

RAPID MIXING OF THE FLIP CHAIN OVER NON-CROSSING SPANNING TREES

KONRAD ANAND, WEIMING FENG, GRAHAM FREIFELD, HENG GUO, MARK JERRUM, JIAHENG WANG

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).

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 [AF96] first proved that it takes at most $2n - 2$ flips to move from one to the other. On the other hand, Hernando, Hurtado, Márquez, Mora and Noy [HHM⁺99] 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 [ABB⁺22, BGNP23, BdMPW24]. Very recently, Bjerkevik, Kleist, Ueckerdt, and Vogtenhuber [BKUV25] 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.

(Mark Jerrum) SCHOOL OF MATHEMATICAL SCIENCES, QUEEN MARY UNIVERSITY OF LONDON, MILE END ROAD, LONDON, E1 4NS, UNITED KINGDOM.

(Weiming Feng) SCHOOL OF COMPUTING AND DATA SCIENCE, THE UNIVERSITY OF HONG KONG, POKFULAM ROAD, HONG KONG, CHINA.

(Konrad Anand, Graham Freifeld, Heng Guo) SCHOOL OF INFORMATICS, UNIVERSITY OF EDINBURGH, INFORMATICS FORUM, EDINBURGH, EH8 9AB, UNITED KINGDOM.

(Jiaheng Wang) FACULTY OF INFORMATICS AND DATA SCIENCE, UNIVERSITY OF REGENSBURG, BAJUWARENSTRAË 4, 93053 REGENSBURG, GERMANY.

E-mail address: konrad.anand@me.com, wfeng@hku.hk, g.freifeld@sms.ed.ac.uk, hguo@inf.ed.ac.uk, m.jerrum@qmul.ac.uk, pw384@hotmail.com.

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). 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 also acknowledges support from the ERC (grant agreement No. 101077083).

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. [KZ03, SS11, SSW11]). Motivations for counting these structures have appeared in other topics, such as quantum gravity [JKP86, KKM85], plane curves [Vir89], algebraic geometry [Dai02, DMRT20], and the theory of discriminants [GKZ08]. 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 #P-hard [Epp20]. On the other hand, counting the number of triangulations of a given set of points remains unresolved (but it is conjectured to be #P-hard, see e.g. [DRS10] or [CP24, §6.7]). Similarly, counting NCSTs over a set of points in general position apparently is also not known to be #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 [ABCR15, ABRS15, KLS18], 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. [Noy98])

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

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 [Sta15] 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 (10).) Reconfiguration results [AF96, HHM⁺99, BKUV25] 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 [JS89, DFK91, JS93]. 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 [MRS97, MT97, EF23], over triangulations on lattices and spheres [CMSS15, Bud17], or over lattice paths and (1-)Dyck paths [MR00, Wil04, CS20]. See Cohen’s thesis [Coh16] 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 [BCM⁺12, MS16, Mov18] or even in algebraic geometry [BDG⁺23].

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 $O(n^8 \log n)$.*

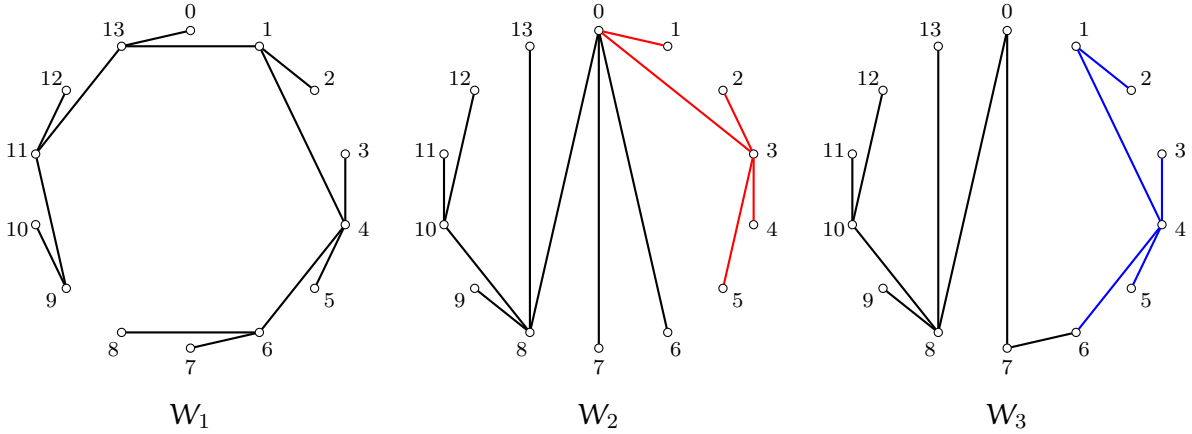


FIGURE 2. NCSTs corresponding to the 2-Dyck paths in Theorem 2.

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 (12)).

To apply the path method for Markov chain comparison [DS93], 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.

