

The Symplectic Geometry of Closed Equilateral Random Walks in 3-Space

Jason Cantarella* and Clayton Shonkwiler*

(Dated: October 22, 2013)

A closed equilateral random walk in 3-space is a selection of unit length vectors giving the steps of the walk conditioned on the assumption that the sum of the vectors is zero. The sample space of such walks with n edges is the $(2n - 3)$ -dimensional Riemannian manifold of equilateral closed polygons in \mathbb{R}^3 . We study closed random walks using the symplectic geometry of the $(2n - 6)$ -dimensional quotient of the manifold of polygons by the action of the rotation group $SO(3)$.

The basic objects of study are the moment maps on equilateral random polygon space given by the lengths of any $(n - 3)$ -tuple of non-intersecting diagonals. The Atiyah-Guillemin-Sternberg theorem shows that the image of such a moment map is a convex polytope in $(n - 3)$ -dimensional space, while the Duistermaat-Heckman theorem shows that the pushforward measure on this polytope is Lebesgue measure on \mathbb{R}^{n-3} . Together, these theorems allow us to define a measure-preserving set of “action-angle” coordinates on the space of closed equilateral polygons. The new coordinate system allows us to make explicit computations of exact expectations for total curvature and for some chord lengths of closed (and confined) equilateral random walks, to give statistical criteria for sampling algorithms on the space of polygons, and to prove that the probability that a randomly chosen equilateral hexagon is unknotted is at least $3/4$.

We then use our methods to construct a new Markov chain sampling algorithm for equilateral closed polygons, with a simple modification to sample (rooted) confined equilateral closed polygons. We prove rigorously that our algorithm converges geometrically to the standard measure on the space of closed random walks, give a theory of error estimators for Markov chain Monte Carlo integration using our method, and analyze the performance of our method. Our methods also apply to open random walks in certain types of confinement, and in general to walks with arbitrary (fixed) edge lengths as well as equilateral walks.

Keywords: closed random walk; statistics on Riemannian manifolds; Duistermaat-Heckman theorem; random knot; random polygon; crankshaft algorithm

1. INTRODUCTION

In this paper, we consider the classical model of a random walk in \mathbb{R}^3 —the walker chooses each step uniformly from the unit sphere. Some of the first results in the theory of these random walks are based on the observation that if a point is distributed uniformly on the surface of a sphere in 3-space and we write its position in terms of the cylindrical coordinates z and θ , then z and θ are independent, uniform random variates. This is usually called Archimedes’ Theorem, and it is the underlying idea in the work of Lord Rayleigh [59], L.R.G. Treloar [68], and many others in the

*University of Georgia, Mathematics Department, Athens GA

theory of random walks, starting at the beginning of the 20th century. In particular, it means that the vector of z -coordinates of the edges of a random walk is uniformly distributed on a hypercube and that the vector of θ -coordinates of the edges is uniformly distributed on the n -torus.

When we condition the walk on closure, it seems that this pleasant structure disappears: the individual steps in the walk are no longer independent random variates, and there are no obvious uniformly distributed random angles or distances in sight. This makes the study of closed random walks considerably more difficult than the study of general random walks. The main point of this paper is that the apparent disappearance of this structure in the case of closed random walks is only an illusion. In fact there is a very similar structure on the space of closed random walks if we are willing to pay the modest price of identifying walks related by translation and rigid rotation in \mathbb{R}^3 . This structure is less obvious, but just as useful.

As it turns out, Archimedes' Theorem was generalized in deep and interesting ways in the later years of the 20th century, being revealed as a special case of the Duistermaat-Heckman theorem [25] for toric symplectic manifolds. Further, Kapovich and Millson [37] and Hausmann and Knutson [31] revealed a toric symplectic structure on the quotient of the space of closed equilateral polygons by the action of the Euclidean group $E(3)$. Together, these theorems define a structure on closed random walk space which is remarkably similar to the structure on the space of open random walks: if we view an n -edge closed equilateral walk as the boundary of a triangulated surface, we will show below that the lengths of the $n - 3$ diagonals of the triangulation are uniformly distributed on the polytope given by the triangle inequalities and that the $n - 3$ dihedral angles at these diagonals of the triangulated surface are distributed uniformly and independently on the $(n - 3)$ -torus. This structure allows us to define a special set of "action-angle" coordinates which provide a measure-preserving map from the product of a convex polytope $P \subset \mathbb{R}^{n-3}$ and the $(n - 3)$ -torus (again, with their standard measures) to a full-measure subset of the Riemannian manifold of closed polygons of fixed edge lengths.

