Mathematische Annalen
https://doi.org/10.1007/s00208-022-02434-4

Mathematische Annalen

Monodromy in Kazhdan–Lusztig cells in affine type A

Michael Chmutov[1] · Joel Brewster Lewis[2] · Pavlo Pylyavskyy[1]

Received: 25 March 2019 / Accepted: 27 June 2022

Abstract

We use the affine Robinson–Schensted correspondence to describe the structure of bidirected edges in the Kazhdan–Lusztig cells in affine type A. Equivalently, we give a comprehensive description of the Knuth equivalence classes of affine permutations.

Contents

M.C. was partially supported by NSF Grant DMS-1503119. J.B.L. was partially supported by NSF Grant DMS-1401792. P.P. was partially supported by NSF Grants DMS-1148634, DMS-1351590, and a Sloan Fellowship.

✉ Joel Brewster Lewis
jblewis@gwu.edu

Michael Chmutov
mchmutov@gmail.com

Pavlo Pylyavskyy
ppylyavs@umn.edu

1 Department of Mathematics, University of Minnesota, Minneapolis, MN 55455, USA

2 George Washington University, Washington, DC 20052, USA

Published online: 05 August 2022

1 Introduction

1.1 Cells in Kazhdan–Lusztig theory

In a groundbreaking paper [14], Kazhdan and Lusztig laid a basis for a new approach to representation theory of Hecke algebras. Since then, this approach has been significantly developed, and is called *Kazhdan–Lusztig theory*. (For a nice introduction to Kazhdan–Lusztig theory, see [13, Ch. 7].) Of particular importance in this theory are the objects called *cells*. Briefly, their definition is as follows. Each Hecke algebra is associated with a Coxeter group W. Kazhdan and Lusztig define a pre-order $\leq_L$ on elements of W. Some pairs v, w of elements of W satisfy both $v \leq_L w$ and $w \leq_L v$, in which case we say that they are left-equivalent, denoted $v \sim_L w$. Similarly one can define right equivalence $\sim_R$. The respective equivalence classes are called the *left cells* and the *right cells*.

Another way to describe cells is via the Kazhdan–Lusztig W-graph; it is a certain directed graph whose vertices are the elements of W. The graph has the property that $v \leq_L w$ precisely when there is a directed path from v to w. Thus the cells are the strongly connected components of the W-graph. Some edges of the W-graph are bidirected, i.e., between a pair of vertices v and w there is an edge $v \to w$ and an edge $w \to v$. In this case, of course, v and w belong to the same Kazhdan–Lusztig cell.

1.2 Type A

In (finite) type A, when W is the symmetric group, the Kazhdan–Lusztig cell structure corresponds to something very familiar to combinatorialists, the *Robinson–Schensted correspondence*. This is a bijective correspondence between elements of the symmetric

group and pairs (P, Q) of standard Young tableaux of the same shape. It is well known [1, 6, 11, 14] that

- two permutations lie in the same left cell if and only if they have the same *recording tableau* Q, and
- two permutations lie in the same right cell if and only if they have the same *insertion tableau* P.

The bidirected edges of the Kazhdan–Lusztig graph in this case are called *Knuth moves*, and one can go between any two permutations with the same insertion tableau (and hence between any two choices of the recording tableau Q) via a series of Knuth moves.

1.3 Affine type A

In affine type A, when W is an affine symmetric group, Chmutov, Pylyavskyy, and Yudovina [8] described, via a combinatorial algorithm called the Affine Matrix Ball Construction (AMBC), a bijection $W \to \Omega_{\text{dom}}$, where Ω_{dom} is the set of triples (P, Q, ρ) such that P and Q are tabloids of the same shape and ρ is an integer vector (called a *dominant weight*) satisfying certain inequalities that depend on P and Q. Relying on the work of Shi on Kazhdan–Lusztig cells in affine type A, they show that this bijection affords a description of cells analogous to the non-affine case: fixing the tabloid Q gives all affine permutations in a left cell while fixing the tabloid P gives all affine permutations in a right cell. The bidirected edges in this case (what Shi called *star operations* [19]) are natural analogues of Knuth moves—see Sect. 3.3.1 for the definition.

1.4 Summary of results

The main accomplishment of this paper is to precisely describe the Knuth equivalence classes of affine permutations, i.e., the equivalence classes of the relation generated by Knuth moves. In the language of [22], we describe the *Kazhdan–Lusztig molecules*. Unlike the case of finite type A, most Kazhdan–Lusztig cells are composed of many molecules. This multiplicity comes in two varieties. First, while Knuth moves preserve the P tabloid, not all Q tabloids can be reached from a given one using Knuth moves; the description of which ones can be reached, given in Sect. 8, is in terms of a variant of the *charge statistic*. Second, even among affine permutations having the same P and Q tabloids, one cannot reach every dominant weight ρ from every other using Knuth moves; the different vectors ρ that can be reached are the subject of Sect. 7, and they depend on the shape of the tabloids.

Along the way, we improve our understanding of some combinatorial aspects of AMBC analogous to the combinatorics of the Robinson–Schensted correspondence. We show how to read off the left and right descent sets of an affine permutation from its image under AMBC (Proposition 3.6), as well as how to read off its sign (Theorem 4.4). We describe precisely how taking the inverse of or doing a Knuth move on an affine permutation affects its image under AMBC (Proposition 3.1 and Theorem 3.11).

In Sect. 6, we briefly discuss how the symmetry of the Dynkin diagram appears in our setting. Finally, we give a better description, in terms of charge, for the inequalities satisfied by an integer vector for it to be dominant (Theorem 5.10). The appearance of charge suggests a connection to crystal graphs, which we describe in Sect. 9.

Section 2 contains the background information, including a summary of the main results of [8], essential to understanding the rest of the paper. Sections 3 and 4 improve our understanding of the various combinatorial aspects of AMBC. Sections 5, 7 and 8 describe the Knuth equivalence classes of permutations in terms of AMBC; these sections rely heavily on the material in Sect. 3 but may be read independently of Sect. 4.

There is a program available to compute AMBC [7]; the reader may want to use it to explore additional examples.

2 Background

2.1 Notational preliminaries

For the duration of this paper, n will be a fixed positive integer. Let $[n] := \{1, \ldots, n\}$. For each $i \in \mathbb{Z}$, denote by $\bar{i}$ the residue class $i + n\mathbb{Z}$, and let $[\overline{n}] := \{\overline{1}, \ldots, \overline{n}\}$.

The *symmetric group* S_n is the Weyl group of type A_{n-1}. We may variously think of its elements as bijections $[n] \to [n]$, as words of length n containing each element of $[n]$ exactly once, or as $n \times n$ permutation matrices. The *extended affine symmetric group* $\widetilde{S}_n$ is the extended affine Weyl group of type $\widetilde{\mathrm{A}}_{n-1}$; it consists of bijections $w : \mathbb{Z} \to \mathbb{Z}$ such that

$$w(i + n) = w(i) + n \qquad \text{for all } i.$$

The elements of $\widetilde{S}_n$ are called *extended affine permutations*. We typically abbreviate this term to *permutations*, and we distinguish the elements of S_n by the name *finite permutations*. Denote by $\widetilde{S}_n^0$ the *affine symmetric group*, i.e., the affine Weyl group of type $\widetilde{\mathrm{A}}_{n-1}$; it consists of permutations $w \in \widetilde{S}_n$ such that

$$\sum_{i=1}^{n} (w(i) - i) = 0.$$

Note that S_n naturally embeds into $\widetilde{S}_n^0 \subsetneq \widetilde{S}_n$: a finite permutation w can be sent to the unique affine permutation that takes the same values as w on $[n]$.

A *partial (extended affine) permutation* is a pair (U, w) where $U \subseteq \mathbb{Z}$ has the property that $(x \in U) \Leftrightarrow (x + n \in U)$ and $w : U \to \mathbb{Z}$ an injection such that $w(i+n) = w(i)+n$. We suppress the explicit mention of the subset U in the notation and just refer to the partial permutation w. Any permutation may be viewed as a partial permutation with $U = \mathbb{Z}$.

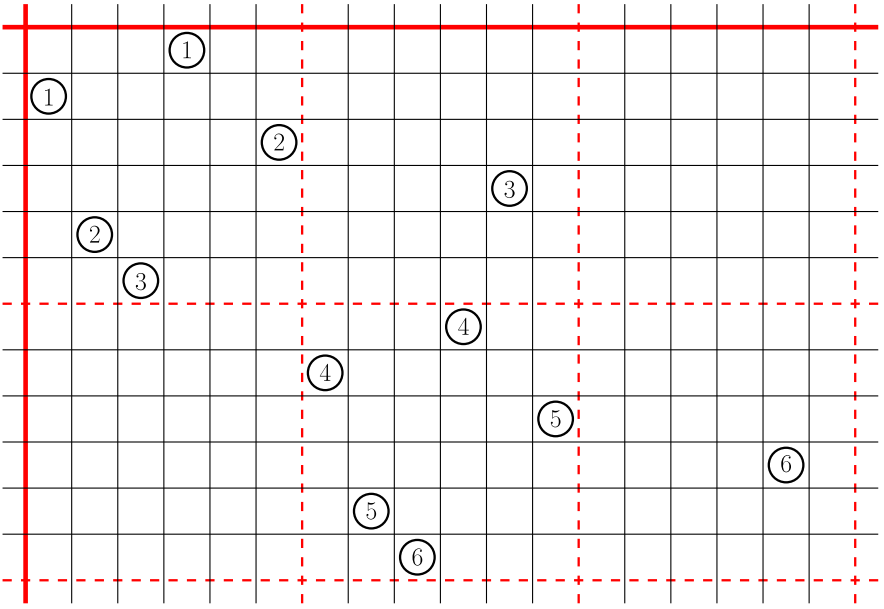

Fig. 1 A proper numbering for the (extended, affine) permutation [4, 1, 6, 11, 2, 3]. The numbering is semi-periodic with period 3

A permutation is determined by its values on $1, \ldots, n$. The *window notation* for a permutation w is $[w(1), \ldots, w(n)]$. A partial permutation w is also determined by its values on $1, \ldots, n$, except it may not be defined on some of them.

We often think of permutations in terms of pictures such as the one in Fig. 1, extending the notion of a permutation matrix to the affine case. More precisely, on the plane we draw an infinite matrix; the rows are labeled by $\mathbb{Z}$, increasing downward, and the columns are labeled by $\mathbb{Z}$, increasing to the right (usual matrix coordinates). The positions in this matrix are called *cells*; there will never be confusion with cells of the Kazhdan–Lusztig variety. To distinguish the 0-th row, figures have a solid red line between the 0-th and 1-st rows, and similarly for columns. We also put dashed red lines every n rows and columns. If $w(i) = j$ then we place a ball in the i-th row and j-th column. For example, the cell $(1, 4)$ in Fig. 1 contains a ball. The balls of a partial permutation will also be referred to by their matrix coordinates. Thus, formally, both balls and cells are just ordered pairs of integers, and we use the word "ball" to indicate that the relevant partial permutation takes a certain value on a certain input. For a partial permutation w, we denote by $\mathcal{B}_w$ the collection of balls of w (a subset of $\mathbb{Z} \times \mathbb{Z}$).

For an integer k and a ball $b = (i, j)$, the ball $b' = b + k(n, n) = (i + kn, j + kn)$ is the $k(n, n)$-*translate* of b. Two balls b and b' are *translates* if for some k one is a $k(n, n)$-translate of the other. The set of all translates of a ball or set of balls is a *translation class*.

We often assign numbers to balls of permutations, as well as to other cells of $\mathbb{Z} \times \mathbb{Z}$. For a partial permutation w, a *numbering* of w is a function $d : \mathcal{B}_w \to \mathbb{Z}$. A numbering

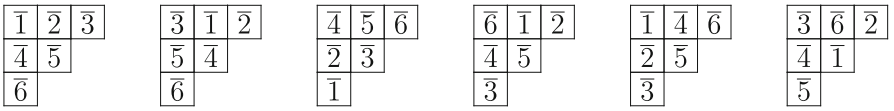

Fig. 2 Several tabloids of shape $\langle 3, 2, 1\rangle$. The first two tabloids are equal, since they differ only by permuting elements within rows

d of w is *semi-periodic* with *period* m if we have $d(b + (n, n)) = d(b) + m$ for every $b \in \mathcal{B}_w$. When referring to a numbering in pictures, we write the number $d(b)$ inside the ball b as in Fig. 1, where we show a semi-periodic numbering of period 3.

We frequently use compass directions (north, east, etc.) to describe relative positions of balls or cells, with north being toward the top of the page (smaller row numbers) and east being toward the right of the page (larger column numbers). Adding the modifier "directly" constrains one of the two coordinates: a cell (i, j) is *directly south* of (i', j') if $i \geqslant i'$ and $j = j'$. By a composite direction (e.g., northeast) we mean north and east. The relations are weak by default: a cell (i, j) is *southwest of* (i', j') if $i \geqslant i'$ and $j \leqslant j'$. Directions define partial orders on $\mathbb{Z} \times \mathbb{Z}$: we say $(i, j) \leqslant_{SW} (i', j')$ if (i, j) is southwest of (i', j').

A *partition* λ is a finite, weakly decreasing sequence of positive integers. We typically treat a partition $\lambda = \langle \lambda_1, \lambda_2, \ldots\rangle$ as equivalent to its *Young diagram*, a left-justified collection of rows of boxes with top row having λ_1 boxes, the row below it having λ_2 boxes, and so on. The number of rows of λ is denoted $\ell(\lambda)$.

Given a partition λ having n boxes, a *tabloid* of *shape* λ is an equivalence class of fillings of the Young diagram of λ with $[\overline{n}]$, where two fillings are considered equivalent when one is obtained from the other by permuting elements within rows.[1] Some examples are shown of Fig. 2; we draw representatives of the equivalence classes and keep in mind that entries within rows can be permuted. Throughout this paper, all tabloids will be filled with distinct residue classes.

Several special tabloids will be important in this paper. The first is the *reverse row superstandard tabloid* of shape $\lambda = \langle \lambda_1, \ldots, \lambda_k\rangle$ with *start* at i: this is the tabloid whose last row has entries

$$\overline{i},\ \overline{i+1},\ \ldots,\ \overline{i+\lambda_k-1},$$

whose next-to-last row has entries

$$\overline{i+\lambda_k},\ \overline{i+\lambda_k+1},\ \ldots,\ \overline{i+\lambda_k+\lambda_{k-1}-1},$$

and so on. Thus the third and fourth tabloids in Fig. 2 are the reverse row superstandard tabloids of shape $\langle 3, 2, 1\rangle$ with starts at 1 and 3, respectively. Similarly, one can define the *column superstandard tabloid* of shape λ with start at i as the tabloid that has $\overline{i}$ in the first row, $\overline{i+1}$ in the second row, …, $\overline{i+\ell(\lambda)-1}$ in the last row, $\overline{i+\ell(\lambda)}$ in the first row, and so on. The last two tabloids in the figure are column superstandard tabloids of shape $\langle 3, 2, 1\rangle$ with starts at 1 and 3, respectively.

[1] Classically, one considers fillings with integers, but nothing is lost by this slightly nonstandard choice.

If T is a tabloid then we denote by T_i the i-th row of T, viewed as a one-row tabloid or, equivalently, as a subset of $[\overline{n}]$. Similarly, we denote by $T_{i,i+1}$ the tabloid consisting of the i-th and $(i+1)$-st rows of T, and by $T_{[i,j]}$ the tabloid consisting of all rows of T with indices between i and j, inclusive.

2.2 An analogue of the Robinson–Schensted correspondence

The paper [8] describes a bijection

$$\Phi : \widetilde{S_n} \to \Omega_{\mathrm{dom}},$$

where Ω_{dom} is the set of triples (P, Q, ρ) such that P and Q are tabloids of the same shape of size n and ρ is an integer vector satisfying certain inequalities depending on P and Q. In this section, we give the relevant definitions and outline the construction of the bijection via an algorithm called the Affine Matrix-Ball Construction (AMBC). A detailed summary of these results occupies the first part of [8]; while the present work is not completely independent of that part of [8], we hope that the short version provided here is sufficient to understand most of the new results in the present paper.

Define Ω to be the collection of triples (P, Q, ρ) where P and Q are tabloids of the same shape λ of size n and $\rho \in \mathbb{Z}^{\ell(\lambda)}$. Suppose $(P, Q, \rho) \in \Omega$ and P and Q have shape λ. For i such that $\lambda_{i-1} = \lambda_i$, there are associated integers $r_i(P, Q)$ called *offset constants* (which are defined precisely in Definition 5.9). If (P, Q, ρ) satisfies the conditions $\rho_i \geqslant \rho_{i-1} + r_i(P, Q)$ then we say that ρ is *dominant* with respect to (P, Q); the pair (P, Q) is usually clear from the context. We define

$$\Omega_{\mathrm{dom}} = \{(P, Q, \rho) \in \Omega : \rho \text{ is dominant}\}.$$

A second algorithm [8, Sect. 4], referred to as the *backward algorithm*, gives a surjection $\Psi : \Omega \to \widetilde{S_n}$. By [8, Thms. 5.11 & 6.3], the restriction of Ψ to Ω_{dom} is equal to Φ^{-1}.

We now briefly describe AMBC, which is closely related to Viennot's geometric construction [23] for the Robinson–Schensted correspondence, called the Matrix-Ball Construction by Fulton [10]. The first step is to produce a special numbering of the balls of the permutation, called a *channel numbering*, as in Fig. 1. This numbering partitions the balls into equivalence classes having the same number; for each class, we form the *zig-zag* having those balls as inner corners, as in Figs. 5 and 6. We get a new partial permutation from the outer corners of the zig-zags, and we iterate. At each step of the iteration, we record basic data about the *back corner-posts* of the zig-zags (illustrated by $*$'s in Figs. 5 and 6): their column-indices are recorded in a row of P, the row-indices in a row of Q, and their *altitude* (a shift parameter distinguishing among the collections occupying the same collection of rows and columns) in an entry of ρ. The remainder of this section is concerned with providing the details behind this summary.

2.2.1 Channel numberings

Definition 2.1 Suppose we have a collection C of cells that is invariant under translation by (n, n) and forms a chain in the partial ordering $\leqslant_{SE}$. Then the *density* of C is the number of distinct translation classes in C.

In Fig. 1, the collection consisting of balls $(2, 1)$, $(5, 2)$, $(6, 3)$, and their translates has density 3.

Definition 2.2 Suppose w is a partial permutation. Then $C \subseteq \mathcal{B}_w$ is a *channel* if all of the following hold:

- C is invariant under translation by (n, n),
- C forms a chain in the partial ordering $\leqslant_{SE}$, and
- the density of C is maximal among all subsets of $\mathcal{B}_w$ satisfying the first two conditions.

In Fig. 1, the collection consisting of balls $(2, 1)$, $(5, 2)$, $(6, 3)$, and their translates is a channel. On the other hand, balls $(1, 4)$ and $(3, 6)$ are not part of any channel, since no translation-invariant chain of the ordering $\leqslant_{SE}$ with density 3 passes through both.

A curious fact about channels is that the collection of southwest-most (or northeast-most) balls of a union of two channels is again a channel (see [8, Prop. 3.13] and the comment following it). Thus there is always a southwest-most channel (in Fig. 1, it is the channel formed by taking the balls $(2, 1)$, $(5, 2)$, $(6, 3)$ and their translates) and a northeast-most channel. In AMBC, we need to pick a distinguished channel at each step; we pick the southwest one.

To perform a step of AMBC requires a special kind of numbering of a permutation, which we define now.

Definition 2.3 For a partial permutation w, a function $d : \mathcal{B}_w \to \mathbb{Z}$ is a *proper numbering* if it is

- monotone: for any $b, b' \in \mathcal{B}_w$, if b lies strictly northwest of b' then $d(b) < d(b')$, and
- continuous: for any $b' \in \mathcal{B}_w$ there exists b northwest of it with $d(b) = d(b') - 1$.

An example of a proper numbering is given in Fig. 1.

Proposition 2.4 [8, Prop. 3.4] *Given a partial permutation w, let $m(w)$ denote the density of the channels of w. All proper numberings of w are semi-periodic with period $m(w)$.*

Definition 2.5 Suppose w is a partial permutation and C is a channel of w. For a ball $b \in \mathcal{B}_w$ and $k \in \mathbb{Z}^{\geqslant 0}$, a *path* of length k from b to C is a sequence of balls $(b_0, b_1, \ldots, b_k)$ such that $b_0 = b$, $b_k \in C$, and b_{i+1} lies strictly northwest of b_i for all i.

Suppose we have a channel C of some partial permutation w. We can ignore all other balls of w and ask for proper numberings of just C itself. It is clear that up to an overall shift there is just one of them, with the balls numbered consecutively by all the integers as one moves from northwest to southeast. We can use this proper numbering of C to produce a proper numbering of all the balls of w.

Definition 2.6 Suppose C is a channel of w and $\tilde{d}$ is a proper numbering of C. Define the *channel numbering* $d_w^C : \mathcal{B}_w \to \mathbb{Z}$ of w by

$$d_w^C(b) := \max_k \max_{(b_0,b_1,\ldots,b_k)} (\tilde{d}(b_k) + k),$$

where the second maximum is taken over all paths from b to C.

Because the number of paths is infinite, it seems *a priori* that $\tilde{d}(b_k) + k$ could be unbounded and so $d_w^C(b)$ undefined; however, by [8, Prop. 3.9] this is not the case. Once we know that $d_w^C(b)$ is finite for any ball b, it is easy to see that d_w^C is, in fact, a proper numbering.

Remark 2.7 Given a channel, there is an infinite family of channel numberings associated with it: they differ by shifting the numbers of all balls by the same amount. In this paper, the distinctions between these numberings never matter, and we will use the phrase "the channel numbering" to refer to a (locally fixed, but) arbitrary choice among them.

2.2.2 Zig-zags

The definition of AMBC involves certain collections of cells called *zig-zags*.

Definition 2.8 A *zig-zag* is a non-empty sequence $(c_1, c_2, \ldots, c_k)$ of cells such that both of the following hold:

- for $1 \leqslant i < k$, c_{i+1} is adjacent to and either directly north or directly east of c_i, and
- if $k \geqslant 2$, then c_2 is directly east of c_1 and c_k is directly north of c_{k-1}.

Given a zig-zag $Z = (c_1, c_2, \ldots, c_k)$, we say that

- the *back corner-post* is the cell in the same column as c_1 and the same row as c_k,
- the *inner corner-posts* are the cells of Z such that no cell directly north or directly west of them is in Z, and

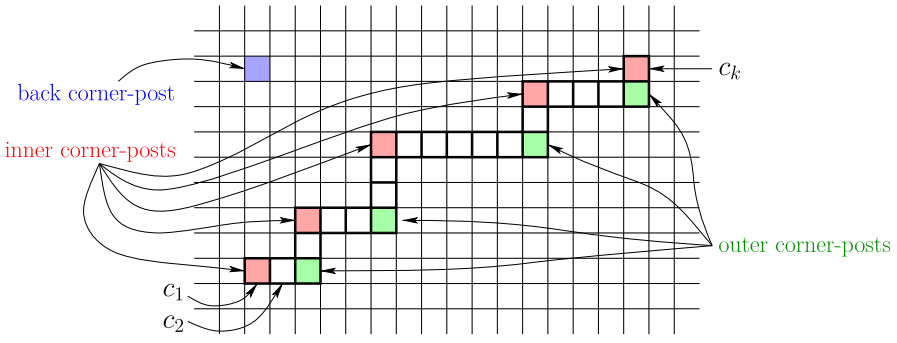

Fig. 3 The different corner-posts of a zig-zag

- if $k \geqslant 2$, the *outer corner-posts* are the cells of Z such that no cell directly south or directly east of them is in Z; if $k = 1$ then there are no outer corner-posts of Z.

These definitions are illustrated in Fig. 3. Notice that the inner and outer corner-posts are always part of the zig-zag. The back-corner post is not usually part of the zig-zag; the only exception is a degenerate zig-zag with one cell, whose back corner-post coincides with its inner corner-post (and which has no outer corner-post).

