

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Lattice Path Matroids: Negative Correlation and Fast Mixing

Emma Cohen* Prasad Tetali†
ecohen@math.gatech.edu tetali@math.gatech.edu

Damir Yeliussizov‡
yeldamir@gmail.com

June 10, 2015

Abstract

Catalan numbers arise in many enumerative contexts as the counting sequence of combinatorial structures. In this work, we consider natural Markov chains on some of the realizations of the Catalan sequence. While our main result is in deriving an $O(n^2 \log n)$ bound on the mixing time in L_2 (and hence total variation) distance for the random transposition chain on Dyck paths, we raise several open questions, including the optimality of the above bound. The novelty in our proof is in establishing a certain negative correlation property among random bases of lattice path matroids, including the so-called *Catalan matroid* which can be defined using Dyck paths.

Keywords: Dyck paths; Balanced matroid; Markov chain Monte Carlo; Transposition shuffle

1 Introduction

There are several combinatorial structures whose enumeration is given by the Catalan sequence: $1, 2, 5, 14, 42, \dots$, where the n th term (for $n \geq 1$) is given by the Catalan number $C_n = \frac{1}{n+1} \binom{2n}{n}$. Some examples include the set of all triangulations of a regular polygon of $n + 2$ sides, the set of non-crossing partitions of an n -set (the lattice on which is of much interest to researchers in free probability), and the set of balanced parentheses, with n left and n right parentheses—or equivalently, described as *Catalan strings* $x \in \{\pm 1\}^{2n}$ of n 1's

*School of Mathematics, Georgia Tech, Atlanta, GA 30332-0160. Corresponding author.

†School of Mathematics and School of Computer Science, Georgia Tech, Atlanta, GA 30332-0160

‡Department of Computer Science and Electrical Engineering, Kazakh-British Technical University, Almaty, Kazakhstan

1
2
3
4
5
6
7
8
9 and n -1 's with non-negative partial sums, $\sum_{i=1}^j x_i \geq 0$, for all j . This last
10 structure (of Catalan strings) is also known as the set of Dyck paths, and is
11 visualized as lattice paths of n up-steps and n down-steps, representing the 1 's
12 and the -1 's respectively.

13 For any given $n \geq 1$, generating a Catalan structure (such as triangulation
14 or a Dyck path) uniformly at random from the set of C_n many, is a straight-
15 forward task, and can be done in time linear in n , not too unlike generating a
16 permutation uniformly at random (out of the $n!$ many permutations) of n dis-
17 tinct letters. However, the study (of convergence to equilibrium) of particular
18 Markov chain Monte Carlo algorithms which yield a random Catalan structure
19 is seemingly much more interesting. The inspiration to this endeavor in part
20 stems from an open problem of David Aldous [1], who conjectured that the ran-
21 dom walk on triangulations of a polygon on n sides, performed using uniform
22 diagonal flips, ought to take time roughly $n^{3/2}$ (up to factors logarithmic in n).
23 Despite much effort by various researchers in this topic over several years, the
24 best known bounds for the mixing time of the chain on triangulations remain
25 those of McShine-Tetali [13], who proved an upper bound of $O(n^4)$ on the re-
26 laxation time (also known as the inverse spectral gap), and Molloy-Reed-Steiger
27 [14] who showed $\Omega(n^{3/2})$ as a lower bound.

28 In the work, we explore a (different) Markov chain on a different realization of
29 the Catalan number: we consider the set of Dyck paths, described using binary
30 strings $x \in \{\pm 1\}^{2n}$, as mentioned above, and consider the natural local move of
31 exchanging x_i and x_j , for a pair (i, j) chosen uniformly at random from among
32 all possible pairs, *provided* that the resulting string is a (valid) Catalan string;
33 else we reject such a move. Note that, without the rejection step, we would have
34 a random transposition chain on the set of $\binom{2n}{n}$ many binary strings with an
35 equal number of 1 's and -1 's. It is well-known (see e.g., [8])—sometimes under
36 the name *Bernoulli-Laplace model*—that $O(n \log n)$ such random transpositions
37 are necessary and sufficient to reach (close to) equilibrium. It is also well-known
38 that the so-called random transposition shuffle of n distinct cards mixes in
39 $O(n \log n)$ time. However, the “sea-level” constraint of having the partial sums
40 of the (Catalan) string stay non-negative seems to thwart any type of straight
41 forward (or otherwise) analyses, despite the best efforts of several experts in the
42 field.
43
44

