset, a directed graph, or a subset of $v[O_\Gamma]$. In this last case, it is understood that the relations are inherited from O_Γ .

Let P be an undirected path in Γ of length- k , i.e., $P = v_0, v_1, \dots, v_{k-1}, v_k$ where $\{v_{i-1}, v_i\} \in e[\Gamma]$ for $i = 1, \dots, k$. Define the function

$$\nu_P: \text{Acyc}(\Gamma) \longrightarrow \mathbb{Z}, \tag{5}$$

where $\nu_P(O_\Gamma)$ is the number of edges in Γ of the form $\{v_{i-1}, v_i\}$ oriented as (v_{i-1}, v_i) in O_Γ , minus the number of edges oriented as (v_i, v_{i-1}) . If P is a cycle (i.e., $v_0 = v_k$), ν_P is preserved under clicks, and thus in this case, it extends to a map $\nu_P^*: \text{Acyc}(\Gamma)/\sim_\kappa \longrightarrow \mathbb{Z}$. In [18], J.-Y. Shi defines this function for Coxeter graphs containing a single cycle, referring to it as Coleman's ν -function (see [5]). The definition given here is more general, and will allow us to extend Shi's characterization of conjugacy classes to include all simply-laced Coxeter groups.

Example 3.2. *Continuing with Example 2.1, consider the three orientations in Figure 1, whose vw -intervals are the following:*

$$\mathcal{I}(O_\Gamma^a) = \{v, w, v', w'\}, \quad \mathcal{I}(O_\Gamma^b) = \{v, w\}, \quad \mathcal{I}(O_\Gamma^c) = \emptyset.$$

Next, consider the undirected (but oriented) path $P = v, v', w', w, v$ (a cycle) and the corresponding map ν_P , as defined in (5). It is easy to check that

$$\nu_P(O_\Gamma^a) = 2, \quad \nu_P(O_\Gamma^b) = 0, \quad \nu_P(O_\Gamma^c) = 2.$$

We conclude that O_Γ^b cannot be κ -equivalent to either O_Γ^a or O_Γ^c . Finally, if we consider the undirected path $Q = v, z, w, v$, we have

$$\nu_Q(O_\Gamma^a) = -1, \quad \nu_Q(O_\Gamma^b) = -1, \quad \nu_Q(O_\Gamma^c) = 1.$$

Therefore, $O_\Gamma^a \not\sim_\kappa O_\Gamma^c$, and hence all three of these orientations lie in distinct κ -classes.

As we will see in Section 4, when taken over all cycles of Γ , the ν -function is actually a complete invariant, i.e., $O \sim_\kappa O'$ if and only iff $\nu_C(O) = \nu_C(O')$ for all cycles C in Γ . First, we need to establish a series of structural results about the vw -interval. Since $\{v, w\} \in e[\Gamma]$, every κ -class contains at least one orientation O_Γ with $v \leq_{O_\Gamma} w$, and thus there is at least one element O_Γ in each κ -class with $\mathcal{I}(O_\Gamma) \neq \emptyset$. As the next result shows, this (non-empty) choice of vw -interval is independent of the choice of representative from $[O_\Gamma]$, meaning that there is a well-defined notion of the vw -interval of a κ -equivalence class. We formalize this by extending the map $\mathcal{I}: \text{Acyc}(\Gamma) \rightarrow \text{Acyc}_{\leq}(\Gamma)$ to a map $\mathcal{I}^*: \text{Acyc}(\Gamma)/\sim_\kappa \rightarrow \text{Acyc}_{\leq}(\Gamma)$.

Proposition 3.3. *The map \mathcal{I} can be extended to a map*

$$\mathcal{I}^* : \text{Acyc}(\Gamma) / \sim_\kappa \longrightarrow \text{Acyc}_{\leq}(\Gamma) \quad \text{by} \quad \mathcal{I}^*([O_\Gamma]) = \mathcal{I}(O_\Gamma^1),$$

where O_Γ^1 is any element of $[O_\Gamma]$ for which $\mathcal{I}(O_\Gamma^1) \neq \emptyset$.

Proof. It suffices to prove that \mathcal{I}^* is well-defined. Consider $O_\Gamma^1 \sim_\kappa O_\Gamma^2$ with $v \leq_{O_\Gamma^i} w$ for $i = 1, 2$. Clearly, $\mathcal{I}(O_\Gamma^1)$ and $\mathcal{I}(O_\Gamma^2)$ contain v and w , so suppose that $a \in \mathcal{I}(O_\Gamma^1) \setminus \{v, w\}$. Then a lies on a directed path P' from v to w in O_Γ^1 , of length $k \geq 2$ (i.e., P' traverses at least 2 edges). Let P be the cycle formed by adding vertex v to the end of P' . Clearly $\nu_P(O_\Gamma^1) = k - 1$ since $O_\Gamma^1(e) = (v, w)$.

By assumption, $O_\Gamma^2 \in [O_\Gamma^1]$ with $v \leq_{O_\Gamma^2} w$. Since ν_P is constant on $[O_\Gamma^1]$ it follows from $\nu_P(O_\Gamma^1) = k - 1 = \nu_P(O_\Gamma^2)$ that every edge of P' is oriented identically in O_Γ^1 and O_Γ^2 , and hence that every directed path P' in O_Γ^1 is contained in O_Γ^2 as well. Therefore, $a \in \mathcal{I}(O_\Gamma^2)$, and the reverse inclusion follows by an identical argument. \square

In light of Proposition 3.3, we define the vw -interval of a κ -class $[O_\Gamma]$ to be $\mathcal{I}^*([O_\Gamma])$. The vw -interval will be central in understanding properties of click-sequences. First, we make a simple observation without proof; it also appears in [19] in the context of admissible sequences in Coxeter theory.

Proposition 3.4. *Let $O_\Gamma \in \text{Acyc}(\Gamma)$, let $\mathbf{c} = c_1 c_2 \cdots c_m$ be an associated click-sequence, and consider any directed edge (v_1, v_2) in O_Γ . Then the occurrences of v_1 and v_2 in \mathbf{c} alternate, with v_1 appearing first.*

Because $\{v, w\} \in e[\Gamma]$, we can say more about the vertices in $\mathcal{I}(O_\Gamma)$ that appear between successive instances of v and w in a click-sequence.

Proposition 3.5. *Let $O_\Gamma \in \text{Acyc}(\Gamma)$, and let $\mathbf{c} = c_1 c_2 \cdots c_m$ be an associated click-sequence that contains every vertex of $\mathcal{I}(O_\Gamma)$ at least once and with $c_1 = v$. Then every vertex of $\mathcal{I}(O_\Gamma)$ appears in \mathbf{c} before any vertex in $\mathcal{I}(O_\Gamma)$ appears twice.*

Proof. The proof is by contradiction. Assume the statement is false, and let $a \in \mathcal{I}(O_\Gamma)$ be the first vertex whose second instance in \mathbf{c} occurs before the first instance of some other vertex $z \in \mathcal{I}(O_\Gamma)$. If $a \neq v$, then a is not a source in O_Γ , and there exists a directed edge (a', a) . By Proposition 3.4, a' must appear in \mathbf{c} before the first instance of a , but also between the two first instances of a . This is impossible, because a was chosen to be the first vertex appearing twice in \mathbf{c} . That only leaves $a = v$, and v must appear twice before the first instance of w . However, this contradicts the statement of Proposition 3.4 because $\{v, w\} \in e[\Gamma]$. \square

The next result shows that for any click-sequence \mathbf{c} that contains every element in $\mathcal{I}(O_\Gamma)$ precisely once, we may assume without loss of generality that the vertices in $\mathcal{I}(O_\Gamma)$ appear consecutively.

