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

Sarah Miracle[*]   Dana Randall[†]   Amanda Pascoe Streib[‡]   Prasad Tetali[§]

Abstract

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 powerful 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. It was shown previously that this chain connects the state space and 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. The triangle-reversing and edge-flipping Markov chains both arise in the context of sampling other combinatorial structures, such as Eulerian orientations and triangulations of planar pointsets, so our results here may shed light on the mixing rate of these related chains as well.

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. In this paper 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 a fixed number of vertices. In the case that the triangulation is fixed, we consider 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. These chains arise in other contexts such as sampling Eulerian orientations and triangulations of fixed planar point sets, so there is additional motivation for understanding their convergence rates.

1.1   3-orientations

Given an undirected graph $G = (V, E)$ and a function $\alpha : V \to \mathbb{Z}^+$, 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 [12, 13, 15] and references therein 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 [22], and more recently Creed [8] showed that it remains $\#P$-complete even when restricted to the class of 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 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 [26, 7] and poset dimension theory [27]. Several intriguing enumeration problems remain open, such as the complexity of enumerating 3-orientations of a planar triangulation (see e.g., [15].) We study the problem of sampling 3-orientations of a fixed (planar) triangulation, as well as sampling 3-orientations of all triangulations with n internal vertices. In particular, we analyze the mixing times of two natural Markov chains for these problems, which were introduced previously but had thus far resisted analysis.

1.2 Results

First, we study the problem of sampling 3-orientations of a fixed triangulation, which was stated as an open problem by Felsner and Zickfeld [15]. 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 [15]). 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 celebrated results of Jerrum, Sinclair and Vigoda [18]), 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 outdegree at each vertex. Brehm [6] showed that for any fixed triangulation T, $\mathcal{M}_{TR}$ connects the state space $\Psi(T)$ of all 3-orientations of T. We also consider a related "cycle-reversing" chain, $\mathcal{M}_{CR}$, that can also 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 [20]. 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 can 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 vertices. An interesting special 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 [8] independently solved the sampling problem in this special case using an approach similar to ours 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 [20] and the 8-vertex model [11]. 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 [6].

In fact, this structure allows us to extend our analysis to certain non-4-connected triangulations that can have vertices of degree greater than six.

Next, we prove that when the maximum degree is unbounded, the chain $\mathcal{M}_{TR}$ 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 [14] 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. It would be interesting to know if there is such a family with bounded degree.

The second problem we study is sampling from Ψ_n, the set of all 3-orientations arising from all possible triangulations on n internal vertices. 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 nth Catalan number. Since exact enumeration is possible, we can sample using the reduction to counting; this was explicitly worked out by Bonichon and Mosbah [5]. 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 [4] showed the the chain $\mathcal{M}_{EF}$ connects the state space 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 other 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 additional 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., [21]). Specifically, we prove:

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

1.3   Techniques

The primary challenge behind the proofs of these results is extracting the right combinatorial insights to understand the dynamics in the context of Schnyder woods and 3-orientations. Fortunately, there is a long history examining the rich structure of Schnyder woods. We were able to extend these results in several new ways, thus allowing us to bound the mixing rates of these chains. The proof of Theorem 1 for $\mathcal{M}_{CR}$ involves a complex coupling argument that is straightforward if T is the triangular lattice, but requires more work to generalize to all triangulations with $\Delta_I \leq 6$. To prove Theorem 3, we produce an intricate triangulation T which is carefully constructed to reveal an exponentially small cut in the state space $\Psi(T)$. Although our choice of T may seem complicated, it was carefully architected using properties of 3-orientations to show that the Markov chain may be slow. Our proof of Theorem 4 involves a detailed application of the comparison method to bound the mixing time of $\mathcal{M}_{EF}$ by relating it to a local Markov chain on Dyck paths, $\mathcal{M}_{DK}$, whose mixing time is known (see [20, 29]). The key obstacle here is decomposing moves of $\mathcal{M}_{DK}$ into moves of $\mathcal{M}_{EF}$ while avoiding congestion. This is especially challenging because although $\mathcal{M}_{DK}$ is local in the setting of Dyck paths, in the context of 3-orientations it can make global changes to a 3-orientation in a single step.

2   Preliminaries

We begin with background on 3-orientations, Schnyder woods, and Markov chains. The *external vertices* and *edges* are those on the infinite face of a planar triangulation. All other vertices and edges are called *internal*. Let $\mathcal{T}_n$ be the set of planar triangulations with n internal vertices. Given a triangulation $T \in \mathcal{T}_n$,

let $\Psi(T)$ be the set of 3-orientations of T. Let $\Psi_n = \cup_{T \in \mathcal{T}_n} \Psi(T)$ denote the set of all 3-orientations over all triangulations on n internal vertices. There are several interesting bijections between 3-orientations of triangulations and other combinatorial objects; for example, in Section 4.1 we present a bijection between Ψ_n and pairs of Dyck paths. Fraysseix and Ossona de Mendez defined a bijection between $\Psi(T)$ and the Schnyder woods of T[16]. A **Schnyder wood** 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 blue, red and green, 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).
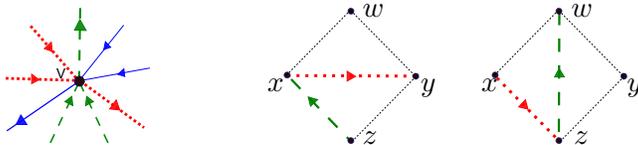
Figure 1: (a) The vertex condition. (b) A red/green swap.

In our figures, we differentiate the colors of edges in the Schnyder woods by dashed lines (green), dotted lines (red), and solid lines (blue). See Figure 4 for an example of a Schnyder wood. 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 at one of the external vertices. See [16] for a proof of the bijection. It is often convenient to consider the Schnyder wood associated with a given 3-orientation, because the vertex condition allows us to infer some of the global information about the 3-orientation by looking locally at each vertex. Throughout the paper, 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, where in this case, n is the number of internal vertices of the triangulations.

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. $\frac{1}{2}$ reverse t to obtain σ_{i+1}.

- Otherwise, $\sigma_{i+1} = \sigma_i$.

Brehm proved that $\mathcal{M}_{TR}$ connects the state space, $\Psi(T)$ [6]. Since all valid moves have the same transition probabilities, $\mathcal{M}_{TR}$ converges to the uniform distribution over $\Psi(T)$.

3.1 Fast mixing 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 auxilliary 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 [20]. 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 2a). We call this edge the *disagreeing edge* of f. Define a tower of length k to be a path of faces $f_1, f_2, \ldots, f_k$ 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).
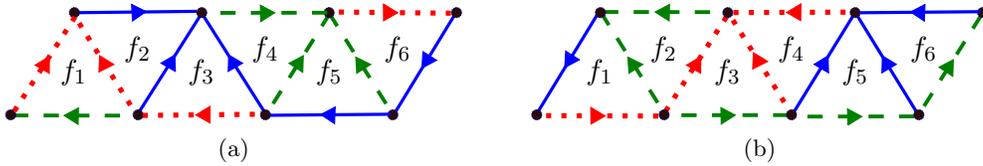
Figure 2: A tower of length 6.

The Tower Markov chain $\mathcal{M}_{CR}$

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 prob. $\begin{cases} \frac{1}{6k} : & k \geq 2 \\ \frac{1}{2} : & k = 1 \end{cases}$
  reverse the orientation of this tower to obtain σ_{i+1}.

- Otherwise, $\sigma_{i+1} = \sigma_i$.

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, which then implies the general result. 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.

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 $\tau(\mathcal{M}_{CR}) = O(n^5)$.*

First we prove Theorem 5 and 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 [6]. The main tool we use to prove Theorem 5 is the path coupling theorem due to Dyer and Greenhill [10].

