

Mixing times of Markov chains on 3-Orientations of Planar Triangulations

Sarah Miracle, Dana Randall, Amanda Pascoe Streib, Prasad Tetali

► **To cite this version:**

Sarah Miracle, Dana Randall, Amanda Pascoe Streib, Prasad Tetali. Mixing times of Markov chains on 3-Orientations of Planar Triangulations. Broutin, Nicolas and Devroye, Luc. 23rd International Meeting on Probabilistic, Combinatorial, and Asymptotic Methods in the Analysis of Algorithms (AofA'12), 2012, Montreal, Canada. Discrete Mathematics and Theoretical Computer Science, DMTCS Proceedings vol. AQ, 23rd Intern. Meeting on Probabilistic, Combinatorial, and Asymptotic Methods for the Analysis of Algorithms (AofA'12), pp.413-424, 2012, DMTCS Proceedings. <hal-01197229>

HAL Id: hal-01197229

<https://hal.inria.fr/hal-01197229>

Submitted on 11 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Mixing times of Markov chains on 3-Orientations of Planar Triangulations

Sarah Miracle^{1†}, Dana Randall^{1‡}, Amanda Pascoe Streib^{2§} and Prasad Tetali^{1,2¶}

¹ College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0765

² School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332-0765

Given a planar triangulation, a 3-orientation is an orientation of the internal edges so all internal vertices have out-degree three. Each 3-orientation gives rise to a unique edge coloring known as a *Schnyder wood* that has proven useful for various computing and combinatorics applications. We consider natural Markov chains for sampling uniformly from the set of 3-orientations. First, we study a “triangle-reversing” chain on the space of 3-orientations of a fixed triangulation that reverses the orientation of the edges around a triangle in each move. We show that (i) when restricted to planar triangulations of maximum degree six, the Markov chain is rapidly mixing, and (ii) there exists a triangulation with high degree on which this Markov chain mixes slowly. Next, we consider an “edge-flipping” chain on the larger state space consisting of 3-orientations of all planar triangulations on a fixed number of vertices. It was also shown previously that this chain connects the state space and we prove that the chain is always rapidly mixing.

Keywords: Markov chains, 3-orientations, Planar triangulations, Schnyder woods

1 Introduction

The 3-orientations of a graph have given rise to beautiful combinatorics and computational applications. A 3-orientation of a planar triangulation is an orientation of the internal edges of the triangulation such that every internal vertex has out-degree three. We study natural Markov chains for sampling 3-orientations in two contexts, when the triangulation is fixed and when we consider the union of all planar triangulations on n vertices. When the triangulation is fixed, we allow moves that reverse the orientation of edges around a triangle if they form a directed cycle. We show that the chain is rapidly mixing (converging in polynomial time to equilibrium) if the maximum degree of the triangulation is six, but can be slowly mixing (requiring exponential time) if the degrees are unbounded. To sample from the set of all 3-orientations of triangulations with n vertices we use a simple “edge-flipping” chain and show it is always rapidly mixing.

[†]Supported in part by a DOE Office of Science Graduate Fellowship, NSF CCF-0830367 and an ARCS Scholar Award.

[‡]Supported in part by NSF CCF-0830367 and CCF-0910584.

[§]Supported in part by a NPSC Fellowship and NSF CCF-0910584.

[¶]Supported in part by NSF DMS-1101447 and CCR-0910584.

These chains arise in contexts such as sampling Eulerian orientations and triangulations of fixed planar point sets, so there is additional motivation for understanding their convergence rates.

More precisely, given an undirected graph $G = (V, E)$ and a function $\alpha : V \rightarrow \mathbb{Z}^+$, an α -orientation is an orientation of E where each vertex v has outdegree $\alpha(v)$. Several fundamental combinatorial structures – spanning trees, bipartite perfect matchings, Eulerian orientations, etc. – can be seen as special instances of α -orientations of planar graphs. We refer the reader to [11, 12, 14] for extensive literature on the subject. Not surprisingly, counting α -orientations is $\#P$ -complete. Namely, consider an undirected Eulerian graph G (with all even degrees); the α -orientations of G , where $\alpha(v) = d(v)/2$, correspond precisely to Eulerian orientations of G . The latter problem has been shown to be $\#P$ -complete by Mihail and Winkler [20], and more recently Creed [7] showed that it remains $\#P$ -complete even when restricted to planar graphs.

The term *3-orientation* refers to an α -orientation of a planar triangulation where all *internal vertices* (vertices not bounding the infinite face) have $\alpha(v) = 3$ and all *external vertices* (the three vertices bounding the infinite face) have $\alpha(v) = 0$. Each 3-orientation gives rise to a unique edge coloring, known as a *Schnyder wood*, whose many combinatorial applications include graph drawing [22, 6] and poset dimension theory [23]. Several intriguing enumeration problems remain open, such as the complexity of enumerating 3-orientations of a planar triangulation (see e.g., [14].) We study the problem of sampling 3-orientations of a fixed (planar) triangulation and sampling 3-orientations of all triangulations with n internal vertices. In particular, we analyze the mixing times of two natural Markov chains, which were introduced previously but had thus far resisted analysis.

First, we study the problem of sampling 3-orientations of a fixed triangulation, which was stated as an open problem by Felsner and Zickfeld [14]. Although there is no known efficient method for counting exactly, there are polynomial-time algorithms for approximately counting and sampling 3-orientations due to a bijection with perfect matchings of a particular bipartite graph (see Section 6.2 in [14]). This bijection allows us to sample 3-orientations in time $O^*(n^7)$ using an algorithm due to Bezáková et al. [1] (improving on the results of Jerrum, Sinclair and Vigoda [17]), but this approach is indirect and intricate.

