Séminaire Lotharingien de Combinatoire **78B** (2017)
Article #10, 12 pp.

Proceedings of the 29th Conference on Formal Power Series and Algebraic Combinatorics (London)

Sweeping up Zeta

Hugh Thomas[1] and Nathan Williams[2]

[1]*LaCIM, Université du Québec à Montéal, Montréal (Québec), Canada*
[2]*University of California, Santa Barbara, Santa Barbara, California, USA*

Abstract. We repurpose the main theorem of [Thomas and Williams, 2014] to prove that modular sweep maps are bijective. We conclude that the general sweep maps defined in [Armstrong, Loehr, and Warrington, 2014] are bijective. As a special case of particular interest, this gives the first proof that the zeta map on rational Dyck paths is a bijection.

Résumé. Nous adaptons le théorème principal de [Thomas et Williams 2014] pour démontrer qu'une version modulaire des applications au balai (« sweep maps ») est bijective. Nous déduisons que les applications au balai générales de [Armstrong, Loehr et Warrington, 2014] sont bijectives. Comme cas d'intérêt particulier, cela donne la première démonstration que l'application zeta sur les chemins de Dyck rationaux est une bijection.

Keywords: zeta map, sweep map, rational Catalan combinatorics, affine Weyl groups

1 Introduction

The sweep map of [2] is a broad generalization of the zeta map on Dyck paths, originally defined by J. Haglund and M. Haiman in the context of the study of diagonal harmonics. Proving bijectivity of the sweep map was an open problem with significant implications in the study of rational Catalan combinatorics (see Section 2). We solve this problem in Theorem 5.1.

We let $m, N \in \mathbb{N}$, and we write $\mathcal{A}$ for the set of words of length N on the alphabet $\{0, 1, 2, \ldots, m-1\}$. For a word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_1 \cdots \mathsf{w}_N \in \mathcal{A}$ and for $1 \leq j \leq N$, define the *modular level* of the letter w_j to be $\ell_j := \sum_{i=1}^{j} \mathsf{w}_i \bmod m$.

The *modular sweep map* is the function $\mathsf{sweep}_m : \mathcal{A} \to \mathcal{A}$ that sorts $\mathsf{w} \in \mathcal{A}$ according to its modular levels as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = m-1, \ldots, 2, 1, 0$, read w from right to left and append to u all letters w_j whose level ℓ_j is equal to k. Define $\mathsf{sweep}_m(\mathsf{w}) := \mathsf{u}$.

Example 1.1. Let $m = 5$ and $N = 7$. We compute the modular levels of the word $\mathsf{w} = 3113214 \in \mathcal{A}$ by summing the initial letters of w modulo m and obtain the image

This is an extended abstract, outlining the main results in the preprint [18].

$\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$ by sorting according to the levels (and then discarding the information about the levels).

$$\begin{array}{cccccccc} \ell: & 3 & 4 & 0 & 3 & 0 & 1 & 0 \\ \mathsf{w}: & 3 & 1 & 1 & 3 & 2 & 1 & 4 \end{array} \xmapsto[\mathsf{sweep}_m]{} \begin{array}{cccccccc} \ell: & 4 & 3 & 3 & 1 & 0 & 0 & 0 \\ \mathsf{u}: & 1 & 3 & 3 & 1 & 4 & 2 & 1 \end{array}.$$

Our main result—proven in Section 4.3—is that sweep_m is invertible.[1]

Theorem 1.2. *The modular sweep map is a bijection $\mathcal{A} \to \mathcal{A}$.*

The remainder of this abstract is organized as follows. We give a brief history in Section 2 by recalling the different contexts in which the modular sweep map has appeared. In Section 3.1, we define the modular presweep map. This map differs from the modular sweep map in that it preserves the additional information of the modular levels. It is easy to invert the modular presweep map, as described in Section 3.2; partitions for which the inverse modular presweep map concludes are called *successful partitions*.

In Section 4.1, we introduce the notion of *equitable partitions* and show that a successful partition is equitable. Using an algorithm communicated to us by F. Aigner, C. Ceballos, and R. Sulzgruber, we construct the rightmost equitable partition in Theorem 4.4. (For the purposes of this abstract, we have preferred to use this algorithm rather than our original algorithm, which is roughly dual to it.) Theorem 4.6 concludes that the rightmost equitable partition and successful partition are the same.