The zig-zags that appear in AMBC are attached to a proper numbering of a permutation.

Definition 2.9 Given a proper numbering d of a partial permutation w, the collection of *zig-zags corresponding to* d is the collection $\{Z_i\}_{i\in\mathbb{Z}}$ where Z_i is the unique zig-zag whose inner corner-posts are precisely the balls of w labeled i by d.

The number of translation classes of zig-zags is the number $m(w)$ appearing in Proposition 2.4.

2.2.3 Streams

The definition of AMBC involves a certain collection of cells called a *stream*: a set that is invariant under translation by (n, n) and forms a chain in the partial ordering $\leqslant_{SE}$. Since streams are translation-invariant, if a stream has a cell with column k, then it has cells in all the columns of $\overline{k}$.

Definition 2.10 For any cell $c = (c_1, c_2)$, let $D(c)$ be the *block diagonal* of c,

$$D(c) := \left\lceil \frac{c_2}{n} \right\rceil - \left\lceil \frac{c_1}{n} \right\rceil.$$

Thus, if one of the translates of c is in $[n] \times [n]$ then $D(c) = 0$; if a translate is in $\{n+1, n+2, \ldots, 2n\} \times [n]$ then $D(c) = -1$; etc.

Definition 2.11 For any translation-invariant collection X of cells, we define $D(X) = \sum_{x\in Y} D(x)$, where Y is a subset of X containing one representative of each translation class. If X is a stream, we call $D(X)$ the *altitude* of X.

In [8, Sect. 3.4], a different definition was given of altitude of a stream; however, it is an easy exercise to see that they are equivalent. The following result shows that the set of rows, set of columns, and altitude uniquely specify a stream.

Proposition 2.12 (Essentially [8, Lem. 3.23 & Prop. 3.24]) *Given two subsets A, B of $[\overline{n}]$ of the same size and an integer r, there is a unique stream of altitude r whose balls lie in rows indexed by A and in columns indexed by B.*

This stream will be denoted $\mathfrak{st}_r(A, B)$; the triple A, B, r is called its *defining data*.

Example 2.13 Let $n = 6$, $A = \{\overline{1}, \overline{3}, \overline{6}\}$, and $B = \{\overline{2}, \overline{4}, \overline{5}\}$. The streams $\mathfrak{st}_{-1}(A, B)$, $\mathfrak{st}_0(A, B)$, and $\mathfrak{st}_1(A, B)$ are shown in Fig. 4. The stream $\mathfrak{st}_0(A, B)$ is the unique choice with all of its cells in translations of $[n] \times [n]$; $\mathfrak{st}_1(A, B)$ is obtained from $\mathfrak{st}_0(A, B)$ by moving the cell in every row from its column east into the next available column in $\bigcup B$.

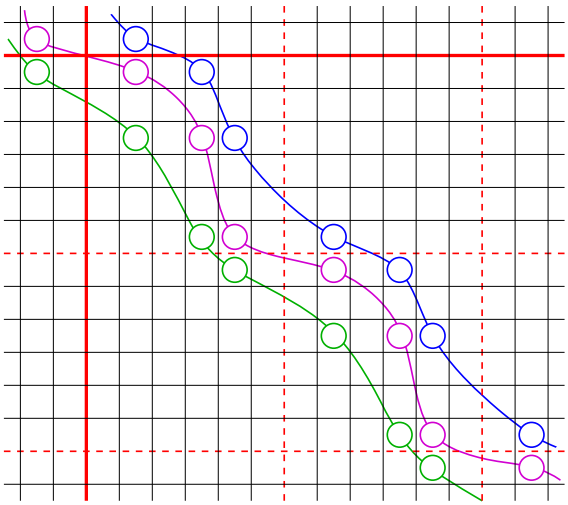

Fig. 4 Streams of altitude -1 (green), 0 (purple), and 1 (blue) for $A = \{\overline{1}, \overline{3}, \overline{6}\}$ and $B = \{\overline{2}, \overline{4}, \overline{5}\}$

The streams that appear in AMBC come from a proper numbering of a permutation.

Definition 2.14 Suppose w is a partial permutation and $d : \mathcal{B}_w \to \mathbb{Z}$ is the southwest channel numbering; let $\{Z_i\}_{i\in\mathbb{Z}}$ be the collection of zig-zags corresponding to d. For each i, let b_i be the back corner-post of Z_i. Then $\mathfrak{st}(w) := \{b_i\}_{i\in\mathbb{Z}}$ is a stream.

2.2.4 The algorithm

We now define the algorithm AMBC, giving the map $\Phi : \widetilde{S}_n \to \Omega$.

Definition 2.15 Suppose w is a partial permutation and $d : \mathcal{B}_w \to \mathbb{Z}$ is the southwest channel numbering. Define $\mathrm{fw}(w)$ to be the permutation whose balls are located at the outer corner-posts of all the zig-zags corresponding to d.

The algorithm is as follows.

- Input $w \in \widetilde{S}_n$.
- Initialize (P, Q, ρ) to $(\varnothing, \varnothing, \varnothing)$.
- Repeat until w is the empty partial permutation:
  - Record the defining data of $\mathfrak{st}(w)$ in the next row of P, Q, and ρ.
  - Reset w to $\mathrm{fw}(w)$.
- Output $(P, Q, \rho) \in \Omega$.

Example 2.16 Let $n = 7$ and consider the permutation $w = [1, 2, 17, 5, 14, 18, 20]$. In the first step of AMBC, shown in Fig. 5, w is numbered in the SW channel numbering; the stream consisting of the back corner-posts for the zig-zags corresponding to this numbering has elements in rows $\overline{3}$, $\overline{6}$, and $\overline{7}$ and columns $\overline{1}$, $\overline{2}$, and $\overline{5}$, and has altitude 3. Thus, the first row of P is $\boxed{\overline{1}}\boxed{\overline{2}}\boxed{\overline{5}}$, the first row of Q is $\boxed{\overline{3}}\boxed{\overline{6}}\boxed{\overline{7}}$, and the first row of ρ

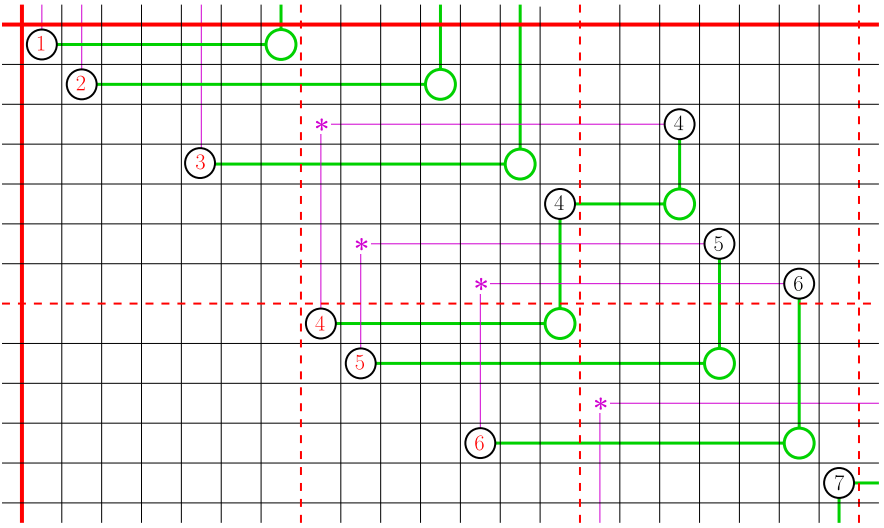

Fig. 5 First step of AMBC for $w = [1, 2, 17, 5, 14, 18, 20]$. The southwest channel numbering is shown with the balls in the channel numbered in red. The balls of $\mathrm{fw}(w)$ are shown in green. The cells of $\mathfrak{st}(w)$ are marked by $*$'s

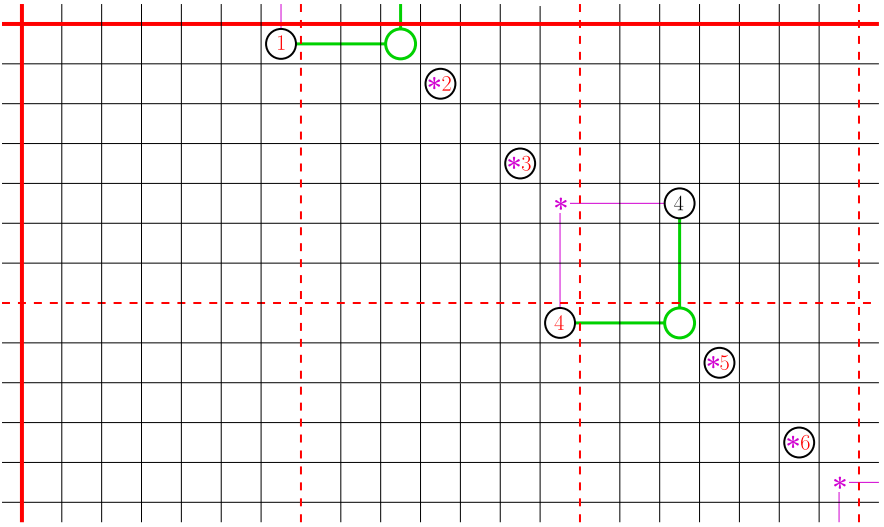

Fig. 6 Second step of AMBC for $w = [1, 2, 17, 5, 14, 18, 20]$

is 3. The second step is shown in Fig. 6, and produces second rows $\boxed{\bar{4}}\boxed{\bar{6}}\boxed{\bar{7}}$, $\boxed{\bar{2}}\boxed{\bar{4}}\boxed{\bar{5}}$ and 3 for P, Q, and ρ. Finally, the third step begins with the outer corner-posts from the second step; these form a partial permutation whose balls are exactly the translates of $(1, 10)$. Thus the stream for this permutation also consists of translates of the single

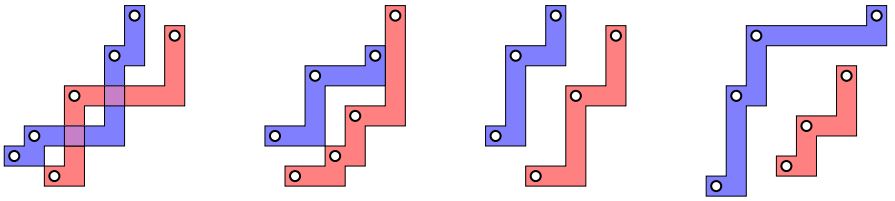

Fig. 7 The blue zig-zag has balls of w with lower value of d than the red zig-zag. The two zig-zag positions shown on the left cannot occur: the first satisfies condition (**b**) of Proposition 2.17 but violates condition (**a**), while the second satisfies condition (**a**) but violates condition (**b**). The two positions on the right are valid

cell $(1, 10)$, and so at the end of this step we set the third row of P to $\boxed{\bar{3}}$, the third row of Q to $\boxed{\bar{1}}$, and the third row of ρ to 1. The resulting triple (P, Q, ρ) is

$$\left(\begin{array}{|c|c|c|} \hline \bar{1} & \bar{2} & \bar{5} \\ \hline \bar{7} & \bar{4} & \bar{6} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \bar{3} \\ \hline \end{array} \quad , \quad \begin{array}{|c|c|c|} \hline \bar{7} & \bar{3} & \bar{6} \\ \hline \bar{2} & \bar{4} & \bar{5} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \bar{1} \\ \hline \end{array} \quad , \quad \begin{pmatrix} 3 \\ 3 \\ 1 \end{pmatrix} \right).$$

2.3 Zig-zags from a proper numbering

The following simple facts about how the zig-zags associated to a proper numbering lie in the plane will occur repeatedly in the arguments that follow. Both halves of the result are illustrated in Fig. 7.

Proposition 2.17 *Suppose that w is a partial permutation, with balls labeled by a proper numbering d. Divide the balls into zig-zags according to d.*

(a) *If $b = (k, w(k))$ and $c = (\ell, w(\ell))$ are consecutive balls in a zig-zag of w and $k < \ell$, then there are no balls of w having larger value of d northwest of the cell $(\ell, w(k))$.*
(b) *If b is the northeast (resp. southwest) ball of one zig-zag and c is the northeast (resp. southwest) ball of another zig-zag and $d(b) < d(c)$, then c lies strictly south (resp. east) of b.*

Proof For part (a), suppose for contradiction that there is such a ball a. By continuity of d (possibly applied many times), there is a ball northwest of a with label $d(b) = d(c) < d(a)$. Since b and c are given as consecutive balls in their zig-zag, this ball cannot be in the rectangle having b and c as vertices; however, by the monotonicity of d it also cannot be northwest of b or of c. But every cell northwest of $(\ell, w(k))$ falls into one of these three sets; this is a contradiction, so no such a exists, as claimed.

For part (b), suppose that c is the northeast ball in its zig-zag and $d(b) < d(c)$. By continuity of d, there is some ball b' of w northwest of c such that $d(b') = d(b)$. The northeast ball in this zig-zag is at least as far north as b', hence is north of c, as claimed. □

3 Knuth moves, descents, and inverses under AMBC

The main result of this section is Theorem 3.11, which describes precisely how making a Knuth move to a permutation affects its image under AMBC. It was already shown in [8, Lem. 12.3] that the P-tabloid does not change; we furthermore show that the Q-tabloid changes by a Knuth move and that ρ changes in a predictable way. This result is crucial to our analysis of the Kazhdan–Lusztig dual equivalence graph in Sects. 7 and 8.

Theorem 3.11 is also interesting in that it provides evidence that AMBC is the "correct" analogue of the Robinson–Schensted correspondence, as it interacts in the appropriate ways with other combinatorial constructions. The first steps towards the proof also are of this form, showing that AMBC respects inverses and descent sets of permutations in the same way as the Robinson–Schensted correspondence.

3.1 Inverses

In the finite case, it is a classical result of Schützenberger [18] (or see [21, Thm. 7.13.1]) that taking the inverse of a permutation w exchanges the insertion and recording tableaux in its image under the Robinson–Schensted correspondence. Here we show the analogous result for AMBC.

Proposition 3.1 *Suppose* $\Phi(w) = (P, Q, \rho)$. *Then* $\Phi(w^{-1}) = (Q, P, (-\rho)')$, *where* $(-\rho)'$ *is the dominant representative of* $-\rho$ *in the fiber under* $\Psi(Q, P, -)$.

Proof Inverting a permutation reflects its matrix across the main diagonal; this switches rows for columns, switches "SW" for "NE", and switches "E" for "S". Moreover, given a stream, reflecting it in the main diagonal negates the altitude of the stream. Thus, when we apply each step of AMBC to the inverse permutation w^{-1}, *using the NE channel numbering* (instead of the usual SW channel numbering) at each step, we recover exactly the same steps as applying AMBC to w using the SW channel numbering, with the following adjustments: first, because the roles of rows and columns are switched, data recorded for w in the P-tabloid is recorded for w^{-1} in the Q-tabloid and vice-versa; and second, because altitudes of streams are negated, each entry of $\rho(w^{-1})$ is the negation of the corresponding entry in $\rho(w)$. It follows from a strong form of the inverse relationship between Φ and Ψ [8, Prop. 5.2] that $\Psi(Q, P, -\rho) = w^{-1}$. Finally, we have by [8, Thm. 6.3] that $\Phi(w^{-1}) = \Phi(\Psi(Q, P, -\rho)) = (Q, P, (-\rho)')$ where $(-\rho)'$ is the dominant representative of $-\rho$. □

Remark 3.2 It is not obvious from the definition of dominance that applying the operation $(P, Q, \rho) \mapsto (Q, P, (-\rho)')$ twice returns the original triple (as it must do, since $(w^{-1})^{-1} = w$): the operation "take the dominant representative" depends on P and Q in a nontrivial way. This oddity is explained in Sect. 5, following the description of the offset constants in Theorem 5.10.

3.2 Descent sets and the τ-invariant

In the finite case, the descent set (appropriately defined) of the Q-tableau of w is equal to the (right) descent set of w (i.e., the set of integers i such that $w(i) > w(i+1)$), and the descent set of the P-tableau is the descent set of w^{-1} [21, Lem. 7.23.1]. Here we show the analogous result for affine permutations and AMBC.

Definition 3.3 For a partial permutation w, the *right descent set* $R(w)$ of w is defined by

$$R(w) = \{\bar{i} \in [\overline{n}] : w(i) > w(i+1)\}.$$

Similarly, the *left descent set* $L(w)$ is defined by

$$L(w) = \{\bar{i} \in [\overline{n}] : w^{-1}(i) > w^{-1}(i+1)\}.$$

So $\bar{i}$ is in $R(w)$ (resp. $L(w)$) precisely when the ball in the $(i+1)$-st row (resp. column) is west (resp. north) of the ball in the i-th row (resp. column). These definitions agree with the usual notions from Coxeter theory.

We call the analogue of the descent set for tabloids the τ-*invariant*, in reference to Vogan's (generalized) τ-invariant [24].

Definition 3.4 For a tabloid T filled with all the elements of $[\overline{n}]$, define the τ-*invariant* by

$$\tau(T) := \{\bar{i} \in [\overline{n}] : \bar{i} \text{ lies in a strictly higher row of } T \text{ than } \overline{i+1}\}.$$

Example 3.5 The tabloid

$$T := \begin{array}{|c|c|c|}\hline \bar{2} & \bar{9} & \bar{5} \\ \hline \bar{8} & \bar{7} & \bar{6} \\ \hline \bar{3} & \bar{1} \\ \cline{1-2} \bar{4} \\ \cline{1-1} \end{array}$$

has $\tau(T) = \{\bar{2}, \bar{3}, \bar{5}, \bar{9}\}$.

Proposition 3.6 *For any permutation w, $L(w) = \tau(P(w))$ and $R(w) = \tau(Q(w))$.*

Proof It is sufficient to show $L(w) = \tau(P(w))$; the other statement follows from Proposition 3.1.

First, suppose $\bar{i} \in L(w)$, so that the ball b in the i-th column is south of the ball c in the $(i+1)$-st column. It follows that every ball strictly northwest of c is also northwest of b, and so by the continuity and monotonicity of the southwest channel numbering d we have $d(c) \leqslant d(b)$.

Consider the application of a single forward step of AMBC to w. The numbering d induces a set of zig-zags. Since b is west of c, we have by Proposition 2.17(b) that c is not the southwest ball in its zig-zag, and so there is a ball of $\mathrm{fw}(w)$ in the $(i+1)$-st column; moreover, by Proposition 2.17(a) this ball is north of b. Either b is the southwest ball of its zig-zag, or not. In the first case, $\bar{i}$ appears in the first row

of $P(w)$ while $\overline{i+1}$ appears in a lower row. In the second case, neither $\bar{i}$ nor $\overline{i+1}$ appears in the first row of $P(w)$, and the ball in the i-th column of $\mathrm{fw}(w)$ is south of the ball in the $(i+1)$-st column; thus we can repeat the argument until we end up in the first case.

Conversely, if $\bar{i} \notin L(w)$ then the ball b in the i-th column is north of the ball c in the $(i+1)$-st column. In this case $d(b) < d(c)$ by monotonicity. If c is the southwest ball of its zig-zag then $\overline{i+1}$ appears in the first row of $P(w)$; thus $\bar{i} \notin \tau(P(w))$, as claimed. Alternatively, c is not the southwest ball of its zig-zag. Then the ball c' southwest of c in its zig-zag is west of b, and by Proposition 2.17(b) b is not the southwest ball of its zig-zag. Moreover, it follows from Proposition 2.17(a) that the ball of $\mathrm{fw}(w)$ in the i-th column must be north of the ball in the $(i+1)$-st column. Thus we can repeat the argument to see that $\overline{i+1}$ cannot lie in a lower row of $P(w)$ than $\bar{i}$, as claimed. □

3.3 Knuth moves

Graphs whose vertices are combinatorial objects such as permutations or tableaux and whose edges are analogues of Knuth moves occur frequently in the study of the symmetric group—see, for example, [3, 5, 12]. In the affine setting, we deal with two kinds of such graphs, one on the set of permutations and one on the set of tabloids.

3.3.1 Knuth moves for permutations and tabloids

In the finite setting, Knuth moves are certain elementary operations on finite permutations, interchanging adjacent entries if one of the neighboring entries is numerically between them—see [21, Ch. 7 App. 1]. The definition of Knuth moves in the affine case is extremely similar.

Definition 3.7 Let $i \in \mathbb{Z}$. Two partial permutations w and w' are *connected by a Knuth move at position* $\bar{i}$ if all of the following hold:

- for all j such that $j \equiv i \pmod{n}$, we have $w'(j) = w(j+1)$ and $w'(j+1) = w(j)$;
- for all j such that $j \not\equiv i \pmod{n}$, $j \not\equiv i+1 \pmod{n}$, we have $w'(j) = w(j)$; and
- at least one of $w(i+2)$ and $w(i-1)$ is numerically between $w(i)$ and $w(i+1)$.

For example, the permutation $w = [3, 1, 2]$ is connected by a Knuth move to $[1, 3, 2]$ (because $w(3) = 2$ has value between $w(1) = 3$ and $w(2) = 1$) and to $[-1, 1, 6]$ (because $w(2) = 1$ has value between $w(0) = -1$ and $w(1) = 3$) but not to $[3, 2, 1]$ (because neither $w(1) = 3$ nor $w(4) = 6$ have value between $w(2) = 1$ and $w(3) = 2$).

We call the translation class of the ball $(i+2, w(i+2))$ or $(i-1, w(i-1))$ mentioned in the last part of the definition a *witness* to the Knuth move.

An alternative way to describe the condition when exchanging the i-th and $(i+1)$-st entries of a finite permutation constitutes a Knuth move is via the right descent set. It is not difficult to see that the condition on the witness precisely means that the right

descent sets of the two permutations are incomparable under the containment partial ordering. Similarly, we can give an alternate condition in the affine case.

Proposition 3.8 *Suppose w is a partial permutation such that $w(i)$, $w(i+1)$, and $w(i+2)$ are defined, and precisely one of $\overline{i}$ and $\overline{i+1}$ is in $R(w)$. Then there exists a unique permutation w' connected to w by a Knuth move whose right descent set contains precisely the other of $\overline{i}$ and $\overline{i+1}$.*

Proof If exactly one of $\overline{i}, \overline{i+1}$ belongs to $R(w)$ then $w(i+1)$ is either the largest or smallest of $w(i), w(i+1), w(i+2)$, and so one of $w(i), w(i+2)$ lies numerically between $w(i+1)$ and the other. Thus a Knuth move that switches $w(i+1)$ with one of $w(i), w(i+2)$ is possible, witnessed by the other of $w(i), w(i+2)$. It is straightforward to verify that this has the desired effect on the descent set, and that no other Knuth move can have this effect. □

In the finite case, Knuth moves preserve the Robinson–Schensted insertion tableau and induce transpositions of entries (which are also called Knuth moves) in the recording tableau. We give the corresponding definition of Knuth moves on tabloids.

Definition 3.9 Two tabloids T and T' are *connected by a Knuth move* if T' is obtained from T by exchanging $\overline{i}$ and $\overline{i+1}$ and $\tau(T)$ and $\tau(T')$ are incomparable with respect to inclusion.

Example 3.10 The tabloid

$$T' := \begin{array}{|c|c|c|}\hline \overline{2} & \overline{1} & \overline{5} \\ \hline \overline{8} & \overline{7} & \overline{6} \\ \hline \overline{3} & \overline{9} \\ \cline{1-2} \overline{4} \\ \cline{1-1} \end{array}$$

is connected by a Knuth move to the tabloid T from Example 3.5: it is obtained by exchanging $\overline{9}$ and $\overline{10} = \overline{1}$, and $\tau(T') = \{\overline{2}, \overline{3}, \overline{5}, \overline{8}\}$ is not contained in and does not contain $\tau(T) = \{\overline{2}, \overline{3}, \overline{5}, \overline{9}\}$.

3.3.2 Knuth moves and AMBC

In this section we describe how a Knuth move on permutations looks after taking images under AMBC.

