A MARKOV CHAIN SAMPLER FOR PLANE CURVES

HARRISON CHAPMAN AND ANDREW RECHNITZER

ABSTRACT. A plane curve is a knot diagram in which each crossing is replaced by a 4-valent vertex, and so are dual to a subset of planar quadrangulations. The aim of this paper is to introduce a new tool for sampling diagrams via sampling of plane curves. At present the most efficient method for sampling diagrams is rejection sampling, however that method is inefficient at even modest sizes. We introduce Markov chains that sample from the space of plane curves using local moves based on Reidemeister moves. By then mapping vertices on those curves to crossings we produce random knot diagrams. Combining this chain with flat histogram methods we achieve an efficient sampler of plane curves and knot diagrams. By analysing data from this chain we are able to estimate the number of knot diagrams of a given size and also compute knotting probabilities and so investigate their asymptotic behaviour.

1. INTRODUCTION

1.1. **Background.** A beautiful aspect of knot theory is that it brings together many different aspects of mathematics and leverages their tools so efficiently. An apt example of this phenomenon are knot diagrams, wherein the entanglement of a smooth string in space is studied as a combinatorial cartoon. From this view, any of the infinite manipulations of a loop of string are simply sequences of a finite number of diagram operations called Reidemeister moves [AB26; Rei48]. This combinatorial regime also serves an essential role in knot identification; many knot polynomials are calculated more naturally for diagrams than for space curves [EM97]. Other invariants, like the important crossing number of a knot, are defined in terms of diagrams.



FIGURE 1. By replacing crossings with vertices, knot diagrams (A) are mapped onto plane curves (B). However, not all 4-valent plane graphs correspond to knots. For example, the Hopf-link (C) maps to a 4-valent graph. To avoid complications caused by symmetries, we study rooted plane curves, in which one edge is selected and assigned an orientation (D).

Knot theory is an important tool in many applied disciplines. Ring polymers exhibit knotting [FW61; Tri+01], which affects their function [BZ04] and their chemical properties [Van95]. Knotted configurations of polymers are studied through *random knot models*,

Date: April 11, 2018.

where knots are sampled from some probability distribution with the aim of modeling physical behavior; for a good overview see [OW07]. A canonical model of random knotting is that of self-avoiding polygons on the cubic lattice \mathbb{Z}^3 [SW88]. A self-avoiding polygon is constructed by embedding a simple closed curve into \mathbb{Z}^3 and by sampling these objects uniformly at random, one obtains a distribution over the space of knot embeddings. A surprising amount is still unknown for all but the most simple of such models. For example, it is still unproven (despite overwhelming evidence) that the exponential growth rates of knotted polygons of fixed knot type are *independent* of knot type [RR11].

As knot diagrams are more naturally suited for the study of invariants, in [CCM16], the first author together with Cantarella and Mastin proposed the random diagram model. In this model, one samples a random knot by, for a fixed number of crossings n, picking one of the finite number of knot diagrams with n crossings uniformly. The diagram model behaves like many other physically motivated models in that unkotted diagrams are exponentially rare [Cha17; Del62; Dia95; FW61; SW88]. At the same time, this diagram model has advantages over other models of knotting. It is possible, for example, to show, via a pattern theorem, that unknot diagrams almost certainly contain slipknots [Cha18], a result which remains a conjecture for unknotted self-avoiding polygons and Gaussian polygons [Mil10].

Another major problem of combinatorial models of knots, is that the underlying objects are difficult to enumerate. The diagram model is no different — the best algorithms for enumerating knot diagrams require time that grows exponentially with the number of crossings [ZZ09]. If an efficient enumeration method were to exist, then this could be readily adapted to give a random sampling method (see, for example, [FS09]). In the absence of such a method, the most obvious approach has been to generate random 4-valent maps and then randomly assign crossings (see Figure 1). However, 4-valent maps corresponding to knots are exponentially rare. Experiments suggest that roughly only 1% of samples are accepted for 60-crossing diagrams and a prohibitively small 0.01% of samples are accepted for 150-crossing diagrams [Cha17]. This makes rejection sampling of 4-valent maps extremely inefficient at even moderate sizes. One way to overcome this is to manipulate a sampled 4-valent map until a knot diagram is obtained [DEZ12; Dun+14], however the resulting space of knots is not uniform.

The absence of an efficient random knot diagram sampler is a major impediment to the study of random knots. The aim of this paper is thus to describe a new efficient method to sample knot diagrams *directly* with uniform probability using Metropolis style Markov chain Monte Carlo (MCMC) sampling [Met+53]. Such Markov chains have been used to study other models of random knotting, particularly on the simple cubic lattice. Foremost among these are the pivot algorithm [Lal69] and the BFACF algorithm [AC83; BF81; CCF83]. The former is extremely efficient [MS88], while the latter has the advantage of conserving topology [RW91b]. Guitter and Orlandini [GO99] augmented the BFACF algorithm with Reidemeister moves to study a model of flat-knots on the square-diagonal lattice.

This paper focuses on sampling *plane curves*, which can be thought of as knot diagrams without crossing-sign information (see Figure 1 and Section 2). Each plane curve of *n*-vertices maps to a unique set of 2^n knot diagrams. Consequently if we can sample plane curves uniformly, then we can also sample knot diagrams uniformly. To avoid complications caused by symmetries, we sample from the space of rooted diagrams (explained below).

The main result of the paper is a Markov Chain over the space of plane curves. By selecting transition probabilities we can sample from this chain with a Boltzmann distribution.

Theorem 1. Let D be a plane curve with $1 \le n$ vertices and let μ is the exponential growth rate of plane curves. The Markov chain described in Section 3.1 has stationary distribution given by

 $\pi(D) \propto z^n$,

provided $0 \le z < \mu^{-1}$. Consequently, plane curves of a fixed size are sampled uniformly.

When running this Markov chain, we found that it was very difficult to obtain a good number of samples at large range of lengths. To overcome this problem we modified our transition probabilities based on Wang Landau density of states estimation [WL01]. This method allows us to sample diagrams nearly uniformly in length (while still sampling uniformly within any given length) and additionally provides estimates of the number of plane diagrams. Let k_{ℓ} be the number of rooted plane curves with ℓ vertices, and let g_{ℓ} be the estimate of k_{ℓ} from the Wang-Landau algorithm. Then we have the following result.

Theorem 2. Let $N \in \mathbb{N}$ and D be a plane curve with $1 \leq n \leq L$ vertices. The Markov chain described in Section 3.2 has stationary distribution given by

$$\pi(D) \propto \frac{1}{g_n}$$

Since $g_n \approx k_n$, plane curves are sampled uniformly within a given size, and approximately uniformly across sizes.

The remainder of Section 1 defines key knot theory concepts and provides some additional context on the problem of enumerating plane curves. Section 2.2 describes the transitions used by our Markov chain. In Section 3, we describe the actual algorithms and prove the main theorem. Section 4 presents the results of experiments which verify the validity of the main theorem as well as explore the structure of large random knot diagrams. Finally, in the concluding Section 5 we discuss progress on additional "diagram Markov chains" for different types of diagram objects, and some obstructions.

2. Preliminaries and Definitions

2.1. **Definitions.** A *knot* is an embedding $K : S^1 \hookrightarrow \mathbb{R}^3$ of a loop into Euclidean 3-space. Typically, knots are considered up to ambient isotopy, wherein two knots are equivalent if one can be manipulated as a closed loop into the other, without self-intersection. For clarity, we call a specific loop embedding a *knot* and an equivalence class of knots a *knot type*. Reidemeister's theorem [AB26; Rei48] transfers this topological theory into a combinatorial one as follows: A *knot diagram* of a knot K is a generic projection of the loop in space to the

FIGURE 2. The three Reidemeister moves, RI, RII, RIII.

sphere, together with extra information at each double point indicating where one piece of the loop passes over the other (called *crossings*), as in Figure 1A, up to oriented homeomorphisms of the sphere. Then two knots K_1 and K_2 are equivalent if and only if their *diagrams* are related by a sequence of Reidemeister moves, depicted in Figure 2. A significant advantage

of the diagram view is that most *knot invariants*, properties of knots which only depend on their knot type, are naturally computed from their diagram representation [Fre+85; Kau87].

There is a natural projection from knot diagrams to a strict subset of 4-valent planar maps called plane curves; simply replace each crossing with a vertex. A *planar map* is a (multi-)graph G, together with an embedding ι into the sphere S^2 so that each component of $S^2 \setminus \iota(G)$ is a topological disk. Necessarily, this means all planar maps are connected. A map is 4-valent if each vertex has degree 4.

Unfortunately, knot diagrams and plane curves are cumbersome to deal with as a result of potential symmetries. To avoid these complications can asymmetrize by marking one edge with a direction. Such objects are called *rooted*. In particular, each rooted *n*-vertex plane curve corresponds to a unique set of exactly 2^n rooted knot diagrams. For a discussion of the techniques and difficulties involved with considering plane curves and knot diagrams with symmetry see [CCM16; CZ16; Val16].

It will be useful for computations to consider the following equivalent view of maps. A 4-valent planar map D with n vertices can be viewed as a *combinatorial map* [Cha17; CZ16], *i.e.* a pair $D = (\sigma, \tau)$ of permutations of 4n flags (sometimes called half-edges or arcs). In this view, σ is a product of n disjoint cycles of length 4 and τ is a product of 2n disjoint cycles of length 2. Each cycle in σ represents a vertex (it permutes the flags attached at each vertex counterclockwise) and each cycle in τ represents an edge (it involutes the two flags that form an edge). The cycles of $\sigma\tau$ correspond to the faces of the map (each permutes the flags of a face clockwise). For 4-valent maps we also have $\sigma^2\tau$, whose cycles correspond to orientations of *link components* or Gauss components. Each flag is contained in exactly one vertex, edge, face and component.

Given an flag a, let e(a) be its edge (*i.e.* cycle in τ), v(a) be its vertex (*i.e.* cycle in σ), and f(a) be its face (*i.e.* cycle in $\sigma\tau$).



FIGURE 3. A plane curve, with its flags marked and a root flag chosen (in blue).

Then the condition that D is a planar map is precisely that product $\sigma\tau$ consists of n+2 cycles by Euler's formula. Relaxing this condition and allowing $\sigma\tau$ to consist of k cycles makes D a map on a surface of genus $g = 1 - \frac{k-n}{2}$. We forbid this for our objects, although in general it is interesting to consider maps on an arbitrary fixed surface Σ . If $\sigma^2\tau$ consists of precisely two cycles, each necessarily of length n, then D is a plane curve and each cycle in $\sigma^2\tau$ corresponds to following the single immersed circle in one of its two possible orientations. Figure 3 shows an example plane curve. Figure 4 shows a random decorated plane curve (knot diagram) and a random decorated 4-valent map (link diagram). The curve condition (that $\sigma^2\tau$ has precisely 2 cycles) makes plane curves exponentially rare within the class of all 4-valent maps.

We note that if we relax either the planarity or the curve condition, the problem greatly simplifies. 4-valent planar maps themselves are well-understood, owing in part to Schaeffer's bijection with blossom trees [Sch97]. This has been used to prove a stunning closed formula

for the growth rate of alternating link types [ST98; Thi98; ZZ02] as well as precise statistics for hyperbolic volumes of random alternating link diagrams [Obe16]. On the other hand, relaxing planarity and considering curves on arbitrary surfaces leads one to the study of *Gauss codes* (usually depicted as signed chord or *Gauss diagrams*), which themselves are counted and well understood [Now09].

Denote by \mathscr{K} the class of all rooted plane curves indexed by number of vertices, and let $K(t) = \sum_{n=1}^{\infty} k_n t^n$ be its generating function. For a plane curve D, let |D| be the number of vertices in D (equivalently, the size of D). Then the generating function can also be written as $K(t) = \sum_{D \in \mathscr{K}} t^{|D|}$. The asymptotic behavior of the coefficients k_n is expected [SZ04; ZZ09] to be

(1)
$$k_n \sim C \mu^n n^{\gamma-2} (1 + O(1/\log n)).$$

Neither a closed formula for k_n nor exact values of μ, γ are known. Conformal field theory arguments suggest [SZ04] that

$$\gamma = -\frac{1+\sqrt{13}}{6}.$$

Additionally, it is known that μ exists [Cha17], with the best numerical estimate [ZZ09] $\mu \approx 11.416 \pm 0.005$. This sort of asymptotic growth is similar to that of self-avoiding walks and polygons in the cubic lattice \mathbb{Z}^3 [Ham61; MS13] (with different constants). Indeed a great many combinatorial objects are known to be counted by sequences which have similar exponential growth with power-law correction; see Flajolet and Sedgewick [FS09] for (a great many) examples. It should be noted, however, that conformal field theory arguments [SZ04] suggest the presence of an inverse logarithmic correction to scaling for plane curves. This is contrast with the observed correction of $n^{-\Delta}$ for many objects (again, see Flajolet and Sedgewick [FS09] for many examples with n^{-1} corrections, and also [CG96] for evidence of a $n^{-3/2}$ correction in self-avoiding walks).