Theorem 6 (Dyer and Greenhill). *Let ϕ be an integer-valued metric on $\Psi \times \Psi$ which takes values in $\{0, ..., B\}$. Let U be a subset of $\Psi \times \Psi$ such that for all $(x_t, y_t) \in \Psi \times \Psi$ there exists a path $x_t = z_0, z_1, ..., z_r = y_t$ between x_t and y_t such that $(z_i, z_{i+1}) \in U$ for $0 \le i < r$ and $\sum_{i=0}^{r-1} \phi(z_i, z_{i+1}) = \phi(x_t, y_t)$. Let M be a Markov chain on state space Ψ with transition matrix P. Consider any random function $f : \Psi \to \Psi$ such that $P[f(x) = y] = P(x, y)$ for all $x, y \in \Psi$, and define a coupling of the Markov chain by $(x_t, y_t) \to (x_{t+1}, y_{t+1}) = (f(x_t), f(y_t))$. If $E[\phi(x_{t+1}, y_{t+1})] \le \phi(x_t, y_t)$, for all $x_t, y_t \in U$, let $\alpha > 0$ satisfy $\Pr[\phi(x_{t+1}, y_{t+1}) \ne \phi(x_t, y_t)] \ge \alpha$ for all t such that $x_t \ne y_t$. Then the mixing time of M is $\tau(M) \le 2\left\lceil \frac{eB^2}{\alpha} \right\rceil$.*

Proof of Theorem 5. Let T be a 4-connected planar triangulation such that $\Delta_I(T) \le 6$. First we prove that $\mathcal{M}_{CR}$ is rapidly mixing on state space $\Psi(T)$. 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, which 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 the $\mathcal{M}_{CR}$ selects the face f and chooses to direct the cycle clockwise or counterclockwise, each of which happens with probability $1/(2(2n+1))$. Moreover, 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.

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, and therefore increase the distance. Any tower that is neither good nor bad does not contain an edge of f, so it is neutral with respect to distance.

Suppose $k \ge 1$, $(f_1, f_2, \ldots, f_k)$ is a good tower in σ, and f_k is adjacent to f. We claim that $(f_1, f_2, \ldots, f_k, f)$ 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 $(f_1, f_2, \ldots, f_k)$ is possible. We must check two things: that $(f_1, f_2, \ldots, f_k, f)$ 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 $\{f_1, f_2, \ldots, f_k\}$, and $(f_1, f_2, \ldots, f_k)$ 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 (see the Appendix Figure 1). 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 $k \ge 1$ 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 $(f_1, f_2, \ldots, f_k)$ is a good towner in σ, then $(f_1, f_2, \ldots, f_k, f)$ is a good tower in τ. On the other hand, it should be clear that if $(f_1, f_2, \ldots, f_k, f)$ is a good tower of length $k + 1 \ge 2$ that ends on f in σ then $(f_1, f_2, \ldots, f_k)$ is a good tower of length k in τ. In either case, if $k \ge 2$ 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.

We claim that 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 $(f_1, f_2, \ldots, f_k)$ 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 will show 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'. Moreover, 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 $\mathbb{E}[\Delta d] \leq \frac{1}{2n+1}\left(-2\left(\frac{1}{2}\right) + 3\left(\frac{1}{3}\right)\right) = 0$. We will see that for any (σ_t, τ_t) (not necessarily at distance 1), $\Pr(d(\sigma_t, \tau_t) \neq d(\sigma_{t+1}, \tau_{t+1})) \geq 1/(4n_2)$. This follows from the connectivity proof of Brehm [6]. For any such pair, there exists a path of transitions of $\mathcal{M}_{TR}$ from σ_t to τ_t of length $d(\sigma_t, \tau_t)$, and the first of these transitions occurs with probability at least $1/(4n+2)$, since each face is selected with probability at least $1/(2n+1)$, and the move succeeds with probability $1/2$. Clearly this move decreases the distance (the number of steps between the configurations) by 1. By Theorem 6 and the bound on the distance between any two 3-orientations given in the Appendix, Lemma A.1a, we see the mixing time $\tau(\mathcal{M}_{CR})$ of $\mathcal{M}_{CR}$ satisfies $\tau(\mathcal{M}_{CR}) \leq \lceil e(2n+1)^5/2 \rceil = O(n^5)$.

Next we use the comparison theorem, Theorem 9 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 [6], 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 Appendix, Section A.1.2.

3.2 $\mathcal{M}_{TR}$ can mix slowly

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

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

The following theorem relates the conductance and mixing time (see, e.g., [17, 28]).

Theorem 7. *For any Markov chain with conductance Φ, $\tau \geq (4\Phi)^{-1} - 1/2$.*

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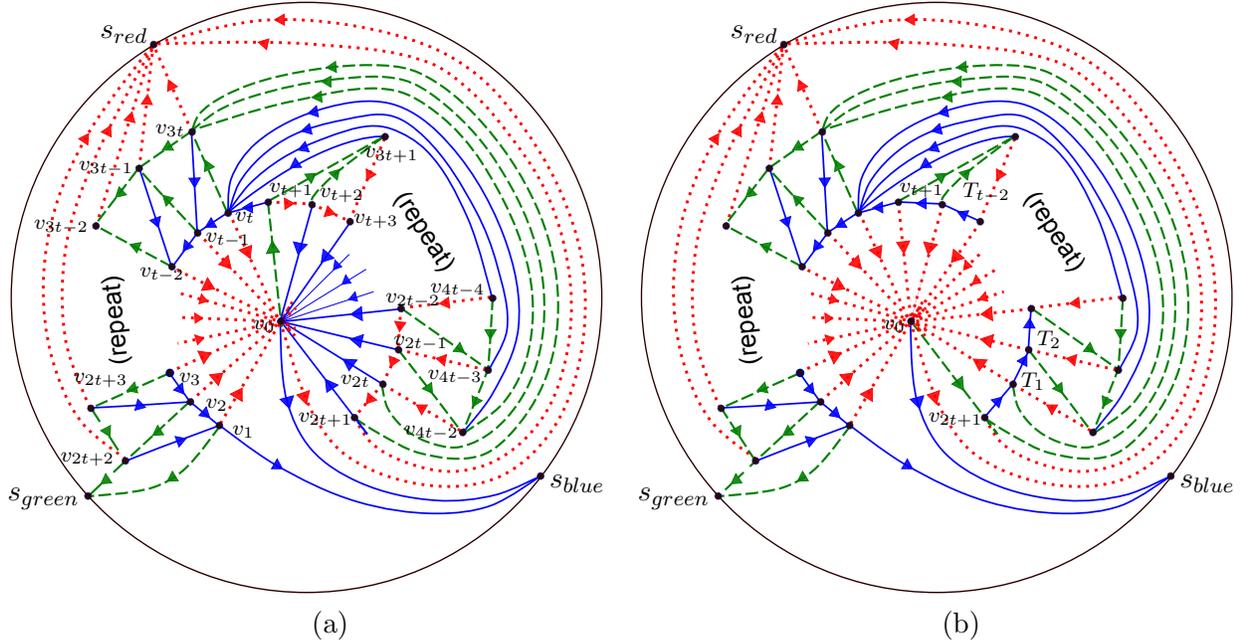
Figure 3: (a) A triangulation for which $\mathcal{M}_{TR}$ mixes slowly. (b) There is an exponential number of 3-orientations with edge (v_0, v_{t+1}) colored red, corresponding to different orientations of $T_1, T_2, \ldots, T_{t-2}$.

First consider the 3-orientation in Figure 3b which is in D. Notice the triangles $T_1, T_2, \ldots T_{t-2}$ do not share any edges and reversing these triangles does not change the color of any edges adjacent to v_0. Each of these triangles has 2 possible orientations and each of these 2^t choices of the orientations of the triangles gives a distinct 3-orientation with edge (v_0, v_{t+1}) colored red implying that $|D| \geq 2^t = 2^{(n-6)/4}$. A similar argument described in Section A.1.3 of the Appendix shows that $|\overline{D}| \geq 2^{t-1} = 2^{(n-2)/4}$.