We consider instead a natural “triangle-reversing” Markov chain, \mathcal{M}_{TR} , that reverses the orientation of a directed triangle in each step, thus maintaining the outdegrees. Brehm [5] showed that for any fixed triangulation T , \mathcal{M}_{TR} connects the state space $\Psi(T)$. We also consider a related “cycle-reversing” chain, \mathcal{M}_{CR} , that can reverse directed cycles containing more than one triangle. The chain \mathcal{M}_{CR} is a non-local version of \mathcal{M}_{TR} based on “tower moves” reminiscent of those in [18]. We show that both of these chains are rapidly mixing. Let $\Delta_I(T)$ denote the maximum degree of any internal vertex of T . We show:

Theorem 1 *If T is a planar triangulation with $\Delta_I(T) \leq 6$, then the mixing time of \mathcal{M}_{CR} on the state space $\Psi(T)$ satisfies $\tau(\mathcal{M}_{CR}) = O(n^5)$.*

We use a standard comparison argument together with Theorem 1 to infer a bound on the mixing time of the triangle-reversing chain \mathcal{M}_{TR} . Thus we prove:

Theorem 2 *If T is a planar triangulation with $\Delta_I(T) \leq 6$, then the mixing time of \mathcal{M}_{TR} on the state space $\Psi(T)$ satisfies $\tau(\mathcal{M}_{TR}) = O(n^8)$.*

Note that the class of planar triangulations with $\Delta_I \leq 6$ is exponentially large in n , the number of internal vertices. An interesting case is when the fixed triangulation is a finite region Λ of the triangular lattice, since sampling 3-orientations on Λ corresponds to sampling Eulerian orientations. Creed [7] independently solved the sampling problem in this special case using a similar approach based on towers; he shows that for certain subsets of the triangular lattice the tower chain can be shown to mix in time

$O(n^4)$. In addition, it was previously shown that similar cycle-reversing chains are rapidly mixing in the context of sampling Eulerian orientations on the Cartesian lattice [18] and the 8-vertex model [10]. Our analysis here bounding the mixing rate of \mathcal{M}_{CR} in the general setting of arbitrary planar graphs with maximum degree 6 requires additional combinatorial insights because we no longer have the regular lattice structure. In particular, we make use of a combinatorial structure outlined by Brehm [5]. In fact, this structure allows us to extend our analysis to certain non-4-connected triangulations that have vertices of degree greater than six. Next, we prove that when Δ_I is unbounded, \mathcal{M}_{TR} may require exponential time. Specifically, we prove:

Theorem 3 *For any (large) n , there exists a triangulation T of size n for which \mathcal{M}_{TR} on the state space $\Psi(T)$ has mixing time $\tau(\mathcal{M}_{TR}) = \Omega(2^{n/4})$.*

Based on the construction we give here, Felsner and Heldt [13] recently constructed another, somewhat simpler, family of graphs for which the mixing rate of \mathcal{M}_{TR} and \mathcal{M}_{CR} is exponentially large. However, we note that their family also has maximum degree that grows with n .

The second problem we study is sampling from the set of all 3-orientations arising from all possible triangulations on n internal vertices. Let Ψ_n be the set of all triangulations of a labelled fixed point set with $n + 3$ vertices, three of which are external vertices, where the edges of the triangulation are not required to be straight and the fixed positions of the points are arbitrary (i.e. all fixed positions result in the same set Ψ_n). The set Ψ_n is known to be in 1-1 correspondence with all pairs of non-crossing Dyck paths, and as such has size $C_{n+2}C_n - C_{n+1}^2$, where C_n is the n th Catalan number. Since exact enumeration is possible, we can sample using the reduction to counting; this was explicitly worked out by Bonichon and Mosbah [4]. We consider a natural Markov chain approach for sampling that in each step selects a quadrangle at random, removes the interior edge, and replaces it with the other diagonal in such a way as to restore the out degree at each vertex. Bonichon, Le Saëc and Mosbah [3] showed that the chain \mathcal{M}_{EF} connects Ψ_n and we present the first bounds showing that the chain is rapidly mixing. Although the exact counting approach already yields a fast approach to sampling, the chain \mathcal{M}_{EF} is compelling because it arises in contexts where we do not have methods to count exactly. For example, it has been proposed as a method for sampling triangulations of a fixed planar point set, a problem that has been open for over twenty years. In addition, there is interest in the mixing rate of this chain precisely because the number is related to the Catalan numbers; there has been extensive work trying to bound mixing rates of natural Markov chains for various families of Catalan structures (see, e.g., [19]). Specifically, we prove:

Theorem 4 *The mixing time of \mathcal{M}_{EF} on the state space Ψ_n satisfies $\tau(\mathcal{M}_{EF}) = O(n^{10} \log n)$.*

2 Preliminaries

We begin with background on 3-orientations, Schnyder woods, and Markov chains. Fraysseix and Ossona de Mendez defined a bijection between $\Psi(T)$ and the Schnyder woods of T [15]. A *Schnyder wood* (see Figure 4) is a 3-coloring and orientation of the edges of T such that for every internal vertex v , v has out-degree exactly 1 in each of the 3 colors, and the clockwise order of the edges incident to v is: outgoing green, incoming blue, outgoing red, incoming green, outgoing blue and incoming red (see Figure 1).