2.2. Shadow Reidemeister moves. In the subsection above, we have defined the set of rooted plane curves that we wish to sample. Unfortunately, as noted above, it is difficult to construct a rooted plane curve. Instead, we will describe a Markov chain that produces new plane curves by performing small local changes. The set of plane curves is closed under these manipulations. Further, any two plane curves are linked by a sequence of these changes.

Reidemeister moves in Figure 2 are an obvious choice of local changes for knot diagrams, and we use a similar set of moves for plane curves. Define the flat (or shadow) Reidemeister moves (also known as homotopy moves [CE17]) to be the same as the Reidemeister moves, except ignoring crossing information [HN10]. Any two plane curves are related by a sequence of flat Reidemeister moves, as;

Theorem 3 (Hass and Scott [HS94], de Graaf and Schrijver [GS97]). Any plane curve can be brought to the trivial figure-eight twist curve seen in Figure 6 by a sequence of flat Reidemeister moves that never increase the size of the plane curve.

This result implies that there exists a sequence of moves between any two plane curves D, N. Furthermore, it implies that at each intermediate state between D and N the curve has no more vertices than the larger of D and N. Chang and Erickson have proven in the arbitrary case that the maximum number of moves required is a small polynomial:

Theorem 4 ([CE17]). The maximum number of non-increasing flat Reidemeister moves required to trivialize a plane curve with n vertices grows as $\Theta(n^{3/2})$. Consequently, it takes



FIGURE 4. A random knot diagram and a random link diagram, each of 100 vertices. Different link components are given different colors. These diagrams were sampled uniformly (using a rejection sampler in the case of the knot diagram) using an interface in plCurve [ACC17] to PlanarMap [Sch99b], and graphics were generated using an orthogonal projection algorithm in pLink, part of SnapPy [Cul+17]. Knot diagrams become exponentially rare as the number of vertices increases [Cha17; SZ04], so are difficult to sample through rejection.



FIGURE 5. The three flat Reidemeister moves, RI, RII, RIII which act on plane curves. These are the natural analogues of Reidemeister moves on knot diagrams.



FIGURE 6. The figure-eight twist curve is the "trivial" plane curve for the sake of implementation; allowing for a trivial curve of 0 vertices is more difficult and provides no benefit. By Theorem 3 all other plane curves can be brought to this shape by a sequence of non-crossing-increasing flat Reidemeister moves.

no more than $\Theta(n^{3/2})$ moves to transform a curve of $m \leq n$ vertices into a curve of n vertices.

A similar pair of results, if flat Reidemeister I moves are forbidden, was proven by Nowik [Now09]. In this case, the path of curves between a curve with n vertices and its trivialization takes no more than $\Theta(n^2)$ moves and involves no intermediate curves with more than n + 2 vertices.

These flat Reidemeister moves form the transitions for our Markov chain on \mathscr{K} . A key property is that they are all reversible. We define each transition in detail, noting their

unique inverses, since such details are necessary for efficient implementation. Let D be a rooted plane curve with root flag a.



FIGURE 7. The flat Reidemeister I moves

2.2.1. Shadow Reidemeister I loop addition, RI^+ . See Figure 7 (left) and consider the edge to be oriented from right to left. This defines two flags, a_1 and a_2 , and we take a_1 to be the root flag. The addition of a loop on a_1 is always possible through the move RI^+ .

Let $(a_1a_2) = e(a_1)$, and let a_3, a_4, a_5, a_6 be four new flags. Then $\mathrm{RI}^+(D, a_1)$ is the rooted map produced by deleting edge $e(a_1)$ from D, then adding the vertex $(a_3a_4a_5a_6)$ and edges $(a_1a_4), (a_5a_6)$, and (a_3a_2) . The new root is the flag a_6 . One can verify that this process is invertible, in particular that:

$$D = \mathrm{RI}^{-}(\mathrm{RI}^{+}(D, a_1), a_6).$$

2.2.2. Shadow Reidemeister I loop deletion, RI⁻. See Figure 7 (right). Loop deletion RI⁻ is possible whenever, given the root flag $a = a_6$, the face $f(a_6)$ is a singleton (*i.e.* the corresponding face is a monogon).

If this is true, then $\operatorname{RI}^-(D, a_6)$ is the rooted map produced from deleting the vertex $v(a_6) = (a_3a_4a_5a_6)$ and the edges $e(a_6) = (a_6a_5)$, $e(a_3) = (a_3a_2)$, $e(a_4) = (a_4a_1)$ from D, and adding the edge (a_1a_2) . The flags a_3, a_4, a_5, a_6 are discarded. The new root is the flag a_1 . This process is invertible:

$$D = \mathrm{RI}^+(\mathrm{RI}^-(D, a_6), a_1)$$



FIGURE 8. The flat Reidemeister II moves

2.2.3. Shadow Reidemeister II bigon addition, RII⁺. See Figure 8 (left). Bigon addition RII⁺ requires a second flag $b \in f(a)$ which is different from the root flag a. Let $a = a_1$ and $b = a_3$.

The rooted map RII⁺(D, a_1 , a_3) is constructed from D as follows. Delete the edges $e(a_1) = (a_1a_2)$ and $e(a_3) = (a_3a_4)$ from D. Add eight new flags a_5 , a_6 , a_7 , a_8 , a_9 , a_{10} , a_{11} , a_{12} . Add the vertices $(a_6a_{10}a_9a_7)$ and $(a_5a_8a_{12}a_{11})$. Insert the edges (a_2a_9) , (a_3a_{10}) , (a_7a_8) , (a_5a_6) , (a_1a_{12}) , and (a_4a_{11}) . The new root is the flag a_6 . This process is invertible:

$$D = \operatorname{RII}^{-}(\operatorname{RII}^{+}(D, a_1, a_3), a_6).$$

2.2.4. Shadow Reidemeister II bigon deletion, RII⁻. See Figure 8 (right). Bigon deletion RII⁻ is possible provided the root flag $a = a_6$ is on a face $f(a_6)$ which consists of precisely two flags (*i.e.* a bigon).

Additionally, it is required that the two exterior faces which are merged by the transition be *distinct*; this is required to preserve connectedness and genus. Indeed, Suppose that D is a curve embedded on an orientable surface of genus g, so that v(D) - e(D) + f(D) = 2(g-1), but that the faces to be merged are not distinct. The number of vertices and edges decrease by 2 and 4 respectively by a RII⁻ operation, as usual, but the number of faces now remains fixed. This implies that the produced curve lives in either a surface of one fewer genus, or if the original map was embedded on the sphere, two disjoint spheres.

The rooted map RII⁻(D, a_6) is constructed from D as follows. Delete the vertex $v(a_6) = (a_6a_{10}a_9a_7)$ and the edges $e(a_9) = (a_2a_9)$, $e(a_{10}) = (a_3a_{10})$, $e(a_7) = (a_7a_8)$ and $e(a_6) = (a_6a_5)$. Delete the vertex $v(a_5) = (a_5a_8a_{12}a_{11})$ and the edges $e(a_{12}) = (a_1a_{12})$ and $e(a_{11}) = (a_4a_{11})$. Add the edges (a_1a_2) and (a_3a_4) . The flags $a_5, a_6, a_7, a_8, a_9, a_{10}, a_{11}, a_{12}$ are discarded. The new root is the flag a_1 . This process is invertible:

$$D = \text{RII}^+(\text{RII}^-(D, a_6), a_1, a_3).$$



FIGURE 9. The flat Reidemeister III move

2.2.5. Shadow Reidemeister III triangle flip, RIII. See Figure 9. Triangle flipping RIII is possible provided the root flag $a = a_8$ lies on a face $f(a_8)$ with precisely three flags, all of whom are contained in *different* vertices (*i.e.* a nondegenerate triangle).

The rooted map RIII(D, a_8) is constructed from D as follows. Say that $v(a_8) = (a_1a_2a_7a_8)$, $e(a_8) = (a_8a_9), e(a_7) = (a_6a_7), v(a_6) = (a_5a_{11}a_{12}a_6)$, and $v(a_9) = (a_3a_9a_{10}a_4)$. Delete vertices $v(a_7), v(a_6), v(a_9)$ and edges $e(a_7), e(a_8)$, and $e(a_{10})$. Insert vertices $(a_6a_8a_{11}a_1), (a_2a_{10}a_7a_3), (a_4a_{12}a_9a_5)$ and edges $(a_6a_9), (a_7a_{12}), and (a_{10}a_{11})$. The new root is the flag a_{11} . This process is invertible:

$$D = \operatorname{RIII}(\operatorname{RIII}(D, a_8), a_{11}).$$

3. MARKOV CHAIN

The five flat Reidemeister moves described in the previous section allow us to define a Markov chain on the space plane curves. That the chain can move between any two given plane curves follows immediately from Theorem 3.

3.1. A Boltzmann Markov chain on plane curves. Our Markov chain sampler for plane curves will have a stationary *Boltzmann distribution*; one which samples curves of different sizes with different probabilities but is *uniform on curves of fixed size*. In other words, for $0 \le z < \mu^{-1}$ an arbitrary parameter, it has stationary distribution such that a curve D has probability

(2)
$$\pi(D) \propto z^{|D|}$$

The parameter z then is then maybe used to control mean size of sampled curves, and smaller values of z will prevent the samples from growing infinitely large.

Let $p_1, p_2, p_3 > 0$. These numbers correspond to the probabilities of performing, respectively, a Reidemeister I, II, or III move. Consequently, we must have $p_i > 0$. Further, we must also have $1 - (p_1 + p_2 + p_3) > 0$, as this quantity will correspond to the probability of selecting a new root flag. Let D_i be the input rooted plane curve with n vertices and root flag a and perform one of the following six subprocedures with different probabilities. If a move fails then set $D_{i+1} := D_i$.

- [0] Re-rooting, with probability $1 (p_1 + p_2 + p_3)$. Given the rooted diagram D_i , forget the root and select a new root b for D_{i+1} from the $(\operatorname{aut} D_i)/(4|D|)$ choices. The probability that this transition succeeds is $1 - (p_1 + p_2 + p_3)$. Note that $\operatorname{aut} D_i$ needn't be calculated for this operation, as it is equivalent to choose one flag from the old rooted curve.
- [1⁺] Loop addition, with probability $p_1/2$. Sample $0 \le \alpha < 1$ and fail immediately if $\alpha > z$. Return $D_{i+1} := \mathrm{RI}^+(D_i, a)$. The probability that this transition succeeds on an *n*-crossing plane curve D is $zp_1/2$.
- [1⁻] Loop deletion, with probability $p_1/2$. Provided f(a) is a loop, return $D_{i+1} := \operatorname{RI}^-(D_i, a)$. The probability that this transition succeeds if root flag a has f(a) a loop in D is $p_1/2$.
- [2⁺] Bigon addition, with probability $p_2/2$. Sample $0 \le \alpha < 1$ and fail immediately if $\alpha > z^2$. The flag *a* lies along a face of *d* edges; provided $d \ne 1$ (otherwise fail), uniformly sample the integer *k* between 1 and d 1. The flag $a' = (\sigma \tau)^k(a)$ is a distinct flag along the same face as *a*. Then return $D_{i+1} := \operatorname{RII}^+(D_i, a, a')$. The probability that this transition succeeds on any given additional flag *a'* along the root *d*-face is

$$\frac{z^2 p_2}{2(d-1)}.$$

[2⁻] Bigon deletion, with probability $p_2/2$. Fail if the flag *a* does not lie along a bigon. The size *d* of the face which would be produced by bigon deletion is the sum $|f(\tau\sigma(a))| + |f(\sigma^3\tau(a))| - 2 = d$. Sample $0 \le \beta < 1$ uniformly and fail if $\beta > (d-1)^{-1}$. Otherwise, return $D_{i+1} := \operatorname{RII}^-(D_i, a)$. The probability that this transition succeeds on a root *a* along a bigon is

$$\frac{p_2}{2(d-1)}.$$

[3] Triangle flipping, with probability p_3 . Fail if the flag *a* does not lie along a nondegenerate triangle. Otherwise, return $D_{i+1} := \text{RIII}(D_i, a)$. The probability that this transition succeeds assuming the root lies along a nondegenerate triangle is p_3 . We will prove that in the limiting distribution the probability that any given n-crossing rooted plane curve D is chosen is

(3)
$$\pi(D) = \frac{z^n}{K(z)},$$

where K(z) is the value of the generating function K(t) at z so if K(z) converges to a number, $\pi(D) \propto z^n$ (this happens provided $z < \mu^{-1}$). It follows that, if we ignore the roots of the sampled diagrams in order to sample *unrooted* diagrams, the probability of an unrooted diagram \overline{D} being sampled is

(4)
$$\pi(\overline{D}) = \frac{4n}{\operatorname{aut}\overline{D}}z^n.$$

We note that as the probability that the automorphism group of a plane curve is trivial tends exponentially quickly to 1 [Cha17], the typical probability of an unrooted diagram \overline{D} will be $\pi(\overline{D}) = 4nz^n$.

