

Inverting the rational sweep map

ADRIANO GARSIA* AND GUOCE XIN†

We present a simple algorithm for inverting the sweep map on rational (m, n) -Dyck paths for a co-prime pair (m, n) of positive integers. This work is inspired by Thomas-Williams work on the modular sweep map. A simple proof of the validity of our algorithm is included.

MSC 2010 SUBJECT CLASSIFICATIONS: 05A19, 05A99, 05E05.

KEYWORDS AND PHRASES: Rational Dyck paths, sweep maps, path diagrams.

1. The algorithm

Inspired by the Thomas-William algorithm [4] for inverting the general modular sweep map, we find a simple algorithm to invert the sweep map for rational Dyck paths. The fundamental fact that made it so difficult to invert the sweep map in this case is that all previous attempts used only the ranks of the vertices of the rational Dyck paths. Moreover the geometry of rational Dyck paths was not consistent with those ranks.

A single picture will be sufficient here to understand the idea. In what follows, we always denote by (m, n) a co-prime pair of positive integers, South end (by letter S) for the starting point of a North step and West end (by letter W) for the starting point of an East step, unless specified otherwise. This is convenient and causes no confusion because we usually talk about the starting points of these steps.

Figure 1 illustrates a rational (m, n) -Dyck path \overline{D} for $(m, n) = (7, 5)$ and its sweep map image D on its right. Recall that the ranks of the starting vertices of an (m, n) -Dyck path \overline{D} are recursively computed starting with rank 0, and adding m after a North step and subtracting n after an East step as shown in Figure 1. To obtain the sweep image D of \overline{D} , we let the

*The first named author is grateful to Nathan Williams for the time and effort that he spent to communicate his pioneering proof of invertibility of the general sweep map. The first named author was supported by NFS grant DMS13–62160.

†The second named author was supported by National Natural Science Foundation of China (11171231).




Figure 1: A rational $(7, 5)$ -Dyck path and its sweep map image.




Figure 2: Transformation of the $(7, 5)$ -Dyck paths in Figure 1.

main diagonal (with slope n/m) sweep from right to left and successively draw the steps of D as follows: i) draw a South end (and hence a North step) when we sweep a South end of \bar{D} ; ii) draw a West end (hence an East step) when we sweep a West end of \bar{D} . The steps of D can also be obtained by rearranging the steps of \bar{D} by increasing ranks of their starting vertices.

The sweep map has become an active subject in the recent 15 years. Variations and extensions have been found, and some classical bijections turn out to be the disguised version of the sweep map. See [1] for detailed information and references.

The open problem was the reconstruction of the path on the left from the path on the right. The idea that leads to the solution of this problem is to draw these two paths as in Figure 2.

That is we first stretch all the arrows so that their lengths correspond to the effect they have on the ranks of the vertices of the path then add an appropriate clockwise rotation to obtain the two path diagrams in Figure 2. The path diagrams are completed by writing an S for each South end in our original path and a W for each West end. On the left we have added a list of each level. The ranks of \overline{D} become visually the levels of the starting points of the arrows. On the right, at each level we count the red (solid) segments and the blue (dashed)¹ segments which traverse that level and record their difference. Of course these differences, called *row counts* or (signed) *row sum*, turn out to be all equal to 0, for obvious reasons. This will be referred to as the *0-row-count property*. Theorem 3 states that this is a characteristic property of rational Dyck paths, which becomes evident when paths are drawn in this manner. This fact is conducive to the discovery of our algorithm for constructing the pre-image of any (m, n) -Dyck path.

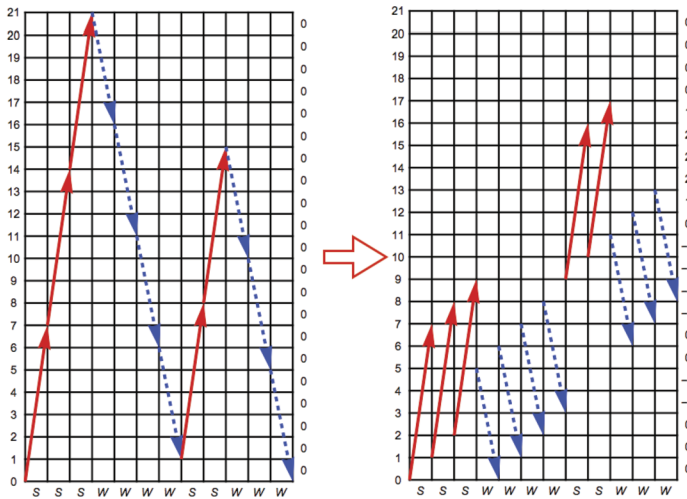


Figure 3: A given rational Dyck path and its starting path diagram on the right.

The first step in our algorithm is to vertically shift the arrows of the path on the left of Figure 3. The resulting path diagram is on the right, whose arrows have their starting ranks minimally strictly increasing. More

¹Suggested by the referee, we have drawn blue dashed arrows for convenience of black-white print. We will only use “red” and “blue” in our transformed Dyck paths, but in our context, red, solid, up and positive slope are equivalent; blue, dashed, down and negative slope are equivalent.

precisely the first three red arrows are lowered in their columns to start at levels 0, 1, 2. To avoid placing part of the first blue arrow below level 0 we lower it to start at level 5. This done all the remaining arrows are successively placed to start at levels 6, 7, 8, 9, 10, 11, 12, 13. Notice the row counts at the right of the resulting path diagram. Our aim is to progressively reduce them all to zeros, which are the row counts characterization of the path diagram we are working to reconstruct.

The miracle is that this can be achieved by a sequence of identical steps. More precisely, at each step of our algorithm we locate the lowest row sum that is greater than 0. We next notice that there is a unique arrow that starts immediately below that row sum. This done we move that arrow one unit upwards. However, to keep the ranks strictly increasing we also shift, when necessary, some of the successive arrows by one unit upwards. In this particular case our MATHEMATICA implementation of the resulting algorithm produced the sequence of 18 path diagrams in Figures 4 and 5. Notice, the green (thick) line has been added in each path diagram to make evident the height of the lowest positive row count. Of course each step ends with an updating of the row counts.