We apply the results of Sections 3 and 4 to prove Theorem 1.2—that the modular sweep map is a bijection—in Section 4.3. In Section 5, we use Theorem 1.2 to invert the sweep map of [2] on words with letters in $\mathbb{Z}$ (rather than $\mathbb{Z}/m\mathbb{Z}$), and we conclude that the zeta map is bijective on Dyck paths and rational Dyck paths.

Acknowledgements

The second author is indebted to Vic Reiner and Dennis Stanton for originally suggesting this area.

We thank Drew Armstrong, Cesar Ceballos, Adriano Garsia, Nick Loehr, and Greg Warrington for their encouragement and enthusiasm, close reading, detection of typos, and suggestions for improvements to the exposition. We thank Marko Thiel for pointing out that the argument of [2, Proposition 3.2] could be extended to arbitrary alphabets in Theorem 5.2. We warmly thank Florian Aigner, Cesar Ceballos, and Robin Sulzgruber for finding Algorithm 2 in the course of their study of an early draft of this paper, and then patiently explaining it to us at FPSAC 2016. We would also like to thank the referees

[1] As G. Warrington pointed out to us at the American Institute of Mathematics in 2012—sorting is not usually an invertible operation!

of this extended abstract and of [18] for their helpful suggestions which improved the exposition.

This research was supported by NSERC and the Canada Research Chairs program.

2 History: Diagonal Harmonics and the Zeta Map

In their study of the space $\mathcal{DH}_n$ of diagonal harmonics [9], A. Garsia and M. Haiman defined a rational function $C_n(q,t)$, symmetric in q and t, with the property that $C_n(1,1) = \frac{1}{n+1}\binom{2n}{n}$. They conjectured that $C_n(q,t)$ was actually a *polynomial* in q and t with nonnegative coefficients—specializing one of the statistics to 1, they gave a combinatorial interpretation of this polynomial using the area statistic on *n-Dyck paths* (lattice paths from $(0,0)$ to (n,n) that stay above the diagonal $y = x$):

$$C_n(q,1) = C_n(1,q) = \sum_{w \text{ an } n\text{-Dyck path}} q^{\mathrm{area}(w)}.$$

The search was on to find a statistic that manifested nonnegativity—an unknown statistic with the property that

$$C_n(q,t) = \sum_{w \text{ an } n\text{-Dyck path}} q^{\mathrm{area}(w)} t^{\mathrm{unknown}(w)}.$$

In [13], "after a prolonged investigation of tables of $C_n(q,t)$," J. Haglund invented the idea of a *bounce path*, which he used to propose exactly such a statistic. Garsia and Haglund subsequently used these ideas to prove nonnegativity of $C_n(q,t)$ in [8].

As the legend goes, Garsia sent a cryptic email to Haiman announcing Haglund's discovery—without providing any specifics as to what the statistic was. Shortly after, Haiman announced that he, too, had produced the desired statistic.[2] Remarkably, Haglund's statistic and Haiman's statistic were *different*. In modern language, Haglund's statistic is known as bounce, while Haiman's is dinv. Haiman and Haglund quickly developed a bijection from n-Dyck paths to themselves—the *zeta map* ζ (see Section 5) [1, 13]—such that

$$(\mathrm{area}(w), \mathrm{bounce}(w)) = (\mathrm{dinv}(\zeta(w)), \mathrm{area}(\zeta(w))).$$

As Dyck paths have been generalized (say, as in Section 5, to lattice paths from $(0,0)$ to (a,b) that stay above the main diagonal), so too have these zeta maps [15, 7, 11, 14]. A modern perspective is that there is only one statistic—area—along with a generalized zeta map [2]. If such a zeta map is bijective on a set of generalized Dyck paths $\mathcal{D}$, one

[2] Garsia subsequently expressed regret that he didn't send this email to Haiman several years earlier.

can combinatorially define polynomials

$$\mathcal{D}(q,t) := \sum_{w \in \mathcal{D}} q^{\mathrm{area}(\mathsf{w})} t^{\mathrm{area}(\zeta(w))},$$

so that (by construction) $\mathcal{D}(q,1) = \mathcal{D}(1,q)$. Surprisingly, these polynomials also often happen to be symmetric in q and t.