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 resulted 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 [MT97] is similar to our result in spirit, which is also built on a comparison argument with Wilson’s lattice path chain [Wil04]. 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 [ALOV19, CGM21]. 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 [ALOV19, CGM21] for analysing Markov chains, most notable among which is the spectral independence technique [ALO20]. 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 [JMPV23]. 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 [AF96, BKUV25] 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 [LP17]. 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

$$(2) \quad d_{\text{TV}}(\pi, \pi') := \frac{1}{2} \sum_{x \in \Omega} |\pi(x) - \pi'(x)|.$$

For a Markov chain P with stationary distribution π , let

$$(3) \quad d(t) = \max_{x \in \Omega} \{d_{\text{TV}}(P^t(x, \cdot), \pi)\}.$$

Then the mixing time of P is defined as $t_{\text{mix}}(P) := \min\{t : d(t) \leq \frac{1}{4}\}$. The constant $1/4$ 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

$$(4) \quad \lambda(P) := \inf \left\{ \frac{\mathcal{E}_P(f, f)}{\text{Var}_{\pi}(f)} \mid f : \Omega \rightarrow \mathbb{R}, \text{Var}_{\pi}(f) \neq 0 \right\}.$$

¹One 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 the reversible Markov chain P can be bounded using the inverse spectral gap (also called the *relaxation time*) by (see e.g. [LP17, Theorem 12.4])

$$(5) \quad t_{\text{mix}}(P) \leq \frac{1}{\lambda(P)} \left(1 + \frac{1}{2} \log \pi_{\min} \right),$$

where $\pi_{\min} := \min_{x \in \Omega} \pi(x)$. This π_{\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.

The relaxation time also provides a lower bound for the mixing time. For any reversible Markov chain, we have (see e.g. [LP17, Theorem 12.5])

$$(6) \quad \frac{1}{\lambda(P)} \leq 1 + 2t_{\text{mix}}.$$

2.1. Path method. The path method [DS93] 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 [JS89].

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

$$(7) \quad B := \max_{(x',y'):\mathbb{P}(x',y')>0} \left(\frac{1}{\pi(x')\mathbb{P}(x',y')} \sum_{\substack{x,y \\ \Gamma_{xy} \ni (x',y')}} \tilde{\pi}(x) \tilde{\mathbb{P}}(x,y) |\Gamma_{xy}| \right).$$

Theorem 3 (Theorem 2.1 of [DS93]). *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. Different authors appear to index Fuss-Catalan numbers differently in the literature. We adopt the one used in Stanley's monograph [Sta15].

3.1. Dyck paths.

Definition 4 (k -Dyck path). Let k be a positive integer. A sequence $a_1, a_2, \dots, a_{(k+1)n} \in \{+1, -k\}$ 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 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_k . 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 5. $\nearrow \nearrow \nearrow \searrow_2 \nearrow \searrow_2 \nearrow \searrow_2 \nearrow \nearrow \nearrow \searrow_2 \nearrow \searrow_2 \searrow_2 \nearrow \searrow_2 \searrow_2 \nearrow \searrow_2$ is a 2-Dyck path.




FIGURE 3. Decomposing Theorem 5 by Theorem 6.

The following simple fact is useful later.

Fact 6. 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.

Theorem 6 immediately implies the recurrence on the number of k -Dyck paths of length $(k + 1)n$, denoted by $C_{k,n}$:

$$(8) \quad C_{k,0} = 1, \quad C_{k,n} = \sum_{\substack{i_1, \dots, i_{k+1} \in \mathbb{Z}_{\geq 0}, \\ \sum_{j=1}^{k+1} i_j = n-1}} C_{k,i_1} C_{k,i_2} \cdots C_{k,i_{k+1}}, \quad \implies \quad C_{k,n} = \frac{1}{kn+1} \binom{(k+1)n}{n}.$$

This is called the k -Fuss-Catalan number [Sta15, 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 [Sta99, 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 zero edge) if $s = t$.

Fact 7. 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 Theorem 6 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 8. *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 Theorem 7, 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 9. *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).*

Proof. We perform an induction on $k := |\{i \geq 1 : ht_U(i) = 0\}|$, the number of times that U hits height zero after the first step.

In the base case where $k = 1$, the decomposition of U is $U = \nearrow U_A \nearrow U_B \searrow U_C$ with U_C being empty. Therefore, the decomposition of UV is $UV = \nearrow U_A \nearrow U_B \searrow V$, which by Theorem 8 gives a decomposition of T_{UV} into (T_{U_A}, T_{U_B}, T_V) , a concatenation of T_U and T_V .

Suppose that this is true for any $k' < k$. Decompose $U = \nearrow U_A \nearrow U_B \searrow U_C =: U'U_C$. On one hand, the number of times that U_C hits height zero is $k - 1$. By induction hypothesis for the case $k' = k - 1$, the tree $T_{U_C V}$ is a concatenation of the tree T_{U_C} and T_V . On the other hand, the decomposition of UV is $UV = \nearrow U_A \nearrow U_B \searrow (U_C V) = U'(U_C V)$, and note that U' hits height zero once. By induction hypothesis for the case $k' = 1$, we have that $T_{UV} = T_{U'(U_C V)}$ is a concatenation of $T_{U'}$ and $T_{U_C V}$. The induction hypothesis is hence true for the case $k' = k$ by putting the above two points together. \square

Inspired by the decomposition, we introduce the following terminology that will be useful later.

Definition 10. In the context of Theorem 7:

- (pivot edge) the edge $(0, t)$ is called the *pivot edge* of T ;
- (gap) the gap after s is 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) .

Example 11. 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 $[6, 7]$.




FIGURE 4. (a)(b) Two ways of representing a non-crossing spanning tree. The tree is corresponded to the 2-Dyck path in Theorem 5. (c) Decomposing the tree by Theorem 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 given by

$$(9) \quad P_{AM}(X, X') = \begin{cases} 1 - \frac{|B(X)|}{6n - 2} & \text{if } X = X'; \\ \frac{1}{6n - 2} & \text{if } X' \in B(X); \\ 0 & \text{otherwise.} \end{cases}$$

Note that $P_{AM}(X, X) \geq 4/5$. 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 12. $t_{\text{mix}}(P_{AM}) = O(n^3 \log n)$.