The final path diagram yields a path that is easily shown to be the desired pre-image. For example, in Figure 6, to obtain the left path from the middle balanced increasing path diagram, we simply start with the leftmost red arrow, and at each step we proceed along the arrow that starts at the rank reached by the previous arrow. Continue until all the arrows have been used. The reason why there always is an arrow that starts at each reached rank, is an immediate consequence of the 0-row-count property of the middle path diagram. Such an arrow is unique in our case, since (m, n) is a coprime pair. On the other hand, to obtain the sweep map image of the left path, we reorder the arrows according to their starting ranks, which corresponds to horizontally shift the arrows (without changing the ranks) so that their starting ranks are increasing. This gives the middle path diagram. Then we read the arrows (ignoring their starting ranks) from left to right. This gives the right path, as desired. This manner of drawing rational Dyck paths makes many needed properties more evident than the traditional manner and therefore also easier to prove. As a case in point, we give a simple visual way of establishing the following nontrivial result (see, e.g., [1]).

Lemma 1. *The sweep image of an (m, n) -Dyck path is an (m, n) -Dyck path.*

Proof. On the left of Figure 6 we have the final path yielded by our algorithm. To obtain the path diagram in the middle we simply rearrange the arrows (by horizontal shifts) so that their starting ranks are increasing. The path on the right is obtained by vertically shifting the successive arrows so

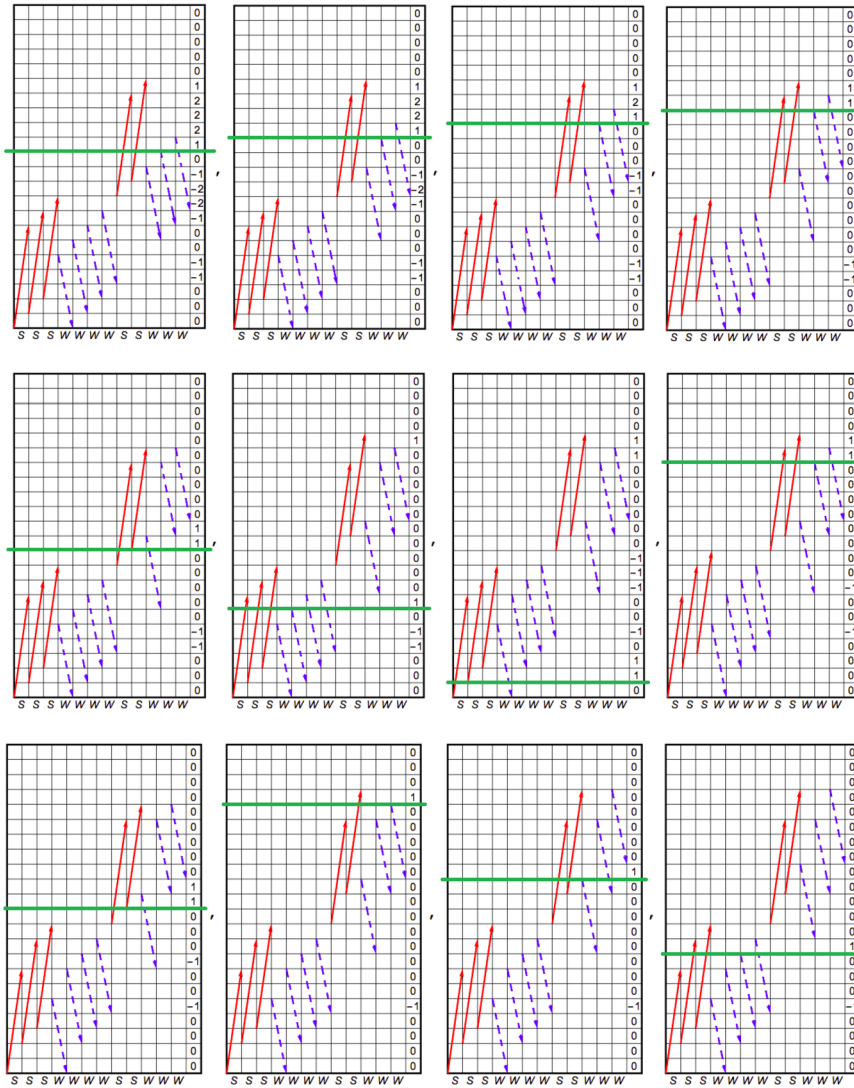


Figure 4: Part 1 of the 18 path diagrams that our algorithm produced.

that they concatenate to a path. To prove that the resulting path is a $(7, 5)$ -Dyck path, we need only show that the successive partial (signed) sums of the segments (i.e., a red segment is counted as 1 and a blue segment is counted as -1 .) of these arrows are all non-negative. This is a consequence of the 0-row-count property. In fact, for example, let us prove that the sum of the segments to the left of the vertical green line v is positive.

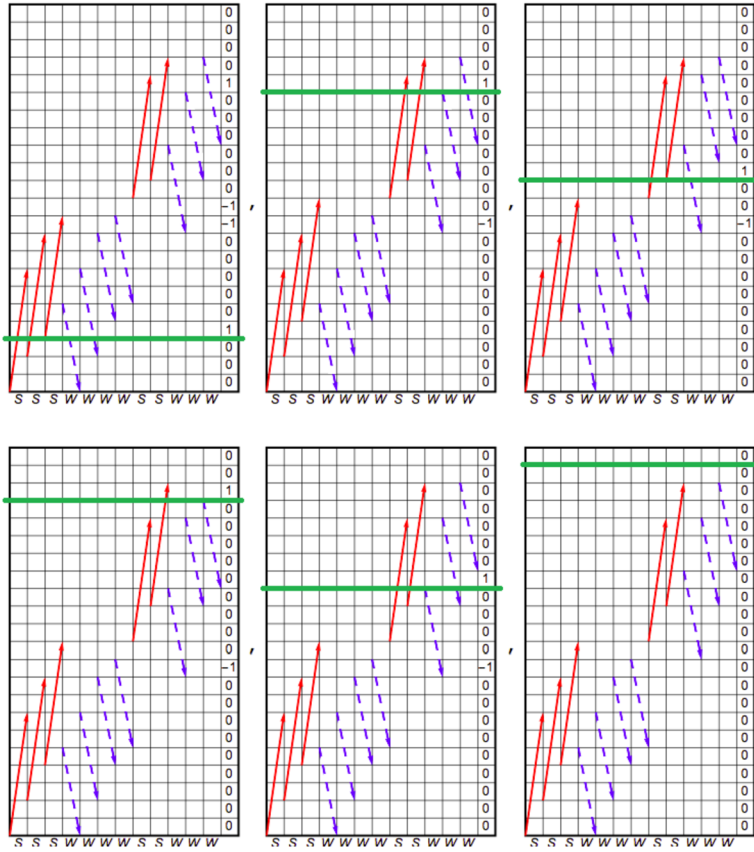


Figure 5: Part 2 of the 18 path diagrams that our algorithm produced.