Next, we show 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 $(v_0, v_1), (v_0, v_2), \ldots, (v_0, v_t)$ must all be directed toward v_0 and colored red; this is because the edge (v_0, s_{blue}) is blue and directed towards s_{blue} in every 3-orientation of G. Similarly, edges $(v_0, v_{t+2}), (v_0, v_{t+3}), \ldots, (v_0, v_{2t+1})$ must all be blue and directed toward v_0. Since v_{t+1} has degree 4 and has the incoming green edge (v_0, v_{t+1}), the other edges incident to v_{t+1} are determined; (v_t, v_{t+1}) is blue, (v_{t+1}, v_{3t+1}) is green and (v_{t+1}, v_{t+2}) is red all directed away from v_{t+1}. Knowing the colors and direction of the edges incident to v_0 and v_{t+1} forces the color and direction of the edges incident to v_{t+2}. Similarly v_{t+1} and v_{t+2} force v_{3t+1} and v_0, v_{t+2} and v_{3t+1} force v_{t+3} and so on until the color and direction of all edges incident to vertices $v_{t+1}, v_{t+2}, \ldots, v_{2t+1}$ and $v_{3t+1}, v_{3t+2}, \ldots, v_{4t-2}$ are forced. Next consider v_t; we know the edges $(v_t, v_{t+1}), (v_t, v_{3t+1}), (v_t, v_{3t+2}), \ldots, (v_t, v_{4t-2})$ are all blue and directed toward v_t which implies (v_t, v_{3t}) and (v_t, v_{t-1}) must be directed outward and green and blue respectively. Now consider v_{3t}; we have already shown that all edges incident to v_{3t} except for $(v_{3t}, v_{t-1}), (v_{3t}, v_{3t-1}), (v_{3t}, v_{s_1})$ have a forced color and are directed inward. Thus, these 3 edges must all be directed outwards with colors blue, green and red respectively. Similarly, knowing the colors and directions of all edges incident to v_{3t}, v_t and v_0 forces the colors and directions of edges incident to v_{t-1} and v_{3t}, v_{t-1} forces v_{3t-1} and so on for the remaining vertices. Since all the colors and directions of the edges are forced there is exactly one 3-orientation with (v_0, v_{t+1}) colored green. 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 also changes the direction.

Finally, given the bound on $|D|$ and $|\overline{D}|$, we derive a bound on the mixing time of $\mathcal{M}_{TR}$. Let g be the single 3-orientation which has edge (v_0, v_{t+1}) colored green. If $\pi(D) \leq 1/2$ then combining the definition of conductance with the bound on $|D|$ yields

$$
\begin{aligned}
\Phi_{\mathcal{M}_{TR}} &\leq \tfrac{1}{\pi(D)} \sum_{d_1 \in D, d_2 \in \overline{D}} \pi(d_1)\mathcal{P}(d_1, d_2) \\
&= \tfrac{1}{\pi(D)} \sum_{d_2 \in \overline{D}} \pi(g)\mathcal{P}(g, d_2) \leq \tfrac{\pi(g)}{\pi(D)} \leq \tfrac{1/Z}{2^{(n-6)/4}/Z} = \tfrac{1}{2^{(n-6)/4}}.
\end{aligned}
$$

If $\pi(D) > 1/2$ then $\pi(\overline{D}) \leq 1/2$ and so by detailed balance and the bound on $|\overline{D}|$,

$$
\begin{aligned}
\Phi_{\mathcal{M}_{TR}} &\leq \tfrac{1}{\pi(\overline{D})} \sum_{d_1 \in \overline{D}, d_2 \in D} \pi(d_1)\mathcal{P}(d_1, d_2) &&= \tfrac{1}{\pi(\overline{D})} \sum_{d_1 \in \overline{D}} \pi(d_1)\mathcal{P}(d_1, g) \\
&= \tfrac{1}{\pi(\overline{D})} \sum_{d_1 \in \overline{D}} \pi(g)\mathcal{P}(g, d_1) &&\leq \tfrac{\pi(g)}{\pi(\overline{D})} \leq \tfrac{1/Z}{2^{(n-2)/4}/Z} = \tfrac{1}{2^{(n-2)/4}}.
\end{aligned}
$$

In both cases, $\Phi_{\mathcal{M}_{TR}} \leq 2^{-(n-6)/4}$. Applying Theorem 7 proves that $\tau(\mathcal{M}_{TR}) \geq 2^{(n-14)/4)} - \frac{1}{2}$. $\qquad\square$

Remark 1. *Combining this result with the comparison argument in Section 3.1 shows that $\tau(\mathcal{M}_{CR})$ can also take exponential time to converge.*

4 Sampling 3-orientations of triangulations on n internal vertices

In this section, we define the local Markov chain $\mathcal{M}_{EF}$ which can be used to sample from the uniform distribution on Ψ_n and then show that $\mathcal{M}_{EF}$ is always rapidly mixing. Our argument relies on a bijection with pairs of Dyck paths to relate the mixing time of a local Markov chain on Dyck paths to $\mathcal{M}_{EF}$ using the comparison method [9]. 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 adjacent facial triangles T_1 and T_2 with a shared edge $\overrightarrow{xy}$ u.a.r.[1]

- Choose an edge $\overrightarrow{zx}$ from $T_1 \cup T_2$ u.a.r., if one exists. With probability 1/2 replace the path $\{\overrightarrow{zx}, \overrightarrow{xy}\}$ by the path $\{\overrightarrow{xz}, \overrightarrow{zw}\}$ 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 $\mathbf{c_j/c_i}$ **swap**. Bonichon, Le Saëc and Mosbah showed in [4] 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 [3]. Dyck paths can be thought of as strings $a_1 a_2 \cdots a_{2n}$ 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 $v_1, v_2, \ldots, v_n$ 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 $1(-1)^{d_2}, 1, (-1)^{d_3}, 1, (-1)^{d_4} \ldots 1(-1)^{d_n} 1(-1)^r$. The structure of the 3-orientation guarantees that the top path will never cross below the bottom path. See Figure 4, and [3] for details.

---

[1]The abbreviation u.a.r. stands for uniformly at random.
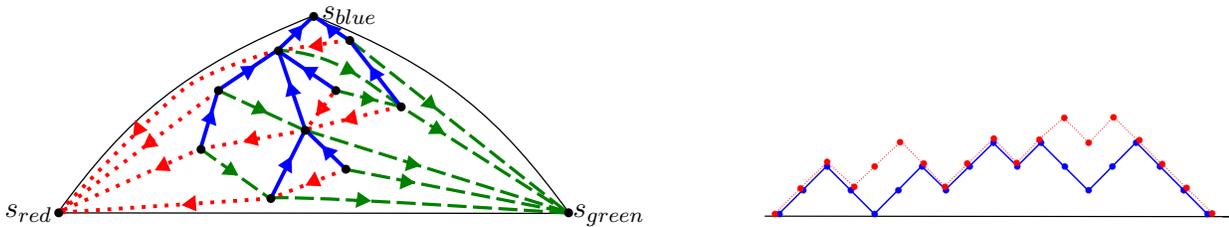
Figure 4: The bijection between 3-orientations and Dyck paths.

We will deduce that $\mathcal{M}_{EF}$ is rapidly mixing by comparing it to $\mathcal{M}_{DK}$, an efficient Markov chain on (pairs of) Dyck paths which was introduced by Luby, Randall and Sinclair [20]. The algorithm proceeds as follows, at each 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 [29] bounds the mixing time of $\mathcal{M}_{DK}$.
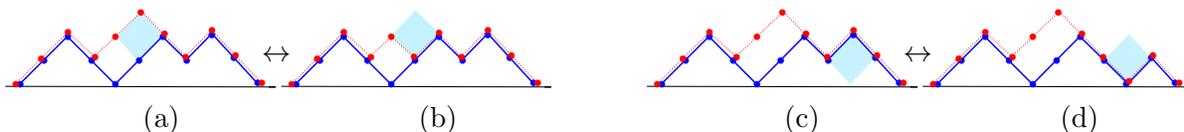
Figure 5: Two moves of the Markov chain $\mathcal{M}_{DK}$.

Theorem 8 (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 8 together with a careful comparison argument.

4.2 Mixing Time of $\mathcal{M}_{EF}$

Next we show that $\mathcal{M}_{EF}$ is an efficient algorithm for sampling from Ψ_n by comparing the moves of $\mathcal{M}_{EF}$ to the moves of $\mathcal{M}_{DK}$. The comparison theorem of Diaconis and Saloff-Coste [9] allows us to compare the mixing times of two reversible Markov chains P and P' on the state space Ψ. Assume they have 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 sets of 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 $X = X_0, X_1, \cdots, X_k = Y$ with $P(X_i, X_{i+1}) > 0$, and let $|\gamma_{XY}|$ denote the length of the path. 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

$$A = \max_{(Z,W) \in E(P)} \left\{ \frac{1}{\pi(Z)P(Z,W)} \sum_{\Gamma(Z,W)} |\gamma_{XY}| \pi(X) P'(X,Y) \right\}.$$

We will use the following version of the comparison Theorem, due to Randall and Tetali [25].