Proving invertibility of these generalized zeta maps has been a traditionally difficult problem; combinatorially proving (q,t)-symmetry has been intractable. Most recently, D. Armstrong, N. Loehr, and G. Warrington have found a very general version of the zeta map, which they called *sweep maps* [2, Section 3.4].

3 Presweeping and Its Inverse

We will factor the modular sweep map as the composition of two maps: the modular presweep map and the forgetful map. In this section, we define the modular presweep map and its inverse.

Adhering to the notation in [17], we prefer to think of the modular levels from the introduction as partitioning the word u into blocks. Define a *partitioned word* for $\mathsf{u} \in \mathcal{A}$ to be a partition u^* of u into m words $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0$—where we use the *block divider* symbol $|$ to separate the blocks—so that $\mathsf{u} = \mathsf{u}^*_{m-1}\cdots\mathsf{u}^*_0$ is their concatenation. We call the word u^*_k the kth *block* and, with apologies to the combinatorics of words community, we write $\mathcal{A}^*$ for the set of all partitioned words of $\mathcal{A}$. We call u the *underlying word* of the partitioned word u^*. We may use either the symbol $\cdot$ or $\emptyset$ to denote an empty block. If the i-th letter u_i of u belongs to the k-th block u^*_k in the partitioned word u^*, we let $\mathsf{block}(\mathsf{u}^*, i) := k$. We fix the notation $|\mathsf{u}| := \sum_{i=1}^{N} \mathsf{u}_i$ and $|\mathsf{u}|_m = \ell_N = |\mathsf{u}| \bmod m$.

3.1 The Modular Presweep Map

The *modular presweep map* is the function $\mathsf{presweep}_m : \mathcal{A} \hookrightarrow \mathcal{A}^*$ that sorts $\mathsf{w} \in \mathcal{A}$ into blocks according to its levels. Precisely, for $k = m-1, \ldots, 2, 1, 0$, first initialize $\mathsf{u}^* := \cdot|\cdot|\cdots|\cdot$ to be the empty partitioned word, then read w from right to left and append to u^*_k all letters w_j whose level $\ell_j = \left(\sum_{i=1}^{j} \mathsf{w}_i \bmod m\right)$ is equal to k. In other words, u^*_k is obtained by extracting all letters of level k from u and reversing their relative order.

Example 3.1. As in Example 1.1, let $m = 5$, $N = 7$, and $\mathsf{w} = 3113214 \in \mathcal{A}$. We compute the modular levels of a word $\mathsf{w} \in A$ by summing the initial letters of w modulo m (below, left). We compute the modular presweep of w by sorting by levels, reading w from right to left. Placing letters with the same level in a block, we obtain the corresponding partitioned word $\mathsf{u}^* := \mathsf{presweep}_m(\mathsf{w})$ in $\mathcal{A}^*$ (below, right).

$$\begin{array}{rccccccc} \ell: & 3 & 4 & 0 & 3 & 0 & 1 & 0 \\ \mathsf{w}: & 3 & 1 & 1 & 3 & 2 & 1 & 4 \end{array} \xmapsto[\mathsf{presweep}_m]{} \begin{array}{rc|cc|c|c|ccc} \ell: & 4 & 3 & 3 & 2 & 1 & 0 & 0 & 0 \\ \mathsf{u}^*: & 1 & 3 & 3 & \cdot & 1 & 4 & 2 & 1 \end{array}$$

3.2 The Inverse Modular Presweep Map

The *inverse modular presweep map* is the function $\mathsf{inverse\_presweep}_m : \mathcal{A}^* \to \mathcal{A}$ such that

$$\mathsf{inverse\_presweep}_m \circ \mathsf{presweep}_m = \mathsf{id}_{\mathcal{A}},$$

where $\mathsf{id}_{\mathcal{A}}(\mathsf{w}) = \mathsf{w}$ is the identity function on $\mathcal{A}$. As explained in [2, Section 5.2] and in [17, Algorithm 2 and Figure 8], if we know how to associate the correct levels to $\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$, it is easy to reconstruct w.

