

Rapid Mixing of the Flip Chain over Non-Crossing Spanning Trees

Konrad Anand  

School of Informatics, Informatics Forum, University of Edinburgh, UK

Weiming Feng  

School of Computing and Data Science, The University of Hong Kong, China

Graham Freifeld  

School of Informatics, Informatics Forum, University of Edinburgh, UK

Heng Guo  

School of Informatics, Informatics Forum, University of Edinburgh, UK

Mark Jerrum  

School of Mathematical Sciences, Queen Mary University of London, UK

Jiaheng Wang  

Faculty of Informatics and Data Science, University of Regensburg, Germany

Abstract

We show that the flip chain for non-crossing spanning trees of $n + 1$ points in convex position mixes in time $O(n^8 \log n)$. We use connections between Fuss-Catalan structures to construct a comparison argument with a chain similar to Wilson’s lattice path chain (Wilson 2004).

2012 ACM Subject Classification Theory of computation $\rightarrow$ Random walks and Markov chains; Theory of computation $\rightarrow$ Computational geometry

Keywords and phrases non-crossing spanning trees, Markov chain, mixing time

Digital Object Identifier 10.4230/LIPIcs.SoCG.2025.8

Related Version Full Version: <https://arxiv.org/abs/2409.07892> [4]

Funding This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 947778). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Weiming Feng: Weiming Feng acknowledges the support of Dr. Max Rössler, the Walter Haefner Foundation, and the ETH Zürich Foundation during his affiliation at ETH Zürich.

Jiaheng Wang: Jiaheng Wang also acknowledges support from ERC (grant agreement No. 101077083, CountHom).

1 Introduction

The reconfiguration of planar objects is an intensively studied topic in computational geometry. Given a set of points, such a planar structure can be a triangulation, a non-crossing perfect matching, or a *non-crossing spanning trees* (NCST), etc. In this paper we focus on the last one. Given $n + 1$ points in a plane, an NCST is a spanning tree such that no two edges intersect or overlap (except on the joint vertex) when all edges are drawn as straight line segments on the plane. The reconfiguration of NCSTs is done by *flips*, which remove an edge of the current tree, and then add an edge back so that the resulting graph is still a valid NCST. See Figure 1 for an example. Given two NCSTs over the same point set in general position, Avid and Fukuda [8] first proved that it takes at most $2n - 2$ flips to move from



© Konrad Anand, Weiming Feng, Graham Freifeld, Heng Guo, Mark Jerrum, and Jiaheng Wang; licensed under Creative Commons License CC-BY 4.0


41st International Symposium on Computational Geometry (SoCG 2025).

Editors: Oswin Aichholzer and Haitao Wang; Article No. 8; pp. 8:1–8:18




Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



one to the other. On the other hand, Hernando, Hurtado, Márquez, Mora and Noy [28] constructed an example where it takes at least $1.5n - O(1)$ flips, even when the points are *in convex position*. These bounds have not been improved until recently [1, 12, 11]. Very recently, Bjerkevik, Kleist, Ueckerdt, and Vogtenhuber [10] improved the upper bound to $5/3n - 3$ and the lower bound to $14/9n - O(1)$ when all points are in convex position.



■ **Figure 1** Illustration of a flip move. The red, dashed edge in the leftmost figure is dropped. This gives 9 possible edges (dotted in the middle figure) that can be added back to form a valid NCST. The blue, thick edge in the rightmost figure is picked.

Another classic topic in computational geometry is to count these planar structures. Traditional combinatorial studies focus on asymptotic upper and lower bounds of these structures (see e.g. [33, 43, 44]). Motivations for counting these structures have appeared in other topics, such as quantum gravity [32, 35], plane curves [47], algebraic geometry [20, 22], and the theory of discriminants [27]. From the computational perspective, the complexity of the counting problem seems to remain elusive in most cases, with an exception of counting triangulations of (not necessarily simple) polygons known to be $\#\mathbf{P}$ -hard [25]. On the other hand, counting the number of triangulations of a given set of points remains unresolved (but it is conjectured to be $\#\mathbf{P}$ -hard, see e.g. [21] or [17, §6.7]). Similarly, counting NCSTs over a set of points in general position apparently is also not known to be $\#\mathbf{P}$ -hard nor polynomial-time solvable, although we conjecture the former to be the case. To cope with the apparent intractability, sub-exponential time and fixed-parameterised approximation algorithms have been studied [2, 3, 34], but the complexity of approximate counting NCSTs does not appear to be known either.

On the other hand, when these $n + 1$ points are in convex position, the corresponding counting problem becomes easy to solve. One can show that the number of NCSTs, denoted by $C_{2,n}$, is subject to the following recursion: (see e.g. [42])

$$C_{2,0} = 1, \quad C_{2,n} = \sum_{\substack{i,j,k \in \mathbb{Z}_{\geq 0}, \\ i+j+k=n-1}} C_{2,i} C_{2,j} C_{2,k}, \quad \implies \quad C_{2,n} = \frac{1}{2n+1} \binom{3n}{n}. \quad (1)$$

Sharing a similar form as the famous *Catalan numbers*, this kind of generalisation is called the *2-Fuss-Catalan numbers*. The $(k$ -Fuss-)Catalan numbers capture more than 200 natural combinatorial structures. For example, the Catalan numbers count the number of triangulations, non-crossing perfect matchings and non-crossing partitions when the points are in convex position. The k -Fuss-Catalan numbers count the number of k -rooted plane trees, $(k+1)$ -ary trees, k -Dyck paths, and many more. We refer interested readers to Stanley's monograph [46] for an exhaustive survey about these Fuss-Catalan structures.

In this paper, we are interested in where reconfiguration and approximate counting meet. We consider a Markov chain whose steps are random flip moves over NCSTs, denoted the *flip chain*. In each step, we drop an edge chosen uniformly at random, and then add back

an edge uniformly at random among all valid choices. (This chain is formally defined in (7).) Reconfiguration results [8, 28, 10] thus imply that the flip chain is *irreducible* (namely, the state space of all NCSTs is connected via flip moves), and when the points are in convex position, the *diameter* of the chain is bounded between $14/9n - O(1)$ and $5/3n - 3$. Markov chains are also the most common approach to approximately count the number of combinatorial structures [30, 24, 31]. While the size of the state space (namely the total number of structures to count) can be exponentially large, the hope is that the Markov chain converges to the uniform distribution within polynomially many steps. The time to converge is called the *mixing time* (defined in Section 2). When the mixing time is bounded by a polynomial, the chain is said to be rapid mixing, and rapid mixing Markov chains usually lead to efficient approximate counting algorithms.