Theorem 9 (Randall and Tetali). *With the above notation, we have $\tau(P) \le 4\log(4/\pi_*) A \tau'$.*

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

10

(respectively, v's c_{i+1} edge, where the subscripts are taken modulo 3). We will often need to move a c_j edge, say $\overrightarrow{vx}$, from some neighbor x of v to another neighbor y of v across several c_i edges. This is achieved through a sequence of c_j/c_i swaps as in Figure 6.
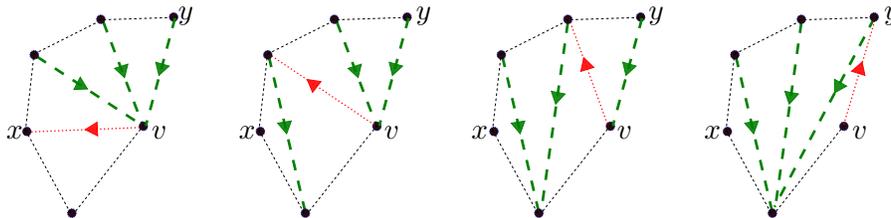
Figure 6: A sequence of red/green swaps.

Proof sketch of Theorem 4. In order to apply Theorem 9 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 $\mathcal{M}_{EF}$ and show that each transition of $\mathcal{M}_{EF}$ is not used by too many canonical paths. We consider four cases, depending on whether the move of $\mathcal{M}_{DK}$ takes a peak to a valley or a valley to a peak and which Dyck path it involves. 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 a transition of $\mathcal{M}_{DK}$ that moves a peak to a valley on the top Dyck path. We will define a canonical path from X to Y using transitions of $\mathcal{M}_{EF}$. Suppose e 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 the peak to valley move, so $\gamma_{XY} = e$. Otherwise, we define two stages in the path γ_{XY}.

Let g_i be the parent of v_i in the green tree. By Lemma A.2, g_i is not s_{green} (this somewhat technical proof is deferred to the appendix, as it does not lend much insight to the reader). Next, 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.
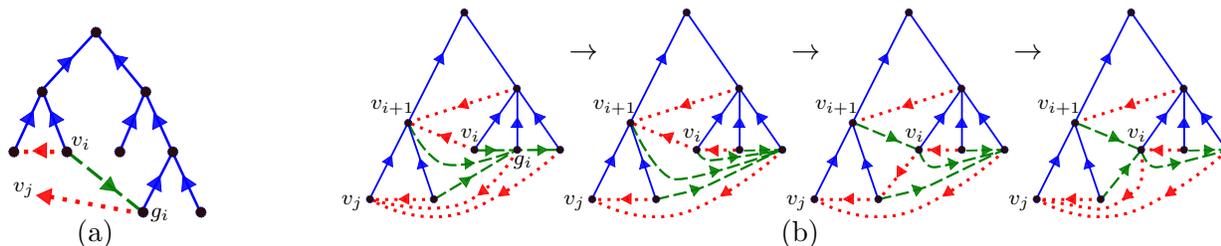
Figure 7: (a) The vertex v_i's red and green edges prevent v_j from satisfying $L(v_{i+1}) \geq L(v_j) \geq L(g_i)$. (b) The canonical path to move a peak down to a valley in the top Dyck path.

In the first stage of the 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}. Recall that $L(v_j) > L(v_{i+1})$, so we do this iteratively by moving an incoming red edge $\overrightarrow{yx}$ either to one of x's neighbors in the blue tree that is larger in L or to x's parent in the red tree, which is also larger in L, if it is a leaf and has no neighbors as shown in Figure 7b, step 2 and 3. We claim one of these moves is always possible. If x has a neighbor y in the blue tree such that $L(y) = L(x) + 1$ then there is a green/red swap centered at x and involving x's green edge that moves an incoming red

edge from x to y as desired. Next, if x is a leaf with red edge $\overrightarrow{xr_x}$ then again there is a green/red swap centered at x and involving x's green edges that moves an incoming red edge from x to r_x as desired. Finally notice that, using this canonical path, we will never bypass v_j because the original red edge $\overrightarrow{g_i v_j}$ blocked any blue leaves between v_{i+1} and g_i from having red parents higher in L than v_j.

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).
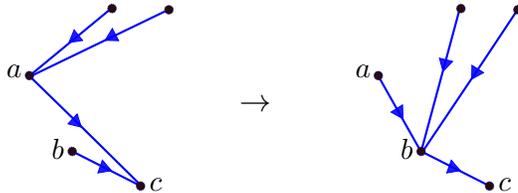
Figure 8: The effect of a valley to peak move of $\mathcal{M}_{DK}$ on the bottom Dyck path

We defer the other cases to Section A.2 of the Appendix. Briefly, the canonical path for a valley to peak move on the top Dyck path is very similar to the above case. However, a valley to peak move on the bottom Dyck path is quite different, and is in fact more involved. The complexity arises because the moves on the bottom path make significant changes to the triangulation. While moves on the top path do not affect the blue tree, moves on the bottom path can make large changes to the blue tree, which in turn significantly alters the red and green trees as well. The effect of such a move on the blue tree in Figure 8 is to replace $\overrightarrow{ac}$ by $\overrightarrow{ab}$ and make all of the blue children of $\overrightarrow{ab}$ point to b. An example of the canonical path is given in Figure 9. In stage 1 (Figure 9(a-c)), the blue edge of a moves from c to b. Then in stage 2 (Figure 9(c-d)), a's red edge moves into position for stage 3 (Figure 9(d-e)), where all incoming blue edges to a move down to point to b. Finally, in stage 4 we repair the red tree. The theorem then follows from the comparison theorem, Theorem 9. $\square$
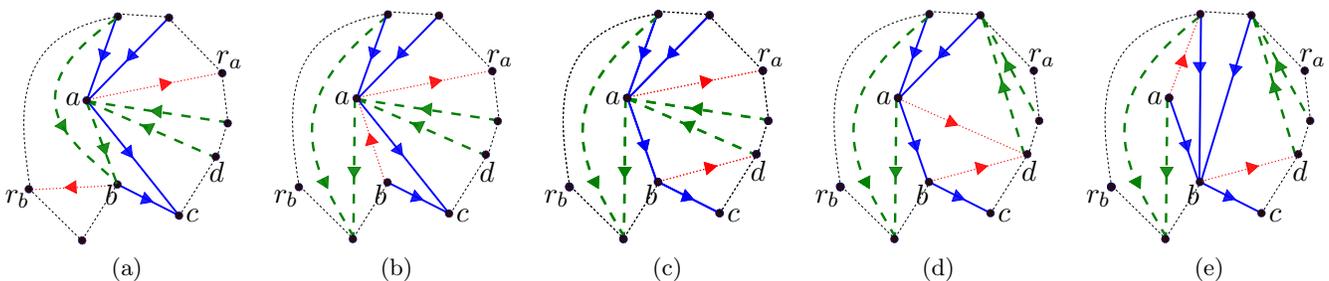
Figure 9: Canonical path to move a valley up to a peak in the blue tree.

5 Concluding Remarks

Several questions remain open. The complexity of enumerating Eulerian orientations in planar graphs of bounded degree is one of the foremost, as raised by [15]. Extending our fast mixing result to triangulations with larger degrees is a natural open problem; perhaps there is an alternate local chain which can sample efficiently from the set of 3-orientations corresponding to any fixed triangulation, without recourse to the bipartite perfect matching sampler of [18]. Finally, our slow mixing example involves vertices of degree $\Omega(n)$, and it would be of interest to find other constructions with constant maximum degree.

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] P. Bhakta, S. Miracle, D. Randall, and A.P. Streib. Mixing Times of Self-Organizing Lists and Biased Permutations. *submitted.*

[3] N. Bonichon. A bijection between realizers of maximal plane graphs and pairs of non-crossing Dyck paths. Discrete Math. **298:** 104–114, 2005.

[4] N. Bonichon, B. Le Saëc and M. Mosbah. Wagner's Theorem on Realizers. Automata, Languages and Programming. **2380:** 776, 2002.

[5] N. Bonichon and M. Mosbah. Watermelon uniform random generation with applications. Theoretical Computer Science. **307(2):** 241–256, 2003.