In our figures, we differentiate the colors of edges in the Schnyder woods by dashed lines (green), dotted lines (red), and solid lines (blue). Notice the orientation of the edges of the Schnyder wood is a 3-orientation and that each of the colors forms a directed tree which spans the internal vertices and is rooted




Fig. 1: (a) The vertex condition. (b) A red/green swap.

at one of the external vertices. Throughout the proofs, when we refer to the colors of the edges of a 3-orientation, we mean the colors of the Schnyder wood associated with that 3-orientation.

Next, we present some background on Markov chains. The time a Markov chain takes to converge to its stationary distribution π is measured in terms of the distance between π and \mathcal{P}^t , the distribution at time t . The *total variation distance* at time t is $\|\mathcal{P}^t, \pi\|_{tv} = \max_{x \in \Psi} \frac{1}{2} \sum_{y \in \Psi} |\mathcal{P}^t(x, y) - \pi(y)|$, where $\mathcal{P}^t(x, y)$ is the t -step transition probability and Ψ is the state space. For all $\epsilon > 0$, the *mixing time* τ of \mathcal{M} is defined as $\tau = \min\{t : \|\mathcal{P}^t, \pi\|_{tv} \leq 1/4, \forall t' \geq t\}$. We say that a Markov chain is *rapidly mixing* if the mixing time is bounded above by a polynomial in n . In this case, n is the number of internal vertices.

3 Sampling 3-orientations of a fixed triangulation

In this section, we consider a Markov chain for sampling the 3-orientations of a given triangulation. Let T be a planar triangulation with n internal vertices. Consider the following natural local Markov chain \mathcal{M}_{TR} on the set of all 3-orientations of T . Select a directed 3-cycle at random and reverse its orientation. We will see that \mathcal{M}_{TR} samples from the uniform distribution, but its efficiency will depend on T . In Section 3.1 we show that if the maximum degree of any internal vertex is at most 6, \mathcal{M}_{TR} is rapidly mixing. In contrast, in Section 3.2 we demonstrate a triangulation T but with unbounded degree for which \mathcal{M}_{TR} takes exponential time to sample from the state space $\Psi(T)$. Define \mathcal{M}_{TR} as follows.

The Markov chain \mathcal{M}_{TR}

Starting at any $\sigma_0 \in \Psi(T)$, iterate the following:

- Choose a triangle t in σ_i u.a.r.
- If t is a directed cycle, then with prob. 1/2 reverse t to obtain σ_{i+1} .
- Otherwise, $\sigma_{i+1} = \sigma_i$.

Brehm proved that \mathcal{M}_{TR} connects the state space, $\Psi(T)$ [5]. Since all valid moves have the same transition probabilities, \mathcal{M}_{TR} converges to the uniform distribution over $\Psi(T)$.

3.1 Fast mixing of \mathcal{M}_{TR} for maximum degree at most 6

In this section we prove that \mathcal{M}_{TR} is rapidly mixing on the state space $\Psi(T)$, if T is a planar triangulation with $\Delta_I(T) \leq 6$. First, we introduce an auxiliary chain \mathcal{M}_{CR} , which we will then use to derive a bound on the mixing time of \mathcal{M}_{TR} . The Markov chain \mathcal{M}_{CR} involves *towers* of moves of \mathcal{M}_{TR} , based on the nonlocal chain introduced in [18]. Notice that if a face f cannot move then two of its edges have the same orientation and the other edge does not (see, e.g., the face f_1 in Figure 2(a)). We call this edge the *disagreeing edge* of f . Define a tower of length k to be a path of faces f_1, f_2, \dots, f_k such that the following three conditions are met: f_k is the only face which is bounded by a directed cycle (i.e. it has a move); for every $1 \leq i < k$, the disagreeing edge of f_i is also incident to f_{i+1} ; and every vertex v is

incident to at most three consecutive faces in the path (see Figure 2). The idea of the tower is that once the edges of f_k are reversed, then the edges of f_{k-1} can be reversed, and so on. We call f_1 the beginning of the tower, and f_k the end. Notice that every face is the beginning of at most one tower (it may be a tower of length 1). The effect of making this sequence of moves is to reverse the edges of the directed cycle surrounding the tower (although the colors on the internal edges also change).




Fig. 2: A tower of length 6.

The Markov chain \mathcal{M}_{CR} operates as follows. Starting at any σ_0 , iterate the following: Choose a (finite) face f u.a.r.; if f is the beginning of a tower of length k , then with probability $\begin{cases} \frac{1}{6k} & : k \geq 2 \\ \frac{1}{2} & : k = 1 \end{cases}$ reverse this tower to obtain σ_{i+1} ; else, $\sigma_{i+1} = \sigma_i$. The moves of \mathcal{M}_{TR} are a subset of the moves of \mathcal{M}_{CR} , so \mathcal{M}_{CR} is connected as well. We first consider the case that T is 4-connected, then we apply the comparison theorem to prove that \mathcal{M}_{TR} is also rapidly mixing and extend the result to non-4-connected triangulations using a result of Brehm [5]. Notice that if T is 4-connected, every 3-cycle is facial, so \mathcal{M}_{TR} selects a face and rotates the edges around that face if possible. The bulk of the work to prove Theorems 1 and 2 is to show that \mathcal{M}_{CR} is rapidly mixing when T is 4-connected. The main tool we use in the case where T is 4-connected is a path coupling theorem due to Dyer and Greenhill [9].