To this end, let A be the arrow that starts on v and ℓ be its starting rank. Let h be the horizontal green (thick) line at level ℓ . Denote by L the region below h , and let L_1, L_2 be the left and right portions of L split by v . Let us also denote by $|rL_1|, |rL_2|$, the red arrow segment counts in the corresponding regions and by $|bL_1|, |bL_2|$ the corresponding blue segment counts. This given, since red segments contribute a 1 and a blue segment contributes -1 to the final count, it follows that

$$i) \quad |rL_1| + |rL_2| = |bL_1| + |bL_2|, \quad ii) \quad |rL_2| = 0.$$

In fact, $i)$ is due to the 0-row-count property and $ii)$ is simply due to the fact that all red arrows to the right of v must start above h . Thus we must have




Figure 6: A $(7, 5)$ -Dyck path on the left; by horizontal shifts of the arrows we obtain the middle path diagram whose starting ranks are increasing; then by vertical shifts of the arrows to obtain a $(7, 5)$ -Dyck path on the right picture.

$$|rL_1| - |bL_1| = |bL_2| \geq 0.$$

This implies that the sum of the arrows to the left of v must be ≥ 0 . \square

A proof of the validity of our algorithm may be derived from the Thomas-Williams result by letting their modulus tend to infinity. However, our algorithm deserves a more direct and simple proof.

Such a proof will be given in the following pages. This proof will be based on the validity of a simpler but less efficient algorithm. To distinguish the above algorithm from our later one, we will call them respectively the **StrongFindRank** and the **WeakFindRank** algorithms, or “strong” and “weak” algorithm for short.

The rest of the paper is organized as follows. Section 2 is devoted to the proof of the **WeakFindRank** algorithms. It also includes all the necessary concepts and concludes with Theorem 8, which asserts the invertibility of the rational sweep map. Theorem 9 is the main result of Section 3. It allows us to analyze the complexity of both the “strong” and the “weak” algorithms. It is also used in Section 4, where we show the validity of the “strong” algorithm. Finally, we discuss the difference between the Thomas-Williams algorithm and our algorithm in Section 5. We also talk about some future plans.




Figure 7: A path diagram for $m = 5, n = 7$ and $N = 21$.

2. The proof

2.1. Balanced path diagrams

A *path diagram* T consists of an ordered set of n red arrows and m blue arrows, placed on a $(m+n) \times N$ lattice rectangle, where N is a large positive integer to be specified. See Figure 7.

A red arrow is the up vector $(1, m)$ and a blue arrow is the down vector $(1, -n)$. The rows of lattice cells will be simply referred to as *rows* and the horizontal lattice lines will be simply referred to as *lines*. On the left of each line we have placed its y coordinate which we will simply refer to as its *level*. The level of the starting point of an arrow is called its *starting rank*, and similarly its *end rank* is the level of its end point. It will be convenient to call *row* i the row of lattice cells delimited by the lines at levels i and $i + 1$. Lattice columns are defined in a similar way.

Given a word Σ with n letters S and m letters W , and a sequence $R = (r_1, \dots, r_{m+n})$ of $n + m$ nonnegative ranks, the path diagram $T(\Sigma, R)$ is obtained by placing the letters of Σ at the bottom of the lattice columns and drawing in the i^{th} column an arrow with starting rank r_i and *red* (solid) if the i^{th} letter of Σ is $\Sigma_i = S$ or *blue* (dashed) if $\Sigma_i = W$. See Figure 7, where $\Sigma = SSSWWWSWW$ and $R = (0, 1, 2, 5, 6, 7, 8, 9, 10, 11, 12, 13)$. The sequence R will be called the *rank sequence* of the path diagram T . Notice that each lattice cell may contain a segment of a red arrow or a segment of a blue arrow or no segment at all. The red segment count of row




Figure 8: The difference $c(j) - c(j - 1)$ is 1 in the left two cases, is -1 in the right two cases, and is 0 in the previous cases.

j will be denoted $c^r(j)$ and the blue segment count is denoted $c^b(j)$. We will set $c(j) = c^r(j) - c^b(j)$ and refer to it as the *count of row j* . In the above display on the right of each row we have attached its row count. The following observation will be crucial in our development.

Lemma 2. *Let $T(\Sigma, R)$ be any path diagram. It holds for every integer $j \geq 1$ that*

$$(1) \quad c(j) - c(j - 1) = \#\{\text{arrows starting at level } j\} - \#\{\text{arrows ending at level } j\}.$$

Proof. Let us investigate the contribution to the difference $c(j) - c(j - 1)$ from a single arrow A . The contribution is 0 if i) A has no segments in rows j and $j - 1$, ii) A has segments in row j and $j - 1$. In both cases, it is clear that A cannot start nor end at level j . Thus the remaining cases are as listed in Figure 8. \square

It will be convenient to say that a path diagram $T(\Sigma, R)$ is *balanced* if all its row counts are equal to 0. The word Σ is said to be the (S, W) -word of a Dyck path D in $\mathcal{D}_{m,n}$, if it is obtained by placing an S when D takes a South end (hence a North step) and a W when D takes a West end (hence an East step).