The mixing times of these geometrically defined Markov chains have received considerable attention from the Markov chain community in the past few decades, even when restricted to points in convex position, such as the chains over convex polygon triangulations [39, 38, 26], over triangulations on lattices and spheres [15, 14], or over lattice paths and (1-)Dyck paths [37, 48, 16]. See Cohen's thesis [18] for a more detailed explanation and comprehensive overview about random walks over 1-Fuss-Catalan structures. Beyond the fields of computational geometry and Markov chain analysis, these chains also find numerous applications, often thanks to the wide connections of Fuss-Catalan numbers, such as for quantum spin systems [13, 41, 40] or even in algebraic geometry [9].

In this paper we give the first polynomial upper bound on the mixing time of the flip chain over NCSTs for points in convex position.

► **Theorem 1.** *For $n + 1$ points on the plane in convex position, the mixing time of the flip chain over their non-crossing spanning trees is $O(n^8 \log n)$.*

The main feature that sets Theorem 1 apart from the results mentioned above is that NCSTs for points in convex position are 2-Fuss-Catalan structures. This renders some tools unavailable and others, such as the path comparison method [23], a lot more intricate to use. We also note that while the very similarly defined bases-exchange chain over spanning trees of a graph is known to have an optimal mixing time [6, 19], NCSTs do not form the bases of a matroid and these results apparently do not apply. We give an overview of our proofs next.

1.1 Technical overview

Our main strategy is to compare the flip chain with another chain that is easier to analyse. Naturally we would choose a chain over structures that have a bijection with NCSTs. There are a few candidates (see e.g. [45]), and we choose 2-Dyck paths. The bijection is defined through the same recursive structure, such as (1), for both 2-Dyck paths and NCSTs. Note that while these structures share the same count and have bijections among them, a local move in one structure may cause drastic changes in another through the bijection. Thus we need to pick a starting point chain so that the distortion caused by the bijection is manageable and the overhead of the comparison can be bounded by a polynomial.

We represent a 2-Dyck path by a string consisting of $2n$ ups ($\nearrow$) and n downs ($\searrow$) such that any prefix of the string contains no fewer ups than twice the number of downs. The formal definition is given in Section 3.1. One can verify that the positions of the down arrows form a basis of a matroid. This generalises the so-called Catalan matroid introduced by Ardila [7]. Since the bases-exchange walk for these matroids has an $O(n \log n)$ mixing time [19], a natural choice would be to compare this chain with the flip chain over NCSTs. This chain randomly chooses a down arrow and moves it to a random valid location. However, as mentioned before, one move in the bases-exchange walk may change NCSTs drastically. See Example 2.

[illegible]

Since the bases-exchange chain appears to be difficult to compare with, we choose a different chain to control the changes in the corresponding NCSTs, and that is the adjacent move chain. Instead of moving a down arrow to a random location, we restrict the movement to be either left or right one position. More formally, each transition of this chain chooses two adjacent coordinates uniformly at random. If swapping them results in a valid path, we do so with probability $1/2$, and otherwise make no change. For example, suppose the current state is $\nearrow \nearrow \searrow \nearrow \searrow$. If, with probability $1/5$, we choose the second and third positions, then the path will not change as swapping them leads to an invalid 2-Dyck path. On the other hand, if (with probability $1/5$) we choose the third and fourth positions, then the chain will move to $\nearrow \nearrow \searrow \nearrow \searrow$ with probability $1/2$ and remain unchanged otherwise.


While this chain is not directly studied before, we verify that the coupling method of Wilson [48] still applies here. Wilson’s original argument is applied to Lattice paths or the Bernoulli-Laplacian model (the uniform distribution over *all* strings with a fixed number of ups and downs). Without too much change, the same argument implies that the adjacent move chain over 2-Dyck paths mixes in $O(n^3 \log n)$ steps.

Our main technical contribution is then to compare the adjacent move chain over 2-Dyck paths with the flip chain over NCSTs. We need to characterise how one adjacent move changes the NCSTs. Consider again W_2 in Example 2, and another 2-Dyck path W_3 that can be reached from W_2 by a single adjacent move:

[illegible]

The adjacent move swaps the position 20 and 21. The corresponding NCST of W_3 is shown in Figure 2. Note that there are still a lot of altered edges, and in fact, one can also construct examples where $\Omega(n)$ edges get changed after an adjacent move. This time, however, the reader may have observed a pattern of the changes: the blue part in the NCST of W_3 is the same as the red part in that of W_2 , but shifted clockwise once. As formally proved in Section 5.1, each adjacent move corresponds to at most two edge flips, plus one shifting in a particular form (see (9)).

To apply the path method for Markov chain comparison [23], we need to simulate an adjacent move with a sequence of flip moves, and to be able to recover the initial and final states of the adjacent move, given the two states of the flip move and at most polynomial amount of additional information. The main task is to simulate the shift. We do so by identifying a hierarchical structure of the shift, and only flip edges following a particular pattern. This design allows us to uniquely recover the initial and final states of the shift, as long as we know the current transition, the recursion depth, and constant amount of extra information. This is the crux of our whole argument and is given in Section 5.



■ **Figure 2** NCSTs corresponding to the 2-Dyck paths in Example 2.

Theorem 1 is shown by combining the comparison argument with the $O(n^3 \log n)$ mixing time of the adjacent move chain. The comparison has an overhead of $O(n^4)$, which, roughly speaking, consists of $O(n)$ for the encoding, $O(n)$ for the path length, and $O(n^2)$ to account for the transition probability difference. Finally, together with the $O(n)$ overhead resulting from spectral gaps, we obtain our mixing time bound of $O(n^8 \log n)$.

1.2 Related work

McShine and Tetali's mixing time argument for convex triangulations [38] is similar to our result in spirit, which is also built on a comparison argument with Wilson's lattice path chain [48]. However, in their case the starting point is 1-Dyck path and the argument is more straightforward. The idiosyncrasies of non-crossing spanning trees warrant a more complicated path construction.

We have mentioned that the down arrows of 2-Dyck paths form bases of a matroid. However it appears difficult to directly relate that matroid and related dynamics with the flip chain. Another very similar structure is the *graphic matroid*, whose bases are spanning trees of a graph. The bases-exchange walk for matroids is known to mix very rapidly [6, 19]. However, with the non-crossing restriction, even when the points are in convex position, NCSTs do *not* form the bases of a matroid if we take all possible edges as the ground set. To see this, we give an example where the basis exchange property of matroids fails. Consider 4 points $\{0, 1, 2, 3\}$ on a circle, a tree T_1 consisting of edges $(0, 1)$, $(1, 3)$ and $(2, 3)$, and another tree $T_2 = \{(0, 2), (1, 2), (2, 3)\}$. Remove the edge $(0, 1)$ from T_1 . The possible choices of the new edge to add are $T_2 \setminus T_1 = \{(0, 2), (1, 2)\}$, but neither of them results in a valid non-crossing spanning tree, contradicting to the basis exchange property. Nonetheless, it is an interesting open problem to utilise techniques for the rapid mixing of graphic spanning trees to obtain better bounds in the non-crossing spanning trees context.