Suppose we have the partitioned word $\mathsf{u}^* := \mathsf{presweep}_m(\mathsf{w})$. Since the letters of w were just rearranged to make u^*, we can determine $\ell_N = |\mathsf{w}|_m = |\mathsf{u}|_m$ from u^*. As we swept w from right to left, the last letter of w is therefore the first letter in $\mathsf{u}^*_{\ell_N}$. Remove this letter from u^*. Subtracting this letter from ℓ_N gives ℓ_{N-1}, and we obtain the $(N-1)$st letter of w as the first remaining letter in block ℓ_{N-1}. In general, for $i = 1, 2, \ldots, N$ we have already computed ℓ_{N-i+1}; subtracting the leftmost remaining letter in $\mathsf{u}^*_{\ell_{N-i+1}}$ from ℓ_{N-i+1} (and removing it from $\mathsf{u}^*_{\ell_{N-i+1}}$) gives ℓ_{N-i}. Pseudo-code for $\mathsf{inverse\_presweep}_m$ is given in Algorithm 1.

Input: A partitioned word $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0 \in \mathcal{A}^*$.
Output: A word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2\cdots\mathsf{w}_N \in \mathcal{A}$ or a subword of u^*.
Let $\ell_N := \left(\sum_{j=1}^N \mathsf{u}_j \bmod m\right)$ and $\mathsf{w} := \emptyset$;
for $i = 1$ **to** N **do**
  if $\mathsf{u}^*_{\ell_{N-i+1}} \neq \emptyset$ **then**
    Remove the first letter of $\mathsf{u}^*_{\ell_{N-i+1}}$ and assign it to w_{N-i+1};
    Prepend w_{N-i+1} to w;
    Let $\ell_{N-i} := (\ell_{N-i+1} - \mathsf{w}_{N-i+1} \bmod m)$;
  end
  else
    Return u^*
  end
end
Return w;

Algorithm 1: $\mathsf{inverse\_presweep}_m : \mathcal{A}^* \hookrightarrow \mathcal{A}$.

We say that Algorithm 1 *succeeds* on a partitioned word u^* if it returns an element of $\mathcal{A}$, and we say that it *fails* if it returns an element of $\mathcal{A}^*$. We call a partitioned word u^*

on which Algorithm 1 succeeds a *successful partition* of the underlying word u. Since Algorithm 1 undoes $\mathsf{presweep}_m$ one step at a time, we conclude that $\mathsf{inverse\_presweep}_m$ is the left inverse of $\mathsf{presweep}_m$.

Example 3.2. To reverse Example 3.1, we compute as follows. We start with the partitioned word u^*:

$$\begin{array}{c|c|cc|c|c|ccc} \ell & 4 & 3 & 3 & 2 & 1 & 0 & 0 & 0 \\ \mathsf{u}^* & 1 & 3 & 3 & \cdot & 1 & 4 & 2 & 1 \end{array}.$$

We find $\ell_N = |\mathsf{u}|_m = 0$. Then we iterate:

i	u^*	ℓ_{N-i}	w
0	$1\|33\|\cdot\|1\|\dot{4}21$	0	$\cdot$
1	$1\|33\|\cdot\|\dot{1}\|421$	1	4
2	$1\|33\|\cdot\|1\|4\dot{2}1$	0	14
3	$1\|\dot{3}3\|\cdot\|1\|421$	3	214
4	$1\|33\|\cdot\|1\|42\dot{1}$	0	3214
5	$\dot{1}\|33\|\cdot\|1\|421$	4	13214
6	$1\|3\dot{3}\|\cdot\|1\|421$	3	113214
7	$1\|33\|\cdot\|1\|421$	0	3113214

Comparing with Example 3.1, we see that we have recovered w.

3.3 Forgetting

We now obtain the modular sweep map from the modular presweep map by forgetting the information of the blocks. The *forgetful map* is the function

$$\mathsf{forget} : \mathcal{A}^* \to \mathcal{A}$$
$$\mathsf{forget}\left(\mathsf{u}^*_{m-1}|\mathsf{u}^*_{m-2}|\cdots|\mathsf{u}^*_0\right) = \mathsf{u}^*_{m-1}\mathsf{u}^*_{m-2}\cdots\mathsf{u}^*_0$$

obtained by concatenating all the blocks of $\mathsf{u}^* \in \mathcal{A}^*$. Thus, the modular sweep map of Section 1 may be written as the composition

$$\mathsf{sweep}_m = (\mathsf{forget} \circ \mathsf{presweep}_m) : \mathcal{A} \to \mathcal{A}.$$

Example 3.3. Continuing with Example 3.1, we forget the partitioning to obtain the modular sweep $\mathsf{u} := \mathsf{sweep}_m(\mathsf{w})$ of w to be

$$(\mathsf{forget} \circ \mathsf{presweep}_m)(3113124) = \mathsf{forget}\,(1|33|\cdot|1|421) = 1331421.$$

Thus, the problem of inverting the modular sweep map has been reduced to showing that there exists a unique successful partition $\mathsf{u}^* \in \mathcal{A}^*$ for each word $\mathsf{u} \in \mathcal{A}$. We do this in the next section.