[6] E. Brehm. 3-orientations and Schnyder 3-tree-decompositions. Master's Thesis. Freie Universität Berlin, 2000.

[7] 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.

[8] P. Creed. Sampling Eulerian orientations of triangular lattice graphs. *Journal of Discrete Algorithms,* **7**: 168–180, 2009.

[9] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible Markov chains. *Annals of Applied Probability,* **3:** 696–730, 1993.

[10] M. Dyer and C. Greenhill. A more rapidly mixing Markov chain for graph colorings. *Random Structures and Algorithms,* **13:** 285–317, 1998.

[11] S. Greenberg and D. Randall. Slow mixing of Markov chains using fault lines and fat contours. *Algorithmica,* **58:** 911–927, 2010.

[12] S. Felsner. Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order,* **18:** 19–37, 2001.

[13] S. Felsner. Geometric Graphs and Arrangements, Vieweg Verlag, 2004.

[14] S. Felsner and D. Heldt. Tower Moves. *Work in progress.*

[15] S. Felsner and F. Zickfeld. On the Number of Planar Orientations with Prescribed Degrees. *Electron. J. Comb.,* **15(1):** Research paper R77, 41, 2008.

[16] H. de Fraysseix and P. Ossona de Mendez. On topological aspects of orientation, *Discrete Mathematics,* **229:** 57–72, 2001.

[17] M. Jerrum and A. Sinclair. Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation,* **82:** 93–133, 1989.

[18] 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.

[19] M. Jerrum, L. Valiant and V. Vazirani. Random generation of combinatorial structures from a uniform distribution. Theoretical Computer Science, **43:** 169–188, 1986.

[20] M. Luby, D. Randall and A. Sinclair. Markov Chains for Planar Lattice Structures. *SIAM Journal on Computing,* **31:** 167–192, 2001.

[21] 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.

[22] M. Mihail and P. Winkler. On the number of Eulerian orientions of a graph. *Algorithmica,* **16:** 402–424, 1996.

[23] M. Molloy, B. Reed and W. Steiger. On the Mixing Rate of the Triangulation walk. *DIMACS Workshop on Rand. Methods in Algorithm Design*, 179-190, 1997.

[24] D. Poulalhon and G. Schaeffer. Optimal Coding and Sampling of Triangulations. *Algorithmica,* **46:** 505-527, 2006.

[25] D. Randall and P. Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics*, **41:** 1598–1615, 2000.

[26] W. Schnyder. Embedding Planar Graphs on the Grid. 1st ACM-SIAM Symposium on Discrete Algorithms (SODA), 138–148, 1990.

[27] W. Schnyder. Planar graphs and poset dimension. Order 5, 323–343, 1989.

[28] A. Sinclair. *Algorithms for random generation & counting: a Markov chain approach.* Birkhäuser, Boston, 1993.

[29] D. Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Annals of Applied Probability*, **14:** 274–325, 2004.

A Appendix

A.1 Sampling 3-orientations of a fixed triangulation

We now consider the second context where we are given a fixed triangulation and would like to sample from its 3-orientations.

A.1.1 Background on 3-orientations of Planar Triangulations

In this section we will provide an overview of several results on 3-orientations of planar triangulations and the Markov chain $\mathcal{M}_{TR}$ which we will use in Section 3.1 to show that $\mathcal{M}_{TR}$ converges quickly when the maximum degree of the triangulation is at most 6. Brehm [6] provides a detailed analysis of the robust structure of 3-orientations of planar graphs. In particular, he constructs a framework which shows that the set of 3-orientations form a distributive lattice and that for a planar triangular graph G, any two 3-orientations of G are connected by a series of moves of the Markov chain $\mathcal{M}_{TR}$. As part of this effort, Brehm examines a potential function on the faces of the graph, which we will see is useful to upper-bound the number of 3-orientations for a given triangulation and the maximum distance between any two 3-orientations.

To show connectivity of $\mathcal{M}_{TR}$, Brehm considers first the case of 4-connected planar triangulations. In this case, every triangle of G is a face and he shows that it is possible to get between any two 3-orientations by a sequence of rotations of directed facial triangles. Suppose now that a planar triangulation G has exactly one triangle t which does not bound a face. For any non-facial cycle, Brehm shows that in any 3-orientation of G, the edges in the region bounded by that cycle that are incident to some vertex v of the cycle must be directed towards v. This implies that no face f contained within t that shares an edge with t can be bounded by a directed triangle; hence such faces can never be rotated. In fact, this implies (see [6] for details) that G can be regarded as the cross product of the triangulation $G|_t$, the restriction of G to the vertices on the boundary of t and the vertices contained within the region bounded by t and the triangulation $G \setminus t$ obtained by removing all vertices and edges contained within the region bounded by t. Thus by allowing $\mathcal{M}_{TR}$ to rotate arbitrary directed triangles, this amounts to extending $\mathcal{M}_{TR}$ to the triangulation $G \setminus t$ which is now 4-connected. The same arguments will hold when G has many non-facial triangles. Thus he obtains the following theorem.

Theorem A.1 (Brehm). *For any planar triangulation G, $\mathcal{M}_{TR}$ connects the set of all 3-orientations of G.*

In our setting, we use the fact that $G|_t$ is independent of $G \setminus t$ to show that the mixing time of $\mathcal{M}_{TR}$ is the maximum of the mixing times of each 4-connected piece of G, subject to the delay which results from the fact that $\mathcal{M}_{TR}$ only attempts to update one 4-connected piece at a time.

Brehm defines a *potential* X of a 4-connected planar triangulation as follows.

Definition A.1. *A potential X of a 4-connected planar triangulation G is a mapping $f \to x_f$ from the interior faces to the natural numbers such that*

- *$x_f = 0$ if the boundary of f contains an exterior edge*

- *$|x_f - x_g| \leq 1$ holds for any two adjacent faces f, g.*

The value of a potential X is defined by $\sum_f x_f$.

It turns out that there is a bijection between 3-orientations of G and a subset of the potentials of G, called *induced* potentials, and that each move of $\mathcal{M}_{TR}$ changes the potential of a face by ± 1 (see [6] for details). Every triangulation $G \in \mathcal{T}_n$ has $2n+1$ faces (not counting the infinite face). It is easy to see that for any potential of G, the maximum value for any face is at most $\lfloor \frac{2n+1}{2} \rfloor$ since each face can only differ

from its neighbors by at most 1 and faces adjacent to the boundary have value 0. Thus, the maximum number of steps of $\mathcal{M}_{TR}$ required to get between two 3-orientations of G is bounded by the maximum value of any potential of G. This implies Lemma A.1(a). Moreover, the number of 3-orientations of a graph G is bounded by the number of induced potentials of G. Since each face in a potential is within 1 from each of its adjacent faces, the number of induced potentials is at most 3^{2n+1}.

Lemma A.1. *Let G be a 4-connected planar triangulation.*

  (a) *The maximum distance between any two 3-orientations of G is at most $(2n+1)^2/2$.*

  (b) *The number of 3-orientations of G is at most 3^{2n+1}.*

A.1.2 Fast mixing for maximum degree at most 6

Let T be a planar triangulation whose maximum degree (of any internal vertex) is at most 6. In this section we prove that $\mathcal{M}_{TR}$ is rapidly mixing on the state space $\Psi(T)$, the set of 3-orientations of T.

Theorem A.2. *Let T be a 4-connected planar triangulation with the degree of all internal vertices at most 6. Then the mixing time $\tau(\mathcal{M}_{TR})$ of $\mathcal{M}_{TR}$ satisfies*

$$\tau(\mathcal{M}_{TR}) = O\left(n^8\right).$$

Proof This proof will reply on the fact that $\mathcal{M}_{CR}$ is rapidly mixing which is proved in Theorem 5.
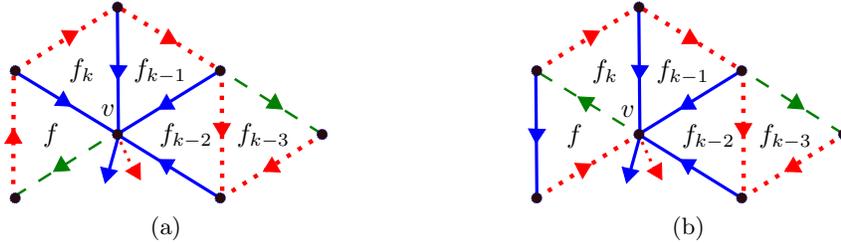
Figure 1: A tower $(f_1, f_2, \ldots, f_k)$ in σ for which $(f_1, f_2, \ldots, f_k, f)$ is not a tower in τ. Notice $\deg(v) \geq 7$.