Theorem 3.11 *Suppose w is a partial permutation and w' differs from w by a Knuth move. Then $P(w) = P(w')$ and $Q(w')$ differs from $Q(w)$ by a Knuth move. Moreover, if the Knuth move on the Q-tabloids exchanges $\overline{i}$ and $\overline{i+1}$ for some $\overline{i} \neq \overline{n}$ then $\rho(w) = \rho(w')$; if instead $\overline{i} = \overline{n}$ is in row k of $Q(w)$ while $\overline{i+1} = \overline{1}$ is in row k', then $\rho(w')$ differs from $\rho(w)$ by subtracting 1 from row k and adding 1 to row k'.*

Example 3.12 Consider the permutation $w = [1, 4, 6, 2, 5, 3]$; under AMBC it corresponds to

$$P = Q = \begin{array}{|c|c|c|}\hline \overline{1} & \overline{2} & \overline{3} \\ \hline \overline{4} & \overline{5} \\ \cline{1-2} \overline{6} \\ \cline{1-1} \end{array}, \quad \rho = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

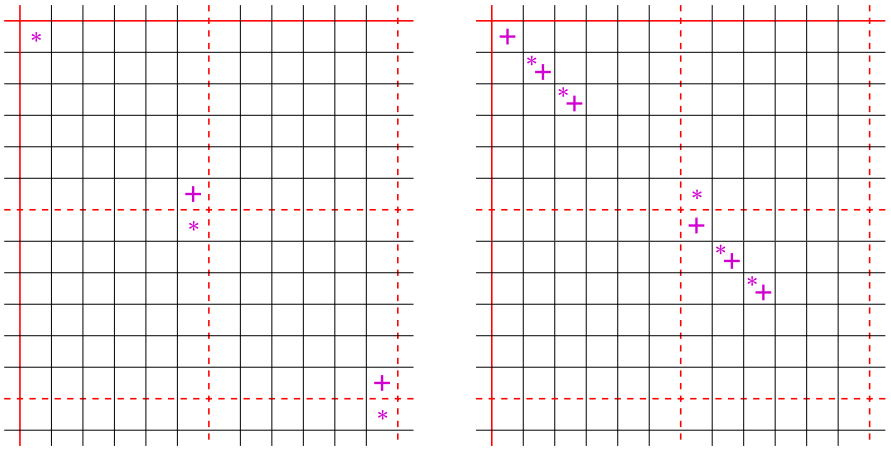

Fig. 8 Streams for the permutation $w = [1, 4, 6, 2, 5, 3]$ (+'s) and its image $[-3, 4, 6, 2, 5, 7]$ (∗'s) under a Knuth move; see Remark 3.14. On the left, the streams S, S' encoded by the third row of the images under AMBC. On the right, the streams T, T' encoded by the first row

Since $w(4) < w(2) < w(3)$, we can apply a Knuth move at position 3. This yields $[1, 4, 2, 6, 5, 3]$, which corresponds under AMBC to the triple

$$P = \begin{array}{|c|c|c|}\hline \overline{1} & \overline{2} & \overline{3} \\ \hline \overline{4} & \overline{5} \\ \cline{1-2} \overline{6} \\ \cline{1-1} \end{array}, \quad Q' = \begin{array}{|c|c|c|}\hline \overline{1} & \overline{2} & \overline{4} \\ \hline \overline{3} & \overline{5} \\ \cline{1-2} \overline{6} \\ \cline{1-1} \end{array}, \quad \rho = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

Example 3.13 Again consider $w = [1, 4, 6, 2, 5, 3]$. Since $w(6) < w(5) < w(7) = 7$, we can apply a Knuth move to w at position 6. This yields $[-3, 4, 6, 2, 5, 7]$, which corresponds under AMBC to the triple

$$P = \begin{array}{|c|c|c|}\hline \overline{1} & \overline{2} & \overline{3} \\ \hline \overline{4} & \overline{5} \\ \cline{1-2} \overline{6} \\ \cline{1-1} \end{array}, \quad Q'' = \begin{array}{|c|c|c|}\hline \overline{6} & \overline{2} & \overline{3} \\ \hline \overline{4} & \overline{5} \\ \cline{1-2} \overline{1} \\ \cline{1-1} \end{array}, \quad \rho' = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}.$$

Remark 3.14 The condition on how the weight changes may be formulated in terms of the streams encoded by the rows. Suppose that Q and Q' differ by a Knuth move exchanging $\overline{i}$ and $\overline{i+1}$, that $\overline{i}$ is in row k of Q, and that and $\overline{i+1}$ is in row k' of Q. Suppose ρ' differs from ρ as described in the theorem. Then the stream S' represented by row k of (P, Q', ρ') differs from the stream S represented by row k of (P, Q, ρ) by shifting the elements with row indices in $\overline{i}$ south by one cell, and similarly the stream T' encoded by row k' of (P, Q', ρ') differs from the stream T represented by row k' of (P, Q, ρ) by shifting the elements with row indices in $\overline{i+1}$ north by one cell.

For example, with w as in Example 3.13, the streams S and S' (shown on the left of Fig. 8) are encoded by the third row of the triples, while the streams T and T' (shown on the right of the figure) are encoded by the first row of the triples. The stream S contains only one translation class of cells, in row $\overline{n}$; to obtain the stream S', one shifts

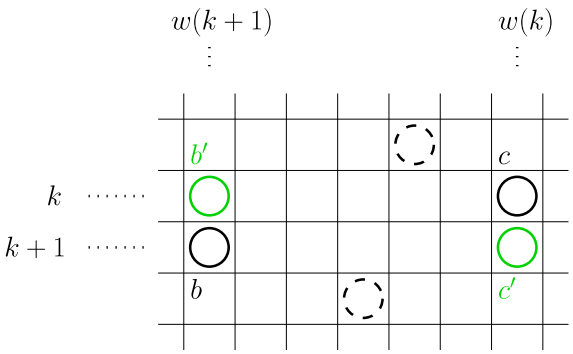

Fig. 9 The positions of the balls in Sect. 3.3.3. At least one of the dashed balls must be present as a witness for the Knuth move

each of those elements one cell south. The stream T contains three translation classes of cells; to obtain the stream T', one shifts the elements in rows indexed by $\overline{1}$ one cell north.

3.3.3 The proof of Theorem 3.11

We fix some notation for the duration of this section. Let w be a partial permutation and let w' be the permutation obtained from it by a Knuth move, affecting balls in rows $\overline{k}$ and $\overline{k+1}$. Without loss of generality, assume $w(k+1) < w(k)$. Let $b = (k+1, w(k+1)) \in \mathcal{B}_w$, $c = (k, w(k)) \in \mathcal{B}_w$, $b' = (k, w(k+1)) \in \mathcal{B}_{w'}$, and $c' = (k+1, w(k)) \in \mathcal{B}_{w'}$. This situation is shown in Fig. 9.

Let d be the southwest channel numbering of w. Since every ball strictly northwest of b is also northwest of c, we have by the properness of the numbering that $d(c) \geqslant d(b)$. Define a semi-periodic numbering d' of w' as follows: for $a \in \mathcal{B}_{w'}$,

$$\begin{aligned} d'(a) &= d(a) && \text{if } a \text{ is not a translate of } b' \text{ or } c', \\ d'(b') &= d(b), && \text{and} \\ d'(c') &= \begin{cases} d(c) & \text{if } d(c) > d(b), \\ d(c)+1 & \text{if } d(c) = d(b), \end{cases} \end{aligned}$$

and extend the numbering to all other balls by semi-periodicity. In other words, to get from w to w', we shift b (together with its translates) one cell north and shift c one cell south; the numbering d' is obtained by keeping the numbering of the balls being shifted, unless $d(b) = d(c)$ (in which case keeping the numbering would yield a non-monotone numbering). As the next lemma shows, this adjustment results in a proper numbering.

Definition 3.15 The proofs that follow are divided into cases. We denote by **A** (for *after*) the case when $w(k+2)$ is between $w(k)$ and $w(k+1)$, and by **B** (for *before*) the case when $w(k-1)$ is between $w(k)$ and $w(k+1)$. We denote by **S** (for *same*) the case when $d(b) = d(c)$, and by **D** (for *different*) the case when $d(c) > d(b)$.

A priori the preceding definition allows four sub-cases. However, the case **BS** is impossible when d is proper: using Proposition 2.17, one can show that for a witness a in

the row before b and c, one must have $d(b) \leqslant d(a) < d(c)$, and this contradicts $d(b) = d(c)$.

Lemma 3.16 *The numbering d' described above is proper.*

Proof We must check the monotonicity and continuity of the function d'. In case **D**, this is straightforward: no ball numbers change and no relation $a \leqslant_{SE} a'$ between two balls is destroyed, so d' is continuous; and the only newly created relations in the SE order are translates of $b' >_{SE} c'$, and these are compatible with monotonicity since $d'(b') = d(b) < d(c) = d'(c')$.

Suppose instead that we are in case **AS**. Let $a = (k+2, w(k+2))$. Every ball strictly northwest of a is strictly northwest of b, strictly northwest of c, or is b itself. Since $d(b) = d(c)$, the largest label of such a ball is $d(b)$, and thus $d'(a) = d(a) = d(b) + 1 = d'(c')$. Then to check continuity and monotonicity of d' it suffices to observe that every ball southeast of c' in w' is also southeast of a. □

Definition 3.17 We denote by f the natural bijection $f : \mathcal{B}_w \to \mathcal{B}_{w'}$ that commutes with translation by (n, n), has $f(b) = b'$ and $f(c) = c'$, and is the identity on balls that do not lie in rows congruent to k or $k+1$ modulo n.

Lemma 3.18 *If C is the southwest channel of w then $f(C)$ is a channel of w' that intersects the southwest channel of w'.*

Proof Let C_1 be any channel of w. For any two balls a, a' in $\mathcal{B}_w$ such that $a <_{SE} a'$, one also has $f(a) <_{SE} f(a')$. Thus $C_1' := f(C_1)$ forms a chain in the southeast ordering. By Lemma 3.16, the numbering d' is proper, so by Proposition 2.4 w' has the same channel density as w. Since C_1 has this density, C_1' does as well, and so C_1' is a channel for w'.

Let C_{SW}' be the southwest channel of w'. By [8, Prop. 3.13] and the surrounding discussion, it follows that every ball of C_{SW}' is weakly southwest of some ball in C_1'. Suppose furthermore that $C_1' \cap C_{SW}' = \varnothing$, so that we may replace the word "weakly" in the preceding sentence with "strictly". We will show that there is a channel of w that contains a ball strictly southwest of some ball of C_1.

Suppose first that C_{SW}' contains at most one of b' and c'. Then $f^{-1}(C_{SW}')$ is a channel of w. The map f^{-1} is order-preserving for the order $<_{SW}$, so in this case every ball of $f^{-1}(C_{SW}')$ is strictly southwest of some ball of $C_1 = f^{-1}(C_1')$, as desired.

Suppose instead that C_{SW}' contains both b' and c'. Then C_1' contains neither, and so $C_1' = C_1$. In case **B**, removing b' and c' from C_{SW}' and replacing them with $(k-1, w(k-1))$ and c gives the desired channel of w. Similarly, in case **A**, removing b' and c' from C_{SW}' and replacing them with b and $(k+2, w(k+2))$ gives the desired channel of w.

We now take the contrapositive: if $C_1 = C$ is the southwest channel of w then (again by [8, Prop. 3.13]) there is no channel of w that contains a ball strictly southwest of some ball of C, and so $C_1' = f(C)$ must have nonempty intersection with C_{SW}', as claimed. □

Recall from Definition 2.5 that a *path* in a permutation w is a sequence of balls of w in which each ball is strictly northwest of the preceding one.

Lemma 3.19 *The numbering d' is the southwest channel numbering of w'.*

Proof Let C be the southwest channel of w. First, we show that for any $x' \in \mathcal{B}_{w'}$, there is a path in $\mathcal{B}_{w'}$ from x' to $f(C)$ along which d' decreases by 1 at each step.

Fix $x' \in \mathcal{B}_{w'}$ and let $x = f^{-1}(x')$. By [8, Rem. 11.7], there is a path $p = (x_0 = x, x_1, \ldots, x_k)$ in w beginning at x and ending at $x_k \in C$ with d decreasing by 1 at each step. By [8, Rem. 11.2], we may choose p so that it does not contain any pair of balls in the same translation class. Then $p' = (f(x_0), \ldots, f(x_k))$ is a path from $x' = f(x)$ to $f(C)$ in w'.

In case **D**, we have $d(x_i) = d'(f(x_i))$ and so d' decreases by 1 at each step of p', as desired.

In case **AS**, if p does not contain a translate of c then $d'(f(x_i)) = d(x_i)$ and so d' decreases by 1 at each step of p, as desired. Otherwise, we assume without loss of generality that c is part of p. In this case, c must be the first entry of p: if $c = x_{i+1}$ were northwest of x_i and $d(c) + 1 = d(x_i)$, then c' would be northwest of $f(x_i)$ and $d'(c') = d(c) + 1 = d(x_i)$; but this contradicts the monotonicity of d' (Lemma 3.16). Finally, we consider the possibility $x' = c'$. In this case, b' is northwest of c' and $d'(b') = d'(c') - 1$, and it follows from the preceding arguments in case **AS** that there is a path from b' to $f(C)$ with d' decreasing by 1 at each step; prepending c' to this path gives the desired one.

By [8, Rem. 11.7], it follows that d' is the channel numbering of w' with respect to $f(C)$. Finally, since (by Lemma 3.18) $f(C)$ intersects the southwest channel of w', we have by [8, Rem. 11.7 & Lem. 3.15] that d' is the southwest channel numbering of w', as claimed. □

For the next result, recall the notation $\mathtt{st}(w)$ from Definition 2.14.

Lemma 3.20 *Precisely one of the following holds:*

- $\mathtt{st}(w) = \mathtt{st}(w')$ *and* $\mathrm{fw}(w)$ *differs from* $\mathrm{fw}(w')$ *by a Knuth move, or*
- *for some* $\ell \in \mathbb{Z}$, $\mathtt{st}(w)$ *differs from* $\mathtt{st}(w')$ *by raising or lowering the cells in rows* $\overline{\ell}$ *by one row, and* $\mathrm{fw}(w)$ *differs from* $\mathrm{fw}(w')$ *by respectively lowering the balls in rows* $\overline{\ell - 1}$ *or raising those in rows* $\overline{\ell + 1}$ *by one row.*

Proof We begin with some preliminaries before moving on to case-analysis. By Lemma 3.19, the numbering d' is the southwest channel numbering of w', so the stream $\mathtt{st}(w')$ and permutation $\mathrm{fw}(w')$ are computed in terms of this numbering. Thus, in all cases, we consider the balls of w and w' to be divided into zig-zags corresponding to the numberings d and d', respectively. The position of every ball of $\mathrm{fw}(w)$ and every cell of $\mathtt{st}(w)$ is determined by two balls of w, and similarly for w'; we call these the *parents*. Finally, since $d(b) \leqslant d(c)$ and c is northeast of b, it follows from Proposition 2.17(b) that b is never the northeast ball in its zig-zag.

Case AD. Since we are in case **D**, for every $x \in \mathcal{B}_w$, we have $d(x) = d'(f(x))$. Therefore, if a zig-zag Z induced by d does not contain a translate of b or c, the same zig-zag (with the same inner-, outer-, and back-corner posts) is induced by d'.

Let Z_b be the zig-zag containing b induced by balls of w and let $Z_{b'}$ be the zig-zag containing b' induced by balls of w'. As noted above, b is not the northeast ball in Z_b. Thus, the east parents of the back-corner posts of Z_b and $Z_{b'}$ are the same, while

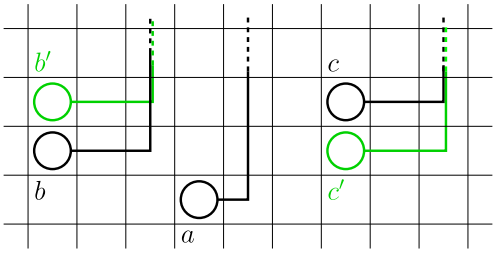

Fig. 10 The case **AD** of the proof of Lemma 3.20 when c is not the northeast ball of its zig-zag

their south parents are either the same (if b is not the southwest ball of Z_b) or lie in the same column; thus the back-corner posts are equal. Similarly, except the ball whose west parent is b, every ball of $\mathrm{fw}(w)$ in Z_b has west parent unchanged by f and north parent whose column is unchanged by f; thus, each of these balls is also a ball of $\mathrm{fw}(w')$. Finally, the balls of $\mathrm{fw}(w)$, $\mathrm{fw}(w')$ with west parents b, b' have the same north parent, so any ball with row index in $\overline{k+1}$ in $\mathcal{B}_{\mathrm{fw}(w)}$ moves up one cell to be a ball of $\mathcal{B}_{\mathrm{fw}(w')}$.

Now consider the zig-zags Z_c and $Z_{c'}$. Suppose first that c is *not* the northeast ball of its zig-zag. Then the same is true of c'. As in the previous paragraph, in this case Z_c and $Z_{c'}$ have the same back-corner post, while any ball with row index in $\overline{k}$ in $\mathcal{B}_{\mathrm{fw}(w)}$ moves down one cell to be a ball of $\mathcal{B}_{\mathrm{fw}(w')}$. Therefore in this case $\mathtt{st}(w) = \mathtt{st}(w')$ and the permutations $\mathrm{fw}(w)$ and $\mathrm{fw}(w')$ differ by moving the balls in rows $\overline{k+1}$ up into row $\overline{k}$ and moving the balls in rows $\overline{k}$ down into row $\overline{k+1}$, in the same way that w and w' differ (see Fig. 10). It remains to show that this is a Knuth move. Let $a = (k{+}2, w(k{+}2)) \in \mathcal{B}_w$ be the witness for the Knuth move between w and w'. Aside from b, every ball northwest of a is also northwest of c; thus by continuity of d we have $d(b) < d(a)$ while by monotonicity we have $d(a) \leqslant d(c)$. Using this last inequality together with Proposition 2.17(b), a is not the northeast ball in its zig-zag. Then it follows from Proposition 2.17(a) that $(\mathrm{fw}(w))(k) < (\mathrm{fw}(w))(k{+}2) < (\mathrm{fw}(w))(k{+}1)$. Thus the ball in row $k+2$ of $\mathrm{fw}(w)$ and $\mathrm{fw}(w')$ is a witness for the fact that $\mathrm{fw}(w)$ differs from $\mathrm{fw}(w')$ by a Knuth move in position $\overline{k}$.

If instead c *is* the northeast ball of its zig-zag, then the outer-corner posts of Z_c and $Z_{c'}$ are equal, so $\mathrm{fw}(w')$ differs from $\mathrm{fw}(w)$ by shifting the ball in row $k+1$ north by one cell, while the back-corner post of $Z_{c'}$ is in row $k+1$ instead of row k, so $\mathtt{st}(w')$ is obtained from $\mathtt{st}(w)$ by shifting the element in row k south by one cell.

Case BD. This case is similar to **AD**; the difference arises when choosing a witness in the case that c is not the northeast ball in its zig-zag. Let $a = (k-1, w(k-1)))$. In this case we have $d(b) \leqslant d(a) < d(c)$. The ball northeast of c in its zig-zag is also north of a, so a must not be the northeast ball in its own zig-zag. This implies $(\mathrm{fw}(w))(k+1) < (\mathrm{fw}(w))(k-1) < (\mathrm{fw}(w))(k)$, and so there is a ball in row $k-1$ of $\mathrm{fw}(w)$ and $\mathrm{fw}(w')$ that witnesses the claimed Knuth move.

Case AS. This case is more intricate because f affects parenthood in a nontrivial way. Let a be the ball (of both w and w') in row $k{+}2$, which witnesses the Knuth move. It was shown in the proof of Lemma 3.16 that in this case $d(a) = d(b){+}1 = d(c){+}1$.

We will now use this to show that $\mathrm{fw}(w)$ and $\mathrm{fw}\big(w'\big)$ differ at most in rows $\overline{k+1}$ and $\overline{k+2}$.

Zig-zags that do not contain a, b, c or their translates are unaffected by the Knuth move. Thus, we may restrict our attention to the two pairs of affected zig-zags. The first pair consists of the zig-zags $Z_{b,c}$, containing b and c induced by balls of w, and $Z_{b'}$, containing b' induced by balls of w'. These two zig-zags differ by removing two consecutive balls $(k, w(k))$ and $(k+1, w(k+1))$ and replacing them with a single ball $(k, w(k+1))$. This operation results in the deletion of the outer-corner post $(k+1, w(k))$, while the back-corner post and every other outer-corner post stays the same.

The second pair consists of the zig-zags Z_a, containing a induced by balls of w, and $Z_{a,c'}$, containing a and c' and induced by balls of w'. We have two cases to consider. If a is *not* the northeast ball in Z_a, suppose that the next ball northeast of a is $(j, w(j))$. Necessarily $j < k$. Thus c' is not the northeast ball in $Z_{a,c'}$, and the back-corner posts of $Z_{a,c'}$ and Z_a coincide. Therefore in this case $\mathfrak{st}(w) = \mathfrak{st}(w')$. Moreover, inserting the new inner-corner post c' between a and $(j, w(j))$ results in the deletion of their child at position $(k+2, w(j))$ and the creation of two new outer-corner posts at positions $(k+2, w(k))$ and $(k+1, w(j))$. Combining this with the previous paragraph, we have that $\mathrm{fw}\big(w'\big)$ differs from $\mathrm{fw}(w)$ by moving the ball in row $k+1$ south one cell into row $k+2$ and moving the ball in row $k+2$ north one cell into row $k+1$. Finally, we must show that this is a Knuth move. Since $(j, w(j)) \in Z_a$, we have by Proposition 2.17(b) that c is not the northeast ball in $Z_{b,c}$, and so also b' is not the northeast ball in $Z_{b'}$. Thus $\mathrm{fw}(w)$ and $\mathrm{fw}\big(w'\big)$ contain some ball (k, y) in row k, and by Proposition 2.17(a) we have $w(k) < y < w(j)$. Thus this ball is a witness to the fact that w and w' differ by a Knuth move in position $\overline{k+1}$.

Alternatively, suppose that a *is* the northeast ball in its zig-zag in $\mathcal{B}_w$. By an analysis similar to the preceding one, we see that $\mathrm{fw}(w) \smallsetminus \mathrm{fw}\big(w'\big)$ consists of all translates of $(k+1, w(k))$, while $\mathrm{fw}\big(w'\big) \smallsetminus \mathrm{fw}(w)$ consists of all translates of $(k+2, w(k))$, and that $\mathfrak{st}(w)$ has an element in row $k+2$ while in $\mathfrak{st}(w')$ this element appears one cell north, in row $k+1$. This is precisely the second situation in the proposition statement. □

With these preliminary results in hand, we move to the proof of the main result of this section.

Proof of Theorem 3.11 Fix a partial permutation w, related to the partial permutation w' by a Knuth move. Begin applying forward moves of AMBC to both w and w'. By Lemma 3.20, the streams of the two permutations may remain equal for several forward steps while the resulting permutations differ by a Knuth move; thus, the corresponding rows of $P(w)$, $Q(w)$ and $\rho(w)$ are respectively equal to those of $P(w')$, $Q(w')$ and $\rho(w')$. Since $w \neq w'$ and AMBC is a bijection, we must at some point reach a forward step where we fall into the second case of Lemma 3.20; without loss of generality, we may assume that this is the first step, and that for some $\ell \in \mathbb{Z}$, $\mathfrak{st}(w)$ differs from $\mathfrak{st}(w')$ by raising the elements in rows $\overline{\ell}$ by one cell, and $\mathrm{fw}(w)$ differs from $\mathrm{fw}\big(w'\big)$ by lowering the balls in rows $\overline{\ell-1}$ by one cell (as in Example 3.12).