4 Equitable Partitions and the Successful Partition

We already solved the problem of constructing the successful partition in [17], where we studied a composition

$$f \circ p,$$

where f is the map forget and p is a map *very slightly* different from $\mathsf{presweep}_m$. In particular, our notions here of a successful partition and the forgetful map coincide with those in [17].

4.1 Equitable Partitions

We expand a partitioned word u^* into an $N \times m$ *balancing array*

$$M^{\mathsf{u}^*} = (M^{\mathsf{u}^*}_{i,j})_{\substack{1 \le i \le N \\ m-1 \ge j \ge 0}}$$

defined by

$$M^{\mathsf{u}^*}_{i,j} := \begin{cases} \blacksquare & \text{if } j \in \{\mathsf{block}(\mathsf{u}^*, i), \mathsf{block}(\mathsf{u}^*, i) - 1, \ldots, \mathsf{block}(\mathsf{u}^*, i) - \mathsf{u}_i + 1\} \bmod m, \\ \cdot & \text{otherwise,} \end{cases}$$

Write $|\mathsf{u}| = \sum_{i=1}^{N} u_i = qm + r$ with $0 \le r < m$. We say that column j (for $m - 1 \ge j \ge 0$) of M^{u^*} is *equitably filled* if:

- $r \ge j \ge 1$ and column j has $q + 1$ copies of the symbol $\blacksquare$, or if
- $j = 0$ or $j > r$, and column j contains q copies of $\blacksquare$.

If a column has (strictly) fewer copies of the symbol $\blacksquare$ than it would to be equitably filled, we say it is *less than equitably filled*; similarly, when a column has (strictly) more copies of $\blacksquare$ we say that it is *more than equitably filled*. In particular, if $r = 0$, then every equitably filled column has q copies of $\blacksquare$. We say that u^* is an *equitable partition* if each of the columns of M^{u^*} is equitably filled.

The motivation for this definition is the following lemma.

Lemma 4.1. *Any successful partition* u^* *is an equitable partition.*

Proof. We can construct all successful partitions as follows [17, Definition 7.5]. Define an infinite complete m-ary tree $\mathcal{T}_m^*$ by

1. The zeroth rank consists of the empty successful partition u^*, given by $\mathsf{u}^*_k = \emptyset$ for $m - 1 \ge k \ge 0$.

2. The children of a successful partition $\mathsf{u}^* = \mathsf{u}^*_{m-1}|\dots|\mathsf{u}^*_0$ are the m successful partitions obtained by prepending $i \pmod m$ to $\mathsf{u}^*_{i+|\mathsf{u}|_m}$.

Then it is easy to see that all partitioned words in $\mathcal{T}^*_m$ are equitable, and that all successful partitions appear in $\mathcal{T}^*_m$ [17, Lemma 7.2]. (It is not yet clear that the images of the words in $\mathcal{T}^*_m$ under the forgetful map are actually distinct.) $\square$

Example 4.2. The equitable partitions $13|31|4|2|1$ and $1|33|\cdot|1|421$ have corresponding balancing arrays

i \ j	4	3	2	1	0
1	■	·	·	·	·
2	■	■	■	·	·
3	·	■	■	■	·
4	·	■	·	·	·
5	■	·	■	■	■
6	·	·	·	■	■
7	·	·	·	·	■

and

i \ j	4	3	2	1	0
1	■	·	·	·	·
2	·	■	■	■	·
3	·	■	■	■	·
4	·	·	·	■	·
5	■	■	■	·	■
6	■	·	·	·	■
7	·	·	·	·	■

.

Since $|\mathsf{u}| = 15 = 3 \cdot 5 + 0$, any equitable filling has three copies of ■ in each column j.

4.2 The Rightmost Partition

Given u, we first prove the existence of a particular equitable partition.

Definition 4.3. A *rightmost equitable partition* is an equitable partition u^* such that any other equitable partition v^* has $\mathsf{block}(\mathsf{v}^*, i) \geq \mathsf{block}(\mathsf{u}^*, i)$ for all i.