Theorem 5 *Let T be a 4-connected planar triangulation with $\Delta_I(T) \leq 6$. Then the mixing time of \mathcal{M}_{CR} on the state space $\Psi(T)$ satisfies $\tau(\mathcal{M}_{CR}) = O(n^5)$.*

Proof. Let T be a 4-connected planar triangulation such that $\Delta_I(T) \leq 6$. First we prove that \mathcal{M}_{CR} is rapidly mixing on $\Psi(T)$. Define the distance d between any two 3-orientations in $\Psi(T)$ to be the minimum number of steps of \mathcal{M}_{TR} from one to the other. Assume $\sigma, \tau \in \Psi(T)$ and τ is obtained from σ by reversing a facial triangle f . We use the trivial coupling that chooses the same face for σ and τ at every step. Suppose without loss of generality that the edges of f are clockwise in σ . There are two obvious moves that decrease the distance, namely when \mathcal{M}_{CR} selects the face f and chooses to direct the cycle either way, each of which happens with probability $1/(2(2n+1))$. Any move of \mathcal{M}_{CR} that does not involve an edge of f occurs with the same probability in σ and τ , and hence is neutral (i.e. does not change the distance). Finally, we call a tower *bad* if it contains a neighbor f' of f that is not the end of the tower. In this case, we say this bad tower is *associated with f'* . On the other hand, a tower is *good* if it ends in f , or if it ends in a face f' adjacent to f and contains no other faces adjacent to f . We will show that the good towers in σ have corresponding good towers in τ , while the bad towers in σ fail in τ , increasing the distance. Any tower that is neither good nor bad does not contain an edge of f .

Suppose $k \geq 1$, (f_1, f_2, \dots, f_k) is a good tower in σ , and f_k is adjacent to f . We claim that $(f_1, f_2, \dots, f_k, f)$ is a good tower in τ . It is clear that in τ , f is the only one of these faces that is bounded by a cycle, and that upon rotating f , the tower (f_1, f_2, \dots, f_k) is possible. We must check two things: that $(f_1, f_2, \dots, f_k, f)$ is a path of faces (i.e. does not contain any cycle of faces), and that every vertex is incident to at most three consecutive faces. The first condition is clear, since f_k is the only

neighbor of f in $\{f_1, f_2, \dots, f_k\}$, and (f_1, f_2, \dots, f_k) is a path of faces. Suppose the second condition does not hold. Then there is a vertex v incident to f, f_k, f_{k-1} , and f_{k-2} . The edges between faces f_{k-2} and f_{k-1} and between f_{k-1} and f_k are either both incoming to v or both outgoing from v . Moreover, since the edge between f_{k-2} and f_{k-1} is the disagreeing edge of f_{k-2} , the two edges of f_{k-2} incident to v are either both incoming to v or both outgoing from v (similarly the two edges of f_k incident to v are either both incoming to v or both outgoing). Hence there are four edges incident to v which are all incoming or all outgoing; a contradiction since a vertex of degree at most 6 with exactly three outgoing edges can have at most three incoming edges as well. Therefore if a good tower of length $k \geq 1$ begins on a face f_1 and ends on a neighbor f_k of f in σ then there is a corresponding tower of length $k+1$ that begins on f_1 and ends on f in τ . Thus we have shown that if (f_1, f_2, \dots, f_k) is a good tower in σ , then $(f_1, f_2, \dots, f_k, f)$ is a good tower in τ . On the other hand, it should be clear that if $(f_1, f_2, \dots, f_k, f)$ is a good tower of length $k+1 \geq 2$ that ends on f in σ then (f_1, f_2, \dots, f_k) is a good tower of length k in τ . In either case, if $k \geq 2$ then the expected change in distance given the choice of these towers is

$$\frac{1}{2n+1} \left(-\frac{1}{6(k+1)} + k \left(\frac{1}{6k} - \frac{1}{6(k+1)} \right) \right) = 0.$$

If $k = 2$ then the expected change in distance is $\frac{1}{2n+1} \left(-\frac{1}{12} + \left(\frac{1}{2} - \frac{1}{12} \right) \right) = 1/3(2n+1)$.

We point out that if σ and τ have good towers using a neighbor f' of f then no bad tower in σ or τ is associated with f' ; that is, if there exists a bad tower containing f' then f' is the end of the tower. Suppose without loss of generality that the good tower is longer in τ than in σ . Then the edge between f and f' is the disagreeing edge of f' in σ so the only way to tower is towards f , so f' is not in a bad tower in σ . On the other hand, in τ , f' is bounded by a cycle, so it must be the end of any tower containing it.

Moreover, σ and τ can have at most two bad towers associated with a given face f' adjacent to f . It is clear that σ (resp., τ) has at most one bad tower that begins in f' , which is defined by the disagreeing edge of f' in σ . However, σ may have a bad tower that uses f' but does not begin in f' . Let (f_1, f_2, \dots, f_k) be such a tower. We will show that $f' = f_2$. Suppose not, so that $f' = f_i$, where $i \geq 3$. Then, as above, there is a vertex v that is incident to f, f_i, f_{i-1} , and f_{i-2} , and the same proof shows that v must have either in-degree at least 4 or out-degree at least 4, which is a contradiction. Therefore bad towers associated with f' must either begin in f' or in a neighbor f_1 of f' . If there is a bad tower in σ (τ) beginning in f_1 then in both σ and τ , the edge between faces f' and f_1 is f_1 's disagreeing edge, which means that τ (resp. σ) cannot have a bad tower beginning in f' . Therefore there are at most two bad towers in σ or τ associated with f' . The expected change in distance given that a bad tower of length $k \geq 2$ is chosen is $2k/(6k(2n+1)) = 1/(3(2n+1))$. Therefore $\mathbb{E}[\Delta d] \leq \frac{1}{2n+1} \left(-2 \left(\frac{1}{2} \right) + 3 \left(\frac{1}{3} \right) \right) = 0$.