Proposition 3.6. *Let $O_\Gamma \in \text{Acyc}(\Gamma)$ be an acyclic orientation with $v \leq_{O_\Gamma} w$. If $\mathbf{c} = c_1 c_2 \cdots c_m$ is an associated click-sequence containing precisely one instance of w , and no subsequent instances of vertices from $\mathcal{I}(O_\Gamma)$, then there exists a click-sequence $\mathbf{c}' = c'_1 c'_2 \cdots c'_m$ such that (i) there exists an interval $[p, q]$ of \mathbb{N} with $c'_j \in \mathcal{I}(O_\Gamma)$ iff $p \leq j \leq q$, and (ii) $\mathbf{c}(O_\Gamma) = \mathbf{c}'(O_\Gamma)$.*

Proof. We prove the proposition by constructing a desired click-sequence \mathbf{c}'' from \mathbf{c} through a series of transpositions where each intermediate click-sequence \mathbf{c}' satisfies $\mathbf{c}(O_\Gamma) = \mathbf{c}'(O_\Gamma)$. Such transpositions are said to have property T .

Let $I = \mathcal{I}(O_\Gamma)$, and let A be the set of vertices in $I^c = \mathbf{v}[\Gamma] \setminus I$ that lie on a directed path in O_Γ to a vertex in I (vertices *above* I), and let B be the set of vertices that lie on a directed path in O_Γ from a vertex in I (vertices *below* I). Let C be the complement of $I \cup A \cup B$. Two vertices $c_i, c_j \in A \cup B$ with $i < j$ for which there is no element $c_k \in A \cup B$ with $i < k < j$ are said to be *tight*. We will investigate when transpositions of tight vertices in a click-sequence \mathbf{c} of O_Γ has property T , and we will see that this is always the case if $c_i \in B$ and $c_j \in A$. Consider the intermediate acyclic orientation after applying successive clicks $c_1 c_2 \cdots c_{i-1}$ to O_Γ . Obviously, c_i is a source. At this point, if c_j were not a source, then there would be an adjacent vertex $a \in A$ with the edge $\{a, c_j\}$ oriented (a, c_j) . For c_j to be clicked as usual (i.e., as a source), a must be clicked first, but this would break the assumption that c_i and c_j are tight. Therefore, c_i and c_j are both sources at this intermediate step, and so the vertices c_i, c_{i+1}, \dots, c_j are an independent set of sources, and may be permuted in any manner without changing the image of the click sequence. Therefore, the transposition of c_i and c_j in \mathbf{c} has property T , as claimed. By iteratively transposing tight pairs in \mathbf{c} , we can construct a click-sequence with the property that every vertex in A comes before every vertex in B . In light of this, we may assume without loss of generality that \mathbf{c} has this property.

The next step is to show that we can move all vertices in A before v , and all vertices in B after w via transpositions having property T . Let a be the first vertex in A appearing after v in the click sequence \mathbf{c} . We claim that the transposition moving a to the position directly preceding v has property T . This is immediate from the observation that when v is to be clicked, a is a source as well, by the definition of A , thus it may be clicked before v , without preventing subsequent clicks of vertices up until the original position of a . Therefore, we may one-by-one move the vertices in A that are between v and w , in front of v . An analogous argument shows that we may move the vertices in B that appear before w to a position directly following w . In the resulting click-sequence \mathbf{c}' , the only vertices between v and w are either in I or C . The subgraph of the directed graph O_Γ induced by C is a disjoint union of weakly connected components, and none of the vertices are adjacent to I . By definition of A and B , there cannot exist a directed edge (c, a) or (b, c) , where $a \in A$, $b \in B$, and $c \in C$. Thus for each weakly connected component of C , the vertices in the component can be moved within \mathbf{c}' , preserving their relative order, to a position either (i) directly after the vertices in A and before v , or (ii) directly after w and before the vertices of B . Call this resulting click-sequence \mathbf{c}'' . As we just argued, all the transpositions occurring in the rearrangement $\mathbf{c} \mapsto \mathbf{c}''$ has property T , and \mathbf{c}'' contains all of the vertices in I in consecutive order, and this proves the result. \square

We remark that the last two results together imply that for the interval $[p, q]$ in the statement of Proposition 3.6, $c_p = v$, $c_q = w$, and the sequence $c_p c_{p+1} \cdots c_q$ contains every vertex in $\mathcal{I}(O_\Gamma)$ precisely once. A simple induction argument implies the following.

Corollary 3.7. *Suppose that $O_\Gamma \in \text{Acyc}(\Gamma)$ with $v \leq_{O_\Gamma} w$, and let $\mathbf{c} = c_1 c_2 \cdots c_m$ be a click-sequence where w appears exactly k times, and no vertex from $\mathcal{I}(O_\Gamma)$ appears in \mathbf{c} after the last instance of w . Then there exists a click-sequence $\mathbf{c}' = c'_1 c'_2 \cdots c'_m$ such that (i) there are k disjoint intervals $[p_i, q_i]$ of \mathbb{N} such that $c_j \in \mathcal{I}(O_\Gamma)$ iff $p_i \leq j \leq q_i$ for some i , and (ii) $\mathbf{c}(O_\Gamma) = \mathbf{c}'(O_\Gamma)$.*

Proof. The argument is by induction on k . When $k = 1$, the statement is simply Proposition 3.6. Suppose the statement holds for all $k \leq N$, for some $N \in \mathbb{N}$, and let \mathbf{c} be a click-sequence containing $N + 1$ instances of w . Let c_ℓ be the second instance of v in \mathbf{c} , and consider the two click-sequences $\mathbf{c}_i := c_1 c_2 \cdots c_{\ell-1}$ and $\mathbf{c}_f := c_\ell c_{\ell+1} \cdots c_m$. By Proposition 3.6, there exists an interval $[p_1, q_1]$ with $p_1 < q_1 < \ell$, and by the induction hypothesis, there exists k intervals $[p_2, q_2], \dots, [p_{k+1}, q_{k+1}]$ with $\ell \leq p_2 < q_2 < \cdots < p_{k+1} < q_{k+1}$ such that if $c_j \in \mathcal{I}(O_\Gamma)$, then $p_i \leq j \leq q_i$ for some $i = 1, \dots, k + 1$. \square