45 Our main result is an $O(n^2 \log n)$ upper bound on the mixing time of the
46 random transposition walk on Dyck paths of length n . The proof relies on a
47 rephrasing of the walk as a basis exchange walk on a balanced matroid. In
48 Section 2 we describe the Catalan matroid (due to Ardila [2]) and the prior
49 work [9, 11, 10] regarding mixing times for the basis exchange walk on balanced
50 matroids. All that remains to bound the mixing time of our Dyck path walk is
51 the proof in Section 3 that the Catalan matroid is balanced. In Section 4 we
52 give a short proof of a mixing time lower bound for Dyck random transpositions
53 and, finally, in Section 5 we explore the possible consequences of this result for
54 other walks on Catalan structures and pose some open questions.
55
56
57
58
59
60
61
62
63
64
65

1.1 Random walks on lattice paths

Throughout, for integers $a \leq b$ we let $[a, b] := \{a, a+1, \dots, b\}$ denote the discrete closed interval and $[a] := [1, a]$.

A *lattice path* of length m is a string $P \in \{\pm 1\}^m$ of *up-steps* (+) and *down-steps* (-). The *height* of P at index i is $h_i(P) = \sum_{j=1}^i P_j$, and we can draw P on the grid as the graph of $f_P(i) = h_i(P)$. That is, we draw a path starting at $(0, 0)$ and taking up-steps $(1, 1)$ and down-steps $(1, -1)$. We will refer to a lattice path P of length m as a lattice path *from* $(0, 0)$ *to* $(m, h_m(P))$.

From this we can define a partial order on the set of lattice paths from $(0, 0)$ to $(m, 2r - m)$ by letting $P \leq Q$ whenever $h_i(P) \leq h_i(Q)$ for all $i \in [m]$. Note that $P \leq Q$ if and only if $q_i \leq p_i$ for all $i \in [r]$, where p_i (resp. q_i) is the index of the i th up-step in P (resp. Q).

Definition 1. A *Dyck path* of length $2n$ is a lattice path P from $(0, 0)$ to $(2n, 0)$ with $h_i(P) \geq 0$ for all $i \in [2n]$.

That is, the Dyck paths are precisely the paths P from $(0, 0)$ to $(0, 2n)$ with $P \geq (+-)^n$. It is a standard result that the number of Dyck paths of length $2n$ is the Catalan number $C_n = \frac{1}{n+1} \binom{2n}{n}$.

A natural class of random walks on lattice paths from $(0, 0)$ to (m, h) is the *transposition walk*, which at each step picks random indices $i, j \in [m]$ and swaps the steps of P at those indices. Indeed, if i and j are chosen uniformly independently at random this is just the Bernoulli-Laplace model on $\binom{[m]}{r}$, where $r = (h+m)/2$ is the number of up-steps in every such path, and has been studied extensively. If we pick i, j to be uniform among *adjacent* pairs of indices, we get a different walk, whose mixing properties were determined in [17].

But what if we consider lattice paths from $(0, 0)$ to $(2n, 0)$ and restrict the walk to Dyck paths? That is, from any Dyck path we pick i, j according to one of these rules but only perform the move if the resulting path is again a Dyck path. (We will call these walks the *Dyck random transpositions walk* and *Dyck adjacent transpositions walk*, respectively.)

Indeed, the analysis of [17] extends also to this case, yielding the same upper bound for the mixing time of the Dyck adjacent transpositions walk. On the other hand, the Dyck random transpositions walk has evaded such precise analysis.