The proof essentially follows Wilson's coupling argument [Wil04]. The core is to define a potential function measuring the difference between two coupled Markov chains. Recall that

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

where each α_i is either 1 (for an up-step) or -2 for a down-step. We say that a path X *dominates* another path Y if $\text{ht}_X(i) \geq \text{ht}_Y(i)$ for any i . Wilson defines a potential $\Phi(X, Y)$ (only when X dominates Y) which is a weighted sum of height differences. When we apply it to 2-Dyck paths (a subset of lattice paths on a lattice with width n and height $2n$), it maintains the following properties:

- (a) $\Phi(X, Y) \geq 0$, and the equality is taken if and only if $X = Y$;
- (b) $\sin(\pi/(3n)) \leq \Phi(X, Y) \leq 4n^2/27$ for any pair of two different 2-Dyck paths of length $3n$.

Consider two instances of Markov chains X_t and Y_t where the former dominates the latter. In the next step, the coupling between X_{t+1} and Y_{t+1} is defined as follows:

- choose the same positions in both chains for the potential swap;
- if, in either X_t or Y_t (or both), the two positions to swap have the same arrow, then perform the transition in the other chain (or neither) as normal;
- otherwise, choose either $\nearrow\searrow$ or $\searrow\nearrow$ with probability $1/2$, respectively, and change the two positions into the pattern chosen in both chains. For any chain resulting in an invalid Dyck path, revert the swap in the invalid chain.

Let Φ_t be the shorthand for $\Phi(X_t, Y_t)$. Wilson also showed that

- (c) Given any pair of X_t, Y_t where X_t dominates Y_t , obtaining X_{t+1} and Y_{t+1} in the above way. Then X_{t+1} still dominates Y_{t+1} , and it holds that

$$\mathbb{E}[\Phi_{t+1} - \Phi_t \mid X_t, Y_t] \leq -\frac{\pi^2}{2n^3} \Phi_t.$$

The coupling here differs from the one in Wilson's in the extra check of whether the swap results in a valid path. This does break some statements in Wilson's paper, such as the mixing time lower bound, but it is easy to verify that the properties a), b), and c) still hold.

Proof of Theorem 12. Consider two instances of the adjacent move chain, whose initial states are X_0 and Y_0 . Here X_0 consists of $2n$ \nearrow s followed by n \searrow s, and Y_0 consists of $\nearrow\nearrow\searrow$ repeated n times. It is easy to verify that any valid 2-Dyck path is dominated by X_0 and dominates Y_0 . By the above properties (b) and (c), we have

$$\sin\left(\frac{\pi}{3n}\right) \Pr[\Phi_t > 0] \leq \mathbb{E}[\Phi_t] \leq \frac{4n^2}{27} \exp\left\{-\frac{\pi^2 t}{2n^3}\right\}.$$

Therefore, after $t > (2/\pi^2)n^3 \log(16n^2/27/\sin(\pi/(3n))) = \Theta(n^3 \log n)$ steps, the probability that Φ_t is non-zero, and hence that X_t and Y_t do not coalesce yet, is at most $1/4$. \square

Corollary 13. $1/\lambda(P_{AM}) = O(n^3 \log n)$.

Proof. This follows immediately from Theorem 12 and (6). \square

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 Theorem 13. 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

$$(10) \quad P_{FM}(S, T) = \begin{cases} \sum_{T': |S \cap T'| = n-1} \frac{1}{n\delta(S, T')} & \text{if } S = T; \\ \frac{1}{n\delta(S, T)} & \text{if } |S \cap T| = n-1; \\ 0 & \text{otherwise,} \end{cases}$$

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

$$(11) \quad P_{FM}(S, S) \geq 1/n, \quad \text{and } P_{FM}(S, T) \geq \Omega(n^{-3}) \text{ if it is non-zero.}$$

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

$$(12) \quad \text{moving from } \begin{array}{c} \text{---} \\ \text{---} \end{array} \text{ to } \begin{array}{c} \text{---} \\ \text{---} \end{array} \text{ or vice versa,}$$

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 14 (shifting). Let T be a non-crossing spanning tree of size m decomposing into (T_A, T_B, T_C) by Theorem 7, 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.

Lemma 15. *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.*

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).

Proof. Without loss of generality, assume F is reached from I with an adjacent move to the left, so $F = (I \setminus \{y\}) \cup \{x\}$ where $y = x + 1$. Let \tilde{I} be the sub-Dyck-path of I from x' to y' , obtained in the following way.

- $y' > y$ is the minimum y' such that $\text{ht}_I(y') = \text{ht}_I(y) - 1$ or $\text{ht}_I(y') = \text{ht}_I(y) - 2$.
- $x' < x$ is the maximum x' such that $\text{ht}_I(x') = \text{ht}_I(y')$.

Note that such y' must exist, because the down-step only changes the height by -2 . The point here is that we want to find x' and y' so that our transition can be run only on the sub-Dyck path (or subtree) made by the interval (x', y') without changing anything on the rest of the structure.

We call the transition $I \rightarrow F$ **Type 1** if $\text{ht}_I(y') = \text{ht}_I(y) - 1$, and **Type 2** if $\text{ht}_I(y') = \text{ht}_I(y) - 2$. See Figure 5 (1a) and (2a) for examples. We will decompose our sub-Dyck path again into blocks of unaffected sub-sub-Dyck paths (A_i 's, B_i 's) which do not change in their interior, and remaining arrows which matter in our transition.