Theorem 3. *Let Σ be the (S, W) -word of $D \in \mathcal{D}_{m,n}$, and let $R = (r_1, r_2, \dots, r_{m+n})$, with $r_1 = 0$, be a weakly increasing sequence of integers. Then R is a rearrangement of the rank sequence of a pre-image \overline{D} (regarded as a path diagram) of D under the sweep map, if and only if the path diagram $T(\Sigma, R)$ is balanced and the sequence R is strictly increasing.*

Proof. Suppose that \overline{D} is a pre-image of D . This given, let $T(\overline{\Sigma}, \overline{R})$ with $\overline{\Sigma}$ the (S, W) -word of \overline{D} , \overline{R} the rank sequence of \overline{D} , and height N chosen to be

a number greater than $nm + \max(\overline{R})$. It is clear that the arrows of $T(\overline{\Sigma}, \overline{R})$ can be depicted by starting at level 0 and drawing a red arrow $(1, m)$ every time \overline{D} takes a South end and a blue arrow $(1, -n)$ every time \overline{D} takes a West end, with each arrow starting where the previous arrow ended. It is obvious that the row counts of $T(\overline{\Sigma}, \overline{R})$ are all 0 since, in each row every red segment is followed by a blue segment. Now let $T(\Sigma, R)$ be the path diagram of same height N obtained by reordering the arrows by their starting ranks. This is achieved by horizontal shifts. Since the co-primality of (m, n) assures that \overline{R} has distinct components, the resulting R is the increasing rearrangement of \overline{R} by a unique permutation. Likewise, the word Σ is obtained by rearranging the letters of $\overline{\Sigma}$ by the same permutation, and it is the sweep image of $\overline{\Sigma}$. Thus we may say that the same permutation can be used to change $T(\overline{\Sigma}, \overline{R})$ into $T(\Sigma, R)$. Since this operation only permutes (or shifts) segments within each row, it follows that all the row counts of $T(\Sigma, R)$ must also be 0. This proves the necessity. See the left two pictures in Figure 6.

For the sufficiency, suppose that the path diagram $T(\Sigma, R)$ is balanced, with D the Dyck path whose word is Σ and R a weakly increasing sequence. Then by Lemma 2 it follows that for every level j , either i) no arrow starts or ends at this level, or ii) if $k > 0$ arrows end (start) at this level then exactly k arrows start (end) at this level. This given, we will construct a Dyck path \overline{D} by the following algorithm. Starting at level 0 we follow the first arrow, which we know is necessarily red and starts at level 0. This arrow ends at level m . Since there is at least one arrow that starts at this level follow the very next arrow that does. Proceeding recursively thereafter, every time we reach a level, we follow the very next arrow that starts at that level. This process stops when we are back at level 0, and we must since in Σ there are n S and m W . Let \overline{D} be the resulting path. Using the colors of the successive arrows of \overline{D} gives us the $\overline{\Sigma}$ word of \overline{D} . Now notice that \overline{D} must be a path in $\mathcal{D}_{m,n}$ since all its starting ranks are nonnegative due to the weakly increasing property of R and therefore they must necessarily be distinct by the co-primality of (m, n) . In particular, if \overline{R} denotes the sequence of starting ranks of \overline{D} we are also forced to conclude that its components are distinct. Since the components of \overline{R} are only a rearrangement of the components of R we deduce that R must have been strictly increasing to start with. This implies that D must be a Sweep map image of \overline{D} since the successive letters of Σ can be obtained by rearranging the letters of $\overline{\Sigma}$ by the same permutation that rearranges \overline{R} to R . This completes the proof of sufficiency. \square

This given, we can easily see that the validity of our “strong” algorithm hinges on establishing that it produces a balanced path diagram after a finite number of steps. Theorem 3 allows us to relax the strictly increasing




Figure 9: Shifting up one unit an arrow from level a to level b will decrease $c(a)$ by 1 and increase $c(b)$ by 1.

requirements on the rank sequences of the successive path diagrams produced by the algorithm. The `WeakFindRank` algorithm, defined below, has precisely that property. This results in a simpler proof of the termination property of both algorithms.

2.2. Algorithm `WeakFindRank` and the justification

Algorithm `WeakFindRank`

Input: A path diagram $T(\Sigma, R^{(0)})$ with Σ the word of a Dyck path $D \in \mathcal{D}_{\downarrow, \setminus}$, a weakly increasing sequence $R^{(0)} = (r_1^{(0)}, r_2^{(0)}, \dots, r_{m+n}^{(0)})$.

Output: A balanced path diagram $T(\Sigma, \bar{R})$.

It will be convenient to keep the common height equal to N for all the successive path diagrams constructed by the algorithm, where $N = U + 2mn$, with $U = \max(R^{(0)}) + m + 1$.

Step 1 Starting with $T(\Sigma, R^{(0)})$ repeat the following step until the resulting path diagram is balanced.

Step 2 In $T(\Sigma, R^{(s)})$, with $R^{(s)} = (r_1^{(s)}, r_2^{(s)}, \dots, r_{m+n}^{(s)})$, find the lowest row j with $c(j) > 0$ and find the rightmost arrow that starts at level j . Suppose that arrow starts at (i, j) . Move up the arrow one level to construct the path diagram $T(\Sigma, R^{(s+1)})$ with $r_i^{(s+1)} = r_i^{(s)} + 1$ and $r_{i'}^{(s+1)} = r_{i'}^{(s)}$ for all $i' \neq i$. If all the row counts are ≤ 0 then stop the algorithm, since all row counts must necessarily vanish.

Figure 9 shows that we are weakly reducing the number of rows with positive row counts in Step 2. It also makes the following key observation evident.

Lemma 4. *If at some point $c(k)$ becomes ≥ 0 then for ever after it will never become negative. In particular, since $c(k) = 0$ with $k > U$ for the*

initial path diagram $T(\Sigma, R^{(0)})$ we will have $c(k) \geq 0$ when $k > U$ for all successive path diagrams produced by the algorithm.

Proof. The lemma holds true because we only decrease a row count when it is positive. \square

We need some basic properties to justify the algorithm.

Lemma 5. *We have the following basic properties.*

- (i) *If row j is the lowest with $c(j) > 0$ then there is an arrow that starts at level j . In this situation, we say that we are working with row j .*
- (ii) *The successive rank sequences are always weakly increasing.*
- (iii) *If $T(\Sigma, R)$ has no positive row counts, then it is balanced. Consequently, if the algorithm terminates, the last path diagram is balanced.*

Proof.

