

UCLA

UCLA Previously Published Works

Title

Move-reduced graphs on a torus

Permalink

<https://escholarship.org/uc/item/8nh1m070>

Authors

Galashin, Pavel

George, Terrence

Publication Date

2024

DOI

10.1090/tran/9168

Copyright Information

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

MOVE-REDUCED GRAPHS ON A TORUS

PAVEL GALASHIN AND TERRENCE GEORGE

ABSTRACT. We determine which bipartite graphs embedded in a torus are move-reduced. In addition, we classify equivalence classes of such move-reduced graphs under square/spider moves. This extends the class of minimal graphs on a torus studied by Goncharov–Kenyon, and gives a toric analog of Postnikov’s and Thurston’s results on a disk.

INTRODUCTION

Let $\mathbb{T} = \mathbb{R}^2/\mathbb{Z}^2$ be a torus, and let Γ be a bipartite graph embedded in \mathbb{T} . We say that two such graphs Γ, Γ' are *move-equivalent* if they are related by the moves (M1)–(M2) shown in Figure 1. We say that Γ is *move-reduced* if there does not exist a graph Γ' move-equivalent to Γ to which we can apply one of the *reduction moves* (R1)–(R3) shown in Figure 2. The goal of this paper is to describe which graphs Γ are move-reduced, and which pairs of move-reduced graphs are move-equivalent. A similar problem has been considered in [GK13] for the class of *minimal graphs*. Each minimal graph is move-reduced, however, the converse is not true; see Figure 3.

We briefly summarize our main results; see Section 1 for more details. It was shown in [GK13] that move-equivalence classes of minimal graphs are classified by their Newton polygons N . The sides of N are obtained by taking the homology classes of strands in Γ . Here, a *strand* is a path making a sharp right (resp., left) turn at each black (resp., white) vertex. A strand of a move-reduced (as opposed to minimal) graph Γ may intersect itself, and this induces a *weak decoration* $\lambda = (\lambda^e)_{e \in E(N)}$ of N , labeling each side $e = (i, j)$ of N by a partition λ^e of $\gcd(i, j)$. Our first main result (Theorem 1.5) gives a characterization of move-reduced graphs in terms of weakly decorated Newton polygons that parallels the results of [GK13, Pos06].

Our second main result concerns move-equivalence classes of move-reduced graphs. The solution to this problem turns out to be more subtle than its counterparts in [GK13, Pos06]. First, we show that in a move-reduced graph, different strands corresponding to the same side of N never cross each other. This induces a *strong decoration* $\alpha = (\alpha^e)_{e \in E(N)}$ of N , labeling each side $e = (i, j)$ of N with a cyclic composition α^e of $\gcd(i, j)$. We associate a *rotation number* $d(\alpha)$ to α , and our second main result (Theorem 1.13) is that the set of all move-reduced graphs with strongly decorated Newton polygon (N, α) is a union of $d(\alpha)$ move-equivalence classes. The classes are distinguished by the value of an explicit *modular invariant* $\mu(\Gamma) \in \mathbb{Z}/d(\alpha)\mathbb{Z}$ associated to each move-reduced graph Γ .

Date: March 1, 2024.

2020 *Mathematics Subject Classification.* Primary: 05C10. Secondary: 13F60.

Key words and phrases. Bipartite graphs on a torus, square moves, affine permutations, conjugation.

P.G. was supported by an Alfred P. Sloan Research Fellowship and by the National Science Foundation under Grants No. DMS-1954121 and No. DMS-2046915.




FIGURE 1. Equivalence moves for bipartite graphs in \mathbb{T} . One can also apply these moves with the roles of white and black swapped. For (M1), the vertices of the square are assumed to have degree at least three. For (M2), the two white vertices are assumed to be distinct and have degree at least two. The shaded area denotes a small open disk inside \mathbb{T} .




FIGURE 2. Reduction moves for bipartite graphs. (R1) removes one of two parallel edges, (R2) removes a leaf together with its single neighbor, and (R3) removes an isolated edge.

Our motivation to study move-reduced graphs arises from the dimer model on Γ and the associated *spectral transform* of [KOS06, KO06]. Each weighted bipartite graph (Γ, wt) with positive real edge weights embedded in \mathbb{T} determines a simple Harnack curve with a distinguished line bundle. It is thus natural to study which limiting objects appear when one sends some edge weights to zero. This corresponds to deleting edges from Γ and then applying reduction moves. Note in particular that the move-reduced graph Γ_2 in Figure 3(b) is obtained from the minimal graph Γ_1 in Figure 3(a) by removing a single edge, which demonstrates that the class of move-reduced graphs is more naturally suited for this problem.

For the case of planar bipartite graphs in a disk, the role of the spectral transform is played by Postnikov's boundary measurement map [Pos06]. This map associates to each weighted planar bipartite graph in a disk a point inside the *totally nonnegative Grassmannian*. The boundary measurement map is continuous with respect to sending some edge weights to zero. Our motivating problem is to find a suitable compactification of the space of Harnack curves with line bundles such that the spectral transform would have similar continuity properties.

As an important step towards Postnikov's results, he characterized move-reduced graphs on a disk and showed that their move-equivalence classes are in bijection with *positroids*. The present manuscript is the first in a series of papers aimed at studying the toric analog of the totally nonnegative Grassmannian and its positroid stratification.

1. MAIN RESULTS

In Section 1.1, we introduce the notions of *weakly* and *strongly decorated polygons*. In Section 1.3, we will associate a weakly decorated polygon with any bipartite graph embedded in the torus, and we will use it to characterize move-reduced graphs. In Section 1.4, we will

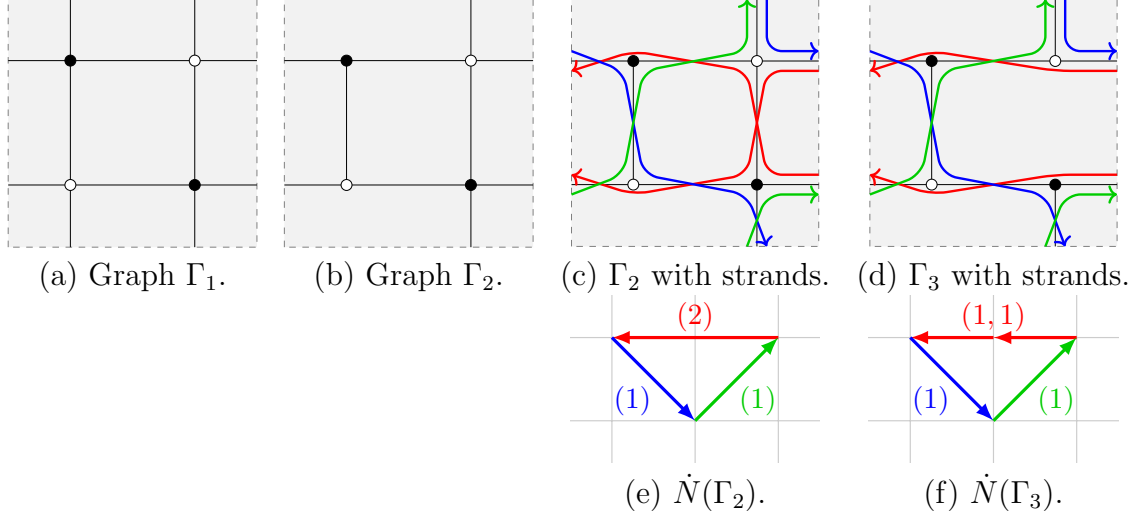


FIGURE 3. The graphs Γ_1 and Γ_3 are minimal in the sense of [GK13] and therefore are move-reduced. The graph Γ_2 is not minimal but is move-reduced. See Section 1.3 for a definition of strands and $\dot{N}(\Gamma)$.

associate a strongly decorated polygon to any move-reduced graph Γ , and will use it to characterize which graphs are move-equivalent to Γ .

1.1. Decorated polygons. A convex polygon N in the plane \mathbb{R}^2 is called *integral* if its vertices are contained in $\mathbb{Z}^2 \subset \mathbb{R}^2$. We denote the set of edges of N by $E(N)$, and orient them counterclockwise around the boundary of N so that each edge is a vector in \mathbb{Z}^2 . For an edge $e = (a, b)$ of N , let $|e|_{\mathbb{Z}} := \gcd(a, b)$ be its *integer length*. For vectors $e, e' \in \mathbb{Z}^2$, let $\det(e, e')$ be the determinant of the 2×2 matrix with columns e, e' .

A *partition* of n with k parts is a tuple $\lambda = (\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0)$ such that $|\lambda| := \lambda_1 + \lambda_2 + \dots + \lambda_k = n$. A *composition* of n with k parts is a tuple $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_k) \in \mathbb{Z}_{>0}^k$ such that $|\alpha| := \alpha_1 + \dots + \alpha_k = n$. A *cyclic composition* of n with k parts is an equivalence class of compositions of n with k parts under cyclic shifts $(\alpha_1, \alpha_2, \dots, \alpha_k) \sim (\alpha_2, \dots, \alpha_k, \alpha_1)$. Thus, forgetting the order of the parts of a (cyclic) composition yields a partition.

Definition 1.1.

- A *weakly decorated polygon* is a pair $\dot{N} = (N, \boldsymbol{\lambda})$, where N is a convex integral polygon, and $\boldsymbol{\lambda} = (\lambda^e)_{e \in E(N)}$, where λ^e is a partition of $|e|_{\mathbb{Z}}$.
- A *strongly decorated polygon* is a pair $\ddot{N} = (N, \boldsymbol{\alpha})$, where N is a convex integral polygon, and $\boldsymbol{\alpha} = (\alpha^e)_{e \in E(N)}$, where α^e is a cyclic composition of $|e|_{\mathbb{Z}}$.

1.2. Minimal graphs. Recall that a *strand* or a *zig-zag path* S is a walk in Γ that turns maximally right at the black vertices and maximally left at the white vertices of Γ . The set of strands of Γ is denoted by $\mathcal{S}(\Gamma)$. Since Γ is finite, a strand S is a (not necessarily simple) closed walk, and we let $[S] \in \mathbb{Z}^2 = H_1(\mathbb{T}, \mathbb{Z})$ denote its homology. Since each edge of Γ is contained in two strands that traverse it in opposite directions, the sum $\sum_{S \in \mathcal{S}(\Gamma)} [S]$ is zero, so we can associate to Γ an integral polygon $N(\Gamma)$, called the *Newton polygon* of Γ , as follows. We let $N(\Gamma)$ be the convex integral polygon N (possibly degenerate, i.e., having

0 area), unique up to translation, whose counterclockwise-oriented boundary consists of the vectors $([S])_{S \in \mathcal{S}(\Gamma)}$ in some order.

A *face* of Γ is a connected component of $\mathbb{T} \setminus \Gamma$. Thus, a face of Γ is contractible if and only if it is homeomorphic to an open disk.

Let $\pi : \mathbb{R}^2 \rightarrow \mathbb{T}$ denote the universal covering map. Let $\tilde{\Gamma} := \pi^{-1}(\Gamma)$ denote the corresponding biperiodic graph in \mathbb{R}^2 . Following Goncharov–Kenyon [GK13], we say that Γ is *minimal* if every face of Γ is contractible, there is no strand with zero homology, there is no strand with a lift in $\tilde{\Gamma}$ with a self-intersection, and there are no *parallel bigons* in $\tilde{\Gamma}$, i.e., pairs of lifts of strands in $\tilde{\Gamma}$ that pass through edges $e_1 \neq e_2$ of $\tilde{\Gamma}$ with both lifts oriented from e_1 to e_2 . The class of minimal graphs is preserved by move-equivalence, and minimal graphs are move-reduced. Therefore, move-equivalence classes of minimal graphs are a subset of move-equivalence classes of move-reduced graphs. We will extend the following two results about minimal graphs to move-reduced graphs.

Lemma 1.2. [GK13, Lemma 3.14] *Let Γ be a minimal bipartite graph embedded in \mathbb{T} with Newton polygon $N(\Gamma) = N$. Then Γ has $2 \text{Area}(N)$ contractible faces.*

We generalize Lemma 1.2 and prove the converse in Theorem 1.5.

Theorem 1.3. [GK13, Theorem 2.5 and Lemma 3.11] *For any convex integral polygon N , there exists a minimal bipartite graph Γ that admits a perfect matching and satisfies $N(\Gamma) = N$. Two minimal graphs Γ, Γ' are move-equivalent if and only if $N(\Gamma) = N(\Gamma')$.*

Theorem 1.3 is generalized in Proposition 1.11 and Theorem 1.13; see Remark 1.14.

1.3. Move-reduced graphs. We say that two strands $S, S' \in \mathcal{S}(\Gamma)$ are *parallel* if $[S], [S'] \neq 0$ and $[S] \in \mathbb{R}_{>0}[S']$. For each edge $e \in E(N)$, we let

$$\mathcal{S}^e(\Gamma) := \{S \in \mathcal{S}(\Gamma) \mid [S] \in \mathbb{R}_{>0}e\}$$

denote the corresponding set of parallel strands. Thus, we have $e = \sum_{S \in \mathcal{S}^e(\Gamma)} [S]$, and we let $\lambda^e := (|[S]|_{\mathbb{Z}})_{S \in \mathcal{S}^e(\Gamma)}$ be the corresponding partition of $|e|_{\mathbb{Z}}$. We call $\dot{N}(\Gamma)$ the *weakly decorated Newton polygon* of Γ . The weakly decorated Newton polygon is invariant under (M1)–(M2) but not under (R1)–(R3).

In Proposition 1.11, we will see that for any weakly decorated polygon \dot{N} , there exists a move-reduced graph Γ satisfying $\dot{N}(\Gamma) = \dot{N}$. On the other hand, it is clear that any graph Γ can be transformed into a move-reduced graph using the moves (M1)–(M2) and (R1)–(R3).

Definition 1.4. For a partition $\lambda = (\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0)$ of n with k parts, the *excess* of λ is defined by $\text{exc}(\lambda) := n - k = \sum_{i=1}^k (\lambda_i - 1)$. If $\boldsymbol{\lambda} = (\lambda^e)_{e \in E}$ is a collection of partitions, we denote $\text{exc}(\boldsymbol{\lambda}) := \sum_{e \in E} \text{exc}(\lambda^e)$.

We are ready to state our first main result.

Theorem 1.5. *Let Γ be a bipartite graph embedded in \mathbb{T} with weakly decorated Newton polygon $\dot{N}(\Gamma) = (N, \boldsymbol{\lambda})$. Assume that Γ has a perfect matching. The following conditions are equivalent.*

- (1) Γ is move-reduced.
- (2) Γ has $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ contractible faces, no contractible connected components, and no leaf vertices.




FIGURE 4. A move-reduced graph with no perfect matchings and whose Newton polygon is a single point.

Remark 1.6. If Γ is minimal, then $\lambda^e = (1, 1, \dots, 1)$ for each edge $e \in E(N)$, so $\text{exc}(\boldsymbol{\lambda}) = 0$. Moreover, Γ has no contractible connected components since we would have a zero-homology strand and no leaf vertices since we would have a strand whose lift in $\tilde{\Gamma}$ has a self-intersection. Therefore, Theorem 1.5 generalizes Lemma 1.2 and also shows the converse.

Remark 1.7. The assumption that Γ has a perfect matching is essential; for example, Theorem 1.5 fails for the graph Γ in Figure 4. This graph is move-reduced and does not have any perfect matchings. Thus, Γ satisfies condition (1) but does not satisfy condition (2) of Theorem 1.5. Alternatively, if Γ has no isolated vertices, the assumption that Γ has a perfect matching can be replaced with either one of the following assumptions:

- the Newton polygon of Γ is not a single point, or
- the number of black and white vertices in Γ is the same;

see part (i) of Theorem 1.16 below.

Remark 1.8. The condition that Γ has $2\text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ contractible faces in (2) is equivalent to a statement that Γ has the minimal possible number of contractible faces among all graphs with weakly decorated Newton polygon $\dot{N}(\Gamma)$.

Example 1.9. For the graphs $\Gamma_1, \Gamma_2, \Gamma_3$ shown in Figure 3, the weakly decorated Newton polygons $\dot{N}(\Gamma_2), \dot{N}(\Gamma_3)$ are computed in Figure 3(e–f). In particular, letting $\dot{N}(\Gamma_2) = (N, \boldsymbol{\lambda})$ and $\dot{N}(\Gamma_3) = (N, \boldsymbol{\lambda}')$, we see that $\text{Area}(N) = 1$, $\text{exc}(\boldsymbol{\lambda}) = 1$, and $\text{exc}(\boldsymbol{\lambda}') = 0$. This is consistent with Theorem 1.5 since Γ_2 has 3 faces, while Γ_3 has 2 faces, all of which are contractible.

1.4. Move-equivalence classes of move-reduced graphs. In this section, each graph is assumed to be bipartite and to have a perfect matching. We view graphs embedded in the torus up to isotopies. Let $\pi : \mathbb{R}^2 \rightarrow \mathbb{T}$ denote the universal covering map.

Proposition 1.10. *Let Γ be a move-reduced graph with Newton polygon N .*

- (1) *There are no zero-homology strands and any lift \tilde{S} of a strand $S \in \mathcal{S}(\Gamma)$ under π does not intersect itself;*
- (2) *Any strand $S \in \mathcal{S}(\Gamma)$ intersects itself $\lfloor |S|_{\mathbb{Z}} - 1 \rfloor$ times;*
- (3) *Any two distinct parallel strands $S, S' \in \mathcal{S}(\Gamma)$ do not share any vertices or edges of Γ .*

Let Γ be a move-reduced graph with Newton polygon N . By part (2) of Proposition 1.10, for $e \in E(N)$, any two strands $S \neq S'$ in $\mathcal{S}^e(\Gamma)$ do not share vertices or edges. Thus, we have a natural cyclic ordering on $\mathcal{S}^e(\Gamma)$ given by the direction of the normal vector to e that points into the interior of N . Let $\alpha^e = (|S|_{\mathbb{Z}})_{S \in \mathcal{S}^e(\Gamma)}$ be the corresponding cyclic

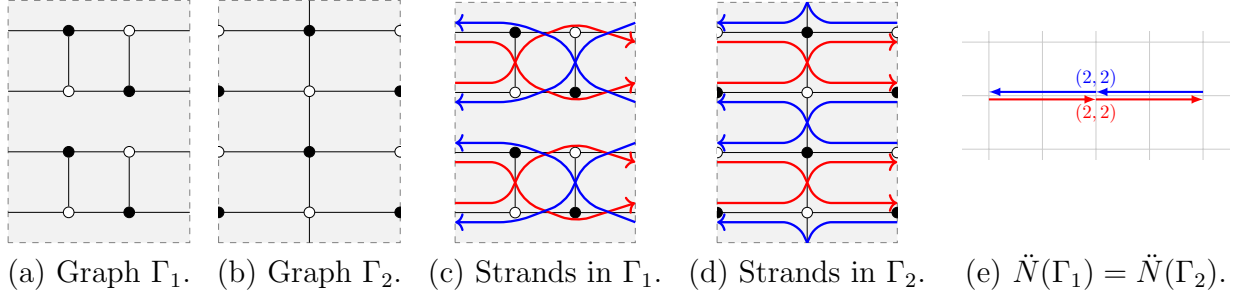


FIGURE 5. Two move-reduced graphs that are not move-equivalent but have the same strongly decorated Newton polygons. The graph Γ_2 has vertices of degree 2 at the vertical boundaries of the rectangle.

composition of $|e|_{\mathbb{Z}}$. We set $\alpha = (\alpha^e)_{e \in E(N)}$, and we refer to $\ddot{N}(\Gamma) := (N, \alpha)$ as the *strongly decorated Newton polygon* of Γ . The following result is shown in Section 6.5.

Proposition 1.11. *For any strongly decorated polygon \ddot{N} , there exists a move-reduced graph Γ that admits a perfect matching and satisfies $\ddot{N}(\Gamma) = \ddot{N}$.*

The moves (M1)–(M2) never change the homology of the strands and preserve the class of move-reduced graphs. Thus, if two move-reduced graphs Γ, Γ' are move-equivalent then we have $\ddot{N}(\Gamma) = \ddot{N}(\Gamma')$. One is tempted to conjecture that the converse is also true, but that is not the case; for instance, the two graphs in Figure 5 have the same strongly decorated Newton polygons, but they are not move-equivalent, since one graph is connected and the other one is not. See Figure 22 for a more subtle example. To remedy this issue, we make the following definition.