For a **Type 1** transition, the sub-Dyck-path \tilde{I} can be written as $\nearrow A_1 \nearrow A_2 \searrow B_1 \nearrow B_2 \searrow$, which becomes $\tilde{F} = \nearrow A_1 \nearrow A_2 \searrow B_1 \nearrow B_2 \searrow$ after the adjacent move. See Figure 5 (1a). In $T_{\tilde{I}}$, we find the subtrees $T_{A_1}, T_{A_2}, T_{B_1}$ and T_{B_2} that correspond to A_1, A_2, B_1 and B_2 respectively. To see this, first notice that $A_1 \nearrow A_2 \searrow B_1$ is the “ A_1 part” in Theorem 6, and B_2 is the “ A_2 part” in Theorem 6. The “ B part” is after y' and is not relevant here, so we do not draw it out. By Theorem 9, we further decompose the former into A_1 and $\nearrow A_2 \searrow B_1$. Notice that A_1 may not be a minimal segment here. Then we use Theorem 6 on $\nearrow A_2 \searrow B_1$ to further decompose it.

There are two distinguished edges a and e in the corresponding tree as labelled in Figure 5 (1a). The edge $a = (a_1, a_2)$ is the pivot edge of $T_{\tilde{I}}$. Denote by $g_a, g_a + 1$ the corresponding gap beneath a . The edge $e = (p, g_e + 1)$ has the overarching edge a . Denote by $g_e, g_e + 1$ the corresponding gap beneath e . We can alternatively say that e is the edge between A_2 and B_1 , and a is its overarching edge. Therefore, the subtree T_{A_1} is $[a_1, p]$, the subtree T_{A_2} is $[p, g_e]$, the subtree T_{B_1} is $[g_e + 1, g_a]$, and the subtree T_{B_2} is $[g_a + 1, a_2]$. In $T_{\tilde{F}}$, the three subtrees $T_{A_1}, T_{B_1}, T_{B_2}$ remain untouched, while T_{A_2} starts at position $p + 1$ instead of p . The two edges a and e are moved to new positions, too.

The above change can be simulated as below by using edge flippings and a shifting. See Figure 5(1b).





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**.

- (S₁) The position of e together with T_{A_2} meets the prerequisite for a right shifting, we perform it so long as T_{A_2} contains at least one edge.
- (M₁) Now the subtree T_{A_2} has been shifted right. Let M_1 move e from $(p, g_e + 1)$ to $(a_1, g_e + 1)$. Since the right endpoint is still connected to the subtree induced by $[p + 1, g_a]$, the result is still a tree, and since there is no edge other than a and e going out of T_{A_1} , the result is still non-crossing.
- (M₂) Let M_2 move a from (a_1, a_2) to $(g_e + 1, a_2)$. Since the left endpoint is still connected to the subtree induced by $[a_1, g_a]$, the result is still a tree, and since there are no edges coming out from beneath the current edge e , the result is still non-crossing.

For a **Type 2** transition, the sub-Dyck-path \tilde{I} can be written as $\nearrow A_1 \nearrow A_2 \nearrow A_3 \searrow B \searrow$, which becomes $\tilde{F} = \nearrow A_1 \nearrow A_2 \nearrow A_3 \searrow B \searrow$ after the adjacent move. See Figure 5(2a). Again, in \tilde{T}_I , we find the corresponding subtrees $T_{A_1}, T_{A_2}, T_{A_3}$ and T_B . They are joined by two edges a and e . The edge $a = (a_1, a_2)$ is the pivot edge of \tilde{T}_I . Denote by $g_a, g_a + 1$ the corresponding gap beneath a . The edge $e = (q, g_e + 1)$ has the overarching edge a . Denote by $g_e, g_e + 1$ the corresponding gap beneath e . Therefore, the subtree T_{A_1} is $[a_1, g_a]$, the subtree T_{A_2} is $[g_a + 1, q]$, the subtree T_{A_3} is $[q, g_e]$, and the subtree T_B is $[g_e + 1, a_2]$. In \tilde{T}_F , the three subtrees T_{A_1}, T_{A_3}, T_B remain untouched, while T_{A_2} starts at position g_a instead of $g_a + 1$. The edge e is moved to a new position.

The above change can be simulated as below by using edge flippings and a shifting. See Figure 5(2b).

- (M₃) Let M_3 move e from $(q, g_e + 1)$ to (g_a, q) . Since e is still has exactly one endpoint connected to the subtree induced by $[g_a + 1, g_e]$, the result is still a tree, and since there is no edge other than e coming out from A_2 , the result is still non-crossing.
- (S₂) The position of e together with T_{A_2} now meets the prerequisite for a left shifting, we perform it so long as T_{A_2} contains at least one edge.
- (M₄) Let M_4 move e from (g_a, q) to (g_a, g_e) . Since there is no edge other than e coming out from the subtree T_{A_3} , the result is still non-crossing.

These validate the lemma. □

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 16. *Let T be a non-crossing spanning tree of m edges decomposing into (T_A, T_B, T_C) by Theorem 7, 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 $T'_B = T_A$.*

To prove this, 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.