Understanding this picture allows us to make some new explicit calculations and prove some new theorems about closed equilateral random walks. For instance, we are able to find an exact formula for the total curvature of closed equilateral polygons, to prove that the expected lengths of chords skipping various numbers of edges are equal to the coordinates of the center of mass of a certain polytope, to compute these moments explicitly for random walks with small numbers of edges, and to give a simple proof that at least $3/4$ of equilateral hexagons are unknotted. Further, we will be able to give a unified theory of several interesting problems about confined random walks, and to provide some explicit computations of chordlengths for confined walks. We state upfront that all the methods we use from symplectic geometry are by now entirely standard; the new contribution of our paper lies in the application of these powerful tools to geometric probability.

We will then turn to sampling for the second half of our paper. Our theory immediately suggests a new Markov chain sampling algorithm for confined and unconfined random walks. We will show that the theory of hit-and-run sampling on convex polytopes immediately yields a sampling algorithm which converges at a geometric rate to the usual probability measure on equilateral

closed random walks (or equilateral closed random walks in confinement). Geometric convergence allows us to apply standard Markov Chain Monte Carlo theory to give error estimators for MCMC integration over the space of closed equilateral random walks (either confined or unconfined). Our sampling algorithm works for any toric symplectic manifold, so we state the results in general terms. We do this primarily because various interesting confinement models for random walks have a natural toric symplectic structure, though our results are presumably applicable far outside the theory of random walks. As with the tools we use from symplectic geometry, hit-and-run sampling and MCMC error estimators are entirely standard ways to integrate over convex polytopes. Again, our main contribution is to show that these powerful tools apply to closed and confined random walks with fixed edgelengths and to lay out some initial results which follow from their use.

2. TORIC SYMPLECTIC MANIFOLDS AND ACTION-ANGLE COORDINATES

We begin with a capsule summary of some relevant ideas from symplectic geometry. A symplectic manifold M is an even-dimensional manifold with a special nondegenerate 2-form ω called the *symplectic form*. The volume form $dm = \frac{1}{n!}\omega^n$ on M is called the *symplectic volume* or *Liouville volume* and the corresponding measure is called *symplectic measure*. A diffeomorphism of a symplectic manifold which preserves the symplectic form is called a *symplectomorphism*; it must preserve symplectic volume as well. An action of the torus T^k on M by symplectomorphisms defines a *moment map* $\mu : M \rightarrow \mathbb{R}^k$, where the torus action preserves the fibers (inverse images of points) of the map. If the action obeys some mild additional technical hypotheses, it is called *Hamiltonian*. Two powerful theorems apply in this case: first, the convexity theorem of Atiyah [3] and Guillemin–Sternberg [30] asserts that the image of the moment map is a convex polytope in \mathbb{R}^k called the *moment polytope*. The fibers over vertices of the moment polytope consist of fixed points of the torus action. Further, if the action is effective, the Duistermaat–Heckman theorem [25] asserts that the pushforward of symplectic measure to the moment polytope is a piecewise polynomial multiple of Lebesgue measure. If $k = n$ the symplectic manifold is called a *toric symplectic manifold* and the pushforward measure on the moment polytope is a constant multiple of Lebesgue measure.

If we can invert the moment map, we can construct a map $\alpha : P \times T^n \rightarrow M^{2n}$ compatible with μ which parametrizes a full-measure subset of M^{2n} by the n coordinates of points in P , which are called the “action” variables, and the n angles in T^n , which are called the corresponding “angle” variables. By convention, we call the action variables d_i and the angle variables θ_i . We have

Theorem 1 (Duistermaat–Heckman [25]). *Suppose M^{2n} is a toric symplectic manifold with moment polytope P , T^n is the n -torus (n copies of the circle) and α inverts the moment map. If we take the standard measure on the n -torus and the uniform (or Lebesgue) measure on $\text{int}(P)$, then the map $\alpha : \text{int}(P) \times T^n \rightarrow M^{2n}$ parametrizing a full-measure subset of M^{2n} in action-angle*

coordinates is measure-preserving. In particular, if $f: M^{2n} \rightarrow \mathbb{R}$ is any integrable function then

$$\int_M f(x) dm = \int_{P \times T^n} f(d_1, \dots, d_n, \theta_1, \dots, \theta_n) d\text{Vol}_{E^n} \wedge d\theta_1 \wedge \dots \wedge d\theta_n \quad (1)$$

and if $f(d_1, \dots, d_n, \theta_1, \dots, \theta_n) = f_d(d_1, \dots, d_n) f_\theta(\theta_1, \dots, \theta_n)$ then

$$\int_M f(x) dm = \int_P f_d(d_1, \dots, d_n) d\text{Vol}_{E^n} \int_{T^n} f_\theta(\theta_1, \dots, \theta_n) d\theta_1 \wedge \dots \wedge d\theta_n. \quad (2)$$