Let $\eta_e: \text{Acyc}(\Gamma) \rightarrow \text{Acyc}(\Gamma')$ be the canonical map that sends O_Γ to $O_{\Gamma'}$. This extends naturally to a map $\eta_e^*: \text{Acyc}(\Gamma)/\sim_\kappa \rightarrow \text{Acyc}(\Gamma')/\sim_\kappa$ between κ -classes. Define

$$\mathcal{I}_e^*: \text{Acyc}(\Gamma')/\sim_\kappa \rightarrow \text{Acyc}_{\leq}(\Gamma)$$

by $\mathcal{I}_e^*([O_{\Gamma'}]) = \mathcal{I}(O_\Gamma^1)$ for any $O_\Gamma^1 \in [O_\Gamma]$ such that $\eta_e^*([O_\Gamma]) = [O_{\Gamma'}]$ with $|\mathcal{I}(O_\Gamma^1)| \geq 3$, and $\mathcal{I}_e^*([O_{\Gamma'}]) = (v, w)$ (that is, the subgraph induced by $\{v, w\}$) if no such acyclic orientation O_Γ^1 exists. The following result relates the vw -intervals of $\text{Acyc}(\Gamma)/\sim_\kappa$ and $\text{Acyc}(\Gamma')/\sim_\kappa$ through a commutative diagram involving \mathcal{I}^* and \mathcal{I}_e^* . An explicit example immediately follows the proof.

Proposition 3.8. *The map \mathcal{I}_e^* is well-defined, and the diagram*

$$\begin{array}{ccc} \text{Acyc}(\Gamma)/\sim_\kappa & \xrightarrow{\mathcal{I}^*} & \text{Acyc}_{\leq}(\Gamma) \\ \eta_e^* \downarrow & \nearrow \mathcal{I}_e^* & \\ \text{Acyc}(\Gamma')/\sim_\kappa & & \end{array}$$

commutes.

Proof. Let $[O_{\Gamma'}] \in \text{Acyc}(\Gamma')/\sim_\kappa$. If there is at most one orientation $O_\Gamma \in \text{Acyc}(\Gamma)$ such that $|\mathcal{I}(O_\Gamma)| \geq 3$ and $\eta_e(O_\Gamma) \in [O_{\Gamma'}]$, or if all orientations of the form O_Γ^1 in the definition of \mathcal{I}_e^* are κ -equivalent, then both statements of the proposition are clear. Assume therefore that there are acyclic orientations $O_\Gamma^\pi, O_\Gamma^\sigma \in \text{Acyc}(\Gamma)$ with $O_\Gamma^\pi \approx_\kappa O_\Gamma^\sigma$, but $\eta_e^*([O_\Gamma^\pi]) = \eta_e^*([O_\Gamma^\sigma])$ and $|\mathcal{I}(O_\Gamma^\pi)|, |\mathcal{I}(O_\Gamma^\sigma)| \geq 3$. It suffices to prove that in this case,

$$\mathcal{I}(O_\Gamma^\pi) = \mathcal{I}(O_\Gamma^\sigma). \tag{6}$$

This is equivalent to showing that the set of vw -paths (directed paths from v to w) in $O_{\Gamma'}^{\pi}$, is the same as the set of vw -paths in $O_{\Gamma'}^{\sigma}$. From this it will also follow that the diagram commutes. By assumption, both of these orientations contain at least one vw -path. We will consider separately the cases when these orientations share or do not share a common vw -path.

Case 1: $O_{\Gamma'}^{\pi}$, and $O_{\Gamma'}^{\sigma}$, share no common vw -path. Let P_1 be a length- k_1 vw -path in $O_{\Gamma'}^{\pi}$, and let P_2 be a length- k_2 vw -path in $O_{\Gamma'}^{\sigma}$. Suppose that in $O_{\Gamma'}^{\pi}$, there are k_2^+ edges along P_2 oriented from v to w , and k_2^- edges oriented from w to v . Likewise, suppose that in $O_{\Gamma'}^{\sigma}$, there are k_1^+ edges along P_1 oriented from v to w , and k_1^- edges oriented from w to v . If $C = P_1P_2^{-1}$ (the cycle formed by traversing P_1 followed by P_2 in reverse), then

$$\nu_C(O_{\Gamma'}^{\pi}) = k_1^+ + k_1^- + k_2^- - k_2^+, \quad \nu_C(O_{\Gamma'}^{\sigma}) = k_1^+ - k_1^- - k_2^- - k_2^+.$$

Equating these values yields $k_1^- + k_2^- = 0$, and since these are non-negative integers, $k_1^- = k_2^- = 0$. We conclude that P_1 is a vw -path in $O_{\Gamma'}^{\sigma}$, and P_2 is a vw -path in $O_{\Gamma'}^{\pi}$, contradicting the assumption that $O_{\Gamma'}^{\pi}$, and $O_{\Gamma'}^{\sigma}$, share no common vw -paths.

Case 2: $O_{\Gamma'}^{\pi}$, and $O_{\Gamma'}^{\sigma}$, share a common vw -path P_1 , say of length k_1 . If these are the only vw -paths, we are done. Otherwise, assume without loss of generality that P_2 is another vw -path in $O_{\Gamma'}^{\pi}$, say of length k_2 . Then if $C = P_1P_2^{-1}$, we have $\nu_C(O_{\Gamma'}^{\pi}) = k_1 - k_2$, and hence $\nu_C(O_{\Gamma'}^{\sigma}) = k_1 - k_2$. Therefore, P_2 is a vw -path in $O_{\Gamma'}^{\sigma}$, as well. Because P_2 was arbitrary, we conclude that $O_{\Gamma'}^{\pi}$, and $O_{\Gamma'}^{\sigma}$, share the same set of vw -paths. Since Case 1 is impossible, we have established (6), and the proof is complete. \square

Example 3.9. Consider the orientations $O_{\Gamma}^a, O_{\Gamma}^b \in \text{Acyc}(\Gamma)$ from our running example (see Figure 1). Since the vw -intervals of O_{Γ}^a and O_{Γ}^b are non-empty (see Example 3.2),

$$\mathcal{I}^*([O_{\Gamma}^a]) = \mathcal{I}(O_{\Gamma}^a) = \{v, w, v', w'\}, \quad \mathcal{I}^*([O_{\Gamma}^b]) = \mathcal{I}(O_{\Gamma}^b) = \{v, w\}.$$

The natural map η_e^* simply removes the edge $\{v, w\}$, i.e.,

$$\eta_e^*([O_{\Gamma}^a]) = [O_{\Gamma'}^a], \quad \eta_e^*([O_{\Gamma}^b]) = [O_{\Gamma'}^b].$$

Finally, Proposition 3.8 guarantees a well-defined map \mathcal{I}_e^* satisfying $\mathcal{I}_e^* \circ \eta_e^* = \mathcal{I}^*$, and thus

$$\mathcal{I}_e^*([O_{\Gamma'}^a]) = \mathcal{I}^*([O_{\Gamma}^a]) = \{v, w, v', w'\}, \quad \mathcal{I}_e^*([O_{\Gamma'}^b]) = \mathcal{I}^*([O_{\Gamma}^b]) = \{v, w\}.$$

This is shown in Figure 2, though note that the domains are the actual κ -classes containing the given orientations, not the orientations themselves.