Definition 1.12. Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$ be a cyclic composition of $n = \alpha_1 + \alpha_2 + \dots + \alpha_m$. Consider a partition $\mathbf{I}(\alpha) = \{I_1, I_2, \dots, I_m\}$ of $\mathbb{Z}/n\mathbb{Z}$ into cyclic intervals of size $|I_j| = \alpha_j$ given by $I_1 = [1, \alpha_1]$, $I_2 = [\alpha_1 + 1, \alpha_1 + \alpha_2]$, etc. The *rotation number* $\text{rot}(\alpha)$ is the smallest integer $r \in [n] := \{1, 2, \dots, n\}$ such that $\sigma^r(\mathbf{I}(\alpha)) = \mathbf{I}(\alpha)$, where $\sigma : \mathbb{Z}/n\mathbb{Z} \rightarrow \mathbb{Z}/n\mathbb{Z}$ is the map sending $i \mapsto i + 1 \pmod{n}$ for all i , and $\sigma(\mathbf{I}(\alpha)) := \{\sigma(I_1), \sigma(I_2), \dots, \sigma(I_m)\}$.

For example, $\text{rot}((1, 1, 1, 1, 1, 1)) = 1$, $\text{rot}((2, 1, 2, 1)) = 3$, and $\text{rot}((2, 2, 1, 1)) = 6$. We have $\text{rot}((n)) = n$ because by convention, we distinguish between cyclic intervals $[j, j + n - 1]$ in $\mathbb{Z}/n\mathbb{Z}$ for different $j \in [n]$.

The *rotation number* of a collection $\alpha = (\alpha^e)_{e \in E}$ of cyclic compositions is given by

$$(1.1) \quad d(\alpha) := \gcd\{\text{rot}(\alpha^e) \mid e \in E\}.$$

The following is our second main result.

Theorem 1.13. *Let $\ddot{N} = (N, \alpha)$ be a strongly decorated polygon. The set of move-reduced graphs Γ satisfying $\ddot{N}(\Gamma) = \ddot{N}$ is a union of $d(\alpha)$ move-equivalence classes. Explicitly, two move-reduced graphs Γ, Γ' are move-equivalent if and only if*

$$(\ddot{N}(\Gamma), \mu(\Gamma)) = (\ddot{N}(\Gamma'), \mu(\Gamma')),$$

where $\mu(\Gamma) \in \mathbb{Z}/d(\alpha)\mathbb{Z}$ is the modular invariant defined in Section 1.5. Moreover, for any $\mu \in \mathbb{Z}/d(\alpha)\mathbb{Z}$, there exists a move-reduced graph Γ that admits a perfect matching and satisfies $(\ddot{N}(\Gamma), \mu(\Gamma)) = (\ddot{N}, \mu)$.




FIGURE 6. Computing the modular invariants (Section 1.5) of graphs Γ_1 and Γ_2 from Figure 5. See Example 1.15.

Remark 1.14. Suppose N is a convex integral polygon and $\alpha^e = (1, 1, \dots, 1)$ for each edge $e \in E(N)$. Then $d(\alpha) = 1$ so there is a unique move-equivalence class of graphs Γ with $\check{N}(\Gamma) = (N, \alpha)$. Since minimal graphs with Newton polygon N have strongly decorated Newton polygon (N, α) , this is the move-equivalence class of minimal graphs with Newton polygon N . Thus, Theorem 1.13 specializes to Theorem 1.3 in this case.

1.5. Modular invariant. We explain the construction of the modular invariant $\mu(\Gamma)$. Let Γ be move-reduced and let $\check{N}(\Gamma) = (N, \alpha)$ be its strongly decorated Newton polygon. Let $e \in E(N)$ and set $r := \text{rot}(\alpha^e)$, $n := |\alpha^e| = |e|_{\mathbb{Z}}$. Thus, r divides n . Let \mathbf{F}^e be the set of connected components of $\mathbb{T} \setminus \bigcup_{S \in \mathcal{S}^e(\Gamma)} S$, which we call e -regions. Construct a labeling $\gamma^e : \mathbf{F}^e \rightarrow \mathbb{Z}/n\mathbb{Z}$ so that for any segment of a strand $S \in \mathcal{S}^e(\Gamma)$ adjacent to e -regions F_- (resp., F_+) to the right (resp., left) of S , the labels $\gamma^e(F_-), \gamma^e(F_+) \in \mathbb{Z}/n\mathbb{Z}$ satisfy $\gamma^e(F_+) \equiv \gamma^e(F_-) + 1 \pmod{n}$.

Since the above conditions only involve differences of γ^e , there are n ways to construct such a labeling γ^e that differ by adding elements of $\mathbb{Z}/n\mathbb{Z}$. We shall choose a particular one as follows. The labeling γ^e induces a partition $\mathbf{I}(\gamma^e)$ of $\mathbb{Z}/n\mathbb{Z}$ into cyclic intervals so that for each strand $S \in \mathcal{S}^e(\Gamma)$, the associated cyclic interval contains $\gamma^e(F)$ for all $F \in \mathbf{F}^e$ appearing immediately to the right of S ; see Figure 6 and Example 1.15. Now, recall that α^e is a cyclic composition. Of all the cyclic shifts of α^e , let $\alpha^e = (\alpha_1^e, \alpha_2^e, \dots, \alpha_m^e)$ be the lexicographically maximal one, and let $\mathbf{I}(\alpha^e)$ be the associated partition of $\mathbb{Z}/n\mathbb{Z}$ into cyclic intervals from Definition 1.12. We say that the labeling γ^e is *lex-maximal* if $\mathbf{I}(\gamma^e) = \mathbf{I}(\alpha^e)$. Since $\sigma^r(\mathbf{I}(\alpha^e)) = \mathbf{I}(\alpha^e)$, we see that there are n/r lex-maximal labelings γ^e . Fix one such labeling and let $\bar{\gamma}^e : \mathbf{F}^e \rightarrow \mathbb{Z}/r\mathbb{Z}$ be obtained by taking the values of γ^e modulo r . Thus, $\bar{\gamma}^e$ does not depend on the choice of γ^e , and is an invariant of Γ .

Repeat the above procedure for all $e \in E(N)$. Let $\mathbf{F}(\Gamma)$ be the set of faces of Γ . We will construct a labeling $\gamma : \mathbf{F}(\Gamma) \rightarrow \mathbb{Z}/d\mathbb{Z}$, where $d := d(\alpha)$. For each face $F \in \mathbf{F}(\Gamma)$, we set $\gamma(F) := \sum_{e \in E(N)} \bar{\gamma}^e(F) \pmod{d}$. This is a well-defined element of $\mathbb{Z}/d\mathbb{Z}$ in view of (1.1). Moreover, any two adjacent faces F, F' of Γ are separated by two strands going in opposite directions, so $\gamma(F) = \gamma(F')$. In other words, the labeling γ is constant. By definition, its value is the *modular invariant* $\mu(\Gamma) \in \mathbb{Z}/d\mathbb{Z}$.

The moves (M1)–(M2) induce bijections between e -regions. Since all the faces involved in (M1)–(M2) except the middle face in (M1) are in the same e -regions, $\mu(\Gamma)$ is invariant under move-equivalence.

Example 1.15. Consider the graphs Γ_1 and Γ_2 from Figure 5. Let $\tilde{N} = (N, \alpha)$ be their strongly decorated Newton polygon shown in Figure 5(e). Thus, N is a line segment of length 4, and let $e = (4, 0)$ and $e' = (-4, 0)$ be the two edges of N . We have $\alpha := \alpha^e = \alpha^{e'} = (2, 2)$, and $\text{rot}(\alpha) = 2$. Examples of lex-maximal labelings γ^e and $\gamma^{e'}$ for Γ_1 and Γ_2 are shown in Figure 6(a–b, d–e). The labeling γ for Γ_1 and Γ_2 is obtained by taking the labeling $\gamma^e + \gamma^{e'}$ shown in Figure 6(c, f) modulo $d(\alpha) = 2$. We see that in fact $\gamma(F_1) = 0 \in \mathbb{Z}/2\mathbb{Z}$ is even for each face F_1 of Γ_1 , while $\gamma(F_2) = 1 \in \mathbb{Z}/2\mathbb{Z}$ is odd for each face F_2 of Γ_2 . Therefore, $\mu(\Gamma_1) = 0 \in \mathbb{Z}/2\mathbb{Z}$ and $\mu(\Gamma_2) = 1 \in \mathbb{Z}/2\mathbb{Z}$, which is consistent with Theorem 1.13 since the graphs Γ_1 and Γ_2 are not move-equivalent.

1.6. Overview of the proof. We shall proceed by relating bipartite graphs embedded in \mathbb{T} to elements of the *double affine symmetric group*, i.e., pairs of affine permutations. In Sections 4 and 5, we show the following result.

Theorem 1.16. *For any move-reduced graph Γ without isolated vertices, exactly one of the following holds:*

- (i) Γ has a single strand that is a simple zero-homology loop. In this case, Γ has no perfect matchings and has a different number of black and white vertices.
- (ii) Γ is move-equivalent to a graph Γ' such that, for a suitable choice of the fundamental domain, each strand $S \in \mathcal{S}(\Gamma')$ with $[S] = (i, j)$ intersects the vertical line $x = 0$ minimally, i.e., exactly $|i|$ times.

In part (i), a *zero-homology loop* is a strand S satisfying $[S] = 0$, and a zero-homology loop S is called *simple* if the lift of S to \mathbb{R}^2 under the covering map $\mathbb{R}^2 \rightarrow \mathbb{T}$ is a simple (i.e., non self-intersecting) closed curve; see e.g. Figure 4(a). In part (ii), choosing a fundamental domain corresponds to the standard $\text{SL}_2(\mathbb{Z})$ -action on the Newton polygon of Γ .

In Section 6.2, we show that if (ii) holds, then Γ' can be put into a particular form called an *affine plabic fence*. Such graphs correspond to shuffles of reduced words of two affine permutations on commuting sets of indices (see Section 6). In Sections 2 and 3, we study the associated conjugation problem for the affine symmetric group, relying on the results of [HN14, Mar20]. Finally, we complete the proof in Sections 6.5–6.6.

1.7. Previous results. The idea of relating bipartite graphs embedded in \mathbb{T} to conjugation of double affine permutations is not new and appears in [LP13, FM16, GSZ20]. A discussion of graphs that are move-reduced but not minimal in the sense of [GK13], and in particular the graph Γ_2 in Figure 3(b), appears in [FM16, Section 8.3]. Graphs that are move-reduced but not minimal also appear in the context of open Toda chains [GSV11, Wil16].

In [GSZ20, Section 4.4], the authors also consider the problem of classifying move-reduced graphs and their move-equivalence classes. They associate a weakly decorated Newton polygon to each graph and prove a lemma classifying conjugacy classes in the double affine symmetric group. However, this classification does not imply a classification of move-reduced bipartite graphs and their move-equivalence classes. The reason is that the moves (M1)–(M2) correspond only to particular kinds of conjugation in the affine symmetric group (see Definition 2.2), not to arbitrary conjugation. This discrepancy leads us to studying strongly decorated Newton polygons and modular invariants.

We also note that in [GSZ20, Section 4.4], the authors rely on Theorem 1.16 and refer to [FM16] for its proof. However, the argument in [FM16, Section 4.1] only applies to graphs whose strands go monotonously from left to right, that is, have strictly increasing




FIGURE 7. Arrow diagrams of affine permutations; see Sections 2.1 and 2.3.

horizontal coordinate. With that assumption, the conclusion of Theorem 1.16 (that each strand intersects the line $x = 0$ minimally) becomes vacuously true.

Acknowledgments. We thank Xuhua He for bringing the paper [Mar20] to our attention. We also thank Timothée Marquis for discussions related to [Mar20] and for updating and extending [Mar20] to the generality suited for our needs (cf. Remarks 2.15 and 3.3). The first author is grateful to Thomas Lam for ideas that originated during the development of [GL21], which were influential for our overall proof strategy and specifically for the arguments in Sections 2 and 3. Finally, we thank the anonymous referee and Niklas Affolter for their valuable suggestions and careful reading of the text.

2. AFFINE PERMUTATIONS, CYCLES, AND SLOPES

As we will explain in Section 6.2, Theorem 1.16 allows one to recast bipartite graphs embedded in \mathbb{T} and their moves into certain conjugation moves on pairs of affine permutations. In this and the next section, we develop the properties of affine permutations needed to complete the proofs of our main results.

Our proof strategy is inspired by that of [Mar20]. The reader familiar with the theory of affine Coxeter groups and their reflection representations is encouraged to consult Remarks 2.15 and 3.3.

2.1. Background and notation. Let $n \geq 1$ and recall that $[n] := \{1, 2, \dots, n\}$. An *affine permutation* is a bijection $f : \mathbb{Z} \rightarrow \mathbb{Z}$ satisfying $f(i + n) = f(i) + n$ for all $i \in \mathbb{Z}$. The group of affine permutations is denoted \tilde{S}_n (where the group operation is given by composition of maps $\mathbb{Z} \rightarrow \mathbb{Z}$). For $f \in \tilde{S}_n$, set

$$n(f) := n, \quad k(f) := \frac{1}{n} \sum_{i \in [n]} (f(i) - i), \quad \text{and} \quad d(f) := \gcd(k(f), n(f)).$$

It is known (see Remark 2.1 below) that $k(f)$ is always an integer. We have $\tilde{S}_n = \bigsqcup_{k \in \mathbb{Z}} \tilde{S}_n^{(k)}$, where $\tilde{S}_n^{(k)} := \{f \in \tilde{S}_n \mid k(f) = k\}$. For $f \in \tilde{S}_n$, let $\bar{f} \in S_n$ be the unique permutation (i.e., bijection $[n] \rightarrow [n]$) satisfying $\bar{f}(i) \equiv f(i) \pmod{n}$ for all $i \in [n]$. For $k \in \mathbb{Z}$, we denote by $f_{k,n} \in \tilde{S}_n^{(k)}$ the affine permutation sending $i \mapsto i + k$ for all $i \in \mathbb{Z}$. The affine permutation f can be written in *window notation* as $[f(1), f(2), \dots, f(n)]$, which completely determines $f(i)$ for all $i \in \mathbb{Z}$.

The group $\tilde{S}_n^{(0)}$ is a Coxeter group with generators $\Pi := \{s_i \mid i \in [n]\}$, where the affine permutation $s_i : \mathbb{Z} \rightarrow \mathbb{Z}$ sends $i \mapsto i + 1$ and $i + 1 \mapsto i$ (and therefore it swaps $i + dn$ with $i + 1 + dn$ for all $d \in \mathbb{Z}$), and $j \mapsto j$ for $j \not\equiv i, i + 1 \pmod{n}$. For $i \in \mathbb{Z}$, we let $s_i := s_{\bar{i}}$ where $\bar{i} \in [n]$ satisfies $\bar{i} \equiv i \pmod{n}$.




FIGURE 8. The two affine permutations on the left are c -equivalent. The affine permutations f, f' on the right satisfy $f \xrightarrow{s_i} f'$ but are not c -equivalent. See Definition 2.2. Figure reproduced from [GL21, Figure 5].

The group $\widetilde{S}_n^{(0)}$ is also known as the *affine Weyl group of type \widetilde{A}_{n-1}* . Let $\Lambda := f_{1,n} \in \widetilde{S}_n^{(1)}$. Thus, $\widetilde{S}_n = \widetilde{S}_n^{(0)} \rtimes \langle \Lambda \rangle$. We will also be interested in the quotient group $\widehat{S}_n := \widetilde{S}_n / \langle \Lambda^n = \text{id} \rangle$, known as the *extended affine Weyl group of type \widetilde{A}_{n-1}* . The group $\widetilde{S}_n^{(0)}$ is a subgroup of both \widetilde{S}_n and \widehat{S}_n . We denote by $\sigma : \widetilde{S}_n \rightarrow \widehat{S}_n$ the *rotation operator* given by $\sigma(f) := \Lambda f \Lambda^{-1}$.

Let $f \in \widetilde{S}_n$. Define

$$\begin{aligned} \text{Inv}(f) &:= \{(i, j) \in \mathbb{Z} \times \mathbb{Z} \mid i < j \text{ and } f(i) > f(j)\}, \\ \ell(f) &:= \#\{(i, j) \in \mathbb{Z} \times \mathbb{Z} \mid i < j, f(i) > f(j), \text{ and } i \in [n]\}. \end{aligned}$$

The *standard arrow diagram* of f is obtained by drawing an arrow $(i/n, 1) \rightarrow (f(i)/n, 0)$ for all $i \in \mathbb{Z}$; see Figure 7(a) for an example when $f = [7, -1, 2, 5, 8, 3, 11]$ in window notation. The set $\text{Inv}(f)$ consists of pairs of crossing arrows, and $\ell(f)$ counts the number of crossing arrows modulo the equivalence relation generated by $(i, j) \sim (i + n, j + n)$ for all $i, j \in \mathbb{Z}$. Alternatively, $\ell(f)$ is the minimal integer l such that f can be written as a product $f = s_{i_1} s_{i_2} \cdots s_{i_l} \Lambda^k$ for some indices i_1, i_2, \dots, i_l, k ; in this case, $s_{i_1} s_{i_2} \cdots s_{i_l} \Lambda^k$ is called a *reduced expression* for f . For the example in Figure 7(a), we have

$$k(f) = 1, \quad \ell(f) = 11, \quad \text{and} \quad f = s_3 s_4 s_6 s_7 s_2 s_5 s_6 s_1 s_4 s_3 s_2 \Lambda.$$

Remark 2.1. In general, the integer $k(f)$ is equal to the signed number of intersections of the arrows with one of the dashed vertical lines.

Following [GP93, GKP00, He07, He10, HN14, Mar20], for $f, f' \in \widetilde{S}_n$, we write $f \xrightarrow{s_i} f'$ if $f' = s_i f s_i$ and $\ell(f') \leq \ell(f)$. We write $f \rightarrow f'$ if there exists a sequence $f = f_0, f_1, \dots, f_m = f'$ of affine permutations such that for each $j \in [m]$, we have $f_{j-1} \xrightarrow{s_i} f_j$ for some $i \in [n]$.

Definition 2.2. We say that $f, f' \in \widetilde{S}_n$ are *c-equivalent* if $f \rightarrow f'$ and $f' \rightarrow f$. In this case, we write $f \overset{c}{\sim} f'$.

This terminology is borrowed from [GL21]. See Figure 8.

When talking about conjugacy classes, we always mean $\widetilde{S}_n^{(0)}$ -conjugacy classes, which we will usually denote by \mathcal{O} . Given a conjugacy class \mathcal{O} , let \mathcal{O}_{\min} be the set of elements of \mathcal{O} of minimal length. We have the following important result.

Theorem 2.3 ([HN14, Theorem 2.9]). *Let $f \in \widetilde{S}_n$ and let \mathcal{O} be the $\widetilde{S}_n^{(0)}$ -conjugacy class containing f . Then there exists $f' \in \mathcal{O}_{\min}$ such that $f \rightarrow f'$.*

Definition 2.4. We say that $f \in \widetilde{S}_n$ is *c-reduced* if for all $f' \in \widetilde{S}_n$ such that $f \rightarrow f'$, we have $\ell(f) = \ell(f')$ (or equivalently, $f' \rightarrow f$).

The following result follows immediately from Theorem 2.3.

Corollary 2.5. *An affine permutation $f \in \tilde{S}_n$ is c -reduced if and only if it has minimal length in its conjugacy class \mathcal{O} (i.e., $f \in \mathcal{O}_{\min}$).*

It is clear that \simeq yields an equivalence relation on the set of c -reduced elements in \tilde{S}_n . The goal of Sections 2 and 3 is to give a solution to the following problem.

Problem 2.6. Determine the structure of c -equivalence classes of c -reduced elements in \tilde{S}_n .

2.2. Cycles and slopes. A set $C \subset \mathbb{Z}$ is called n -periodic if for all $i \in \mathbb{Z}$, we have $i \in C$ if and only if $i + n \in C$.

Definition 2.7. Let $f \in \tilde{S}_n$. A set $C \subset \mathbb{Z}$ is called f -closed if it is nonempty, n -periodic, and for all $i \in \mathbb{Z}$, we have $i \in C$ if and only if $f(i) \in C$.

Definition 2.8. Let C be an f -closed set. Because it is n -periodic, the set $C \cap [n]$ is nonempty. Let $n_f(C) := \#(C \cap [n])$. There exists a unique order-preserving bijection $r_C : C \rightarrow \mathbb{Z}$ sending $\min(C \cap [n])$ to $1 \in \mathbb{Z}$. The restriction $f|_C \in \tilde{S}_{n_f(C)}$ is an affine permutation defined by

$$(2.1) \quad f|_C := r_C \circ f \circ r_C^{-1}.$$

Given an f -closed set C , we let

$$n_f(C) = n(f|_C), \quad k_f(C) := k(f|_C), \quad d_f(C) := d(f|_C), \quad \text{and} \quad \nu_f(C) := \frac{k_f(C)}{n_f(C)}.$$

The rational number $\nu_f(C)$ is called the *slope* of C . Thus, we have $f|_C \in \tilde{S}_{n'}^{(k')}$ for $n' = n_f(C)$ and $k' = k_f(C)$.