In this case, the first rows of $P(w)$ and $P(w')$ are equal, the first rows of $Q(w)$ and $Q(w')$ differ by swapping $\overline{\ell-1}$ with $\overline{\ell}$, and $\mathrm{fw}(w)$ differs from $\mathrm{fw}\big(w'\big)$ by moving

balls between rows $\overline{\ell-1}$ and $\overline{\ell}$ (leaving the other row empty). If $\overline{\ell} \neq \overline{1}$ then the elements of $\mathfrak{st}(w)$ have the same block diagonals as the corresponding elements of $\mathfrak{st}(w)$, and so in this case the first rows of $\rho(w)$ and $\rho(w')$ are equal. If instead $\overline{\ell} = \overline{1}$ then the element of the stream that moves also decreases its block diagonal by 1, and so the first rows of $\rho(w)$ and $\rho(w')$ differ by 1 in this case (as in Example 3.13).

It is straightforward to check from the definition of AMBC that if u and u' are partial permutations such that u' is obtained from u by shifting the balls in rows $\overline{\ell-1}$ down by one cell into an empty row, then $P(u) = P(u')$ and $Q(u')$ is obtained from $Q(u)$ by replacing $\overline{\ell-1}$ with $\overline{\ell}$. If $\overline{\ell} \neq \overline{1}$ then the corresponding streams of both permutations have the same altitudes and so $\rho(u') = \rho(u)$; if instead $\overline{\ell} = \overline{1}$ then at the forward step that places $\overline{\ell}$ into $Q(u')$ and $\overline{\ell-1}$ into $Q(u)$, the stream coming from u' has altitude one less than the stream coming from u.

Finally, it follows from the preceding discussion that $P(w') = P(w)$, $Q(w')$ differs from $Q(w)$ be exchanging the two entries $\overline{\ell}$ and $\overline{\ell-1}$, and $\rho(w')$ differs from $\rho(w)$ as described in the theorem. The fact that the change from $Q(w)$ to $Q(w')$ is a Knuth move is a consequence of Proposition 3.6. □

3.4 Covering

In this short section, we prove that any Knuth move on tabloids can be realized as the change in Q-tabloids under a Knuth move on permutations.

Definition 3.21 For a partition λ, the *Kazhdan–Lusztig dual equivalence graph* (KL DEG) $\mathcal{A}_\lambda$ of shape λ is the graph whose vertices are the tabloids of shape λ and whose edges are the Knuth moves. Say that a permutation w has *shape* λ if the tabloids in its image under AMBC are of shape λ.

The KL DEG $\mathcal{A}_{\langle 4,1,1\rangle}$ is shown in Fig. 11.

Lemma 3.22 *Consider the graph on permutations of shape λ whose edges are Knuth moves. The map $w \mapsto Q(w)$ is a graph covering of $\mathcal{A}_\lambda$.*

Proof By Theorem 3.11, this map is a surjective graph morphism; we only need to check that it is an isomorphism around each vertex. Consider a permutation w. Suppose for some i, exactly one of $\overline{i}$ and $\overline{i+1}$ is in $R(w)$. By Proposition 3.8, there exists a unique Knuth move to a permutation w' with exactly the other of $\overline{i}$ and $\overline{i+1}$ is in $R(w')$. By considering the various possible relative locations of $\overline{i-1}$, $\overline{i}$, $\overline{i+1}$, and $\overline{i+2}$, it is easy to show that the same is true for tabloids with respect to τ. Combined with the fact that $R(w) = \tau(Q(w))$ (Proposition 3.6), this finishes the proof. □

In fact, the following stronger statement follows from the proof of Lemma 3.22.

Proposition 3.23 *If tabloids Q and Q' are related by a Knuth move, then for any tabloid P and any weight ρ, the permutation $w = \Psi(P, Q, \rho)$ is related by a Knuth move to a permutation w' with $Q(w') = Q'$.*

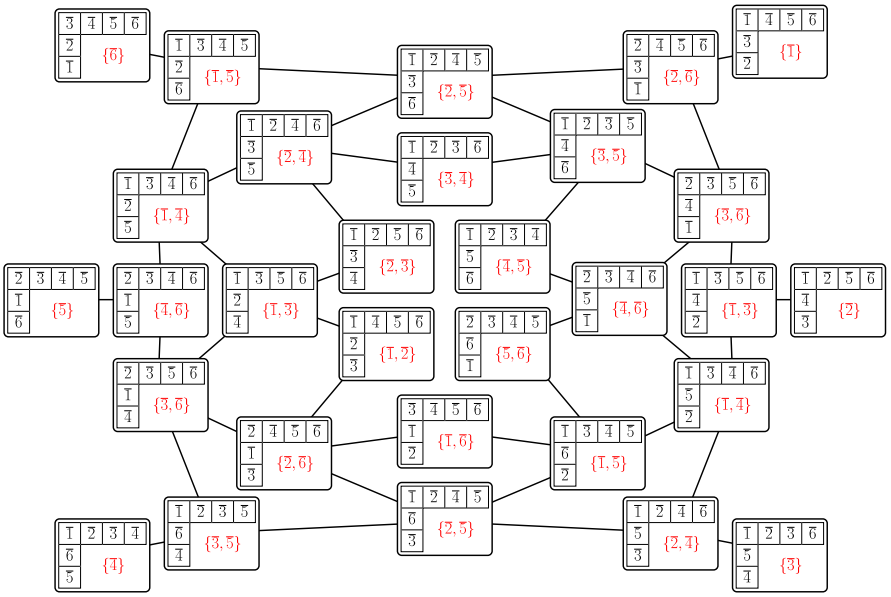

Fig. 11 The KL DEG $\mathcal{A}_{\langle 4,1,1\rangle}$. The τ-invariants are shown in red

4 Signs

It is folklore that recovering the length (i.e., number of inversions) of a finite permutation from its image under the Robinson–Schensted correspondence is hard. However, Reifegerste showed how to recover the *sign* of a permutation from its image [17]. In this section, we prove an analogous result in the affine case.

Reifegerste's theorem statement requires two definitions. First, given a standard Young tableau T of size n, define the *inversion number* $\mathrm{inv}(T)$ to be the number of pairs (i, j) such that $i < j$ and i appears strictly below j in T. Second, given a partition λ, define its *inversion number* $\mathrm{inv}(\lambda)$ to be the sum $\lambda_2 + \lambda_4 + \cdots$ of its even-indexed parts.

Theorem 4.1 (Reifegerste) *Given a finite permutation $w \in S_n$, let (P, Q) be the pair of standard Young tableaux associated to w by the Robinson–Schensted correspondence, both having shape λ. The sign of w is given by*

$$\mathrm{sgn}(w) = (-1)^{\mathrm{inv}(P)+\mathrm{inv}(Q)+\mathrm{inv}(\lambda)}.$$

Define the *length* $\ell(w)$ of a permutation w in $\widetilde{S}_n$ to be the number of *inversions*, that is, translation classes of pairs (a, b) of balls of w with a northeast of b. Define the *sign* $\mathrm{sgn}(w)$ of w to be $(-1)^{\ell(w)}$. If w belongs to the (non-extended) affine Weyl group $\widetilde{S}_n^0$, these definitions agree with the usual Coxeter length and sign of w; moreover, they extend naturally to give the sign and the length of any partial permutation.

To state our generalization of Theorem 4.1, we need some additional definitions.

Definition 4.2 There is a natural cyclic order on $[\overline{n}]$. In this section, it is convenient to "break" this order into a linear order, as follows:

$$\overline{1} < \overline{2} < \cdots < \overline{n}.$$

We refer to this total ordering as the *broken order*. (See Sect. 6 for more on this symmetry-breaking.)

Definition 4.3 Given a tabloid T, define the *inversion number* $\mathrm{inv}(T)$ to be the number of pairs $(\overline{i}, \overline{j})$ such that $\overline{1} \leqslant \overline{i} < \overline{j} \leqslant \overline{n}$ in the broken order and $\overline{i}$ appears strictly below $\overline{j}$ in T. Given a shape λ and a weight $\rho \in \mathbb{Z}^{\ell(\lambda)}$, define the *inversion number* $\mathrm{inv}_\lambda(\rho)$ to be the sum of the entries of ρ that correspond to rows of λ of odd length. The inversion number $\mathrm{inv}(\lambda)$ of a partition continues to have the same meaning as above.

Theorem 4.4 *Given $w \in \widetilde{S}_n$, let (P, Q, ρ) be the triple associated to w by AMBC, with P and Q of shape λ. The sign of w is given by*

$$\mathrm{sgn}(w) \cdot (-1)^{\sum_{i=1}^{n}(w(i)-i)} = (-1)^{\mathrm{inv}(P)+\mathrm{inv}(Q)+\mathrm{inv}(\lambda)+\mathrm{inv}_\lambda(\rho)}. \tag{1}$$

Remark 4.5 In the case that w belongs to the (non-extended) affine Weyl group $\widetilde{S}_n^0$, we have that $\sum_{i=1}^{n}(w(i) - i) = 0$ and that $\mathrm{sgn}(w)$ agrees with the usual notion of sign in a Coxeter group. In this case, Theorem 4.4 becomes

$$\mathrm{sgn}(w) = (-1)^{\mathrm{inv}(P)+\mathrm{inv}(Q)+\mathrm{inv}(\lambda)+\mathrm{inv}_\lambda(\rho)}.$$

In the case that w belongs to the finite symmetric group S_n, we have that ρ is the all-zero vector, so $\mathrm{inv}_\lambda(\rho) = 0$, and we recover Theorem 4.1.

The proof makes use of a variation of the following result from [8]. Denote by $\sim$ the translation equivalence relation.

Lemma 4.6 *For a cell $c = (c_1, c_2)$, let $\overline{D}(c)$ denote the diagonal of c, i.e., $\overline{D}(c) = c_2 - c_1$. If $w \in \widetilde{S}_n$ with $\Phi(w) = (P, Q, \rho)$, P and Q having shape λ, then*

$$\sum_{i=1}^{n}(w(i) - i) = \sum_{b \in \mathcal{B}_w/\sim} \overline{D}(b) = \sum_{i=1}^{\ell(\lambda)} \left(\sum_{c \in \mathfrak{st}_{\rho_i}(P_i, Q_i)/\sim} \overline{D}(c) \right) = n \left(\sum_{i=1}^{\ell(\lambda)} \rho_i \right).$$

Proof The first equality is the definition of $\overline{D}$. The second equality is a repeated use of [8, Lem. 10.3], as described in the paragraph following its proof. The third equality is [8, Lem. 10.4]. □

The same exact arguments can be applied to block diagonals (the reader may wish to recall the relevant definition from Sect. 2.2.3) to get the following result.

Lemma 4.7 *Suppose $w \in \widetilde{S}_n$ and $\Phi(w) = (P, Q, \rho)$, where P and Q have shape λ. Then*

$$D(w) = \sum_{i=1}^{\ell(\lambda)} D(\mathfrak{st}_{\rho_i}(P_i, Q_i)) = \sum_{i=1}^{\ell(\lambda)} \rho_i.$$

That is, the sum of block diagonals of translate classes of balls of a permutation is equal to the sum of block diagonals of translate classes of cells of all the streams involved in the application of AMBC to the permutation.

We now give the proof of the main theorem of this section. The reader may wish to consult Example 4.8 below while reading the proof.

Proof of Theorem 4.4 In order to make an argument by induction, we carefully track how the two sides of (1) change in the application of a single forward step of AMBC.

A monotone numbering of a partial permutation corresponds to a division of the balls into zig-zags. This also induces a division of the inversions of the partial permutation into two classes: those that occur between balls in the same zig-zag and those that occur between balls in different zig-zags. We give the corresponding counts names, as follows: for the partial permutation w with numbering $d : \mathcal{B}_w \to \mathbb{Z}$ and corresponding zig-zags $\{Z_i\}_{i\in\mathbb{Z}}$, the *internal inversion number* is

$$\mathrm{int}(w) := \sum_{Z/\sim} |\{a, b \in Z \cap \mathcal{B}_w : a \text{ is strictly northeast of } b\}|,$$

a sum over translation classes of zig-zags, and the *external inversion number* is

$$\mathrm{ext}(w) := \sum_{(Z,Z')/\sim} \big|\{a \in Z \cap \mathcal{B}_w, b \in Z' \cap \mathcal{B}_w : a \text{ is strictly northeast of } b\}\big|,$$

a sum over translation classes of pairs of zig-zags.

Fix a partial permutation w. The balls of w will be numbered by its southwest channel numbering d. Let $w' = \mathrm{fw}(w)$ be the partial permutation that results from a single forward step of AMBC. Let d' be the numbering of balls of w' induced from this step, i.e., the balls of w' numbered i are precisely the ones contained in Z_i. (This numbering is monotone by Proposition 2.17(a).) We count internal and external inversions of w relative to d, and of w' relative to d'.

First, we consider the internal inversions of w. Choose a zig-zag Z of w, and suppose it includes k balls in $\mathcal{B}_w$. Any pair of these balls forms an inversion. Thus Z contributes $\binom{k}{2}$ internal inversions to w. Obviously Z includes exactly $k-1$ balls in $\mathcal{B}_{w'}$, and so the number of inversions of w' in Z is $\binom{k-1}{2}$. The difference between these two numbers is $k - 1 = |Z \cap \mathcal{B}_w| - 1$. Let λ be the shape of w. The number of translation classes of balls in w is $|\lambda|$ and the number of translation classes of zig-zags is λ_1, so

$$\mathrm{int}(w) = \sum_{Z/\sim} \binom{|Z \cap \mathcal{B}_w|}{2}$$

$$
\begin{aligned}
&= \sum_{Z/\sim} (|Z \cap \mathcal{B}_w| - 1) + \sum_{Z/\sim} \binom{|Z \cap \mathcal{B}_{w'}|}{2} \\
&= |\lambda| - \lambda_1 + \mathrm{int}(w') \\
&= \mathrm{int}(w') + \sum_{i>1} \lambda_i . \qquad (2)
\end{aligned}
$$

Next, we consider external inversions. For a pair of integers $i < j$, we partition the collection of inversions of w between the i-th and j-th zig-zags into four parts. Recall that, by definition, to say that (a, b) is an inversion is to say that a is strictly northeast of b, and recall also the notation $a <_{SW} b$ means that a is strictly southwest of b. Define

$$
\begin{aligned}
E_1^{i,j} &= \{(a,b) \in \mathcal{B}_w^2 : b <_{SW} a, a \in Z_i, b \text{ is the southwest ball of } Z_j\}, \\
E_2^{i,j} &= \{(a,b) \in \mathcal{B}_w^2 : b <_{SW} a, a \in Z_i, b \text{ is not the southwest ball of } Z_j\}, \\
E_3^{i,j} &= \{(b,a) \in \mathcal{B}_w^2 : a <_{SW} b, a \in Z_i, b \text{ is the northeast ball of } Z_j\}, \text{ and} \\
E_4^{i,j} &= \{(b,a) \in \mathcal{B}_w^2 : a <_{SW} b, a \in Z_i, b \text{ is not the northeast ball of } Z_j\}.
\end{aligned}
$$

Note that this really is a partition, in that $\bigcup_{k=1}^4 E_k^{i,j}$ is the collection of all inversions between the i-th and j-th zig-zags and the $E_k^{i,j}$ pairwise do not intersect.

Suppose that (b', a') is an inversion of w' with $a' \in Z_i$, $b' \in Z_j$. Let a, b be the balls of w in the same rows as a', b', respectively, so a is directly west of a' and b directly west of b'. It follows from Proposition 2.17(a) that (b, a) is an inversion in w. Moreover, b is not the northeast ball of Z_j, so $(b, a) \in E_4^{i,j}$. Conversely, it follows from Proposition 2.17 that every inversion $(b, a) \in E_4^{i,j}$ arises from an inversion of w' in this way.

By the same argument, inversions (a', b') of w' with $a' \in Z_i$, $b' \in Z_j$ are in bijection with $E_2^{i,j}$. Thus the number of inversions of w' with one ball in Z_i and the other one in Z_j is $\left|E_4^{i,j}\right| + \left|E_2^{i,j}\right|$. Summing this over all translation classes of pairs (Z_i, Z_j) of zig-zags with $i < j$ gives

$$
\mathrm{ext}(w') = \sum \left(\left|E_4^{i,j}\right| + \left|E_2^{i,j}\right|\right) = \mathrm{ext}(w) - \sum \left(\left|E_1^{i,j}\right| + \left|E_3^{i,j}\right|\right). \qquad (3)
$$

We turn our attention to the summands in this last expression.

If a, b are balls of w and $d(a) < d(b)$ then, by monotonicity of d, a cannot lie southeast of b. Consequently (a, b) is an inversion if and only if a lies to the east of b. When b is the southwest ball of its zig-zag Z_j, a ball a lies to the east of b if and only if it lies to the east of the back-corner post c_j of Z_j. Thus, $\left|E_1^{i,j}\right|$ is equal to the number of balls in $Z_i \cap \mathcal{B}_w$ lying strictly east of c_j. Similarly, when b is the northeast ball of Z_j, (b, a) is an inversion if and only if a lies to the south of b, and so $\left|E_3^{i,j}\right|$ is equal to the number of balls of $Z_i \cap \mathcal{B}_w$ lying strictly south of c_j. Also note that

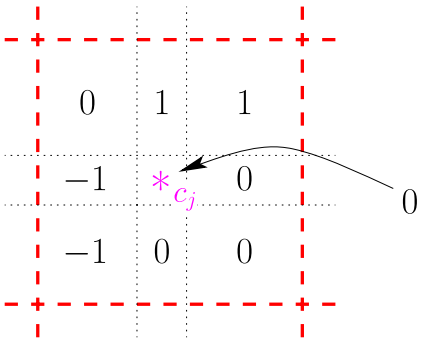

Fig. 12 The dependence of $\mathrm{fix}(b, j)$ on the congruence classes of rows and columns of b relative to c_j

Z_i has no balls which share a row or a column with c_j (since $i \neq j$). Combining the above observations, we have that $\left|E_1^{i,j}\right| + \left|E_3^{i,j}\right|$ is congruent modulo 2 to the number of balls of $Z_i \cap \mathcal{B}_w$ that lie either strictly northeast or strictly southwest of c_j.

Now we sum this expression over all translation classes of pairs (Z_i, Z_j) with $i < j$. We choose as representatives the pairs $i < j$ such that $1 \leqslant j \leqslant m$, where m is the number of translation classes of zig-zags (i.e., the number appearing in Proposition 2.4). By the preceding paragraph, $\sum \left(\left|E_1^{i,j}\right| + \left|E_3^{i,j}\right|\right)$ is congruent modulo 2 to the number of pairs (j, b) such that $1 \leqslant j \leqslant m$, $b \in \mathcal{B}_w$, $d(b) < j$, and b is either strictly northeast or strictly southwest of c_j. By Proposition 2.17(b), if $i \geqslant j$ then Z_i contains no balls either strictly northeast or strictly southwest of c_j. Hence we may drop the condition $d(b) < j$ and, using the notation $\equiv$ for congruence modulo 2, conclude that

$$\sum \left(\left|E_1^{i,j}\right| + \left|E_3^{i,j}\right|\right) \equiv \left|\{(j, b) \in [m] \times \mathcal{B}_w : b <_{SW} c_j \text{ or } b >_{SW} c_j\}\right|.$$

Now fix j and a ball $b = (b_1, b_2)$ of w. Let the coordinates of c_j be given by $c_j = (c_{j,1}, c_{j,2})$. We claim that the number of translates of b that lie either strictly northeast or strictly southwest of c_j is equal to $\left|D(b) - D(c_j) + \mathrm{fix}(b, j)\right|$, where

$$\mathrm{fix}(b, j) := \left(\begin{cases} 1 & \text{if } \overline{b_1} < \overline{c_{j,1}} \\ 0 & \text{if } \overline{b_1} \geqslant \overline{c_{j,1}} \end{cases}\right) - \left(\begin{cases} 1 & \text{if } \overline{b_2} < \overline{c_{j,2}} \\ 0 & \text{if } \overline{b_2} \geqslant \overline{c_{j,2}} \end{cases}\right) \tag{4}$$

and the inequalities between equivalence classes are in the broken order. (This is illustrated in Fig. 12: the coordinates of b and c_j are reduced modulo n so that they lie in the same $n \times n$ square, and then the value of $\mathrm{fix}(b, j)$ is determined by the relative positions of the two.) Since each row and column congruence class of b can be either smaller, equal, or greater than the corresponding class for c_j, in principle there are nine cases of this claim to consider. The cases are very similar, so we only discuss two of them in detail.

First, consider the case when $\overline{b_1} > \overline{c_{j,1}}$ and $\overline{b_2} > \overline{c_{j,2}}$ in broken order. Without loss of generality, assume that $D(b) \geqslant D(c_j)$, so that translates of b can be northeast but not southwest of c_j (see Fig. 13 for examples with $D(b) - D(c_j)$ being 1, 2, and

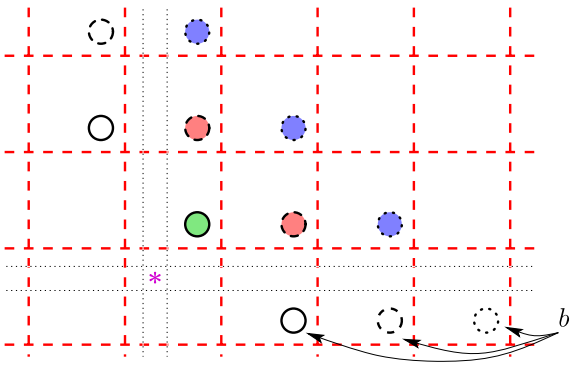

Fig. 13 Counting translates of b lying northeast of c_j when $D(b) - D(c_j) = 1$ (solid), $D(b) - D(c_j) = 2$ (dashed), and $D(b) - D(c_j) = 3$ (dotted)

3). The translate $b + k(n, n)$ is northeast of c_j if and only if $c_{j,1} > b_1 + kn$ and $c_{j,2} < b_2 + kn$ (where now we are comparing integers, not residue classes). Thus, the number of translates of b northeast of c_j is equal to the number of integer solutions k to the inequalities

$$\frac{c_{j,1} - b_1}{n} > k > \frac{c_{j,2} - b_2}{n}.$$

Since $\overline{b_1} > \overline{c_{j,1}}$, the largest integer smaller than $(c_{j,1} - b_1)/n$ is $\lceil c_{j,1}/n \rceil - \lceil b_1/n \rceil - 1$, while since $\overline{b_2} > \overline{c_{j,2}}$ the smallest integer larger than $(c_{j,2} - b_2)/n$ is $\lceil c_{j,2}/n \rceil - \lceil b_2/n \rceil$; thus, the displayed inequalities are equivalent to

$$\left\lceil \frac{c_{j,1}}{n} \right\rceil - \left\lceil \frac{b_1}{n} \right\rceil - 1 \geqslant k \geqslant \left\lceil \frac{c_{j,2}}{n} \right\rceil - \left\lceil \frac{b_2}{n} \right\rceil.$$

The number of solutions k is

$$\begin{aligned} &\left(\left\lceil \frac{c_{j,1}}{n} \right\rceil - \left\lceil \frac{b_1}{n} \right\rceil - 1 \right) - \left(\left\lceil \frac{c_{j,2}}{n} \right\rceil - \left\lceil \frac{b_2}{n} \right\rceil \right) + 1 = D(b) - D(c_j) \\ &= \left| D(b) - D(c_j) + 0 \right|. \end{aligned}$$

Finally, in this case $\mathrm{fix}(b, c_j) = 0$, as needed.

Second, consider the case that $\overline{b_1} < \overline{c_{j,1}}$ in broken order and $\overline{b_2} = \overline{c_{j,2}}$. Since $\overline{b_2} = \overline{c_{j,2}}$, the ball b is a translate of the ball in the same column as c_j. By the definition of AMBC, this ball lies directly south of c_j, so necessarily $D(b) < D(c_j)$ and we seek translates of b that are strictly southwest of c_j. By the same analysis as in the previous case, these translates correspond to integer solutions of

$$\frac{c_{j,1} - b_1}{n} < k < \frac{c_{j,2} - b_2}{n}.$$

By the hypotheses on the relative values of the b_i and $c_{j,i}$, this is equivalent to

$$\left\lceil \frac{c_{j,1}}{n} \right\rceil - \left\lceil \frac{b_1}{n} \right\rceil + 1 \leqslant k \leqslant \left\lceil \frac{c_{j,2}}{n} \right\rceil - \left\lceil \frac{b_2}{n} \right\rceil - 1.$$

Thus in this case the number of solutions is $D(c_j) - D(b) - 1 = \left|D(b) - D(c_j) + 1\right|$. Since $\mathrm{fix}(b, c_j) = 1$, this gives the desired result.