There has been a flurry of new tools building upon [6, 19] for analysing Markov chains, most notable among which is the spectral independence technique [5]. However, the results utilising spectral independence or its variants mostly focus on spin systems, and non-crossing trees are not spin systems. The trees all have the same size, which is a constraint usually not present in spin systems. There are some exceptions, such as the optimal mixing time of the down-up walk of independent sets of the same size [29]. These analyses rely heavily on the relationship between the fixed size distribution with the one without size restrictions. There is no similar connection to be utilised for non-crossing trees.

1.3 Open problems

The most straightforward open problem is to establish the correct order of the mixing time for the NCST flip chain. The $O(n^8 \log n)$ bound in Theorem 1 is unlikely to be tight. Our starting point, the adjacent move chain for 2-Dyck paths, has diameter $\Omega(n^2)$ by a simple potential argument,¹ and thus does not mix very fast. On the other hand, reconfiguration results [8, 10] imply that the flip chain has diameter $O(n)$. Therefore, we expect the flip chain to mix at least no slower than the adjacent move chain, instead of the polynomial slow-down in our current bound. On the other hand, there is no known lower bound on the mixing time other than the trivial $\Omega(n)$ obtained via the diameter bound. It would be interesting to close or shrink this gap.

Another interesting problem is to study the flip chain when points are *not necessarily* in convex position. Recall that the complexity of counting NCSTs with input points in general position, either exactly or approximately, is not known. Studying the flip chain would be helpful in resolving the approximate counting complexity. Constructing a set of points such that the flip chain mixes torpidly (i.e., not mixing in polynomial time) would rule out direct MCMC approaches, and often torpid-mixing instances help people construct gadgets for hardness reductions. On the other hand, if this chain always mixes in polynomial time, then efficient approximate counting would be possible.

2 Markov chain comparison

In this section we review some notions of discrete-time Markov chains over discrete state spaces. For detailed backgrounds we refer the reader to the book [36]. Let Ω be a finite discrete state space and π a distribution on Ω . Let $P : \Omega \times \Omega \rightarrow \mathbb{R}_{\geq 0}$ be the transition matrix of a Markov chain with stationary distribution π . We say P is reversible if $\pi(x)P(x, y) = \pi(y)P(y, x)$ for any $x, y \in \Omega$.

Given two distributions π and π' over Ω , the total variation distance between them is defined as $d_{\text{TV}}(\pi, \pi') := \frac{1}{2} \sum_{x \in \Omega} |\pi(x) - \pi'(x)|$. For a Markov chain P with stationary distribution π , let $d(t) = \max_{x \in \Omega} \{d_{\text{TV}}(P^t(x, \cdot), \pi)\}$. Then the mixing time of P is defined as $t_{\text{mix}}(P) := \min\{t : d(t) \leq \frac{1}{2e}\}$. The constant $1/(2e)$ here is not usually important, as the error decays exponentially as t increases.

Let f, g be two real valued functions on Ω . The Dirichlet form of P with the stationary distribution π is

$$\mathcal{E}_P(f, g) := \frac{1}{2} \sum_{x, y \in \Omega} [f(x) - f(y)][g(x) - g(y)]\pi(x)P(x, y).$$

Define the variance of a function $f : \Omega \rightarrow \mathbb{R}$ by $\text{Var}_\pi(f) := \mathbb{E}[f^2] - (\mathbb{E}[f])^2$. Then the *spectral gap* of the Markov chain P is

$$\lambda(P) := \inf \left\{ \frac{\mathcal{E}_P(f, f)}{\text{Var}_\pi(f)} \mid f : \Omega \rightarrow \mathbb{R}, \text{Var}_\pi(f) \neq 0 \right\}. \quad (2)$$

Equivalently, the spectral gap is the difference between the largest and second largest eigenvalues of the transition matrix of P , but the definition in (2) will be more convenient for us later.

¹ One can also consider how the area below the path evolves and show that the mixing time is at least $\Omega(n^3 / \log n)$, so Wilson's bound is almost tight.

The mixing time of any reversible Markov chain P can be bounded using the inverse spectral gap (also called the *relaxation time*) (see e.g. [36, Theorem 12.4, 12.5]):

$$\frac{1}{\lambda(P)} - 1 \leq t_{\text{mix}}(P) \leq \frac{1}{\lambda(P)} \left(1 + \frac{1}{2} \log \pi_{\min} \right), \quad (3)$$

where $\pi_{\min} := \min_{x \in \Omega} \pi(x)$. This $\pi_{\min}$ term is usually a single inverse exponential, which means the mixing time bound obtained this way typically differs from the relaxation time by a linear factor.

2.1 Path method

The path method [23] is a way to compare two Markov chains P and $\tilde{P}$, with stationary distributions π and $\tilde{\pi}$, by simulating any transition in $\tilde{P}$ using a series of transitions in P . This is a generalisation of the canonical path method by Jerrum and Sinclair [30].

For every pair of $x, y \in \Omega$, let $\Gamma_{xy} = (x = x_0, x_1, \dots, x_{\ell-1}, x_{\ell} = y)$ be a path of length $\ell = |\Gamma_{xy}|$ in the state space where $P(x_{i-1}, x_i) > 0$ for all $i \in [\ell]$. The congestion ratio B for a collection of paths $\Gamma = \{\Gamma_{xy}\}_{x, y \in \Omega}$ is defined as

$$B := \max_{(x', y') : P(x', y') > 0} \left(\frac{1}{\pi(x')P(x', y')} \sum_{\substack{x, y \\ \Gamma_{xy} \ni (x', y')}} \tilde{\pi}(x) \tilde{P}(x, y) |\Gamma_{xy}| \right). \quad (4)$$

► **Theorem 3** (Theorem 2.1 of [23]). *Let P and $\tilde{P}$ be two reversible Markov chains. Let B be the congestion ratio for paths $\Gamma = \{\Gamma_{xy}\}$ using transitions in P . Then for all $f : \Omega \rightarrow \mathbb{R}$,*

$$\mathcal{E}_{\tilde{P}}(f, f) \leq B \mathcal{E}_P(f, f).$$

Roughly speaking, the congestion bounds how much of a slowdown the path method can cause. It gives an upper bound on how much of a bottleneck can result from our simulated transitions.

3 Fuss-Catalan numbers and structures

This section introduces the combinatorial structures that we frequently utilise. These structures are closely related to the famous *Fuss-Catalan number*.


► **Remark 4.** Different authors appear to index Fuss-Catalan numbers differently in the literature. We adopt the definition used in Stanley's monograph [46].