Definition 2.9. A *cycle* of f is a minimal by inclusion f -closed set C . The set of cycles of f is denoted \mathcal{C}_f .

Thus, the cycles of f are in bijection with the cycles of \bar{f} , and a nonempty subset of \mathbb{Z} is f -closed if and only if it is a disjoint union of cycles of f . For $i \in \mathbb{Z}$, we write $\nu_f(i) := \nu_f(C)$, where C is the cycle of f containing i .

Example 2.10. Let $f = [7, -1, 2, 5, 8, 3, 11]$ in window notation be the affine permutation in Figure 7(a). Then f has two cycles: \mathcal{C} (resp., \mathcal{C}') consists of all $i \in \mathbb{Z}$ congruent to 1, 4, 5, 7 (resp., to 2, 3, 6) modulo $n = 7$. We have

$$\begin{array}{llll} n_f(\mathcal{C}) = 4, & k_f(\mathcal{C}) = 2, & d_f(\mathcal{C}) = 2, & \nu_f(\mathcal{C}) = 1/2, \\ n_f(\mathcal{C}') = 3, & k_f(\mathcal{C}') = -1, & d_f(\mathcal{C}') = 1, & \nu_f(\mathcal{C}') = -1/3. \end{array}$$

Given $f \in \tilde{S}_n$ and $\nu \in \mathbb{Q}$, we set

$$\begin{aligned} \mathcal{C}_f(\nu) &:= \{C \in \mathcal{C}_f \mid \nu_f(C) = \nu\}, & \mathcal{C}_{f,\nu} &:= \bigsqcup_{C \in \mathcal{C}_f(\nu)} C, \\ n_f(\nu) &:= \sum_{C \in \mathcal{C}_f(\nu)} n_f(C), & k_f(\nu) &:= \sum_{C \in \mathcal{C}_f(\nu)} k_f(C), & d_f(\nu) &:= \sum_{C \in \mathcal{C}_f(\nu)} d_f(C). \end{aligned}$$




FIGURE 9. The arrow diagram $D_f(x)$ is topologically equivalent to the standard arrow diagram of $g_x f g_x^{-1}$; see Example 2.12.

For $f \in \widetilde{S}_n$, we set

$$\boldsymbol{\nu}_f := \{\nu \in \mathbb{Q} \mid C_{f,\nu} \text{ is nonempty}\}; \quad \text{therefore,} \quad \bigsqcup_{\nu \in \boldsymbol{\nu}_f} C_{f,\nu} = \mathbb{Z}.$$

For $\nu \in \boldsymbol{\nu}_f$, we have $\nu = k_f(\nu)/n_f(\nu)$ and $\gcd(k_f(\nu), n_f(\nu)) = d_f(\nu)$.

Definition 2.11. Let $\boldsymbol{\lambda}^f := (\lambda^{f,\nu})_{\nu \in \boldsymbol{\nu}_f}$, where $\lambda^{f,\nu}$ is the integer partition of $d_f(\nu)$ induced by $(d_f(C))_{C \in C_{f,\nu}}$.

2.3. Arrow diagrams. Let

$$\mathcal{L}'_n := \{x : \mathbb{Z} \rightarrow \mathbb{R} \mid x_{i+n} = x_i + 1 \text{ for all } i \in \mathbb{Z}\}.$$

We may identify \mathcal{L}'_n with \mathbb{R}^n via a map $x \mapsto (x_1, \dots, x_n)$. Let \mathcal{L}_n be the quotient of the vector space $\mathcal{L}'_n \cong \mathbb{R}^n$ by the linear span of the vector $(1, 1, \dots, 1) \in \mathbb{R}^n$. Thus, two elements $x, y \in \mathcal{L}'_n$ represent the same element of \mathcal{L}_n if and only if there is a constant c such that $x_i - y_i = c$ for all $i \in \mathbb{Z}$.

For any $g \in \widetilde{S}_n$, we have a point $\frac{1}{n}g \in \mathcal{L}_n$ sending $i \mapsto \frac{1}{n}g(i)$. To a point $x \in \mathcal{L}_n$, we associate a *labeled point configuration* $D(x)$, that is, a collection of labeled points on the real line: a point labeled $i \in \mathbb{Z}$ is located at coordinate x_i . We denote $\text{Im}(x) := \{x_i \mid i \in \mathbb{Z}\} \subset \mathbb{R}$.

Recall the notion of a standard arrow diagram from Section 2.1. More generally, to each $f \in \widetilde{S}_n$ and $x \in \mathcal{L}_n$ one can associate an *arrow diagram* $D_f(x)$ obtained by drawing an arrow $(x_i, 1) \rightarrow (x_{f(i)}, 0)$ for each $i \in \mathbb{Z}$. For example, the standard arrow diagram of f is just $D_f(x)$ for $x = \frac{1}{n}\text{id}$, where $\text{id} \in \widetilde{S}_n^{(0)}$ is the identity map.

We say that $x \in \mathcal{L}_n$ is *generic* if $x_i \neq x_j$ for all $i \neq j \in \mathbb{Z}$. We denote by \mathcal{L}_n° the set of generic elements of \mathcal{L}_n . The *cutoff point* for $x \in \mathcal{L}_n^\circ$ is the midpoint of the interval of all $c \in \mathbb{R} \setminus \text{Im}(x)$ such that

$$(2.2) \quad \#\{i \leq 0 \mid x_i > c\} = \#\{i \geq 1 \mid x_i < c\}.$$

For $x \in \mathcal{L}_n^\circ$, we let g_x be the affine permutation in $\widetilde{S}_n^{(0)}$ such that for all $i, j \in \mathbb{Z}$, $g_x(i) < g_x(j)$ if and only if $x_i < x_j$. Explicitly, if $c \in \mathbb{R}$ is the cutoff point for x and $i_1, i_2, \dots, i_n \in \mathbb{Z}$ are such that $c < x_{i_1} < x_{i_2} < \dots < x_{i_n} < c + 1$ then we have $g_x^{-1} = [i_1, i_2, \dots, i_n]$ in window notation; cf. Example 2.12 below.

For $f \in \widetilde{S}_n$ and $x \in \mathcal{L}_n^\circ$, the arrow diagram $D_f(x)$ is topologically equivalent to the standard arrow diagram of $g_x f g_x^{-1}$. That is, we have an order-preserving bijection $\phi_x := x \circ g_x^{-1} : \mathbb{Z} \rightarrow \text{Im}(x)$ such that $(i, j) \in \text{Inv}(g_x f g_x^{-1})$ if and only if the arrows starting at $(\phi_x(i), 1)$ and $(\phi_x(j), 1)$ cross in $D_f(x)$.

Example 2.12. Let $f = [7, -1, 2, 5, 8, 3, 11]$ be the affine permutation shown in Figure 7(a). An example of the arrow diagram $D_f(x)$ for some $x \in \mathcal{L}_n^\circ$ is shown in Figure 7(b) and also in Figure 9(a). We have $g_x^{-1} = [2, 3, 1, 4, 6, 5, 7]$ in window notation, which is obtained by reading off the labels between the two vertical dashed lines. These dashed lines are located at positions c and $c+1$, where c is the cutoff point of x . We find $g_x = s_5 s_2 s_1 \in \widetilde{S}_n^{(0)}$, and thus $g_x f g_x^{-1} = [-2, 1, 7, 6, 2, 10, 11]$ in window notation. The standard arrow diagram of $g_x f g_x^{-1}$ is shown in Figure 9(b). Comparing it with $D_f(x)$ shown in Figure 9(a), we see that indeed the two arrow diagrams are topologically equivalent (modulo a relabeling of the points given by the map ϕ_x).

We think of an arrow diagram $D_f(x)$ for $(f, x) \in \widetilde{S}_n \times \mathcal{L}_n^\circ$ as a “geometric realization” of the affine permutation $g_x f g_x^{-1}$, and extend our definitions and notation to this case. For example, we denote by $\ell_f(x) := \ell(g_x f g_x^{-1})$ the number of crossings in $D_f(x)$ modulo the shift $\frac{1}{n}(i, j) \mapsto \frac{1}{n}(i+n, j+n)$. For $f \in \widetilde{S}_n$ and $x, x' \in \mathcal{L}_n^\circ$, write $D_f(x) \rightarrow D_f(x')$ if $g_x f g_x^{-1} \rightarrow g_{x'} f g_{x'}^{-1}$, and $D_f(x) \overset{c}{\sim} D_f(x')$ if $g_x f g_x^{-1} \overset{c}{\sim} g_{x'} f g_{x'}^{-1}$. We say that $D_f(x)$ is c -reduced if so is $g_x f g_x^{-1}$.

We say that x is *almost generic* if there exist $(i_0, j_0) \in \mathbb{Z}^2$ such that for all $i \neq j$, we have $x_i \neq x_j$ unless $\{i, j\} = \{i_0 + dn, j_0 + dn\}$ for some $d \in \mathbb{Z}$. Thus, $D_f(x) \rightarrow D_f(x')$ if there exists a continuous curve $x(t) \in \mathcal{L}_n$, $t \in [0, 1]$, such that $x(0) = x$, $x(1) = x'$, $x(t)$ is almost generic for t in some finite set B and generic for $t \in [0, 1] \setminus B$, and $\ell_f(x(t))$ is a weakly decreasing function on $[0, 1] \setminus B$.

2.4. ϵ -straight arrow diagrams. Fix $f \in \widetilde{S}_n$. Recall that for $C \in \mathcal{C}_f$ and $i \in C$, we write $\nu_f(i) := \nu_f(C)$. For $\epsilon > 0$ and $x \in \mathcal{L}_n$, we say that the arrow diagram $D_f(x)$ is ϵ -straight if for all $i \in \mathbb{Z}$, $x_{f(i)}$ is ϵ -close to $x_i + \nu_f(i)$. For example, the arrow diagram $D_f(x)$ shown in Figure 7(b) is ϵ -straight for some $0 < \epsilon < 0.15$.

Denote by $\text{Str}_\epsilon(f)$ the set of ϵ -straight elements in \mathcal{L}_n :

$$\text{Str}_\epsilon(f) := \{x \in \mathcal{L}_n : |x_{f(i)} - (x_i + \nu_f(i))| \leq \epsilon \text{ for all } i \in \mathbb{Z}\}.$$

We set $\text{Str}_\epsilon^\circ(f) := \text{Str}_\epsilon(f) \cap \mathcal{L}_n^\circ$.

The following result is an analog of [Mar20, Lemma 6.8(1)]; see also [Mar18, Lemma 5.4] and [Mar14, Proposition 3.4]. See Remark 2.15 below for the relation between our results and those of Marquis.

Proposition 2.13. *For any $f \in \widetilde{S}_n$, $x \in \mathcal{L}_n^\circ$, and $\epsilon > 0$, there exists $y \in \text{Str}_\epsilon^\circ(f)$ such that $D_f(x) \rightarrow D_f(y)$.*

Example 2.14. The diagram in Figure 7(a) can be continuously deformed into the diagram in Figure 7(b). During the deformation, the point labeled 1 passes to the right through the points labeled 2, 3 while the point labeled 5 passes to the right through the point labeled 6. The resulting sequence of swaps is recorded in the reduced word for $g_x = s_5 s_2 s_1$; cf. Example 2.12.

Proof. We will find a smooth curve $x(t), t \in \mathbb{R}_{\geq 0}$ in \mathcal{L}_n such that $x(0) = x$ and such that we can take $y := x(t)$ for t sufficiently large. The curve will be defined via the following linear system of first order ordinary differential equations (ODEs):

$$(2.3) \quad \partial x_i(t) / \partial t = x_{f(i)}(t) - x_i(t), \quad \text{for all } i \in \mathbb{Z}.$$

Rewriting each $x_i(t)$ in terms of $(x_j(t))_{j \in [n]}$, we obtain an $n \times n$ inhomogeneous linear system of ODEs. It splits into independent systems for each cycle of f .

Fix a single cycle C of f , and let $m := n_f(C)$. We have an $m \times m$ system of ODEs of the form $\partial z(t)/\partial t = Az(t) + b$, for a constant $m \times m$ matrix A and a constant vector $b \in \mathbb{R}^m$. Let $w := \tilde{f}|_C \in S_m$ be the permutation obtained by taking $f|_C$ modulo m . The permutation matrix P_w of w has eigenvalues $e^{2\pi ir/m}$ for $r = 0, 1, \dots, m-1$. We have $A = P_w - I_m$, where I_m is an $m \times m$ identity matrix. Thus, the eigenvalues of A are $\lambda_r := e^{2\pi ir/m} - 1$ for $r = 0, 1, \dots, m-1$. (In particular, they are all distinct and have nonpositive real part.) A general solution to the homogeneous system $\partial z(t)/\partial t = Az(t)$ is then a linear combination of vector-valued functions of the form $\exp(\lambda_r t)z_r$, where z_r is the eigenvector of A corresponding to λ_r .

One eigenvalue of A is $\lambda_0 = 0$, and the corresponding eigenvector is $z_0 := (1, 1, \dots, 1)^T$. The vector b is a 0, 1-vector with 1's in positions corresponding to $i \in [n] \cap C$ such that $f(i) > n$. In particular, the sum of coordinates of b is $k_f(C)$, and thus $\nu_f(C)z_0 - b$ belongs to the image of A . Letting \tilde{z}_0 be one of its preimages under A , we see that $z(t) = \nu_f(C)tz_0 - \tilde{z}_0$ is a solution to the inhomogeneous system $\partial z(t)/\partial t = Az(t) + b$. Thus, an arbitrary solution differs from it by a linear combination of the functions $\exp(\lambda_r t)z_r$, each of which is constant (for $r = 0$) or decays exponentially (for $r \neq 0$).

It follows that for t large enough and $i \in \mathbb{Z}$, we have $x_i(t) = \nu_f(i)t + o(t)$, and $\partial x_i(t)/\partial t = \nu_f(i) + o(1)$. By (2.3), we get $x(t) \in \text{Str}_\epsilon(f)$ for all t sufficiently large. It is also clear that for t outside a discrete set, we have $x(t) \in \text{Str}_\epsilon^\circ(f)$.

Since $x = x(0)$ was generic, we can change it slightly so that each point $x(t)$ is almost generic for t in some discrete set B and generic for $t \in [0, \infty) \setminus B$. We claim that $\ell_f(x(t))$ is weakly decreasing for $t \in [0, \infty) \setminus B$. Indeed, let $t_0 \in B$ be such that $x_i(t_0) = x_j(t_0)$ for some $i, j \in \mathbb{Z}$. Since $x(t_0)$ is almost generic, we have $x_{f(i)}(t_0) \neq x_{f(j)}(t_0)$.¹ Thus, $\partial x_i(t)/\partial t \neq \partial x_j(t)/\partial t$ at $t = t_0$. Suppose that $\partial x_i(t)/\partial t > \partial x_j(t)/\partial t$ at $t = t_0$, so $x_{f(i)}(t_0) > x_{f(j)}(t_0)$. Then $x_i(t^-) < x_j(t^-)$ and $x_i(t^+) > x_j(t^+)$ for some $t^- < t_0 < t^+$ very close to t_0 . We still have $x_{f(i)}(t^-) > x_{f(j)}(t^-)$. Thus, the arrows starting at $x_i(t^-)$ and $x_j(t^-)$ form a crossing in $D_f(x(t^-))$ but do not form a crossing in $D_f(x(t^+))$. Therefore $\ell_f(x(t^-)) \geq \ell_f(x(t^+))$. Note that the arrows starting at $x_{f^{-1}(i)}(t^\pm)$ and $x_{f^{-1}(j)}(t^\pm)$ may or may not form a crossing in $D_f(x(t^\pm))$. Depending on that, the difference $\ell_f(x(t^-)) - \ell_f(x(t^+))$ is either zero or one. \square

Remark 2.15. Our constructions can be translated into the well-studied geometric setup as we now explain. The group $\tilde{S}_n^{(0)}$ acts simply transitively on the set Σ of chambers of an infinite hyperplane arrangement $\{x_i = x_j + k \mid i \neq j \in [n], k \in \mathbb{Z}\}$ in $\mathbb{R}^n / \langle (1, 1, \dots, 1) \rangle$. Choosing a distinguished fundamental chamber C_0 , the map $g \mapsto gC_0$ yields a bijection $\tilde{S}_n^{(0)} \xrightarrow{\sim} \Sigma$. Identifying $\mathcal{L}_n \xrightarrow{\sim} \mathbb{R}^n / \langle (1, 1, \dots, 1) \rangle$ by a linear isomorphism sending $x \mapsto (x_1, \dots, x_n)$, the \tilde{S}_n -action on \mathcal{L}_n coincides with its action on $\mathbb{R}^n / \langle (1, 1, \dots, 1) \rangle$. For $g \in \tilde{S}_n^{(0)}$, the point $\frac{1}{n}g$ gets mapped to the barycenter of the corresponding (simplicial) chamber gC_0 . An element $x \in \mathcal{L}_n$ is generic if and only if it belongs to the interior of a chamber, and almost generic if and only if it belongs to the interior of a facet of a chamber. The set $\text{Str}_\epsilon(f)$ for $\epsilon = 0$ equals the set denoted $\text{Min}(f)$ in [Mar20]. The map sending f to the tuple $(f|_\nu)_{\nu \in \nu_f}$ of its restrictions is the map denoted π_{Σ^n} in [Mar20] (whose image is an element of finite order; cf. Lemma 3.6). Our proof strategy may be considered an adaptation of [Mar20, Proof of

¹This statement is true unless $f(i) = i + kn$ and $f(j) = j + kn$ for some k . But in that case, we have $x_i \neq x_j$ (because x was generic) and $x_i(t') = i + kt'$, $x_j(t') = j + kt'$ for all $t' \geq 0$, so $x_i(t') \neq x_j(t')$.

Proposition 6.20]: given an arbitrary chamber C , construct a walk from C to a chamber intersecting $\text{Min}(f)$, and then use the projection π_{Σ^n} to obtain an element of finite order. The notion of a modular invariant was inspired by [Mar20, Part (A $_{\ell}^{(1)}$) of Theorem 10.12].

Remark 2.16. One key point that allows for a significant simplification in our approach in type A (compared to the approach of [Mar20] for arbitrary Coxeter groups) is a new proof of Proposition 2.13 using ODEs. We hope that this argument can be of independent interest. It appears to generalize to affine Weyl groups but not to arbitrary Coxeter groups.

2.5. Vector configurations and conjugacy. We return to Problem 2.6. Our first goal is to describe $\tilde{S}_n^{(0)}$ -conjugacy classes in \tilde{S}_n .

Let $f \in \tilde{S}_n$. For $\nu \in \nu_f$ and $C \in \mathbf{C}_f$, set

$$e_f(\nu) := (n_f(\nu), k_f(\nu)) \quad \text{and} \quad e_f(C) := (n_f(C), k_f(C)).$$

Clearly, $e_f(\nu) = \sum_{C \in \mathbf{C}_f(\nu)} e_f(C)$ is a sum of collinear vectors, and their integer lengths are given by $|e_f(C)|_{\mathbb{Z}} = d_f(C)$, so $|e_f(\nu)|_{\mathbb{Z}} = d_f(\nu)$. We let $E_f := (e_f(\nu))_{\nu \in \nu_f}$ be the *vector configuration* associated to f . By analogy with Definition 1.1, we call $\dot{E}_f := (E_f, \boldsymbol{\lambda}^f)$ the *weakly decorated vector configuration* associated to f , where $\boldsymbol{\lambda}^f$ was introduced in Definition 2.11.