The other seven cases are extremely similar to these two.

We now plug the claim in to (3). Since $|x| \equiv x \equiv -x$, we can drop absolute values and extra negative signs to get

$$\begin{aligned}
&\mathrm{ext}(w) - \mathrm{ext}(\mathrm{fw}(w)) \\
&\quad \equiv \sum_{b \in (\mathcal{B}_w/\sim)} \sum_{j=1}^{m} \left|D(b) - D(c_j) + \mathrm{fix}(b, j)\right| \\
&\quad = \lambda_1 \cdot \sum_{b \in (\mathcal{B}_w/\sim)} D(b) + |\lambda| \cdot \sum_{j=1}^{m} D(c_j) + \sum_{b \in (\mathcal{B}_w/\sim)} \left(\sum_{j=1}^{m} \mathrm{fix}(b, j) \right).
\end{aligned}$$

By Lemma 4.7, the first sum is equal to $\sum_i \rho_i$. The second sum is the definition of the altitude of $\mathtt{st}(w)$, and so is equal to ρ_1. For the third sum, we break it into two pieces as in (4), one involving the row indices and one involving the column indices. The sum involving the row indices is equal to the number of pairs of a ball b and a stream cell c_j such that $\overline{b_1} < \overline{c_{j,1}}$. Note that the numbers $\overline{c_{j,1}}$ constitute the first row of Q, while the other rows of balls will necessarily appear in the lower rows of Q. Thus the desired number of pairs is precisely equal to $\mathrm{inv}_1(Q)$, where $\mathrm{inv}_i(T)$ is the number of inversions of a tabloid T such that the top row involved is row i. The same analysis applies to the other piece, and so the last sum in the previous displayed equation simplifies to $\mathrm{inv}_1(Q) - \mathrm{inv}_1(P)$. Therefore

$$\mathrm{ext}(w) - \mathrm{ext}(\mathrm{fw}(w)) \equiv |\lambda|\, \rho_1 + \lambda_1 \sum_{i \geqslant 1} \rho_i + \mathrm{inv}_1(Q) + \mathrm{inv}_1(P).$$

Combining this with (2) yields

$$\ell(w) - \ell(\mathrm{fw}(w)) \equiv \sum_{i>1} \lambda_i + \lambda_1 \sum_{i \geqslant 1} \rho_i + |\lambda|\, \rho_1 + \mathrm{inv}_1(Q) + \mathrm{inv}_1(P). \tag{5}$$

Now suppose $w \in \widetilde{S}_n$. The congruence (5) holds at each step of AMBC, with appropriate adjustments to the indexing (i.e., the first row of $P(\mathrm{fw}(w))$ should be numbered 2, not 1); after the last step of AMBC, we are left with the empty permutation, which has 0 inversions. Thus the sum of the left side of (5) over all the steps gives

$\ell(w)$. By Lemma 4.6, we have $n\sum_{i\geqslant 1}\rho_i = \sum_{i=1}^{n}(w(i)-i)$, and so

$$\begin{aligned}
\ell(w) &\equiv \sum_{j=1}^{\ell(\lambda)}\left(\sum_{i>j}\lambda_i + \lambda_j\sum_{i\geqslant j}\rho_i + \rho_j\sum_{i\geqslant j}\lambda_i + \mathrm{inv}_j(Q) + \mathrm{inv}_j(P)\right)\\
&= \sum_{i=1}^{\ell(\lambda)}(i-1)\lambda_i + \sum_{j=1}^{\ell(\lambda)}\left(\lambda_j\sum_{i\geqslant j}\rho_i + \rho_j\sum_{i\geqslant j}\lambda_i\right) + \mathrm{inv}(Q) + \mathrm{inv}(P)\\
&\equiv \mathrm{inv}(\lambda) + \sum_{j=1}^{\ell(\lambda)}\left(\lambda_j\rho_j + \lambda_j\sum_{i\geqslant 1}\rho_i\right) + \mathrm{inv}(Q) + \mathrm{inv}(P)\\
&= \mathrm{inv}(\lambda) + \sum_{j=1}^{\ell(\lambda)}\lambda_j\rho_j + n\sum_{i\geqslant 1}\rho_i + \mathrm{inv}(Q) + \mathrm{inv}(P)\\
&\equiv \mathrm{inv}(\lambda) + \mathrm{inv}_\lambda(\rho) + \sum_{i=1}^{n}(w(i)-i) + \mathrm{inv}(Q) + \mathrm{inv}(P).
\end{aligned}$$

Raising -1 to the power of each side yields the desired result. □

Example 4.8 The permutation $w = [7, 2, 4, 1]$ in $\widetilde{S}_4$ is shown in Fig. 14. For readability, we label inversions using only the column index, so that the seven inversions of w are $(7, 2)$, $(7, 4)$, $(7, 1)$, $(7, 6)$, $(7, 5)$, $(2, 1)$, and $(4, 1)$. Under AMBC, w corresponds to the triple

$$P = \begin{array}{|c|c|}\hline \bar{1} & \bar{4} \\ \hline \bar{2} \\ \cline{1-1} \bar{3} \\ \cline{1-1}\end{array}\,, \qquad Q = \begin{array}{|c|c|}\hline \bar{1} & \bar{3} \\ \hline \bar{2} \\ \cline{1-1} \bar{4} \\ \cline{1-1}\end{array}\,, \qquad \rho = \begin{pmatrix}0\\0\\1\end{pmatrix},$$

having $\mathrm{inv}(P) = 2$, $\mathrm{inv}(Q) = 1$, $\mathrm{inv}(\lambda) = 1$, and $\mathrm{inv}_\lambda(\rho) = 1$. Thus in this case, Theorem 4.4 asserts that $(-1)^7\cdot(-1)^4 = (-1)^5$. Of the seven inversions, $(7, 2)$, $(7, 1)$, and $(2, 1)$ are internal, while $w' = \mathrm{fw}(w)$ has a single internal inversion $(7, 2)$; the difference $3 - 1$ is equal to $|\lambda| - \lambda_1 = 2$. The four external inversions of w are partitioned as follows: $E_1^{1,2} = \{(7, 4)\}$, $E_2^{1,3} = \{(7, 6)\}$, $E_1^{1,3} = \{(7, 5)\}$, and $E_3^{1,2} = \{(4, 1)\}$. The inversion $(7, 6)$ corresponds to the unique external inversion $(7, 6)$ in w'. Finally, we note that $\mathrm{fix}(b, j) = 0$ except for the case $b = (4, 1)$ and $j = 2$ (with $c_2 = (3, 4)$).

Remark 4.9 Reifegerste proves Theorem 4.1 by induction, using a result of Beissinger [4] to show that it holds for an element in each Knuth equivalence class and then proving that its validity is preserved under Knuth moves. The second half of this argument (the inductive step) is straightforward using Theorem 3.11; however, it is not clear what Knuth class representatives could play the role of Beissinger and Reifegerste's involutions for the base case.

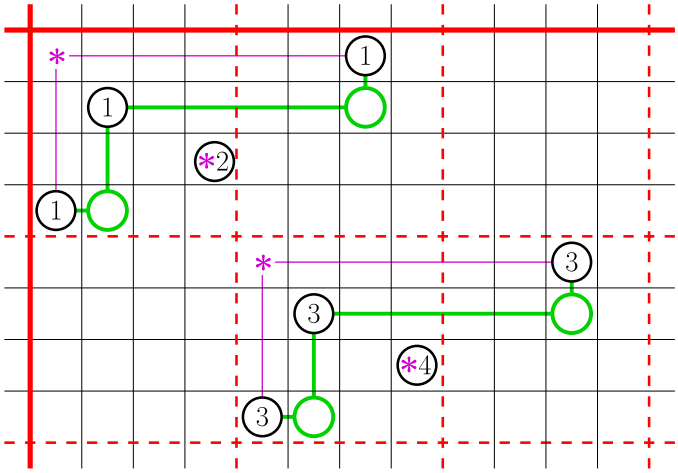

Fig. 14 One forward step of AMBC, starting with the permutation [7, 2, 4, 1]

5 Charge

The purpose of this section is to introduce a statistic on tabloids we call *local charge*. Charge (originally defined in [15]; see also [16] for an exposition) is a classical statistic for tableaux, which generalizes naturally to tabloids. What is most useful to us is a "local" version that depends on a pair of adjacent rows, rather than the entire tabloid, and in this section we only define it in this context. This statistic arises in two ways in our theory. First, the instance we deal with in this section, is that it plays a crucial role in understanding the offset constants the definition of dominance (see Proposition 3.1 and Remark 3.2). This is captured in Theorem 5.10, the main result of this section. Second, as will be described in Sect. 8, it appears in the description of the connected components of the KL DEG $\mathcal{A}_\lambda$.

5.1 Definitions and basic properties

Definition 5.1 Suppose T is a tabloid of shape $\langle m, m\rangle$. An *activation ordering* for T is a bijection $o : [m] \to T_1 \subsetneq [\overline{n}]$. We think of activation orderings as linear orderings of the entries of the top row of T. The *standard activation ordering* is the activation ordering in which the entries of the top row are arranged to increase in the broken order (see Definition 4.2).

Whenever an activation ordering is needed but not specified, we assume that the standard activation ordering is used.

Definition 5.2 Suppose T is a tabloid whose k-th and $(k+1)$-st rows have the same size m, and o is an activation ordering for $T_{k,k+1}$ (the tabloid consisting of the k-th and $(k+1)$-st rows of T). Then the *charge matching* between rows k and $k+1$ of T with ordering o is the matching of entries in row k with entries in row $k+1$ defined

as follows. Suppose

$$T_{k,k+1} = \begin{array}{|c|c|c|c|}\hline \overline{a_1} & \overline{a_2} & \cdots & \overline{a_m} \\ \hline \overline{b_1} & \overline{b_2} & \cdots & \overline{b_m} \\ \hline \end{array}, \tag{6}$$

where $\overline{a_i} = o(i)$ for all i. For $i = 1, 2, \ldots, m$, match $\overline{a_i}$ with the smallest (in broken order) unmatched $\overline{b_j}$ such that $\overline{b_j} > \overline{a_i}$, if such a $\overline{b_j}$ exists; otherwise, match $\overline{a_i}$ to the smallest unmatched $\overline{b_j}$.

Definition 5.3 Suppose T is a tabloid whose k-th and $(k+1)$-st rows have the same size m, and o is an activation ordering of $T_{k,k+1}$. The *local charge in row k* of T, denoted $\mathrm{lch}_k^o(T)$, is the number of elements $\overline{a}$ of the k-th row of T such that $\overline{a}$ is matched to $\overline{b}$ and $\overline{a} > \overline{b}$. We may say that such an entry $\overline{a}$ (or the corresponding $\overline{b}$, or the pair $(\overline{a}, \overline{b})$) *contributes* to the charge. If o is the standard ordering, we omit o from the notation and write $\mathrm{lch}_k(T)$.

Example 5.4 Consider the tabloid

$$T = \begin{array}{|c|c|c|c|}\hline \overline{2} & \overline{4} & \overline{6} & \overline{10} \\ \hline \overline{3} & \overline{7} & \overline{8} \\ \cline{1-3} \overline{1} & \overline{5} & \overline{9} \\ \cline{1-3} \end{array}.$$

The local charge $\mathrm{lch}_1(T)$ is not defined since the first and second rows of T have different lengths. With the standard activation ordering, the charge matching between the second and third rows is $\overline{3} \leftrightarrow \overline{5}$, $\overline{7} \leftrightarrow \overline{9}$, $\overline{8} \leftrightarrow \overline{1}$. The only pair that contributes to charge is $(\overline{8}, \overline{1})$, so $\mathrm{lch}_2(T) = 1$.

Suppose now we use a non-standard activation ordering between the second and third rows, say, $o(1) = \overline{8}$, $o(2) = \overline{7}$, $o(3) = \overline{3}$. Then the charge matching is $\overline{8} \leftrightarrow \overline{9}$, $\overline{7} \leftrightarrow \overline{1}$, $\overline{3} \leftrightarrow \overline{5}$. For this o, the only pair that contributes to the charge is $(\overline{7}, \overline{1})$, and so $\mathrm{lch}_2^o(T) = 1$.

We now show that the coincidence in the example above is not an accident: local charge does not depend on the choice of activation ordering.

Lemma 5.5 *Suppose T is a tabloid whose k-th and $(k+1)$-st rows have the same size m. Then $\mathrm{lch}_k^o(T) = \mathrm{lch}_k(T)$ for every activation ordering o.*

Proof Suppose $T_{k,k+1}$ is as in (6) with $o(i) = \overline{a_i}$ for $1 \leqslant i \leqslant m$. Since any permutation is a product of adjacent transpositions, it is sufficient to prove that $\mathrm{lch}_k^o(T) = \mathrm{lch}_k^{o'}(T)$, where o' is the ordering defined by

$$\begin{array}{ll} o'(j) = \overline{a_{j+1}}, & \\ o'(j+1) = \overline{a_j}, & \text{and} \\ o'(i) = \overline{a_i} & \text{if } i \neq j \text{ or } j+1. \end{array}$$

Furthermore, with this choice, the entries $\overline{a_1}, \ldots, \overline{a_{j-1}}$ match to the same elements of the second row in the charge matchings with orderings o, o'. Thus, we may as well remove these $j-1$ entries from both the first and second rows, and consider the case

$j = 1$. As a final simplification, we may assume that $\overline{a_1} < \overline{a_2}$: if not, switch the roles of o and o'.

Let $\overline{b_{i_1}}$ and $\overline{b_{i_2}}$ be the elements of the second row with which $\overline{a_1}$ and $\overline{a_2}$ match in the charge matching with ordering o. We seek to show that in the charge matching with ordering o', $\overline{a_1}$ and $\overline{a_2}$ still match (in some order) to $\overline{b_{i_1}}$ and $\overline{b_{i_2}}$ in a way that preserves the number of contributing pairs.

A priori, there are 12 possible orders for the four values $\overline{a_1}, \overline{a_2}, \overline{b_{i_1}}, \overline{b_{i_2}}$ in broken order; however, six of these are incompatible with the fact that $\overline{a_1}$ matches to $\overline{b_{i_1}}$ in the matching with ordering o. This leaves six cases to check; we present the two most complicated here, and leave the other four to the reader.

Suppose $\overline{a_1} < \overline{b_{i_1}} < \overline{b_{i_2}} < \overline{a_2}$. Since $\overline{a_2}$ matches to $\overline{b_{i_2}}$ in the matching with ordering o, there are no values in T_{k+1} larger than $\overline{a_2}$, and the only value in T_{k+1} smaller than $\overline{b_{i_2}}$ is $\overline{b_{i_1}}$. It follows that in the matching with ordering o', $\overline{a_2}$ matches to $\overline{b_{i_1}}$ and $\overline{a_1}$ matches to $\overline{b_{i_2}}$. Then $\overline{a_1}$ does not contribute to charge in either matching and $\overline{a_2}$ contributes to the charge in both matchings.

Suppose $\overline{b_{i_2}} < \overline{a_1} < \overline{a_2} < \overline{b_{i_1}}$. Since $\overline{a_1}$ matches to $\overline{b_{i_1}}$ in the matching with ordering o, $\overline{b_{i_1}}$ is the smallest value in T_{k+1} larger than $\overline{a_1}$, and hence is the smallest value in T_{k+1} larger than $\overline{a_2}$. Moreover, since $\overline{a_2}$ matches to $\overline{b_{i_2}}$, there can be no other values in T_{k+1} larger than $\overline{a_2}$, and therefore none larger than $\overline{a_1}$; and $\overline{b_{i_2}}$ is the smallest value in T_{k+1}. Thus in the charge matching with ordering o', $\overline{a_2}$ matches to $\overline{b_{i_1}}$ and $\overline{a_1}$ matches to $\overline{b_{i_2}}$. Thus $\overline{b_{i_2}}$ contributes to the charge in both matchings, while $\overline{b_{i_1}}$ does not contribute to the charge in either.

Since the same two elements in the second row are paired to $\{a_1, a_2\}$ in the matchings corresponding to o and o', it follows that all pairs involving $\overline{a_i}$ for $i > 2$ are also the same in both matchings. Thus the two local charges are equal, as claimed. □

5.2 A formula for the offset constants

Recall the discussion of the offset constants in Sect. 2.2. The goal of this section is to give a formula for these constants in terms of local charge. Before we can do that, we need to precisely define the constants. We use the following convention: whenever we have a stream S and a proper numbering of its cells, we let $S^{(i)}$ denote the cell numbered i.

Proposition-Definition 5.6 *Suppose S and T are streams of the same density m such that no cell of S shares a row or a column with a cell of T, and S is properly numbered. Then there is a unique proper numbering of T with the following properties:*

(1) *for all i, $S^{(i)}$ is northwest of $T^{(i)}$, and*
(2) *for some j, $S^{(j+1)}$ is not northwest of $T^{(j)}$.*

This numbering is called the backward numbering *of T with respect to S.*

Proof A different definition of the backward numbering is given in [8]; the proof proceeds by showing they are equivalent. The proof may be skipped without harming understanding of the rest of this section.

Let d be the backward numbering of T constructed in [8, Sect. 4.1]. By construction, d is monotone and it satisfies point (1). As described in [8, Rem. 13.3], it also satisfies

point (2). Thus such numberings exist. Let d' be any numbering satisfying (1) and (2). By [8, Rem. 13.1], $d'(x) \leqslant d(x)$ for every $x \in T$. Since d' satisfies (2), it must coincide with d on at least one translation class of cells. However, [8, Prop. 13.2] states that d is, in fact, proper (i.e., T is numbered by consecutive integers). Thus any monotone numbering that coincides with it on a translate class must coincide with it everywhere, so $d' = d$. □

Condition (2) in the definition of the backward numbering may be met in two different ways: $T^{(j)}$ might be either north or west of $S^{(j+1)}$. The situation when both occur simultaneously is special.

Definition 5.7 [8, Def. 5.4] Suppose S and T are streams of the same density m, and no cell of S shares a row or a column with a cell of T. Number S with a proper numbering and number T with the backward numbering with respect to S. Then T is said to be *concurrent to* S if there exist i and j (possibly equal) such that $T^{(i)}$ is north of $S^{(i+1)}$ and $T^{(j)}$ is west of $S^{(j+1)}$.

Proposition 5.8 [8, Prop. 5.6] *Suppose A, B, A', B' are equinumerous subsets of $[\overline{n}]$ with $A \cap A' = B \cap B' = \varnothing$. Then there exists a unique integer r such that $\mathfrak{st}_r(A', B')$ is concurrent to $\mathfrak{st}_0(A, B)$.*

Definition 5.9 [8, Def. 5.8] Suppose P and Q are tabloids of the same shape λ, and that $\lambda_i = \lambda_{i+1}$. The *dominance constant* $r_{i+1}(P, Q)$ is the unique integer r such that $\mathfrak{st}_r(P_{i+1}, Q_{i+1})$ is concurrent to $\mathfrak{st}_0(P_i, Q_i)$. The weight ρ is *dominant* for P, Q if and only if

$$\rho_{i+1} \geqslant \rho_i + r_{i+1}(P, Q)$$

for every i such that $\lambda_i = \lambda_{i+1}$.

Theorem 5.10 *Suppose that P and Q are two tabloids of shape λ and $\lambda_i = \lambda_{i+1}$. Then the dominance constants are given by*

$$r_{i+1}(P, Q) = \mathrm{lch}_i(P) - \mathrm{lch}_i(Q).$$

Before proceeding with the proof we need a technical lemma.

Lemma 5.11 *Fix $m \in \mathbb{Z}^{>0}$. Suppose we have two sets of integers $A = \{a_1, a_2, \ldots, a_m\}$ and $B = \{b_1, b_2, \ldots, b_m\}$ with $1 \leqslant a_1 < \cdots < a_m$ and $1 \leqslant b_1 < \cdots < b_m$. Define $\ell(A, B)$ to be the maximal integer such that there exist indices $1 \leqslant i_1 < i_2 < \cdots < i_{\ell(A,B)} \leqslant m$ satisfying*

$$a_1 < b_{i_1}, \quad a_2 < b_{i_2}, \quad \ldots, \quad a_{\ell(A,B)} < b_{i_{\ell(A,B)}}.$$

Choose $b' > \max(a_m, b_m)$, and let $B' = \{b_2, \ldots, b_m, b'\}$. If $\ell(A, B) < m$ then $\ell(A, B') = \ell(A, B) + 1$.

Proof The proof is illustrated in Fig. 15. Let k be the largest integer such that $i_k = k$ (and if $i_1 > 1$ then take $k = 0$). For $1 \leqslant j \leqslant k$, define $b'_{i_j} = b_{i_j+1} (= b_{j+1})$. For $k+1 \leqslant j \leqslant \ell(A,B)$, define $b'_{i_j} = b_{i_j}$. Define $b'_{i_{\ell(A,B)+1}} = b'$. Since $i_{k+1} > k+1 = i_k + 1$, the b'_{i_j} are distinct. For $1 \leqslant j \leqslant \ell(A,B)$, we have $a_j < b_{i_j} \leqslant b'_{i_j}$, and also $a_{\ell(A,B)+1} \leqslant a_m < b'_{i_{\ell(A,B)+1}}$. Thus $\ell(A,B') \geqslant \ell(A,B)+1$.

Conversely, if for some integer k there are inequalities

$$a_1 < b'_{i_1}, \quad a_2 < b'_{i_2}, \quad \ldots, \quad a_k < b'_{i_k}$$

such that the right sides are distinct elements of B', then removing the inequality involving $b'_m = b'$, if it exists, leaves $k-1$ inequalities between elements of A and elements of B, so $\ell(A,B') \leqslant \ell(A,B)+1$. This completes the proof. □

We are now prepared for the proof of the main result of this section.

Proof of Theorem 5.10 Both sides of the desired equality depend only on a pair of adjacent rows of the same size. Hence, in the remainder of the proof we assume that P and Q are tabloids of shape $\langle m, m\rangle$ with entries in $[\overline{n}]$ for some $m \leqslant n/2$.

Consider $S := \mathfrak{st}_0(P_1, Q_1)$. Number S properly so that the north-most ball south-east of the origin is numbered 1, and therefore $S^{(1)}, \ldots, S^{(m)}$ lie in $[n] \times [n]$. Let $r = \mathrm{lch}_1 P - \mathrm{lch}_1 Q$ and $T = \mathfrak{st}_r(P_2, Q_2)$. We must show that T is concurrent to S. First, we give a description of the backward numbering of T.

Denote $P_1 = \{\overline{a_1}, \ldots, \overline{a_m}\}$, $P_2 = \{\overline{b_1}, \ldots, \overline{b_m}\}$, $Q_1 = \{\overline{a'_1}, \ldots, \overline{a'_m}\}$, and $Q_2 = \{\overline{b'_1}, \ldots, \overline{b'_m}\}$, with representatives chosen so that $1 \leqslant a_1 < a_2 < \cdots < a_m \leqslant n$, and similarly for the other three rows. Extend the indices to all of $\mathbb{Z}$ via $a_{i+m} = a_i + n$, so that $\bigcup P_1 = \{\cdots < a_{-1} < a_0 < a_1 < \cdots\}$ and $a_0 \leqslant 0 < a_1$, and similarly for the other rows. With this choice, we have $S^{(i)} = (a'_i, a_i)$ for all i. Define a numbering of T as follows: the ball numbered i is

$$\left(b'_{i+\mathrm{lch}_1(Q)}, b_{i+\mathrm{lch}_1(P)}\right) =: T^{(i)}.$$

This is clearly a proper numbering of T; we claim that this is the backward numbering of T with respect to S. To prove this, we must show that $S^{(i)}$ is northwest of $T^{(i)}$ for every i, and that $S^{(j+1)}$ is not northwest of $T^{(j)}$ for some j.