2 Lattice path matroids

2.1 Matroids and the basis exchange walk

Recall that a nonempty set $\mathcal{B} \subseteq 2^U$ is the set of *bases* of a *matroid* $M = (U, \mathcal{B})$ if the following *basis exchange axiom* holds:

Matroid Basis Exchange Axiom. For any bases $A, B \in \mathcal{B}$ and every $e \in A \setminus B$ there exists $f \in B \setminus A$ such that $A \setminus \{e\} \cup \{f\} \in \mathcal{B}$ is a basis.

Among other things, this axiom guarantees that all bases have the same cardinality, which is the *rank* of M .

We will also make use of two dual operations on matroids: contraction and deletion.

Definition 2. For a matroid $M = (U, \mathcal{B})$ and an element $e \in M$, the matroid M *contract* e is $M_e = (U, \mathcal{B}_e)$, where $\mathcal{B}_e = \{B \in \mathcal{B} : e \in B\}$. Similarly, M *delete* e is $M^e = (U, \mathcal{B}^e)$, where $\mathcal{B}^e = \{B \in \mathcal{B} : e \notin B\}$.¹

A *minor* of M is any matroid which can be obtained from M through a series of contractions and deletions.

The order in which contractions and deletions are performed does not matter, so we will write M_I^J for the matroid obtained from M by contracting the elements in I and deleting the elements in J , and \mathcal{B}_I^J will denote the set of bases of M_I^J .

Given a matroid M , Feder and Mihail [9] study the following *basis exchange walk* on the state space \mathcal{B} of bases:

From state $B \in \mathcal{B}$, pick uniformly and independently at random elements $a \in U$ and $b \in B$ and move in the next step to $B' = B \cup \{a\} \setminus \{b\}$ if $B' \in \mathcal{B}$ (else remain at B).

The matroid basis exchange axiom guarantees that this walk is ergodic, and since it is symmetric its stationary distribution is uniform. It is tentatively conjectured that the basis exchange walk is fast for any matroid (i.e., the mixing time is bounded by some polynomial in $m = |U|$), but there is little evidence in favor of this.

On the other hand, [9] introduced the notion of *balanced* matroids, which capture the notion that for a randomly chosen basis (of the matroid or any of its minors), conditioning on the occurrence of one element (in the basis) makes the occurrence of any other less probable. They show that for the case of balanced matroids the walk is indeed rapidly mixing, by using decomposition techniques to bound its spectral gap.

Definition 3. A matroid $M = (U, \mathcal{B})$ is *negatively correlated* if for every pair of distinct elements $e, f \in U$

$$\frac{|\mathcal{B}_e|}{|\mathcal{B}|} \geq \frac{|\mathcal{B}_{ef}|}{|\mathcal{B}_f|}. \quad (1)$$

The matroid M is *balanced* if M and all of its minors are negatively correlated.

Negative correlation is equivalent to the very natural condition that for a uniform random basis $B \in \mathcal{B}$, $\Pr[e \in B] \geq \Pr[e \in B | f \in B]$.

¹Note that this definition differs slightly from the usual one, in which the element e being contracted or deleted is removed from the ground set. In our case it will be convenient to leave e in place, so as to more easily identify $\mathcal{B}_e \subseteq \mathcal{B}$ and preserve identities such as $\mathcal{B} = \mathcal{B}_e \cup \mathcal{B}^e$ (with $\mathcal{B}_e \cap \mathcal{B}^e = \emptyset$).




Figure 1: The Dyck path $(++++-+-++---+----)$ above corresponds to the basis $\{1, 2, 3, 4, 6, 8, 9, 13\}$ of the Catalan matroid of order 8.

Indeed, many common classes of matroids are balanced, including uniform matroids (whose bases are all size- r subsets of the ground set), graphic matroids (with the ground set being the edges of a connected graph and the bases being spanning trees of the graph), matroids of rank ≤ 3 , and regular matroids (matroids which can be represented over every field) [9, 6]. We will rely on the following extension of Feder and Mihail’s result for balanced matroids, due to Jerrum and Son.