Proposition 2.17. *Let $f, f' \in \tilde{S}_n$. Then f is $\tilde{S}_n^{(0)}$ -conjugate to f' if and only if $\dot{E}_f = \dot{E}_{f'}$.*

Proof. Since \dot{E}_f depends only on the cycles of f and their slopes, it is clearly invariant under conjugation, which shows the “only if” direction. Suppose now that $f, f' \in \tilde{S}_n$ are such that $\dot{E}_f = \dot{E}_{f'}$. Because the permutations $\bar{f}, \bar{f}' \in S_n$ have the same cycle type, they are conjugate in S_n . We may therefore apply S_n -conjugation to f' (permuting the cycles along the way) to obtain an element f'' such that $\bar{f} = \bar{f}''$ (in particular, f and f'' have the same sets of cycles), and such that for each cycle C of f , we have $n_f(C) = n_{f''}(C)$ and $k_f(C) = k_{f''}(C)$. Let $t_{e_i} \in \tilde{S}_n$ be the affine permutation sending $i \mapsto i + n$ and $j \mapsto j$ for $j \not\equiv i \pmod{n}$, called a *translation element*. Thus, $t_{e_i - e_j} := t_{e_i} t_{e_j}^{-1}$ belongs to $\tilde{S}_n^{(0)}$, and we see that f can be obtained from f'' via conjugations by such elements $t_{e_i - e_j}$ for i, j belonging to the same cycle of f . \square

2.6. A characterization of minimal-length elements. Our next goal is to give an explicit characterization of c -reduced affine permutations; see Corollary 2.20.

Given two subsets $A, B \subset \mathbb{R}^2$, define their *Minkowski sum* by $A + B := \{a + b \mid a \in A, b \in B\}$. Given a vector configuration $E = \{e_1, e_2, \dots, e_m\} \subset \mathbb{Z}^2$, the associated *zonotope* $\mathcal{Z}(E)$ is the convex polygon in \mathbb{R}^2 obtained as the Minkowski sum of line segments

$$\mathcal{Z}(E) := [0, e_1] + [0, e_2] + \dots + [0, e_m].$$

For $e_1, e_2 \in \mathbb{R}^2$, recall that $\det(e_1, e_2)$ is the determinant of the 2×2 matrix with columns e_1, e_2 . The following formula for the area of $\mathcal{Z}(E)$ is well known [McM84]:

$$\text{Area}(\mathcal{Z}(E)) = \sum_{1 \leq i < j \leq m} |\det(e_i, e_j)|.$$

Recall the notion of $\text{exc}(\boldsymbol{\lambda})$ from Definition 1.4.

Lemma 2.18. *Let $f \in \tilde{S}_n$. Then f is c -reduced if and only if*

$$(2.4) \quad \ell(f) = \text{Area}(\mathcal{Z}(E)) + \text{exc}(\boldsymbol{\lambda}), \quad \text{where } \dot{E}_f = (E, \boldsymbol{\lambda}).$$

In the proof of the lemma, we will count inversions $(j, j') \in \text{Inv}(f)$ according to the cycles containing j and j' .

Definition 2.19. Given two cycles $C, C' \in \mathbf{C}_f$, their *ordered crossing number* is defined as

$$\text{xing}(C, C') := \#\{(j, j') \in \text{Inv}(f) \mid j \in [n] \cap C \text{ and } j' \in C'\}.$$

Thus, we have $\sum_{C, C' \in \mathbf{C}_f} \text{xing}(C, C') = \ell(f)$.

Proof of Lemma 2.18. Denote the right-hand side of (2.4) by $\ell(\dot{E}_f)$. First, we show that for any $f \in \tilde{S}_n$, we have $\ell(f) \geq \ell(\dot{E}_f)$. Observe that if $g \in \tilde{S}_n^{(k)}$ has a single cycle then $\ell(g) \geq d(g) - 1$ (where $d(g) = \gcd(k, n)$), because the map $f_{k,n} = \Lambda^k$ has $\gcd(k, n)$ cycles, and for each $i \in [n]$, $s_i g$ has either one more or one less cycle than g . Thus, each cycle C of f contributes at least $d_f(C) - 1$ to $\ell(f)$:

$$(2.5) \quad \text{xing}(C, C) \geq d_f(C) - 1.$$

It follows that for each $\nu \in \nu_f$, we have

$$\sum_{C \in \mathbf{C}_f(\nu)} \text{xing}(C, C) \geq \text{exc}(\lambda^{f, \nu}).$$

To each cycle C we can associate a piecewise-linear curve $P^{(f)}(C)$ in \mathbb{R}^2 obtained by choosing some $i \in C$ and joining the points $p_d := (d, \frac{1}{n} f^d(i))$ for $d = 0, 1, \dots, n_f(C)$; cf. [GL21, Section 4]. We have $p_0 = (0, \frac{i}{n})$ and $p_{n_f(C)} = (n_f(C), \frac{i}{n} + k_f(C))$, thus $P^{(f)}(C)$ gives rise to a closed curve on $\mathbb{T} = \mathbb{R}^2/\mathbb{Z}^2$ with homology $e_f(C) = (n_f(C), k_f(C))$. It is well known that given integers n', k', n'', k'' with $k'/n' > k''/n''$, a curve in \mathbb{T} with homology (n', k') intersects a curve with homology (n'', k'') from below at least $\left| \det \begin{pmatrix} n' & k' \\ n'' & k'' \end{pmatrix} \right|$ times. Thus, given cycles $C \neq C'$, we have

$$(2.6) \quad \text{xing}(C, C') \geq \begin{cases} 0, & \text{if } \nu_f(C) \leq \nu_f(C'); \\ |\det(e_f(C), e_f(C'))|, & \text{otherwise.} \end{cases}$$

We have shown that $\ell(f) \geq \ell(\dot{E}_f)$.

Conversely, consider a weakly decorated vector configuration $\dot{E} = (E, \boldsymbol{\lambda})$. By Proposition 2.17, $\mathcal{O} := \{f \in \tilde{S}_n \mid \dot{E}_f = \dot{E}\}$ is an $\tilde{S}_n^{(0)}$ -conjugacy class. By Corollary 2.5, $f \in \mathcal{O}$ is c-reduced if and only if $f \in \mathcal{O}_{\min}$. We have shown above that for any $f \in \mathcal{O}$, $\ell(f) \geq \ell(\dot{E})$. It remains to construct $g \in \mathcal{O}$ such that $\ell(g) = \ell(\dot{E})$. Such an affine permutation will be constructed in Section 3.2. \square

Corollary 2.20. *Let $f \in \tilde{S}_n$. Then f is c-reduced if and only if all of the following conditions are satisfied.*

- (1) For each $C \in \mathbf{C}_f$, $\text{xing}(C, C) = d_f(C) - 1$.
- (2) For each $C \neq C'$ in \mathbf{C}_f , we have

$$\text{xing}(C, C') = \begin{cases} 0, & \text{if } \nu_f(C) \leq \nu_f(C'); \\ |\det(e_f(C), e_f(C'))|, & \text{otherwise.} \end{cases}$$

Proof. We have lower bounds on $\text{xing}(C, C)$ and $\text{xing}(C, C')$ given by (2.5)–(2.6). Moreover, we showed in Lemma 2.18 that f is c-reduced if and only if all of these inequalities are equalities. \square

Remark 2.21. Corollary 2.20 was obtained jointly with Thomas Lam during the development of [GL21].

Corollary 2.22. *If $f \in \tilde{S}_n$ is c -reduced and $C \subset \mathbb{Z}$ is f -closed then $f|_C$ is c -reduced.*

3. THE STRUCTURE OF C -EQUIVALENCE CLASSES

The goal of this section is to give a complete description of c -equivalence classes of c -reduced affine permutations; see Theorem 3.2.

3.1. Cyclic compositions. Let $f \in \tilde{S}_n$ be c -reduced. Fix a slope $\nu \in \nu_f$. By Corollary 2.20, we have $\text{xing}(C, C') = 0$ for all $C \neq C'$ in $\mathbf{C}_f(\nu)$. We thus get a natural cyclic order on the set $\mathbf{C}_f(\nu)$ induced by the cyclic order on $[n] \cong \mathbb{Z}/n\mathbb{Z}$. Recall that $\sum_{C \in \mathbf{C}_f(\nu)} d_f(C) = d_f(\nu)$. In other words, the cyclic order on $\mathbf{C}_f(\nu)$ yields a cyclic composition $\alpha^{f, \nu}$ of $d_f(\nu)$. Letting $\alpha_f := (\alpha^{f, \nu})_{\nu \in \nu_f}$, we consider the *strongly decorated vector configuration* $\ddot{E}_f := (E_f, \alpha_f)$.

Lemma 3.1. *Let $f, f' \in \tilde{S}_n$ be c -reduced. If $f \stackrel{\circ}{\sim} f'$ then $\ddot{E}_f = \ddot{E}_{f'}$.*

Proof. By assumption, $\ell(f) = \ell(f')$. It suffices to consider the case $f' = s_i f s_i$ for some $i \in [n]$. By Proposition 2.17, we have $\dot{E}_f = \dot{E}_{f'}$. Thus, we only need to check that the relative order on $\mathbf{C}_f(\nu)$ is preserved for each slope $\nu \in \nu_f$. Let C_i, C_{i+1} be the cycles passing through $i, i+1$, respectively, and suppose that $\nu_f(C_i) = \nu_f(C_{i+1}) = \nu$. In order to apply a c -equivalence $f \mapsto f'$, there must be a crossing between the arrows $i \mapsto g(i)$ and $i+1 \mapsto g(i+1)$ for either $g = f$ or $g = f'$; see Figure 8. But because f, f' are c -reduced, by Corollary 2.20, they cannot have crossings between different cycles of the same slope. Thus we must have $C_i = C_{i+1}$, so the relative order on $\mathbf{C}_f(\nu)$ is preserved. \square

To give the converse to Lemma 3.1, we need to consider *modular invariants* discussed in Section 1.5. Recall from Definition 1.12 that for a cyclic composition α , we have the rotation number $\text{rot}(\alpha)$, and for a family α of cyclic compositions, $d(\alpha)$ is the greatest common divisor of their rotation numbers.

Given a conjugacy class \mathcal{O} and a strongly decorated vector configuration \ddot{E} , let

$$\mathcal{O}_{\min}[\ddot{E}] := \{f \in \mathcal{O}_{\min} \mid \ddot{E}_f = \ddot{E}\}.$$

The goal of this section is to prove the following result.

Theorem 3.2. *Let $f \in \tilde{S}_n$ be c -reduced. Let \mathcal{O} be the $\tilde{S}_n^{(0)}$ -conjugacy class of f . Then $\mathcal{O}_{\min}[\ddot{E}_f]$ is a union of $d(\alpha_f)$ -many c -equivalence classes. Moreover, for any two c -reduced $f, f' \in \mathcal{O}_{\min}$, we have*

$$(3.1) \quad f \stackrel{\circ}{\sim} f' \iff (\ddot{E}_f, \mu(f)) = (\ddot{E}_{f'}, \mu(f')),$$

where $\mu(f) \in \mathbb{Z}/d(\alpha_f)\mathbb{Z}$ is the modular invariant defined in (3.2).

Remark 3.3. Alternatively, Theorem 3.2 may be deduced from the recently updated version of [Mar20, Theorem B].

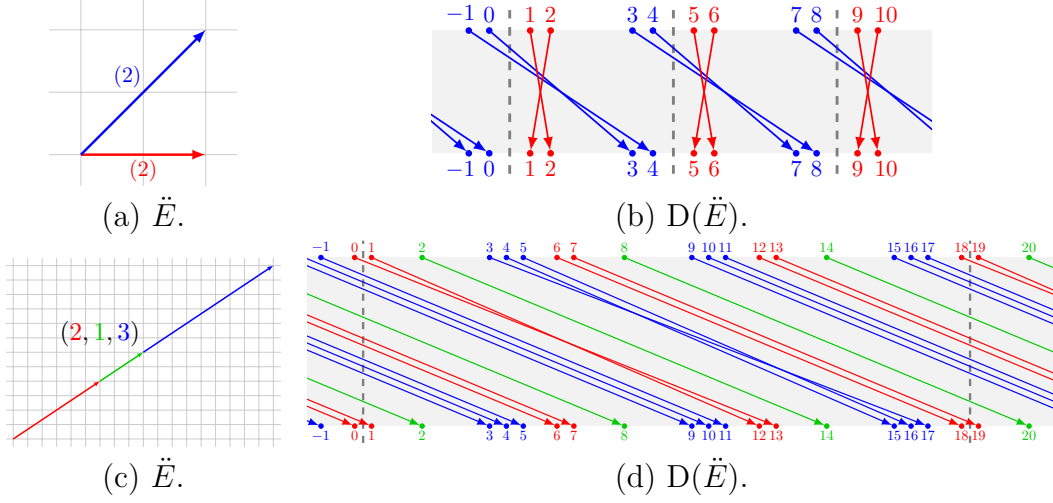


FIGURE 10. A strongly decorated vector configuration (left) and an ϵ -straight arrow diagram (right); see Section 3.2.

3.2. Constructing ϵ -straight diagrams explicitly. Let $\ddot{E} = (E, \alpha)$ be a strongly decorated vector configuration and fix a small $\epsilon > 0$. Our goal is to construct an ϵ -straight arrow diagram $D(\ddot{E}) = D_g(x)$ for some $x \in \mathcal{L}_n^\circ$ and c -reduced $g \in \tilde{S}_n$ with $\ddot{E}_g = \ddot{E}$. We start with an example and then proceed with a formal description.

Example 3.4. Let $\ddot{E} = (E, \alpha)$ denote the strongly decorated vector configuration shown in Figure 10(a). Thus, the vectors in E are $e_1 = (2, 0)$, $e_2 = (2, 2)$, and $\alpha^{e_1} = \alpha^{e_2} = (2)$. An ϵ -straight arrow diagram $D(\ddot{E})$ is shown in Figure 10(b). On the other hand, if $\ddot{E} = (E, \alpha)$ is the strongly decorated vector configuration shown in Figure 10(c), then E consists of a single vector $e = (18, 12)$ decorated by a cyclic composition $\alpha^e = (2, 1, 3)$. The associated ϵ -straight arrow diagram $D(\ddot{E})$ is constructed in Figure 10(d).

For a vector $e = (a, b) \in \mathbb{Z}^2$, we denote $n(e) := a$ and $k(e) := b$. For $e \in E$, let $\nu(e) = k(e)/n(e)$ denote its slope. Assume that $n(e) > 0$ for all $e \in E$. Let $\alpha = (\alpha^e)_{e \in E}$ and $\alpha^e = (\alpha_1^e, \dots, \alpha_{m_e}^e)$. Consider the circle \mathbb{R}/\mathbb{Z} and choose a collection of *starting points* $\mathbf{p} = (\bar{p}_i^e)_{e \in E, i \in [m_e]}$, where $\bar{p}_i^e \in \mathbb{R}/\mathbb{Z}$. Let $\bar{P}_i^e := \{\bar{p}_i^e + r\nu(e) \mid r \in \mathbb{Z}\} \subset \mathbb{R}/\mathbb{Z}$ be the set containing \bar{p}_i^e and consisting of $n(e)/|e|_{\mathbb{Z}}$ equally spaced points on a circle. We choose \mathbf{p} so that we additionally have:

- (1) $\text{dist}_{\mathbb{R}/\mathbb{Z}}(\bar{P}_i^e, \bar{P}_{i'}^{e'}) > \epsilon$ for all $(e, i) \neq (e', i')$; and
- (2) the points $(\bar{p}_1^e, \bar{p}_2^e, \dots, \bar{p}_{m_e}^e)$ are cyclically ordered in \mathbb{R}/\mathbb{Z} .

Now, for each fixed $e \in E$ and $i \in [m_e]$, we construct an arrow diagram D_i^e . Let $P_i^e \subset \mathbb{R}$ be the preimage of \bar{P}_i^e under the projection $\mathbb{R} \rightarrow \mathbb{R}/\mathbb{Z}$, and choose $p' \in P_i^e$. Set $d := \alpha_i^e$. For each $r \in [d]$, set $p'_r := p' + \frac{r\epsilon}{d}$. We refer to the points $(p'_r)_{r \in [d]}$ as the *block* associated to p' , and denote by $P'_{e,i} := P_i^e + \frac{\epsilon}{d}[d]$ the set of points in all such blocks. Let $\bar{p}' \in \bar{P}_i^e$ be the image of p' in \mathbb{R}/\mathbb{Z} . If $\bar{p}' \not\equiv \bar{p}_i^e \pmod{\mathbb{Z}}$ then we draw an arrow $(p'_r, 1) \rightarrow (p'_r + \nu(e), 0)$ for each $r \in [d]$. Otherwise, we draw an arrow $(p'_r, 1) \rightarrow (p'_{\sigma(r)} + \nu(e), 0)$ for each $r \in [d]$, where $\sigma = (1\ 2 \ \dots \ d) \in S_d$ is a d -cycle. The resulting arrow diagram is denoted D_i^e .

Let $\mathbf{P} := \bigsqcup_{e \in E, i \in [m_e]} P'_{e,i} \subset \mathbb{R}$ be the resulting set of points, and let $D(\ddot{E}) := \bigcup_{i \in [m_e]} D_i^e$ be the corresponding arrow diagram. Let $x : \mathbb{Z} \rightarrow \mathbf{P}$ be an order-preserving map. Then

there exists a unique affine permutation $g \in \widetilde{S}_n$ such that $D(\ddot{E}) = D_g(x)$. By construction, $\ddot{E}_g = \ddot{E}$ and $\ell(g) = \ell(\ddot{E})$, which completes the proof of Lemma 2.18. By Lemma 2.18, g is c -reduced.

3.3. Affine permutations of constant slope.

Definition 3.5. Let $f \in \widetilde{S}_n$ and $\nu \in \mathbb{Q}$. We say that f is of *constant slope* ν if $\nu_f = \{\nu\}$. (That is, if all cycles of f are of the same slope ν .)

It is clear that if $f \in \widetilde{S}_n^{(k)}$ is of constant slope ν then we must have $\nu = k/n$.

Recall that \widehat{S}_n is a quotient of \widetilde{S}_n by Λ^n . We denote the quotient map $\widetilde{S}_n \rightarrow \widehat{S}_n$ by $f \mapsto \widehat{f}$.

Lemma 3.6. *Let $f \in \widetilde{S}_n$. Then $\widehat{f} \in \widehat{S}_n$ has finite order if and only if f is of constant slope.*

Proof. Let N be the least common multiple of $n_f(C)$ for all $C \in \mathbf{C}_f$. Then f^N is a translation element; that is, $f^N(i) = i + d_i n$ for all $i \in \mathbb{Z}$, where $(d_i)_{i \in \mathbb{Z}}$ is some sequence of integers. Explicitly, if $i \in C$ then $d_i = N\nu_f(C) \in \mathbb{Z}$. This implies the result. \square

Let $f \in \widetilde{S}_n^{(k)}$ be c -reduced and of constant slope, and set $d := \gcd(k, n)$. By Corollary 2.20, the arrows between different cycles of f do not cross. Therefore, for each cycle $C \in \mathbf{C}_f$, we have $C = C + d$ as subsets of \mathbb{Z} . Denoting by $I_C \subset \mathbb{Z}/d\mathbb{Z}$ the image of C under the map $\mathbb{Z} \mapsto \mathbb{Z}/d\mathbb{Z}$, we get a partition $\mathbf{I}(f) = \{I_C \mid C \in \mathbf{C}_f\}$ of $\mathbb{Z}/d\mathbb{Z}$ into cyclic intervals.² It is clear that $\mathbf{I}(f)$ is invariant under c -equivalence.

Proposition 3.7 ([Mar20, Proposition A]). *Let $f, f' \in \widetilde{S}_n^{(k)}$ be c -reduced and of constant slope. Then*

$$f \stackrel{c}{\sim} f' \quad \text{if and only if} \quad \mathbf{I}(f) = \mathbf{I}(f').$$

We say that a cyclic composition $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$ is written in *normal form* if the sequence $(\alpha_1, \alpha_2, \dots, \alpha_m)$ is lexicographically maximal out of all sequences obtained by rotating α , i.e., $(\alpha_r, \alpha_{r+1}, \dots, \alpha_m, \alpha_1, \dots, \alpha_{r-1})$ for $r \in [m]$. As in Definition 1.12, we associate to α a partition $\mathbf{I}(\alpha) = (I_1, I_2, \dots, I_m)$ of $\mathbb{Z}/d\mathbb{Z}$ (where $d = \alpha_1 + \alpha_2 + \dots + \alpha_m$) into cyclic intervals given by $I_1 = [1, \alpha_1]$, $I_2 = [\alpha_1 + 1, \alpha_1 + \alpha_2]$, etc.

Note that if $\alpha = \alpha^{f, \nu}$ then we have $d = \alpha_1 + \alpha_2 + \dots + \alpha_m = \gcd(k, n)$, and therefore we have two partitions $\mathbf{I}(\alpha)$ and $\mathbf{I}(f)$ of $\mathbb{Z}/d\mathbb{Z}$ into cyclic intervals. These partitions are related by a rotation σ^r of $\mathbb{Z}/d\mathbb{Z}$ for some r ; however, this rotation is only defined up to a symmetry of $\mathbf{I}(\alpha)$, i.e., up to $\sigma^{\text{rot}(\alpha)}$. (Here, $\text{rot}(\alpha)$ divides d .)