This Markov chain is ergodic as it satisfies the following three properties;

- (1) It is connected: It is possible to get from any one plane curve to another in a finite number of transitions. Provided $p_1, p_2, p_3 > 0$ it is possible to get between any two pairs of unrooted curves (by Theorem 3); provided $p_1 + p_2 + p_3 < 1$ it is guaranteed that any flag may be chosen as the root.
- (2) It is aperiodic: Since at each step there is a non-zero probability that the transition failes, there is no periodicity in the Markov chain.
- (3) The chain satisfies detailed balance: For any two curves D and N, the transition probabilities P and curve probabilities π satisfy,

$$P(D \to N)\pi(D) = P(N \to D)\pi(N).$$

This last point requires the most care and we discuss it below, with more details in Appendix A.

The fundamental theorem of Markov chains then yields the following result.

Theorem 1 (Slightly restated). This Markov chain satisfies detailed balance. Furthermore, if $p_1, p_2, p_3 \neq 0$, $(p_1 + p_2 + p_3) \neq 1$, and $z < \mu^{-1}$ the chain is ergodic with stationary distribution,

$$\pi(D) = \frac{z^{|D|}}{\sum_{\ell} k_{\ell} z^{\ell}} \propto z^{|D|}.$$

The proof of this result is primarily routine. We have omitted some details which can be found in Appendix A.

Proof. We will begin by assuming that

(5)
$$\pi(D) = \frac{z^{|D|}}{\sum_{\ell} k_{\ell} z^{\ell}} \propto z^{|D|},$$

and proving that detailed balance holds with this hypothesis. Notice that z has been chosen sufficiently small so that the denominator converges [Cha17]. In all cases the denominator is the same and a common factor in the calculations that follow, so we omit it.

Let D be a rooted plane curve of n vertices and a be the root flag in D. Observe first that the three pairs of reversing transitions (RI⁺, RI⁻), (RII⁺, RII⁻), and (RIII, RIII) all change the number of vertices by distinct complementary amounts; hence any two diagrams can be related by at most one pair of these transitions. Notice that if the vertex counts agree, then a pair of diagrams may be related by re-rooting. The main concern now is to prove the detailed balance equations for all possible transitions. Full details are in Appendix A; we demonstrate that detailed balance holds in the case of the RII moves, as an example.

Suppose that $N = \text{RII}^+(D, a, a')$ with root flag b. This means that N is unique in that $D = \text{RII}^-(N, b)$. The flags a, a' lie along a face in D of degree d. Then

(6)
$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$

(7)
$$\frac{z^2 p_2}{2(d-1)} z^n = \frac{p_2}{2(d-1)} z^{n+2}.$$

In all other cases, the transition probabilities are symmetrically zero. Hence we conclude that detailed balance holds with the hypothesized probability distribution.

Provided $p_1, p_2, p_3 \neq 0$ and $(p_1 + p_2 + p_3) \neq 1$, the Markov chain can reach all rooted plane curves as all flat Reidemeister moves and the rerooting move have nonzero transition probability. Hence in this case, the Markov chain is ergodic.

Notice that, by Theorem 3, any plane curve can be reduced to the trivial plane curve by a sequence of flat Reidemeister moves which *never* increases the number of vertices. This property implies that if we restrict our Markov chain to the set of plane curves of size at most L, then the chain remains ergodic. This is definitely not true for *knot diagrams* [KL06]. There exist diagrams of knots which are locally minimal in the sense that they can be reduced in crossing number, but only through sequences of moves which increase the crossing number at some point.

Corollary 5. Imposing the further restriction on the loop addition and bigon addition transitions that we fail if the input plane curve D would transition to have more than L vertices (i.e. if D has L or L-1 vertices, respectively) yields a Markov chain which explores all of $\mathscr{K}_{i\leq L}$ and is ergodic. Furthermore, the probability of sampling any plane curve D with $m \leq L$ crossings is,

$$\pi(D) = \frac{z^m}{\sum_{\ell=1}^L k_\ell z^\ell} \propto z^m$$

Proof. This follows from Theorem 3 and the proof of the previous theorem.

At this stage we are free to choose p_1, p_2, p_3 . It is not obvious, how they should be chosen to produce the most efficient sampling method. In practice, we simplified the problem by choosing $p_1 + p_2 + p_3 = 1$ and then applying a rerooting move after each step. The ergodicity of this chain follows by very similar arguments.

Corollary 6. For $p = (p_1, p_2, p_3)$ with $p_1 + p_2 + p_3 = 1$, consider the Markov chain which differs from the prior by, before and after each transition step, uniformly randomly re-rooting the state curve. Setting $p_1+p_2+p_3 = 1$ eliminates the re-rooting transition. Then this Markov chain is ergodic.

The proof is quite standard, we give it in Appendix A. Notice that in practice, there is no need between consecutive transitions to re-root more than once. In our experiments, we choose the five flat Reidemeister moves by taking $p_1 = 2/5$, $p_2 = 2/5$ and $p_3 = 1/5$.

3.2. Flat histogram sampling by Wang-Landau state density estimation. Our simulations in Section 4.1.1 show that the output curve sizes have high variance and change rather drastically with the parameter z. There are several different approaches that one might use to overcome this — such as multiple Markov chain Monte Carlo [Gey91; Orl98]. We have, instead, chosen the flat histogram method invented by Wang and Landau [WL01]. The transition probabilities are chosen so that the number of objects sampled at size n is approximately equal to the number of objects sampled at size m. This requires that

(8)
$$\frac{\pi(D)}{\pi(N)} \approx \frac{k_{|N|}}{k_{|D|}}$$

Wang and Wang and Landau's algorithm [WL01] achieves this by estimating the density of states — *i.e.* by estimating k_n . Consequently we must fix a maximum size L and set the probabilities of transitions to diagrams with more than L vertices to zero. The algorithm begins with a tuning phase which runs through the Markov chain adjusting transition properties until the above condition approximately holds. These final transition properties are then used for a MCMC sampler.

There are three advantages to this strategy over Boltzmann MCMC sampling:

- (1) Mixing of the Markov chain is more efficient: Rather than get caught up in one small range of sizes (in part due to the *local minimum* bottleneck; see Section 4.1.1), with Wang-Landau tuning the Markov chain moves frequently between all sizes. This improves the independence of two samples of the same size.
- (2) Guarantee of sampling objects of given size: The Markov chain produces objects of sizes up to L with nonzero probability, so given enough tuning and enough random samples, a sufficient number of objects of a desired size will be sampled.
- (3) Approximate enumeration: Transition probabilities found during the tuning step are directly related to the counts of objects of each size. We use this data to estimate the number k_n in Section 4.1.2.

Rather than depending on a single parameter z, sampling via a Wang-Landau implementation requires a data structure (ℓ, L, G) , where:

- ℓ is the minimum size of plane curves. To ensure ergodicity, we always have $\ell = 1$, although further results on plane curves may allow ℓ to vary (possibly in terms of L) while still guaranteeing ergodicity.
- L is the maximum size of plane curves.
- $G = (G_i)_{i=\ell}^L$ is a vector approximate enumeration data in the following sense: If $g_n = e^{G_n}$, then $g_n/g_{n-1} \approx k_n/k_{n-1}$. Since we do not know k_n for $n \ge 28$), this data is gathered via a tuning phase (described in Section 3.3 below).

Given the approximate enumeration data $G_n = \log(g_n)$, define probabilities

(9)
$$\operatorname{tp}(n,m) = \min\{1, \exp(G_n - G_m)\} = \min\{1, g_n/g_m\} \approx k_n/k_m,$$

unless $m < \ell$ or m > L in which case $\operatorname{tp}(n, m) = 0$. Wang-Landau flattened MCMC sampling then works as follows. The implementation which follows is largely the same as the Boltzmann implementation, with each transition instead required to pass a check of probability min $\{1, g_n/g_{n+k}\}$, where n is the size of the input curve and k is the change in size of the transition operation. Let $p_1, p_2, p_3 > 0$ and $(p_1 + p_2 + p_3) < 1$ and let D_i be the input rooted plane curve with n_i vertices and root flag a and perform one of the following six moves with different probabilities. If a move fails then set $D_{i+1} := D_i$.

- [0] Re-rooting, with probability $1 (p_1 + p_2 + p_3)$. Given the rooted diagram D_i , forget the root and select a new root b for D_{i+1} from the $(\operatorname{aut} D_i)/(4|D|)$ choices. The probability that this transition is chosen and succeeds is $1 (p_1 + p_2 + p_3)$.
- [1⁺] Loop addition, with probability $p_1/2$. Sample $0 \le \alpha < 1$ and fail immediately if $\alpha > \operatorname{tp}(n_i, n_{i+1})$. Return $D_{i+1} := \operatorname{RI}^+(D_i, a)$. The probability that this transition is chosen and succeeds on an *n*-crossing rooted plane curve is

$$\operatorname{tp}(n, n+1)\frac{p_1}{2} = \min\left\{1, \frac{g_n}{g_{n+1}}\right\}\frac{p_1}{2}$$

[1⁻] Loop deletion, with probability $p_1/2$. Sample $0 \le \alpha < 1$ and fail immediately if $\alpha > \operatorname{tp}(n_i, n_{i-1})$. Provided f(a) is a loop, return $D_{i+1} := \operatorname{RI}^-(D_i, a)$. The probability that this transition is chosen and succeeds is f(a) is a loop is

$$\operatorname{tp}(n, n-1)\frac{p_1}{2} = \min\left\{1, \frac{g_n}{g_{n-1}}\right\}\frac{p_1}{2}.$$

[2⁺] Bigon addition, with probability $p_2/2$. Sample $0 \leq \alpha < 1$ and fail immediately if $\alpha > \operatorname{tp}(n_i, n_{i+2})$. The flag *a* lies along a face of *d* edges; provided $d \neq 1$ (otherwise fail), uniformly sample the integer *k* between 1 and d-1. The flag $a' = (\sigma \tau)^k(a)$ is a distinct flag along the same face as *a*. Then return $D_{i+1} := \operatorname{RII}^+(D_i, a, a')$. The probability that this transition is chosen and succeeds on an *n*-crossing plane curve with an given cofacial flag a' is

$$\operatorname{tp}(n, n+2)\frac{p_2}{2(d-1)} = \min\left\{1, \frac{g_n}{g_{n+2}}\right\}\frac{p_2}{2(d-1)}$$

[2⁻] Bigon deletion, with probability $p_2/2$. Sample $0 \le \alpha < 1$ and fail immediately if $\alpha > \operatorname{tp}(n_i, n_{i-2})$. Fail if the flag *a* does not lie along a bigon. The size *d* of the face which would be produced by bigon deletion is the sum $|f(\tau\sigma(a))| + |f(\sigma^3\tau(a))| - 2 = d$. Sample $0 \le \beta < 1$ uniformly and fail if $\beta > (d-1)^{-1}$. Otherwise, return $D_{i+1} := \operatorname{RII}^{-}(D_i, a)$. The probability that this transition is chosen and succeeds if f(a) is a bigon and d > 1 is,

$$\operatorname{tp}(n, n-2)\frac{p_2}{2} = \min\left\{1, \frac{g_n}{g_{n-2}}\right\}\frac{p_2}{2(d-1)}$$

[3] Triangle flipping, with probability p_3 . Fail if the flag *a* does not lie along a nondegenerate triangle. Otherwise, return $D_{i+1} := \text{RIII}(D_i, a)$. The probability that this transition is chosen and succeeds if *a* lies along a nondegenerate triangle is p_3 .

3.3. Wang Landau tuning. Before sampling, we have to gather data for g_n (the approximate enumeration) via a tuning algorithm with parameters ℓ , the smallest size diagram to allow in the sample space (always in this article $\ell = 1$, as otherwise it is not necessarily clear if the Markov chain is ergodic), L, the largest size diagram in the sample space, ϵ , which describes the desired flatness of the sampling histogram, and Δ , a threshold for flatness of a histogram of occurrences.

A starting point D_0 in the sample space of diagrams is chosen; our algorithm starts with the figure-eight diagram in Figure 6. The values g_n are initialized to 0. Finally, a scaling factor f is initialized; we start it at f = 1.