Let $O_{\Gamma} \in \text{Acyc}(\Gamma)$ and assume $I = \mathcal{I}(O_{\Gamma})$ is non-empty. We write Γ_I for the graph formed from Γ by contracting all vertices in I to a single vertex, which we denote by V_I . Note that if I only contains v and w then $\Gamma_I = \Gamma_e''$. Moreover, O_{Γ} gives rise to an orientation O_{Γ_I} of Γ_I , and this orientation is clearly acyclic.

Proposition 3.10. Let $O_{\Gamma}^1, O_{\Gamma}^2 \in \text{Acyc}(\Gamma)$ and assume $\mathcal{I}(O_{\Gamma}^1) = \mathcal{I}(O_{\Gamma}^2)$. If $O_{\Gamma}^1 \approx_{\kappa} O_{\Gamma}^2$ then $[O_{\Gamma_I}^1] \approx_{\kappa} [O_{\Gamma_I}^2]$.

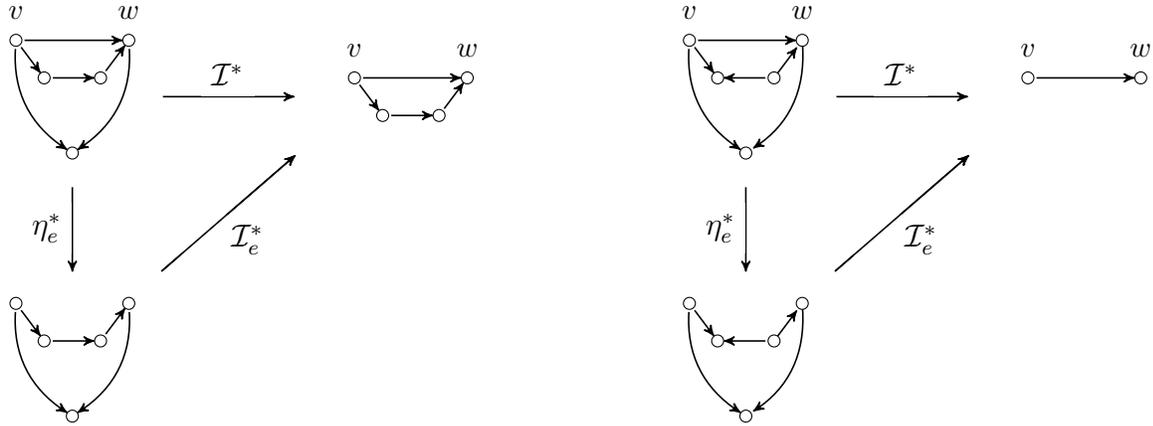


Figure 2: An explicit example of the commutative diagram relating the maps η_e^* , \mathcal{I}^* , and \mathcal{I}_e^* from Proposition 3.8. The domain of these maps are actually the sets of κ -equivalence classes, e.g., $[O_\Gamma^a]$ (left) and $[O_\Gamma^b]$ (right), but they are shown acting on actual orientations (O_Γ^a and O_Γ^b) for clarity.

Proof. We prove the contrapositive statement. Set $I = \mathcal{I}(O_\Gamma^1)$, suppose $|I| = k$, and let $v_1 v_2 \cdots v_k$ be a linear extension of I . For any click-sequence \mathbf{c}_I between two acyclic orientations $O_{\Gamma_I}^1$ and $O_{\Gamma_I}^2$ in $\text{Acyc}(\Gamma_I)$, let \mathbf{c} be the click-sequence formed by replacing every occurrence of V_I in \mathbf{c}_I by the sequence $v_1 \cdots v_k$. Then $\mathbf{c}(O_\Gamma^1) = O_\Gamma^2$ and $O_\Gamma^1 \sim_\kappa O_\Gamma^2$ as claimed. \square

We can now utilize the results on poset structure just developed to establish a bijection

$$\Theta: \text{Acyc}(\Gamma)/\sim_\kappa \longrightarrow (\text{Acyc}(\Gamma'_e)/\sim_\kappa \cup \text{Acyc}(\Gamma''_e)/\sim_\kappa),$$

valid for any cycle-edge e . For $[O_\Gamma] \in \text{Acyc}(\Gamma)/\sim_\kappa$, let O_Γ^π denote an orientation in $[O_\Gamma]$ such that $\pi = v\pi_2 \cdots \pi_n$ and $w = \pi_i$ for i minimal. We define Θ by

$$[O_\Gamma] \xrightarrow{\Theta} \begin{cases} [O_{\Gamma''}^\pi], & \exists O_\Gamma^\pi \in [O_\Gamma] \text{ with } \pi = vw\pi_3 \cdots \pi_n \\ [O_{\Gamma'}^\pi], & \text{otherwise.} \end{cases} \quad (7)$$

Note that $[O_\Gamma]$ is mapped into $\text{Acyc}(\Gamma'')/\sim_\kappa$ if and only if the only vertices in $\mathcal{I}_e^*([O_\Gamma])$ are v and w . Since κ -equivalence over Γ implies κ -equivalence over Γ' , Θ does not depend on the choice of π , and thus is well-defined. We continue our running example below to illustrate this.

Example 3.11. *The orientations O_Γ^a and O_Γ^b from our running example are shown in Figure 3. Since $\mathcal{I}_e^*([O_\Gamma^a]) = \{v, v', w', w\}$, the map Θ removes the edge $e = \{v, w\}$. However, since $\mathcal{I}_e^*([O_\Gamma^b]) = \{v, w\}$, the map Θ contracts e .*

The results we have derived for the vw -interval now allow us to establish the following:

Theorem 3.12. *The map Θ is a bijection.*

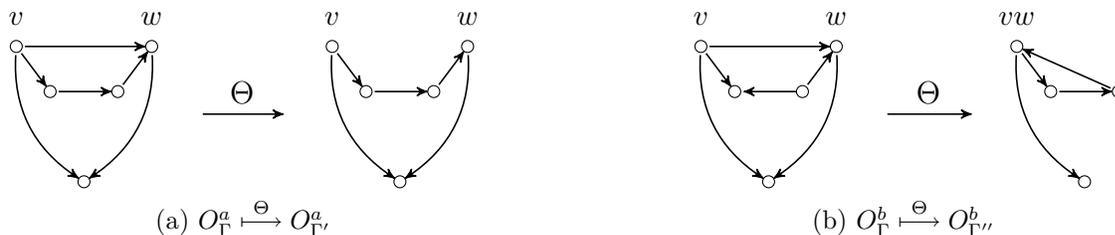


Figure 3: An example of the map Θ applied to the orientations O_{Γ}^a and O_{Γ}^b from Example 2.1 and Figure 1. If contracting the edge $e = \{v, w\}$ would introduce a directed cycle (as in O_{Γ}^a), then we must delete it. Otherwise (as in O_{Γ}^b), contract it. Note that this happens precisely when (v, w) is the only directed path from v to w .