Theorem 4.4. *Any* $\mathsf{u} \in \mathcal{A}$ *admits a unique rightmost partition* $\mathsf{rightmost}(\mathsf{u})$.

Proof Outline. We claim that Algorithm 2 constructs the unique rightmost equitable partition.[3] One first shows that Algorithm 2 does not attempt any illegal moves, and so returns an equitable partition. One then shows that Algorithm 2 outputs the unique rightmost equitable partition u^*. $\square$

Example 4.5. We illustrate Algorithm 2 applied to the word $\mathsf{u} = 1331421$. At each step, the rightmost column with less than its equitable filling is highlighted.

[3] Algorithm 2 was communicated to us by F. Aigner, C. Ceballos, and R. Sulzgruber.

Input: A word $\mathsf{u} \in \mathcal{A}$.
Output: The rightmost equitable partition $\mathsf{u}^* \in \mathcal{A}^*$.
Set $\mathsf{u}^* = \cdot|\cdot|\cdots|\mathsf{u}$;
while u^* *is not an equitable partition* **do**
  Let j be the rightmost column of M^{u^*} that is less than equitably filled;
  Delete the leftmost letter of u^*_{j-1} and append it to u^*_j;
end
Return u^*;

Algorithm 2: $\mathsf{rightmost} : \mathcal{A} \to \mathcal{A}^*$.

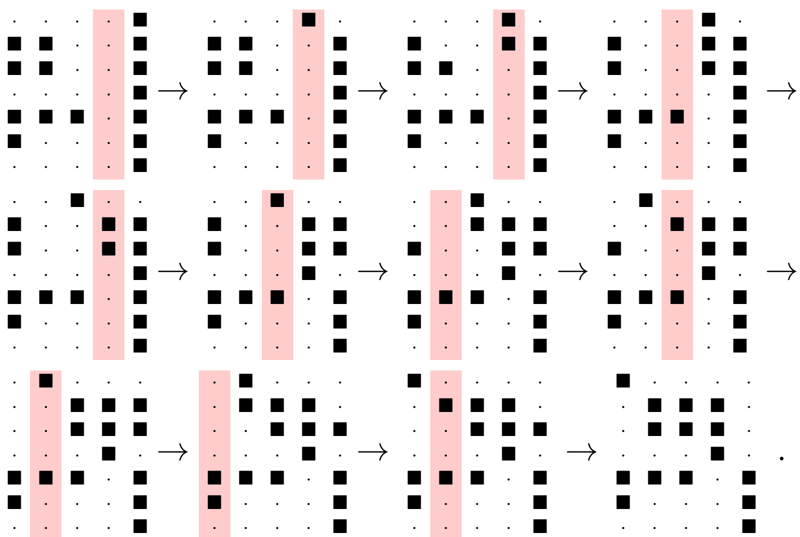

Thus, the rightmost equitable partition of u is $\mathsf{u}^* = 1|33|\cdot|1|421$.

4.3 The Successful Partition and Inverting the Modular Sweep Map

We can now state the following theorem:

Theorem 4.6. *For* $\mathsf{u} \in \mathcal{A}$, *the rightmost equitable partition* $\mathsf{rightmost}(\mathsf{u})$ *is the unique successful partition of* u.

Proof. Apply Algorithm 1 to $\mathsf{rightmost}(\mathsf{u})$. Suppose it does not succeed. Leave all the letters that were visited in place, and shift all the other letters right one block. One checks that this yields an equitable partition which is further to the right than $\mathsf{rightmost}(\mathsf{u})$. The key point here is that whenever Algorithm 1 finishes, the copies of ■ corresponding to the *remaining* letters are equally distributed among the columns. The existence of an equitable partition obtained by moving letters of $\mathsf{rightmost}(\mathsf{u})$ to the right contradicts the fact that $\mathsf{rightmost}(\mathsf{u})$ is rightmost, so our assumption that Algorithm 1 did not succeed must have been wrong.

We conclude uniqueness of the successful partition using the cardinality argument from [17]. We have already shown that every word in $\mathcal{A}$ has a successful partition. Since the tree $\mathcal{T}_m^*$ in the proof of Lemma 4.1 contains every successful partition of words in $\mathcal{A}$, and since its Nth level has m^N elements, we conclude that every word in $\mathcal{A}$ has a unique successful partition. □

From this, the main theorem follows.