We use Theorem 9, introduced in Section 4.2, to derive a bound on the mixing time of $\mathcal{M}_{TR}$. To do so we need to bound the constant A given in that theorem. For each edge (σ, τ) in $\mathcal{M}_{TR}$ that takes a counterclockwise cycle f and makes it clockwise, let $\Gamma(\sigma, \tau)$ denote the set of edges (x, y) of $\mathcal{M}_{CR}$ such that the tower $f_1, f_2, \ldots, f_k$ that takes x to y contains the face f. Since a tower is uniquely determined by its direction and its length, $|\Gamma(\sigma, \tau)| \leq 3n^2$. Define $A := \max_{(\sigma, \tau)} (|\Gamma(\sigma, \tau)|) \leq 3n^2$. Since each 3-orientation has the same stationary probability, Lemma A.1b implies that the minimum weight of any state is $\pi_* \geq 3^{-(2n+1)}$. Therefore, by Theorem 9, the mixing time $\tau(\mathcal{M}_{TR})$ of $\mathcal{M}_{TR}$ over 4-connected planar triangulations satisfies

$$\tau(\mathcal{M}_{TR}) \leq 4 \log(4/\pi_*) A \tau(\mathcal{M}_{CR}) = O(n^8). \qquad \square$$

Finally, we can extend this to non-4-connected planar triangulations, where $\mathcal{M}_{TR}$ may select non-facial triangles. Brehm [6] proves that if T has a non-facial triangle C, the edges on its interior that are incident to C must point towards C. This implies that for all $\sigma \in \Psi(T)$ every face on the interior of C that contains an edge of C is not bounded by a directed cycle, so cannot rotate, regardless of the orientation of C. Thus $\mathcal{M}_{TR}$ acts completely independently on the interior of C from the exterior of C, and its mixing time is independent of the mixing time of the exterior as well.

Let T be a planar triangulation and assume that $C_1, C_2, \ldots, C_k$ are all the non-facial triangles of T. Let T_i be the triangulation consisting of all faces contained within C_i and not within any other non-facial

triangle contained within C_i. Let τ_i be the mixing time of $\mathcal{M}_{TR}$ on T_i and let F_i be the number of faces within C_i. Then the mixing time $\tau(\mathcal{M}_{TR})$ of $\mathcal{M}_{TR}$ on T will be

$$\tau(\mathcal{M}_{TR}) = \max_i \left\{ \tau_i \frac{2n+1}{F_i} \right\}. \tag{1}$$

This follows since the Markov chain acts independently on each 4-connected subtriangulation; see [2], e.g., for a proof or the final version of this paper. Using equation 1, it is easy to derive Theorem 2. In fact, this shows that $\mathcal{M}_{TR}$ is rapidly mixing on state space $\Psi(T)$ for any planar triangulation T whose 4-connected triangulations $T_1, T_2, \ldots, T_k$ each have maximum degree (of any internal vertex) 6. Moreover, the same proof can be applied to extend equation 5 to non-4-connected triangulations, proving Theorem 1.

A.1.3 $\mathcal{M}_{TR}$ can mix slowly

(a)   (b)
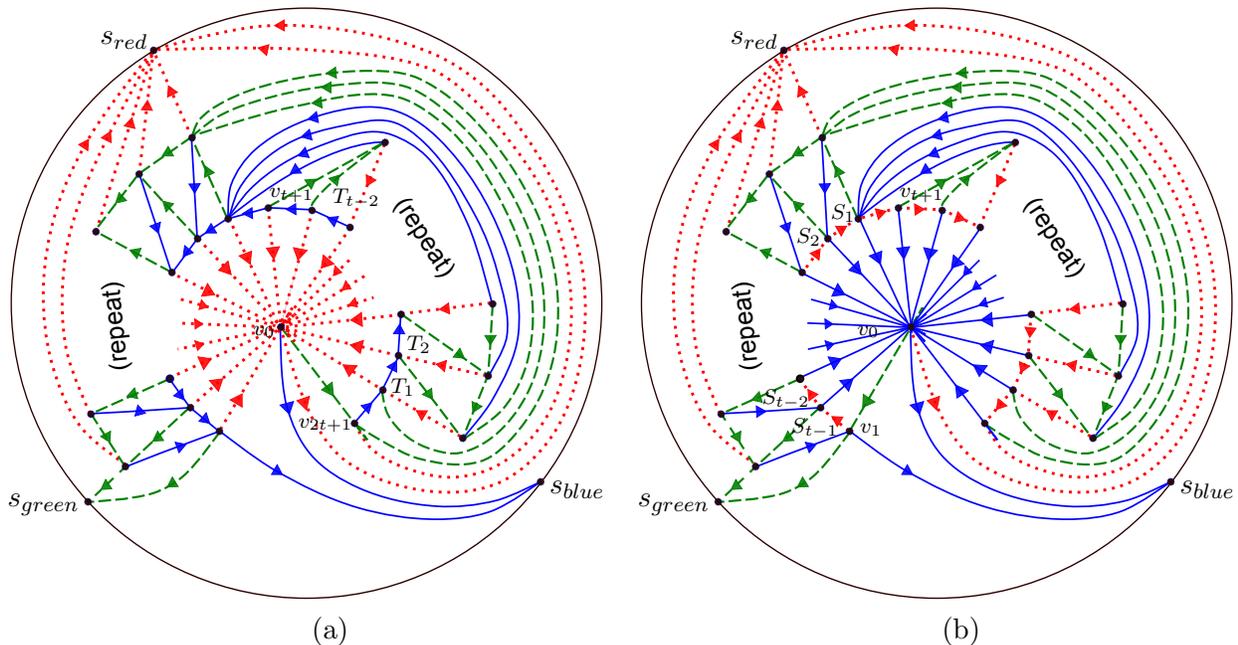
Figure 2: There is an exponential number of 3-orientations with edge (v_0, v_{t+1}) colored blue or red corresponding to the different orientations of triangles $T_1, T_2, \ldots, T_{t-2}$ and $S_1, S_2, \ldots, S_{t-2}$ respectively.

A.2   Mixing Time of $\mathcal{M}_{EF}$

We start by showing that $\mathcal{M}_{EF}$ is an efficient algorithm for sampling from Ψ_n by comparing the moves of $\mathcal{M}_{EF}$ to the moves of $\mathcal{M}_{DK}$, using the comparison theorem, due to Diaconis and Saloff-Coste [9]. We need to define for each transition of $\mathcal{M}_{DK}$ a canonical path using transitions of $\mathcal{M}_{EF}$ and show that each transition is not used by too many canonical paths. We begin with Lemma A.2, which ensures that, for any peak to valley transition (X, Y) of $\mathcal{M}_{DK}$ on the top path, the vertex v_i associated with that move cannot have its green edge point to the green source s_{green}.

Lemma A.2. *Let (X, Y) be a transition of $\mathcal{M}_{DK}$ that moves the i^{th} one on the top path to the right by one position. Then $\overrightarrow{v_i s_{green}}$ is not an edge in X.*

Proof. For the sake of contradiction, suppose $\overrightarrow{v_i s_{green}}$ is an edge in X. Let P be the path from s_{blue} to v_i in the blue tree combined with the edge $\overrightarrow{v_i s_{green}}$. This path partitions the vertices of X into two sets

according to which side of the path they are on. Let S be the set of vertices on the same side of P as v_{i+1} (not including the vertices along P) and let $\overline{S}$ be the remaining vertices. Let $x \in V(X)$ and define x's red edge by $\overrightarrow{xr_x}$. We claim that $L(x) < L(r_x)$; this is a direct result of a bijection between Dyck paths and 3-orientations given in [3] (this is a different bijection from the one given in Section 2). Now follow the blue path from x to s_{blue} until you reach the first vertex z' that has an incoming blue edge from a vertex z such that $L(z) > L(x)$. Define $z_x = z$. We claim that in all 3-orientations with this same blue tree, $L(r_x) \geq L(z_x)$. To prove this, we need some facts from [3]. Specifically, they show that any 3-orientation can be obtained by starting from a fixed 3-orientation, called a *star realizer*, which is unique to each blue tree, and then applying a series of transformations. The star realizer satisfies the property that $L(r_x) \geq L(z_x)$ for any vertex x, and each transformation maintains this property.