- (i) By the choice of j , we have $c(j) > 0$ and $c(j - 1) \leq 0$. Thus $c(j) - c(j - 1) > 0$, which by Lemma 2, shows that there is at least one arrow starting at rank j .
- (ii) Our choice of i in step (2) assures that the next rank sequence remains weakly increasing.
- (iii) Since each of our path diagrams $T(\Sigma, R)$ has n red arrows of length m and m blue arrows of length n , the total sum of row counts of any $T(\Sigma, R)$ has to be 0. Thus if $T(\Sigma, R)$ has no positive row counts, then it must have no negative row counts either, and is hence balanced. \square

Justification of Algorithm WeakFindRank. By Lemma 5, we only need to show that the algorithm terminates. To prove this we need the following auxiliary result.

Lemma 6. *Suppose we are working with row k , that is $c(k) > 0$ and $c(i) \leq 0$ for all $i < k$. If row ℓ has no segments for some $\ell < k$, then the current path diagram $T(\Sigma, R')$ has no segments below row ℓ .*

Proof. Suppose to the contrary that $T(\Sigma, R')$ has a segment below row ℓ , then let V be the right most arrow that contains such a segment and say that it starts at column i . Since row ℓ has no segments, the starting rank of V must be $\leq \ell$. This implies that $i + 1 < m + n$ since the arrow that starts at level k must be to the right of V (by the increasing property of R'). This given, the current path diagram could look like in Figure 10, where the two green (thick) lines divide the plane into 4 regions, as labelled in the display.

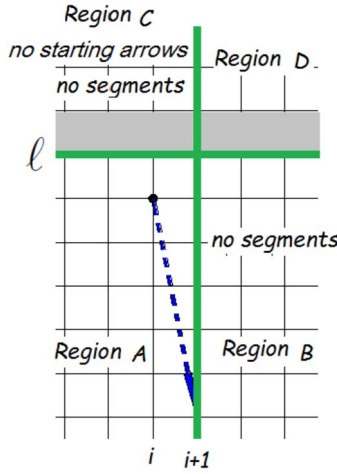


Figure 10: When row ℓ has no segments, there will be no segments in regions B and C .

The weakly increasing property of R' forces no starting ranks in C , therefore there are no segments there. By the choice of V there cannot be any segments in B . Thus the (gray) empty row ℓ forces no segments within both B and C .

Now notice that since Σ is the word of a path $D \in \mathcal{D}_{m,n}$ the number of red segments to the left of column $i + 1$ minus the number of blue segments to the left of that column must result in a number $s > 0$. However, since $c(j) \leq 0$ for all $j \leq \ell$ it follows that $c(0) + \dots + c(\ell - 1) \leq 0$. But since regions B and C have no segments it also follows that $s = c(0) + \dots + c(\ell - 1) \leq 0$, a contradiction. \square

Next observe that since each step of the algorithm increases one of the ranks by one unit, after M steps we will have $|R^{(M)} - R^{(0)}| = \sum_{i=1}^{m+n} r_i^{(M)} - \sum_{i=1}^{m+n} r_i^{(0)} = M$. This given, if the algorithm iterates Step 2 forever, then the maximum rank will eventually exceed any given integer. In particular, we will end up working with row k with k so large that $k - U$ exceeds the total number mn of red segments. At that point we will have $c(k) > 0$ and $c(j) = 0$ for all the $k - U$ values $j = U, U + 1, \dots, k - 1$. The reason for this is that we must have $c(j) \leq 0$ for all $0 \leq j < k$ and by Lemma 4 we must also have $c(j) \geq 0$ for $j \geq U$. Now, by the pigeon hole principle, there must also be some $U \leq \ell < k$ for which $c^r(\ell) = 0$. But then it follows that $c^b(\ell) = c^r(\ell) - c(\ell) = 0$, too. That means that row ℓ contains no segments.

Then Lemma 6 yields that there cannot be any segments below row ℓ either. This implies that the total row count is $\sum_{j \geq 0} c(j) = \sum_{j \geq U} c(j) \geq c(k) > 0$, a contradiction. \square

Thus the WeakFindRank algorithm terminates and we can draw the following important conclusion.

Theorem 7. *Given any (m, n) -Dyck path D with (S, W) -word Σ and any initial weakly increasing rank sequence $R^{(0)} = (r_1^{(0)}, r_2^{(0)}, \dots, r_{m+n}^{(0)})$, let $T(\Sigma, \bar{R})$ be the balanced path diagram produced by the WeakFindRank algorithm. Then the rank sequence*

$\bar{R} = (\bar{r}_1, \bar{r}_2, \dots, \bar{r}_{m+n})$ *will be strictly increasing. Moreover, the sequence*

$$\tilde{R} = (0, \bar{r}_2 - \bar{r}_1, \bar{r}_3 - \bar{r}_1, \dots, \bar{r}_{m+n} - \bar{r}_1)$$

is none other than the increasing rearrangement of the rank sequence of a pre-image \bar{D} of D under the sweep map.

Proof. Clearly, the path diagram $T(\Sigma, \tilde{R})$ of height $N = \tilde{r}_{m+n} + m + 1$ will also be balanced. Thus, by Theorem 3, \tilde{R} must be the increasing rearrangement of the rank sequence of a pre-image \bar{D} of D . In particular not only \tilde{R} but also \bar{R} itself must be strictly increasing. \square

This result has the following important corollary.

Theorem 8. *For any co-prime pair (m, n) the rational (m, n) -sweep map is invertible.*

Proof. Lemma 1 shows that the rational (m, n) -sweep map is into. Theorem 7 gives a proof (independent of the Thomas-Williams proof) that it is onto. Since the collection of (m, n) -Dyck paths is finite, the sweep map must be bijective. \square

Figure 11 depicts the entire history of the WeakFindRank algorithm applied to a $(7, 5)$ -Dyck path D paired with initial rank sequence $R^{(0)} = (0, 0, 5, 5, 5, 5, 5, 5, 5, 5, 5)$. Both D (on the left) and its pre-image \bar{D} (on the right) are depicted below.