Define the local charge on a pair of equinumerous sets of integers (rather than equivalence classes modulo n) in the obvious way, so that $\mathrm{lch}(\{a_1, a_2, \ldots, a_m\}, \{b_1, b_2, \ldots,$

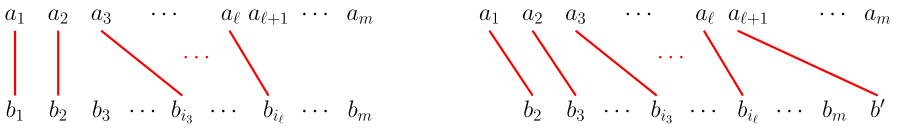

Fig. 15 A figure illustrating the proof of Lemma 5.11. The red lines represent inequalities between an element a_j of A and a larger element B (left) or B' (right). In this example, $i_1 = 1$ and $i_2 = 2$, while $i_3 > 3$

$b_m\}) = \mathrm{lch}_1(P)$. We apply Lemma 5.11 repeatedly, each time replacing the smallest entry b_i in the second row with $b_{i+m} = b_i + n$; this yields

$$\mathrm{lch}(\{a_1, a_2, \ldots, a_m\}, \{b_{1+\mathrm{lch}_1(P)}, \ldots, b_{m+\mathrm{lch}_1(P)}\}) = 0.$$

Therefore $a_i < b_{i+\mathrm{lch}_1(P)}$ for all i, and so $S^{(i)}$ is west of $T^{(i)}$ for all i. Applying the same argument to Q (*mutatis mutandis*) yields that $S^{(i)}$ is north of $T^{(i)}$ for all i.

For the second half of the claim, we have two cases: first, if the original local charge is positive, so in the previous paragraph Lemma 5.11 is applied a positive number of times, then applying the lemma one fewer time produces

$$\mathrm{lch}(\{a_1, \ldots, a_m\}, \{b_{\mathrm{lch}_1(P)}, \ldots, b_{m-1+\mathrm{lch}_1(P)}\}) = 1.$$

Therefore there is some j such that $b_{j-1+\mathrm{lch}_1(P)} < a_j$, and so $S^{(j)}$ is not west (and *a fortiori* not northwest) of $T^{(j-1)}$. Second, if the original local charge is 0, we need a variant of Lemma 5.11 in which the *largest* element in the bottom row is replaced by a new *minimal* element; the proof is no harder, we apply the variant lemma once, and the conclusion is the same. This completes the proof that the numbering described is the backward numbering of T with respect to S.

Moreover, applying the same argument as in the previous paragraph to Q instead of P yields that there is some j such that $S^{(j)}$ is not north of $T^{(j-1)}$. Together, these two paragraphs establish that T is concurrent to S. □

5.3 Resolving a conundrum

As described in Sect. 2.2, the composite map $\Phi \circ \Psi$ is a surjection from Ω to Ω_{dom}, and the image of (P, Q, ρ) is a triple (P, Q, ρ') where ρ' is the *dominant representative* of ρ. In this section, we use Theorem 5.10 to resolve a mystery mentioned in Remark 3.2: why does applying the operation $(P, Q, \rho) \mapsto (Q, P, (-\rho)')$ twice return the original triple?

In this context, it is convenient to define the *symmetrized offset constants*

$$s_i(P, Q) := \sum_{j=i'+1}^{i} r_j(P, Q),$$

where i' is the smallest integer such that $\lambda_{i'} = \lambda_i$. We may reformulate Definition 5.9 in terms of these constants: ρ is dominant if for each $1 \leqslant i < \ell(\lambda)$ one of the following holds:

- $\lambda_i > \lambda_{i+1}$, or
- $\lambda_i = \lambda_{i+1}$ and $(\rho - \mathbf{s})_{i+1} \geqslant (\rho - \mathbf{s})_i$,

where $\mathbf{s} = (s_i(P, Q))_{1 \leqslant i \leqslant \ell(\lambda)}$. Suppose that λ has k distinct part sizes, occurring with multiplicities $m_1, \ldots, m_k$; let $m'_i := m_1 + \cdots + m_i$. According to [8, Sect. 6], the dominant representative of ρ can be computed as follows:

- Let $\mathbf{x} = \rho - \mathbf{s}$.

- For each $1 \leqslant j \leqslant k$, let $(y_{m'_{j-1}+1}, y_{m'_{j-1}+2}, \ldots, y_{m'_j})$ be the increasing rearrangement of $(x_{m'_{j-1}+1}, x_{m'_{j-1}+2}, \ldots, x_{m'_j})$.
- Let $\rho' = \mathbf{y} + \mathbf{s}$.

We now resolve the mystery: by Theorem 5.10, $r_i(Q, P) = -r_i(P, Q)$ for all i, and so also $s_k(Q, P) = -s_k(P, Q)$. If $\lambda_i = \lambda_{i+1} = \cdots = \lambda_j$ is a maximal string of equal part sizes then for any $i \leqslant k \leqslant j$, there exists $i \leqslant \ell \leqslant j$ such that $(-\rho)'_k = -\rho_\ell - s_\ell(P, Q) + s_k(P, Q)$. Therefore $(-(-\rho)')_k = \rho_\ell + s_\ell(P, Q) - s_k(P, Q)$. To take the dominant representative of this vector, we subtract $s_k(Q, P) = -s_k(P, Q)$, giving a vector whose k-th entry is $\rho_\ell + s_\ell(P, Q)$; since ρ is dominant, this is the ℓ-th largest of these numbers, so after rearranging we add $s_\ell(Q, P) = -s_\ell(P, Q)$, returning the value ρ_ℓ, as desired.

6 Symmetries

The definition of the extended affine symmetric group (as in Section 2.1) gives no special standing to the interval $[n]$ over any other interval of n consecutive integers. This indifference can be seen in the expression $\widetilde{S}_n = \widetilde{S}_n^0 \rtimes \mathbb{Z}$ for $\widetilde{S}_n$ as a semidirect product of the non-extended affine Weyl group $\widetilde{S}_n^0$ and the infinite cyclic subgroup generated by the "shift" permutation $s = [2, 3, \ldots, n+1]$. It may also be seen in the Dynkin diagram of the group $\widetilde{S}_n^0$, which is a cycle on n vertices, each vertex representing a simple transposition: conjugation of $\widetilde{S}_n^0$ by s is equivalent to rotating the Dynkin diagram by one position.

The broken order of Sects. 4 and 5 breaks this symmetry by giving special standing to the link from $\overline{n}$ to $\overline{1}$. This corresponds to other broken symmetries in our setting: the definition of altitude of a stream (and so everything that follows from it) gives special standing to the square $[n] \times [n]$.

The shift permutation s has a natural action on tabloids, sending $\overline{i}$ to $\overline{i+1}$ for all i.

Proposition 6.1 *Fix a partition λ. The action of s is a graph automorphism of the KL DEG $\mathcal{A}_\lambda$.*

Proof It is immediate from the definitions that s commutes with Knuth moves. □

7 Monodromy

In this section and the next, we study the Kazhdan–Lusztig molecules of $\widetilde{S}_n$: given a permutation w corresponding to a triple (P, Q, ρ) under AMBC, we specify exactly which triples (P', Q', ρ') correspond to a permutation w' in the same Knuth equivalence class as w. By Theorem 3.11, it is necessary that $P' = P$. In this section, we ask which values of ρ' are possible if also $Q' = Q$; in the next section, we consider which tabloids Q' are possible. Together, the answers to these questions provide a complete classification.

Our approach is to use Lemma 3.22: if there is a sequence $(w = w_1, w_2, \ldots, w_k = w')$ of permutations in which each consecutive pair is connected by a Knuth move

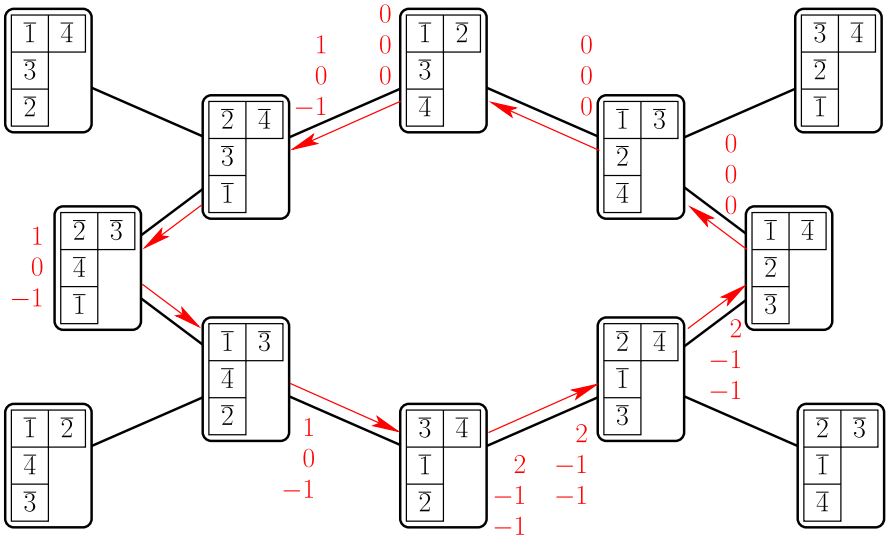

Fig. 16 A loop in $\mathcal{A}_{\langle 2,1,1\rangle}$ and the corresponding weight change vector

and $Q(w) = Q(w')$, we may project this sequence via $w_i \mapsto Q(w_i) =: T_i$ to a sequence $L = (T_1, \ldots, T_k)$ of tabloids in the KL DEG $\mathcal{A}_\lambda$ for the associated shape λ. By Lemma 3.22, consecutive elements of L are connected by Knuth moves, and by hypothesis $T_1 = T_k$, so this sequence is a loop (i.e., a closed walk on the graph). Moreover, by Theorem 3.11, the *weight change vector* $\rho(w') - \rho(w)$ depends only on the loop L. This suggests the following definition.

Definition 7.1 Given a permutation $w \in \widetilde{S_n}$, the *monodromy group* G_w based at w is the set

$$G_w := \{\rho(w_k) - \rho(w)\}_{(w=w_1,\ldots,w_k)},$$

where the indexing tuple $(w_1, \ldots, w_k)$ runs over all sequences of permutations such that $Q(w_k) = Q(w)$ and w_{i+1} is connected by a Knuth move to w_i for all i.

Implicit in this definition is the fact that G_w is an (abelian) group under vector addition. This is easy to see using Theorem 3.11: the concatenation of two loops in a KL DEG has a lift whose monodromy group element is the sum of the lifts of the two individual loops.

Example 7.2 Figure 16 shows a loop in the KL DEG of shape $\langle 2, 1, 1\rangle$, beginning and ending at the column superstandard tabloid with start at 1. The red vectors give the weight change at each step of the loop. The monodromy element of any lift of the loop is $(2, -1, -1)$.

Remark 7.3 If $Q(w)$ and $Q(w')$ are connected in the KL DEG, then it follows by the standard topological argument (again, using Theorem 3.11) that $G_w = G_{w'}$ as subsets

of $\mathbb{Z}^{\ell(\lambda)}$. Thus monodromy is base-point independent on connected components of KL DEGs. (In fact, one consequence of Theorem 7.28, the main result of this section, is that G_w only depends on the shape of $Q(w)$.)

The next lemma shows it suffices to understand G_w when $Q(w)$ is column superstandard.

Lemma 7.4 *Every connected component of a KL DEG contains a reverse row superstandard tabloid and a column superstandard tabloid.*

Proof The Shi combing procedure (originally from [19], and explained in our language in [8, Sect. 9]) produces a sequence of Knuth moves from any permutation to a permutation whose right descent set is a singleton. By Theorem 3.11 and Proposition 3.6, this can be used to produce a sequence of Knuth moves from any tabloid to a tabloid whose τ-invariant is a singleton. The tabloids with a unique descent are exactly the reverse row superstandard tabloids. This completes the first claim.

By the first claim, each column superstandard tabloid lies in the same connected component as some reverse row superstandard tabloid. The automorphism s of Proposition 6.1 acts transitively on the classes of column superstandard and reverse row superstandard tabloids; applying it to our sequence of Knuth move shows that, in fact, every reverse row superstandard tabloid has a column superstandard tabloid in its connected component. This completes the second claim. □

7.1 Bounding the monodromy group from above

In this section, we establish some restrictions on the monodromy group which, in particular, imply that the monodromy group for rectangular shapes in trivial. In the next section, we will prove that these are the only restrictions.

Proposition 7.5 *Suppose λ is a partition and w is a permutation of shape λ. Suppose $\alpha \in G_w$. If $\lambda_i = \lambda_{i+1}$ then $\alpha_i = \alpha_{i+1}$.*

For example, in Example 7.2 the coordinates of the monodromy element corresponding to the two rows of length 1 are equal.

Proof If this were not the case then repeating the monodromy or its inverse would eventually lead to a weight ρ for which the difference $\rho_i - \rho_{i+1}$ is arbitrarily large. From Definition 5.9, such a weight is non-dominant. However, since the image of AMBC is Ω_{dom}, it follows from Theorem 3.11 that Knuth moves preserve dominance. □

This result has an interesting corollary.

Corollary 7.6 *If λ is a rectangle then for any w of shape λ, the monodromy group G_w is trivial and the Knuth equivalence class of w is finite.*

Proof Each Knuth move preserves the sum of the entries of ρ. Thus, every element in G_w has sum of coordinates 0. By Proposition 7.5, if λ has parts of a unique size, then

every element in G_w has all coordinates equal. Putting the two together gives the first half of the claim. For the second half, there are only finitely many choices of a tabloid Q of a given shape, and by the first half of the claim there is at most one choice of ρ in the equivalence class once P, Q are fixed. □

We will see as part of Theorem 7.28 that for any non-rectangular shape, there exists a non-trivial monodromy, so the corresponding Kazhdan–Lusztig molecule is infinite.

Definition 7.7 Suppose $\lambda = \langle r_1^{m_1}, \ldots, r_k^{m_k} \rangle$ is a partition whose distinct row-lengths are $r_1 > \cdots > r_k > 0$, of multiplicities $m_1, \ldots, m_k$, and let $m'_i = m_1 + \cdots + m_i$ be the distinct part-sizes of λ'. For $1 \leqslant i \leqslant k$, let $e_i \in \mathbb{Z}^{\ell(\lambda)}$ be the vector with 1's in its first m'_i rows and 0's afterward. Define

$$G_\lambda^{\mathrm{up}} = \left\{ \sum_{i=1}^{k} a_i e_i \;\middle|\; a_1, \ldots, a_k \in \mathbb{Z} \text{ and } \sum_{i=1}^{k} a_i m'_i = 0 \right\}. \tag{7}$$

Example 7.8 For the shape $\lambda = \langle 3^2, 2^3, 1 \rangle$, one has $e_1 = (1, 1, 0, 0, 0, 0)$, $e_2 = (1, 1, 1, 1, 1, 0)$, and $e_3 = (1, 1, 1, 1, 1, 1)$. The space G_λ^{up} consists of vectors in the $\mathbb{Z}$-span of the e_i whose sum of coordinates is 0, such as $3e_1 - 6e_2 + 4e_3 = (1, 1, -2, -2, -2, 4)$.

Proposition 7.9 *We have* $G_w \subseteq G_\lambda^{\mathrm{up}}$.

Proof Taking $e_0 = 0$ by convention, the vector $e_i - e_{i-1}$ has 1's exactly in those indices for which the corresponding row of λ has length r_i. Thus, G_λ^{up} contains all integer vectors α with entry sum 0 such that $\alpha_i = \alpha_{i+1}$ whenever $\lambda_i = \lambda_{i+1}$. Then the result follows from Proposition 7.5. □

The next section is devoted to proving the reverse inclusion $G_w \supseteq G_\lambda^{\mathrm{up}}$, and thus equality.

7.2 Explicit description of monodromy generators

In this section, we construct an explicit set of monodromies that generate all of G_λ^{up}. We start by describing a $\mathbb{Z}$-basis for the lattice of solutions of a linear Diophantine equation in many variables.

7.2.1 Linear Diophantine equations in many variables

Let $m'_1, \ldots, m'_k \in \mathbb{Z}$. Consider the lattice (in $\mathbb{Z}^k$) of solutions to the equation

$$\sum_{i=1}^{k} m'_i x_i = 0. \tag{8}$$

Some solutions of this equation are obvious.

Definition 7.10 For $1 \leqslant i < j \leqslant k$, define $\mathbf{x}^{(i,j)} \in \mathbb{Z}^k$ by

$$x_l^{(i,j)} = \begin{cases} -\frac{m'_j}{\gcd(m'_i, m'_j)} & \text{if } l = i, \\ \frac{m'_i}{\gcd(m'_i, m'_j)} & \text{if } l = j, \\ 0 & \text{otherwise.} \end{cases}$$

The main result of this section is that these obvious solutions generate the full lattice. (We suspect that this result is known, but have not been able to find it in the literature.)

Proposition 7.11 *The lattice of solutions of* (8) *is the* $\mathbb{Z}$*-span of* $\mathbf{x}^{(i,j)}$.

Example 7.12 Consider the equation

$$3x + 2y + 6z = 0.$$

The three obvious solutions are $(-2, 3, 0)$, $(0, -3, 1)$, and $(-2, 0, 1)$. The claim is that these three vectors span the (2-dimensional) lattice of solutions.

We begin the proof of Proposition 7.11 with a technical lemma.

Lemma 7.13 *For any integers* $m'_1, \ldots, m'_j$, *we have*

$$\gcd\left(\frac{m'_1}{\gcd(m'_1, m'_j)}, \ldots, \frac{m'_{j-1}}{\gcd(m'_{j-1}, m'_j)}\right) = \frac{\gcd(m'_1, \ldots, m'_{j-1})}{\gcd(m'_1, \ldots, m'_j)}.$$

Proof Consider a prime p. For each $1 \leqslant i \leqslant j$, denote by v_i the maximal power of p dividing m'_i. First, suppose that $v_j > v_i$ for some i. Then p does not divide $\frac{m'_i}{\gcd(m'_i, m'_j)}$, and similarly p does not divide $\frac{\gcd(m'_1, \ldots, m'_{j-1})}{\gcd(m'_1, \ldots, m'_j)}$. Thus p does not divide either side of the claimed equality.

Second, suppose $v_j \leqslant \min\{v_1, \ldots, v_{j-1}\}$. Then the maximal power of p that divides the right-hand side is $\min\{v_1, \ldots, v_{j-1}\} - v_j$, while the maximal power of p that divides the left-hand side is $\min\{v_1 - v_j, \ldots, v_{j-1} - v_j\} = \min\{v_1, \ldots, v_{j-1}\} - v_j$. Since these are equal and p is arbitrary, the result follows. □

Proof of Proposition 7.11 Suppose we have some solution $\mathbf{x}^{(0)} = (x_1^{(0)}, \ldots, x_k^{(0)})$ of (8). Clearly $m'_k x_k^{(0)}$ is divisible by $\gcd(m'_1, \ldots, m'_{k-1})$. Thus $x_k^{(0)}$ is divisible by $\frac{\gcd(m'_1, \ldots, m'_{k-1})}{\gcd(m'_1, \ldots, m'_k)}$. By Lemma 7.13 and the Euclidean algorithm, there are integers $c_1, \ldots, c_{k-1}$ such that

$$\sum_{i=1}^{k-1} c_i \cdot \frac{m'_i}{\gcd(m'_i, m'_k)} = x_k^{(0)}.$$

Then

$$\mathbf{x}^{(1)} := \mathbf{x}^{(0)} - \sum_{i=1}^{k-1} c_i \cdot \mathbf{x}^{(i,k)}$$

is a solution of (8) with the last entry $x_k^{(1)} = 0$. Repeating the argument to successively eliminate coordinates produces an expression for $\mathbf{x}^{(0)}$ as an integer linear combination of the $\mathbf{x}^{(i,j)}$, as desired. □

7.2.2 Column rebases

In this section and the next, we provide the building blocks necessary to produce an explicit set of paths in $\mathcal{A}_\lambda$ that, when lifted to permutations, correspond to elements in the monodromy group G_w and generate G_λ^{up}. By Remark 7.3 and Lemma 7.4, it suffices to understand the situation in which we start with a column superstandard tabloid. Thus, our constructions are adapted to this case.

It is convenient to work not with tabloids but with associated tableaux, filled with equivalence classes modulo n, that arise by ordering the entries of each row into columns. We use usual matrix coordinates to refer to entries of tableaux, so that $(1, 3)$ is the entry in the first (= top = longest) row and third column.

Our first construction, called a *column rebase*, is defined below; the definition is illustrated in Example 7.15 and Fig. 17.

Definition 7.14 Suppose that T is a tabloid and $\widetilde{T}$ is a corresponding tableau such that, for some column indices c_1, c_2 and some nonnegative integers i, k, l, column c_1 of $\widetilde{T}$ is filled (in order from top to bottom) with the elements $\overline{i+1}, \overline{i+2}, \ldots, \overline{i+k+l}$ and column c_2 is filled with the elements $\overline{i+k+l+1}, \overline{i+k+l+2}, \ldots, \overline{i+k+2l}$. Then the *forward column rebase* of *type* (i, k, l) of $\widetilde{T}$ is the sequence of tableaux $(\widetilde{T} = \widetilde{T}_0, \widetilde{T}_1, \ldots, \widetilde{T}_{(k+l)\cdot l})$, where $\widetilde{T}_j$ is formed from $\widetilde{T}_{j-1}$ by switching the elements in positions $((k+l)m - j + 1, c_1)$ and (m, c_2), where $m = \lceil j/(k+l) \rceil$ is the unique integer such that

$$(k+l) \cdot (m-1) + 1 \le j \le (k+l) \cdot m.$$

We say that the reversed sequence is the *backward column rebase* of $\widetilde{T}' := \widetilde{T}_{(k+l)\cdot l}$. The *forward column rebase* of T is the sequence of tabloids that we get by forgetting the row orders in every tableau in the forward column rebase of $\widetilde{T}$, and similarly for the *backward column rebase* of the tabloid T' associated to $\widetilde{T}'$.

Example 7.15 Below is a forward column rebase of type $(6, 2, 2)$. The affected columns are the first and third ones.