Theorem 4 ([10]). *The spectral gap λ and log-Sobolev constant α for the basis exchange walk on a balanced matroid M of rank r on a ground set of size m are lower bounded by*

$$\lambda \geq \frac{2}{mr} \quad \text{and} \quad \alpha \geq \frac{1}{2mr}.$$

In particular, via standard mixing time bounds (see [7]), the bound on the log-Sobolev constant implies that the mixing time of the basis exchange walk on a balanced matroid is at most $O(mr \log \log |\mathcal{B}|)$.

Remark 5. Due to these standard bounds and the result for balanced matroids the proof of our main result never explicitly refers to the definitions of λ and α . These definitions can, however, be found in Section 4, where they are used to give a lower bound on the mixing time for the Dyck random transposition chain.

2.2 The Catalan matroid and other lattice path matroids

To aid our setting of the Catalan transposition walk in terms of matroids, we refer to the observation (made independently by Ardila [2] and Bonin and De Mier [4]) that the set of Dyck paths can be thought of as a matroid.

Definition 6. The *Catalan matroid* of order n is $\mathcal{C}(n) = ([2n], \mathcal{B}(n))$, where the elements of $\mathcal{B}(n)$ are the index-sets of up-steps in Dyck paths of length $2n$.

This matroid is precisely the transversal matroid for the set system $\mathcal{S} = \{[1], [3], \dots, [2n-1]\}$ (i.e., its bases are the systems of distinct representatives of \mathcal{S}). In short, the representative of the set $[2i-1]$ will be the index of the i th up-step in the corresponding Dyck path (although this assignment of distinct representatives need not be unique). Indeed,

Observation 7. *The basis exchange walk on the Catalan matroid is exactly the random transposition walk on Dyck paths.*

In other words, given Theorem 4, to obtain a mixing time bound for the Dyck transposition walk it suffices to show that the Catalan matroids are balanced.

To this end, it would be convenient if the Catalan matroids belonged to some known class of balanced matroids. As noted above, the main class of matroids known to be balanced is *regular matroids*, but Ardila notes that $\mathcal{C}(n)$ is not representable over any \mathbb{F}_q for $q \leq n - 2$ and thus is not regular. Transversal matroids also need not be balanced in general: Choe and Wagner [6] give a transversal matroid of rank 4 which is not balanced.

Thus the bulk of our work here will be to show that the Catalan matroid is balanced, and from this the main mixing result will follow immediately. For our later discussion it will help to work with a minor-closed class of matroids (which the Catalan matroids certainly are not). Bonin and de Mier ([3]) discuss the following class of *lattice path matroids*, which generalize the Catalan matroids by allowing *any* pair of bounding paths.

Definition 8. For two lattice paths $A \leq B$ from $(0, 0)$ to $(m, 2r - m)$, consider the set L of lattice paths P from $(0, 0)$ to $(m, 2r - m)$ with $A \leq P \leq B$. The *lattice path matroid* $\mathcal{L}[A, B]$ (which is of rank r on ground set $[m]$) has as its bases the index sets of up-steps of paths in L .

Although it is not immediately obvious that $\mathcal{L}[A, B]$ is a matroid, [3] observes that in fact the lattice path matroid $\mathcal{L}[A, B]$ is the transversal matroid of the set system $\{[a_1, b_1], [a_2, b_2], \dots, [a_r, b_r]\}$, where a_i (resp. b_i) is the index of the i th up-step in A (resp. B). Indeed, it is shown in that paper that the transversal matroids of a set systems $\{[a_1, b_1], [a_2, b_2], \dots, [a_r, b_r]\}$ with $a_1 \leq a_2 \leq \dots \leq a_r$ and $b_1 \leq b_2 \leq \dots \leq b_r$ are precisely the lattice path matroids. (The set system of this form corresponding to a given lattice path matroid may not be unique; for example, the Catalan matroid is the transversal matroid for both $\{[1], [3], \dots, [2n - 1]\}$ and $\{[1], [2, 3], \dots, [n, 2n - 1]\}$.) The class of basis exchange walks on lattice path matroids also includes the (unconstrained) Bernoulli-Laplace model as the basis exchange walk on a uniform matroid.