Definition 3.8. Let $f \in \widetilde{S}_n^{(k)}$ be c -reduced of constant slope $\nu = k/n$, and let $\alpha := \alpha^{f, \nu}$ be written in normal form. The *modular invariant* $\mu(f) \in \mathbb{Z}/\text{rot}(\alpha)\mathbb{Z}$ is the unique element such that $\sigma^{\mu(f)}(\mathbf{I}(\alpha)) = \mathbf{I}(f)$.

Corollary 3.9. *Let $f, f' \in \widetilde{S}_n^{(k)}$ be c -reduced and of constant slope $\nu = k/n$. Then*

$$f \stackrel{c}{\sim} f' \quad \text{if and only if} \quad (\alpha^{f, \nu}, \mu(f)) = (\alpha^{f', \nu}, \mu(f')).$$

²The case where f is a single cycle requires special care. As mentioned after Definition 1.12, we distinguish between different cyclic intervals $[j, j+d-1]$ of $\mathbb{Z}/d\mathbb{Z}$. Topologically, the standard arrow diagram of f (viewed as a union of arrows) will be disconnected, and we choose $\mathbf{I}(f) := \{[j, j+d-1]\}$ for $j \in \mathbb{Z}/d\mathbb{Z}$ such that the points $(j, 1)$ and $(j-1, 1)$ belong to different connected components.

Proof. The \implies direction is clear since both $\alpha^{f,\nu}$ and $\mu(f)$ are invariant under c-equivalence. Conversely, having $\alpha^{f,\nu} = \alpha^{f',\nu}$ implies that $\mathbf{I}(f)$ and $\mathbf{I}(f')$ coincide up to cyclic shift, and $\mu(f) = \mu(f')$ guarantees that $\mathbf{I}(f) = \mathbf{I}(f')$. The result then follows from Proposition 3.7. \square

3.4. Finishing the proof. For $f \in \widetilde{S}_n$ and $\nu \in \boldsymbol{\nu}_f$, let $f|_\nu := f|_{C_{f,\nu}}$. Thus, $f|_\nu$ has constant slope ν . If in addition f is c-reduced then by Corollary 2.22, so is $f|_\nu$. In this case, recall from Definition 3.8 that the modular invariant $\mu(f|_\nu)$ is an element of $\mathbb{Z}/\text{rot}(\alpha^{f,\nu})\mathbb{Z}$. By (1.1), $d(\boldsymbol{\alpha}_f)$ is defined as the greatest common divisor of the numbers $d(\alpha^{f,\nu})$ over all $\nu \in \boldsymbol{\nu}_f$.

Definition 3.10. For c-reduced $f \in \widetilde{S}_n$, define the *modular invariant* $\mu(f) \in \mathbb{Z}/d(\boldsymbol{\alpha}_f)\mathbb{Z}$ by

$$(3.2) \quad \mu(f) := \sum_{\nu \in \boldsymbol{\nu}_f} \mu(f|_\nu) \pmod{d(\boldsymbol{\alpha}_f)}.$$

Lemma 3.11. *Let $f, f' \in \widetilde{S}_n$ be c-reduced. If $f \stackrel{\circ}{\sim} f'$ then $\mu(f) = \mu(f')$.*

Proof. Suppose that $f \xrightarrow{s_i} f'$ for some $i \in [n]$. If $\nu_f(i) = \nu_f(i+1) = \nu$ then we have $f|_\nu \xrightarrow{s_i} f'|_\nu$, so $f|_\nu$ and $f'|_\nu$ are c-equivalent and thus by Corollary 3.9, we have $\mu(f|_\nu) = \mu(f'|_\nu)$. Suppose now that $\nu_f(i) \neq \nu_f(i+1)$. If $0 < i < n$ then we have $f|_\nu = f'|_\nu$ and thus $\mu(f|_\nu) = \mu(f'|_\nu)$.

It remains to consider the case $i = n$ and $\nu_f(0) \neq \nu_f(1)$. Let $\nu_0 := \nu_f(0)$, $\nu_1 := \nu_f(1)$. Since $\nu_0 \neq \nu_1$, by the definition of $f|_{\nu_0}$ in (2.1), we see that $f'|_{\nu_0} = \sigma(f|_{\nu_0})$ and $f'|_{\nu_1} = \sigma^{-1}(f|_{\nu_1})$. Here, $\sigma(g) = \Lambda g \Lambda^{-1}$ is the rotation operator introduced in Section 2.1. Thus, $\mu(f'|_{\nu_0}) = \mu(f|_{\nu_0}) + 1$ and $\mu(f'|_{\nu_1}) = \mu(f|_{\nu_1}) - 1$, and the sum in (3.2) remains the same. \square

We will need one more tool for working with ϵ -straight diagrams from Section 2.4. Fix c-reduced $f \in \widetilde{S}_n$ and small $\epsilon > 0$. For $x \in \text{Str}_\epsilon^\circ(f)$ such that $D_f(x)$ is c-reduced, recall from Corollary 2.20 that $D_f(x)$ contains no crossings between distinct cycles of the same slope.

Definition 3.12. Let $x \in \text{Str}_\epsilon^\circ(f)$ be c-reduced and let $\mathbf{a} := (a_C)_{C \in \mathcal{C}_f}$ be a family of real numbers associated to the cycles of f . Consider a curve $x(t)$, $t \geq 0$, given for each $i \in \mathbb{Z}$ by $x_i(t) = x_i + ta_C$, where C is the cycle containing i . Let $T > 0$ be such that for $t \in [0, T]$, $x_i(t) \neq x_j(t)$ for any $i \neq j$ such that $\nu_f(i) = \nu_f(j)$. In this case, we say that $x' := x(T)$ is obtained from $x = x(0)$ by *block-shifting*.

In other words, block-shifting allows us to move the collections of points $(x_i)_{i \in C}$ independently for each cycle C , subject to the condition that two cycles of the same slope never collide. It is clear that for ϵ sufficiently small, if $x \in \text{Str}_\epsilon^\circ(f)$ is c-reduced and $x' \in \text{Str}_\epsilon^\circ(f)$ is obtained from x by block-shifting then x' is c-reduced and $D_f(x) \rightarrow D_f(x')$.

Proof of Theorem 3.2. The \implies direction follows from Lemmas 3.1 and 3.11.

For the \impliedby direction, let $f, f' \in \mathcal{O}_{\min}$. Thus, $f' = f_0 f f_0^{-1}$ for some $f_0 \in \widetilde{S}_n^{(0)}$. Let $x, x' \in \text{Str}_\epsilon^\circ(f)$ be obtained from $\frac{1}{n} \text{id}, \frac{1}{n} f_0 \in \mathcal{L}_n^\circ$ via Proposition 2.13 so that $D_f(\frac{1}{n} \text{id}) \rightarrow D_f(x)$ and $D_{f'}(\frac{1}{n} \text{id}) = D_f(\frac{1}{n} f_0) \rightarrow D_f(x')$.

Set $h := g_x f g_x^{-1}$ and $h' := g_{x'} f g_{x'}^{-1}$, where g_x was defined after (2.2). We have $f \stackrel{\circ}{\sim} h$ and $f' \stackrel{\circ}{\sim} h'$. Since f, f' are c-reduced, so are $D_f(x), D_f(x')$ and h, h' . Since $\ddot{E}_f = \ddot{E}_{f'}$, and thus $\ddot{E}_h = \ddot{E}_{h'}$, we see that the partitions $\mathbf{I}(h|_\nu)$ and $\mathbf{I}(h'|_\nu)$ of $\mathbb{Z}/d_f(\nu)\mathbb{Z}$ into cyclic intervals differ by rotation for all $\nu \in \boldsymbol{\nu}_f$. Our goal is to apply block-shifting to x with the aim of

achieving $\mathbf{I}(h|_\nu) = \mathbf{I}(h'|_\nu)$ for all $\nu \in \nu_f$. To do so, consider the following operation on the partitions $(\mathbf{I}(h|_\nu))_{\nu \in \nu_f}$:

$$(3.3) \quad \text{for some } \nu \neq \nu' \text{ in } \nu_f, \text{ replace } \mathbf{I}(h|_\nu) \mapsto \sigma(\mathbf{I}(h|_\nu)) \text{ and } \mathbf{I}(h|_{\nu'}) \mapsto \sigma^{-1}(\mathbf{I}(h|_{\nu'})).$$

We first explain how to obtain (3.3) via block-shifting.

Applying block-shifting to $D_f(x)$ corresponds to applying a sequence $h \xrightarrow{s_{i_1}} h_1 \xrightarrow{s_{i_2}} \dots \xrightarrow{s_{i_l}} h'$ of c-equivalences. In order to control how each restriction $h|_\nu$ changes under such operations, we need to distinguish between the cases $i_j = n$ and $i_j \neq n$ as we did in the proof of Lemma 3.11.

Recall the notion of the cutoff point from (2.2). Suppose that applying block-shifting to x switches the positions of adjacent points x_j and x_k for some j, k . If the cutoff point of x is between $x_j + d$ and $x_k + d$ for some $d \in \mathbb{Z}$ then the corresponding c-equivalence corresponds to s_n , otherwise it corresponds to s_i for $i \in [n - 1]$.

Consider slopes $\nu \neq \nu'$ in ν_f . We may apply block-shifting to move $C_{f,\nu}$ (resp., $C_{f,\nu'}$) to the right (resp., left) so that no point in $\text{Im}(x)$ passes through the cutoff point c of x , until c is located in an interval of $\mathbb{R} \setminus \text{Im}(x)$ between a point of $C_{f,\nu}$ and a point of $C_{f,\nu'}$. We may then shift $C_{f,\nu}$ (resp., $C_{f,\nu'}$) further to the right (resp., left) until these two points swap places. This corresponds to replacing $h|_\nu$ with $\sigma(h|_\nu)$ and $h|_{\nu'}$ with $\sigma^{-1}(h|_{\nu'})$, which results in applying (3.3) to $\mathbf{I}(h|_\nu)$ and $\mathbf{I}(h|_{\nu'})$.

Recall that for $\nu \in \nu_f$, by the definition of $\text{rot}(\alpha^{f,\nu})$, we have $\sigma^{\text{rot}(\alpha^{f,\nu})}(\mathbf{I}(h|_\nu)) = \mathbf{I}(h|_\nu)$. Let $d := d(\alpha_f) = \gcd\{\alpha^{f,\nu} \mid \nu \in \nu_f\}$. Write $d = \sum_{\nu' \in \nu_f} a_{\nu'} \text{rot}(\alpha^{f,\nu'})$ for some integers $a_{\nu'}$. Then, for each fixed $\nu \in \nu_f$, we have $(a_\nu \text{rot}(\alpha^{f,\nu}) - d) + \sum_{\nu' \neq \nu} a_{\nu'} \text{rot}(\alpha^{f,\nu'}) = 0$. Consider a vector $\mathbf{c} \in \mathbb{Z}^{\nu_f}$ with coordinates $c_\nu := a_\nu \text{rot}(\alpha^{f,\nu}) - d$ and $c_{\nu'} := a_{\nu'} \text{rot}(\alpha^{f,\nu'})$ for $\nu' \neq \nu$. Since the coordinates of \mathbf{c} sum up to zero, \mathbf{c} may be written as a \mathbb{Z} -linear combination of vectors of the form $\mathbf{e}_{\nu_1} - \mathbf{e}_{\nu_2}$, where $(\mathbf{e}_{\nu'})_{\nu' \in \nu_f}$ is the standard basis of \mathbb{Z}^{ν_f} . We can therefore use (3.3) to rotate each $\mathbf{I}(h|_{\nu'})$, $\nu' \in \nu_f$, by the corresponding coefficient $c_{\nu'}$. The result of this operation is

$$(3.4) \quad \text{replace } \mathbf{I}(h|_\nu) \mapsto \sigma^{-d}(\mathbf{I}(h|_\nu)) \text{ and preserve } \mathbf{I}(h|_{\nu'}) \text{ for all } \nu' \neq \nu.$$

Fix $\nu \in \nu_f$. Applying (3.3), we can achieve $\mathbf{I}(h|_{\nu'}) = \mathbf{I}(h'|_{\nu'})$ for all $\nu' \neq \nu$. Recall that $\ddot{E}_h = \ddot{E}_{h'}$, and thus $\mu(h) = \mu(h')$. We see that $\mathbf{I}(h|_\nu)$ and $\mathbf{I}(h'|_\nu)$ differ by rotation by a multiple of d , so applying (3.4), we achieve $\mathbf{I}(h|_\nu) = \mathbf{I}(h'|_\nu)$.

By definition, we have $h = g_x f g_x^{-1}$ and $h' = g_{x'} f g_{x'}^{-1}$. Thus, for $r \in \mathbb{Z}$, we have $\sigma^r(h) = g_{x \circ \Lambda^r} \sigma^r(f) g_{x \circ \Lambda^r}^{-1}$ and $\sigma^r(h') = g_{x' \circ \Lambda^r} \sigma^r(f) g_{x' \circ \Lambda^r}^{-1}$. As labeled point configurations, x and $x \circ \Lambda^r$ differ by adding r to all labels. Thus, the cutoff point of $x \circ \Lambda^r$ is to the left of the cutoff point of x , and there are precisely r points of $\text{Im}(x) = \text{Im}(x \circ \Lambda^r)$ between them. Assuming ϵ is small enough, we can choose r such that the cutoff points of $x \circ \Lambda^r$ and $x' \circ \Lambda^r$ are not ϵ -close to any point in $\text{Im}(x \circ \Lambda^r) \cup \text{Im}(x' \circ \Lambda^r)$. We then replace h, h', x, x' with $\sigma^r(h), \sigma^r(h'), x \circ \Lambda^r, x' \circ \Lambda^r$, respectively. Clearly, showing $h \stackrel{\circ}{\sim} h'$ is equivalent to showing $\sigma^r(h) \stackrel{\circ}{\sim} \sigma^r(h')$.

Applying block-shifting to x so that all points in $\text{Im}(x)$ are far away from the cutoff point of x results in a sequence of c-equivalences of the form $h \mapsto s_i h s_i$ for i not equal to 0 modulo n . Therefore such block-shifting preserves the condition that $\mathbf{I}(h|_\nu) = \mathbf{I}(h'|_\nu)$. We apply such block-shifting to both x and x' so that for $\nu \neq \nu'$, no point in $C_{f,\nu}$ is ϵ -close to a point in $C_{f,\nu'}$.

By Proposition 3.7, we have $h|_\nu \stackrel{\circ}{\sim} h'|_\nu$ for all $\nu \in \boldsymbol{\nu}_f$. Since $h|_\nu$ and $h'|_\nu$ are c-reduced, they have no crossings between different cycles. Thus, each c-equivalence in $h|_\nu \stackrel{\circ}{\sim} h'|_\nu$ swaps points x_i and x_j from the same cycle. Such points are ϵ -close together in x because the arrows incident to them form a crossing and have slopes ϵ -close to ν . Thus, the interval between x_i and x_j contains neither the cutoff point of x nor any point of $C_{f,\nu'}$ for $\nu' \neq \nu$. It cannot contain any point of $C_{f,\nu}$ because x_i and x_j are adjacent in $h|_\nu$. We can therefore lift the c-equivalence $h|_\nu \stackrel{\circ}{\sim} h'|_\nu$ to h and assume that $h|_\nu = h'|_\nu$ for all $\nu \in \boldsymbol{\nu}_f$. Thus, the only difference between h and h' is the relative ordering of the blocks of ϵ -close points in x, x' corresponding to the different slopes. Applying further block-shifting, we can permute such blocks until their orders coincide, which gives the desired c-equivalence $h \stackrel{\circ}{\sim} h'$. \square

4. PLABIC GRAPHS AND TRIPLE-CROSSING DIAGRAMS

We discuss the properties of bipartite graphs embedded in \mathbb{T} and explain how to recast them in the equivalent languages of *plabic graphs* [Pos06] and *triple-crossing diagrams* [Thu17].

4.1. Triple-crossing diagrams in the disk. The results of this section were independently discovered by [Pos06] and [Thu17]. We state the results in terms of Thurston's notion of triple-crossing diagrams.

Definition 4.1. A *triple-crossing diagram* D in the disk $\mathbb{D} := [0, 1]^2$ is a smooth immersion of a disjoint union of oriented circles and closed intervals into \mathbb{D} , defined up to isotopy. The image of a connected component is called a *strand*. The image of a circle is called a *loop* and the image of a closed interval is called an *arc*. The immersion is required to satisfy the following conditions:

- (1) Three strands cross at each intersection point. We call these intersection points *triple crossings*.
- (2) The endpoints of the arcs are distinct points on the boundary of \mathbb{D} , and there are no other points of D on the boundary of \mathbb{D} .
- (3) The orientations of the strands induce consistent orientations on the boundaries of the faces of D .

Here, a *face* of D is a connected component of $\mathbb{D} \setminus D$. Property (3) implies that around every triple crossing, the orientations of strands alternate in and out, and that the orientations of the end points alternate in and out along the boundary of \mathbb{D} . If D has n arcs, then it has $2n$ boundary points, and the connectivity of the arcs induces a matching of the in-boundary points with the out-boundary points, called the *trip permutation* in [Pos06].

Definition 4.2. A triple-crossing diagram D in the disk \mathbb{D} is said to be *reduced* if it has the fewest number of triple crossings among all triple-crossing diagrams with the same trip permutation.

Definition 4.3. Two triple-crossing diagrams are said to be *move-equivalent* if they are related by move (M1)' in Figure 11. A triple-crossing diagram D is called *move-reduced* if it is not move-equivalent to a triple-crossing diagram D' to which one of the reduction moves (R1)'–(R2)' can be applied.

Remark 4.4. Postnikov's reduction move (R1)' in Figure 11 differs from Thurston's 1 – 0 move (see Figure 16). Postnikov's move will be more important for our eventual goal




FIGURE 11. Equivalence move $(M1)'$ and reduction moves $(R1)'$ – $(R2)'$ for triple-crossing diagrams. Each move has two possible strand orientations. $(R2)'$ removes a strand that is a simple loop.

of understanding the behavior of the dimer model under taking limits, since it preserves dimer partition functions (cf. [Pos06, Theorem 12.1]). On the other hand, Thurston's move preserves the trip permutation which allows for inductive arguments; e.g., in the proof of Theorem 1.5.. It also appears naturally in connection with double affine permutations (see Section 6.2 and Remark 6.2).

A *monogon* in \mathbb{D} is a strand with a self-intersection. A *parallel bigon* in \mathbb{D} is a pair of strands with two intersection points $x \neq y$, with both strands oriented from x to y .

Theorem 4.5 ([Pos06, Theorem 13.2 and Lemma 13.6] and [Thu17, Theorem 7]). *Let D be a triple-crossing diagram in \mathbb{D} . The following are equivalent.*

- (1) D is move-reduced;
- (2) D is reduced;
- (3) D contains no loops, monogons, or parallel bigons.

Theorem 4.6 ([Pos06, Corollary 14.7] and [Thu17, Theorem 3]). *All $n!$ matchings of in- and out-boundary points are realizable as trip permutations of move-reduced triple-crossing diagrams.*

Theorem 4.7 ([Pos06, Theorem 13.4] and [Thu17, Theorem 5]). *Any two move-reduced triple-crossing diagrams with the same trip permutation are move-equivalent.*

Each pair of in- and out-endpoints in the matching divides the boundary of \mathbb{D} into two intervals. Suppose that I is a minimal such interval with respect to inclusion. We say that a strand S whose endpoints are the endpoints of I is *boundary-parallel* if there are no triple crossings within the region between S and I .

Proposition 4.8 ([Pos06, Proof of Theorem 13.4 and Figure 13.4] and [Thu17, Lemma 12]). *Suppose I is an inclusion-minimal interval of the boundary matching of a move-reduced triple-crossing diagram D , and let S be the strand in D whose endpoints are the endpoints of I . Then, D is move-equivalent to a triple-crossing diagram D' in which S is boundary-parallel.*

4.2. Plabic graphs and triple-crossing diagrams on the torus. A *plabic graph* $\Gamma = (B \sqcup W, E)$ on a torus \mathbb{T} is a (finite) graph embedded in \mathbb{T} such that:

- (1) The vertices of Γ are colored black or white. The set of black vertices (resp., white vertices) is denoted by B (resp., W).
- (2) The set of edges of Γ is denoted by E . Each edge is incident to two vertices of opposite colors or incident to two white vertices.
- (3) The black vertices are trivalent.




FIGURE 12. The procedure to convert plabic graphs into triple-crossing diagrams and vice versa.