Proof. We first prove that Θ is surjective. Let $I = \{v, w\}$ and consider an element $[O_{\Gamma''}] \in \text{Acyc}(\Gamma'') / \sim_{\kappa}$ with $O_{\Gamma''}^{\pi} \in [O_{\Gamma''}]$ where $\pi = v_I \pi_2 \cdots \pi_{n-1}$. Let $\pi^+ = vw \pi_2 \cdots \pi_{n-1} \in S_{\Gamma}$. Clearly $[O_{\Gamma}^{\pi^+}] \in \text{Acyc}(\Gamma) / \sim_{\kappa}$ is mapped to $[O_{\Gamma''}]$ by Θ .

Next, consider an element $[O_{\Gamma'}] \in \text{Acyc}(\Gamma') / \sim_{\kappa}$. If there is no element $O_{\Gamma'}^{\pi}$ of $[O_{\Gamma'}]$ such that $\pi = vw \pi_3 \cdots \pi_n$, then no elements of $[O_{\Gamma}]$ are of this form either, and by definition $[O_{\Gamma}]$ has a preimage under Θ . We are left with the case where $[O_{\Gamma'}]$ contains an element $O_{\Gamma'}^{\pi}$ such that $\pi = vw \pi_3 \cdots \pi_n$, and we must show that there exists $O_{\Gamma'}^{\sigma}$ of $[O_{\Gamma'}]$ such that $[O_{\Gamma}^{\sigma}]$ contains no element of the form O_{Γ}^{σ} with $\sigma = vw \sigma_3 \cdots \sigma_n$. Note that if $\sigma = vw \sigma_3 \cdots \sigma_n$, then the vertices in $\mathcal{I}(O_{\Gamma}^{\sigma})$ are precisely v and w . If the orientation $O_{\Gamma'}$ had a directed path from v to w , then the corresponding orientation $O_{\Gamma} \in \text{Acyc}(\Gamma)$ formed by adding the edge e with orientation (v, w) has vw -interval of size at least 3, so by Proposition 3.3, the acyclic orientation O_{Γ} cannot be κ -equivalent to any orientation O_{Γ}^{σ} such that $\sigma = vw \sigma_3 \cdots \sigma_n$.

Thus it remains to consider the case when $[O_{\Gamma'}]$ contains no acyclic orientation with a directed path from v to w . Pick any simple undirected path P' from v to w in Γ' , which is possible since e is a cycle-edge. Choose an orientation in $[O_{\Gamma'}]$ for which $\nu_{P'}$ is maximal. Without loss of generality we may assume that $O_{\Gamma'}$ is this orientation. Let $O_{\Gamma} \in \text{Acyc}(\Gamma)$ be the orientation that agrees with $O_{\Gamma'}$, and with e oriented as (w, v) . Since we have assumed that there is no directed path from v to w this orientation is acyclic. We claim that for any $\sigma = vw \sigma_3 \cdots \sigma_n$ one has $O_{\Gamma}^{\sigma} \notin [O_{\Gamma}]$. To see this, assume the statement is false. Let P be the undirected cycle in Γ formed by adding the edge e to the path P' . Because e is oriented as (v, w) in O_{Γ}^{σ} and as (w, v) in O_{Γ} , we have $\nu_P(O_{\Gamma}^{\sigma}) = \nu_{P'}(O_{\Gamma'}^{\sigma}) - 1$ and $\nu_P(O_{\Gamma}) = \nu_{P'}(O_{\Gamma'}) + 1$. If O_{Γ} and O_{Γ}^{σ} were κ -equivalent, then

$$\nu_{P'}(O_{\Gamma'}^{\sigma}) - 1 = \nu_P(O_{\Gamma}^{\sigma}) = \nu_P(O_{\Gamma}) = \nu_{P'}(O_{\Gamma'}) + 1,$$

and thus $\nu_{P'}(O_{\Gamma'}^{\sigma}) = \nu_{P'}(O_{\Gamma'}) + 2$. Any click sequence mapping O_{Γ} to O_{Γ}^{σ} is a click-sequence from $O_{\Gamma'}$ to $O_{\Gamma'}^{\sigma}$. Therefore, $O_{\Gamma'}^{\sigma} \in [O_{\Gamma'}]$, which contradicts the maximality of $\nu_{P'}(O_{\Gamma'})$. We therefore conclude that $O_{\Gamma}^{\sigma} \notin [O_{\Gamma}]$, that $\Theta([O_{\Gamma}]) = [O_{\Gamma'}]$, and hence that Θ is surjective.

We next prove that Θ is an injection. By Proposition 3.10 (with $I = \{v, w\}$), Θ is injective when restricted to the preimage of $[O_{\Gamma''}]$ under Θ . Thus it suffices to show

that every element in $\text{Acyc}(\Gamma')/\sim_\kappa$ has a unique preimage under Θ . By Proposition 3.8, every preimage of $[O_{\Gamma'}]$ must have the same vw -interval I , containing $k > 2$ vertices. Suppose there were preimages $[O_\Gamma^\pi] \neq [O_\Gamma^\sigma]$ of $[O_{\Gamma'}]$. By Proposition 3.10, it follows that $O_{\Gamma_I}^\pi \approx_\kappa O_{\Gamma_I}^\sigma$. We will now show that this leads to a contradiction.

Assume that $\mathbf{c} = c_1 \cdots c_m$ is a click-sequence from $O_{\Gamma'}^\pi$ to $O_{\Gamma'}^\sigma$. If one of π or σ is not κ -equivalent to a permutation with vertices v and w in succession, then their corresponding κ -classes would be unchanged by the removal of edge e . In light of this, we may assume that $\pi = v\pi_2 \dots \pi_{n-1}w$ and $\sigma = v\sigma_2 \dots \sigma_{n-1}w$, and thus that $c_1 = v$ and $c_m = w$. By Proposition 3.6, we may assume that the vertices in I appear in \mathbf{c} in some number of disjoint consecutive ‘‘blocks,’’ i.e., subsequences of the form $c_i \cdots c_{i+k-1}$. Replacing each of these blocks with V_I yields a click-sequence from $O_{\Gamma_I}^\pi$ to $O_{\Gamma_I}^\sigma$, contradicting the fact that $O_{\Gamma_I}^\pi \approx_\kappa O_{\Gamma_I}^\sigma$. Therefore, no such click sequence \mathbf{c} exists, and Θ must be an injection, and the proof is complete. \square

The main result in [10] is a recurrence relation for $\kappa(\Gamma)$ under edge deletion and edge contraction. This is an easy consequence of Theorem 3.12.

Corollary 3.13 ([10]). *Let Γ be a finite undirected graph with $e \in e[\Gamma]$, and let Γ'_e be the graph obtained from Γ by deleting e , and let Γ''_e be the graph obtained from Γ by contracting e . Then*

$$\kappa(\Gamma) = \begin{cases} \kappa(\Gamma_1)\kappa(\Gamma_2), & e \text{ is a bridge linking } \Gamma_1 \text{ and } \Gamma_2, \\ \kappa(\Gamma'_e) + \kappa(\Gamma''_e), & e \text{ is a cycle-edge.} \end{cases} \quad (8)$$

The first part involving a bridge is straightforward, while the second part is a direct consequence of Theorem 3.12.