For any pair (σ_t, τ_t) in $\Psi(T)$, by a connectivity proof of Brehm [5], there exists a path of transitions of \mathcal{M}_{TR} from σ_t to τ_t of length $d(\sigma_t, \tau_t)$. The first of these transitions occurs with probability at least $1/(4n+2)$ and decreases the distance by 1. Thus, $\Pr(d(\sigma_t, \tau_t) \neq d(\sigma_{t+1}, \tau_{t+1})) \geq 1/(4n+2)$. By the path coupling theorem [9] and a bound on the distance between any two 3-orientations, we see the mixing time $\tau(\mathcal{M}_{CR})$ of \mathcal{M}_{CR} satisfies $\tau(\mathcal{M}_{CR}) \leq \lceil e(2n+1)^5/2 \rceil = O(n^5)$. \square

Finally, we use the comparison theorem, Theorem 8 introduced in Section 4.2, to derive a bound on the mixing time $\tau(\mathcal{M}_{TR})$ of \mathcal{M}_{TR} for 4-connected planar triangulations. We can then extend both these results to general planar triangulations with $\Delta_I \leq 6$ using a result of Brehm [5], proving Theorems 1 and 2. In fact, we show that \mathcal{M}_{TR} is rapidly mixing for a certain class of planar triangulations that can have vertices of degree greater than 6 (informally, planar triangulations where each 4-connected sub-triangulation has $\Delta_I \leq 6$). More details can be found in the full version of the paper.

3.2 Slow mixing of \mathcal{M}_{TR} for unbounded degree

We now exhibit a triangulation on which \mathcal{M}_{TR} takes exponential time to converge. A key tool is *conductance*, which for an ergodic Markov chain with stationary distribution π is

$$\Phi_{\mathcal{M}} = \min_{\substack{S \subseteq \Psi \\ \pi(S) \leq 1/2}} \frac{1}{\pi(S)} \sum_{s_1 \in S, s_2 \in \bar{S}} \pi(s_1) \mathcal{P}(s_1, s_2).$$

The following theorem relates the conductance and mixing time (see, e.g., [16, 24]).

Theorem 6 For any Markov chain with conductance Φ , $\tau \geq (4\Phi)^{-1} - 1/2$.

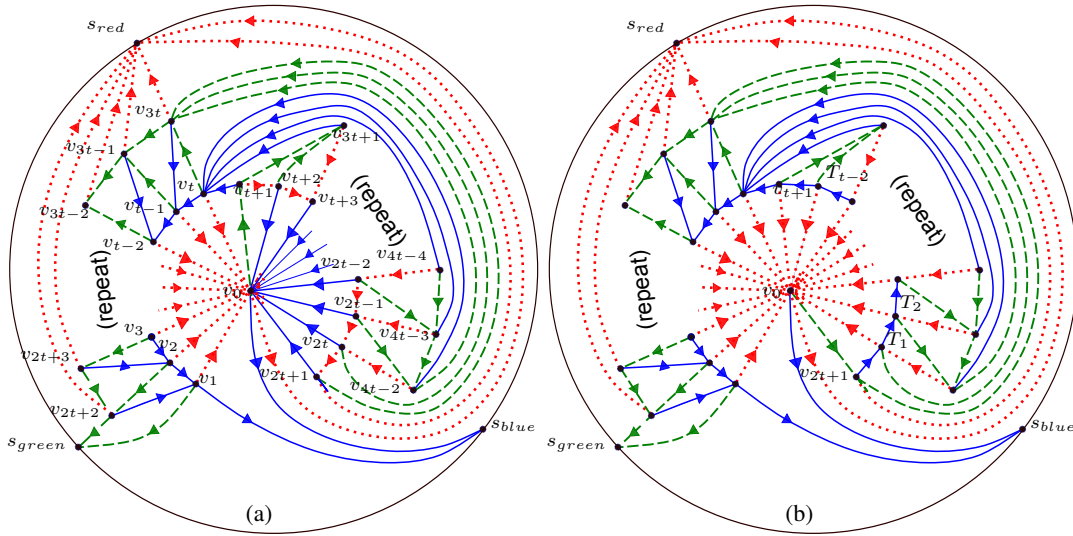


Fig. 3: (a) A triangulation for which \mathcal{M}_{TR} mixes slowly. (b) There is an exponential number of 3-orientations with edge (v_0, v_{t+1}) colored red, corresponding to different orientations of T_1, T_2, \dots, T_{t-2} .