We identify plabic graphs that are related by contracting an edge incident to two distinct white vertices into a single white vertex. Therefore, we can assume that each white-white edge is a loop based at a white vertex.

Remark 4.9. The term *plabic graph* was originally introduced to denote *planar bicolored graphs* [Pos06]. We still use it for graphs embedded on a torus, with the main emphasis on the properties that such graphs (i) are drawn on \mathbb{T} without self-intersections, and (ii) are bicolored but not necessarily bipartite. Our definition of a plabic graph is more restrictive than that of [Pos06]. Such plabic graphs were previously studied under the name *white-partite* [GPW22, Definition 7.14] or *black-trivalent* [Gal21, Definition 8.1, Remark 8.2].

Definition 4.10. A *triple-crossing diagram* D on the torus \mathbb{T} is a smooth immersion of a disjoint union of oriented circles into \mathbb{T} . The image of a circle is called a *strand*, and the set of strands of D is denoted $\mathcal{S}(D)$. The immersion is required to satisfy the following conditions:

- (1) Three strands cross at each intersection point. We call these intersection points *triple crossings*.
- (2) The orientations of the strands induce consistent orientations on the boundaries of the faces of D .

Similarly to Definition 4.1, a *face* of D is a connected component of $\mathbb{T} \setminus D$. The property (2) implies that around every triple crossing, the orientations of the strands alternate in and out. (However, the converse need not hold if D has a non-contractible face.) Each strand S in D determines a homology class $[S] \in H_1(\mathbb{T}, \mathbb{Z}) \cong \mathbb{Z}^2$.

Lemma 4.11. *The sum of the homology classes of all strands is 0 in $H_1(\mathbb{T}, \mathbb{Z})$.*

Proof. Let R_+ (resp., R_-) denote the union of the faces of D such that the induced orientation is counterclockwise (resp., clockwise). Then, by property (2), we have that $\sum_{S \in \mathcal{S}(D)} S = \partial R_+ = -\partial R_-$ as 1-cycles in \mathbb{T} , and $\overline{R_+} \cup \overline{R_-} = \mathbb{T}$. Therefore,

$$2 \sum_{S \in \mathcal{S}(D)} [S] = [\partial R_+] - [\partial R_-] = [\partial \mathbb{T}] = 0. \quad \square$$

Remark 4.12. A triple-crossing diagram can be converted into a plabic graph and vice versa using the local procedure shown in Figure 12. When converting a triple-crossing diagram into a plabic graph, the ambiguity in the third case in Figure 12(b) is irrelevant since we identify plabic graphs related by contracting white-white edges incident to distinct white vertices. The notions of *move-reduced* and *move-equivalent* plabic graphs are given by Definition 4.3.

Remark 4.13. If Γ is a bipartite graph in \mathbb{T} with all black vertices of degree at least three, one can convert Γ into a plabic graph by applying a sequence of black uncontraction

moves (M2). When Γ has black vertices of degree zero, one, or two,³ extra care needs to be taken; see Appendix A. Conversely, any plabic graph can be converted into a bipartite graph by placing a black vertex of degree two in the middle of each white-white edge.

The notions of weakly/strongly decorated Newton polygons and modular invariants introduced in Section 1.3 for bipartite graphs extend to plabic graphs in an obvious way. Using Remark 4.12, we can transfer them to triple-crossing diagrams.

In what follows, we will prove the versions of our main results translated into the language of triple-crossing diagrams and plabic graphs. The proof for bipartite graphs follows from the results in Appendix A. For instance, in Section 6.3, we will prove the following version of Theorem 1.5 and show that it implies Theorem 1.5 in Appendix A.1.

Theorem 4.14. *Let D be a triple-crossing diagram with weakly decorated polygon $\dot{N} = (N, \boldsymbol{\lambda})$. Assume that N is not a single point. The following are equivalent.*

- (1) D is move-reduced;
- (2) D has no connected components that are contractible in \mathbb{T} and contains $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ triple crossings.

Similarly to Remark 1.8, $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ is the minimal possible number of triple crossings for a triple-crossing diagram with weakly decorated Newton polygon $(N, \boldsymbol{\lambda})$.

Let $\pi : \mathbb{R}^2 \rightarrow \mathbb{T}$ denote the universal covering map. The following result will be proved in Section 6.4.

Proposition 4.15. *Let D be a move-reduced triple-crossing diagram with Newton polygon N .*

- (1) *The preimage \tilde{D} of D under π contains no closed loops, and any lift \tilde{S} of a strand $S \in \mathcal{S}(D)$ does not intersect itself;*
- (2) *Any strand $S \in \mathcal{S}(D)$ intersects itself $|\llbracket S \rrbracket|_{\mathbb{Z}} - 1$ times;*
- (3) *Any two distinct parallel strands $S, S' \in \mathcal{S}(D)$ do not intersect, and there is no face of D that contains portions of both strands in its boundary.*

By part (3) of Proposition 4.15, there is a natural cyclic order on each set of parallel strands, so the strongly decorated Newton polygon \ddot{N} is well-defined.

5. REDUCTION TO THE CYLINDER

The goal of this section is to prove the triple-crossing diagram version of Theorem 1.16. Consider a move-reduced triple-crossing diagram D on the torus \mathbb{T} . Let \mathbb{D} be a fundamental rectangle for \mathbb{T} , and let u, d, l, r denote the up, down, left and right sides of \mathbb{D} , respectively. Identifying the u and d sides, we get a cylinder \mathbb{A} , and further identifying the l and r sides, we get a torus \mathbb{T} . We have quotient maps $\mathbb{D} \rightarrow \mathbb{A} \rightarrow \mathbb{T}$. The images of the u and d sides in \mathbb{A} or in \mathbb{T} coincide and are referred to as the $u - d$ side. Similarly, the images of the l and r sides in \mathbb{T} are referred to as the $l - r$ side.

We say that triple-crossing diagrams D and D' are *isotopic in \mathbb{T}* if there is an ambient isotopy of \mathbb{T} taking D to D' . When applying such isotopies, we fix the fundamental rectangle \mathbb{D} . Using an isotopy in \mathbb{T} if necessary, we assume that the intersections of strands with the sides of \mathbb{D} are transverse. A strand S with homology class (i, j) must intersect the $l - r$ side at least i times and the $u - d$ side at least j times. The preimage of a strand under the map $\mathbb{A} \rightarrow \mathbb{T}$ is either a union of arcs with endpoints on the boundary of \mathbb{A} or a closed loop in \mathbb{A} .

³This applies especially to the case of a degree-two black vertex connected to the same white vertex by both edges.




FIGURE 13. The move (P), pushing a boundary-parallel strand past the d -side of \mathbb{D} removing the intersection points with the d -side, and the move (T), interchanging the relative order of the endpoints of the red and blue strands along the d -side of \mathbb{D} , thereby removing the triangular region bounded by the strands and the d -side of \mathbb{D} .

5.1. Pushing strands through the boundary.

Lemma 5.1. *Suppose we have a strand S in \mathbb{D} with both endpoints on a side s of \mathbb{D} . Then, using moves and isotopy in \mathbb{T} , we can remove the endpoints of S in s without increasing the number of intersections of any other strands with the boundary of \mathbb{D} .*

Proof. Let I denote the interval in s between the endpoints of S . Suppose I is minimal with respect to inclusion among all intervals on the boundary of \mathbb{D} between endpoints of strands. Using Proposition 4.8, we make S boundary-parallel, and then apply an isotopy in \mathbb{T} pushing the strand S past the s -side of \mathbb{D} (Figure 13(a)).

If I is not minimal, we use induction on the number of intervals contained in I . Suppose I' is an inclusion-minimal interval contained in I . Using the above procedure, we can remove the endpoints of I' and thereby reduce the number of intervals contained in I . \square

Definition 5.2. We call the above procedure *move (P)*; see Figure 13(a).

Lemma 5.3 ([GK13, Figure 12]). *Suppose we have a pair of strands in \mathbb{D} that have consecutive in- or out-endpoints on a side of \mathbb{D} , and moreover, suppose that these two strands cross in \mathbb{D} . Then, the relative order of the two endpoints along s can be reversed using moves and isotopy in \mathbb{T} without increasing the number of intersections of any other strands with the boundary of \mathbb{D} .*

We remove the consecutiveness assumption from Lemma 5.3.

Lemma 5.4. *Suppose we have a pair (S, S') of strands with endpoints on a side s of \mathbb{D} with the same orientation (i.e., both in or both out), and suppose that S, S' cross in \mathbb{D} . Then, the relative order of the endpoints of S, S' in s can be reversed using moves and isotopy in \mathbb{T} , thereby removing the triangular region bounded by S, S' and the side s , without increasing the number of intersections of any other strands with the boundary of \mathbb{D} .*

Proof. Assume without loss of generality that both S, S' have an out-endpoint in s . Let I be the interval in s between the endpoints of S, S' . Use move (P) to remove any strands that have both endpoints in I . Then, any strand with an out-endpoint in I must cross at least one of S, S' . Let ℓ be the number of crossings formed by pairs of strands having an out-endpoint in I . Repeatedly using Lemma 5.3, we can decrease ℓ until it becomes equal to 1 and use Lemma 5.3 once more to swap the endpoints of S, S' . \square

Definition 5.5. We call the procedure in Lemma 5.4 *move (T)*; see Figure 13(b).

5.2. **Affine matchings.** Fix $n \geq 1$. Consider an infinite vertical strip \mathbb{S} with points labeled

$$(5.1) \quad \dots, A_0, A_{\bar{1}}, A_1, A_{\bar{2}}, A_2, \dots \quad \text{and} \quad \dots, B_0, B_{\bar{1}}, B_1, B_{\bar{2}}, B_2, \dots$$

on the left and the right boundary of \mathbb{S} from bottom to top, so that the points A_i, B_i are at the same height, and the points $A_{\bar{i}}, B_{\bar{i}}$ are at the same height, for each $i \in \mathbb{Z}$. Let $A := \{A_i \mid i \in \mathbb{Z}\}$, $\bar{A} := \{A_{\bar{i}} \mid i \in \mathbb{Z}\}$, $B := \{B_i \mid i \in \mathbb{Z}\}$, $\bar{B} := \{B_{\bar{i}} \mid i \in \mathbb{Z}\}$. Let $X_i + n := X_{i+n}$ and $X_{\bar{i}} + n := X_{\overline{i+n}}$ for $X \in \{A, B\}$ and $i \in \mathbb{Z}$.

Definition 5.6. An *affine matching with period n* is a bijection $\pi : A \sqcup \bar{B} \rightarrow \bar{A} \sqcup B$ such that $\pi(A_{i+n}) = \pi(A_i) + n$ and $\pi(B_{\overline{i+n}}) = \pi(B_{\bar{i}}) + n$ for all $i \in \mathbb{Z}$.

This notion is closely related to the classical notion of *affine permutations* discussed in Section 2.1. An affine matching is represented by drawing an arrow from x to $\pi(x)$ inside \mathbb{S} for all $x \in A \sqcup \bar{B}$.

A triple-crossing diagram D in \mathbb{T} gives rise to an affine matching π_D as follows. Let \mathbb{D} be a fundamental rectangle. Using an $\mathrm{SL}_2(\mathbb{Z})$ transformation, we can assume that there are no strands with homology classes in $\{0\} \times \mathbb{Z}$ other than zero-homology loops. Let n denote half the number of intersection points of strands with the $l - r$ side of \mathbb{D} (so there are $2n$ endpoints of strands on the $l - r$ side of \mathbb{D} , half of which are oriented in and half of which are oriented out). Let \mathbb{S} denote the infinite vertical strip that is the universal cover of \mathbb{A} . Then, \mathbb{S} consists of \mathbb{Z} -many copies of \mathbb{D} glued along the $u - d$ sides, which we label $\dots, \mathbb{D}_{-1}, \mathbb{D}_0, \mathbb{D}_1, \dots$ from bottom to top. Applying an isotopy in \mathbb{T} , we may assume that the bottom-most intersection point of a strand with the left side of \mathbb{D} is oriented in. Label the intersection points of strands with the boundary of \mathbb{S} as in (5.1). Thus, the points in $A \sqcup \bar{B}$ are in-endpoints and the points in $\bar{A} \sqcup B$ are out-endpoints. The connectivity of strands in \mathbb{S} determines an affine matching π_D with period n , which, moreover, has total signed number of crossings through any horizontal line equal to 0 (since the total homology of all strands is 0 by Lemma 4.11).

Remark 5.7. Each strand S in \mathbb{S} determines a word w_S in the alphabet $\{u, d, l, r\}$ whose letters from left to right record the crossings of S with the sides of \mathbb{D} as we move along the strand. Using move (P), we can assume that there are no occurrences of ud or du in w_S ; thus, we have $w_S = xy^kz$ for some $x, z \in \{l, r\}$, $y \in \{u, d\}$, and $k \geq 0$. For a strand S such that $w_S = xy^kz$, we denote by S_1, \dots, S_{k+1} the corresponding strands in \mathbb{D} .

Lemma 5.8. *If the strands S and T emanating from A_i and A_{i+1} cross in \mathbb{S} , then we can swap their endpoints A_i and A_{i+1} using moves and isotopy in \mathbb{T} without increasing the number of intersections of any other strands with the boundary of \mathbb{A} .*

Proof. By translating the fundamental rectangle, we can assume that $i = 1$. Without loss of generality, we can assume that $w_S = ru^kz$ where $k \geq 0$ and $z \in \{l, r\}$. If $w_T = rr$ or $w_T = rdv$ for some word v , then the segments S_1 and T_1 cross in \mathbb{D} and we can use move (T). Suppose $w_T = ru^mz$ with $m \geq 0$ and $z \in \{l, r\}$. Consider a crossing of S and T in \mathbb{S} . If this crossing belongs to S_j then it must also belong to T_j . Let $j \geq 1$ be the minimal index such that S_j crosses T_j . If $j \geq 2$, then applying move (T) at the $u - d$ side, we can swap the bottom endpoints of S_j and T_j so that the strands S_{j-1}, T_{j-1} would cross. We continue this process until $j = 1$ and apply move (T) at the $l - r$ side. \square

5.3. **Proof of Theorem 1.16.** We are now ready to prove Theorem 1.16.




FIGURE 14. The cases in the proof of Theorem 1.16; see Section 5.3.

Lemma 5.9. *The number of intersections of strands in D with the sides of \mathbb{A} can be made either the minimum possible (i.e., equal to $\sum_{S \in \mathcal{S}(D)} |i|$, where $[S] = (i, j)$) or equal to 2 using moves and isotopy in \mathbb{T} .*

Proof. Consider the affine matching π_D with period n equal to half the number of intersection points of strands with the $l - r$ side of \mathbb{D} and consider a strand S from A_i to $A_{\bar{j}}$ with the smallest value of $\text{dist}(A_i, A_{\bar{j}})$. Assume $i < j$. By minimality of $\text{dist}(A_i, A_{\bar{j}})$, the strand T emanating from A_{i+1} must cross the strand S . By Lemma 5.8, we can swap the endpoints A_i and A_{i+1} , decreasing $\text{dist}(A_i, A_{\bar{j}})$. Eventually, we force $\text{dist}(A_i, A_{\bar{j}})$ to be less than the height of \mathbb{D} , in which case we apply move (P) and decrease n . We proceed until either $n = 1$ or when there are no arcs from A_i to $A_{\bar{j}}$, in which case there will be also no arcs from $B_{\bar{i}}$ to B_j (for any $i, j \in \mathbb{Z}$). \square

We now study the case when $n = 1$. Since the total signed number of crossings through any horizontal line is equal to 0, we have either:

- (1) $\pi_D(A_1) = B_{k+1}$ and $\pi_D(B_{\bar{1}}) = A_{\overline{-k+1}}$ for some $k \in \mathbb{Z}$; or
- (2) $\pi_D(A_1) = A_{\overline{k+1}}$ and $\pi_D(B_{\bar{1}}) = B_{-k+1}$ for some $k \in \mathbb{Z}$.

See Figure 14. If π_D satisfies (1), then the number of intersections of strands with the sides of \mathbb{A} is minimal and equal to 2. Suppose that π_D satisfies (2). If $k \in \{0, -1\}$, we can use move (P) to remove the two intersection points, so the number of intersections of strands with sides of \mathbb{A} is minimal and equal to 0. However, if $k \notin \{0, -1\}$, the number of intersections of strands with the sides of \mathbb{A} is not minimal (since we have $\sum_{S \in \mathcal{S}(D)} |i| = 0$), and we call such a triple-crossing diagram *exceptional*. In this case, D consists of a single strand S in \mathbb{T} that is a zero-homology loop (see Figure 4 for the associated bipartite graph when $k = 1$). It is not hard to see that the strand S is simple, that is, lifts to a non self-intersecting closed curve in \mathbb{R}^2 . This is case (i) of Theorem 1.16. In order to complete the proof, we need to show that in this case, the associated bipartite graph Γ has no perfect matchings.

Proposition 5.10. *Let Γ be a move-reduced bipartite graph in \mathbb{T} . Suppose that Γ has a single strand which is a simple zero-homology loop. Then Γ has a different number of black and white vertices, and, in particular, has no perfect matchings.*

Proof. Given a closed immersed curve $\rho : S^1 \rightarrow \mathbb{R}^2$ with non-vanishing differential, we let $\text{wind}(\rho) \in \mathbb{Z}$ denote its *winding number*, which is the number counterclockwise turns made by the tangent vector of ρ . For a collection $\boldsymbol{\rho}$ of such curves such that no three curves intersect

at the same point, we let $\text{wind}(\boldsymbol{\rho}) := \sum_{\rho \in \boldsymbol{\rho}} \text{wind}(\rho)$ denote their total winding number. One can check that the total winding number $\text{wind}(\boldsymbol{\rho})$ is invariant under the skein relation

$$(5.2) \quad \begin{array}{c} \text{X} \\ \text{---} \end{array} \rightarrow \begin{array}{c} \text{---} \\ \text{---} \end{array}.$$

Let $\tilde{\Gamma}$ be the lift of Γ to the universal cover \mathbb{R}^2 of $\mathbb{T} = \mathbb{R}^2/\mathbb{Z}^2$. Let S be the unique strand of Γ . We apply a small perturbation to S so that each triple crossing is replaced by three double crossings arranged in a small counterclockwise cycle around the corresponding black vertex of Γ . After this perturbation, each white vertex of Γ is surrounded by a clockwise cycle of strands, and the strand directions alternate around each face of Γ ; see [Pos06, Lemma 14.4]. Let \tilde{S} be some lift of S to \mathbb{R}^2 . Thus, \tilde{S} is a simple closed curve, and therefore $\text{wind}(\tilde{S}) = \pm 1$, depending on whether \tilde{S} is oriented counterclockwise or clockwise. Any two lifts of S differ by a shift in \mathbb{Z}^2 . Let $N \gg 1$ be a large positive integer, and let $\boldsymbol{\rho}$ be the collection of all \mathbb{Z}^2 -shifts of \tilde{S} that are contained inside the square $[0, N]^2 \subset \mathbb{R}^2$. Because \tilde{S} has bounded size, there are at least cN^2 such shifts, for some fixed constant $c > 0$, and therefore $|\text{wind}(\boldsymbol{\rho})| \geq cN^2$. On the other hand, resolving all crossings in $\boldsymbol{\rho}$ using the skein relation (5.2), we obtain a collection $\boldsymbol{\rho}'$ of simple closed curves satisfying $\text{wind}(\boldsymbol{\rho}) = \text{wind}(\boldsymbol{\rho}')$. Each of these curves will contain a single vertex of $\tilde{\Gamma}$ inside of it. Moreover, if the vertex of $\tilde{\Gamma}$ inside $\rho' \in \boldsymbol{\rho}'$ is black (resp., white), then ρ' is oriented counterclockwise (resp., clockwise). Therefore, the difference between the numbers of black and white vertices of $\tilde{\Gamma}$ contained inside $[0, N]^2$ is of size at least cN^2 . This implies that Γ must have a different number of black and white vertices: otherwise, the difference between the numbers of black and white vertices of $\tilde{\Gamma}$ contained inside $[0, N]^2$ would be at most linear in N . \square

6. RELATING AFFINE PERMUTATIONS TO BIPARTITE GRAPHS ON A TORUS

The goal of this section is to apply the results of Sections 2 and 3 to bipartite graphs embedded in \mathbb{T} and to finish the proof of our main results, Theorems 1.5 and 1.13.