A boxed lattice square in column i with an integer k inside indicates that arrow A_i was processed at the k^{th} step of the algorithm. As a result A_i was lifted from the level of the bottom of the square to its top level. For instance the square with 63 inside indicates that the red arrow A_7 was lifted at the 63rd step of the algorithm from starting at level 9 to starting at level 10. We also see that the last time that the arrow A_8 reached its final starting level at step 87. The successive final starting levels of arrows A_1, A_2, \dots, A_{12} give

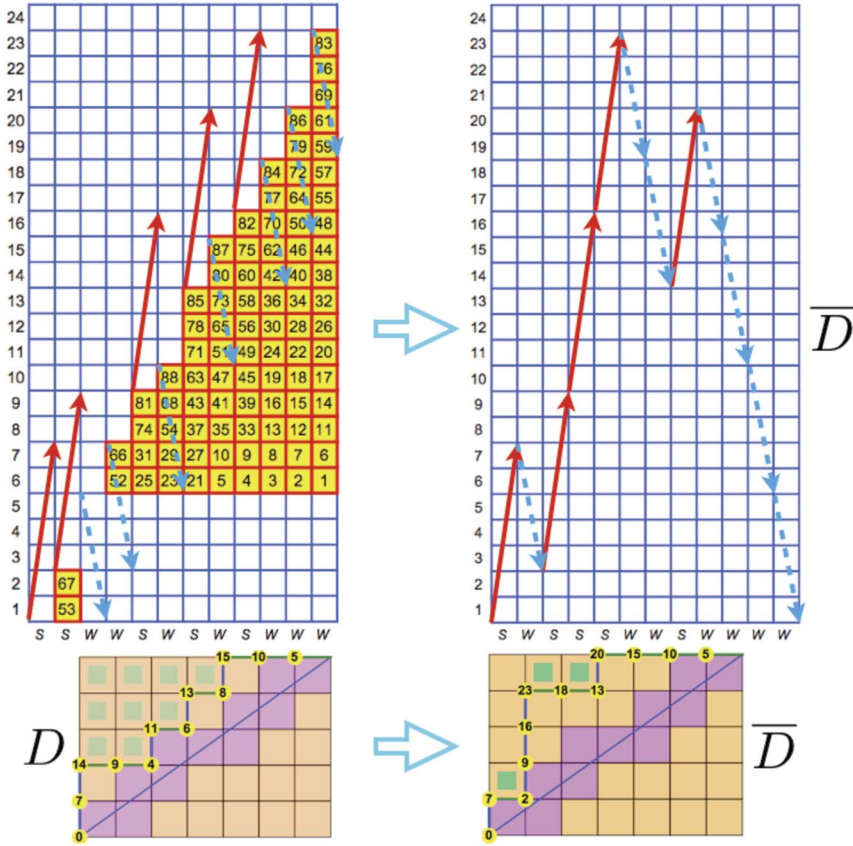


Figure 11: The top left picture is the balanced path diagram produced by the WeakFindRank algorithm. The top right one is obtained by Theorem 3.

the increasing rearrangement of the ranks of the path \bar{D} . Notice, arrows A_1 and A_3 were never lifted.

3. Tightness of algorithm WeakFindRank

Following the notations in Theorem 7, the number of steps needed for Algorithm WeakFindRank is $|\bar{R}| - |R^{(0)}| = |\tilde{R}| - |R^{(0)}| + (m+n)\bar{r}_1$. We will show that a specific starting path diagram can be chosen so that $\bar{r}_1 = 0$.

For two rank sequences $R = (r_1, r_2, \dots, r_{n+m})$ and $R' = (r'_1, r'_2, \dots, r'_{n+m})$ let us write $R \preceq R'$ if and only if we have $r_i \leq r'_i$ for all $1 \leq i \leq m+n$, if $r_i < r'_i$ for at least one i we will write $R \prec R'$. The distance of R from R' ,

will be expressed by the integer

$$|R' - R| = \sum_{i=1}^{m+n} (r'_i - r_i) = \sum_{i=1}^{m+n} r'_i - \sum_{i=1}^{m+n} r_i.$$

Given the (S, W) -word Σ of an (m, n) -Dyck path D , we will call the initial starting sequence $R^{(0)}$ *canonical for Σ* if it is obtained by replacing the first string of S in Σ by 0's and all the remaining letters by n , and call the balanced path diagram $T(\Sigma, \tilde{R})$ yielded by Theorem 7 *canonical for Σ* . Clearly $R^{(0)} \preceq \tilde{R}$. This given, we can prove the following remarkable result.

Theorem 9. *Let Σ be the (S, W) -word of a Dyck path $D \in \mathcal{D}_{m,n}$. If $R^{(0)}$ and $T(\Sigma, \tilde{R})$ are canonical for Σ and R is any increasing sequence which satisfies the inequalities*

$$R^{(0)} \preceq R \preceq \tilde{R}$$

then the WeakFindRank algorithm with starting path diagram $T(\Sigma, R)$ will have as output the rank sequence \tilde{R} .

Proof. Notice if $|\tilde{R} - R| = 0$ there is nothing to prove. Thus we will proceed by induction on the distance of R from \tilde{R} . Now assume the theorem holds for $|\tilde{R} - R| = K$. We need to show that it also holds for $|\tilde{R} - R| = K + 1$. Suppose one application of step (2) on R gives R' . We need to show that $R' \preceq \tilde{R}$. This is done since R and R' only differ from one unit we will have $|\tilde{R} - R'| = K$ and then the inductive hypothesis would complete the proof.

Thus assume if possible that this step (2) cannot be carried out because it requires increasing by one unit an $r_i = \tilde{r}_i$. Suppose further that under this step (2) the level k was the lowest with $c(k) > 0$ and thus the arrow A_i was the right most that started at level k . In particular this means that $r_i = \tilde{r}_i = k$. Since $|\tilde{R} - R| \geq 1$, there is at least one i' such that $r_{i'} < \tilde{r}_{i'}$. If $r_{i'} = k'$ let i' be the right most with $r_{i'} = k'$. Define R'' to be the rank sequence obtained by replacing $r_{i'}$ by $r_{i'} + 1$ in R . The row count $c(k')$ is decreased by 1 and another row count is increased by 1, so that $c(k)$ in R'' is still positive. Since $|\tilde{R} - R''| = K$ the induction hypothesis assures that the WeakFindRank algorithm will return \tilde{R} . But then in carrying this out, we have to work on row k , sooner or later, to decrease the positive row count $c(k)$. But there is no way the arrow A_i can stop being the right most starting at level k , since arrows to the right of A_i start at a higher level than A_i and are only moving upwards. Thus the fact that the WeakFindRank algorithm outputs \tilde{R} contradicts that fact that the application of step (2) to R cannot be carried out.