In addition to showing this correspondence, Bonin and de Mier show that the class of lattice path matroids is closed under taking minors and duals. They also define a smaller (in fact, minimal) minor-closed class of matroids containing \mathcal{C}_n , which they call *generalized Catalan matroids*, consisting of the lattice path matroids $M(A, B)$ where $A = (+)^r (-)^{m-r}$ is maximal among all paths from $(0, 0)$ to $(m, 2r - m)$. They give several nice properties of generalized Catalan matroids, which we will not go into here as our analysis covers all lattice path matroids. Indeed, Sohoni [15] has actually already shown the balanced property of generalized Catalan matroids, although he calls them *Schubert matroids* and does not link his result to rapid mixing of the Dyck random transposition chain (although he does mention the Dyck adjacent transposition chain). Our proof applies to the more general class of lattice path matroids.

Given that the class of lattice path matroids is minor-closed, it suffices for our main result to show that lattice path matroids are negatively correlated, but it is worth noting that the same method can also be used to show negative correlation directly for any minor of a lattice path matroid.

3 Mixing bound for random transpositions

To show negative correlation for lattice path matroids, it will be convenient to use an equivalent formulation of (1), which is easily obtained by repeatedly applying the identity $|\mathcal{B}| = |B_e| + |B^e|$,

$$|\mathcal{B}_{ef}||\mathcal{B}^{ef}| \leq |\mathcal{B}_e^f||\mathcal{B}_f^e|. \quad (2)$$

Now we are ready to prove

Theorem 9. *For every pair of lattice paths $A \leq B$ from $(0, 0)$ to $(m, 2r - m)$ the lattice path matroid $\mathcal{L}[A, B]$ is negatively correlated.*

Proof. Let $\mathcal{L}[A, B] = ([m], \mathcal{B})$ be any lattice path matroid.

To prove inequality (2) for every pair $e < f \in [m]$, we will construct an injective map

$$\varphi_{ef} : \mathcal{B}_{ef} \times \mathcal{B}^{ef} \rightarrow \mathcal{B}_e^f \times \mathcal{B}_f^e.$$

Note that we can associate members of these sets with lattice paths in $\mathcal{L}[A, B]$ by

- \mathcal{B}_{ef} : paths with up-steps at indices e, f ,
- \mathcal{B}^{ef} : paths with down-steps at indices e, f ,
- \mathcal{B}_e^f : paths with an up-step at index e and a down-step at index f , and
- \mathcal{B}_f^e : paths with a down-step at index e and an up-step at index f .

Let $P \in \mathcal{B}_{ef}$ and $Q \in \mathcal{B}^{ef}$ be lattice paths and consider the following cases. (The figures shown are for the Catalan matroid.)

Case 1. Suppose the paths P, Q intersect (without necessarily crossing) in the region (I) between e and f . Note that this includes all cases where the path P is below Q at e and above Q at f or vice versa. Take the first such intersection point x in (I) and switch the paths P, Q after x to obtain new paths $P' \in \mathcal{B}_e^f, Q' \in \mathcal{B}_f^e$ as shown in Figure 2. Set $\varphi_{ef}(P, Q) = (P', Q')$.

Case 2. Suppose the paths do not meet in region (I) and consider the paths P, Q in the region (II) after position f . Imagine translating the fragment of path Q after f to P so that their initial points in f coincide. If the imaginary fragment intersects P , let x be the first such point of intersection. Construct new paths P', Q' by swapping the segments of P and Q between f and x (including f itself). See Figure 3. If the new paths P', Q' are both between A and B (so that $P' \in \mathcal{B}_e^f$ and $Q' \in \mathcal{B}_f^e$), then we set $\varphi_{ef}(P, Q) = (P', Q')$.




Figure 2: The injection for Case 1.




Figure 3: The injection for Case 2.