The algorithm then proceeds as follows.

(1) If $f < \epsilon$, terminate.

- (2) A histogram $H = (H_n)_{n=\ell}^L$ of bins ℓ to L inclusive is initialized empty. This histogram will track the occurrences of diagrams of size n at each step of the Markov chain.
- (3) Step, via the Wang-Landau weighted algorithm described above, producing D_{i+1} from the current D_i .
- (4) Every S_0 steps, increment $H_{|D_{i+1}|}$ by 1, and increment $g_{|D_{i+1}|}$ by f.
- (5) Every S_1 steps, check if the histogram is Δ -flat, *i.e.*, check if

(10)
$$\frac{\min H}{1-\Delta} > \frac{\sum_{i=\ell}^{L} H_i}{L-\ell} > \frac{\max H}{1+\Delta}.$$

If so, let f := f/2 and proceed with step (1). Otherwise repeat step (3).

The numbers S_0 , S_1 are somewhat arbitrary, however S_1 should be chosen to allow sufficient time for the chain to diffuse over all sizes. We note that, especially for large spreads of ℓ and L, the tuning phase may be time-consuming. As the algorithm is tuning and g is being updated, the chain is *not* ergodic as it does not satisfy the detailed balance condition. However, the tuning phase approaches a limiting state where $g_{m+1}/g_m \approx k_{m+1}/k_m$. At that point, one may cease updating g and use those now fixed probabilities. This chain will satisfy detailed balance and so be ergodic. This tuned g data may then be reused for a number of further simulations without recalculation.

Theorem 2. Let $N \in \mathbb{N}$ and D be a plane curve with $1 \leq n \leq L$ vertices and let g_n be fixed. The Markov chain described in Section 3.2 has stationary distribution given by

$$\pi(D) \propto \frac{1}{g_n}$$

Since $g_n \approx k_n$, plane curves are sampled uniformly within a given size, and approximately uniformly across sizes.

Proof. We will begin by assuming that

(11)
$$\pi(D) \propto \frac{1}{g_{|D|}}.$$

where g_n data is fixed from the tuning phase of the algorithm.

Let D be a rooted plane curve of n vertices and a be the root flag in D. Observe first that the three pairs of reversing transitions (RI⁺, RI⁻), (RII⁺, RII⁻), and (RIII, RIII) all change the number of vertices by distinct complementary amounts; hence any two diagrams can be related by at most one pair of transitions, or, if their vertex counts agree, a re-rooting. The detailed balance proofs are nearly identical to the Boltzmann case, substituting in the new value of π and the new transition probabilities. Additional details can be found in Appendix A. Provided $p_1, p_2, p_3 > 0$, $(p_1 + p_2 + p_3) < 1$, and $\ell = 1$, the Markov chain can reach all plane curves (as all plane curves can be changed through a crossing-non-increasing pathway to the curve with one crossing). Hence in this case, the Markov chain is ergodic.

By Corollary 6, we are able to simplify the Markov chain by re-rooting between steps, and only using the five Reidemeister transition operations as we did in the Boltzmann case. Again, in our simulations (detailed in the next section) we chose $(p_1, p_2, p_3) = (2/5, 2/5, 1/5)$.

4. SIMULATIONS AND DATA

We implemented both MCMC samplers in c++. Plane curves are stored as combinatorial maps; a collection of vertices, edges, and flags with bidirectional references between flags and their vertices, as well as flags and their edges. At each step, an flag is selected from the diagram at random; this takes O(1) time, and the parameter $0 \le \alpha < 1$ is sampled uniformly at random. Both of the RI moves, as well as the RIII move take constant time. The RII⁺ move requires an extra random number $0 \le \gamma < 1$ which determines the second flag for the transition and is performed in constant time. The RII⁻ move requires both the sampling of the additional number $0 \le \beta < 1$ as well as a count of the face sizes diagonal to the bigon. A plane curve has average face degree strictly increasing and limiting on 4, so counting a face size requires a constant number of operations on average.

For low numbers n of crossings, it is also possible to sample plane curves uniformly through rejection sampling. A single sample is produced by sampling 4-valent maps uniformly until a plane curve is obtained. The maps are sampled via Gilles Schaeffer's bijection with blossom trees [Sch97] using his PlanarMap software[Sch99a; Sch99b]. The rejection step is simple, but plane curves are exponentially rare [SZ04], making this approach ineffective even for relatively small sizes: On a quad core 3.4Ghz Intel i5-7500 machine, sampling 10⁵ 10-crossing curves takes 4.9 seconds, but sampling the same number of 100-crossing curves takes 712.6 seconds. For comparison, it only takes 21.9 seconds to sample 10⁵ 100-vertex 4-valent maps.

In this section, we examine data from our simulations. First, we examine the distributions of plane curve sizes that our implementations produce. Then, we check how well our Wang-Landau implementation converges to the uniform distribution across fixed sizes by comparing statistics to the rejection sampler. The data of these sections are based on the following sampler runs:

- (1) Wang-Landau (WL). After tuning a Wang-Landau Markov chain to $f < 10^{-8}$ with histogram flatness threshold $\Delta = 0.99$, we drew a total of 2×10^7 samples for sizes $1 \le n \le 500$ with 10^3 steps between each sample. The tuning phase took approximately 500 minutes, and the sampling took 160 minutes. The sampled size probability distribution for this run is presented in Figure 12. The minimum number of samples for any size is 3.8933×10^4 .
- (2) Rejection. Using a rejection sampler for plane curves, we gathered 10^4 samples of plane curves with n vertices, for each n = 5m from 5 to 100. This took approximately 40 minutes.
- (3) All 4-valent maps. Using a uniform sampler for all 4-valent maps, we gathered precisely 10^4 samples of 4-valent maps with n vertices, for each n = 5m from 5 to 100. This took approximately 30 seconds.

Finally, we examine how our Wang-Landau sampler augmented to sample knot diagrams compares to the rejection sampler, as another check on the theoretical limiting distribution and our implementation.

4.1. Size distributions. We first examine the sample histograms of the Boltzmann and Wang-Landau samplers. This serves several purposes: First, as we are unable to sample diagrams of given fixed size, we would like to know how frequently we will sample a particular size using these methods. Second, in order to avoid correlated samples we would like our Markov chain to explore the full range of lengths frequently. Last, we can use comparisons of the sampled size distribution to better understand the counting sequence of plane curves.

4.1.1. Boltzmann sampler implementation and the Boltzmann parameter z. The MCMC sampler approximating the Boltzmann distribution on plane curves (described in Section 3.1) samples from a distribution,

(12)
$$\pi(D) \propto z^{|D|}$$

where the parameter z affects the size of plane curves produced. Hence the probability of sampling any plane curve of size n is proportional to $k_n z^n$. As noted above, we choose $z < 1/\mu$ and so by Equation 1 this probability is asymptotic to $(\mu z)^n n^{\gamma-2}$. Note that γ is expected to be *negative* so any choice of $z < \mu^{-1} \approx 0.0876$ will not, a priori produce a finite local maximum.

To observe this we ran a number of experiments sampling 10^6 plane curves of a maximum size of 200 crossings with varying z. These data produce the approximate size distributions of Figure 10. This figure implies that it is difficult to pick z to obtain samples at large size while still sampling many objects of small size.



FIGURE 10. Histogram of samples from various MCMC runs, with different z-values. The maximum-size plane curve was 200 vertices. The mixing time was 10^3 steps, and a total of 10^6 samples were drawn. The run for z = 0.05 took 65 seconds to complete, the run for z = 0.1 took 93 seconds, and run time generally increased with z.

We note that we want our algorithm to return to small sizes on a regular basis. First, returning to small sizes "erases" the entire object before producing a new sample. Second, the best known result on ergodicity relies on paths through small plane curves. Not enough is known about the connectivity of the space of plane curves to alleviate these concerns, providing strong justification for using the Wang-Landau sampler variant instead.

4.1.2. Wang-Landau implementation. In comparison with the MCMC Boltzmann sampler, the tuning phase of the Wang-Landau sampler ensures that the size distribution sampled is actually approximately flat. Figure 11 demonstrates how values of g_n approach the exact



FIGURE 11. As the tuning phase proceeds, the values $G_n = \log g_n$ converge to G_n^{final} . We also observe that $G_n^{final} \approx \log k_n$ for available exact enumeration data, $n \leq 27$.



FIGURE 12. Distribution of plane curve sizes from a run of the Wang-Landau sampler. The sample histogram flatness is $\max\{0.973325, 1-(1.030075-1)\} = 96.9925\%$.

numbers of rooted plane curves, k_n . Figure 12 presents the size distribution of a run of 2×10^7 samples of size between 1 and 500 produced after tuning to $f < 10^{-8}$ with $\Delta = 0.99$.

As mentioned, the Wang-Landau algorithm tuning step provides an estimation of ratios in the counting sequence for plane curves. We can use this data to provide approximations to the numbers of curves of given sizes. Using the tuning data for the run above ($\ell = 1, L =$ $500, f < 10^{-8}, \Delta = 0.99$) we obtained approximate counts for n = 1 to n = 27 and can compare the approximations to the precise counts from [ZZ09] in Table 1. All approximate counts are within 0.004 of their exact value.

n	WL estimated k_n	Exact k_n	% Error
1	2	2	0
2	8	8	$-4.07 \cdot 10^{-2}$
3	42.01	42	$3.49 \cdot 10^{-2}$
4	260.15	260	$5.61 \cdot 10^{-2}$
5	1,796.5	1,796	$2.79 \cdot 10^{-2}$
6	13,404.48	13,396	$6.33 \cdot 10^{-2}$
7	$1.06\cdot 10^5$	$1.06 \cdot 10^{5}$	$5.25 \cdot 10^{-2}$
8	$8.71 \cdot 10^{5}$	$8.71 \cdot 10^{5}$	$6.43 \cdot 10^{-2}$
9	$7.42 \cdot 10^{6}$	$7.42 \cdot 10^{6}$	$3.99 \cdot 10^{-2}$
10	$6.5 \cdot 10^{7}$	$6.5 \cdot 10^{7}$	$6.25 \cdot 10^{-2}$
11	$5.82 \cdot 10^{8}$	$5.83 \cdot 10^{8}$	$-4.42 \cdot 10^{-2}$
12	$5.32 \cdot 10^{9}$	$5.32 \cdot 10^{9}$	$7.57 \cdot 10^{-2}$
13	$4.95 \cdot 10^{10}$	$4.94 \cdot 10^{10}$	0.12
14	$4.66 \cdot 10^{11}$	$4.65 \cdot 10^{11}$	0.16
15	$4.44 \cdot 10^{12}$	$4.43 \cdot 10^{12}$	0.16
16	$4.28 \cdot 10^{13}$	$4.27 \cdot 10^{13}$	0.12
17	$4.16 \cdot 10^{14}$	$4.16 \cdot 10^{14}$	$8.45 \cdot 10^{-2}$
18	$4.08 \cdot 10^{15}$	$4.08 \cdot 10^{15}$	0.1
19	$4.04 \cdot 10^{16}$	$4.03 \cdot 10^{16}$	0.19
20	$4.02 \cdot 10^{17}$	$4.02 \cdot 10^{17}$	0.18
21	$4.04 \cdot 10^{18}$	$4.02 \cdot 10^{18}$	0.3
22	$4.07 \cdot 10^{19}$	$4.06 \cdot 10^{19}$	0.25
23	$4.12 \cdot 10^{20}$	$4.11 \cdot 10^{20}$	0.24
24	$4.19 \cdot 10^{21}$	$4.18 \cdot 10^{21}$	0.19
25	$4.29 \cdot 10^{22}$	$4.28 \cdot 10^{22}$	0.3
26	$4.41 \cdot 10^{23}$	$4.39 \cdot 10^{23}$	0.38
27	$4.54 \cdot 10^{24}$	$4.53 \cdot 10^{24}$	0.25

TABLE 1. Comparison of counts of rooted plane curves from the Wang-Landau tuning step with $\ell = 1, L = 500, \varepsilon = 10^{-8}, \Delta = 0.99$, versus exact values gathered using another method [ZZ09].

We are able to obtain estimates for the unknowns μ and γ in the predicted asymptotic growth formula $k_n \sim C\mu^n n^{\gamma-2}$ from Wang-Landau g_n data. We attempted to use simple ratio estimates $r_n = g_{n+1}/g_n \sim \mu \left(1 + \frac{\gamma-2}{n}\right)$, however the results are extremely noisy. Instead we used linear regression to fit to the model

(13)
$$\log g_n = \log C + n \log \mu + (\gamma - 2) \log n$$

Since we expect the g_n data to be noisier for larger n we fitted the above linear form to a subset of the data $\{g_n \mid 10 \le n \le n_{max}\}$. We then varied n_{max} to get a rough estimate of corrections to the above asymptotic form. The resulting estimates (as functions of n_{max}) are shown in Figure 13. These results are consistent with earlier estimates [ZZ09] of $\mu \approx 11.416$ and $\gamma = -\frac{1+\sqrt{13}}{6} \approx -0.768$.