4 A Complete Invariant of $\text{Acyc}(\Gamma)/\sim_\kappa$

Theorem 3.12 is more than just an alternative proof of the enumeration of $\kappa(\Gamma)$. We can utilize the explicit bijection to derive an additional interesting and useful corollary: When taken over all cycles C in a graph Γ , ν_C is a complete invariant of $\text{Acyc}(\Gamma)/\sim_\kappa$. This result is originally due to Pretzel [16], though the techniques are much different than the ones here.

Theorem 4.1. *If $\nu_C(O_\Gamma^1) = \nu_C(O_\Gamma^2)$ for every cycle C , then $O_\Gamma^1 \sim_\kappa O_\Gamma^2$.*

Proof. Assume the statement is false and let Γ be a graph for which it fails, minimal with respect to $|e[\Gamma]|$. Fix a cycle-edge $e = \{v, w\}$, and for any $[O_\Gamma] \in \text{Acyc}(\Gamma)/\sim_\kappa$, call an orientation $O_\Gamma^\pi \in [O_\Gamma]$ distinguished with respect to e if $\pi = \pi_1\pi_2 \cdots \pi_n$ such that (i) $\pi_1 = v$, and (ii) $\pi_k = w$ where k is minimal given that $\pi_1 = v$. By assumption, there exists $O_\Gamma^1 \approx_\kappa O_\Gamma^2$ with $\nu_C(O_\Gamma^1) = \nu_C(O_\Gamma^2)$ for every cycle C in Γ . Without loss of generality, we may assume that O_Γ^1 and O_Γ^2 are distinguished orientations with respect to $e = \{v, w\}$. Define a vw -path to be a directed path from v to w that does not traverse e . There are three cases to consider:

Case 1: Both O_Γ^1 and O_Γ^2 contain a vw -path. By definition of Θ , both $\Theta(O_\Gamma^1)$ and $\Theta(O_\Gamma^2)$ are contained in $\text{Acyc}(\Gamma'_e)/\sim_\kappa$. Moreover, because Θ is a bijection, they are distinct elements in $\text{Acyc}(\Gamma'_e)/\sim_\kappa$. Every cycle in Γ'_e is also a cycle in Γ , and therefore, $\nu_{C'}(\Theta(O_\Gamma^1)) = \nu_{C'}(\Theta(O_\Gamma^2))$ for every cycle C' in Γ'_e , contradicting the minimality of $|e[\Gamma]|$.

Case 2: Neither O_Γ^1 nor O_Γ^2 contain a vw -path. Since O_Γ^1 and O_Γ^2 are distinguished with respect to e , $\Theta([O_\Gamma^i]) = [O_{\Gamma''}^i] \in \text{Acyc}(\Gamma'')/\sim_\kappa$ for $i = 1, 2$. Again, these two orientations are distinct because Θ is a bijection. Any cycle C'' in Γ''_e beginning and ending at the vertex V (the image of v and w under contraction) can be canonically extended to a cycle C of Γ . Therefore, $\nu_{C''}(\Theta(O_{\Gamma''}^1)) = \nu_{C''}(\Theta(O_{\Gamma''}^2))$, again contradicting the minimality of $|e[\Gamma]|$.

Case 3: Precisely one of O_Γ^1 and O_Γ^2 contain a vw -path. Without loss of generality, suppose that P is a length- k vw -path in O_Γ^1 , and let C be the cycle formed by adding vertex v to the end of P . Clearly, $\nu_C(O_\Gamma^1) = k - 1$, and by assumption, $\nu_C(O_\Gamma^2) = k - 1$ as well. However, this means that every edge in P is oriented from v to w in O_Γ^2 , contradicting the assumption that O_Γ^2 did not contain a vw -path. Therefore, Case 3 is impossible, and the proof is complete. \square

Example 4.2. *It is immediate from the recurrence (8) in Corollary 3.13 that $\kappa(\Gamma) = 1$ if and only if Γ is a forest. Let Circle_n be a chordless n -cycle, and Line_n the line graph on n vertices. Deleting or contracting any edge of Circle_3 leaves a tree, and so $\kappa(\text{Circle}_3) = 2$. For $n > 3$, deleting an edge from Circle_n leaves Line_n , and contracting an edge leaves Circle_{n-1} . By Corollary 3.13, if $n > 3$, then*

$$\kappa(\text{Circle}_n) = \kappa(\text{Line}_n) + \kappa(\text{Circle}_{n-1}) = 1 + \kappa(\text{Circle}_{n-1}). \quad (9)$$

From the base case of $\kappa(\text{Circle}_3) = 2$, we immediately deduce that $\kappa(\text{Circle}_n) = n - 1$.

Example 4.3. *Let Γ be the undirected version of the orientations from our running example, with edge $e = \{v, w\}$ as before. Deleting e leaves Γ'_e , a 5-cycle. Contracting e leaves Γ''_e , a 3-cycle with an extra edge hanging off (see Figure 3 in Example 3.11). By the recurrence (8) in Corollary 3.13 along with (9),*

$$\kappa(\Gamma) = \kappa(\Gamma'_e) + \kappa(\Gamma''_e) = 4 + 2 = 6.$$

Representatives from the six distinct κ -classes of Γ are shown in Figure 4. This particular transversal was chosen so that v is a source, and so the vw -intervals can be identified immediately, and they are (from left-to-right)

$$\{v, v', w', w\}, \quad \{v, w\}, \quad \{v, v', w', z, w\}, \quad \{v, z, w\}, \quad \{v, z, w\}, \quad \{v, w\}.$$

Note that the first two orientations in Figure 4 are O_Γ^a and O_Γ^b , and the third is κ -equivalent to O_Γ^c . Letting $P = v, v', w', w, v$ and $Q = v, z, w, v$ be the paths as defined in Example 3.2, the pair (ν_P, ν_Q) is a complete invariant of $\text{Acyc}(\Gamma)/\sim_\kappa$. For each of the six orientations, the values of $\nu_P(O_\Gamma)$ (top) and $\nu_Q(O_\Gamma)$ (bottom) are shown in Figure 4.

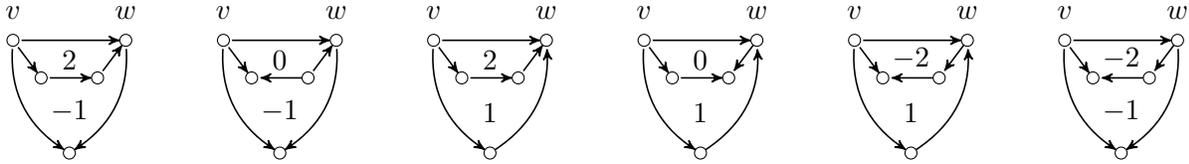


Figure 4: A transversal for $\text{Acyc}(\Gamma)/\sim_\kappa$, with the values of the complete invariant (ν_P, ν_Q) .