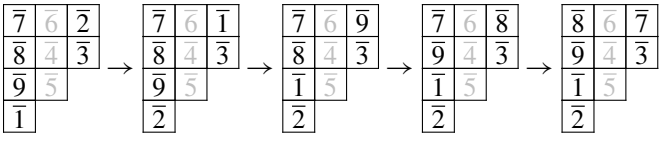

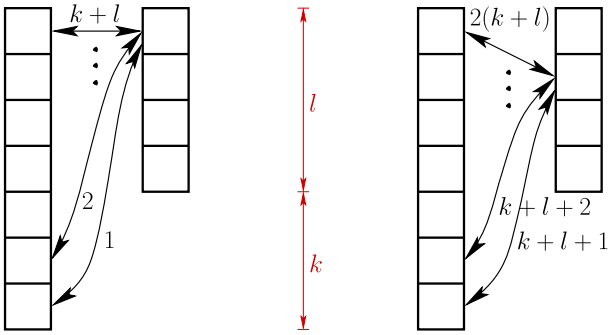

Fig. 17 Sequence in which exchanges are to be performed

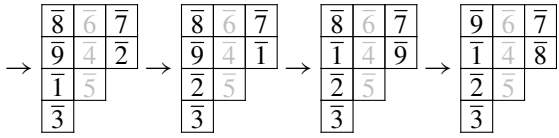

Proposition 7.16 *In the forward column rebase* $(T_0, \ldots, T_{(k+l)\cdot l})$ *of type* (i, k, l) *of the tabloid* T_0*, every two consecutive tabloids are either equal or connected by a Knuth move.*

Proof Let $(\widetilde{T}_0, \ldots, \widetilde{T}_{(k+l)\cdot l})$ be the associated sequence of tableaux, and let c_1 and c_2 be the indices of the relevant columns. Choose $j \in [1, (k+l)\cdot l]$ and write $j = (k+l)\cdot(m-1) + j'$ for integers m, j' with $1 \leqslant m \leqslant l$ and $1 \leqslant j' \leqslant k+l$. It is straightforward to show by induction that column c_1 of tableau $\widetilde{T}_j$ consists of the entries

$$\begin{aligned}&\overline{i+m}, \overline{i+m+1}, \ldots, \overline{i+m+k+l-j'-1},\\ &\overline{i+m+k+l-j'+1}, \overline{i+m+k+l-j'+2}, \ldots, \overline{i+m+k+l}\end{aligned}$$

in order from top to bottom, while column c_2 consists of the entries

$$\begin{aligned}&\overline{i+1}, \overline{i+2}, \ldots, \overline{i+m-1}, \qquad \overline{i+m+k+l-j'},\\ &\overline{i+m+k+l+1}, \overline{i+m+k+l+2}, \ldots, \overline{i+k+2l}.\end{aligned}$$

The last swap taken was between the entries $\overline{i+m+k+l-j'}$ and $\overline{i+m+k+l-j'+1}$ in positions $(k+l-j'+1, c_1)$ and (m, c_2). If $k+l-j'+1 = m$ then the corresponding tabloids are equal. If $k+l-j'+1 > m$ then $\tau(T_{j-1})$ includes $\overline{i+m+k+l-j'-1}$ but not $\overline{i+m+k+l-j'}$ while $\tau(T_j)$ includes $\overline{i+m+k+l-j'}$ but not $\overline{i+m+k+l-j'-1}$, and so the swap is a Knuth move. Similarly, if $k+l-j'+1 < m$ then $\tau(T_{j-1})$ includes $\overline{i+m+k+l-j'}$ but not $\overline{i+m+k+l-j'+1}$ while $\tau(T_j)$ includes $\overline{i+m+k+l-j'+1}$ but not $\overline{i+m+k+l-j'}$, and so the swap is a Knuth move. This completes the proof. □

It follows from the preceding proof that in the final tabloid $T_{(k+l)\cdot l}$ of a forward column rebase, the shorter column c_2 is filled (in order from top to bottom) with the elements $\overline{i+1}, \overline{i+2}, \ldots, \overline{i+l}$ while the longer column c_1 is filled with the elements $\overline{i+l+1}, \overline{i+l+2}, \ldots, \overline{i+k+2l}$.

We now describe how weight changes under a column rebase. We will need two pieces of notation.

Definition 7.17 Suppose we have a partition λ and a positive integer $l \leqslant \ell(\lambda)$. Let $\mathbf{u}^l_{\text{col}} \in \mathbb{Z}^{\ell(\lambda)}$ be the vector with 1 in the first l rows and 0 in the remaining rows. (There will never be ambiguity about the shape λ, so we suppress it from the notation.)

The second piece is somewhat more involved. Consider a tableau $\widetilde{T}$ of shape λ such that one can apply a column rebase of type (i, k, l). Informally, the vector $\mathbf{u}^{\widetilde{T},l}_{\text{disp}}$ keeps track of how far 1 shifts in the process of a column rebase. If $\overline{1}$ does not lie in the two columns of interest, define $\mathbf{u}^{\widetilde{T},l}_{\text{disp}} := \mathbf{u}^l_{\text{col}}$. If $\overline{1}$ does lie in one of these columns, consider the cyclic order on the columns shown in Fig. 18. Let F be the set of l cells following the cell containing $\overline{1}$ in this cyclic order, not including the cell itself, and define $\mathbf{u}^{\widetilde{T},l}_{\text{disp}} \in \mathbb{Z}^{\ell(\lambda)}$ by

$$\left(\mathbf{u}^{\widetilde{T},l}_{\text{disp}}\right)_m := \text{ number of cells of } F \text{ in row } m.$$

Example 7.18 Consider the column rebase from Example 7.15, so $l = 2$. In this case the collection F consists of the second and third cell in the first column. Thus

$$\mathbf{u}^l_{\text{col}} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} \quad \text{and} \quad \mathbf{u}^{\widetilde{T},l}_{\text{disp}} = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}.$$

Proposition 7.19 *Suppose T is a tabloid of shape λ to which one can apply a forward column rebase of type (i, k, l), with associated tableau $\widetilde{T}$. Then the weight change vector of the rebase is*

$$\mathbf{u}^l_{\text{col}} - \mathbf{u}^{\widetilde{T},l}_{\text{disp}}.$$

Proof If $\overline{1}$ is not in either of the two columns involved in the rebase, then by Theorem 3.11 the weight change is 0 at each step, as claimed. Otherwise, $\overline{1}$ lies somewhere in the two columns of interest. There are three cases, depending on the position of $\overline{1}$; these are pictured in Fig. 19. The three cases involve similar considerations, so we provide a detailed analysis of just one of them.

Suppose that $\overline{1}$ is in the shorter of the two relevant columns in $\widetilde{T}$, i.e., the column that does not contain $\overline{i+1}$, and let r be the index of the row in which it appears. (In Fig. 19, this is the blue region; in Fig. 20, it corresponds to the right panel.) The first $(k+l)\cdot(r-1)$ steps of the column rebase do not involve the cell containing

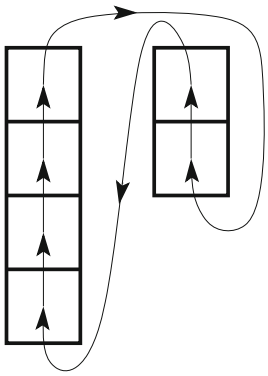

Fig. 18 A cyclic order on the cells of two columns

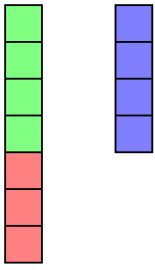

Fig. 19 The cases for the position of $\overline{1}$ in the proof of Proposition 7.19

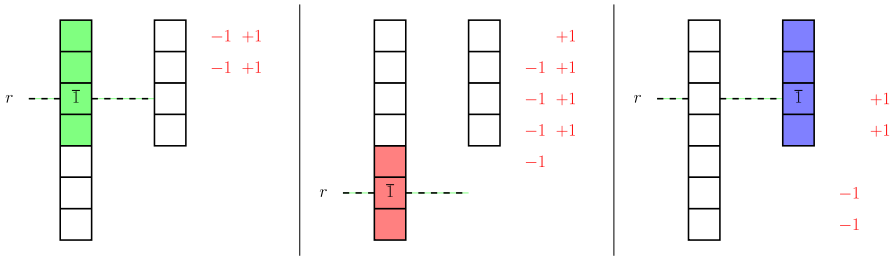

Fig. 20 The action of a forward column rebase on weights, depending on the initial position of $\overline{1}$

$\overline{1}$, and so by Theorem 3.11 do not change the weight. In step $(k+l)\cdot(r-1)+1$, the entry $\overline{1}$ in row r is swapped with the entry $\overline{n}$ in row $k+l$; this adds 1 to ρ_r and subtracts 1 from ρ_{k+l}. The next $k+l-1$ swaps to do not involve the entry $\overline{1}$. Finally, for $m=1,\ldots,l-r$, we have that in the mth remaining set of $k+l$ swaps, the entry $\overline{1}$ begins in row $k+l-m+1$, is swapped with the entry $\overline{2}$ in row $r+m$, and then is swapped with entry $\overline{n}$ in row $k+l-m$. This last swap is the only one that affects the weight ρ: it adds 1 to ρ_{m+r} and subtracts 1 from ρ_{k+l-m}. (This is illustrated in Fig. 20.) The net effect of these additions is to add 1 to ρ in rows $r, r+1,\ldots,l$ and to subtract 1 from ρ in rows $k+l, k+l-1,\ldots,k+r$. This coincides exactly with $\mathbf{u}^l_{\mathrm{col}} - \mathbf{u}^{\widetilde{T},l}_{\mathrm{disp}}$. □

7.2.3 Column cycles

In this section, we make a second step towards explicit monodromy generators by stitching column rebases together into longer paths in the KL DEG. Again, for the terminological convenience of being able to talk about columns, we often deal with tableaux.

Definition 7.20 Suppose λ is a partition and $i \in \mathbb{Z}$. Let λ' be the conjugate partition of λ. The *column superstandard tableau* of shape λ with *start* i is the tableau with whose first column contains in order from top to bottom the entries $\overline{i}, \overline{i+1}, \ldots, \overline{i+\lambda'_1-1}$, whose second column contains $\overline{i+\lambda'_1}, \overline{i+\lambda'_1+1}, \ldots, \overline{i+\lambda'_1+\lambda'_2-1}$, and so on.

Note that our definition differs slightly from the usual one, in which there is only a single column superstandard tableau of each shape (with start 1). For every column superstandard tabloid, it is possible to assign the entries of each row to columns in order to produce a column superstandard tableau, and forgetting the row assignments in a column superstandard tableau leaves a column superstandard tabloid. An example of a column superstandard tableau is given in Fig. 22.

Definition 7.21 Suppose T is the column superstandard tabloid of shape λ with start $\overline{i}$, and $\widetilde{T}$ is the corresponding tableau. The *forward column cycle* of $\widetilde{T}$ with respect to column j is the following sequence of tableaux: it is the concatenation of the forward column rebase of $\widetilde{T}$ on the $(j-1)$-st and j-th columns, the forward column rebase of the resulting tableau on the $(j-2)$-nd and j-th columns, etc., through a forward column rebase on the first and j-th columns, followed by the reverse column rebase on the j-th and λ_1-th (last) columns, then the reverse column rebase on the j-th and (λ_1-1)-st columns, etc., through the reverse column rebase on the j-th and $(j+1)$-st columns.

The *forward column cycle* on T is the sequence of tabloids that one gets by forgetting the row orders in every tableau in the forward column cycle of $\widetilde{T}$. The reversed sequences of tableaux and tabloids are called *reverse column cycles*.

Example 7.22 A forward column cycle of a tableau with respect to the second column is shown below; the first step is a forward column rebase, and the latter two are reverse column rebases:

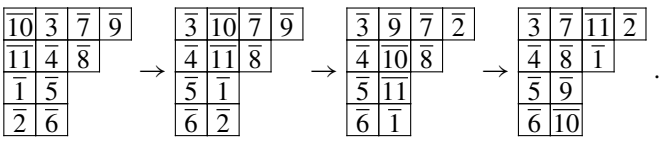

Proposition 7.23 *Forward column cycles are well defined. Furthermore, if $\widetilde{T}$ is the column superstandard tableau of shape λ and start $\overline{i}$, then the last tableau in the forward column cycle of $\widetilde{T}$ is the column superstandard tableau of shape λ and start $\overline{i+\lambda'_j}$.*

Proof It is conceptually straightforward but notationally cumbersome to check by induction that, after each rebase, every column consists of an interval in $[\overline{n}]$ and the two columns that are relevant for the next rebase have entries that are related in the appropriate way. We leave the details to the reader. □

Now we describe the effect of a column cycle on the weight. As before, we start by introducing a vector tracking the displacement of $\overline{1}$. Suppose $\widetilde{T}$ is a column superstandard tableau of shape λ and k is an integer. Consider the cyclic order on the cells of λ illustrated in Fig. 21: every cell not in the first row is followed by the cell directly above it, the top cell in the first column is followed by the bottom cell in the last column, and the top cell in every other column is followed by the bottom cell in the previous column. Denote by F the collection of k cells following the cell of $\widetilde{T}$ containing $\overline{1}$ in this order. Define $\mathbf{v}_{\mathrm{disp}}^{\widetilde{T},k} \in \mathbb{Z}^{\ell(\lambda)}$ by

$$\left(\mathbf{v}_{\mathrm{disp}}^{\widetilde{T},k}\right)_m := \text{ number of cells of } F \text{ in row } m.$$

Example 7.24 Suppose $\widetilde{T}$ is the column superstandard tableau of shape $\langle 4, 3, 2^2\rangle$ with start $\overline{10}$; it is shown in Fig. 22. The four cells following the cell with $\overline{1}$ are shown in the figure with red squares. Thus $\mathbf{v}_{\mathrm{disp}}^{\widetilde{T},4} = \begin{pmatrix} 2 \\ 2 \\ 0 \\ 0 \end{pmatrix}$.

Proposition 7.25 *Suppose T is a column superstandard tabloid of shape λ, with corresponding tableau $\widetilde{T}$. Consider performing a forward column cycle with respect to*

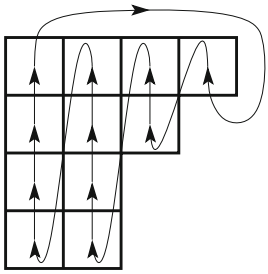

Fig. 21 The cyclic order on the partition $\langle 4, 3, 2^2\rangle$

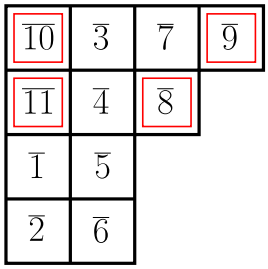

Fig. 22 The column superstandard tableau of shape $\langle 4, 3, 2^2\rangle$ with start at $\overline{10}$. The four cells following $\overline{1}$ are boxed in red

the j-th column. Then the weight change vector is

$$\Delta_{\text{wt}} = \mathbf{u}_{\text{col}}^{\lambda_{j'}} - \mathbf{v}_{\text{disp}}^{\widetilde{T},\lambda'_j}.$$

Example 7.26 Consider the column cycle of Example 7.22. Computing the weight change for the three column rebases via Proposition 7.19, the total weight change vector is

$$\Delta_{\text{wt}} = \begin{pmatrix}0\\0\\0\\0\end{pmatrix} + \begin{pmatrix}-1\\0\\1\\0\end{pmatrix} + \begin{pmatrix}0\\-1\\0\\1\end{pmatrix} = \begin{pmatrix}-1\\-1\\1\\1\end{pmatrix}.$$

On the other hand, the vector $\mathbf{v}_{\text{disp}}^{\widetilde{T},4}$ was computed in Example 7.24. Since we are dealing with the longest column, the vector $\mathbf{u}_{\text{col}}^{\lambda'_1}$ consists of all 1's. Thus by Proposition 7.25, the weight change is

$$\Delta_{\text{wt}} = \begin{pmatrix}1\\1\\1\\1\end{pmatrix} - \begin{pmatrix}2\\2\\0\\0\end{pmatrix} = \begin{pmatrix}-1\\-1\\1\\1\end{pmatrix}.$$

Proof of Proposition 7.25 There are five cases, depending on where $\overline{1}$ is in the tabloid: whether it is in the first λ'_j rows or not, and if so whether it is in the first column, between column 1 and column j, in column j, or in a larger-numbered column. These are illustrated in Fig. 23; we refer to them below by their color in this figure. (If $j = 2$, the blue case does not occur; if $j = 1$, the blue and purple cases do not occur.)

The green case, when $\overline{1}$ is in a row lower than λ'_j, is the easiest one; in this case $\overline{1}$ is only involved in a single column rebase, and it is easy to see that $\mathbf{u}_{\text{disp}}^{\widetilde{T},\lambda'_j} = \mathbf{v}_{\text{disp}}^{\widetilde{T},\lambda'_j}$. This gives the desired result.

In the blue case, when $\overline{1}$ is in the first λ'_j rows in a column between column 1 and column j, there are two column rebases affecting $\overline{1}$. The first one is of the kind portrayed on the left of Fig. 20; in such a rebase, the vectors $\mathbf{u}_{\text{col}}^{\lambda'_j}$ and $\mathbf{u}_{\text{disp}}^{\widetilde{T},\lambda'_j}$ are equal, and so the weight does not change. In the second rebase, the shorter column involved has length λ'_j. In addition, in this rebase the longer column is directly to the right of the one in which $\overline{1}$ originally started; thus, $\mathbf{u}_{\text{disp}}^{\widetilde{T},\lambda'_j}$ for the second rebase is equal to $\mathbf{v}_{\text{disp}}^{\widetilde{T},\lambda'_j}$. This gives the desired result.

If $\overline{1}$ is in the purple region (so $j > 1$) then only one rebase affects the position of $\overline{1}$ and it is exactly the same as the second rebase of the blue case.

In the red and yellow cases, the position of $\overline{1}$ can be affected by more than two column rebases, making these more complicated to analyze. These cases are very similar, so we will only consider the yellow case in detail. Suppose that $\overline{1}$ begins in the yellow region, say in column $j' > j$ and in row r. Then the weight does not change

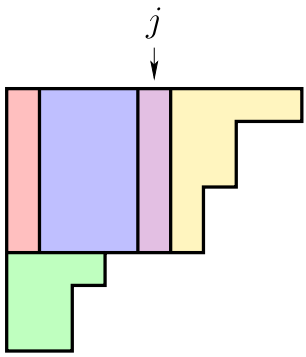

Fig. 23 The cases in the proof of Proposition 7.25

until the reverse rebase between column j and column j'; during this rebase, $\overline{1}$ moves from column j' to column j and remains in row r. What follows is several (possibly none) reverse rebases during which $\overline{1}$ moves lower in column j: after the first one it will be in row $r+\lambda'_{j'-1}$, then in row $r+\lambda'_{j'-1}+\lambda'_{j'-2}$, etc. Let j'' be the last column for which the reverse rebase affects the weight. This can happen in one of two ways: either $j''=j+1$ and the algorithm ends with $\overline{1}$ in column j, or the rebase transfers $\overline{1}$ into column j'', where it stays until the end. Let r'' be the final row of $\overline{1}$.

If $\overline{1}$ ends up in column j, then $\lambda'_j \geqslant r''=r+\lambda_{j+1}+\cdots+\lambda_{j'-1}$ and so by direct examination of the cyclic order on cells of λ we have

$$\begin{aligned}\mathbf{u}_{\mathrm{col}}^{\lambda'_j}-\mathbf{v}_{\mathrm{disp}}^{\widetilde{T},\lambda'_j}&=\mathbf{u}_{\mathrm{col}}^{\lambda'_j}-\left(\mathbf{u}_{\mathrm{col}}^{r-1}+\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'-1}}+\cdots+\mathbf{u}_{\mathrm{col}}^{\lambda'_{j+1}}+\left(\mathbf{u}_{\mathrm{col}}^{\lambda'_j}-\mathbf{u}_{\mathrm{col}}^{r''-1}\right)\right)\\&=\mathbf{u}_{\mathrm{col}}^{r''-1}-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j+1}}-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j+2}}-\cdots-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'-1}}-\mathbf{u}_{\mathrm{col}}^{r-1}.\end{aligned}$$

Otherwise,

$$\mathbf{u}_{\mathrm{col}}^{\lambda'_j}-\mathbf{v}_{\mathrm{disp}}^{\widetilde{T},\lambda'_j}=\mathbf{u}_{\mathrm{col}}^{\lambda'_j}-(\mathbf{u}_{\mathrm{col}}^{\lambda'_{j''}}-\mathbf{u}_{\mathrm{col}}^{r''-1})-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j''+1}}-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j''+2}}-\cdots-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'-1}}-\mathbf{u}_{\mathrm{col}}^{r-1}.$$

On the other hand, we can track the weight change at each stage of the column cycle using Proposition 7.19. Since we are doing reverse column rebases, the contributions from the col-vectors are negated relative to the proposition, and they sum to

$$-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'}}-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'-1}}-\cdots-\mathbf{u}_{\mathrm{col}}^{\lambda'_{j''}}.$$

By tracking the position of $\overline{1}$ after each column rebase, we have that if $\overline{1}$ ends up in column j, then the total contribution from the disp-vectors is

$$\mathbf{u}_{\mathrm{col}}^{r''-1}+\left(\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'}}-\mathbf{u}_{\mathrm{col}}^{r-1}\right).$$

If instead $\overline{1}$ ends up in column $j'' \neq j$, this contribution is $\mathbf{u}_{\mathrm{col}}^{r''-1} + \left(\mathbf{u}_{\mathrm{col}}^{\lambda'_{j'}} - \mathbf{u}_{\mathrm{col}}^{r-1}\right) + \mathbf{u}_{\mathrm{col}}^{\lambda'_j}$. In both cases, the resulting weight change matches the computed value $\mathbf{u}_{\mathrm{col}}^{\lambda'_j} - \mathbf{v}_{\mathrm{disp}}^{\widetilde{T},\lambda'_j}$. □

Corollary 7.27 *The weight change vectors associated to column cycles with respect to different columns of the same length are equal.*

7.2.4 Building monodromy group generators from column cycles

Consider the partition $\lambda = \langle r_1^{m_1}, \ldots, r_k^{m_k} \rangle$. For $1 \leqslant i \leqslant k$, let $m'_i = m_1 + \cdots + m_i$ be the distinct part-sizes of λ'. Recall the space G_λ^{up} from Definition 7.7.

Theorem 7.28 *Suppose λ is a partition and w is a permutation of shape λ. Then $G_w = G_\lambda^{\mathrm{up}}$.*

Proof By Proposition 7.9, $G_w \subseteq G_\lambda^{\mathrm{up}}$. We now prove the opposite containment. By (7), there is a natural bijection between elements of G_λ^{up} and integer solutions $(a_i)_{1 \leq i \leq k}$ of the equation $\sum_{i=1}^k a_i m'_i = 0$. By Proposition 7.11, the solution set of this equation is the $\mathbb{Z}$-span of the vectors $\mathbf{x}^{(i,j)} \in \mathbb{Z}^k$ defined in Definition 7.10. The map associating an element of G_λ^{up} to a solution is obviously $\mathbb{Z}$-linear, so it is sufficient to show that

$$\mathbf{v}^{(i,j)} := \frac{m'_i}{\gcd(m'_i, m'_j)} e'_j - \frac{m'_j}{\gcd(m'_i, m'_j)} e'_i$$

is an element of G_w.

As explained in Remark 7.3, by Lemma 7.4, we may restrict ourselves to the case when $Q(w)$ is column superstandard. Given such a tabloid T, consider the following concatenation of column cycles in $\mathcal{A}_\lambda$:

(1) a sequence of $\frac{m'_j}{\gcd(m'_i, m'_j)}$ reverse column cycles with respect to a column of size m'_i, followed by

(2) a sequence of $\frac{m'_i}{\gcd(m'_i, m'_j)}$ forward column cycles with respect to a column of size m'_j.

(By Corollary 7.27, it doesn't matter which columns of the given lengths are selected.) By Proposition 7.23, the final tabloid of this sequence of tabloids is equal to the initial tabloid T, i.e., this sequence is a loop in $\mathcal{A}_\lambda$. By Proposition 7.25, the weight change vector associated to this loop is equal to

$$\Delta_{\mathrm{wt}} = -\frac{m'_j}{\gcd(m'_i, m'_j)} \mathbf{u}_{\mathrm{col}}^{m'_i} + \sum_{k=0}^{m'_j/\gcd(m'_i,m'_j)-1} \mathbf{v}_{\mathrm{disp}}^{\widetilde{T}+km'_i,m'_i}$$

$$+\left(\frac{m_i'}{\gcd(m_i', m_j')}\mathbf{u}_{\mathrm{col}}^{m_j'} - \sum_{k=0}^{m_i'/\gcd(m_i',m_j')-1} \mathbf{v}_{\mathrm{disp}}^{\widetilde{T}+m_i'm_j'/\gcd(m_i',m_j')-km_j',m_j'}\right),$$

where by $\widetilde{T}+m$ we mean the result of adding m to every entry of the tableau $\widetilde{T}$. Observe that the two sums of disp-vectors cancel exactly: the reverse cycles and the forward cycles deal with the same (multi)set of $\frac{m_i'm_j'}{\gcd(m_i',m_j')}$ consecutive cells in the cyclic order. Thus

$$\begin{aligned}\Delta_{\mathrm{wt}} &= \frac{m_i'}{\gcd(m_i', m_j')}\mathbf{u}_{\mathrm{col}}^{m_j'} - \frac{m_j'}{\gcd(m_i', m_j')}\mathbf{u}_{\mathrm{col}}^{m_i'} = \frac{m_i'}{\gcd(m_i', m_j')}e_j' - \frac{m_j'}{\gcd(m_i', m_j')}e_i' \\ &= \mathbf{v}^{(i,j)}.\end{aligned}$$

Lifting this cycle from $\mathcal{A}_\lambda$ to the graph on permutations shows that $\mathbf{v}^{(i,j)} \in G_w$ for every $1 \leqslant i < j \leqslant k$. This completes the proof. □

8 Components of KL DEGs

In this section, we finish the description of the Knuth equivalence classes of permutations by describing which tabloids can be reached from a given one by Knuth moves. Using the machinery from the previous section, it is easy to see that certain column superstandard tabloids are connected by Knuth moves; together with the fact that every Knuth class of tabloids contains a superstandard tabloid, this gives an upper bound on the number of connected components. Then we use the charge statistic to produce a matching lower bound.