4.2. Face degrees in plane curves. In this section we seek statistic which distinguish plane curves from all 4-valent maps. A simple class of statistics to gather from random maps



FIGURE 13. Estimates of μ and γ using least squares fits on $\{g_n \mid 10 \leq n \leq n_{max}\}$ plotted as a function of n_{max} .

are vertex and face degrees. Plane curves, as a subclass of 4-valent planar maps, only ever have vertex degree 4, so only face statistics are nontrivial. Euler's formula implies that the average face degree for any 4-valent map of n crossings is 4n/(n+2), so this also cannot distinguish plane curves from its superclass. However, we will see that the *distribution of* face degrees differs.

4.2.1. Face degree probabilities. We check the counts of faces of fixed degree which appear (for a curve of n crossings, this takes O(n) time to compute as all plane curves have 4n flags, each of which needs only be visited once). These quantities are expected to exhibit linear growth, in agreement with the results for a large number of map classes [Lis99]. In fact, it is known:

Theorem 7. Let $k \ge 1$, and let $p_{n,k}$ denote the probability that an arbitrary face of a random plane curve of n crossings has degree k. Namely, notice that $(n+2)p_{n,k}$ is the expected number of degree k faces in a random plane curve of n crossings. Then

(14)
$$1 > \limsup_{n \to \infty} p_{n,k} \ge \liminf_{n \to \infty} p_{n,k} > 0.$$

Proof. Certainly, $0 \le p_{n,k} \le 1$ for all n, k as they denote probabilities; furthermore, $\sum_{k=1}^{n} p_{n,k} = 1$. The pattern theorem for plane curves [Cha17] says that, for any prime substructures T_k containing a k-gon, there are constants $c_k > 0$, $1 > d_k > 0$ and $N \ge 0$ so that for all $n \ge N$ the probability that an arbitrary plane curve of size n contains more than $c_k n$ copies of T_k is at least $1 - d_k^n$.

So for $n \ge N$, $(n+2)p_{n,k} > (1-d^n)c_k n$. Solving for $p_{n,k}$ and passing to limit yields,

(15)
$$\liminf_{n \to \infty} p_{n,k} > \liminf_{n \to \infty} (1 - d^n) c_k \frac{n}{n+2} = c_k.$$

That $1 > \limsup_{n \to \infty} p_{n,k}$ for any k follows from that for $\ell \neq k$ the existence of a ℓ -gon lowers the number of faces which may be k-gons. That all $\liminf_{n \to \infty} p_{n,\ell} > 0$ yields the result.

The following proposition summarizes the results we will use for planar 4-valent maps.

Theorem 8 (Follows from [GR94; Lis99]). The limiting expected number (as maps grow large) of faces of degree k in a random 4-valent map is,

(16)
$$\frac{n+2}{k}[y^k]\left((1/3)(1+y/2)^{-\frac{1}{2}}(1-5y/6)^{-\frac{3}{2}}\right).$$

Proof. This is a rephrasing of selected results in [GR94; Lis99] in the language of 4-valent maps. Theorem 1 of [GR94] says that the generating series of limiting probabilities q_k (as maps grow large) that the root vertex in an arbitrary map has degree k is,

(17)
$$\sum_{k\geq 0} q_k y^k = (1/12)(1+y/2)^{-\frac{1}{2}}(1-5y/6)^{-\frac{3}{2}}.$$

Duality of the class of rooted planar maps says the same result holds for faces. The bijection between m-edged rooted planar maps and m-faced rooted planar quadrangulations then says the same result holds for vertices in quadrangulations (see for instance the proof of Proposition 12 in [BC13]).

Section 2.5 in [Lis99] relates, for a quadrangulation of m vertices, the probability $q_{k,m}$ that the root vertex has degree k to the probability $p_{k,m}$ that an arbitrary vertex has degree k by,

(18)
$$p_{k,m} = \frac{4(m-2)}{m} \frac{q_{k,m}}{k} \sim 4 \frac{q_{k,m}}{k}$$

whence $p_k = \lim_{m \to \infty} p_{k,m}$ and $q_k = \lim_{m \to \infty} q_{k,m}$ are related by $p_k = 4q_k/k$. Noting that 4-valent maps with n vertices are dual to quadrangulations with n + 2 vertices yields the result.

We have computed linear regressions using least squares for 4-valent maps sampled using the Schaeffer bijection, plane curves sampled using rejection, and plane curve sampled using our Wang-Landau sampler for face degrees from 1 to 9 and presented these data in Table 2 alongside precise densities obtained via Taylor series expansion on the result of Theorem 8. We present these data for small faces sizes in Figure 14. In all cases, the data of the Wang-Landau sampler resides within the error bars of the rejection curve sampler data, and away from the arbitrary 4-valent map data.

We note that: Random curve diagrams have fewer bigons and quadrangles than generic 4-valent maps. This phenomenon is specific to degrees 2 and 4 (at least for face degrees at most 9); every other degree face is more common in random plane curves.

It is furthermore expected of classes of random maps that, for fixed size, number of faces of fixed degree k is a normally distributed statistic [DP13]. In Figure 15 below, we compare distributions of different k-gon ratios for n = 40 and n = 100 crossings. As expected, the curves show a close similarity between the uniformly sampled curves and the Wang-Landau MCMC sampled curves. In the cases of 1- and 2-gons, it is easy to see a difference from the all 4-valent map sampler. In the case of larger faces, the differences in averages are on a much smaller order, and the curves are no longer possible to distinguish. In all cases, it seems that as the number of crossings n grows large, the distributions are approximately normal.

We note here that the face degree probabilities are closely related to the probability that a given Markov chain transition succeeds on a given diagram, although it is actually the quantity $q_{k,n} = \frac{kp_{k,n}(n+2)}{4n} \approx \frac{kp_{n,k}}{4}$ discussed prior which is at play (transitions by definition occur at the root along the root face). Namely, the probability that an RIII move can succeed

k	4-valent p_k , theor.	4-valent p_k	Rejection p_k	WL p_k
1	$\frac{1}{3} = 0.\overline{3}$	$0.3328 \pm 2 \cdot 10^{-4}$	$0.3496 \pm 2 \cdot 10^{-4}$	$0.35036 \pm 4 \cdot 10^{-5}$
2	$\frac{1}{6} = 0.1\overline{6}$	$0.1662 \pm 2 \cdot 10^{-4}$	$0.1407 \pm 1 \cdot 10^{-4}$	$0.14056 \pm 3 \cdot 10^{-5}$
3	$\frac{13}{108} = 0.12\overline{037}$	$0.1200 \pm 2 \cdot 10^{-4}$	$0.1222 \pm 1 \cdot 10^{-4}$	$0.12257 \pm 3 \cdot 10^{-5}$
4	$\frac{55}{648} \approx 0.08488$	$0.0846 \pm 1\cdot 10^{-4}$	$0.0831 \pm 1 \cdot 10^{-4}$	$0.08298 \pm 2 \cdot 10^{-5}$
5	$\frac{83}{1296} \approx 0.06404$	$0.0641 \pm 2 \cdot 10^{-4}$	$0.0663 \pm 2 \cdot 10^{-4}$	$0.06624 \pm 2 \cdot 10^{-5}$
6	$\frac{377}{7776} \approx 0.048448$	$0.0484 \pm 2 \cdot 10^{-4}$	$0.0500 \pm 1 \cdot 10^{-4}$	$0.04977 \pm 2 \cdot 10^{-5}$
7	$\frac{1751}{46656} \approx 0.03753$	$0.03769 \pm 8 \cdot 10^{-5}$	$0.0395 \pm 1 \cdot 10^{-4}$	$0.03941 \pm 1 \cdot 10^{-5}$
8	$\frac{101}{3456} \approx 0.02922$	$0.0292 \pm 1 \cdot 10^{-4}$	$0.0305 \pm 1 \cdot 10^{-4}$	$0.03059 \pm 1 \cdot 10^{-5}$
9	$\frac{115825}{5038848} \approx 0.02299$	$0.0232 \pm 1 \cdot 10^{-4}$	$0.02419 \pm 8 \cdot 10^{-5}$	$0.02419 \pm 1 \cdot 10^{-5}$

TABLE 2. Limiting face degree densities. Theoretical fits for 4-valent p_k come from Taylor series expansion of the result in Theorem 8. Experimental columns are the slopes of linear functions fit to data using least squares.

is $q_{3,n+2} \approx 0.09193$, the probability that an RI⁻ operation will succeed is $q_{1,n+2} = 0.08759$, and the probability that an RII⁻ operation will succeed is $q_{2,n+2} \approx 0.07028$ (not taking into account the extra Metropolis-Hastings step required to create large faces or sites which would create monogons, see the description of move RII⁻ in Section 3).

We reinterpret briefly the probabilities $p_{k,n+2}$ in the context of random knot diagrams. We note that in the case of prime alternating diagrams, face degrees are related to the hyperbolic volume of the resulting knot [Obe16]. We further note that a random curve has $\approx 0.35(n+2) > (n+2)/3$ monogons says that a random knot diagram has, on average, at least .35(n+2) vertices that have *no* impact on the knot type and could be immediately reduced by a Reidemeister I move. One half of all crossing assignments for bigon vertices can be reduced by Reidemeister II moves, so a random knot diagram will have around 0.14(n+2)/2bigons that can be removed by Reidemeister II moves.

4.2.2. Maximum face degree. Another quantity of interest in the study of planar maps is the maximum vertex degree Δ_n and the maximum face degree Δ_n^* . As noted above, $\Delta_n = 4$ because all vertices are 4-valent, so we examine expectation of Δ_n^* . It is expected that this quantity exhibits $\Theta(\log n + \log \log n)$ growth as it does for general maps. A result of [GW00] has that for general maps the expected maximum face degree is precisely:

(19)
$$\mathbb{E}(\Delta_n^*) = \frac{\ln(n) - \frac{1}{2}\ln(\ln(n))}{\ln(6/5)}$$

We present the difference between expectations and that of general maps in Figure 16. From this, it does not appear that plane 4-valent maps exhibit the same behavior as the general map case: Indeed, as the bijection between arbitrary maps and 4-valent maps makes both faces and vertices into 4-valent map faces, the expected maximum vertex degree Δ_n^* is in fact an expectation of a maximum $\mathbb{E}(\max(\Delta_n, \Delta_n^*))$ over both vertex and face degrees of arbitrary maps. We have plotted histograms of the maximum face degree distribution



FIGURE 14. Average counts $p_{k,n+2}(n+2)$ of faces of degrees 1–6.

in curves and 4-valent maps for fixed size in Figure 17. These distributions are clearly not Gaussian. We note that the histograms of all cases look similar even though they have differing means (see Figure 16), this can be explained by noting that the difference between the means is small and growing very slowly in n.

We present the data for $\mathbb{E}(\Delta_n^*)$ of the Wang-Landau MCMC sampler up to 500 crossings it is impractical to gather samples of this size from the rejection sampler—in Figure 18. The



FIGURE 15. Distribution of k-gon counts for n = 40 and n = 100 crossings, for various k.



FIGURE 16. Average size of largest face Δ_n^* .



FIGURE 17. Distribution of largest face degree Δ_n^* for n = 40 and n = 100 crossings.

 $\Theta(\log n)$ trend continues, as hypothesized. We note that the quantity $\mathbb{E}(\Delta_n^*)$ is related to the success rate of the RII⁻ transition, and suggests that the rejection rate for creating large faces through this transition is roughly bounded by $1/\Theta(\log n)$.

4.3. Average Casson invariant. Our results above suggest that objects we sample using our Wang-Landau algorithm are giving the same statistics as those generated using rejection sampling. To further test this idea, we can compare further plane curve statistics.

Plane curves are equivalent to Arnol'd's spherical curves [Arn95]. Thus, we can check \mathbb{Z} -valued spherical curve invariants for plane curves (these are not defined in the case of a general 4-valent map). These are both simple to compute and interesting knot theoretically: By following the crossings in order around the plane curve and counting those which are "interlaced", we are able to compute $-\frac{1}{2}(2St + J^+)$ in $O(n^2)$ time (we direct the reader to [Arn95] for more details). This curve invariant is in fact related to a knot invariant, the



FIGURE 18. Expected maximum face valence in plane curves appears to grow logarithmically.

finite type invariant v_2 , also called the Casson invariant: $-\frac{1}{8}(2St + J^+)$ is the expected value of v_2 over all possible over-under crossing sign assignments to the plane curve. A comparison of $\mathbb{E}(v_2)$ data for Wang-Landau MCMC and rejection samples are presented in Figure 19; the data are remarkably consistent.