This implies that any red edge in X directed toward vertices in S must be directed toward vertices in S for all 3-orientations with the same blue tree, which implies that there is no configuration Y which has the same blue tree as X and the same incoming red degree sequence except for increasing v_i by one and decreasing v_{i+1} by one. This is a contradiction since we have assumed (X, Y) is a valid move. Hence we have shown that g_i is not s_{green}.

$\square$

We are now ready to prove Theorem 4.

Proof of Theorem 4. In order to apply Theorem 9 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 $\mathcal{M}_{EF}$ and show that each transition is not used by too many canonical paths. We consider four cases, depending on whether the move of $\mathcal{M}_{DK}$ takes a peak to a valley and which Dyck path it involves. If the transition $e = (X, Y)$ of $\mathcal{M}_{DK}$ involves both the top and bottom paths, then our path will do each move separately; that is, if e is a peak to valley move, then we first make the peak to valley move on the bottom path, and then make the peak to valley move on the top path (and vice versa for a valley to peak move). In order to bound the number of paths using a transition (Z, W) of $\mathcal{M}_{EF}$, we need a bit to record whether or not the transition (X, Y) affects both paths. Notice that if e does not involve the bottom path, then the move affects only the red and green trees, and leaves the blue tree unaltered.

A Valley to Peak Move on the Top Dyck Path: Next consider the case where $e = (X, Y)$ changes a valley to a peak. Assume that e moves the ith 1 on the top Dyck path to the left one position. From the bijection we know that this move corresponds to, in the red tree, decreasing the incoming degree of v_i by one and increasing the incoming degree of v_{i+1} by one. If v_i and v_{i+1} are physically adjacent in the blue tree then there is a green/red swap centered at v_i and involving v_i's green edge that performs exactly these changes so we will define γ_{XY} to be this single move. If v_i and v_{i+1} are not adjacent then we will define two stages in γ_{XY}. Let r_i be v_i's parent in the red tree. In the first stage we transfer an incoming red edge from v_i to r_i using the green/red swap centered at v_i and involving v_i's green edge. In the second stage we will repeatedly use the canonical path defined for peak to valley moves to transfer an incoming red edge from r_i to v_{i+1}.

Again given a transition (Z, W) of $\mathcal{M}_{EF}$ we must upper bound the number of canonical paths γ_{XY} that use this edge in this case. First, we will record a bit to determine which stage we are in. For both stages we will record the vertex v_i and the vertex r_i. If v_i and v_{i+1} are adjacent we will record v_{i+1} instead of r_i. For the second stage if we are in the process of making a peak to valley move which affects the vertices v_p and v_{p+1}, we will record the vertex v_p. In addition we will record the extra vertex v_j that is needed when we are in the middle of a peak to valley move, as explained in the peak to valley case above. Given this information we can uniquely recover X and Y. This implies that in this case there are at most $2n^4$ canonical paths which use any edge (Z, W).

A Valley to Peak Move on the Bottom Dyck Path: Let $e = (X, Y)$ be a transition of the tower chain $\mathcal{M}_{CR}$ on the bottom Dyck path, and suppose the edge e is a move that takes a valley of the bottom Dyck path and moves it to a peak. This affects the blue tree as follows (see Figure 3): a's blue
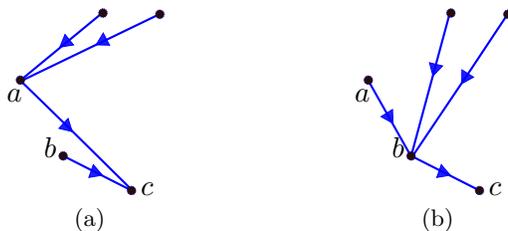
Figure 3: A move of $\mathcal{M}_{DK}$ that moves the bottom Dyck path from a valley to a peak takes the blue tree from (a) to (b).

edge moves from c to b, and all the (blue) children of $\overrightarrow{ac}$ (if any exist) become children of $\overrightarrow{bc}$. However, to define the canonical path between these two configurations, it is necessary to also know what the top Dyck path looks like, as it determines the red (and therefore the green) tree. Figure 4 shows how the red and green tree might look. Our path will first update the blue tree from X to match the blue tree of Y, and then update the red tree (and therefore the green tree) to match the red tree of Y using the steps outlined in the previous two cases. For a vertex $v \in T$ let r_v denote the head of v's red edge in X. We will go through 4 distinct stages in the canonical path. In stage 1 the blue edge of a moves from c to b. Then in stage 2, a's red edge moves into position for stage 3, where all incoming blue edges to a move down to point to b. Finally, in stage 4 we repair the red tree.

Stage 1. Given the vertex condition (Figure 1a) for b and the bijection between the bottom Dyck path and the blue tree, there is no edge in the angle $\angle acb$ and thus a, b, and c form a triangle in our triangulation. Vertex b may have some green edges coming in between its blue edge $\overrightarrow{bc}$ and its red edge $\overrightarrow{br_b}$. If so, then $\overrightarrow{ab}$ is an edge, as in Figure 4a. The first step along the canonical path from X to Y is to rotate b's red edge to point to a (see Figure 4(a-b)). This is accomplished through a sequence of red/green swaps, one for each green edge coming into b, as in Figure 6. Now we are ready to move the blue edge $\overrightarrow{ac}$ to $\overrightarrow{ab}$ by swapping with b's red edge, as in Figure 4(b-c).
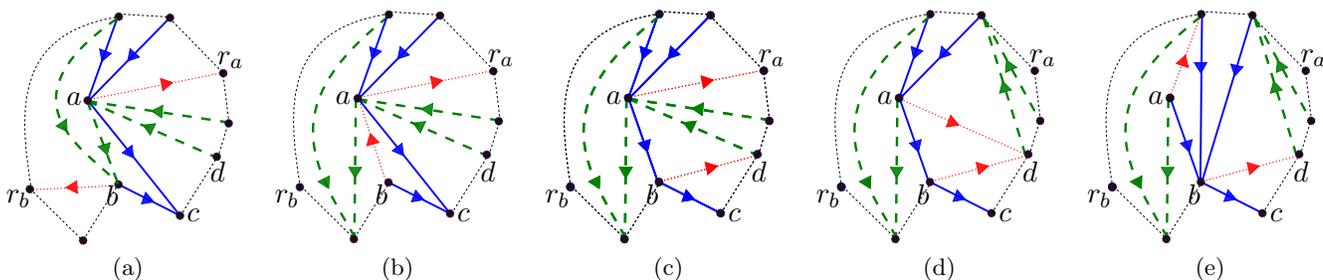
Figure 4: Canonical path to move a valley up to a peak in the blue tree.

Stage 2. Consider vertex a. Counterclockwise from a's blue edge, there may be some green edges coming into a, followed by the red edge $\overrightarrow{ar_a}$. If there are green edges, let d be the tail of the last incoming green edge to a; however, if there are no green edges, then $r_a = d$. In this stage of our canonical path, a's red edge moves from r_a to d (if $r_a \neq d$) by a series of red/green swaps. See Figure 6(c-d).

Stage 3. Next, we use a's red edge to move all the blue children of $\overrightarrow{ab}$ to point to b, one at a time, in a clockwise manner. See Figure 6(d-e). Now the blue tree is completely fixed.

Stage 4. Finally, we must repair the red tree. Notice that the red edges of a and b are the only red edges that we moved; we must move them to their proper place. We must increase the degree of r_b and r_a back to match their indegree in X. To do this without affecting the blue tree, we first make a's red edge point to r_b (we'll call this **Stage 4a**) and then make b's red edge point to r_a (**Stage 4b**). These moves can each be accomplished by a sequence of red/green swaps without affecting the blue tree.