3.1 Dyck paths

► **Definition 5** (k -Dyck path). *Let k be a positive integer. A sequence $(a_1, a_2, \dots, a_{(k+1)n}) \in \{+1, -k\}^{(k+1)n}$ forms a k -Dyck path of length $(k+1)n$, if*

$$\forall j \in [(k+1)n], \quad \sum_{i=1}^j a_i \geq 0, \quad \text{and} \quad \sum_{i=1}^{(k+1)n} a_i = 0.$$

Imagine that there is a hiking frog starting from the sea level. Each time it takes a step forward, it either lifts itself up by 1 meter, or drops by k meters. The two conditions are interpreted as (1) at any point it cannot go below the sea level, and (2) after $(k+1)n$ steps



■ **Figure 3** Decomposing Example 6 by Fact 7.

it must land on the sea level. This view leads to the following terminology for k -Dyck paths. We replace $+1$ with an *up-step* $\nearrow$, and $-k$ with a *down-step* $\searrow$. The partial sum

$$\text{ht}(j) := \sum_{i=1}^j a_i$$

indicates the *height* of the frog after j steps. Below is an example of a valid 2-Dyck path of length 21.

► **Example 6.** $\nearrow \nearrow \nearrow \searrow \nearrow \searrow \nearrow \searrow \nearrow \nearrow \nearrow \searrow \nearrow \searrow \searrow \nearrow \searrow$ is a 2-Dyck path.

The following simple fact is useful later.

► **Fact 7.** *Any non-empty k -Dyck path W can be uniquely written in the form*

$$W = \nearrow A_1 \nearrow A_2 \cdots \nearrow A_k \searrow B,$$

where A_i 's and B are (possibly empty) k -Dyck paths.

The decomposition can be obtained as follows. The B part begins at the position where the height *first* reaches 0 (except the starting point), namely $\text{Start}(B) = \min\{j : \text{ht}_W(j) = 0, j \geq 1\}$. Each A_i part begins at the position one step after where the height reaches $i-1$ for the *last* time before the B part, namely $\text{Start}(A_i) = 1 + \max\{j : \text{ht}_W(j) = i-1, 0 \leq j < \text{Start}(B)\}$. See Figure 3 for an example of such a decomposition for a 2-Dyck path.

Fact 7 immediately implies the recurrence on the number of k -Dyck paths of length $(k+1)n$, denoted by $C_{k,n}$:

$$C_{k,0} = 1, \quad C_{k,n} = \sum_{\substack{i_1, \dots, i_{k+1} \in \mathbb{Z}_{\geq 0}, \\ \sum_{j=1}^{k+1} i_j = n-1}} C_{k,i_1} C_{k,i_2} \cdots C_{k,i_{k+1}}, \implies C_{k,n} = \frac{1}{kn+1} \binom{(k+1)n}{n}. \quad (5)$$

This is called the k -Fuss-Catalan number [46, Chapter 4, A14], a generalisation of the Catalan number. It is well-known that the number of 1-Dyck paths equals the Catalan number.

3.2 Non-crossing spanning trees

Non-crossing spanning trees for points in convex position have a one-to-one correspondence with 2-Dyck paths. It will be more convenient to consider the following alternative way to draw the trees (see [45, Exercise 5.46]). Instead of putting points in convex position, we

arrange them on a line. Each edge is represented by a curve, drawn above the line segment. A tree is *non-crossing* if and only if it can be drawn this way without any two edges intersecting (except at the endpoints). See Figure 4(a)(b) for the two drawings of the same tree. In other words, each edge is represented by a pair of integers (a, b) where $0 \leq a < b \leq n$, and there does not exist any two edges (a, b) and (c, d) that $a < c < b < d$. As a result, any subtree must consist of points labeled by consecutive numbers. For this reason, if a subtree is spanned by $s, s+1, \dots, t$, we use a shorthand $[s, t]$ for it. A subtree is empty (consists of no edges) if $s = t$.

► **Fact 8.** Any n -edge non-empty non-crossing spanning tree T on points labeled $0, 1, \dots, n$ can be uniquely decomposed into a tuple of three non-crossing trees (T_A, T_B, T_C) , with the edge counts summing up to $n-1$, in the following way:

- Let t be the maximum index that there is an edge $(0, t)$.
- Let s be the index such that, after removing the edge $(0, t)$, the subtree $[0, s]$ is connected, but $[0, s+1]$ is not. In other words, there is a gap between s and $s+1$.
- The first tree T_A is the subtree $[0, s]$. It is empty if $s = 0$.
- The second tree T_B is the subtree $[s+1, t]$. It is empty if $t = s+1$.
- The last tree T_C is the subtree $[t, n]$. It is empty if $n = t$.

Resembling Fact 7 for 2-Dyck paths, the above decomposition not only suggests exactly the same recurrence (1) for the number of n -edge non-crossing spanning trees, but also a *one-to-one correspondence* between 2-Dyck paths and non-crossing spanning trees. Compare Figure 4(c) and Figure 3 for instance.

► **Corollary 9.** The following ruleset, when applied inductively, gives a one-to-one correspondence between 2-Dyck paths of length $3n$ and non-crossing spanning trees containing n edges.

- An empty path corresponds to an empty tree.
- If a tree T is decomposed into (T_A, T_B, T_C) by Fact 8, with each part corresponding to paths A , B and C respectively, then T corresponds to the path $\nearrow A \nearrow B \searrow C$.


The decomposition has the following useful property.

► **Proposition 10.** Let U and V be two 2-Dyck paths of length $3n_U$ and $3n_V$ respectively, and let T_U and T_V be their corresponding non-crossing spanning trees. Then, the non-crossing spanning tree T_{UV} corresponding to the concatenated Dyck path UV is a concatenation of T_U and T_V , meaning that the subtree $[0, n_U]$ of T_{UV} is T_U , and the subtree $[n_U, n_U + n_V]$ of T_{UV} is T_V (up to a shift in indexing).

The proof of Proposition 10 can be found in the full version [4]. Inspired by the decomposition, we introduce the following terminology that will be useful later.

► **Definition 11.** In the context of Fact 8:

- (pivot edge) the edge $(0, t)$ is called the pivot edge of T ;
- (gap) the vertices between s and t are called the gap beneath the pivot edge $(0, t)$; the pivot edge is called the overarching edge of the gap;
- (overarching edge) an edge (a, b) is the overarching edge of another edge (c, d) , if $a \leq c < d \leq b$, and there does not exist any other edge (a', b') that $a \leq a' \leq c < d \leq b' \leq b$;
- (minimal segment) a subtree $[c, d]$ is a minimal segment under an edge (a, b) , if the edge (c, d) exists, and (a, b) is the overarching of (c, d) ; a subtree $[c, d]$ is an outmost minimal segment, if the edge (c, d) exists, and there is no overarching edge of (c, d) .



■ **Figure 4** (a)(b) Two ways of representing a non-crossing spanning tree. The tree is corresponded to the 2-Dyck path in Example 6. (c) Decomposing the tree by Fact 8.

► **Example 12.** Consider the tree in Figure 4.