Figure 12: A $(7, 5)$ -Dyck path with area 4. Removing the black cell changes the rank 18 to $18 - 7 - 5 = 6$.

Thus we will be able to lift A_i one level up as needed and obtain the sequence R' by replacing r_i by $r_i + 1$ in R . But now we will have $|\tilde{R} - R'| = K$ with $R' \prec \tilde{R}$ and the inductive hypothesis will assure us that the `WeakFindRank` algorithm starting from R will return \tilde{R} as asserted. \square

It is clear now that the complexity of the `WeakFindRank` algorithm is $O(|\tilde{R}|)$. Recall that reordering \tilde{R} gives the rank set $\{r_1, r_2, \dots, r_{m+n}\}$ of \bar{D} . It is known and easy to show that

$$\text{area}(\bar{D}) = \frac{1}{m+n} \left(\sum_{i=1}^{m+n} r_i - \binom{m+n}{2} \right),$$

where $\text{area}(\bar{D})$ is the number of lattice cells between \bar{D} and the diagonal. Indeed, from Figure 12, it should be evident that reducing the area by 1 corresponds to reducing the sum $r_1 + \dots + r_{m+n}$ by $m+n$.

It follows that

$$(2) \quad |\tilde{R}| = (m+n)\text{area}(\bar{D}) + \binom{m+n}{2} = O((m+n)\text{area}(\bar{D})).$$

Theorem 9 together with (2) gives the following result.

Corollary 10. *Given any (m, n) -Dyck path D , its pre-image \bar{D} can be produced in $O((m+n)\text{area}(\bar{D}))$ running time.*

Proof. Let Σ be the (S, W) -word of D . We first construct the path diagram $T = (\Sigma, R^{(0)})$ with $R^{(0)}$ being canonical for Σ and compute the row counts of the path diagram. Next we use the `WeakFindRank` algorithm to update T until we get the balance path diagram (Σ, \tilde{R}) by Theorem 9. Finally we use Theorem 7 to find the pre-image \bar{D} .

Iteration only appears in the middle part, where the `WeakFindRank` algorithm performs $|\widehat{R}| - |R^{(0)}|$ times of Step 2. In each Step 2, we search for the lowest positive row count $c(j)$, then search for the rightmost arrow r_i that is equal to j , and finally update r_i by $r_i + 1$ and the row counts at *only* two rows (see Figure 9). Therefore, the total running time is $O(|\widehat{R}|)$, and the corollary follows by (2). \square

4. Validity of algorithm `StrongFindRank`

Let $R = (r_1 \leq r_2 \leq \dots \leq r_l)$ be a sequence of nonnegative integers of length l . The *strict cover* $R' = \mathbf{sc}(R) = (r'_1 < r'_2 < \dots < r'_l)$ of R is recursively defined by $r'_1 = r_1$ and $r'_i = \min(r_i, r'_{i-1} + 1)$ for $i \geq 2$. It is the unique minimal strictly increasing sequence satisfying $R' \succeq R$. The following principle is straightforward.

If $R \prec \overline{R}$ with R weakly increasing and \overline{R} strictly increasing, then $\mathbf{sc}(R) \preceq \overline{R}$.

A direct consequence is that $\widehat{R}^{(0)} = \mathbf{sc}(R^{(0)}) \preceq \widetilde{R}$ if $R^{(0)}$ is canonical for Σ . This sequence is exactly the starting rank sequence of our “strong” algorithm (see Figure 3). It will be good to review our definitions before we proceed.

Algorithm `StrongFindRank`

Input: A path diagram $T(\Sigma, \widehat{R}^{(0)})$ with Σ the word of a Dyck path $D \in \mathcal{D}_{m,n}$, the nonnegative strictly increasing rank sequence $\widehat{R}^{(0)}$ as above.

Output: The balanced path diagram $T(\Sigma, \widehat{R})$.

It will be convenient to keep the common height equal to N for all the successive path diagrams constructed by the algorithm, where $N = U + 2mn$, with $U = \max(R^{(0)}) + m + 1$.

Step 1 Starting with $T(\Sigma, \widehat{R}^{(0)})$ repeat the following step until the resulting path diagram is balanced.

Step 2 In $T(\Sigma, \widehat{R}^{(s)})$, with $\widehat{R}^{(s)} = (\widehat{r}_1^{(s)}, \widehat{r}_2^{(s)}, \dots, \widehat{r}_{m+n}^{(s)})$ find the lowest row j with $c(j) > 0$ and find the unique arrow that starts at level j . Suppose that arrow starts at (i, j) . Define R' to be the rank sequence obtained from $\widehat{R}^{(s)}$ by the replacement $\widehat{r}_i^{(s)} \rightarrow \widehat{r}_i^{(s)} + 1$, and set $\widehat{R}^{(s+1)} = \mathbf{sc}(R')$. Construct the path diagram $T(\Sigma, \widehat{R}^{(s+1)})$ and update the row counts. If all the row counts of $T(\Sigma, \widehat{R}^{(s)})$ are ≤ 0 then stop the algorithm and return $\widehat{R}^{(s)}$, since all these row counts must vanish.

This given, the validity of the **StrongFindRank** algorithm is an immediate consequence of the following surprising result.

Theorem 11. *Let $D \in \mathcal{D}_{m,n}$ with (S, W) -word Σ , and let the balanced path diagram (Σ, \tilde{R}) be canonical for Σ . Then all the successive rank sequences $\widehat{R}^{(s)}$ produced by the **StrongFindRank** algorithm satisfy the inequality*

$$(3) \quad \widehat{R}^{(s)} \preceq \tilde{R}$$

and since the successive rank sequences satisfy the inequalities

$$(4) \quad \widehat{R}^{(0)} \prec \widehat{R}^{(1)} \prec \widehat{R}^{(2)} \prec \widehat{R}^{(3)} \prec \dots \prec \widehat{R}^{(s)},$$

there will necessarily come a step when $T(\Sigma, \widehat{R}^{(s)}) = T(\Sigma, \tilde{R})$. At that time the algorithm will stop and output \tilde{R} .

Proof. The inequality (4) clearly holds since we always shift arrows upwards.

We prove the inequality (3) by induction on s . The basic fact that will play a crucial role is that the output \tilde{R} is strictly increasing. See Theorem 3.