5 Conjugacy of Coxeter Elements

Our analysis of $\text{Acyc}(\Gamma)/\sim_\kappa$ also gives a straightforward solution to the conjugacy problem for Coxeter elements in simply-laced Coxeter groups. Before stating the theorem and proof, we will briefly review the connection between κ -equivalence and Coxeter theory, as described in [10]. A *Coxeter group* is a generalized reflection group, generated by n distinguished involutions s_1, \dots, s_n by the presentation

$$W = \langle s_1, \dots, s_n \mid (s_i s_j)^{m_{ij}} \rangle,$$

where $m_{ij} = 1$ iff $i = j$, and $m_{ij} \geq 2$ otherwise. If $s_i s_j$ has infinite order, then we say that $m_{ij} = \infty$. The pair (W, S) of the group W with the generating set S is called a *Coxeter system*, which is uniquely encoded by its Coxeter graph Γ , with vertex set S and edge set $\{s_i, s_j\}$ for which $m_{ij} \geq 3$, where each edge has weight m_{ij} . A Coxeter group is *simply-laced* if each $m_{ij} \leq 3$, thus every undirected graph is the Coxeter graph of a simply-laced Coxeter group.

A *Coxeter element* is the product of the generators in some order, and there is a natural bijection between the set $C(W, S)$ of Coxeter elements $\prod_i s_{\pi(i)}$ of a Coxeter group (see, e.g. [2]) with generators s_i for $1 \leq i \leq n$ and Coxeter graph Γ (ignoring bond strengths), and the set of α -equivalence classes $[\pi]_\Gamma$. This is clear since the commuting generators are precisely those that are not connected in Γ . Thus, there is a natural bijection

$$C(W, S) \longrightarrow \text{Acyc}(\Gamma). \quad (10)$$

Moreover, conjugating a Coxeter element $c = \prod s_{\pi(i)}$ by $s_{\pi(1)}$ corresponds to a cyclic shift, i.e.,

$$s_{\pi(1)}(s_{\pi(1)}s_{\pi(2)} \cdots s_{\pi(n)})s_{\pi(1)} = s_{\pi(2)} \cdots s_{\pi(n)}s_{\pi(1)},$$

since each generator s_i is an involution. Therefore, κ -equivalence naturally carries over to an equivalence relation on $C(W, S)$. The ν -function carries over as well – define $\nu_P(c)$ to be $\nu_P(O_\Gamma)$, where O_Γ is the acyclic orientation of Γ corresponding to c . It is now elementary to see that $c, c' \in C(W, S)$ are conjugate if $c \sim_\kappa c'$. However, the converse of this statement was not proven until 2009.

Theorem 5.1 ([6]). *Let (W, S) be a Coxeter system. Then two Coxeter elements $c, c' \in C(W, S)$ are conjugate if and only if $c \sim_\kappa c'$.*

It follows immediately that the number of distinct conjugacy classes containing Coxeter elements is exactly $\kappa(\Gamma)$. Until the Erikssons' proof, the result was known only for

the special case of $C(W, S)$ when Γ was simply-laced and unicyclic, established by Shi in 2001 [18]. It is elementary to weaken the simply-laced condition to the bond strengths being multiples of three or infinite, which Shi mentions in [18]. The bijection in Theorem 3.12 applied to Shi's result for unicyclic Coxeter graphs yields a simple and elegant proof of the result for all simply-laced systems, which we present below.

Proof. Suppose for sake of contradiction that the Coxeter elements $c = c_1 c_2 \cdots c_n$ and $c' = c'_1 c'_2 \cdots c'_n$ are conjugate with $c \approx_\kappa c'$, and that $w c w^{-1} = c'$ for some $w = w_1 \cdots w_k \in W$ with each $w_i \in S$. By Theorem 4.1, there is some simple chordless cycle $P = v_0, v_1, \dots, v_{m-1}, v_m$ (i.e., $v_0 = v_m$) in Γ such that $\nu_P(c) \neq \nu_P(c')$. Let $S_P = \{s_i \mid i \in P\}$, and let C_m be the (circular) Coxeter graph induced by the vertices in P . The Coxeter group generated by S_P is the affine Weyl group \tilde{A}_{m-1} , and there is a natural homomorphism $W \xrightarrow{\varphi} \tilde{A}_{m-1}$ defined on the generators by

$$s_i \xrightarrow{\varphi} \begin{cases} s_i & i \in P \\ 1 & i \notin P. \end{cases}$$

Since c and c' are conjugate in W , $\varphi(c)$ and $\varphi(c')$ are conjugate in \tilde{A}_{m-1} . By choice of P , $\nu_P(c) \neq \nu_P(c')$, and thus $\varphi(c) \not\approx_\kappa \varphi(c')$. However, since the statement holds for unicyclic graphs, we must have $\varphi(c) \sim_\kappa \varphi(c')$, which is the desired contradiction. \square

Theorems 4.1 and 5.1 give us an easy way to verify whether two Coxeter elements c and c' in any Coxeter group W are conjugate. Pick a cycle basis of the Coxeter graph Γ , and for each cycle C , compute $\nu_C(c)$ and $\nu_C(c')$. By Theorem 4.1, $c \sim_\kappa c'$ iff $\nu_C(c) = \nu_C(c')$ for each C . By Theorem 5.1, this is equivalent to c and c' being conjugate in W . Therefore, conjugacy of Coxeter elements can be verified in $O(n^2)$ steps, where $n = |v[\Gamma]|$. One application of this is seeing how the conjugacy classes split as an edge $\{s_i, s_j\}$ is added to Γ (or equivalently, as the relation $(s_i s_j)^{m_{ij}}$ is added to the group presentation).

6 Discrete Dynamical Systems, Node-firing Games, and Quiver Representations

We conclude with a brief discussion of how the equivalence relation studied in this paper arises in various areas of mathematics. The original motivation came from both authors' interest in *sequential dynamical systems* (SDSs). The equivalence relation \sim_α arises naturally in the study of functional equivalence of these systems. This can be seen as follows. Given a graph Γ with vertex set $\{1, 2, \dots, n\}$ as above, a state $x_v \in K$ is assigned to each vertex v of Γ for some finite set K . The *system state* is the tuple consisting of all the vertex states, and is denoted by $x = (x_1, \dots, x_n) \in K^n$. The sequence of states associated to the 1-neighborhood $B_1(v; \Gamma)$ of v in Γ (in some fixed order) is denoted by $x[v]$. A sequence of vertex functions $(f_i)_i$ with $f_i: K^{d(i)+1} \rightarrow K$ induces Γ -local functions $F_i: K^n \rightarrow K^n$ of the form

$$F_i(x_1, \dots, x_n) = (x_1, \dots, x_{i-1}, f_i(x[i]), x_{i+1}, \dots, x_n).$$