- The pivot edge of the whole tree is $(0, 6)$. The pivot edge of the subtree $[0, 3]$ is $(0, 2)$.
- The gap beneath the edge $(0, 6)$ is after 3. The gap beneath the edge $(4, 6)$ is after 5.
- The overarching edge of the edge $(0, 1)$ is $(0, 2)$, whose overarching edge is $(0, 6)$, which does not have any overarching edge.
- The minimal segments under the edge $(0, 6)$ are $[0, 2]$, $[2, 3]$ and $[4, 6]$. The outmost minimal segments are $[0, 6]$ and $[6, 7]$.

4 Spectral gap of the adjacent move chain

Over the state space of 2-Dyck paths of length $3n$, define the following *adjacent move chain*:

- (1) starting from the current path X , choose a position $i \in [3n - 1]$;
- (2) let X' be the same as X except that the steps at positions i and $i + 1$ are swapped;
- (3) if X' is invalid, then the chain stays at X . Otherwise, it moves to X' with probability $1/2$ and stays at X with probability $1/2$.

Formally, let $B(X)$ be the set of 2-Dyck paths that can be obtained from X by swapping two adjacent coordinates. Trivially, $|B(X)| < n$. The transition probability of this chain is

$$P_{\text{AM}}(X, X') = \begin{cases} 1 - \frac{|B(X)|}{6n - 2} & \text{if } X = X'; \\ \frac{1}{6n - 2} & \text{if } X' \in B(X); \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Note that $P_{\text{AM}}(X, X) \geq 4/5$. It is easy to see that P_{AM} is irreducible and verify that $P(X, X') = P(X', X)$, so P_{AM} is reversible.

We have the following bound on the mixing time of this chain.

► **Theorem 13.** $t_{\text{mix}}(P_{\text{AM}}) = O(n^3 \log n)$.

The proof essentially follows Wilson's coupling argument [48]. We use the same potential function and verify that the key properties for the mixing time upper bound still hold. (The properties in [48] for the mixing time lower bound no longer hold.) More details are given in the full version [4]. Combining Theorem 13 and (3) has an immediate corollary.

► **Corollary 14.** $1/\lambda(P_{\text{AM}}) = O(n^3 \log n)$.

5 Adjacent move versus flip move

In this section we analyse the mixing time of the flip chain by comparing it with the adjacent move chain and using Corollary 14. The flip chain P_{FM} is a random walk on non-crossing spanning trees (or non-crossing trees), where in each step we remove an edge uniformly at random, and add back a non-crossing edge between the two components. Such a move is called an (edge) flip. Formally, given two non-crossing spanning trees S and T , the probability of moving from S to T is

$$P_{\text{FM}}(S, T) = \begin{cases} \sum_{T': |S \cap T'| = n-1} \frac{1}{n\delta(S, T')} & \text{if } S = T; \\ \frac{1}{n\delta(S, T)} & \text{if } |S \cap T| = n-1; \\ 0 & \text{otherwise,} \end{cases} \quad (7)$$

where $\delta(S, T)$ is the number of edges that one can add to $S \cap T$ to make it a non-crossing spanning tree. Note that

$$P_{\text{FM}}(S, S) \geq 1/n, \quad \text{and } P_{\text{FM}}(S, T) \geq \Omega(n^{-3}) \text{ if it is non-zero.} \quad (8)$$

It is easy to see that P_{FM} is irreducible. Note when $|X \cap Y| = n-1$, $P(X, Y) = \frac{1}{n\delta(S, T)} = \frac{1}{n\delta(T, S)} = P(Y, X)$, and it is trivial to see that $P(X, Y) = P(Y, X)$ in the other two cases, so P_{FM} is symmetric, thus is reversible.

5.1 Characterise the adjacent move via flip moves

A single adjacent move on non-crossing spanning trees still requires altering an $\Omega(n)$ number of edges. Yet, we will show that it consists of a constant number of flip moves, plus one subtree *shifting*, taking the form of


$$\text{moving from } \begin{array}{c} \text{---} \text{---} \text{---} \end{array} \text{ to } \begin{array}{c} \text{---} \text{---} \end{array} \text{ or vice versa,} \quad (9)$$

where the bubble does not change. Note that the overarching edge of the subtree remains unchanged. This is formally defined as follows, using the non-crossing tree decomposition.


► **Definition 15** (shifting). *Let T be a non-crossing spanning tree of size m decomposing into (T_A, T_B, T_C) by Fact 8, where T_B and T_C are empty. Let T' be another non-crossing spanning tree of size m decomposing into (T'_A, T'_B, T'_C) , where both T'_A and T'_C are empty, and $T'_B = T_A$. A shifting is the transformation from T to T' (right shifting) or from T' to T (left shifting).*

This particular form of shifting may look strange at first sight. A more intuitive version would be just moving a subtree left or right, without caring about the overarching edge or the other side of the gap. However, as it will become clear later when we bound the congestion of the paths, it is this form of shifting that allows an efficient encoding.

Next is the characterisation of an adjacent move via flip moves and shifting.



■ **Figure 5** (*a) The 2-Dyck path representations of $\tilde{I}$. In the adjacent move, the blue part is changed into red. (*b) The corresponding flip move sequences. (1*) Corresponding to **Type 1**. (2*) Corresponding to **Type 2**.



■ **Figure 6** An example of a shift sequence. The black longer bubbles represent minimal segments, and there may be 0 or more of them. The red shorter bubbles represent valid sub-trees and are not necessarily minimal. The minimal segments $S_{i+1}, \dots, S_k$ have finished moving, and $S_1, \dots, S_{i-1}$ are yet to be moved, with S_i being the current one moving. M_5, M_6, M_7 are edge flips, and R represents a recursive move.

► **Lemma 16.** *Let $I \rightarrow F$ be a transition in the adjacent move P_{AM} such that $I \neq F$. This transition can be simulated from T_I to T_F by at most two edge flippings in P_{FM} , plus at most one shifting.*

The proof of Lemma 16 can be found in the full version [4]. There are two cases in this lemma, depending on whether one or two steps the first time the 2-Dyck path goes below the end of the adjacent move, which we call Type-1 and Type-2 moves. See Figure 5 (1a) and (2a) for illustrations of the two types of transitions. The flip moves to simulate the adjacent move are also intuitive. See the illustrations in Figure 5 (1b) and (2b).

The next step is to resolve the shifting, by designing a sequence of flippings that simulates it. We call this sequence of flip moves a *shift sequence*. The design of the shift sequence is the trickiest part of the whole comparison argument, because we have to be able to recover the initial and final states by just looking at a transition with some succinct extra encoding. We only need to deal with right shifting, as for left shifting we just reverse the sequence.