FIGURE 19. (Left) Average $-\frac{1}{8}(2St + J^+) = \mathbb{E}(v_2)$. (Right) Same statistic with approximate leading linear behavior subtracted to amplify detail. As this statistic is not well-defined for an arbitrary 4-valent map, that data is not present.

We present the data for $\mathbb{E}(v_2)$ up to 500 vertices generated from the Wang-Landau algorithm in Figure 21. Our data suggest that the average v_2 invariant grows linearly with the number of crossings with a limiting slope of 0.0268 ± 0.001 . As a comparison, Even-Zohar et al. [Eve+16] prove that for the Petaluma model of random knots, the expectation of v_2 grows quadratically with size with leading coefficient $1/24 \approx 0.0417$. As the petal diagrams of size n of the Petaluma model can be viewed as "star diagrams" with $\frac{n^2-3n}{2}$ crossings, size in the Petaluma model is related quadratically to that of our model. Thus, it is reasonable that our average v_2 data should grow linearly. It is also expected that the distribution of $\mathbb{E}(v_2)$ over curves of a fixed size tends to be normal; we present histograms of this in Figure 20 that appears to agree with this hypothesis.



FIGURE 20. Distribution of mean v_2 invariant for n = 40 and n = 100 crossings.

4.4. Knotting probabilities. The main aim of constructing these Markov chains was to provide an efficient mechanism to sample large random knot diagrams. We have used the Wang Landau chain to sample random plane curves which we then use to construct random knot diagrams. This last step is done by mapping each vertex to an over or under crossing uniformly at random. We then use knot invariants to determine the topology of the resulting diagram and we compare the resulting knotting probabilities to those found using the rejection sampler.

In [Cha17] the first author gathered data for knotting probabilities in knot diagrams of size up to 100 crossings. We have used our Wang-Landau sampler to produce random knot diagrams and classify them by HOMFLY-PT polynomial using lmpoly [EM91; EM97] of size up to 230 crossings. There are two reasons that we have not investigated beyond this size. First, the version of lmpoly included with plCurve [ACC17] will not compute HOMFLY-PT polynomials of diagrams of more than 255 crossings. Second, the difficulty of computing the HOMFLY-PT polynomial increases dramatically with the number of crossings; the algorithm employed by lmpoly has exponential run time[EM91] and computation is known to be NP-hard. We gathered data using two runs (from the same tuning data for $1 \le n \le 230$ with $f < 10^{-8}$ and $\Delta = 0.99$). In each run we took 10^3 steps between samples, and a total of 1.15×10^7 diagrams were sampled. Each run took approximately 9 hours to compute, with HOMFLY-PT polynomial calculation being the primary bottleneck. The histograms for sizes sampled are presented in Figure 22. We note some caveats for these experiments:

(1) The HOMFLY-PT polynomial is *not* sufficient to distinguish all knot types, and indeed there are infinite families of knots demonstrating this [Kan86]. However, it



FIGURE 21. The average v_2 invariant appears to have linear dominant order growth. Here we present its ratio by number of crossings for finer resolution, where a naïve least squares linear regression suggests a slope of $(2.68 \pm 0.1) \times 10^{-2}$.

is still unknown whether the unknot 0_1 is the only knot with trivial HOMFLY-PT polynomial; this is related to the still open Jones Conjecture [Kau87].

- (2) To facilitate a greater number of samples, we imposed a timeout on lmpoly's HOMFLY-PT calculation of 10ms. See Figure 23. Knots whose HOMFLY-PT failed to be calculated before this cutoff are still counted and categorized as unclassified. Hence it is possible that probabilities presented in our data which follows are smaller than the actual. There is evidence that these failures are rare for simple knots such as the trefoil and the unknot see Figure 24. We note that the uniform data presented alongside our MCMC sampled data was gathered at a larger timeout.
- (3) We ignore chirality in these data. A chiral knot is a knot which is different than its mirror, such as the trefoil 3_1 , while an achiral knot like the figure-eight knot 4_1 is equivalent to its mirror image. By symmetry (*i.e.* by flipping all crossing signs simultaneously), it can be seen that the probability of drawing a chiral knot is equivalent to that for its mirror image. For the case of the chiral granny $3_1#3_1$ and achiral square $3_1#3_1^*$ composite knots, our data suggests that the square knot is as likely as either chiral image of the granny knot; we hence suppress data for the granny knot.

It is believed that for lattice models of random knots the number $\kappa_n(K)$ of knot diagrams with fixed knot type K has asymptotic growth rate

(20)
$$\kappa_n(K) \sim C_K \tau_K^n n^{\alpha_K + N_K},$$

where C_K depends on the knot type and N_K is the number of prime components making up the knot type K. It is believed that the constant τ_K does not depend on the knot type [DT97; Orl+98; RW91a]. It has been proved that τ_0 exists for many lattice models (this follows from standard supermultiplicativity arguments) and also for random knot diagrams [Cha17],



FIGURE 22. Histograms of sample sizes for two datasets used to examine random knotting.



FIGURE 23. The time taken to compute the HOMFLY-PT polynomial increases dramatically with the complexity of the knot diagram. Accordingly, we imposed a time-cutoff on these computations and if the invariant was not computed before the indicated time then it was left unclassified. In this figure we show the proportion of knot diagrams that failed to be classified as a function of the number of crossing and the cutoff time.



FIGURE 24. The probability of a knot being classified as an unknot and a trefoil as a function of the number of crossings and the time-cutoff of the HOMFLY-PT computation. This data shows that increasing the cutoff does not significantly affect this probability. This is consistent with the hypothesis that most HOMFLY-PT computation failures are for complicated knots.



FIGURE 25. A random trefoil knot diagram of 50 crossings; the knotted portion of the curve is highlighted. We see that the knotted portion is quite small and this is expected to be typical.

however it is still an open problem to prove the existence of τ_K for any other knot type. It is known, however, that τ_0 is strictly smaller than 2μ for random knot diagrams [Cha17]; which is comparable to a similar result for self-avoiding polygons [Pip89; SW88]. There is strong numerical evidence for self-avoiding polygons that the exponent α_K is independent of knot type [Orl+98; RR11]. Consequently we conjecture that for random knot diagrams that

(21)
$$\kappa_n(K) \sim C_K \tau_0^n n^{\alpha + N_K},$$

where $\alpha = \gamma - 2$, where γ is the same "universal" critical exponent [SZ04] in the asymptotic formula for plane curves. This asymptotic form is consistent with the idea that the knotted portion of a random knot diagram is localized — see Figure 25.



FIGURE 26. Plots for a WL sample against the data for a rejection sampler in [Cha17]. Plots for other knot types are similar, and rejection data is consistent with WL sampled data throughout.

Under this assumption, the probability that a random knot diagram of size n exhibits knot type K scales as

(22)
$$p_n(K) = \kappa_n(K) / \kappa_n \sim D_K \rho^n n^{N_K},$$

where $0 < \rho < 1$. We plot knot probability data from both our Wang Landau sampler and rejection sampling in Figure 26. We see that both sampling methods agree and that the data is consistent with the scaling form in Equation 22. In particular, in Figure 27, we plot the logarithm of the knotting probabilities divided by n^{N_K} and see that the resulting slopes are extremely similar. A simple linear regression of this data shows that $\rho \approx 0.95$. This is evidence that the growth rate of random knot diagrams of fixed knot type K is independent of K and that

(23)
$$\tau_K = \tau_0 \approx 2 \times 11.41 \times 0.95 = 21.7.$$

The authors intend to test this hypothesis further in future work.

As a final comparison to the uniform sampler data, we compute ratios of knot probabilities as in [RR11]. Namely, the expected growth rates of knot probabilities has that, for two knot types K and L, the ratio of probabilities should obey,

(24)
$$\frac{p_n(K)}{p_n(L)} \sim \frac{D_K \rho^n n^{N_K}}{D_L \rho^n n^{N_L}} = \frac{D_K}{D_L} n^{N_K - N_L}$$

Hence we expect $p_n(K)/p_n(L) \cdot n^{N_L - n_K}$ to tend to a constant as n increases. We plot this data for ratios of prime knots in Figure 28, and for ratios of unknots to prime knots Figure 29. We also show the ratios of square knots to trefoils and unknots in Figure 30. While this ratio data is noisy for larger n it is consistent with the above scaling form.

5. CONCLUSION

We have described a new Markov chain Monte Carlo method for sampling random plane curves efficiently. This then trivially extends to sample random knot diagrams by mapping vertices to crossings. This enables us to sample of large knot diagrams which are otherwise simply too rare to sample by rejection sampling methods. Due to the difficulty of tuning



FIGURE 27. After normalizing by a factor of n^{-N_K} (where N_K is the number of prime components of the knot type K, knot probabilities are all approximately log-linear.



FIGURE 28. Ratios of probabilities of prime knot types in diagrams. Probabilities within each ratio are taken from independent runs of the MCMC sampler. Legend entries are sorted by their values at 1/n = 0.05.

the Boltzmann MCMC to sample diagrams of a wide range of sizes, we modified our original Markov chain to use the flat histogram methods of Wang and Landau. This results in a chain that samples approximately uniformly across a wide range of sizes, and additionally gives estimates of the numbers of plane curves and knot diagrams. These estimated counts are in agreement with previously conjectured asymptotic forms. Hence we conclude that our MCMC implementation can attain similar accuracy of the uniform rejection methods in far less time. We have then tested the data from the Wang-Landau chain against data from a rejection sampler and find strong agreement across a wide range of statistics.



FIGURE 29. Ratios of probabilities of the unknot with prime knots, with correction factor n. Probabilities within each ratio are taken from independent runs of the MCMC sampler. Legend entries are sorted by their values at 1/n = 0.05.



FIGURE 30. Ratios of probabilities of the square knots to trefoils and unknots (with corrections of n and n^2 respectively. Probabilities within each ratio are taken from independent runs of the MCMC sampler. Legend entries are sorted by their values at 1/n = 0.05.

Plane curves are a subset of 4-valent maps and it is not well undestood how different these two sets of objects are. With this in mind, we computed the average maximum face degree and face degree degree distributions and find significant differences. One of the main aims of sampling random plane curves is to study random knotting. Our Wang-Landau sampler allows us to generate significant numbers of random knot diagrams of a range of sizes. We have then classified the knot types of those diagrams using HOMFLY-PT polynomials. This data allows us to conjecture that random knot diagrams have very similar asymptotic behaviour to other models of random knotting, such as self-avoiding polygons. These asymptotic forms are consistent with the idea that the knotted portion of a random knot diagram is quite localized. We also give evidence that the HOMFLY-PT polynomial software [EM97] that we used struggles to compute invariants of complicated knots, but does succeed for simple knots, even when their embeddings might be very large.

5.1. Future work. The Markov transitions presented here are based on "shadows" of the Reidemeister moves on knot diagrams. It is thus natural to consider a Markov chain on knot diagrams generated by similar transitions corresponding to the proper Reidemeister moves, taking into account over-under signing of the diagram. This would produce a Markov process on the so-called Reidemeister graph [BC18]. Indeed, we expect this Markov chain to be ergodic (with ergodicity classes being fixed knot types), and we expect transitions to be of similar computational complexity. It would also have the advantage of not requiring a knot-classification step. We note, however, that Corollary 5 fails in this case; it is known that there are diagrams who represent the same knot type, but whose transition paths all involve an increase in the number of crossings [KM06]. To make matters worse, unlike the spaces of plane curves where the diameter has $n^{3/2}$ growth, the upper bound on the diameters of the spaces of knot diagrams is far larger [HL01; Lac15; Now09]. Hence in the case of knot diagrams, care must be taken to ensure that there are satisfactory parameters for the Markov chain to converge to the uniform distribution in a reasonable amount of time. It should also be noted that rejection sampling becomes even less efficient for sampling fixed knot types, since not only are knot diagrams exponentially rare in the space of 4-valent maps, but knot diagrams of a specific type are exponentially rare in the space of knot diagrams.

Beyond sampling knot diagrams with fixed knot type, the flat Reidemeister moves discussed in this paper also apply to planar immersions of any fixed number of circles. These diagrams are called *link shadows*; the smallest such object is the unique 2-crossing 4-valent planar map of 2 link components. In this case, the Markov chain is still ergodic (the proof of Theorem 1 is not affected by the immersion having a different number of link components). Hence this technique could be used to sample large immersions of any fixed number of link components. We could also restrict or alter the transitions; for instance we could remove the shadow Reidemeister I move, whence the ergodicity classes of the Markov chain would be spherical curves of fixed spherical Whitney number [Arn95; Now09]. Further, by using Reidemeister moves instead of flat Reidemeister moves, one could also sample link diagrams of fixed link types.