Proof. We construct the sequence by induction on the number of edges of the tree. Let T be a non-crossing spanning tree where T_B and T_C are empty. Denote by $a = (a_1, a_2)$ the pivot edge of T . Let S_1, \dots, S_k be the minimal segments of T_A from left to right. We perform the following one by one for each minimal segment, such that after finishing the first t minimal segments, we obtain a non-crossing spanning tree T'' such that T''_A contains S_1, \dots, S_{k-t} , and T''_B contains S_{k-t+1}, \dots, S_k .

For each current segment $S = S_i$, let e be its pivot edge, and we perform a sequence of three flip moves, M_5, M_6, M_7 , with recursively constructed shift sequences under e after M_5 and M_6 , respectively, to shift the subtrees of S one by one. Note that the overarching edge of e is a . Decompose S into (S_A, S_B, \emptyset) ; the last component is empty because S is a minimal segment. Let p be the initial left endpoint of its pivot edge e , and fix g_e to be the start of the gap before any flip moves have been made beneath e . Let g_a be the start of the current gap beneath a . Therefore, the initial position of e is (p, g_a) , S_A is given by $[p, g_e]$, and S_B is given by $[g_e + 1, g_a]$.




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.

First assume that neither of S_A or S_B is empty. The simulation of the shifting is given by the following.

- (M₅) Let M_5 move e from (p, g_a) to $(p, g_e + 1)$. Since the left endpoint of e remains connected to the subtree $S_B = [g_e + 1, g_a]$, the result is still a tree, and since there are no edges other than e going from S_A to S_B , the result is still non-crossing, and thus M_5 is a valid flip move.
- (R) Now the position of S_A and the edge e meet the prerequisite for a right shifting. We perform a right shifting.
- (M₆) Let M_6 move e from $(p, g_e + 1)$ to $(g_e + 1, g_a + 1)$. The resulting graph is still a tree, because one endpoint of e that joins S_A and S_B remains unchanged, and the other endpoint is still outside the $S_A S_B$ part after the move. Moreover, since the overarching edge of e is α , there is no edge connecting any point in S_B with another point to the left of p . And g_a is the gap beneath the edge α , meaning there is no edge joining any point in S_B with another point to the right of $g_a + 1$. This means the resulting tree is still non-crossing, and hence M_6 is a valid flip move.
- (R) Now the position of S_B and the edge e meet the prerequisite for a right shifting. We perform a right shifting.
- (M₇) Let M_7 move e from $(g_e + 1, g_a + 1)$ to $(p + 1, g_a + 1)$. A similar argument as M_5 also gives that M_7 is a valid flip move.

We tweak the above procedure a bit when at least one of S_A and S_B is empty:

- (1) When both subtrees are empty, we only use M_6 . Note that since S_A is empty, we meet the assumption for M_6 , and since $p = g_e = g_a - 1$, e is in the correct initial position for M_6 and S has shifted right. This is our base case.

- (2) When S_A is empty but S_B is not, we use moves M_5 and M_6 , and then recursively construct a shift sequence under e . Since S_A is empty, we satisfy the assumption for M_6 . Note $p = g_e$, so M_6 puts e in a spot shifted right from its initial position, and by the induction S_A is shifted right as well, so S has shifted right.
- (3) When S_A is not empty but S_B is, we recursively construct a shift sequence under e , and then use M_6 and M_7 . Since S_B is empty and we begin by shifting A , we satisfy the assumptions for M_6 and M_7 , and since $g_e + 1 = g_a$, e is in the correct initial position for M_6 . Note M_7 puts e in a spot shifted right from its initial position, and by the induction, we have shifted S_A right, so S has shifted right.

We concatenate the sequences for each S_i to generate a sequence to shift T_A right. Note that each edge except the pivot edge of T is moved one to three times, so the total length of the shift sequence is at most $3m$. \square

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 have already described how to simulate an adjacent move via flip moves and shifting in Theorem 15, and how to simulate shifting via flip moves in Theorem 16. We summarise the whole construction in Algorithm 1, where we expand out the Type-1 move and omit the details of the Type-2 move to avoid redundancy. (The flip moves are already described in Theorem 15 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 $P_{FM}(Z_i, Z_{i+1}) > 0$ for any $i \leq \ell - 1$. Its validity, namely that each transition is valid, is implied by Theorem 15 and Theorem 16. 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 17. *Let \mathcal{E}_{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}_{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 $f : \mathcal{B} \rightarrow \mathbb{R}$,*²

$$\mathcal{E}_{FM}(f, f) \geq \Omega(n^{-4}) \cdot \mathcal{E}_{AM}(f, f).$$

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 that we can uniquely recover the (I, F) pair by looking at 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 18 (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 $P_{AM}(I, F) > 0$, 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 $d \in [n]$.

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} . Theorem 18 immediately implies the following.

²With a slight notation abuse, we assume $f(X) = f(T_X)$ due to the one-to-one correspondence.