Proof of Theorem 3. We show that for the generalized triangulation G given in Figure 3 with $n = 4t - 2$ internal vertices, \mathcal{M}_{TR} takes exponential time to converge. Specifically, we show that although there is an exponential number of 3-orientations where edge (v_0, v_{t+1}) is colored blue or red, all paths between these 3-orientations with (v_0, v_{t+1}) colored differently must include a 3-orientation where (v_0, v_{t+1}) is colored green. There is only a single 3-orientation that satisfies this property (namely, the one pictured in Figure 3a), which creates a bottleneck in the state space. Let D be the set of 3-orientations of G with (v_0, v_{t+1}) colored red or green and \bar{D} , the complement of D , be the set of 3-orientations with (v_0, v_{t+1}) colored blue. In order to show that both D and \bar{D} are exponentially large we produce a triangulation in each set which contains roughly t directed triangles which do not share any edges and reversing these triangles does not change the colors of the edges adjacent to v_0 . Hence each of the 2^t choices of the orientations of these triangles gives a distinct 3-orientation with edge (v_0, v_{t+1}) colored appropriately (see Figure 3b and triangles T_1 thru T_{t-2}). Moreover, there is only one 3-orientation with (v_0, v_{t+1})

colored green, corresponding to Figure 3a. By the Vertex Condition, if edge (v_0, v_{t+1}) is green then edges $(v_0, v_1), (v_0, v_2), \dots, (v_0, v_t)$ must all be directed toward v_0 and colored red; this is because edge (v_0, s_{blue}) is blue and directed toward s_{blue} in every 3-orientation of G . Using a similar argument, one can check that there is a unique way to extend this coloring of the edges. To go from a configuration where edge (v_0, v_{t+1}) has color red (blue) to blue (resp., red) one must go through a coloring where the edge is green. This is because the only choices for edge (v_0, v_{t+1}) are red directed toward v_0 , blue directed toward v_0 , and green directed away, and any move that changes the color must also change the direction. Using this bad cut, together with Theorem 6 we show $\tau(\mathcal{M}_{TR}) \geq 2^{(n-14)/4} - \frac{1}{2}$. \square

4 Sampling 3-orientations of triangulations on n internal vertices

We consider a local Markov chain \mathcal{M}_{EF} for sampling uniformly from Ψ_n and show \mathcal{M}_{EF} is always rapidly mixing. Our argument relies on a bijection with pairs of Dyck paths to relate the mixing time of a chain on Dyck paths to \mathcal{M}_{EF} using the comparison method [8]. Define \mathcal{M}_{EF} as follows (see Figure 1).

The Markov chain \mathcal{M}_{EF}

Starting at any $\sigma_0 \in \Psi_n$, iterate the following:

- Choose two facial triangles T_1 and T_2 with a shared edge \overrightarrow{xy} u.a.r.⁽ⁱ⁾
- Choose an edge \overrightarrow{zx} from $T_1 \cup T_2$ u.a.r., if one exists. With prob. $1/2$ replace the path $\{\overrightarrow{zx}, \overrightarrow{xy}\}$ by $\{\overrightarrow{xz}, \overrightarrow{zw}\}$ where w is the remaining vertex of $T_1 \cup T_2$.
- Otherwise, $\sigma_{i+1} = \sigma_i$.

If the edge \overrightarrow{zx} with color c_i is replaced by the edge \overrightarrow{xz} with color c_j , we call this a **c_j/c_i swap**. Bonichon, Le Saëc and Mosbah showed in [3] that \mathcal{M}_{EF} connects the state space Ψ_n . Since all valid moves have the same transition probabilities, this implies that \mathcal{M}_{EF} converges to the uniform distribution over Ψ_n .

4.1 The bijection between Ψ_n and pairs of Dyck paths

The key to bounding the mixing time of \mathcal{M}_{EF} is a bijection between Ψ_n and pairs of nonoverlapping Dyck paths of length $2n$, introduced by Bonichon [2]. Dyck paths can be thought of as strings $a_1 a_2 \dots a_{2n}$ containing an equal number of 1's and -1 's, where for any $1 \leq k \leq 2n$, $\sum_{i=1}^k a_i \geq 0$. Recall that a 3-orientation of a triangulation can be viewed as a union of three trees, one in each color. In the bijection, the bottom Dyck path corresponds to the blue tree and the top Dyck path indicates the degree of each vertex in the red tree. The green tree is determined uniquely by the blue and red trees. More specifically, given $\sigma \in \Psi_n$, to determine the bottom Dyck path, start at the root of the blue tree and trace along the border of the tree in a clockwise direction until you end at the root. The first time you encounter a vertex, insert a 1 in the Dyck path, the second time you encounter the vertex insert a -1 . Let v_1, v_2, \dots, v_n be the order of the vertices as they are encountered by performing this DFS traversal of the blue tree in a clockwise direction and define L to be the resulting linear order on the vertices. Let d_i be number of incoming red edges incident to v_i . Let r be the number of incoming red edges incident to s_{red} . The top Dyck path is as follows $1(-1)^{d_2}, 1, (-1)^{d_3}, 1, (-1)^{d_4} \dots 1(-1)^{d_n} 1(-1)^r$. The structure of the 3-orientation guarantees that the top path will never cross below the bottom path. See Figure 4, and [2] for details.

We bound the mixing rate of \mathcal{M}_{EF} by comparing it to \mathcal{M}_{DK} , an efficient Markov chain on (pairs of) Dyck paths introduced by Luby, Randall and Sinclair [18]. The algorithm proceeds as follows. At each

⁽ⁱ⁾ The abbreviation u.a.r. stands for uniformly at random.