The sequential dynamical system map with update order $\pi = (\pi_i)_i \in S_\Gamma$ is the function composition

$$F_\pi = F_{\pi_n} \circ F_{\pi_{n-1}} \circ \cdots \circ F_{\pi_2} \circ F_{\pi_1}. \quad (11)$$

By construction, if $\pi \sim_\alpha \pi'$ holds, then F_π and $F_{\pi'}$ are identical as functions, independent of the choice of state space K or vertex functions. Thus, $\alpha(\Gamma)$ is a general upper bound for the number of functionally non-equivalent SDS maps that can be generated over the graph Γ for a fixed sequence of Γ -local functions. Moreover, for any graph Γ , there exist Γ -local functions for which this bound is sharp [14]. A weaker form of equivalence is *cycle equivalence*, which means that the dynamical system maps are conjugate (using the discrete topology) when restricted to their sets of periodic points. In the language of graph theory, this means their periodic orbits are isomorphic as directed graphs. For an update order $\pi = \pi_1 \cdots \pi_n$, define $\text{shift}(\pi) = \pi_2 \cdots \pi_n \pi_1$. The following result shows how κ -equivalent update orders yield dynamical system maps that are cycle equivalent.

Theorem 6.1. *For any finite set K of vertex states, and for any $\pi \in S_\Gamma$, the SDS maps F_π and $F_{\text{shift}(\pi)}$ are cycle equivalent.*

We refer to [11] for the proof of this result, as well as additional background on equivalences of sequential dynamical systems, and applications of κ -equivalence to the structural properties of their phase spaces. It is interesting to note that for the class of *update sequence independent* (see [9]) sequential dynamical systems with binary states, there is an additional equivalence on acyclic orientations that governs cycle equivalence: reversal of all edge orientations.

The *chip-firing* game was introduced by Björner, Lovász, and Shor [3]. It is played over an undirected graph Γ , and each vertex is given some number of (but possibly zero) chips. If vertex i has degree d_i , and at least d_i chips, then a legal move (or a “click”) of vertex i is a transfer of one chip to each neighboring vertex. This may be viewed as a generalization of a source-to-sink move for acyclic orientations where the out-degree of a vertex plays the role of the chip count. The chip-firing game is closely related to the *numbers game* [2]. In the numbers game over a graph Γ , the legal sequences of moves are in 1–1 correspondence with the reduced words of the Coxeter group with Coxeter graph Γ . For an excellent summary and comparison of these games, see [7].

A quiver is a finite directed graph (loops and multiple edges are allowed), and appears primarily in the study of representation theory. A quiver Q with a field K gives rise to a path algebra KQ , and there is a natural correspondence between KQ -modules and K -representations of Q . In fact, there is an equivalence between the categories of quiver representations, and modules over path algebras. A path algebra is finite-dimensional if and only if the quiver is acyclic, and the modules over finite-dimensional path algebras form a reflective subcategory. A *reflection functor* maps representations of a quiver Q to representations of a quiver Q' , where Q' differs from Q by a source-to-sink operation [12]. We note that while the composition of n source-to-sink operations (one for each vertex) maps a quiver back to itself, the corresponding composition of reflection functors is not

the identity, but rather a *Coxeter functor*. In fact, the same result in [19] about powers of Coxeter elements being reduced was proven previously using techniques from the representation theory of quivers [8].

We hope this paper will motivate further explorations of the connections between these topics. We are particularly curious about any implications to the representation theory of quivers. This is a field which the both authors of this paper are quite unfamiliar with, yet it motivated Kleiner and Pelley to study admissible sequences and apply these tools from quiver representations to Coxeter groups. Without this work, the aforementioned papers of Speyer and the Erikssons would likely not have materialized.

Acknowledgments. Both authors are grateful to the NDSSL group at Virginia Tech for the support of this research. Special thanks to Ed Green and Ken Millett for helpful discussions and feedback, and to William Y. C. Chen for valuable advice regarding the preparation and structuring of this paper.

References

- [1] M. Auslander, I. Reiten, and S. O. Smalø. *Representation Theory of Artin Algebras*. Cambridge University Press, 1997.
- [2] A. Björner and F. Brenti. *Combinatorics of Coxeter Groups*. Springer-Verlag, New York, 2005.
- [3] A. Björner, L. Lovász, and P. Shor. Chip-firing games on graphs. *European J. Combin.*, 12:283–291, 1991.
- [4] P. Cartier and D. Foata. *Problemes combinatoires de commutation et réarrangements*, volume 85 of *Lecture Notes in Mathematics*. Springer Verlag, 1969.
- [5] A. J. Coleman. Killing and the Coxeter transformation of Kac-Moody algebras. *Invent. Math.*, 95:447–477, 1989.
- [6] H. Eriksson and K. Eriksson. Conjugacy of Coxeter elements. *Electron. J. Combin.*, 16(2):#R4, 2009.
- [7] K. Eriksson. Node firing games on graphs. *Contemp. Math.*, 178:117–127, 1994.
- [8] M. Kleiner and A. Pelley. Admissible sequences, preprojective representations of quivers, and reduced words in the Weyl group of a Kac-Moody algebra. *Internat. Math. Res. Notices*, 2007, May 2007.
- [9] M. Macauley, J. McCammond, and H. S. Mortveit. Dynamics groups of asynchronous cellular automata. *J. Algebraic Combin.*, 33:31–55, 2011.
- [10] M. Macauley and H. S. Mortveit. On enumeration of conjugacy classes of Coxeter elements. *Proc. Amer. Math. Soc.*, 136:4157–4165, 2008.

- [11] M. Macauley and H. S. Mortveit. Cycle equivalence of graph dynamical systems. *Nonlinearity*, 22:421–436, 2009.
- [12] R. Marsh, M. Reineke, and A. Zelevinsky. Generalized associahedra via quiver representations. *Trans. Amer. Math. Soc.*, 355:4171–4186, 2003.
- [13] E. Miller, V. Reiner, and B. Sturmfels, editors. *Geometric Combinatorics*, volume 13 of *IAS/Park City Mathematics Series*. AMS and IAS/Park City Mathematics Institute, 2007.
- [14] H. S. Mortveit and C. M. Reidys. Discrete, sequential dynamical systems. *Discrete Math.*, 226:281–295, 2001.
- [15] I. Novik, A. Postnikov, and B. Sturmfels. Syzygies of oriented matroids. *Duke Math. J.*, 111:287–317, 2002.
- [16] O. Pretzel. On reorienting graphs by pushing down maximal vertices. *Order*, 3(2):135–153, 1986.
- [17] J.-Y. Shi. The enumeration of Coxeter elements. *J. Algebraic Combin.*, 6:161–171, 1997.
- [18] J.-Y. Shi. Conjugacy relation on Coxeter elements. *Adv. Math.*, 161:1–19, 2001.
- [19] D. E. Speyer. Powers of Coxeter elements in infinite groups are reduced. *Proc. Amer. Math. Soc.*, 137:1295–1302, 2009.
- [20] J. R. Stembridge. On the fully commutative elements of Coxeter groups. *J. Algebraic Combin.*, 5:353–385, 1996.