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 Theorem 15;
2 if It is a Type-1 move then
3   | return Transition1( $I, F$ );
4 else
5   | return Transition2( $I, F$ );

6 procedure Transition1( $I, F$ ):
7   | // It is required that  $I \rightarrow F$  is a Type-1 move.
8   | Let  $Z_0 \leftarrow I$ ;
9   | Find  $a = (a_1, a_2)$ ,  $e = (p, g_e + 1)$  and  $g_a$  as in Theorem 15;
10  | if  $p \leq g_e - 1$  then
11  |   | Let  $(\ell - 2, Z_1, \dots, Z_{\ell-2}) \leftarrow \text{ShiftRight}(p, g_e + 1, 1, 1)$ ;
12  |   | Perform an  $M_1$  move of  $e$  to obtain  $Z_{\ell-1}$ ;
13  |   | Let  $Z_\ell \leftarrow F$ , which is obtained by an  $M_2$  move of  $e$  from  $Z_{\ell-1}$ ;
14  |   | return  $Z_0, \dots, Z_\ell$ ;

14 procedure ShiftRight( $s, t, \text{depth}, \text{step}$ ):
15  | Let  $k \geq 1$  be the number of minimal segments  $S_1, \dots, S_k$  beneath the edge  $(s, t)$ ;
16  |  $\text{last} \leftarrow \text{step}$ ;
17  | for  $i$  from  $k$  downto 1 do
18  |   | Let  $e = (p, q)$  be the pivot edge of  $S_i$ , and  $g_e, g_{e+1}$  be the gap beneath  $e$ ; // (#0)
19  |   | if  $q \geq g_e + 2$  then
20  |   |   | Perform an  $M_5$  move of  $e$  to obtain  $Z_{\text{last}}$ , and update  $\text{last} \leftarrow \text{last} + 1$ ; // (#1)
21  |   |   | if  $p \leq g_e - 1$  then
22  |   |   |   |  $(\ell', Z_{\text{last}}, \dots, Z_{\text{last}+\ell'-1}) \leftarrow \text{ShiftRight}(p, g_e + 1, \text{depth} + 1, \text{last})$ ;
23  |   |   |   | Update  $\text{last} \leftarrow \text{last} + \ell'$ ; // (#2)
24  |   |   |   | Perform an  $M_6$  move of  $e$  to obtain  $Z_{\text{last}}$ , and update  $\text{last} \leftarrow \text{last} + 1$ ; // (#3)
25  |   |   |   | if  $q \geq g_e + 2$  then
26  |   |   |   |   |  $(\ell'', Z_{\text{last}}, \dots, Z_{\text{last}+\ell''-1}) \leftarrow \text{ShiftRight}(g_e + 1, q + 1, \text{depth} + 1, \text{last})$ ;
27  |   |   |   |   | Update  $\text{last} \leftarrow \text{last} + \ell''$ ; // (#4)
28  |   |   |   |   | if  $p \leq g_e - 1$  then
29  |   |   |   |   |   | Perform an  $M_7$  move of  $e$  to obtain  $Z_{\text{last}}$ , and update  $\text{last} \leftarrow \text{last} + 1$ ; // (#5)
30  |   | return  $(\text{last} - \text{step}, Z_{\text{step}}, \dots, Z_{\text{last}})$ ;