Case 3. Otherwise, perform the mirror image of the operation from Case 2 in the part (III) up to position e , as shown in Figure 4 and again set $\varphi_{ef}(P, Q) = (Q', P')$.




Figure 4: The injection for Case 3. Case 2 fails because the translated segment of Q leaves the valid region before intersecting Q .

Case 1 covers all pairs of paths P, Q such that P is above Q at e and below it at f or vice versa. Case 2 covers all remaining pairs of paths with P above Q at f , since in this case the fragment of Q must intersect P and both resulting paths (P', Q') are between P and Q . Finally, Case 3 covers all still remaining pairs of paths with P below Q at e . Hence every pair of paths P, Q is covered by one of these cases.

Finally we must argue that φ_{ef} is injective. Given P', Q' we must be able to determine which case above was applied to produce them, and from this it is simple to recover P and Q . Since we know e, f it is easy to identify the regions (I), (II) and (III) in the three cases. If P', Q' intersect in region (I) then they must have come from Case 1, as neither of the later cases can produce such an intersection. If they do not, first try to apply the inverse operation for Case 2, which is the same as the forwards operation; if this was not possible for the starting paths P, Q then it is not possible for P', Q' either (as performing the transformation of Case 3 cannot cause the transformation of Case 2 to become valid if it was not already), so if we have a pair of paths P', Q' for which the

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

second move is possible and does result in lattice paths between A and B then we know we must have arrived at it through Case 2. Otherwise, we were in Case 3. See Figure 5.




Figure 5: An example of recovering the case from the result. Since the paths do not intersect in region (I) we must have come from Case 2 or 3. Case 2 fails because the translated segment of P' leaves the valid region before rejoining Q' , so the paths must have come from Case 3. Note that the transformation ϕ_{ef} is usually not onto.

□

Remark 10. Since Case 1 can be applied whenever P is above Q at e and below at f or vice versa, Case 2 can be applied whenever P is above Q at both e and f , and Case 3 can be applied whenever P is below Q at both e and f , it is tempting to try to divide the cases more naturally according to the relative heights of P and Q at e and f . However, note that in Cases 2 and 3 (and sometimes Case 1) the resulting paths P', Q' always have P' above Q' at both e and f , and this more natural division of cases does not allow us to recover uniquely which transformation was applied.

Remark 11. While the result of [3] that lattice path matroids are minor-closed shows that it suffices for us to prove negative correlation, it is worth noting that the same injection shows negative correlation directly for any minor of a lattice path matroid. In particular, if we assume that P and Q above are both in the minor \mathcal{B}_I^J (i.e., have up-steps at the indices of I and down-steps at the indices in J) then the resulting paths P' and Q' again have this property. Moreover, the injection also implies that lattice path matroids are Rayleigh matroids (a class of matroids studied in [6]).

4 Lower bound

To derive a lower bound on the mixing time of the Dyck random transposition chain we will need to refer to the definitions of the spectral gap and log-Sobolev constant. Recall that for a Markov chain with state space Ω , transition matrix P , and stationary distribution π

$$\lambda := \inf_f \frac{\mathcal{E}(f, f)}{\text{Var}_\pi(f)} \quad \text{and} \quad \alpha := \inf_f \frac{\mathcal{E}(f, f)}{\mathcal{L}_\pi(f)},$$

where

$$\mathcal{E}(f, f) = \frac{1}{2} \sum_{x, y} (f(x) - f(y))^2 P(x, y) \pi(x) \quad \text{and}$$

$$\mathcal{L}_\pi(f) = \sum_x f(x)^2 (\log f(x)^2 - \log \mathbb{E}_\pi(f^2)) \pi(x)$$

and the infima are taken over non-constant functions $f : \Omega \rightarrow \mathbb{R}$.

Using standard results we can (asymptotically) lower bound the mixing time by the relaxation time $\frac{1}{\lambda}$.