6.1. The double affine symmetric group. The *double affine symmetric group* \ddot{S}_n is generated by $S \sqcup \bar{S} \sqcup \{\Lambda\}$, where $S := \{s_i \mid i \in \mathbb{Z}/n\mathbb{Z}\}$ and $\bar{S} := \{\bar{s}_i \mid i \in \mathbb{Z}/n\mathbb{Z}\}$, subject to the relations

$$(6.1) \quad \begin{aligned} s_i s_{i+1} s_i &= s_{i+1} s_i s_{i+1}, & \Lambda s_{i+1} &= s_i \Lambda, & s_i^2 &= 1, & s_i s_j &= s_j s_i & \text{if } |i-j| > 1, \\ \bar{s}_i \bar{s}_{i+1} \bar{s}_i &= \bar{s}_{i+1} \bar{s}_i \bar{s}_{i+1}, & \Lambda \bar{s}_{i+1} &= \bar{s}_i \Lambda, & \bar{s}_i^2 &= 1, & \bar{s}_i \bar{s}_j &= \bar{s}_j \bar{s}_i & \text{if } |i-j| > 1, \\ \Lambda^n &= 1, & s_i \bar{s}_j &= \bar{s}_j s_i. \end{aligned}$$

In other words, we have an isomorphism $\ddot{S}_n \cong (\tilde{S}_n^{(0)} \times \tilde{S}_n^{(0)}) \rtimes \langle \Lambda \rangle / \langle \Lambda^n \rangle$, where Λ acts on each copy of $\tilde{S}_n^{(0)}$ by conjugation. Any element $w \in \ddot{S}_n$ can be written as a product $w = s_{i_1} s_{i_2} \cdots s_{i_l} \Lambda^k \bar{s}_{j_m} \bar{s}_{j_{m-1}} \cdots \bar{s}_{j_1}$ for some $k \in \{0, 1, \dots, n-1\}$ and $l, m \geq 0$. If $l+m$ is minimal among all such ways of writing w as a product, then $s_{i_1} s_{i_2} \cdots s_{i_l} \Lambda^k \bar{s}_{j_m} \bar{s}_{j_{m-1}} \cdots \bar{s}_{j_1}$ is called a *reduced expression* for w , and $l+m$ is called the *length* of w and denoted $\ell(w)$. Note that

$$f := s_{i_1} s_{i_2} \cdots s_{i_l} \Lambda^k \quad \text{and} \quad \bar{f} := \bar{s}_{n-j_1} \bar{s}_{n-j_2} \cdots \bar{s}_{n-j_m} \Lambda^k$$

are then reduced expressions for affine permutations $f, \bar{f} \in \tilde{S}_n$. We denote $\phi(w) := (f, \bar{f})$ and call (f, \bar{f}) the *pair of affine permutations* associated to w . We have $\ell(w) = \ell(f) + \ell(\bar{f})$. We explain the reasoning behind the formula for \bar{f} in Remark 6.3.




FIGURE 15. Plabic graphs and triple-crossing diagrams in \mathbb{A} associated to generators.




FIGURE 16. $(R1)''$ Thurston's $1 - 0$ move.

Remark 6.1. For any $k \in \mathbb{Z}$ and $f, \bar{f} \in \tilde{S}_n^{(k)}$, there exists $w \in \check{S}_n$ satisfying $\phi(w) = (f, \bar{f})$.

6.2. Relating triple-crossing diagrams in \mathbb{A} to double affine permutations. Let w be a double affine permutation and let $w_1 w_2 \cdots w_l$ be an expression for w , where $w_i \in S \sqcup \bar{S} \sqcup \{\Lambda\}$. Following Fock and Marshakov [FM16], we associate to the expression $w_1 w_2 \cdots w_l$ a triple-crossing diagram in \mathbb{A} as follows. Each generator $s \in S \sqcup \bar{S} \sqcup \{\Lambda\}$ is assigned a triple-crossing diagram $D(s)$ in \mathbb{A} as shown in Figure 15. The triple-crossing diagram $D(w_1 w_2 \cdots w_l)$ for the expression $w_1 w_2 \cdots w_l$ is obtained by concatenating the diagrams $D(w_1), D(w_2), \dots, D(w_l)$ from left to right, so that the right boundary of $D(w_i)$ is glued to the left boundary of $D(w_{i+1})$ for $i \in [l - 1]$. Further gluing the right boundary of $D(w_l)$ to the left boundary of $D(w_1)$, we obtain a triple-crossing diagram in \mathbb{T} . The corresponding plabic graph in \mathbb{T} is called an *affine plabic fence*. As explained in [FM16, Appendix D], each relation in (6.1) can be realized using isotopy and moves on the corresponding triple-crossing diagrams, except for the relations $s_i^2 = 1$ and $\bar{s}_i^2 = 1$, which are realized using Thurston's $1 - 0$ move $(R1)''$ (Figure 16). Note that the left-hand side of $(R1)''$ is the same as $(R1)'$ (but the right-hand side is not), and therefore a triple-crossing diagram D is move-reduced if and only if it is not move-equivalent to a triple-crossing diagram D' to which either $(R1)''$ or $(R2)'$ can be applied.

Remark 6.2. Postnikov's reduction $(R1)'$ leads to the relations $s_i^2 = s_i$ and $\bar{s}_i^2 = \bar{s}_i$ of the 0 -Hecke monoid.




FIGURE 17. Uncrossing a triple crossing near the left boundary of \mathbb{A} (dashed).

Remark 6.3. Rotation by 180 degrees acts on the triple-crossing diagrams by

$$D(s_i) \mapsto D(s_{n-i}), \quad D(s_{\bar{i}}) \mapsto D(s_{n-i}), \quad D(\Lambda) \mapsto D(\Lambda),$$

and induces an antiautomorphism of \tilde{S}_n sending $s_i \mapsto s_{n-i}$, $s_{\bar{i}} \mapsto s_{n-i}$, and $\Lambda \mapsto \Lambda$. We have defined $\phi(w) = (f, \bar{f})$ so that rotation of $D(w)$ by 180 degrees translates under ϕ into an automorphism of $\tilde{S}_n \times \tilde{S}_n$ sending $(f, \bar{f}) \mapsto (\bar{f}, f)$.

Lemma 6.4. *Suppose D is a move-reduced triple-crossing diagram in \mathbb{T} which is not a single strand that is a simple zero-homology loop. There is a double affine permutation $w = w(D)$ such that D is move-equivalent to $D(w)$, and such that f, \bar{f} are both c -reduced, where $\phi(w) = (f, \bar{f})$.*

Proof. Since D is not a single strand that is a simple zero-homology loop, after applying a move-equivalence using Theorem 1.16, we may assume that the number of intersections of strands in D with the sides of \mathbb{A} is minimal and that there are no strands with homology class in $\{0\} \times \mathbb{Z}$. Let $\pi := \pi_D$ denote the affine matching of D ; cf. Section 5.2. Then, we have $\pi(A) = B$ and $\pi(\bar{B}) = \bar{A}$. As in Remark 5.7, for any strand S in D , the word w_S is given by $w_S = xy^kx$ for $x \in \{l, r\}$ and $y \in \{u, d\}$.

We first show that S cannot intersect itself in \mathbb{A} . Without loss of generality, assume $w_S = ru^k r$ for $k \geq 0$. If $k = 0$, $S = S_1$ cannot intersect itself since D is move-reduced. Let $k > 0$ and let S_1, \dots, S_{k+1} denote the corresponding strands in \mathbb{D} in order as we move along S so that S_1 has word ru , S_2, \dots, S_k have words uu and S_{k+1} has word ur . Since D is move-reduced, no S_i has a self-intersection, so any self-intersection of S must be between S_i and S_j for $i \neq j$. If $j \neq k+1$, using move (T) on the $u-d$ side, we can make the intersection to be between S_{i+1} and S_{j+1} . Therefore, applying move (T) repeatedly, we may assume that the only intersections are between S_{k+1} and some of the other strands. Then, since S_1, \dots, S_k do not cross, the endpoints of S_1, \dots, S_k on the u -side of \mathbb{D} appear in order from left to right, so the endpoints of S_2, \dots, S_{k+1} on the d -side of \mathbb{D} also appear in order from left to right. If S_{k+1} intersects S_i , they form a parallel bigon contradicting that D is move-reduced.

We construct $w = w(D)$ by induction on the number of triple crossings in D . Suppose D contains no triple crossings. Then, the affine matching is of the form $\pi(A_i) = B_{i+m}$ and $\pi(B_{\bar{i}}) = A_{\bar{i}-m}$ for some $m \in \mathbb{Z}$. We assign the double affine permutation $w := \Lambda^m$ to D .

Suppose the number of triple crossings in D is nonzero. Since no strand has a self-intersection, there must be three distinct strands in \mathbb{A} at every triple crossing, so two of them must have their in-endpoints on the same side of \mathbb{A} . Any strand whose in-endpoint is between them must cross at least one of them (since no strand starts and ends on the same side of \mathbb{A}), so there exists $i \in [n]$ such that the strands S and T emanating respectively from either A_i and A_{i+1} or from $B_{\bar{i}}$ and $B_{\bar{i}+1}$ cross in \mathbb{A} . The proofs of Lemmas 5.4 and 5.8 show that we can create a triple crossing between S_1 and T_1 near the boundary of \mathbb{A} (i.e., such




FIGURE 18. (a) Uncrossing the strands S_1 (blue) and S_2 (green), while the strand S_3 (red) is unaffected. (b) The uncrossing move applied to the two strands participating in both triple crossings on the left-hand side of $(M1)'$.

that there are no other triple crossings in the triangular region bounded by S_1, T_1 and the side of \mathbb{A} containing their in-endpoints). Let D' be the triple-crossing diagram in \mathbb{A} obtained by uncrossing this triple crossing (Figure 17). We let $w := s_i w(D')$ (resp., $w := w(D') s_{\bar{i}}$) if the two strands emanate from A_i and A_{i+1} (resp., $B_{\bar{i}}$ and $B_{\bar{i}+1}$).

Clearly, $D(w)$ is move-equivalent to D , hence it is move-reduced. Let $\phi(w) = (f, \bar{f})$. We show that f and \bar{f} are c-reduced. Suppose not. By Theorem 2.3, there is a c-reduced pair (f', \bar{f}') such that $f \rightarrow f'$ and $\bar{f} \rightarrow \bar{f}'$, and we must have used either $s_i^2 = 1$ or $s_{\bar{i}}^2 = 1$ at least once. This implies that $D(w)$ is not move-reduced, a contradiction. \square

6.3. Proof of Theorem 4.14. (1) \implies (2): Suppose D is move-reduced. Since N is not a point, by Lemma 6.4, D is move-equivalent to a triple-crossing diagram $D(w)$, where w is a double affine permutation. Therefore, D and $D(w)$ have the same number of triple crossings. By Lemma 2.18, the number of triple crossings in $D(w)$ is $\ell(f) + \ell(\bar{f}) = \text{Area}(\mathcal{Z}(E_f)) + \text{Area}(\mathcal{Z}(E_{\bar{f}})) + \text{exc}(\boldsymbol{\lambda})$, where $\phi(w) = (f, \bar{f})$. By (6.2) below, we have $\ell(f) + \ell(\bar{f}) = 2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$. If D had a contractible connected component D' , then D' must have a loop strand. By Theorem 4.5, D' , and therefore D is not move-reduced.

For the converse implication, we will need the following result.

Lemma 6.5. *Let D be a triple-crossing diagram with weakly decorated Newton polygon \dot{N} . If D is not move-reduced, then there is a move-reduced triple-crossing diagram D' with weakly decorated Newton polygon \dot{N} containing strictly fewer triple crossings than D .*

Proof. Recall the reduction move $(R1)''$ shown in Figure 16. The move $(R1)''$ preserves the connectivity of the strands, and therefore does not change \dot{N} . If D is not move-reduced, then we can use moves $(M1)'$, $(R1)''$ and $(R2)'$ to get a move-reduced D' with weakly decorated Newton polygon \dot{N} . Since D has no contractible components, $(M1)'$ cannot create contractible components, and therefore we must use $(R1)''$ at least once before we can use $(R2)'$. Since we decrease the number of triple crossings when we apply $(R1)''$, D' contains strictly fewer triple crossings than D . \square

(2) \implies (1): Suppose that D has $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ triple crossings and that D has no contractible connected components. If D is not move-reduced, there is a move-reduced D' with weakly decorated Newton polygon \dot{N} with fewer than $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ triple crossings by Lemma 6.5, contradicting (1) \implies (2). \square

6.4. Proof of Proposition 4.15. Let p be a triple-crossing at which three strands S_1, S_2, S_3 meet. We call the variant of the skein relation shown in Figure 18(a) *uncrossing S_1 and S_2 at p* .




FIGURE 19. There is no way to complete the red and blue strands so that they do not cross without creating a self-intersection.

By Lemma 6.4, D is move-equivalent to a triple-crossing diagram $D(w)$, where $w \in \check{S}_n$ for some n and the associated affine permutations f, \bar{f} are c-reduced.

To show part (1), suppose there is a closed loop \tilde{S} in \tilde{D} . Then, the projection $S := \pi(\tilde{S})$ of this closed loop is a strand with $[S] = (0, 0)$. Since move-equivalence preserves homology classes of strands, S becomes a zero-homology strand in $D(w)$. Since every strand in $D(w)$ moves monotonously to the left or to the right, there are no zero-homology strands in $D(w)$, a contradiction. If \tilde{D} contains a strand \tilde{S} with a self-intersection, then uncrossing $S := \pi(\tilde{S})$ at the triple point with the self-intersection yields a triple-crossing diagram with the same weakly decorated Newton polygon but with fewer triple crossings, contradicting Theorem 4.14.

We now show part (2). By Corollary 2.20(1), part (2) is true for $D(w)$. Suppose part (2) is false for D . Since D is move-equivalent to $D(w)$, there is an intermediate triple-crossing diagram D' for which part (2) is false, but upon applying (M1)' to D' , it becomes true. Then there is a strand S in D' that intersects itself more than $\lfloor [S] \rfloor_{\mathbb{Z}} - 1$ times, but upon applying (M1)', the number of intersections becomes $\lfloor [S] \rfloor_{\mathbb{Z}} - 1$. Since (M1)' only removes crossings between the two anti-parallel strands T_1 and T_2 that cross on the left-hand side of (M1)', T_1 and T_2 should both be portions of S . Upon uncrossing T_1 and T_2 at both the triple crossings (see Figure 18(b)), the Newton polygon is unchanged, and the strand S splits into a loop and at most two other strands, so $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda})$ can decrease by at most one, but the number of triple crossings decreases by two, contradicting Theorem 4.14.

To show part (3), we will need the following lemma.

Lemma 6.6. *Suppose $S, S' \in \mathcal{S}(D)$ are two distinct parallel strands that do not intersect. Let R be a closed topological disk in \mathbb{T} whose interior contains some portion of S and S' . Let a and b (resp., c and d) denote the in- and out-endpoints of S (resp., S') around the boundary of R . Then, the cyclic order of the endpoints around the boundary of R cannot be $abcd$ or $dcba$.*

Proof. Let \tilde{D} denote the preimage of D in \mathbb{R}^2 . Let \tilde{R} be a lift of R in \mathbb{R}^2 . Let \tilde{S}, \tilde{S}' denote the lifts of S, S' to \tilde{D} that intersect \tilde{R} . Let $N \gg 1$ be a large positive integer, and consider a circle of radius N centered at \tilde{R} . Then, either S or S' has a self-intersection (Figure 19) which contradicts part (1) of Proposition 4.15. \square

Suppose there is a face F of D with portions of S, S' in its boundary. Recall from Definition 4.10 that the strands in D induce a consistent orientation around the boundary of F . We let R be a disk that contains a portion of F together with parts of S and S' , and get a contradiction with Lemma 6.6.




FIGURE 20. Proof of (6.2): the dashed line subdivides N into two polygons whose areas are $\frac{1}{2} \text{Area}(\mathcal{Z}(E_+))$ and $\frac{1}{2} \text{Area}(\mathcal{Z}(E_-))$.

Finally, we prove the statement in part (3) that parallel strands do not intersect. Similarly to part (2), by Corollary 2.20(2), part (3) is true for $D(w)$. Therefore, there is a D' move-equivalent to D such that part (3) is false for D' , but upon applying $(M1)'$ to D' , it becomes true. The two anti-parallel strands that cross on the left-hand side of $(M1)'$ should be portions of S, S' respectively. Upon uncrossing S and S' at both triple crossings, the union of S and S' becomes the union of a loop and a strand T with homology class $[T] = [S] + [S']$. Therefore, N is unchanged and $2 \text{Area}(N) + \text{exc}(\lambda)$ decreases by one, but the number of triple crossings decreases by two, again contradicting Theorem 4.14. \square

6.5. Proof of Proposition 1.11. Suppose $\ddot{N} = (N, \alpha)$ is a strongly decorated Newton polygon and $\mu \in \mathbb{Z}/d(\alpha)\mathbb{Z}$. Recall from Section 3.2 that for $e = (a, b) \in \mathbb{Z}^2$, we denote $n(e) := a$ and $k(e) := b$, and $\nu(e) = k(e)/n(e)$. Using an $\text{SL}_2(\mathbb{Z})$ transformation, we can assume that $n(e) \neq 0$ for all $e \in E(N)$. We assign to \ddot{N} the pair (\ddot{E}_+, \ddot{E}_-) of strongly decorated vector configurations, consisting of edges of N oriented to the right and left, respectively, as follows. We define:

- (1) $E_+ := \{e \in E(N), n(e) > 0\}$ and $\alpha_+ = (\alpha^e)_{e \in \ddot{E}_+}$; and
- (2) $E_- := \{-e \mid e \in E(N), n(e) < 0\}$ and $\alpha_- = (\text{rev}(\alpha^e))_{-e \in \ddot{E}_-}$, where for a cyclic composition $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$, $\text{rev}(\alpha) := (\alpha_m, \alpha_{m-1}, \dots, \alpha_1)$ is the cyclic composition with the cyclic order reversed.

Similarly to Remark 6.3, we have rotated the vectors in E_- by 180 degrees. We have the following basic relation between the area of N and the areas of the zonotopes $\mathcal{Z}(E_+), \mathcal{Z}(E_-)$:

$$(6.2) \quad 2 \text{Area}(N) = \text{Area}(\mathcal{Z}(E_+)) + \text{Area}(\mathcal{Z}(E_-)).$$

To see this, observe that the lower boundary of $\mathcal{Z}(E_+)$ coincides with the lower boundary of N (given by the vectors in E_+), and the upper boundary of $\mathcal{Z}(E_+)$ is obtained by rotating its lower boundary by 180 degrees. A similar statement holds for $\mathcal{Z}(E_-)$, from which the result follows; see Figure 20.

Let f and \bar{f} be a pair of c-reduced affine permutations with $\ddot{E}_f = \ddot{E}_+$ and $\ddot{E}_{\bar{f}} = \ddot{E}_-$ constructed as in Section 3.2. Observe that $\sum_{e \in E_+} k(e) = \sum_{e \in E_-} k(e)$, and thus by Remark 6.1, there exists $w \in \ddot{S}_n$ satisfying $\phi(w) = (f, \bar{f})$. By (6.2), the triple-crossing diagram $D := D(w)$ has the correct number of triple crossings, so it is move-reduced by Theorem 4.14. Let $\Gamma := \Gamma(D)$ be the associated bipartite graph (cf. Appendix A). By construction, $\ddot{N}(\Gamma) = \ddot{N}$.

Finally, we show that Γ has a perfect matching. Let $w = s_{i_1} s_{i_2} \cdots s_{i_k} \Lambda^k s_{\bar{j}_1} s_{\bar{j}_2} \cdots s_{\bar{j}_m}$ be a reduced expression. Omitting all generators s_{i_k} and $s_{\bar{i}_k}$ such that the corresponding vertical




FIGURE 21. The triple-crossing diagram $D(w)$ with strongly decorated Newton polygon $\ddot{N} = (N, \alpha)$ from Example 6.7.




FIGURE 22. Two plabic graphs $\Gamma(w), \Gamma(w')$ from Example 6.8 having the same strongly decorated Newton polygons but different modular invariants. According to Theorem 1.13, these graphs are not move-equivalent.

edge in Γ is traversed by the same strand in the opposite directions (i.e., yields a self-intersection in $D(w)$), we get a triple-crossing diagram D' with strongly decorated Newton polygon (N, α') satisfying $(\alpha')^e = (1, 1, \dots, 1)$ for all $e \in E(N)$. By Theorem 4.14 and part (2) of Proposition 4.15, D' is move-reduced, so it is minimal in the sense of [GK13]. The corresponding bipartite graph $\Gamma' := \Gamma(D')$ has a perfect matching by [GK13, Lemma 3.11], and since Γ' is obtained from Γ by deleting a subset of edges, so does Γ . \square

Example 6.7. Let $\ddot{N} = (N, \alpha)$ be the strongly decorated Newton polygon with edges $e_1 = (2, 0)$, $e_2 = (2, 2)$ and $e_3 = (-4, -2)$, and $\alpha^{e_1} = \alpha^{e_2} = \alpha^{e_3} = (2)$ shown in Figure 22(c). The strongly decorated vector configuration \ddot{E}_+ and its ϵ -straight arrow diagram $D(\ddot{E}_+)$ are shown in Figure 10(a–b). From $D(\ddot{E}_+)$, we find the reduced expression $f = s_1 s_3 s_4 s_3 s_1 s_4 \Lambda^2$. Similarly, we have $\bar{f} = s_1 \Lambda^2$, so that $w = s_1 s_3 s_4 s_3 s_1 s_4 \Lambda^2 s_{\bar{3}}$. The corresponding triple-crossing diagram $D(w)$ is shown in Figure 21.