31 procedure Transition2( $I, F$ ):
32  | // It is required that  $I \rightarrow F$  is a Type-2 move.
33  | // Omitted
```

Corollary 19. For each (Z, Z') with $P_{\text{FM}}(Z, Z') > 0$, there are at most $12n$ pairs (I, F) such that $(Z, Z') \in \Gamma_{\text{IF}}$.

We first conclude the proof of Theorem 17 by using the above encoding.

Proof of Theorem 17. Since P_{AM} and P_{FM} are reversible, Theorem 3 applies. The key quantity, namely the congestion, is bounded as follows:

$$\begin{aligned} B &= \max_{\substack{(Z, Z'):\\ P_{\text{FM}}(Z, Z') > 0}} \sum_{\substack{(I, F): \\ (Z, Z') \in \Gamma_{\text{IF}}}} \frac{\ell_{\text{IF}} \cdot P_{\text{AM}}(I, F)}{P_{\text{FM}}(Z, Z')} \leq \max_{\substack{(Z, Z'):\\ P_{\text{FM}}(Z, Z') > 0}} \sum_{\substack{(I, F): \\ (Z, Z') \in \Gamma_{\text{IF}}}} \frac{3n + 2}{(6n - 2)P_{\text{FM}}(Z, Z')} \\ &\leq \frac{5}{4} \max_{\substack{(Z, Z'):\\ P_{\text{FM}}(Z, Z') > 0}} \frac{1}{P_{\text{FM}}(Z, Z')} |\{(I, F) : (Z, Z') \in \Gamma_{\text{IF}}\}| = O(n^4), \end{aligned}$$

where the last bound is due to (11) and Theorem 19. \square

Proof of Theorem 18. Recall that the edge flips defined in Theorem 15, which we shall refer to as *the top level*, invoke once the subroutine of constructing the shift sequence if the part to shift is non-empty. The shift sequence construction subroutine ShiftLeft or ShiftRight (Theorem 16) might make further recursive calls. This inspires us to define the mapping

$$\varphi_{Z \rightarrow Z'}(\mathbb{T}_I, \mathbb{T}_F) = (\text{dir}, M, d),$$

where

- $\text{dir} \in \{\text{Left}, \text{Right}\}$ indicates the direction of adjacent move from I to F ,
- $M \in \{M_1, M_2, M_3, M_4, S_1, S_2\}$ indicates where the current transition $Z \rightarrow Z'$ is constructed as in Theorem 15, and
- $d \in [n]$ indicates the depth of the recursion of the shift sequence construction. Formally, it is the parameter depth of the recursive call where Z' is obtained in ShiftRight (or ShiftLeft) in Algorithm 1. It is defined only when $M \in \{S_1, S_2\}$, and otherwise arbitrary.

The mapping defined above is total because the level of recursion cannot go beyond the number of edges n . It is then left for us to show that it is injective.

Without loss of generality, we only consider the case $\text{dir} = \text{Left}$; otherwise, we can just swap Z and Z' to obtain the unique reconstruction of I and F using the argument for the case $\text{dir} = \text{Left}$, and then swap I and F .

If $M \in \{M_1, M_2, M_3, M_4\}$, the initial state I and the final state F can be very easily determined. Consider $M = M_1$; the other cases have a similar argument. This case corresponds to the **Type 1** move. We can find the edge being moved by looking at $Z \rightarrow Z'$, which is the edge e as in the statement of M_1 move in Theorem 15. We can also determine the edge α therein because it is the overarching edge of e . The initial I is determined by shifting the subtree beneath e in Z back to the left, and the final F is determined by performing M_2 from Z' (as we already know what the two edges are).

Otherwise, $M = S_1$, which is a right shifting in the **Type 1** move, or $M = S_2$, which is a left shifting in the **Type 2** move. We prove the former case, as the argument for the latter one is identical.

Let e_d be the edge being moved when the move $Z \rightarrow Z'$ is performed. From $i = d - 1$ to $i = 0$, we can determine the overarching edge of e_{i+1} from Z (or Z'), denoted by e_i , whose current two endpoints in Z (or Z') are (s_i, t_i) . Finally, find the overarching edge α of e_0 . The recursive calls that involve the current edge e_d are therefore ShiftRight($s_i, t_i, i + 1, \text{step}_i$) for $i = 0, 1, \dots, d - 1$. We recover the initial on each level from bottom to top to obtain $Z_{\text{step}_{d-1}-1}, Z_{\text{step}_{d-2}-1}, \dots, Z_{\text{step}_0-1} = Z_1$ as follows.

At recursion level d , the transition $Z \rightarrow Z'$ must be an edge move M_5, M_6 or M_7 at this level. This can be determined by the following argument. (See Figure 6 as a reference.)

- If e_d in Z starts left before the gap beneath e_{d-1} , and if the left end of e_d remains unchanged in Z' , then it is an M_5 move with a non-empty S_B part.
- Otherwise, if e_d in Z' ends right after the gap beneath e_{d-1} , and if the right end of e_d is the same in Z and Z' , then it is an M_7 move with a non-empty S_A part.
- Otherwise, it is an M_6 move.

Knowing which move it is allows us to determine the S_A and S_B part of the current minimal segment S_i that is being processed. To recover the initial state beneath e_{d-1} , we first recover the current minimal segment S_i (obtaining Figure 6(#0)), and then move everything beneath e_{d-1} after the gap beneath e_{d-1} to the left by one position. In this way, we revert the subroutine $\text{ShiftRight}(s_{d-1}, t_{d-1}, d, \text{step}_{d-1})$ and obtain $Z_{\text{step}_{d-1}-1}$.

Now suppose that we are at recursion depth $d' < d$. The current transition $Z \rightarrow Z'$ must be inside a recursive call on this level, which means Z and Z' are the same except the part with coordinates from $s_{d'+1}$ to $t_{d'+1}$ (namely the part beneath $e_{d'}$). There are two possible recursive calls in ShiftRight . To determine which kind of recursive call it is (either the first one, between (#1) and (#2) in Algorithm 1, or the second one, between (#3) and (#4)), we look at the gap beneath $e_{d'-1}$. If it is on the right of $e_{d'}$, then the recursive call is the first one, or otherwise it is the second one. In any case, we first recover the part beneath $e_{d'}$ and obtain $Z_{\text{step}_{d'}-1}$. At the current level, this ends up in either the state (#1) or (#3), which further allows us to go back to state (#0) by recovering the initial state of the current minimal segment, and then move everything beneath $e_{d'-1}$ after the gap beneath $e_{d'-1}$ to the left by one position. In this way, we revert the subroutine $\text{ShiftRight}(s_{d'-1}, t_{d'-1}, d', \text{step}_{d'-1})$ and obtain $Z_{\text{step}_{d'-1}-1}$.

The above induction provides us a way to reconstruct the initial state beneath e_0 , which gives Z_0 .

To recover the final state T_F , one only needs to modify the above procedure a bit, by moving every minimal segments to the right in each call of ShiftRight , and perform the M_1 and M_2 move on e_0 and a on the top level.