Acknowledgements

The authors would like to thank the Pacific Institute of Mathematics Sciences for funding the collaborative research group (CRG) on Applied Combinatorics. Part of this work started at a summer school on applied combinatorics funded by that CRG. The second author acknowledges funding from NSERC in the form of a Discovery Grant. The authors would also like to thank Chris Soteros, Jason Cantarella, and Stu Whittington for many helpful discussions.

References

- [AB26] J. W. Alexander and G. B. Briggs. "On Types of Knotted Curves". In: The Annals of Mathematics 28.1/4 (1926), p. 562. ISSN: 0003-486X. DOI: 10.2307/1968399. [AC83] C. Aragão de Carvalho and S. Caracciolo. "A new Monte-Carlo approach to the critical properties of self-avoiding random walks". In: Journal de Physique 44.3 (1983), pp. 323–331. ISSN: 0302-0738. DOI: 10.1051/jphys:01983004403032300. [ACC17] T. Ashton, J. Cantarella, and H. Chapman. plCurve: Fast polygon library. Available at http://www.jasoncantarella.com/wordpress/software/plcurve/. 2017. [Arn95] V. I. Arnol'd. "The geometry of spherical curves and the algebra of quaternions". In: Russian Mathematical Surveys 50.1 (1995), p. 1. DOI: 10.1070/RM1995v050n01ABEH001662. I. Benjamini and N. Curien. "Simple random walk on the uniform infinite [BC13] planar quadrangulation: subdiffusivity via pioneer points". In: Geometric and Functional Analysis 23.2 (Mar. 2013), pp. 501–531. ISSN: 1420-8970. DOI: 10.1007/s00039-013-0212-0. [BC18] A. Barbensi and D. Celoria. "The Reidemeister Graph Is a Complete Knot Invariant". In: ArXiv Preprints (2018). arXiv: 1801.03313 [math.GT]. URL: http://arxiv.org/abs/1801.03313v1. [BF81] B. Berg and D. Foerster. "Random paths and random surfaces on a digital computer". In: *Physics Letters B* 106.4 (Nov. 1981), pp. 323–326. ISSN: 0370-2693. DOI: 10.1016/0370-2693(81)90545-1. [BZ04] G. R. Buck and E. L. Zechiedrich. "DNA Disentangling By Type-2 Topoisomerases". In: Journal of Molecular Biology 340.5 (2004), pp. 933–939. DOI: 10.1016/j.jmb.2004.05.034. [CCF83]C. de Carvalho, S. Caracciolo, and J. Fröhlich. "Polymers and $g|\phi|^4$ theory in four dimensions". In: Nuclear Physics B 215.2 (Jan. 1983), pp. 209–248. ISSN: 0550-3213. DOI: 10.1016/0550-3213(83)90213-4. [CCM16] J. Cantarella, H. Chapman, and M. Mastin. "Knot probabilities in random diagrams". In: Journal of Physics A: Mathematical and Theoretical 49.40 (2016), p. 405001. DOI: 10.1088/1751-8113/49/40/405001. [CE17] H.-C. Chang and J. Erickson. "Untangling Planar Curves". In: Discrete & Computational Geometry 58.4 (July 2017), pp. 889–920. ISSN: 1432-0444. DOI: 10.1007/s00454-017-9907-6. A. Conway and A. J. Guttmann. "Square lattice self-avoiding walks and [CG96]corrections to scaling". In: Physical Review Letters 77.26 (1996), p. 5284. [Cha17]H. Chapman. "Asymptotic laws for random knot diagrams". In: Journal of Physics A: Mathematical and Theoretical (2017). DOI: 10.1088/1751-8121/aa6e45. arXiv: 1608.02638 [math.GT]. [Cha18]H. Chapman. "Slipknotting in the knot diagram model". In: (Mar. 2018). Submitted. arXiv: 1803.07114. [Cul+17]M. Culler et al. SnapPy, a computer program for studying the geometry and
 - topology of 3-manifolds. Available at http://snappy.computop.org. 2017.

[CZ16]	R. Coquereaux and JB. Zuber. "Maps, Immersions and Permutations". In: J. Knot Theory Ramifications 25.08 (2016), p. 1650047. DOI:
[Del62]	 10.1142/s0218216516500474. M. Delbrück. "Knotting Problems in Biology". In: Mathematical Problems in the Biological Sciences. Ed. by R. Bellman. American Mathematical Society, 1062. Chan 5, no. 55, 62, part 10, 1000 (again (014)).
[DEZ12]	 Y. Diao, C. Ernst, and U. Ziegler. "Generating Large Random Knot Projections". In: <i>Physical and Numerical Models in Knot Theory</i>. World Scientific. 2012. Chap. 22, pp. 473–404, poi: 10, 1142/0720212702460, 0022
[Dia95]	Y. Diao. "The Knotting of Equilateral Polygons in \mathbb{R}^3 ". In: Journal of Knot Theory and Its Ramifications 04.02 (1995), pp. 189–196. DOI: 10. 1142/S0218216595000090
[DP13]	M. Drmota and K. Panagiotou. "A Central Limit Theorem for the Number of Degree-k Vertices in Random Maps". In: <i>Algorithmica</i> 66.4 (Feb. 2013), pp. 741–761. ISSN: 1432-0541. DOI: 10.1007/s00453-013-9751-x
[DT97]	T. Deguchi and K. Tsurusaki. "Universality of random knotting". In: <i>Physical Review E</i> 55.5 (May 1997), pp. 6245–6248. ISSN: 1095-3787. DOI: 10.1103/physreve. 55.6245
[Dun+14]	N. Dunfield et al. Random Knots: A preliminary report. 2014. URL:
[EM91]	B. Ewing and K. C. Millett. "A Load Balanced Algorithm for the Calculation of the Polynomial Knot and Link Invariants". In: <i>The Mathematical Heritage of</i>
[EM97]	 C.F. Gauss (Sept. 1991), pp. 225–266. DOI: 10.1142/9789814503457_0017. B. Ewing and K. C. Millett. "Computational algorithms and the complexity of link polynomials". In: Progress in knot theory and related topics. Vol. 56. Travaux en Cours Hermann Paris 1997 pp. 51–68.
[Eve+16]	C. Even-Zohar et al. "Invariants of Random Knots and Links". In: Discrete & Computational Geometry 56.2 (2016), pp. 274–314. DOI: 10.1007/s00454-016-9798-y.
[Fre+85]	P. Freyd et al. "A New Polynomial Invariant of Knots and Links". In: <i>Bulletin of the American Mathematical Society</i> 12.2 (1985), pp. 239–247. DOI: 10.1090/s0273-0979-1985-15361-3.
[FS09]	P. Flajolet and R. Sedgewick. <i>Analytic Combinatorics</i> . Cambridge University Press, 2009.
[FW61]	H. L. Frisch and E. Wasserman. "Chemical Topology 1". In: <i>Journal of the American Chemical Society</i> 83.18 (Sept. 1961), pp. 3789–3795. DOI: 10.1021/jc014702015
[Gey91]	C. J. Geyer. "Markov Chain Monte Carlo Maximum Likelihood, Computing Science and Statistics". In: <i>Proceedings of the 23rd Symposium Interface, 1991.</i> 1991.
[GO99]	E. Guitter and E. Orlandini. "Monte Carlo Results for Projected Self-Avoiding Polygons: a Two-Dimensional Model for Knotted Polymers". In: <i>Journal of</i> <i>Physics A: Mathematical and General</i> 32.8 (1999), pp. 1359–1385. DOI: 10.1088/0305-4470/32/8/006.

36	REFERENCES
[GR94]	Z. Gao and L. Richmond. "Root Vertex Valency Distributions of Rooted Maps and Rooted Triangulations". In: <i>European Journal of Combinatorics</i> 15.5 (Sept.
[GS97]	1994), pp. 483–490. ISSN: 0195-6698. DOI: 10.1006/eujc.1994.1050. M. de Graaf and A. Schrijver. "Making Curves Minimally Crossing by Reidemeister Moves". In: <i>Journal of Combinatorial Theory, Series B</i> 70.1 (May
[GW00]	1997), pp. 134–156. ISSN: 0095-8956. DOI: 10.1006/jctb.1997.1754. Z. Gao and N. C. Wormald. "The Distribution of the Maximum Vertex Degree in Random Planar Maps". In: <i>Journal of Combinatorial Theory, Series A</i> 89.2 (Feb. 2000), pp. 201–230, ISSN: 0007-3165, DOI: 10.1006/jctb.1999.3006
[Ham61]	J. M. Hammersley. "The number of polygons on a lattice". In: <i>Proceedings of the Cambridge Philosophical Society</i> 57 (1961), p. 516. DOI: 10.1017/S030500410003557X
[HL01]	J. Hass and J. Lagarias. "The number of Reidemeister moves needed for unknotting". In: Journal of the American Mathematical Society 14.2 (2001), pp. 399–428
[HN10]	A. Henrich and S. Nelson. "Semiquandles and flat virtual knots". In: <i>Pacific Journal of Mathematics</i> 248.1 (2010), pp. 155–170. DOI: 10.2140/pim.2010.248.155
[HS94]	J. Hass and P. Scott. "Shortening curves on surfaces". In: <i>Topology</i> 33.1 (Jan. 1994), pp. 25–43. ISSN: 0040-9383. DOI: 10.1016/0040-9383(94)90033-7.
[Kan86]	T. Kanenobu. "Infinitely many knots with the same polynomial invariant". In: Proceedings of the American Mathematical Society 97.1 (Jan. 1986), pp. 158–158. ISSN: 0002-9939. DOI: 10.1090/s0002-9939-1986-0831406-7.
[Kau87]	L. H. Kauffman. "State Models and the Jones Polynomial". In: <i>Topology</i> 26.3 (1987), pp. 395–407. DOI: 10.1016/0040-9383(87)90009-7.
[KL06]	L. H. Kauffman and S. Lambropoulou. "Hard unknots and collapsing tangles". In: Introductory Lectures on Knot Theory—Selected Lectures presented at the Advanced School and Conference on Knot Theory and its Applications to Physics and Biology ICTP, Trieste, Italy, 11–29 May 2009. 2006.
[KM06]	L. H. Kauffman and V. O. Manturov. "Virtual knots and links". In: <i>Proceedings</i> of the Steklov Institute of Mathematics 252.1 (Jan. 2006), pp. 104–121. ISSN: 1531-8605. DOI: 10.1134/s0081543806010111.
[Lac15]	M. Lackenby. "A polynomial upper bound on Reidemeister moves". In: Annals of Mathematics 182.2 (2015), pp. 491–564.
[Lal69]	M. Lal. "Monte Carlo computer simulation of chain molecules. I". In: <i>Molecular</i> physics 17.1 (1969), pp. 57–64.
[Lis99]	V. A. Liskovets. "A Pattern of Asymptotic Vertex Valency Distributions in Planar Maps". In: <i>Journal of Combinatorial Theory, Series B</i> 75.1 (1999), pp. 116–133. DOI: 10.1006/jctb.1998.1870.
[Met+53]	N. Metropolis et al. "Equation of State Calculations by Fast Computing Machines". In: <i>The Journal of Chemical Physics</i> 21.6 (June 1953), pp. 1087–1092 JSSN: 1080–7600 pol: 10.1063/1.1699114
[Mil10]	K. C. Millett. "Knots, Slipknots, and Ephemeral Knots In Random Walks and Equilateral Polygons". In: <i>Journal of Knot Theory and Its Ramifications</i> 19.05 (2010), pp. 601–615. DOI: 10.1142/S0218216510008078.