Fig. 4: The bijection between 3-orientations and Dyck paths.

step select a point on one of the two Dyck paths uniformly at random. If the point is a local maximum (or minimum) then push it down (or up) with probability 1/2 as shown in Figure 5(a-b). If this move is blocked by a local maximum (or minimum) in the bottom (or top) Dyck path as shown in Figure 5c then push both Dyck paths down (or up) with probability 1/2 as shown in Figure 5(c-d). The following theorem due to Wilson [25] bounds the mixing time of \mathcal{M}_{DK} .




Fig. 5: Two moves of the Markov chain \mathcal{M}_{DK} .

Theorem 7 (Wilson) *The chain \mathcal{M}_{DK} has mixing time $\tau(\mathcal{M}_{DK}) = \Theta(n^3 \log n)$.*

Using the above bijection, the Markov chain \mathcal{M}_{DK} on Dyck paths can be translated into a Markov chain on 3-orientations of triangulations, but its moves are quite unnatural in that setting. We obtain a bound on the mixing time of \mathcal{M}_{EF} using Theorem 7 together with a careful comparison argument.

4.2 Fast mixing of \mathcal{M}_{EF}

Next we show that \mathcal{M}_{EF} is efficient for sampling from Ψ_n by comparing \mathcal{M}_{EF} and \mathcal{M}_{DK} . The comparison theorem of Diaconis and Saloff-Coste [8] relates the mixing times of two reversible Markov chains P and P' on the state space Ψ , with the same stationary distribution π , and mixing times τ and τ' respectively. Let $E(P) = \{(X, Y) : P(X, Y) > 0\}$ and $E(P') = \{(X, Y) : P'(X, Y) > 0\}$ denote the transitions of the two Markov chains, viewed as directed graphs. For each $X, Y \in \Psi$ with $P'(X, Y) > 0$, define a canonical path γ_{XY} using a sequence of states $X = X_0, X_1, \dots, X_k = Y$ with $P(X_i, X_{i+1}) > 0$, and let $k = |\gamma_{XY}|$ denote the length. Let $\Gamma(Z, W) = \{(X, Y) \in E(P') : (Z, W) \in \gamma_{XY}\}$ be the set of canonical paths that use the transition (Z, W) of P . Let $\pi_* = \min_{X \in \Psi} \pi(X)$. Finally, define $A = \max_{(Z, W) \in E(P)} \{(\sum_{\Gamma(Z, W)} |\gamma_{XY}| \pi(X) P'(X, Y)) / (\pi(Z) P(Z, W))\}$. We will use a version of the comparison theorem in terms of mixing times, due to Randall and Tetali [21].

Theorem 8 (Randall and Tetali) *With the above notation, we have $\tau(P) \leq 4 \log(4/\pi_*) A \tau'$.*

Next we introduce some notation. Let c_1 be blue, c_2 be red, and c_3 be green. Given a vertex v and $i \in \{1, 2, 3\}$, the unique outgoing edge of v with color c_i is called v 's c_i edge. We also define the *first* (last) incoming c_i -edge of v to be the incoming c_i -edge of v that is in a facial triangle with v 's c_{i-1} edge (respectively, v 's c_{i+1} edge, where the subscripts are taken modulo 3).




Fig. 6: A sequence of red/green swaps.

Proof sketch of Theorem 4. In order to apply Theorem 8 to relate the mixing time of \mathcal{M}_{EF} with the mixing time of \mathcal{M}_{DK} we need to define for each transition of \mathcal{M}_{DK} a canonical path using transitions of \mathcal{M}_{EF} . There are several cases to consider, depending on whether a move affects the top path, the bottom path or both, and whether it inverts a valley or a peak. If the move $e = (X, Y)$ affects both paths, we view the move as two separate moves (X, Z) and (Z, Y) , one on each path, and we concatenate the canonical paths as follows: $\gamma_{X,Y} = (\gamma_{X,Z}, \gamma_{Z,Y})$. Hence in the following, we assume that the transitions of \mathcal{M}_{DK} affect only one Dyck path. Let $e = (X, Y)$ be such a transition which moves the i^{th} 1 (where $i > 1$) on the top path to the right one position (i.e. the Dyck path move swaps the i^{th} 1 with a -1 on it's right, changing a peak to a valley). From the bijection, we know this move does not affect the blue tree and corresponds to, in the red tree, increasing the incoming degree of v_i by one and decreasing the incoming degree of v_{i+1} by one. If v_i and v_{i+1} are adjacent in the blue tree (there is a blue edge $\overrightarrow{v_{i+1}v_i}$) this implies that there is a red/green swap involving v_i 's green edge and v_{i+1} 's first incoming red edge. This swap is exactly the peak to valley move, so $\gamma_{XY} = e$. Otherwise, we define two stages in the path γ_{XY} .

To assist in defining the paths, let g_i be the parent of v_i in the green tree. Let v_j to be the parent of g_i in the red tree. Notice that $L(v_i) < L(v_{i+1}) \leq L(v_j)$, since $L(v_j) > L(g_i)$ and v_i 's red and green edges prevent v_j from satisfying $L(v_{i+1}) \geq L(v_j) \geq L(g_i)$ as shown in Figure 7a. In the first stage of the path




Fig. 7: (a) The vertex v_i 's red and green edges prevent v_j from satisfying $L(v_{i+1}) \geq L(v_j) \geq L(g_i)$. (b) The canonical path to move a peak down to a valley in the top Dyck path.

γ_{XY} we make the sequence of red/green swaps centered at g_i that move the red edge $\overrightarrow{g_i v_j}$ to $\overrightarrow{g_i v_i}$ without affecting any other red edges as shown in Figure 7b, step 1 (see Figure 6 for detail on the sequence of swaps). In the second stage we transfer an incoming red edge from v_{i+1} to v_j , completing γ_{XY} .