The encoding is unique because, if there are two pairs (I, F) , (I', F') such that they receive the same encoding $\varphi_{Z \rightarrow Z'}(T_I, T_F)$, the above procedure correctly recovers the initial and final states. Because the procedure is deterministic and only depends on $Z \rightarrow Z'$ and the encoding, it generates the same output as containing the initial and final states, which means $(I, F) = (I', F')$. \square

We end this section by concluding the proof of the main theorem.

Proof of Theorem 1. We combine our comparison of the adjacent move and flip chains to get

$$\begin{aligned}
t_{\text{mix}}(\mathbb{P}_{\text{FM}}) &\leq \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{1}{\lambda(\mathbb{P}_{\text{FM}})} \\
\text{(by Equation (4))} \quad &= \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{1}{\inf \left\{ \frac{\mathcal{E}_{\text{FM}}(f, f)}{\text{Var}_{\pi}(f)} \mid f : \Omega \rightarrow \mathbb{R}, \text{Var}_{\pi}(f) \neq 0 \right\}} \\
\text{(by Theorem 17)} \quad &\leq \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{\mathcal{O}(n^4)}{\inf \left\{ \frac{\mathcal{E}_{\text{AM}}(f, f)}{\text{Var}_{\pi}(f)} \mid f : \Omega \rightarrow \mathbb{R}, \text{Var}_{\pi}(f) \neq 0 \right\}} \\
\text{(by Equation (4))} \quad &= \left(1 + \frac{1}{2} \log \pi_{\min}\right) \frac{\mathcal{O}(n^4)}{\lambda(\mathbb{P}_{\text{AM}})} \\
\text{(by Theorem 13)} \quad &= \left(1 + \frac{1}{2} \log \pi_{\min}\right) \mathcal{O}(n^7 \log n) \\
&= \mathcal{O}(n^8 \log n). \quad \square
\end{aligned}$$

REFERENCES

- [ABB⁺22] 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. 1
- [ABCR15] 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. 2
- [ABRS15] Victor Alvarez, Karl Bringmann, Saurabh Ray, and Raimund Seidel. Counting triangulations and other crossing-free structures approximately. *Comput. Geom.*, 48(5):386–397, 2015. 2
- [AF96] David Avis and Komei Fukuda. Reverse search for enumeration. *Discret. Appl. Math.*, 65(1-3):21–46, 1996. 1, 2, 5
- [ALO20] 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. 5

- [ALOV19] 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. [3](#), [4](#), [5](#)
- [Ard03] Federico Ardila. The Catalan matroid. *J. Comb. Theory, Ser. A*, 104(1):49–62, 2003. [3](#)
- [BCM⁺12] 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. [2](#)
- [BDG⁺23] 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. [2](#)
- [BdMPW24] 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. [1](#)
- [BGNP23] 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. [1](#)
- [BKUV25] Håvard Bakke Bjerkevik, Linda Kleist, Torsten Ueckerdt, and Birgit Vogtenhuber. Flipping non-crossing spanning trees. In *SODA*, pages 2313–2325. SIAM, 2025. [1](#), [2](#), [5](#)
- [Bud17] 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. [2](#)
- [CGM21] Mary Cryan, Heng Guo, and Giorgos Mousa. Modified log-Sobolev inequalities for strongly log-concave distributions. *Ann. Probab.*, 49(1):506–525, 2021. [3](#), [4](#), [5](#)
- [CMSS15] Pietro Caputo, Fabio Martinelli, Alistair Sinclair, and Alexandre Stauffer. Random lattice triangulations: structure and algorithms. *Ann. Appl. Probab.*, 25(3):1650–1685, 2015. [2](#)
- [Coh16] Emma Cohen. *Problems in Catalan mixing and matchings in regular hypergraphs*. PhD thesis, Georgia Institute of Technology, 2016. [2](#)
- [CP24] Swee Hong Chan and Igor Pak. Computational complexity of counting coincidences. *Theor. Comput. Sci.*, 1015:114776, 2024. [2](#)
- [CS20] Alessandra Caraceni and Alexandre Stauffer. Polynomial mixing time of edge flips on quadrangulations. *Probab. Theory Related Fields*, 176(1-2):35–76, 2020. [2](#)
- [Dai02] 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. [2](#)
- [DFK91] 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. [2](#)
- [DMRT20] Mehmet Demirtas, Liam McAllister, and Andres Rios-Tascon. Bounding the Kreuzer-Skarke landscape. *Fortschr. Phys.*, 68(11-12):2000086, 13, 2020. [2](#)
- [DRS10] 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. [2](#)
- [DS93] Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3):696–730, 1993. [3](#), [4](#), [6](#)
- [EF23] 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. [2](#)
- [Epp20] David Eppstein. Counting polygon triangulations is hard. *Discret. Comput. Geom.*, 64(4):1210–1234, 2020. [2](#)
- [GKZ08] 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. [2](#)
- [HHM⁺99] 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. [1](#), [2](#)
- [JKP86] J. Jurkiewicz, A. Krzywicki, and B. Petersson. A numerical study of discrete Euclidean Polyakov surfaces. *Physics Letters B*, 168(3):273–278, 1986. [2](#)
- [JMPV23] 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. [5](#)
- [JS89] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989. [2](#), [6](#)
- [JS93] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993. [2](#)
- [KKM85] 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. [2](#)
- [KLS18] Marek Karpinski, Andrzej Lingas, and Dzmityr 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. [2](#)
- [KZ03] 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. [2](#)
- [LP17] David A. Levin and Yuval Peres. *Markov Chains and Mixing Times*. American Mathematical Society, 2017. [5](#), [6](#)
- [Mov18] Ramis Movassagh. The gap of Fredkin quantum spin chain is polynomially small. *Ann. Math. Sci. Appl.*, 3(2):531–562, 2018. [2](#)
- [MR00] 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. [2](#)

- [MRS97] 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. [2](#)
- [MS16] 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. [2](#)
- [MT97] 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. [2](#), [4](#)
- [Noy98] Marc Noy. Enumeration of noncrossing trees on a circle. *Discret. Math.*, 180(1-3):301–313, 1998. [2](#)
- [SS11] Micha Sharir and Adam Sheffer. Counting triangulations of planar point sets. *Electron. J. Comb.*, 18(1), 2011. [2](#)
- [SSW11] 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. [2](#)
- [Sta99] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. [3](#), [7](#)
- [Sta15] Richard P. Stanley. *Catalan numbers*. Cambridge University Press, New York, 2015. [2](#), [6](#), [7](#)
- [Vir89] O. Ya. Viro. Real plane algebraic curves: constructions with controlled topology. *Algebra i Analiz*, 1(5):1–73, 1989. [2](#)
- [Wil04] David Bruce Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Ann. Appl. Probab.*, 14(1):274–325, 2004. [2](#), [3](#), [4](#), [9](#)