► **Lemma 17.** *Let T be a non-crossing spanning tree of m edges decomposing into (T_A, T_B, T_C) by Fact 8, where T_B and T_C are empty, and T_A is not empty. Then we can use $\ell \leq 3m$ flip moves to simulate a right shifting, i.e., reach a non-crossing spanning tree T' decomposing into (T'_A, T'_B, T'_C) where both T'_A and T'_C are empty, and $T'_B = T_A$.*

The proof of Lemma 17 can be found in the full version [4]. Essentially, we design a recursive way of simulating shifts. We decompose the subtree to shift into minimal segments, and then move them one at a time. To move a minimal segment, we perform edge flips so that subtrees of the minimal segment are in the form of a shift, and finish moving it using recursion. See Figure 6 for an illustration of an example.

■ **Algorithm 1** Constructing the canonical path.

Input: two 2-Dyck paths $I \neq F$ of length $3n$ that $P_{AM}(I, F) > 0$
Output: a canonical path $I = Z_0, Z_1, \dots, Z_\ell$

- 1 Decide if I to F is a Type-1 or Type-2 move as in Lemma 16;
- 2 **if** *it is a Type-1 move* **then** **return** Transition1(I, F) **else** **return** Transition2(I, F);
- 3 **procedure** Transition1(I, F): // Requiring $I \rightarrow F$ to be a Type-1 move.
- 4 Let $Z_0 \leftarrow I$;
- 5 Find $a = (a_1, a_2)$, $e = (p, g_e + 1)$ and g_a as in Lemma 16;
- 6 **if** $p \leq g_e - 1$ **then** let $(\ell - 2, Z_1, \dots, Z_{\ell-2}) \leftarrow \text{ShiftRight}(p, g_e + 1, 1, 1)$;
- 7 Perform an M_1 move of e to obtain $Z_{\ell-1}$;
- 8 Let $Z_\ell \leftarrow F$, which is obtained by an M_2 move of e from $Z_{\ell-1}$;
- 9 **return** $Z_0, \dots, Z_\ell$;
- 10 **procedure** ShiftRight($s, t, \text{depth}, \text{step}$):
- 11 Let $k \geq 1$ be the number of minimal segments $S_1, \dots, S_k$ beneath the edge (s, t) ;
- 12 last $\leftarrow \text{step}$;
- 13 **for** i from k down to 1 **do**
- 14 Let $e = (p, q)$ be the pivot edge of S_i and g_e, g_{e+1} be the gap beneath e ;
 // (#0)
- 15 **if** $q \geq g_e + 2$ **then**
- 16 Perform an M_5 move of e to obtain Z_{last} , and last $\leftarrow \text{last} + 1$; // (#1)
- 17 **if** $p \leq g_e - 1$ **then**
- 18 $(\ell', Z_{\text{last}}, \dots, Z_{\text{last}+\ell'-1}) \leftarrow \text{ShiftRight}(p, g_e + 1, \text{depth} + 1, \text{last})$;
- 19 Update last $\leftarrow \text{last} + \ell'$; // (#2)
- 20 Perform an M_6 move of e to obtain Z_{last} , and update last $\leftarrow \text{last} + 1$; // (#3)
- 21 **if** $q \geq g_e + 2$ **then**
- 22 $(\ell'', Z_{\text{last}}, \dots, Z_{\text{last}+\ell''-1}) \leftarrow \text{ShiftRight}(g_e + 1, q + 1, \text{depth} + 1, \text{last})$;
- 23 Update last $\leftarrow \text{last} + \ell''$; // (#4)
- 24 **if** $p \leq g_e - 1$ **then**
- 25 Perform an M_7 move of e to obtain Z_{last} , and last $\leftarrow \text{last} + 1$; // (#5)
- 26 **return** (last - step, $Z_{\text{step}}, \dots, Z_{\text{last}}$);

5.2 Path construction

We are going to apply the path method, namely Theorem 3, to compare P_{AM} and P_{FM} . Because of the bijection between 2-Dyck paths and NCSTs, we simply treat P_{AM} and P_{FM} as having the same stationary distribution. To compare the two chains, we will need to construct a path (namely, a sequence of transitions) for any two 2-Dyck paths I and F where $P_{AM}(I, F) > 0$ using only transitions in P_{FM} . We summarise the construction in Algorithm 1, with half of the procedure omitted due to redundancy. This construction follows closely the simulation of an adjacent move via two flip moves and one shifting in Lemma 16, and then simulating the shifting via flip moves in Lemma 17. The edges $a = (a_1, a_2)$, $e = (p, g_e + 1)$ and g_a in Transition1(I, F) are defined in the proof of Lemma 16. Intuitively, a and e are the distinguished edges marked in Figure 5, and $g_a, g_a + 1$ is the unique gap beneath a . The markers (#0) to (#5) are used in the proof of Lemma 19. We expand out the Type-1 move

and omit the details of the Type-2 move to avoid redundancy. (The flip moves for Type-2 moves are described in the proof of Lemma 16 and Figure 5(2*), and the `ShiftLeft` sequence is the reverse of `ShiftRight`.)

Let T_I and T_F be the non-crossing spanning trees corresponding to I and F . The construction yields a path $T_I = Z_0 \rightarrow Z_1 \rightarrow \dots \rightarrow Z_\ell = T_F$ where $P_{\text{FM}}(Z_i, Z_{i+1}) > 0$ for any $i \leq \ell - 1$. Its validity, namely that each transition is valid, is implied by Lemma 16 and Lemma 17. These two lemmas also imply an upper bound on the length of the path $\ell \leq 3n + 2$ ($3n$ for the shift sequence and 2 for the two simple flips).

5.3 Bound the congestion

The goal of the Markov chain comparison argument is to show the following lemma:

► **Lemma 18.** *Let $\mathcal{E}_{\text{AM}}$ be the Dirichlet form of the adjacent move chain P_{AM} over the uniform distribution of 2-Dyck paths of length $3n$, and $\mathcal{E}_{\text{FM}}$ be the Dirichlet form of the flip chain P_{FM} over the uniform distribution of non-crossing spanning trees containing n edges. Then for all functions f , it holds that $\mathcal{E}_{\text{FM}}(f, f) \geq \Omega(n^{-4}) \cdot \mathcal{E}_{\text{AM}}(f, f)$.*

Lemma 18 is a straightforward application of Theorem 3 where we bound the congestion using Equation (8) and Corollary 20. The proof can be found in the full version [4].

The next lemma shows that our canonical path construction ensures that each transition in the flip move P_{FM} is used by a limited number of (I, F) pairs in the adjacent move P_{AM} . This is proved by designing an encoding such that we can uniquely recover the (I, F) pair from the current transition and the encoding. The number of (I, F) pairs is then bounded by the number of possible encodings, which we ensure to be linear in n .