Definition 8.1 For a partition λ, define

$$d_\lambda := \gcd(\lambda_1', \lambda_2', \dots),$$

where λ' is the conjugate partition of λ. Equivalently, d_λ is the greatest common divisor of the multiplicities of parts of λ.

Recall that $\mathcal{A}_\lambda$ denotes the KL DEG of shape λ (see Definition 3.21).

Proposition 8.2 *The number of connected components of $\mathcal{A}_\lambda$ is no larger than d_λ.*

Proof Column cycles are walks in $\mathcal{A}_\lambda$ that take a column superstandard tabloid to another one, related by adding λ_j' to each entry for some j. Since d_λ is a $\mathbb{Z}$-linear combination of the λ_j''s, two column superstandard tabloids that differ by d_λ are in the same connected component. By Lemma 7.4, every connected component of $\mathcal{A}_\lambda$ contains a column superstandard tabloid, so there are at most d_λ connected components. □

We seek to prove that this upper bound is in fact the correct value. Our proof makes use of the charge statistic, which was partially introduced in Sect. 5.

Definition 8.3 Given a tabloid T of shape λ, extend Definition 5.3 of local charge by defining $\mathrm{lch}_i(T) = 0$ if $\lambda_i \neq \lambda_{i+1}$. The *charge* of T is defined by

$$\mathrm{charge}(T) := \sum_{i=1}^{\ell(\lambda)-1} i \cdot \mathrm{lch}_i(T).$$

Remark 8.4 The definition of charge given here is not the conventional one (as in, e.g., [9, Sect. 5]). However, it is straightforward to see using Lemma 5.5 that our definition of the charge of a tabloid T differs by a multiple of d_λ from the conventional charge of the word whose i-th letter is the index of the row of T containing $\overline{i}$. In light of the statement of Theorem 8.6 below, this is sufficient for our purposes, and the present definition is easier to work with.

Before finishing the description of the connected components of $\mathcal{A}_\lambda$, we need a lemma regarding the charges of column superstandard tabloids.

Lemma 8.5 *Suppose λ is a partition of n and $1 \leqslant i \leqslant n$. Let T be the column superstandard tabloid of shape λ starting at $\overline{i}$, and T' be the column superstandard tabloid of shape λ starting at $\overline{i+1}$. Then* $\mathrm{charge}(T') \equiv \mathrm{charge}(T) - 1 \pmod{d_\lambda}$.

Proof By construction, all the local charges for a column superstandard tabloid are 0, except if $\overline{n}$ is located in the row directly above the row containing $\overline{1}$ and the two rows have the same size, in which case the local charge is 1. If $\overline{n}$ is in row r of T and $\lambda_{r-1} = \lambda_r = \lambda_{r+1}$ then $\mathrm{charge}(T) = r + 1$ and $\mathrm{charge}(T') = r$, and the result holds. If $\lambda_{r-1} = \lambda_r > \lambda_{r+1}$ or if $r = \ell(\lambda)$ then $\mathrm{charge}(T) = 0$ and $\mathrm{charge}(T') = r - 1$. In this case, $r = \lambda'_i$ for some i, so $r \equiv 0 \pmod{d_\lambda}$ and the result holds. Similar considerations apply if $\lambda_{r-1} < \lambda_r = \lambda_{r+1}$ or if $r = 1$. Finally, if $\lambda_{r-1} < \lambda_r < \lambda_{r+1}$ then $d_\lambda = 1$ and the result is trivial. □

Theorem 8.6 *Let λ be a partition and suppose T and T' are tabloids of shape λ. Then T and T' are in the same connected component of $\mathcal{A}_\lambda$ if and only if*

$$\mathrm{charge}(T) \equiv \mathrm{charge}(T') \pmod{d_\lambda}.$$

Proof We will show that if T and T' are two tabloids of shape λ connected by a Knuth move, then $\mathrm{charge}(T) \equiv \mathrm{charge}(T') \pmod{d_\lambda}$. The "only if" direction of the claimed result follows immediately by induction. The "if" direction then follows using Proposition 8.2 and Lemma 8.5.

Suppose the Knuth move between T and T' exchanges the entries $\overline{i}$ and $\overline{i+1}$; without loss of generality, we may assume that $\overline{i}$ is in a higher row of T than $\overline{i+1}$. We consider two cases, depending on the value of $\overline{i}$.

Case $\overline{i} \neq \overline{n}$. In this case, $\overline{i}$ and $\overline{i+1}$ are either both larger or both smaller in broken order than each other entry in the tabloid. Thus, if these entries lie in non-adjacent rows then the local charge between each pair of adjacent rows is the same in T and T',

and so $\mathrm{charge}(T) = \mathrm{charge}(T')$. Suppose instead that $\overline{i}$ and $\overline{i+1}$ lie in adjacent rows $r, r+1$. The only local charge that could differ between T and T' is lch_r, and so the only interesting case is when $\lambda_r = \lambda_{r+1}$. Moreover, the only way the local charge lch_r could differ between T and T' is if $\overline{i}$ is matched with $\overline{i+1}$ in the charge matching in T. By Lemma 5.5, the local charge does not depend on the activation ordering, so it suffices to produce an activation ordering for T in which $\overline{i}$ is not matched with $\overline{i+1}$.

In order for the swap of $\overline{i}$ and $\overline{i+1}$ to be a Knuth move, the descent sets $\tau(T)$ and $\tau(T')$ must be incomparable. This is possible only if T_r contains $\overline{i-1}$ or T_{r+1} contains $\overline{i+2}$ (or both). Suppose first that T_r contains $\overline{i-1}$. In this case, take an activation ordering that begins with $\overline{i-1}$: if $\overline{i} \neq \overline{1}$, then $\overline{i+1}$ is the smallest value larger than $\overline{i-1}$ in broken order, so $(\overline{i-1}, \overline{i+1})$ belongs to the charge matching; if $\overline{i} = \overline{1}$ then $\overline{i-1} = \overline{n}$ matches to the smallest element of T_2, which is $\overline{2} = \overline{i+1}$. Since $\overline{i-1}$ matches with $\overline{i+1}$ in both cases, $\overline{i}$ does not, as needed.

Now suppose instead that T_{r+1} contains $\overline{i+2}$. In this case, take an activation ordering that ends with $\overline{i}$: if $\overline{i} \neq \overline{n-1}$ then $\overline{i+2}$ gets matched strictly after $\overline{i+1}$ does, so $\overline{i+1}$ is matched to something earlier that $\overline{i}$; if $\overline{i} = \overline{n-1}$ then the second row contains $\overline{i+2} = \overline{1}$, which again gets matched strictly after $\overline{i+1} = \overline{n}$. In both cases, $\overline{i}$ does not match to $\overline{i+1}$, as needed.

Case $\overline{i} = \overline{n}$. Suppose that $\overline{n}$ is in row r_1 in T and $\overline{1}$ is in row $r_2 > r_1$. Then $\mathrm{lch}_j(T) = \mathrm{lch}_j(T')$ for every $j \notin \{r_1-1, r_1, r_2-1, r_2\}$, and so the charge difference $\mathrm{charge}(T) - \mathrm{charge}(T')$ depends only on those four or fewer local charges. For convenience, define $\Delta_j := j\,\mathrm{lch}_j(T) - j\,\mathrm{lch}_j(T')$. We have numerous possible cases to consider: whether or not $\lambda_j = \lambda_{j+1}$ for each $j \in \{r_1-1, r_1, r_2-1, r_2\}$, and whether $r_2 = r_1+1$ or not.

Suppose that $\lambda_{r_1-1} = \lambda_{r_1}$. Then we have by Lemma 5.11 that $\mathrm{lch}_{r_1-1}(T) = \mathrm{lch}_{r_1-1}(T') - 1$, and so $\Delta_{r_1-1} = 1 - r_1$. If instead $\lambda_{r_1-1} < \lambda_{r_1}$ or $r_1 = 1$ then $r_1 \equiv 1 \pmod{d_\lambda}$ and so $\Delta_{r_1-1} = 0 \equiv 1 - r_1 \pmod{d_\lambda}$ in this case as well.

It is not hard to modify the proof of Lemma 5.11 to the situation in which the largest and smallest elements are swapped out of the *top*, rather than bottom, row; it follows from this modified lemma and an argument identical to the preceding paragraph that $\Delta_{r_2} \equiv -r_2 \pmod{d_\lambda}$.

Suppose now that rows r_1 and r_2 are not adjacent, so that $T_{r_2-1} = T'_{r_2-1}$. If $\lambda_{r_2-1} = \lambda_{r_2}$ then by Lemma 5.11 we have $\mathrm{lch}_{r_2-1}(T') = \mathrm{lch}_{r_2-1}(T) - 1$, and so $\Delta_{r_2-1} = r_2 - 1$. As before, if $\lambda_{r_2-1} > \lambda_{r_2}$ then $r_2 \equiv 1 \pmod{d_\lambda}$ and so $\Delta_{r_2-1} = 0 \equiv r_2 - 1 \pmod{d_\lambda}$ in this case as well. Similarly, as in the preceding paragraph we have $\Delta_{r_1} \equiv r_1 \pmod{d_\lambda}$.

Combining the results of the previous three paragraphs, when $r_1 < r_2 - 1$ we have that

$$\begin{aligned}\mathrm{charge}(T) - \mathrm{charge}(T') = \Delta_{r_1-1} + \Delta_{r_1} + \Delta_{r_2-1} + \Delta_{r_2} &\equiv (1 - r_1) \\ &+ r_1 - r_2 + (r_2 - 1) \equiv 0 \pmod{d_\lambda},\end{aligned}$$

as claimed.

In the remainder of the proof, we consider the case that $r := r_1$ is equal to $r_2 - 1$, so that $\overline{1}$ and $\overline{n}$ lie in adjacent rows in T and T'. In this case, there are (at most) three contributions to the charge difference: $\mathrm{charge}(T) - \mathrm{charge}(T') = \Delta_{r-1} + \Delta_r + \Delta_{r+1}$.

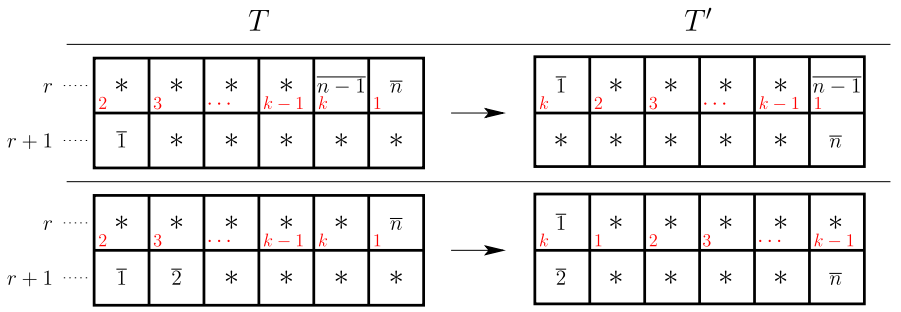

Fig. 24 Activation orders used in the proof of Theorem 8.6 in the case that $\overline{n}$ and $\overline{1}$ are swapped and lie in adjacent rows. Black numbers are tabloid entries; black $*$'s represent entries whose values are irrelevant; and red numbers represent the activation order

The analyses of the second and third paragraphs of the case $\overline{i} = \overline{n}$ still apply in this setting, giving $\Delta_{r-1} \equiv 1 - r$ and $\Delta_{r+1} \equiv -1 - r$ modulo d_λ.

First, suppose that $\lambda_r > \lambda_{r+1}$. Then $\Delta_r = 0$, and so $\Delta_{r-1} + \Delta_r + \Delta_{r+1} \equiv -2r \equiv 0 \pmod{d_\lambda}$, as needed.

Second (and last), suppose that $\lambda_r = \lambda_{r+1} =: k$. We claim that $\mathrm{lch}_r(T) - \mathrm{lch}_r(T') = 2$. Since the swap of $\overline{n}$ and $\overline{1}$ is a Knuth move, either $\overline{n-1}$ is in row r of T or $\overline{2}$ is in row $r+1$. We use different activation orders depending on which situation we are in; they are illustrated in Fig. 24. We analyze the case when $\overline{n-1}$ is in row r (the top half of the figure). Since the other case is similar, we omit the details, but we do provide the relevant illustration in the bottom half of the figure.

In T, we use the activation order that begins with $\overline{n}$ first, then proceeds through the other values in some order, then ends with $\overline{n-1}$. With this ordering, $\overline{n}$ matches with $\overline{1}$ in row $r+1$ and contributes 1 to the local charge. The next $k-2$ entries match in some order to $k-2$ of the elements in row $r+1$. Finally, in the last activation in T, $\overline{n-1}$ contributes 1 to the local charge (since the only entry of T larger than $\overline{n-1}$ is $\overline{n}$, and it is not in row $r+1$). In T', we use the order that begins with $\overline{n-1}$, then proceeds through the other values in the same order as in T, then ends with $\overline{1}$. With this activation ordering, $\overline{n-1}$ matches to $\overline{n}$ in row $r+1$ and does not contribute to the local charge. At this point, the unmatched entries in row $r+1$ of T' are exactly the same as the unmatched entries in row $r+1$ of T after the first pair are matched; hence the next $k-2$ edges of the charge matching are exactly the same in T' as in T, with the same contribution to charge. Finally, the last match is between $\overline{1}$ and something larger in broken order, with no contribution to local charge. Thus indeed the local charge between these rows is 2 less for T'.

Putting everything together, we have in this case that $\Delta_r = 2r$ and so

$$\begin{aligned}\mathrm{charge}(T) - \mathrm{charge}(T') &= \Delta_{r-1} + \Delta_r + \Delta_{r+1} \equiv (1-r) + 2r + (-1-r) \\ &= 0 \pmod{d_\lambda},\end{aligned}$$

as claimed. This completes all cases. □

One interesting corollary is that we can determine exactly when a Kazhdan–Lusztig cell consists of a single Knuth equivalence class.

Corollary 8.7 *The Kazhdan–Lusztig cell containing the permutation w consists of a single Knuth equivalence class if and only if the shape λ of w has all rows distinct.*

Proof Choose a permutation $w \in \widetilde{S}_n$ of shape λ. If λ has a repeated part, say $\lambda_r = \lambda_{r+1}$, then by Theorem 7.28, the difference $\rho(w')_r - \rho(w')_{r+1}$ is the same for every w' in the Knuth equivalence class of w. However, from the definition of Ω_{dom}, this is not the case for every w' in the same Kazhdan–Lusztig cell as w. Thus distinct rows is a necessary condition. On the other hand, if all rows of λ are distinct then by Theorem 8.6, the Knuth class contains permutations with every possible Q-tabloid of shape λ, and by Theorem 7.28, for each Q-tabloid the Knuth class contains permutations with every possible weight vector ρ. Thus in this case the Knuth class is equal to the Kazhdan–Lusztig cell. □

Remark 8.8 The results of this section combined with Proposition 6.1 imply that the action of the shift permutation $s = [2, 3, \ldots, n, n+1]$ on tabloids of shape λ provides graph isomorphisms between the d_λ connected components of $\mathcal{A}_\lambda$.

9 Crystals

In this section, we briefly describe a connection to affine crystals that is suggested by the importance of the charge statistic; the finite version of the connection has been developed by Assaf [2]. We use notation for affine crystals from [20, Sect. 4], except we label crystal operators f_i and e_i by residue classes modulo n as opposed to by the representatives of these classes from $\{0, 1, \ldots, n-1\}$.

There is a *reading word map* RW from the set of tabloids of shape $\lambda = \langle \lambda_1, \lambda_2, \ldots, \lambda_\ell \rangle$ to the set of vertices of $B := B^{1,\lambda_\ell} \otimes B^{1,\lambda_{\ell-1}} \otimes \cdots \otimes B^{1,\lambda_1}$: it reads the rows from bottom to top and takes representatives in $[n]$ for all residue classes. For example,

$$RW\left(\begin{array}{|c|c|c|}\hline \overline{2} & \overline{3} & \overline{7} \\ \hline \overline{1} & \overline{4} \\ \cline{1-2} \overline{5} \\ \cline{1-1} \overline{6} \\ \cline{1-1}\end{array}\right) = \boxed{6} \otimes \boxed{5} \otimes \boxed{1}\boxed{4} \otimes \boxed{2}\boxed{3}\boxed{7}.$$

The image of this map is precisely the set of vertices of B where each element of $[n]$ is used exactly once; call this set of vertices V. For $\overline{i} \in [\overline{n}]$, let $K_{\overline{i}}$ be the composition of crystal operators $e_{\overline{i}} e_{\overline{i+1}} f_{\overline{i}} f_{\overline{i+1}}$.

Proposition 9.1 *Suppose T and T' are tabloids of the same shape. Then T and T' are connected by a Knuth move if and only if $RW(T) \neq RW(T')$ and the two reading words are connected by some $K_{\overline{i}}$.*

Proof sketch Suppose T and T' are connected by a Knuth move that exchanges $\overline{i}$ and $\overline{i+1}$; without loss of generality $\overline{i} \notin \tau(T)$. Suppose $\overline{i}$ is in row k of T and $\overline{i+1}$ is in

row ℓ; so $k > \ell$. Since the exchange of $\overline{i}$ and $\overline{i+1}$ is a Knuth move, we know that either $\overline{i-1}$ lies in a row weakly below ℓ and strictly above k, or $\overline{i+2}$ lies in a row strictly below ℓ and weakly above k. We treat the first of these cases; the second is identical. One easily checks by definition of the crystal operators that $K_{\overline{i-1}}(RW(T)) = RW(T')$ (the computation depends on whether $i = 1$, $i = n$, or neither).

Now suppose that $RW(T) \neq RW(T')$ and $K_{\overline{i}}(RW(T)) = RW(T')$. One needs to check that whatever the positions of (the representatives in $[n]$ of) $\overline{i}$, $\overline{i+1}$ and $\overline{i+2}$ in $RW(T)$, the result of the action by $K_{\overline{i}}$ yields either the reading word of a tabloid related to T by a Knuth move, or $RW(T)$ again, or some vertex of B that is not in V. The computation depends on whether $\overline{i} = \overline{n}$, $\overline{i} = \overline{n-1}$, or neither. □

Finally, we give a simple application of the crystal connection. Notice that the definition of KL DEGs $\mathcal{A}_\mu$ does not rely on the shape μ being a partition.

Proposition 9.2 *Suppose μ is a composition and λ is the partition formed from μ by sorting. Then there is a graph isomorphism $\varphi : \mathcal{A}_\lambda \to \mathcal{A}_\mu$.*

Proof It is sufficient to show that A_μ is isomorphic to $A_{\mu'}$ where μ' is obtained from μ by switching two adjacent parts. By [20, Thm. 4.8], there is an isomorphism between $B^{1,r} \otimes B^{1,s}$ and $B^{1,s} \otimes B^{1,r}$ for any r and s (it is called the *combinatorial R-matrix*). Moreover, the explicit description of this morphism in the same section shows that it preserves the number of times each letter appears in the tableaux. Applying this morphism to two adjacent factors of B and pulling back the results using Proposition 9.1 (which also does not rely on the shape being a partition) finishes the proof. □

References

1. Ariki, S.: Robinson-Schensted correspondence and left cells. arXiv:math/9910117, (1999)
2. Assaf, S. H.: A combinatorial realization of Schur-Weyl duality via crystal graphs and dual equivalence graphs. Discrete Mathematics & Theoretical Computer Science (FPSAC 2008)
3. Assaf S. H.: Dual equivalence graphs I: A new paradigm for Schur positivity. Forum Math. Sigma, 3:e12, 33, (2015)
4. Beissinger, J.S.: Similar constructions for Young tableaux and involutions, and their application to shiftable tableaux. Discrete Math. **67**(2), 149–163 (1987)
5. Blasiak, J., Fomin, S.: Noncommutative Schur functions, switchboards, and Schur positivity. Selecta Math. (N.S.) **23**(1), 727–766 (2017)
6. Barbasch, D., Vogan, D.: Primitive ideals and orbital integrals in complex classical groups. Math. Ann. **259**, 153–199 (1982)
7. Chmutov, M.: Affine Matrix-Ball Construction Java program. https://sites.google.com/site/affiners/ (also included in the arXiv distribution), (2015)
8. Chmutov, M., Pylyavskyy, P., Yudovina, E.: Matrix-ball construction of affine Robinson-Schensted correspondence. Selecta Math. (N.S.) **24**(2), 667–750 (2018)
9. Dalal, A. J., Morse, J.: A t-generalization for Schubert representatives of the affine Grassmannian. In: 25th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2013), Discrete Math. Theor. Comput. Sci. Proc., AS, pages 1125–1136. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2013
10. Fulton, W.: Young tableaux: with applications to representation theory and geometry, volume 35 of London Mathematical Society Student Texts. Cambridge University Press, (1997)
11. Garsia, A.M., McLarnan, T.J.: Relations between Young's natural and the Kazhdan-Lusztig representations of S_n. Advances in Mathematics **69**(1), 32–92 (1988)

12. Housley, M., Russell, H.M., Tymoczko, J.: The Robinson-Schensted correspondence and A_2-web bases. J. Algebraic Combin. **42**(1), 293–329 (2015)
13. Humphreys, J.E.: Reflection groups and Coxeter groups, Cambridge Studies in Advanced Mathematics, vol. 29. Cambridge University Press, Cambridge (1990)
14. Kazhdan, D., Lusztig, G.: Representations of Coxeter groups and Hecke algebras. Invent. Math. **53**, 165–184 (1989)
15. Lascoux, A., Schützenberger, M.-P.: Croissance des polynômes de Foulkes-Green. CR Acad. Sci. Paris **228**, 95–98 (1979)
16. Manivel, L.: Symmetric functions, Schubert polynomials and degeneracy loci, volume 6 of *SMF/AMS Texts and Monographs*. American Mathematical Society, Providence, RI; Société Mathématique de France, Paris, 2001. Translated from the 1998 French original by John R. Swallow, Cours Spécialisés [Specialized Courses], 3
17. Reifegerste, A.: Permutation sign under the Robinson-Schensted correspondence. Annals of Combinatorics **8**(1), 103–112 (2004)
18. Schützenberger, M.P.: Quelques remarques sur une construction de Schensted. Math. Scand. **12**, 117–128 (1963)
19. Shi, J.Y.: The generalized Robinson-Schensted algorithm on the affine Weyl group of type A_{n-1}. Journal of Algebra **139**(2), 364–394 (1991)
20. Shimozono, M.: Crystals for Dummies. http://www.aimath.org/WWN/kostka/crysdumb.pdf, (2005)
21. Stanley, R. P.: Enumerative combinatorics. Vol. 2, volume 62 of Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin
22. Stembridge, J.R.: More W-graphs and cells: Molecular components and cell synthesis. Palo Alto. The Atlas of Lie Groups, In Notes from the AIM workshop (2008)
23. Viennot, G.: Une forme géométrique de la correspondance de Robinson-Schensted. Combinatoire et représentation du groupe symétrique, pages 29–58, (1977)
24. Vogan, D.A., Jr.: A generalized τ-invariant for the primitive spectrum of a semisimple Lie algebra. Math. Ann. **242**(3), 209–224 (1979)