Given a transition (Z, W) of \mathcal{M}_{EF} we must bound the number of canonical paths γ_{XY} using this edge. To do so, we analyze the amount of information needed in addition to (Z, W) to determine X and Y uniquely. We record the vertex v_i and the vertex v_j . If v_i and v_{i+1} are adjacent we record v_{i+1} instead of v_j . Notice in this case the canonical path only involves red/green and green/red swaps. If we are moving a red edge to a higher vertex in L then we are in stage 2 and otherwise we are in stage 1. Given this information we can uniquely recover X and Y . We only need to record two vertices, so in this

case there are at most n^2 canonical paths which use any edge (Z, W) .

We defer the other cases to the full version. Briefly, the canonical path for a valley to peak move on the top Dyck path is very similar to the above case. However, a move on the bottom Dyck path is quite different and more complex because these moves significantly change the triangulation. While above, the moves do not affect the blue tree, moves on the bottom path can make large changes to the blue tree, which in turn significantly alter the red and green trees as well. The effect of such a move on the blue tree in Figure 8 is to replace \vec{ac} by \vec{ab} and make all of the blue children of \vec{ab} point to b . An example of the canonical path is given in Figure 8. In stage 1 (Figure 8(a-c)), the blue edge of a moves from c to b . Then in stage 2 (Figure 8(c-d)), a 's red edge moves into position for stage 3 (Figure 8(d-e)), where all incoming blue edges to a move down to point to b . Finally, in stage 4 we repair the red tree. \square




Fig. 8: Canonical path to move a valley up to a peak in the blue tree.

Acknowledgments

The last author thanks Stefan Felsner for introducing him to the problem of sampling 3-orientations and for providing several useful links to the relevant literature.

References

- [1] I. Bezakova, D. Stefankovic, V. Vazirani, and E. Vigoda. Accelerating Simulated Annealing for the Permanent and Combinatorial Counting Problems. *SIAM Journal of Computing*, **37**: 1429-1454, 2008. Preliminary version in SODA 2006.
- [2] N. Bonichon. A bijection between realizers of maximal plane graphs and pairs of non-crossing Dyck paths. *Discrete Math.* **298**: 104–114, 2005.
- [3] N. Bonichon, B. Le Saëc and M. Mosbah. Wagner’s Theorem on Realizers. *Automata, Languages and Programming*. **2380**: 776, 2002.
- [4] N. Bonichon and M. Mosbah. Watermelon uniform random generation with applications. *Theoretical Computer Science*. **307(2)**: 241–256, 2003.
- [5] E. Brehm. 3-orientations and Schnyder 3-tree-decompositions. Master’s Thesis. Freie Universität Berlin, 2000.
- [6] Y. Chiang, C. Chi and H. Lu. Orderly spanning trees with applications to graph encoding and graph drawing. 12th ACM-SIAM Symp. on Discrete Algorithms (SODA), 506–515, 2001.
- [7] P. Creed. Sampling Eulerian orientations of triangular lattice graphs. *Journal of Discrete Algorithms*, **7**: 168–180, 2009.
- [8] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *Annals of Applied Probability*, **3**: 696–730, 1993.

- [9] M. Dyer and C. Greenhill. A more rapidly mixing Markov chain for graph colorings. *Random Structures and Algorithms*, **13**: 285–317, 1998.
- [10] S. Greenberg and D. Randall. Slow mixing of Markov chains using fault lines and fat contours. *Algorithmica*, **58**: 911–927, 2010.
- [11] S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order*, **18**: 19–37, 2001.
- [12] S. Felsner. *Geometric Graphs and Arrangements*, Vieweg Verlag, 2004.
- [13] S. Felsner and D. Heldt. Tower Moves. *Work in progress*.
- [14] S. Felsner and F. Zickfeld. On the Number of Planar Orientations with Prescribed Degrees. *Electron. J. Comb.*, **15(1)**: Research paper R77, 41, 2008.
- [15] H. de Fraysseix and P. Ossona de Mendez. On topological aspects of orientation, *Discrete Mathematics*, **229**: 57–72, 2001.
- [16] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation*, **82**: 93–133, 1989.
- [17] M. Jerrum, A. Sinclair and E. Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries. *JACM*, **51**: 671–697, 2004.
- [18] M. Luby, D. Randall and A. Sinclair. Markov Chains for Planar Lattice Structures. *SIAM Journal on Computing*, **31**: 167–192, 2001.
- [19] L. McShine and P. Tetali. On the mixing time of the triangulation walk and other Catalan structures. *Randomization Methods in Algorithm Design*, **43**: 147–160, 1999.
- [20] M. Mihail and P. Winkler. On the number of Eulerian orientations of a graph. *Algorithmica*, **16**: 402–424, 1996.
- [21] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics*, **41**: 1598–1615, 2000.
- [22] W. Schnyder. Embedding Planar Graphs on the Grid. 1st ACM-SIAM Symposium on Discrete Algorithms (SODA), 138–148, 1990.
- [23] W. Schnyder. Planar graphs and poset dimension. *Order* 5, 323–343, 1989.
- [24] A. Sinclair. *Algorithms for random generation & counting: a Markov chain approach*. Birkhäuser, Boston, 1993.
- [25] D. Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Annals of Applied Probability*, **14**: 274–325, 2004.