► **Lemma 19 (encoding).** *Let $Z \rightarrow Z'$ be a transition of P_{FM} over non-crossing spanning trees containing n edges. Let $I \neq F$ be two 2-Dyck paths such that $P_{\text{AM}}(I, F) > 0$, and that the path constructed according to Algorithm 1 contains the transition $Z \rightarrow Z'$. Then there exists an injective mapping*

$$\varphi_{Z \rightarrow Z'} : \mathcal{T}_n \times \mathcal{T}_n \rightarrow \{\text{Left}, \text{Right}\} \times \{M_1, M_2, M_3, M_4, S_1, S_2\} \times [n]$$

where $\mathcal{T}_n$ is the set of all non-crossing spanning trees of n edges. In other words, given $Z \rightarrow Z'$, and (dir, M, d) , we can uniquely determine I, F such that $\varphi_{Z \rightarrow Z'}(T_I, T_F) = (\text{dir}, M, d)$, where $\text{dir} \in \{\text{Left}, \text{Right}\}$, $M \in \{M_1, M_2, M_3, M_4, S_1, S_2\}$, and $d \in [n]$.

Proof sketch. We provide a proof sketch here, and the full proof can be found in the full version [4]. The mapping is defined as $\varphi_{Z \rightarrow Z'}(T_I, T_F) = (\text{dir}, M, d)$, where

- dir indicates the direction of adjacent move from I to F ,
- M indicates where the current transition $Z \rightarrow Z'$ is located in the construction of the canonical path, as in Algorithm 1, and
- d indicates the depth of the recursion of the shift sequence construction. It is defined only when $M \in \{S_1, S_2\}$, and otherwise arbitrary.

In particular, the six possibilities of M correspond to the six types of moves in Figure 5 (1b) and (2b), where $M_1, \dots, M_4$ denote edge flips, and S_1 and S_2 denote shifting sequences.

To show the mapping is injective, we just need to be able to recover T_I and T_F given (dir, M, d) . The transition $Z \rightarrow Z'$ gives the edge which is being moved in this step, and dir indicates if it is a left shifting or right shifting. First, we use M to determine where the current move is in the construction of the canonical path. If it is neither S_1 nor S_2 , then the initial and final states can be reconstructed easily. Otherwise, the current step is in the

recursive construction of the canonical path. Starting from this edge being moved, together with the number d in the mapping above, we can find the sequence of overarching edges which the initial shifting involves. Furthermore, by looking at the current edge and the gap beneath its overarching edge, we can determine which step in the subroutine **ShiftRight** (or **ShiftLeft**) the current transition is in. This information allows us to rewind all the edge flips that have been performed and recover the initial state T_I . The final state T_F is recovered similarly. ◀

Let Γ_{IF} be the path from I to F . We write $(Z, Z') \in \Gamma_{IF}$ if the (Z, Z') transition is in the path Γ_{IF} . Lemma 19 immediately implies the following.

► **Corollary 20.** *For each (Z, Z') with $P_{\text{FM}}(Z, Z') > 0$, there are at most $12n$ pairs (I, F) such that $(Z, Z') \in \Gamma_{IF}$.*

Proof of Theorem 1. We have

$$\begin{aligned}
 t_{\text{mix}}(P_{\text{FM}}) &\leq \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{1}{\lambda(P_{\text{FM}})} = \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{1}{\inf \left\{ \frac{\mathcal{E}_{\text{FM}}(f, f)}{\text{Var}_{\pi}(f)} \right\}} && \text{(by (2))} \\
 &\leq \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{O(n^4)}{\inf \left\{ \frac{\mathcal{E}_{\text{AM}}(f, f)}{\text{Var}_{\pi}(f)} \right\}} && \text{(by Lemma 18)} \\
 &= \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{O(n^4)}{\lambda(P_{\text{AM}})} && \text{(by (2))} \\
 &= \left(1 + \frac{1}{2} \log \pi_{\min}\right) O(n^7 \log n) && \text{(by Corollary 14)} \\
 &= O(n^8 \log n). && \blacktriangleleft
 \end{aligned}$$

References

- 1 Oswin Aichholzer, Brad Ballinger, Therese Biedl, Mirela Damian, Erik D. Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Josef Tkadlec, and Yushi Uno. Reconfiguration of non-crossing spanning trees. *arXiv*, abs/2206.03879, 2022. doi:10.48550/arXiv.2206.03879.
- 2 Victor Alvarez, Karl Bringmann, Radu Curticapean, and Saurabh Ray. Counting triangulations and other crossing-free structures via onion layers. *Discret. Comput. Geom.*, 53(4):675–690, 2015. doi:10.1007/S00454-015-9672-3.
- 3 Victor Alvarez, Karl Bringmann, Saurabh Ray, and Raimund Seidel. Counting triangulations and other crossing-free structures approximately. *Comput. Geom.*, 48(5):386–397, 2015. doi:10.1016/J.COMGEO.2014.12.006.
- 4 Konrad Anand, Weiming Feng, Graham Freifeld, Heng Guo, Mark Jerrum, and Jiaheng Wang. Rapid mixing of the flip chain over non-crossing spanning trees. *arXiv*, abs/2409.07892, 2024. doi:10.48550/arXiv.2409.07892.
- 5 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *FOCS*, pages 1319–1330. IEEE, 2020. doi:10.1109/FOCS46700.2020.00125.
- 6 Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: high-dimensional walks and an FPRAS for counting bases of a matroid. In *STOC*, pages 1–12. ACM, 2019. doi:10.1145/3313276.3316385.
- 7 Federico Ardila. The Catalan matroid. *J. Comb. Theory, Ser. A*, 104(1):49–62, 2003. doi:10.1016/S0097-3165(03)00121-3.
- 8 David Avis and Komei Fukuda. Reverse search for enumeration. *Discret. Appl. Math.*, 65(1-3):21–46, 1996. doi:10.1016/0166-218X(95)00026-N.