Theorem 12. *The spectral gap λ and log-Sobolev constant α for the Dyck random transposition chain satisfy*

$$2\alpha \leq \lambda \leq \frac{4}{n}.$$

Proof. It is always the case that $\alpha \leq \lambda/2$ (see, e.g., [10]), so it suffices to show that $1/\lambda \geq n/4$.

For $x \in \Omega$ a Dyck path of length $2n$, consider the function $f(x)$ giving the number of down-steps of x at even indices. The number of Dyck paths x of length $2n$ with $f(x) = k$ is precisely the Narayana number $N(n, k) = \frac{1}{n} \binom{n}{k} \binom{n}{k-1}$ (see [16]). Note that f is 1-Lipschitz, in that if x, y differ by a single transposition (i.e. $P(x, y) > 0$) then $|f(x) - f(y)| \leq 1$. We can therefore bound λ by noting that for a 1-Lipschitz function we have $\mathcal{E}(f, f) \leq 1/2$, and so it suffices to give a lower bound on $\text{Var}_\pi(f)$, which is precisely the variance of the Narayana distribution with p.m.f. $p(k) = N(n, k)/C_n$. This distribution is hypergeometric, and its variance is $\frac{(n+1)(n-1)}{4(2n-1)} \geq n/8$ (see, e.g., [12]). \square

Corollary 13. *The (total variation) mixing time of the Dyck random transposition walk is at least $\Omega(n)$.*

5 Open questions

Of course, the most immediate open question is whether the $O(n^2 \log n)$ mixing time bound for Dyck transpositions is tight. The general result for lattice path matroids is clearly not tight in every case (as it gives only order $n^2 \log n$ mixing time for Bernoulli-Laplace, where we know that the true bound is order $n \log n$), but there are lattice path matroids for which the $O(n^2 \log n)$ bound is tight². We believe that the Catalan matroid falls closer to the Bernoulli-Laplace side of the spectrum. Either way, we hope in the future to work towards a more complete characterization of the mixing rates for various lattice path transposition chains.

It is also not clear what consequences the new bound for Dyck paths might have for mixing on other Catalan structures. Many well-known Catalan structures suggest natural Markov chains, but it is striking how different they can

²One example is the lattice path matroid $\mathcal{L}[A, B]$ where $A = (-+)^n$ and $B = (+-)^n$, which is equivalent to the random walk on the n -cube slowed down by a factor of n .

1
2
3
4
5
6
7
8
9 be—even though there are natural bijections between classes of Catalan structures, they frequently do not preserve natural notions of distance, making standard comparison arguments tricky. There seems to be no obvious way to leverage even a hypothetical $O(n \log n)$ mixing time result for Dyck transpositions to yield improved bounds for any other Catalan chain. It is worth noting that the classical $(2n + 1$ -to-one many) Chung-Feller cyclic lemma suggests projecting the chain consisting of all transpositions on the (Bernoulli-Laplace) space of $\{\pm 1\}^{2n+1}$, with $(n + 1)$ 1’s and $n - 1$ ’s, onto the set of Dyck paths. This projection chain is in fact a Markov chain, as was observed by [5], and as such inherits a lower bound on the spectral gap of order $1/n$, from that of the Bernoulli-Laplace chain; however, many of the moves in the resulting projection chain are not particularly natural when viewed from the point of Dyck paths, and once again comparisons do not seem to help very much.

10
11
12
13
14
15
16
17
18
19
20
21
22 An alternate approach would be to come up with a general scheme to bound mixing times which could be modified for any Catalan structure. For example, it is well-known that the Catalan numbers satisfy the recurrence

$$C_n = \sum_{i=1}^n C_{i-1} C_{n-i},$$

23
24
25
26
27
28
29
30
31
32
33
34
35
36 and this is evidenced by a recursive structure in almost every combinatorial realization of the Catalan sequence. It is common to exploit recursive structure in bounding mixing times, and indeed both the triangulation bound of [14] and the balanced matroid bounds of [9, 10, 11] exploit different recursive structures of those objects. Unlike the recurrence above, the recursive structures used in those papers seem very particular to the Catalan realizations at hand, and do not generalize well to other Catalan chains.