Theorem 1.2. *The modular sweep map is a bijection $\mathcal{A} \to \mathcal{A}$.*

Proof. We have inverted $\mathsf{presweep}_m$ in Section 3.2, forget in Theorem 4.6, and the modular sweep map may be written as the composition

$$\mathsf{sweep}_m = \mathsf{forget} \circ \mathsf{presweep}_m : \mathcal{A} \to \mathcal{A}.$$ □

5 Applications

By taking m sufficiently large, the modular sweep map emulates the sweep map introduced in [2], as we now explain.

Fix $a := (a_1, \ldots, a_n) \in \mathbb{Z}^n$, let $e := (e_1, \ldots, e_n) \in \mathbb{N}^n$, and define $\mathcal{A}_{\mathbb{Z}}$ to be the set of words containing e_j copies of a_j for $1 \leq j \leq n$. For a word $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2 \cdots \mathsf{w}_N \in \mathcal{A}_{\mathbb{Z}}$, define the *level* of w_j to be the integer $\ell_j := \sum_{i=1}^{j} \mathsf{w}_i$ for $1 \leq j \leq N$.

The *sweep map* is the function $\mathsf{sweep} : \mathcal{A}_{\mathbb{Z}} \to \mathcal{A}_{\mathbb{Z}}$ that sorts $\mathsf{w} \in \mathcal{A}_{\mathbb{Z}}$ according to its levels as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = -1, -2, -3, \ldots$ and then $k = \ldots, 3, 2, 1, 0$, read w from right to left and append to u all letters w_j whose level ℓ_j is equal to k. Define $\mathsf{sweep}(\mathsf{w}) := \mathsf{u}$.

Theorem 5.1 ([2, Conjecture 3.3 (a)])**.** *The sweep map is a bijection $\mathcal{A}_{\mathbb{Z}} \to \mathcal{A}_{\mathbb{Z}}$.*

Proof. Since the modular sweep map only *permutes* its input, it restricts to a bijection on words with a specified content. We claim that by choosing m large enough, the modular sweep map agrees with the sweep map when the letters a_j with multiplicities e_j are replaced by their natural representatives $a_j \pmod{m}$ in $\{0, 1, 2, \ldots, m-1\}$ (and all other elements are given multiplicity 0), and similarly for the levels ℓ_j. □

Let $\mathcal{A}_{\mathbb{N}} \subseteq \mathcal{A}_{\mathbb{Z}}$ be the subset of words in $\mathcal{A}_{\mathbb{Z}}$ whose levels are all nonnegative. Following [2], we call $\mathcal{A}_{\mathbb{N}}$ the set of *Dyck words*. An argument generalizing [2, Proposition 3.2], suggested to us by M. Thiel, allows one to deduce that sweeping is also bijective on Dyck words.

Theorem 5.2 ([2, Conjecture 3.3 (b)])**.** *The sweep map is a bijection $\mathcal{A}_{\mathbb{N}} \to \mathcal{A}_{\mathbb{N}}$.*

Finally, we consider the special case of Dyck words for an alphabet $\{a, b\}$ of size $n = 2$, such that $a > 0$ and $b < 0$ and where the letter a occurs $-b$ times and the letter b occurs a times. We shall write this set of Dyck words as $\mathcal{D}_{a,b}$—these paths are of fundamental importance for the study of rational (type A) Catalan combinatorics [4, 3, 12, 5, 6, 16, 19, 10].

By [2, Table 1] and [2, Theorem 4.8, Lemma 4.10, Theorem 4.12], the *zeta map* may be defined as a variant of the sweep map $\zeta : \mathcal{D}_{a,b} \to \mathcal{D}_{a,b}$ that sorts $\mathsf{w} \in \mathcal{D}_{a,b}$ as follows: initialize $\mathsf{u} = \emptyset$ to be the empty word. For $k = 0, 1, 2, \ldots$ and then $k = \ldots, -3, -2, -1$, read w from left to right and append to u all letters w_j whose level ℓ_j is equal to k. Define $\zeta(\mathsf{w}) := \mathsf{u}$.

We have the following corollary of Theorem 5.1, which is of independent interest.