- 9 Elia Bisi, Piotr Dyszewski, Nina Gantert, Samuel G. G. Johnston, Joscha Prochno, and Dominik Schmid. Random planar trees and the Jacobian conjecture. *arXiv*, abs/2301.08221, 2023. [arXiv:2301.08221](#).
- 10 Håvard Bakke Bjerkevik, Linda Kleist, Torsten Ueckerdt, and Birgit Vogtenhuber. Flipping non-crossing spanning trees. In *SODA*, pages 2313–2325. SIAM, 2025. [doi:10.1137/1.9781611978322.77](#).
- 11 Nicolas Bousquet, Lucas de Meyer, Théo Pierron, and Alexandra Wesolek. Reconfiguration of plane trees in convex geometric graphs. In *SoCG*, volume 293 of *LIPICs*, pages 22:1–22:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. [doi:10.4230/LIPICs.SOCG.2024.22](#).
- 12 Nicolas Bousquet, Valentin Gledel, Jonathan Narboni, and Théo Pierron. A note on the flip distance between non-crossing spanning trees. *Comput. Geom. Topol.*, 2(1):8:1–8:7, 2023. URL: <https://www.cgt-journal.org/index.php/cgt/article/view/36>.
- 13 Sergey Bravyi, Libor Caha, Ramis Movassagh, Daniel Nagaj, and Peter W. Shor. Criticality without frustration for quantum spin-1 chains. *Physical review letters*, 109(20):207202, 2012.
- 14 Thomas Budzinski. On the mixing time of the flip walk on triangulations of the sphere. *C. R. Math. Acad. Sci. Paris*, 355(4):464–471, 2017. [doi:10.1016/j.crma.2017.02.011](#).
- 15 Pietro Caputo, Fabio Martinelli, Alistair Sinclair, and Alexandre Stauffer. Random lattice triangulations: structure and algorithms. *Ann. Appl. Probab.*, 25(3):1650–1685, 2015. [doi:10.1214/14-AAP1033](#).
- 16 Alessandra Caraceni and Alexandre Stauffer. Polynomial mixing time of edge flips on quadrangulations. *Probab. Theory Related Fields*, 176(1-2):35–76, 2020.
- 17 Swee Hong Chan and Igor Pak. Computational complexity of counting coincidences. *Theor. Comput. Sci.*, 1015:114776, 2024. [doi:10.1016/J.TCS.2024.114776](#).
- 18 Emma Cohen. *Problems in Catalan mixing and matchings in regular hypergraphs*. PhD thesis, Georgia Institute of Technology, 2016.
- 19 Mary Cryan, Heng Guo, and Giorgos Mousa. Modified log-Sobolev inequalities for strongly log-concave distributions. *Ann. Probab.*, 49(1):506–525, 2021.
- 20 Dimitrios I. Dais. Resolving 3-dimensional toric singularities. In *Geometry of toric varieties*, volume 6 of *Sémin. Congr.*, pages 155–186. Soc. Math. France, Paris, 2002.
- 21 Jesús A. De Loera, Jörg Rambau, and Francisco Santos. *Triangulations*, volume 25 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2010. Structures for algorithms and applications.
- 22 Mehmet Demirtas, Liam McAllister, and Andres Rios-Tascon. Bounding the Kreuzer-Skarke landscape. *Fortschr. Phys.*, 68(11-12):2000086, 13, 2020. [doi:10.1002/prop.202000086](#).
- 23 Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993.
- 24 Martin E. Dyer, Alan M. Frieze, and Ravi Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. [doi:10.1145/102782.102783](#).
- 25 David Eppstein. Counting polygon triangulations is hard. *Discret. Comput. Geom.*, 64(4):1210–1234, 2020. [doi:10.1007/S00454-020-00251-7](#).
- 26 David Eppstein and Daniel Frishberg. Improved mixing for the convex polygon triangulation flip walk. In *ICALP*, volume 261 of *LIPICs*, pages 56:1–56:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. [doi:10.4230/LIPICs.ICALP.2023.56](#).
- 27 I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, resultants and multidimensional determinants*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, 2008. Reprint of the 1994 edition.
- 28 M. Carmen Hernando, Ferran Hurtado, Alberto Márquez, Mercè Mora, and Marc Noy. Geometric tree graphs of points in convex position. *Discret. Appl. Math.*, 93(1):51–66, 1999. [doi:10.1016/S0166-218X\(99\)00006-2](#).

- 29 Vishesh Jain, Marcus Michelen, Huy Tuan Pham, and Thuy-Duong Vuong. Optimal mixing of the down-up walk on independent sets of a given size. In *FOCS*, pages 1665–1681. IEEE, 2023. doi:10.1109/FOCS57990.2023.00101.
- 30 Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989. doi:10.1137/0218077.
- 31 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993. doi:10.1137/0222066.
- 32 J. Jurkiewicz, A. Krzywicki, and B. Petersson. A numerical study of discrete Euclidean Polyakov surfaces. *Physics Letters B*, 168(3):273–278, 1986. doi:10.1016/0370-2693(86)90978-0.
- 33 Volker Kaibel and Günter M. Ziegler. Counting lattice triangulations. In *Surveys in combinatorics, 2003 (Bangor)*, volume 307 of *London Math. Soc. Lecture Note Ser.*, pages 277–307. Cambridge Univ. Press, Cambridge, 2003.
- 34 Marek Karpinski, Andrzej Lingas, and Dmitriy Sledneu. A QPTAS for the base of the number of crossing-free structures on a planar point set. *Theor. Comput. Sci.*, 711:56–65, 2018. doi:10.1016/J.TCS.2017.11.003.
- 35 V.A. Kazakov, I.K. Kostov, and A.A. Migdal. Critical properties of randomly triangulated planar random surfaces. *Physics Letters B*, 157(4):295–300, 1985. doi:10.1016/0370-2693(85)90669-0.
- 36 David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2017.
- 37 Russell A. Martin and Dana Randall. Sampling adsorbing staircase walks using a new Markov chain decomposition method. In *FOCS*, pages 492–502. IEEE Computer Society, 2000. doi:10.1109/SFCS.2000.892137.
- 38 Lisa McShine and Prasad Tetali. On the mixing time of the triangulation walk and other Catalan structures. In *Randomization Methods in Algorithm Design*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 147–160. DIMACS/AMS, 1997. doi:10.1090/DIMACS/043/09.
- 39 Michael Molloy, Bruce Reed, and William Steiger. On the mixing rate of the triangulation walk. In *Randomization Methods in Algorithm Design*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 179–190. DIMACS/AMS, 1997. doi:10.1090/DIMACS/043/11.
- 40 Ramis Movassagh. The gap of Fredkin quantum spin chain is polynomially small. *Ann. Math. Sci. Appl.*, 3(2):531–562, 2018. doi:10.4310/AMSA.2018.v3.n2.a5.
- 41 Ramis Movassagh and Peter W. Shor. Supercritical entanglement in local systems: Counterexample to the area law for quantum matter. *Proceedings of the National Academy of Sciences*, 113(47):13278–13282, 2016. doi:10.1073/pnas.1605716113.
- 42 Marc Noy. Enumeration of noncrossing trees on a circle. *Discret. Math.*, 180(1-3):301–313, 1998. doi:10.1016/S0012-365X(97)00121-0.
- 43 Micha Sharir and Adam Sheffer. Counting triangulations of planar point sets. *Electron. J. Comb.*, 18(1), 2011. doi:10.37236/557.
- 44 Micha Sharir, Adam Sheffer, and Emo Welzl. On degrees in random triangulations of point sets. *J. Comb. Theory, Ser. A*, 118(7):1979–1999, 2011. doi:10.1016/J.JCTA.2011.04.002.
- 45 Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999.
- 46 Richard P. Stanley. *Catalan numbers*. Cambridge University Press, New York, 2015.
- 47 O. Ya. Viro. Real plane algebraic curves: constructions with controlled topology. *Algebra i Analiz*, 1(5):1–73, 1989.
- 48 David Bruce Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Ann. Appl. Probab.*, 14(1):274–325, 2004. doi:10.1214/aoap/1075828054.