37
38
39
40
41
42
43 Finally, we have mentioned above the conjecture that the basis exchange walk is rapidly mixing for *all* matroids. As evidence, it would be interesting to show rapid mixing for some class of matroids which is not a subclass of balanced matroids. One promising candidate might be the class of transversal matroids, for which the basis exchange walk seems closely related to the well-studied problem of walks on bipartite graph matchings.

44 45 Acknowledgements

46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65 The second author is grateful for many fruitful discussions he has had (over the years) on “Catalan shuffles” with Pietro Caputo, Christine Heitsch, Fabio Martinelli, Yuval Peres, David Wilson, and his former students, Lisa McShine, Antonio Blanca, and Ricardo Restrepo; in fact, the idea of the transposition chain analyzed in the present work was conceived in conversations with Heitsch. The authors thank the Institute for Mathematics and its Applications (IMA) for its excellent research atmosphere, generous support, and hospitality, which were all crucial to the collaboration resulting in the present work. The first two authors also acknowledge additional NSF support by way of the grant DMS-1407657.

References

- [1] D. Aldous. Triangulating the circle, at random. *The American Mathematical Monthly*, 101(3):pp. 223–233, 1994.
- [2] F. Ardila. The Catalan matroid. *Journal of Combinatorial Theory, Series A*, 104(1):49–62, Oct. 2003.
- [3] J. E. Bonin and A. de Mier. Lattice path matroids: structural properties. *European Journal of Combinatorics*, 27(5):701–738, July 2006.
- [4] J. E. Bonin, A. de Mier, and M. Noy. Lattice path matroids: enumerative aspects and Tutte polynomials. *Journal of Combinatorial Theory, Series A*, 104(1):63–94, Oct. 2003.
- [5] P. Caputo, R. Restrepo, and P. Tetali. Random transpositions on cyclic permutations of lattice paths. *Unpublished notes*, 2010.
- [6] Y. Choe and D. G. Wagner. Rayleigh matroids. *Combinatorics, Probability and Computing*, 15(05):765–781, 2006.
- [7] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *The Annals of Applied Probability*, 6(3):695–750, 1996.
- [8] P. Diaconis and M. Shahshahani. Time to reach stationarity in the Bernoulli–Laplace diffusion model. *SIAM Journal on Mathematical Analysis*, 18(1):208–218, 1987.
- [9] T. Feder and M. Mihail. Balanced matroids. In *ACM Symposium on Theory of Computing*, pages 26–38, New York, New York, USA, 1992. ACM Press.
- [10] M. Jerrum and J.-B. Son. Spectral gap and log-Sobolev constant for balanced matroids. In *Foundations of Computer Science*, pages 721–721, 2002.
- [11] M. Jerrum, J.-B. Son, P. Tetali, and E. Vigoda. Elementary bounds on Poincaré and log-Sobolev constants for decomposable Markov chains. *The Annals of Applied Probability*, 14(4):1741–1765, Nov. 2004.
- [12] N. L. Johnson, A. W. Kemp, and S. Kotz. *Univariate discrete distributions*. John Wiley & Sons, 2005.
- [13] L. McShine and P. Tetali. On the mixing time of the triangulation walk and other Catalan structures. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science: Methods in Algorithm Design*, 43:147–160, 1999.
- [14] M. S. Molloy, B. Reed, and W. Steiger. On the mixing rate of the triangulation walk. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 43–179, 1998.

- 1
2
3
4
5
6
7
8
9 [15] M. Sohoni. Rapid mixing of some linear matroids and other combinatorial
10 objects. *Graphs and Combinatorics*, 15(1):93–107, 1999.
11
12 [16] R. A. Sulanke. The Narayana distribution. *Journal of Statistical Planning*
13 *and Inference*, 101(1-2):311–326, 2002.
14
15 [17] D. B. Wilson. Mixing times of lozenge tiling and card shuffling Markov
16 chains. *The Annals of Applied Probability*, 14(1):274–325, 02 2004.
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65