The case $s = 0$ of (3) is obviously true since $\widehat{R}^{(0)}$ is the strict cover of $R^{(0)} \preceq \tilde{R}$. Assume $\widehat{R}^{(s)} \preceq \tilde{R}$ and we need to show (3) holds true for $s + 1$. Now $\widehat{R}^{(s+1)}$ is the strict cover of R' , where R' is the auxiliary rank sequence used by Step 2 of the **StrongFindRank** algorithm. Since R' is precisely the successor of $\widehat{R}^{(s)}$ by Step 2 of the **WeakFindRank** algorithm, it will necessarily satisfy the inequality $R' \prec \tilde{R}$ by Theorem 9. Our principle then guarantees that we will also have

$$\widehat{R}^{(s+1)} \preceq \tilde{R}$$

unless $\widehat{R}^{(s)} = \tilde{R}$ and the **StrongFindRank** algorithm terminates. □

Remark 12. *This proof makes it evident that, to construct the pre-image of an (m, n) Dyck path, the **StrongFindRank** algorithm will be more efficient than the **WeakFindRank** algorithm. This is partly due to the fact that the distances $|\widehat{R}^{(s+1)} - \widehat{R}^{(s)}|$ do turn out bigger than one unit most of the time, as we can see in the following display.*

In the middle of Figure 13 we have a Dyck path D , and below it, its pre-image \overline{D} . To recover \overline{D} from D we applied to D the **WeakFindRank** algorithm (on the left) and the **StrongFindRank** algorithm (on the right). The display shows that the “weak” algorithm required about 3 times more steps than the “strong” algorithm. The numbers in the Cyan squares reveal that, in several steps, two or more arrows were lifted at the same time. For instance, in step 13, as many as 4 arrows were lifted. The other step saving feature of the

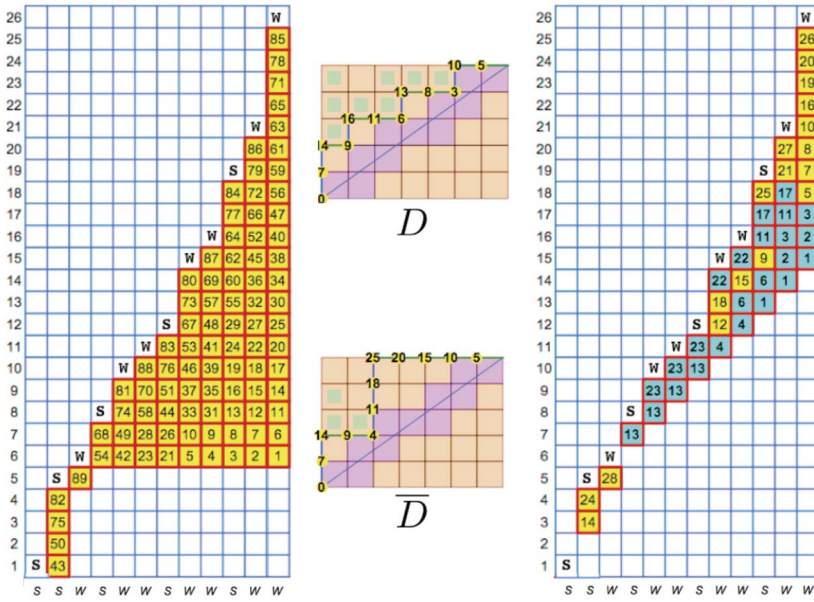


Figure 13: Comparison of the WeakFindRank algorithm and the StrongFindRank algorithm by an example.

“strong” algorithm is due to starting from the strict cover of the canonical starting sequence. This is evidenced by the difference between the number of white cells below the colored ones on the left and on the right diagrams.

5. Discussion and future plans

This work is done after the authors read [4] version 1, especially after the first named author talked with Nathan Williams. The concept “balanced path diagram” is a translation of “equitable partition” in [4]. The intermediate object “increasing balanced path diagram” is what we missed in our early attempts: The obvious 0-row-count property of Dyck paths gives the necessary part of Theorem 3, but we never considered the 0-row-count property to be sufficient until we read the paper [4].

Once Theorem 3 is established, inverting the sweep map is reduced to searching for the corresponding increasing balanced path diagram. Our algorithm is similar to the Thomas-Williams algorithm in the sense that both algorithms proceed by picking an initial candidate and then repeat an identical updating process until terminates. In the rational Dyck path model, our updating process is natural and has more freedom than the Thomas-

Williams algorithm. For instance, we can start with any weakly increasing rank sequence.

The precise relation between our algorithm and the Thomas-Williams one will be addressed in an upcoming paper, where we will extend the arguments in this paper to a more general class of sweep maps. These sweep maps have been defined in [1]. Though the invertibility of these sweep maps can be deduced from the modular sweep map model [4], they deserve direct proofs.

Even the rational sweep map needs further studied. The (m, n) -rational sweep map on $\mathcal{D}_{m,n}$ is known to take the `dinv` statistic to the `area` statistic. This result is proved combinatorially by Gorsky and Mazin in [3], but the proof is indirect. Our view of Dyck paths leads to visual description of the `dinv` statistics and a simple proof of the `dinv` and `area` result. See [2] for detailed information and references.

References

- [1] D. Armstrong, N. A. Loehr, and G. S. Warrington, Sweep maps: A continuous family of sorting algorithms, *Adv. Math.* 284 (2015), 159–185. [MR3391074](#)
- [2] A. Garsia and G. Xin, Dinv and Area, *Electron. J. Combin.*, 24 (1) (2017), P1.64. [MR3651946](#)
- [3] E. Gorsky and M. Mazin, Compactified Jacobians and q, t -Catalan Numbers, *J. Combin. Theory Ser. A*, 120 (2013), 49–63. [MR2971696](#)
- [4] H. Thomas and N. Williams, Sweeping up zeta, preprint, arXiv:1512.01483. [MR3678592](#)

ADRIANO GARSIA
DEPARTMENT OF MATHEMATICS
UCSD
USA
E-mail address: garsiaadriano@gmail.com

GUOCE XIN
SCHOOL OF MATHEMATICAL SCIENCES
CAPITAL NORMAL UNIVERSITY
BEIJING 100048
PR CHINA
E-mail address: guoce.xin@gmail.com

RECEIVED 20 MARCH 2016