REFERENCES

[MS13]	N. Madras and G. Slade. <i>The Self-Avoiding Walk</i> . Probability and Its
	Applications. Birkhäuser Boston, 2013. ISBN: 9781461241324. URL:
5	https://books.google.com/books?id=JsoFCAAAQBAJ.
[MS88]	N. Madras and A. D. Sokal. "The pivot algorithm: a highly efficient Monte
	Carlo method for the self-avoiding walk". In: Journal of Statistical Physics
_	50.1-2 (1988), pp. 109–186.
[Now09]	T. Nowik. "Complexity of plane and spherical curves". In: Duke Mathematical
	Journal 148.1 (May 2009), pp. 107–118. DOI: 10.1215/00127094–2009–022.
[Obe16]	M. Obeidin. "Volumes of Random Alternating Link Diagrams". In: (2016).
$\left[\bigcirc \mathbf{r} \right] + \bigcirc 8$	F Orlandini et al "Asymptotics of Knotted Lattice Polygons" In: <i>Journal of</i>
[011+90]	<i>Physics A: Mathematical and General</i> 31.28 (1998), pp. 5953–5967. DOI:
	10.1088/0305-4470/31/28/010.
[Orl98]	E. Orlandini. "Monte Carlo study of polymer systems by multiple Markov
[]	chain method". In: Numerical Methods for Polymeric Systems. Springer, 1998.
	pp. 33–57.
[OW07]	E. Orlandini and S. G. Whittington. "Statistical Topology of Closed Curves:
	Some Applications in Polymer Physics". In: Reviews of Modern Physics 79.2
	(2007), pp. 611-642. DOI: 10.1103/revmodphys.79.611.
[Pip89]	N. Pippenger. "Knots in Random Walks". In: Discrete Applied Mathematics
	25.3 (1989), pp. 273–278. DOI: 10.1016/0166-218x(89)90005-x.
[Rei48]	K. Reidemeister. <i>Knotentheorie</i> . Ergebnisse der Mathematik und Ihrer
	Grenzgebiete. Chelsea Pub. Co., 1948. URL:
	https://books.google.com/books?id=W8s-AAAAIAAJ.
[RR11]	E. J. J. van Rensburg and A. Rechnitzer. "On the universality of knot
	probability ratios". In: Journal of Physics A: Mathematical and Theoretical
	44.16 (2011), p. 162002. URL:
	http://stacks.iop.org/1751-8121/44/i=16/a=162002.
[RW91a]	E. J. J. v. Rensburg and S. G. Whittington. "The dimensions of knotted
	polygons". In: Journal of Physics A: Mathematical and General 24.16 (Aug.
	1991), pp. 3935–3948. ISSN: 1361-6447. DOI: 10.1088/0305-4470/24/16/028.
[RW91b]	E. J. J. van Rensburg and S. G. Whittington. "The BFACF Algorithm and
	Knotted Polygons". In: Journal of Physics A: Mathematical and General 24.23
5 m 1	(1991), pp. 5553–5567. DOI: 10.1088/0305-4470/24/23/021.
[Sch97]	G. Schaeffer. "Bijective census and random generation of Eulerian planar maps
	with prescribed vertex degrees." eng. In: The Electronic Journal of
	Combinatorics 4.1 (1997), Research paper R20, Research paper R20, 14 p. URL:
	http://eudml.org/doc/119255.
[Sch99a]	G. Schaeffer. <i>PlanarMap</i> . Available at http:
	//www.lix.polytechnique.fr/~schaeffe/PagesWeb/PlanarMap/index-
	en.html. 1999.
[Sch99b]	G. Schaetter. "Random Sampling of Large Planar Maps and Convex
	Polyhedra". In: Proceedings of the Thirty-first Annual ACM Symposium on
	Theory of Computing. STOC '99. Atlanta, Georgia, USA: ACM, 1999,
	pp. 760–769. ISBN: 1-58113-067-8. DOI: 10.1145/301250.301448.

00	
[ST98]	C. Sundberg and M. Thistlethwaite. "The rate of growth of the number of prime alternating links and tangles". In: <i>Pacific Journal of Mathematics</i> 182.2 (Feb. 1998), pp. 329–358, ISSN: 0030-8730, DOI: 10.2140/pim.1998.182.329.
[SW88]	D. W. Sumners and S. G. Whittington. "Knots in self-avoiding walks". In: Journal of Physics A: Mathematical and General 21.7 (Apr. 1988),
[SZ04]	 pp. 1689–1694. DOI: 10.1088/0305-4470/21/7/030. G. Schaeffer and P. Zinn-Justin. "On the asymptotic number of plane curves and alternating knots". In: <i>Experimental Mathematics</i> 13.4 (2004), pp. 483–493 JSSN: 1058-6458
[Thi98]	M. Thistlethwaite. "On The Structure and Scarcity of Alternating Links and Tangles". In: Journal of Knot Theory and Its Ramifications 07.07 (Nov. 1998), pp. 981–1004. ISSN: 1793-6527. poi: 10.1142/s021821659800053x
[Tri+01]	S. Trigueros et al. "Novel display of knotted DNA molecules by two-dimensional gel electrophoresis". In: <i>Nucleic Acids Research</i> 29.13 (2001), e67–e67.
[Val16]	G. Valette. "A Classification of Spherical Curves Based on Gauss Diagrams". In: Arnold Mathematical Journal (2016), pp. 1–23. ISSN: 2199-6806. DOI: 10.1007/s40598-016-0049-3.
[Van95]	C. Vanderzande. "On knots in a model for the adsorption of ring polymers". In: Journal of Physics A: Mathematical and General 28.13 (1995), p. 3681.
[WL01]	F. Wang and D. P. Landau. "Efficient, Multiple-Range Random Walk Algorithm To Calculate the Density of States". In: <i>Physical Review Letters</i> 86.10 (2001), pp. 2050–2053. DOI: 10.1103/physrevlett.86.2050.
[ZZ02]	P. Zinn-Justin and JB. Zuber. "Matrix integrals and the counting of tangles and links". In: <i>Discrete Mathematics</i> 246.1-3 (Mar. 2002), pp. 343–360. ISSN: 0012-365X. DOI: 10.1016/s0012-365x(01)00267-9.
[ZZ09]	P. Zinn-Justin and JB. Zuber. "Knot theory and matrix integrals". In: Random Matrix Theory. Eds Akemann, Baik and Di Francesco. 2009.

REFERENCES

APPENDIX A. COMPLETE DETAILED BALANCE PROOFS

Detailed balance equations for Boltzmann Markov chain; Theorem 1. We check the detailed balance equations for each transition:

(1) Suppose that $N = \mathrm{RI}^+(D, a)$ with root flag b. This means that N is unique in that $D = \mathrm{RI}^-(N, b)$. Then

$$\begin{split} P(D \rightarrow N) \pi(D) &= P(N \rightarrow D) \pi(N) \\ \frac{z p_1}{2} z^n &= \frac{p_1}{2} z^{n+1}, \end{split}$$

so the equation holds.

(2) Suppose that $N = \operatorname{RII}^+(D, a, a')$ with root flag b. This means that N is unique in that $D = \operatorname{RII}^-(N, b)$. The flags a, a' lie along a face in D of degree d. Then

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$\frac{z^2 p_2}{2(d-1)} z^n = \frac{p_2}{2(d-1)} z^{n+2}.$$

(3) Suppose that $N \neq \text{RIII}(D, a)$ and that N is a re-rooting of D. Then,

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$(1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} D}{4n}z^n = (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} N}{4n}z^n.$$

Because D and N differ only by a re-rooting, their underlying number of automorphisms are the same; aut D = aut N. Hence equality follows.

(4) Suppose that N = RIII(D, a) has root b and that N is not a re-rooting of D. Then N is unique in that D = RIII(N, b), so

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$z^{n}p_{3} = z^{n}p_{3}.$$

(5) If N = RIII(D, a) has root b and N is a re-rooting of D, then the transition probabilities of the previous two cases are summed (as the different transitions are independent), so that

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$

$$\left(p_3 + (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} D}{4n}\right)z^n = \left(p_3 + (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} N}{4n}\right)z^n. \qquad \Box$$

Proof of Corollary 6. As this Markov chain can perform all flat Reidemeister transitions and achieve all curve rootings, this Markov chain explores the space of curves. It remains to show that detailed balance holds.

For a pair of diagrams D, N, let $P_p(D \to N)$ denote the probability of transitioning from D to N under this modified Markov chain. Let $P'_p(D \to N)$ be the probability of transitioning from D to N under the original Markov chain (no interstitial re-rooting). Finally, let $P'_0(D \to N)$ be the probability of transitioning from D to N under the original Markov chain (no interstitial re-rooting). Finally, let $P'_0(D \to N)$ be the probability of transitioning from D to N under the original Markov chain (no interstitial re-rooting). Finally, let $P'_0(D \to N)$ be the probability of transitioning from D to N under the original Markov chain with all $p_i = 0$ (only re-roots are performed).

Notice that $P_p(D \to N) = \sum_B \sum_C P'_0(C \to N) P'_p(B \to C) P'_0(D \to B)$, where the sums are over all rooted curves. Then,

(25a)
$$P_p(D \to N)\pi(D) = \sum_B \sum_C P'_0(C \to N)P'_p(B \to C)P'_0(D \to B)\pi(D)$$

(25b)
$$= \sum_{B} \sum_{C} P'_0(C \to N) P'_p(B \to C) P'_0(B \to D) \pi(B)$$

(25c)
$$= \sum_{B} \sum_{C} P'_0(C \to N) P'_p(C \to B) P'_0(B \to D) \pi(C)$$

(25d)
$$= \sum_{B} \sum_{C} P'_{0}(N \to C) P'_{p}(C \to B) P'_{0}(B \to D) \pi(N)$$

(25e)
$$= P_p(N \to D)\pi(N),$$

so detailed balance holds for the modified Markov chain, and hence it is ergodic.

Detailed balance equations for Wang-Landau Markov chain; Theorem 2. The main concern now is that transitions must pass an additional Metropolis-Hastings check of min $\{1, g_n/g_m\}$. Note that for any $g_n, g_m > 0$,

(26)
$$\frac{\min\{1, g_n/g_m\}}{\min\{1, g_m/g_n\}} = \frac{g_n}{g_m}$$

REFERENCES

We check the detailed balance equations for each transition:

(1) Suppose that $N = \mathrm{RI}^+(D, a)$ with root flag b. This means that N is unique in that $D = \mathrm{RI}^-(N, b)$. Then

$$\begin{split} P(D \to N)\pi(D) &= P(N \to D)\pi(N) \\ \min\left\{1, \frac{g_n}{g_{n+1}}\right\} \frac{p_1}{2} \frac{1}{Rg_n} = \min\left\{1, \frac{g_{n+1}}{g_n}\right\} \frac{p_1}{2} \frac{1}{Rg_{n+1}} \\ \frac{\min\left\{1, g_n/g_{n+1}\right\}}{\min\left\{1, g_{n+1}/g_n\right\}} \frac{p_1}{2} \frac{1}{Rg_n} = \frac{p_1}{2} \frac{1}{Rg_{n+1}} \\ \frac{g_n}{g_{n+1}} \frac{p_1}{2} \frac{1}{Rg_n} = \frac{p_1}{2} \frac{1}{Rg_{n+1}}, \end{split}$$

so the equation holds.

(2) Suppose that $N = \operatorname{RII}^+(D, a, a')$ with root flag b. This means that N is unique in that $D = \operatorname{RII}^-(N, b)$. The flags a, a' lie along a face in D of degree d. Then

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$\max\left\{1, \frac{g_n}{g_{n+2}}\right\} \frac{p_2}{2(d-1)} \frac{1}{Rg_n} = \max\left\{1, \frac{g_{n+2}}{g_n}\right\} \frac{p_2}{2(d-1)} \frac{1}{Rg_{n+2}}$$
$$\frac{g_n}{g_{n+2}} \frac{p_2}{2(d-1)} \frac{1}{Rg_n} = \frac{p_2}{2(d-1)} \frac{1}{Rg_{n+2}}.$$

(3) Suppose that $N \neq \text{RIII}(D, a)$ and that N is a re-rooting of D. Then,

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$(1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} D}{4n}\frac{1}{Rg_n} = (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} N}{4n}\frac{1}{Rg_n}$$

Because D and N differ only by a re-rooting, their underlying number of automorphisms are the same; aut D = aut N. Hence equality follows.

(4) Suppose that N = RIII(D, a) has root b and that N is not a re-rooting of D. Then N is unique in that D = RIII(N, b), so

$$\begin{split} P(D \to N) \pi(D) &= P(N \to D) \pi(N) \\ \frac{p_3}{Rg_n} &= \frac{p_3}{Rg_n}. \end{split}$$

(5) If N = RIII(D, a) has root b and N is a re-rooting of D, then the transition probabilities of the previous two cases are summed (as the different transitions are independent), so that

$$P(D \to N)\pi(D) = P(N \to D)\pi(N)$$
$$\left(p_3 + (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} D}{4n}\right)\frac{1}{Rg_n} = \left(p_3 + (1 - (p_1 + p_2 + p_3))\frac{\operatorname{aut} N}{4n}\right)\frac{1}{Rg_n}.$$

In all other cases, the transition probabilities are symmetrically zero. Hence we conclude that detailed balance holds with the hypothesized probability distribution. \Box

E-mail address: hchaps@gmail.com

REFERENCES

DEPARTMENT OF MATHEMATICS, COLORADO STATE UNIVERSITY, FORT COLLINS CO

E-mail address: andrewr@math.ubc.ca

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA, VANCOUVER BC