Example 6.8. Let $w' = s_1 s_3 s_4 s_3 s_1 s_4 \Lambda^2 s_{\bar{2}}$ be obtained from w in Example 6.7 by replacing $s_{\bar{3}}$ with $s_{\bar{2}}$. The associated plabic⁴ graphs $\Gamma(w), \Gamma(w')$ shown in Figure 22 have the same strongly decorated Newton polygons but different modular invariants in $\mathbb{Z}/d(\alpha)\mathbb{Z}$, where $d(\alpha) = 2$.

⁴Strictly speaking, the graphs shown in Figure 22 are not plabic in the language of Section 4.2 since they have edges with both endpoints black. To convert them into plabic graphs, one has to add a degree two white vertex in the middle of each such edge.




FIGURE 23. Moves and reductions on graphs that correspond to moves and reductions on triple-crossing diagrams. Under $\Gamma \mapsto D(\Gamma)$, (M1b) and (M2b-iii) become (M1)', and (R1-iii) and (R2w-iii) become (R1)'.

6.6. Proof of Theorem 1.13. We show first that two move-reduced graphs Γ, Γ' are move-equivalent if and only if

$$(\ddot{N}(\Gamma), \mu(\Gamma)) = (\ddot{N}(\Gamma'), \mu(\Gamma')).$$

The \implies direction is clear, since both \ddot{N} and μ are invariant under move-equivalence; see Sections 1.4 and 1.5.

For the \impliedby direction, using Lemma 6.4, we assume that the triple-crossing diagram D (resp., D') associated to Γ (resp., Γ') is of the form $D(w)$ (resp., $D(w')$) for some double affine permutations w, w' . Let (f, \bar{f}) (resp., (f', \bar{f}')) be the pair of affine permutations associated to w (resp., w'). Let $\sigma(w) = \Lambda w \Lambda^{-1}$ be the rotation operator, and let $\sigma(\Gamma)$ be the bipartite graph associated to the triple-crossing diagram $D(\sigma(w))$. Note that $\mu(\sigma(\Gamma)) = \mu(\Gamma)$, but $\mu(\sigma(f)) = \mu(f) + 1$ and $\mu(\sigma(\bar{f})) = \mu(\bar{f}) - 1$. Therefore, replacing Γ with $\sigma^{\mu(f') - \mu(f)}(\Gamma)$, we can assume that $\mu(f) = \mu(f')$.

We will show that there is an $r \in \mathbb{Z}$ such that $\sigma^r(f) \stackrel{\circ}{\sim} f'$ and $\sigma^r(\bar{f}) \stackrel{\circ}{\sim} \bar{f}'$. Since $\ddot{N}(\Gamma) = \ddot{N}(\Gamma')$ implies that $\ddot{E}_f = \ddot{E}_{f'}$ and $\ddot{E}_{\bar{f}} = \ddot{E}_{\bar{f}'}$, by Theorem 3.2, it suffices to show that there is an $r \in \mathbb{Z}$ such that $\mu(\sigma^r(f)) = \mu(f')$ and $\mu(\sigma^r(\bar{f})) = \mu(\bar{f}')$, or equivalently, such that $r \equiv 0 \pmod{d(\alpha_f)}$ and $r \equiv \mu(\bar{f}) - \mu(\bar{f}') \pmod{d(\alpha_{\bar{f}})}$. Note that $d(\alpha) = \gcd(d(\alpha_f), d(\alpha_{\bar{f}}))$ and $\mu(\Gamma) \equiv \mu(f) + \mu(\bar{f}) \pmod{d(\alpha)}$. Since $\mu(\Gamma) = \mu(\Gamma')$ and $\mu(f) = \mu(f')$, we have $\mu(\bar{f}) - \mu(\bar{f}') \equiv 0 \pmod{d(\alpha)}$. The existence of such an r follows from Lemma 6.9.

Lemma 6.9. *Let d_1, d_2 be positive integers, and let $d = \gcd(d_1, d_2)$. Then, there exists $r \in \mathbb{Z}$ such that $r \equiv 0 \pmod{d_1}$ and $r \equiv d \pmod{d_2}$.*

Proof. Let $x, y \in \mathbb{Z}$ be such that $xd_1 + yd_2 = d$. Take $r := xd_1$. □

Finally, we show that there are $d(\alpha)$ nonempty move-equivalence classes, i.e., that for any $\mu \in \mathbb{Z}/d(\alpha)\mathbb{Z}$, there is a move-reduced graph Γ with $(\ddot{N}(\Gamma), \mu(\Gamma)) = (\ddot{N}, \mu)$. In the proof of Proposition 1.11, we constructed Γ with $\ddot{N}(\Gamma) = \ddot{N}$ as the graph associated to a triple-crossing diagram $D(w)$. Let (f, \bar{f}) be the pair of affine permutations associated to w . Let w' be the double affine permutation associated to $(\sigma^r(f), \bar{f})$ and let Γ' be the bipartite graph associated to the triple-crossing diagram $D(w')$. Then

$$(\ddot{N}(\Gamma'), \mu(\Gamma')) = (\ddot{N}(\Gamma), \mu(\Gamma) + r),$$

so by varying r , we can get any value of the modular invariant.

APPENDIX A. FROM BIPARTITE GRAPHS TO TRIPLE-CROSSING DIAGRAMS

The goal of this section is to give a relation (Lemmas A.3 and A.5) between bipartite graphs and triple-crossing diagrams; cf. Remark 4.13. We will use these results to deduce




FIGURE 24. Moves and reductions on graphs that are not compatible with triple-crossing diagrams. (R1-ii), (R2b) and (R3) contain black leaves. Both sides of (R2w-ii) correspond to the same triple-crossing diagram.

Theorem 1.5 from Theorem 4.14. Unless otherwise stated, all graphs in this section are bipartite.

Let (M1b) (resp., (M1w)) denote the version of (M1) with black (resp., white) trivalent interior vertices, and let (M2b) (resp., (M2w)) denote the version of (M2) contracting/uncontracting black (resp., white) vertices. Note that (M1w) can be realized using (M1b) and (M2b) — we uncontract all the black vertices using (M2b), apply (M1b) and then contract using (M2b).

We say that Γ is (M2w)-*reduced* if contraction using (M2w) cannot be applied to Γ .

Lemma A.1. *Two (M2w)-reduced graphs Γ and Γ' are move-equivalent if and only if they are related by (M1b) and (M2b). An (M2w)-reduced graph is move-reduced if and only if it is not move-equivalent to an (M2w)-reduced graph to which one of (R1)–(R3) can be applied.*

Proof. By inspection, we see that no move or reduction, except possibly (M2b), involves a degree-two black vertex that can be contracted using (M2w). (Moves (R1) and (R2) might have degree-two black vertices but they cannot be contracted using (M2w).) Applying (M2b) with a degree-two black vertex is the same as applying (M2w). \square

For the rest of this section, we assume that our graphs are (M2w)-reduced. A *white-based loop* in Γ is a parallel edge in which the black vertex has degree two (see the left-hand side of (R1-ii) in Figure 24). In (R1), if the black vertex has degree greater than three, then we can uncontract using (M2b) to make it degree three, and denote this case of (R1) by (R1-iii). Otherwise, we have a white-based loop and we denote this case by (R1-ii) (Figure 24).

Let (R2b) (resp., (R2w)) denote black (resp., white) leaf removal. If the white leaf in (R2w) is incident to a black vertex of degree greater than three, we can uncontract the black vertex using (M2b) to get a white leaf incident to a black vertex of degree three and call this (R2w-iii). If we have a black vertex of degree two, we call it (R2w-ii) (Figure 24).

Let Γ be a graph in \mathbb{T} . Use (R2b), (R2w-ii) and (R3) to remove all black leaves and white leaves incident to degree-two black vertices. Use (M2b) to uncontract black vertices with degree greater than three until every black vertex has degree either zero, two or three. We call such a graph *partially reduced*. Remove isolated black vertices and omit all degree-two black vertices, converting the two incident edges into a single edge to get a plabic graph. Use the procedure in Figure 12 to obtain a triple-crossing diagram $D(\Gamma)$. Under this procedure, zig-zag paths in Γ become strands of $D(\Gamma)$. The choices in applying (M2b) lead to move-equivalent triple-crossing diagrams. Let (M2b-iii) denote the resplit move (Figure 23), which consists of two applications of (M2b).

Lemma A.2. *Two partially reduced graphs Γ and Γ' are move-equivalent if and only if they are related by (M1b) and (M2b-iii). A partially reduced graph Γ is move-reduced if and only*




FIGURE 25. Reductions on triple-crossing diagrams that do not correspond to reductions on graphs under $D \mapsto \Gamma(D)$. Here, (a) and (b) are the two versions of $(R2)'$, and (c) and (d) the two versions of $(R1)''$. In (d), we further contract the degree-two black vertex to get an $(M2w)$ -reduced graph.

if it is not move-equivalent to a partially reduced graph Γ' to which either $(R1-ii)$, $(R1-iii)$ or $(R2w-iii)$ can be applied.

Proof. Since Γ is partially reduced, no black vertices involved in $(M2b)$ have degree two. Any applications of $(M2b)$ involving black vertices of degree greater than three can be decomposed into multiple applications of $(M2b-iii)$.

Clearly, $(M1b)$ preserves partial reducedness. Contracting/uncontracting using $(M2b)$ does not change whether any of the moves $(R2b)$, $(R2w-ii)$ or $(R3)$ can be applied. Therefore, if Γ is related to Γ' using $(M2b)$, then we can further apply $(M2b)$ to make Γ' partially reduced. The reductions $(R1-ii)$, $(R1-iii)$ and $(R2w-iii)$ are the only ones that can be applied to a partially reduced graph. \square

Conversely, we obtain a graph $\Gamma(D)$ from a triple-crossing diagram D as follows. Use the procedure in Figure 12 to obtain a plabic graph, contract any white-white edges incident to distinct white vertices and place a black vertex at the midpoint of each white-white edge to obtain a bipartite graph $\Gamma(D)$. The different choices in applying Figure 12(b) all lead to the same plabic graph when we contract any white-white edges incident to distinct white vertices.

Note that $\Gamma(D)$ is partially reduced and has no isolated black vertices. Therefore, $\Gamma \mapsto D(\Gamma)$ and $D \mapsto \Gamma(D)$ are inverse functions between partially reduced graphs without isolated black vertices and triple-crossing diagrams.

Lemma A.3. *The functions $\Gamma \mapsto D(\Gamma)$ and $D \mapsto \Gamma(D)$ between partially reduced graphs without isolated black vertices and triple-crossing diagrams respect move-equivalence.*

Proof. Under the correspondence, $(M1)'$ becomes either $(M1b)$ or $(M2b-iii)$, so the result follows from Lemma A.2. \square

Remark A.4. The functions $\Gamma \mapsto D(\Gamma)$ and $D \mapsto \Gamma(D)$ do not commute with reductions; see Figures 24 and 25.

Lemma A.5. *The function $\Gamma \mapsto D(\Gamma)$ is a bijection between move-equivalence classes of move-reduced graphs without isolated vertices and move-equivalence classes of move-reduced triple-crossing diagrams, with inverse $D \mapsto \Gamma(D)$.*

Proof. If Γ is move-reduced and has no isolated black vertices, then it can be transformed using $(M2b)$ into a partially reduced graph. Therefore, move-equivalence classes of move-reduced graphs without isolated black vertices are in bijection with move-equivalence classes of move-reduced partially reduced graphs without isolated black vertices.

By Lemma A.3, $\Gamma \mapsto D(\Gamma)$ is a bijection between move-equivalence classes of move-reduced graphs and some subset T of move-equivalence classes of triple-crossing diagrams that we need to identify.

Let Γ be move-reduced and partially reduced without isolated vertices. Γ has no white-based loops; otherwise (R1–ii) can be applied. Since isolated loops in $D := D(\Gamma)$ correspond to isolated white vertices or white-based loops in Γ , D contains no isolated loops. Since (M1)' cannot create isolated loops, D is move-reduced if and only if it is not move-equivalent to a D' to which (R1)' can be applied. Under the functions $\Gamma \mapsto D(\Gamma)$ and $D \mapsto \Gamma(D)$, (R1–iii) and (R2w–iii) become (R1)' (Figure 23), so D is move-reduced by Lemma A.2. Therefore, T is contained in the set of move-equivalence classes of move-reduced triple-crossing diagrams.

Let D be a move-reduced triple-crossing diagram, and let $\Gamma := \Gamma(D)$. Since

- (1) (R1–iii) and (R2w–iii) become (R1)' (Figure 23);
- (2) An isolated white vertex becomes the left-hand side of (R2)' (Figure 25(a)); and
- (3) The left-hand side of (R1–ii) becomes the left-hand side of (R2)' (Figure 25(b)),

Γ has no isolated white vertices, and Γ is move-reduced by Lemma A.2. Therefore, T contains the set of move-equivalence classes of move-reduced triple-crossing diagrams. \square

A.1. Proof of Theorem 1.5. We need the following lemma.

Lemma A.6. *Let Γ be a move-reduced graph without isolated vertices. Assume $N(\Gamma)$ is not a single point. The number of contractible faces of Γ is equal to the number of degree-three black vertices of Γ .*

Proof. If Γ is the affine plabic fence associated to Λ^k , then both numbers are zero. Each s_i and $s_{\bar{i}}$ increases both numbers by one. Therefore, the result holds for affine plabic fences. If Γ is move-reduced, then Γ is move-equivalent to the bipartite graph associated with an affine plabic fence by Lemma 6.4 and Lemma A.5, and move-equivalence does not change the number of contractible faces. \square

If Γ and Γ' are related by (M2w) or (M2b), then each of the conditions (1)–(2) holds for Γ if and only if it holds for Γ' . Each of (1)–(2) imply that Γ is leafless. Therefore, we can assume that Γ is partially reduced. Moreover, since Γ has a perfect matching, Γ has no isolated vertices.

(1) \implies (2): Since Γ has a perfect matching, N is not a single point by Theorem 1.16. The implication follows from Lemma A.5, Theorem 4.14 and Lemma A.6.

(2) \implies (1): Suppose N is a single point. Then, $2 \text{Area}(N) + \text{exc}(\boldsymbol{\lambda}) = 0$ so Γ has no contractible faces. Since (M1)–(M2) cannot create leaves or contractible faces, none of the reductions (R1)–(R3) can be applied to any graph move-equivalent to Γ , so Γ is move-reduced.

Assume N is not a single point. By Lemma A.2, Γ is not move-reduced if and only if it is move-equivalent to a partially reduced Γ'' to which either (R1–ii), (R1–iii) or (R2w–iii) can be applied. Since Γ has no leaves and (M1)–(M2) cannot create leaves, either (R1–ii) or (R1–iii) can be applied to Γ' . Let $D := D(\Gamma)$ be the associated triple-crossing diagram. Then, either (R1)'' or (R2)' can be applied to D' . We decrease the number of contractible faces in Γ when we apply either reduction (see Figure 25 (b) and (c)). Transform D' into a move-reduced D'' by further using (M1)', (R1)'' and (R2)'. The graph $\Gamma'' = \Gamma(D'')$ has strictly fewer contractible faces than Γ , no isolated vertices, and satisfies $\dot{N}(\Gamma'') = \dot{N}(\Gamma)$. Since N is

not a single point, Γ'' has $2 \text{Area}(N) + \text{exc}(\lambda)$ contractible faces by Lemma A.5, Theorem 4.14 and Lemma A.6, a contradiction.

REFERENCES

- [FM16] Vladimir V. Fock and Andrey Marshakov. Loop groups, clusters, dimers and integrable systems. In *Geometry and quantization of moduli spaces*, Adv. Courses Math. CRM Barcelona, pages 1–66. Birkhäuser/Springer, Cham, 2016.
- [Gal21] Pavel Galashin. Critical varieties in the Grassmannian. [arXiv:2102.13339v2](https://arxiv.org/abs/2102.13339v2), 2021.
- [GK13] Alexander B. Goncharov and Richard Kenyon. Dimers and cluster integrable systems. *Ann. Sci. Éc. Norm. Supér. (4)*, 46(5):747–813, 2013.
- [GKP00] Meinolf Geck, Sungsoo Kim, and Götz Pfeiffer. Minimal length elements in twisted conjugacy classes of finite Coxeter groups. *J. Algebra*, 229(2):570–600, 2000.
- [GL21] Pavel Galashin and Thomas Lam. Positroid Catalan numbers. [arXiv:2104.05701v1](https://arxiv.org/abs/2104.05701v1), 2021.
- [GP93] Meinolf Geck and Götz Pfeiffer. On the irreducible characters of Hecke algebras. *Adv. Math.*, 102(1):79–94, 1993.
- [GPW22] Pavel Galashin, Alexander Postnikov, and Lauren Williams. Higher secondary polytopes and regular plabic graphs. *Adv. Math.*, 407:Paper No. 108549, 52, 2022.
- [GSV11] Michael Gekhtman, Michael Shapiro, and Alek Vainshtein. Generalized Bäcklund-Darboux transformations for Coxeter-Toda flows from a cluster algebra perspective. *Acta Math.*, 206(2):245–310, 2011.
- [GSZ20] Pavlo Gavrylenko, Mykola Semenyakin, and Yegor Zenkevich. Solution of tetrahedron equation and cluster algebras. [arXiv:2010.15871v2](https://arxiv.org/abs/2010.15871v2), 2020.
- [He07] Xuhua He. Minimal length elements in some double cosets of Coxeter groups. *Adv. Math.*, 215(2):469–503, 2007.
- [He10] Xuhua He. Minimal length elements in conjugacy classes of extended affine Weyl group. [arXiv:1004.4040v1](https://arxiv.org/abs/1004.4040v1), 2010.
- [HN14] Xuhua He and Sian Nie. Minimal length elements of extended affine Weyl groups. *Compos. Math.*, 150(11):1903–1927, 2014.
- [KO06] Richard Kenyon and Andrei Okounkov. Planar dimers and Harnack curves. *Duke Math. J.*, 131(3):499–524, 2006.
- [KOS06] Richard Kenyon, Andrei Okounkov, and Scott Sheffield. Dimers and amoebae. *Ann. of Math. (2)*, 163(3):1019–1056, 2006.
- [LP13] Thomas Lam and Pavlo Pylyavskyy. Total positivity for loop groups II: Chevalley generators. *Transform. Groups*, 18(1):179–231, 2013.
- [Mar14] Timothée Marquis. Conjugacy classes and straight elements in Coxeter groups. *J. Algebra*, 407:68–80, 2014.
- [Mar18] Timothée Marquis. Cyclically reduced elements in Coxeter groups. [arXiv:1812.02543v4](https://arxiv.org/abs/1812.02543v4), 2018.
- [Mar20] Timothée Marquis. Structure of conjugacy classes in Coxeter groups. [arXiv:2012.11015v2](https://arxiv.org/abs/2012.11015v2), 2020.
- [McM84] Peter McMullen. Volumes of projections of unit cubes. *Bull. London Math. Soc.*, 16(3):278–280, 1984.
- [Pos06] Alexander Postnikov. Total positivity, Grassmannians, and networks. Preprint, <http://math.mit.edu/~apost/papers/tpgrass.pdf>, 2006.
- [Thu17] Dylan P. Thurston. From dominoes to hexagons. In *Proceedings of the 2014 Maui and 2015 Qinhuaodao conferences in honour of Vaughan F. R. Jones’ 60th birthday*, volume 46 of *Proc. Centre Math. Appl. Austral. Nat. Univ.*, pages 399–414. Austral. Nat. Univ., Canberra, 2017.
- [Wil16] Harold Williams. Toda systems, cluster characters, and spectral networks. *Comm. Math. Phys.*, 348(1):145–184, 2016.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, CA 90095, USA
Email address: galashin@math.ucla.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48103, USA
Email address: georgete@umich.edu