Corollary 5.3 (Zeta for Rational Dyck Paths)**.** *The zeta map is a bijection $\mathcal{D}_{a,b} \to \mathcal{D}_{a,b}$.*

Proof. For $\mathsf{w} = \mathsf{w}_1\mathsf{w}_2 \cdots \mathsf{w}_N$, let $\mathsf{rev}(\mathsf{w}) := \mathsf{w}_N \cdots \mathsf{w}_2\mathsf{w}_1$ and $-\mathsf{w} := (-\mathsf{w}_1)(-\mathsf{w}_2)\cdots(-\mathsf{w}_N)$. Then the zeta map may be computed as

$$\zeta(\mathsf{w}) = -\left(\mathsf{rev} \circ \mathsf{sweep} \circ \mathsf{rev}\right)(-\mathsf{w}).$$

Since sweep is a bijection, we conclude that ζ is a bijection. □

References

[1] G. Andrews, C. Krattenthaler, L. Orsina, and P. Papi. "*ad*-nilpotent $\mathfrak{b}$-ideals in $sl(n)$ having a fixed class of nilpotence: combinatorics and enumeration". *Trans. Amer. Math. Soc.* **354** (2002), pp. 3835–3853. DOI.

[2] D. Armstrong, N. Loehr, and G. Warrington. "Sweep maps: A continuous family of sorting algorithms". *Adv. Math.* **284** (2015), pp. 159–185. DOI.

[3] D. Armstrong, N. Loehr, and G. Warrington. "Rational Parking Functions and Catalan Numbers". *Ann. Combin.* **20** (2016), pp. 21–58. DOI.

[4] D. Armstrong, B. Rhoades, and N. Williams. "Rational Associahedra and Noncrossing Partitions". *Electronic J. Combin.* **20** (2013), Art. P54. URL.

[5] M. Bodnar and B. Rhoades. "Cyclic Sieving and Rational Catalan Theory". *Electronic J. Combin.* **23** (2016), Art. P2.4. URL.

[6] C. Ceballos, T. Denton, and C. Hanusa. "Combinatorics of the zeta map on rational Dyck paths". *J. Combin. Theory Ser. A* **141** (2016), pp. 33–77. DOI.

[7] E. Egge, J. Haglund, K. Killpatrick, and D. Kremer. "A Schröder generalization of Haglund's statistic on Catalan paths". *Electronic J. Combin.* **10** (2003), Art. R16. URL.

[8] A. Garsia and J. Haglund. "A proof of the q,t-Catalan positivity conjecture". *Discrete Math.* **256** (2002), pp. 677–717. DOI.

[9] A. Garsia and M. Haiman. "A remarkable q, t-Catalan sequence and q-Lagrange inversion". *J. Algebraic Combin.* **5** (1996), pp. 191–244. DOI.

[10] A. Garsia and G. Xin. "Inverting the Rational Sweep Map". 2016. arXiv:1602.02346.

[11] E. Gorsky and M. Mazin. "Compactified Jacobians and q, t-Catalan numbers, II". *J. Algebraic Combin.* **39** (2014), pp. 153–186. DOI.

[12] E. Gorsky, M. Mazin, and M. Vazirani. "Affine permutations and rational slope parking functions". *Trans. Amer. Math. Soc.* **368** (2016), pp. 8403–8445. DOI.

[13] J. Haglund. "Conjectured statistics for the q, t-Catalan numbers". *Adv. Math.* **175** (2003), pp. 319–334. DOI.

[14] K. Lee, L. Li, and N. Loehr. "Combinatorics of certain higher q, t-Catalan polynomials: chains, joint symmetry, and the Garsia–Haiman formula". *J. Algebraic Combin.* **39** (2014), pp. 749–781. DOI.

[15] N. Loehr. "Multivariate analogues of Catalan numbers, parking functions, and their extensions". PhD thesis. University of California, San Diego, 2003.

[16] R. Sulzgruber. "Rational Shi tableaux and the skew length statistic". 2015. arXiv:1512.04320.

[17] H. Thomas and N. Williams. "Cyclic symmetry of the scaled simplex". *J. Algebraic Combin.* **39** (2014), pp. 225–246. DOI.

[18] H. Thomas and N. Williams. "Sweeping up Zeta". 2016. arXiv:1512.01483.

[19] G. Xin. "An efficient search algorithm for inverting the sweep map on rational Dyck paths". 2015. arXiv:1505.00823.