Given a transition (Z, W) of $\mathcal{M}_{EF}$, we must upper bound the number of canonical paths $\gamma_{X,Y}$ that use this edge. If (Z, W) is in stage 1, we need to remember vertices a and r_b, and a bit to tell us whether

19

or not we have moved $\overrightarrow{ac}$ yet. Given this information, we can recover b and c and can undo all red/green swaps in order to get back to X. Given X we can find Y since we know which valley to flip up. If (Z, W) is in stage 2, then we only need to record r_b, since (Z, W) moves a's red edge, so we know a. To get back to the last configuration in stage 1, we just need to move a's red edge counter-clockwise until it can't make any more red/green swaps. Thus we can get back to the last configuration in stage 1, and using a and r_b we can recover X. If (Z, W) is in stage 3, we need to record r_b. Each move in stage 3 takes a child of $\overrightarrow{ab}$ and moves it to point to b. Hence, we know a. Notice that since $\triangle abc$ was facial in X, all blue edges coming into b in σ_1 before $\overrightarrow{ab}$ (in the counterclockwise direction) were children of $\overrightarrow{ac}$ in X. Thus, given a, we know that we must use a's red edge to move each of these children back up to a. This brings us back to the last configuration in stage 2; using a and r_b, we can recover X. If (Z, W) is in stage 4a, then we know that the blue tree agrees with Y, and the red edges of a and b are the only red edges in a different position in Z than in Y. Since (Z, W) moves a's red edge, we know a, and b is the vertex that shares a's blue edge. Given b, it is easy to recover r_a, since to find r_a, just move b's red edge counter-clockwise until it can't make any more red/green swaps. We need to record r_b and then it is easy to get to Y. If (Z, W) is in stage 4b, then we know that the blue tree agrees with Y and b's red edge is the only red edge that is in a different position in Z than in Y. Since (Z, W) moves b's red edge, we know b. As described in stage 4a, it is easy to then find r_a; move b's red edge all the way there to get to Y. In each of the four stages we need to record a maximum of 2 vertices and a single bit. This implies that in this case there are $O(n^2)$ canonical paths which use any edge (Z, W).

A Peak to Valley Move on the Bottom Dyck Path: Next suppose $e = (X, Y)$ is a move that takes a peak of the bottom Dyck path and moves it to a valley. This moves the blue tree in Figure 3b to Figure 3a. In general, a's blue edge moves from b to c, and all the (blue) children of $\overrightarrow{bc}$ clockwise from $\overrightarrow{ab}$ become children of $\overrightarrow{ac}$. It is important to note that in order for the peak to valley move to be possible, a must be a leaf of the blue tree. Figure 5 shows how the red and green trees might look. Using the vertex condition (Figure 1a) for b, going clockwise around b there are $k_1 \geq 0$ blue edges coming into b after $\overrightarrow{ab}$, followed by b's red edge, then there are $k_2 \geq 0$ green edges coming in, and finally b has its blue edge $\overrightarrow{bc}$. As above, r_v is the head of v's red edge. We will go through 5 stages in the canonical path.
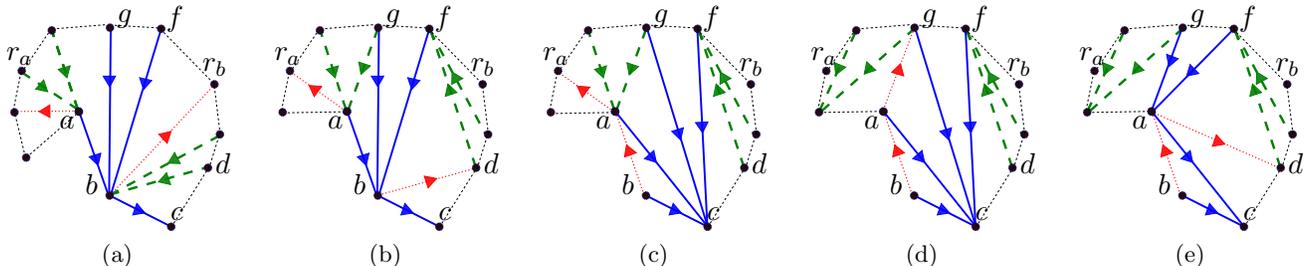
Figure 5: Canonical path to move a peak down to a valley in the blue tree.

Stage 1. If there are green edges coming into b, let d be the tail of the last incoming green edge; if b has no green edges coming in, then set $d = r_b$. During the first stage of the canonical path, b's red edge must move from r_b to d (if $r_b \neq d$) by making k_2 red/green swaps.

Stage 2. Let f be the last blue incoming edge to b. Now we can move $k_1 + 1$ blue edges pointing to b (namely, all blue edges between $\overrightarrow{ab}$ and $\overrightarrow{fb}$, inclusive) to point to c. Do this sequentially and counterclockwise, using b's red edge to swap. See Figure 5(b-c).

Stage 3. Vertex a may have green edges entering between the edges $\overrightarrow{ar_a}$ and $\overrightarrow{ac}$. Let g be the tail of a's last incoming green edge, if it exists; otherwise, $g = r_a$. If $g \neq r_a$, we move a's red edge from r_a to g.

Stage 4. Using a's red edge, we can move the k_1 blue edges from c back to a, correcting the blue tree.

Stage 5. To repair the red tree, first we need to move a's red edge to point to r_b (**Stage 5a**). Notice that b's red edge now points toward a. Thus we need to pass a red edge from a to r_a (**Stage 5b**). Since a is a leaf of the blue tree, this can be done using a sequence of red/green swaps.

Given a transition (Z, W) of $\mathcal{M}_{EF}$ we must upper bound the number of canonical paths $\gamma_{X,Y}$ that may use this edge. If (Z, W) is in stage 1, then it moves b's red edge, so we know b. To get back to X, move b's red edge counter-clockwise until it can't make any more red/green swaps. Hence we don't need to remember any vertices. Given X and b, we can find Y since we know which peak to flip down. If (Z, W) is in stage 2, then it moves some blue edge from b to c, so we know both b and c. We must record f, but then we can get back to the last configuration in stage 1, and using b we can get back to X. If (Z, W) is in stage 3, then it moves a's red edge, so we know a. We need to remember r_a to get back to stage 2, and remember f to get back to stage 1. Given a we know b, so we can then get back to X. If (Z, W) is in stage 4, then it moves a's red edge, so we know a. To get back to stage 3, just move a's red edge as far counter-clockwise as possible. Hence to get back to X we need to record r_a and f. If (Z, W) is in stage 5a, then we are moving a's red edge, so we know a. Rotate it as far counter-clockwise as possible to find r_b. To recover Y, we just need to know r_a so we can pass an edge from a to r_a. If (Z, W) is in stage 5b, then we need to remember a and r_a. Move a's red edge all the way to r_a to obtain Y. In each of the five stages we need to record a maximum of 2 vertices. This implies that in this case there are $O(n^2)$ canonical paths which use any edge (Z, W).

We have shown that in each of the four cases above there is a maximum of $O(n^4)$ canonical paths which use any edge (Z, W). If the move of $\mathcal{M}_{DK}$ affects both the top and bottom paths, we can think of this move as two moves, each of which affects only the top or bottom path; hence, we concatenate the paths for each of those moves. Therefore, if we record a bit to decide if the move of $\mathcal{M}_{DK}$ affects both the top and bottom paths, as well as a bit to decide which of the stages we are in, this implies that across all cases there is a maximum of $O(n^4)$ canonical paths which use any edge (Z, W). Notice that the maximum length of any path γ_{XY} is $O(n^2)$. We can now upper bound the quantity A which is needed to apply Theorem 9 as follows:

$$
\begin{aligned}
A &= \max_{(Z,W) \in E(P)} \left\{ \frac{1}{\pi(Z)P(Z,W)} \sum_{\Gamma(Z,W)} |\gamma_{XY}| \pi(X) P'(X, Y) \right\} \\
&\leq \max_{(Z,W) \in E(P)} \left\{ 4n \sum_{\Gamma(Z,W)} \frac{O(n^2)}{2n} \right\} \leq O(n^6).
\end{aligned}
$$

Moreover, we can bound π_* as follows:

$$
\pi_* = \min_{X \in \Psi} \pi(X) = \frac{1}{C_{n+2}C_n - C_{n+1}^2} \geq \frac{1}{30^n}.
$$

Applying Theorems 8 and 9 , proves that $\tau(\mathcal{M}_{EF}) = O(n^{10} \log n)$. Therefore $\mathcal{M}_{EF}$ is an efficient sampling algorithm for sampling from the set of all 3-orientations over any triangulation on n internal vertices.