All this seems forbiddingly abstract, so we give a specific example which will prove important below. The 2-sphere is a symplectic manifold where the symplectic form ω is the ordinary area form, and the symplectic volume and the Riemannian volume are the same. Any area-preserving map of the sphere to itself is a symplectomorphism, but we are interested in the action of the circle (or 1-torus) on the sphere given by rotation around the z -axis. This action is by area-preserving maps, and hence by symplectomorphisms, and in fact it is Hamiltonian. We can see that the action preserves the fibers of the moment map, which are just horizontal circles on the sphere. Since the dimension of the torus (1) is half the dimension of the sphere (2), the sphere is then a toric symplectic manifold. The image of the moment map is the closed interval $[-1, 1] \subset \mathbb{R}$, which is certainly a convex polytope. And as the Duistermaat-Heckman theorem claims, *the pushforward of Lebesgue measure on the sphere to this interval is a constant multiple of the Lebesgue measure on the line*. This, of course, is just Archimedes' Theorem, restated in a very sophisticated form.

In particular, it means that one can sample points on the sphere uniformly by choosing their z and θ coordinates independently from uniform distributions on the interval and the circle. The Duistermaat-Heckman theorem extends a similar sampling strategy to any toric symplectic manifold. The best way to view this sampling strategy, we think, is as a useful technique in the theory of intrinsic statistics on Riemannian manifolds (cf. [56]) which applies to a special class of manifolds. In principle, one can sample the entirety of any Riemannian manifold by choosing charts for the manifold explicitly and then sampling appropriate measures on a randomly chosen chart. Since the charts are maps from balls in Euclidean space to the manifold, this reduces the problem to sampling a ball in \mathbb{R}^n with an appropriate measure. Of course, this point of view is so general as to be basically useless in practice: you rarely have explicit charts for a nontrivial manifold, and the resulting measures on Euclidean space could be very exotic and difficult to sample accurately.

Action-angle coordinates, however, give a single "chart" with a simple measure to sample – the product of Lebesgue measure on the convex moment polytope and the uniform measure on the torus. There is a small price to pay here. We cannot sample *all* of the toric symplectic manifold this way. The boundary of P corresponds to a sort of skeleton inside the toric symplectic manifold M , and we cannot sample this skeleton in any very simple way using action-angle coordinates. Of course, if we are using the Riemannian (or symplectic) volume of M to define the probability measure, this is a measure zero subset, so it is irrelevant to theorems in probability. The benefit is that by deleting this skeleton, we remove most of the topology of M , leaving us with the topologically very simple sample space $P \times T^{n-3}$.

3. TORIC SYMPLECTIC STRUCTURE ON RANDOM WALKS OR POLYGONAL “ARMS”

We now consider the classical space of random walks of fixed step length in \mathbb{R}^3 and show that the arguments underlying the historical application of Archimedes’ theorem (e.g. in Rayleigh [59]) can be viewed as arguments about action-angle coordinates on this space as a toric symplectic manifold. We denote the space of open “arm” polygons with n edges of lengths $\vec{r} = (r_1, \dots, r_n)$ in \mathbb{R}^3 by $\text{Arm}_3(n; \vec{r})$. In particular, the space of equilateral n -edge arms (with unit edges) is denoted $\text{Arm}_3(n; \vec{1})$. If we consider polygons related by a translation to be equivalent, the space $\text{Arm}_3(n; \vec{r})$ is a product $S^2(r_1) \times \dots \times S^2(r_n)$ of round 2-spheres with radii given by the r_i . The standard probability measure on this space is the product measure on these spheres; this corresponds to choosing n independent points distributed according to the uniform measure on S^2 to be the edge vectors of the polygon.

Proposition 2. *The space of fixed edgelenh open polygonal “arms” $\text{Arm}_3(n; \vec{r})$ is the product of n spheres of radii $\vec{r} = (r_1, \dots, r_n)$. This is a $2n$ -dimensional toric symplectic manifold where the Hamiltonian torus action is given by rotating each sphere about the z -axis, and the symplectic volume is the standard measure. The moment map $\mu: \text{Arm}_3(n; \vec{r}) \rightarrow \mathbb{R}^n$ is given by the z -coordinate of each edge vector, and the image of this map (the moment polytope) is the hyperbox $\Pi_{i=1}^n[-r_i, r_i]$. There is a measure-preserving map*

$$\alpha: \Pi_{i=1}^n[-r_i, r_i] \times T^n \rightarrow \text{Arm}_3(n; \vec{r}).$$

given explicitly by $\vec{e}_i = (\cos \theta_i \sqrt{1 - z_i^2}, \sin \theta_i \sqrt{1 - z_i^2}, z_i)$.

Proof. Recall that the moment polytope is the convex hull of the images of the fixed points of the Hamiltonian torus action. The only polygonal arms fixed by the torus action are those where every edge is in the $\pm z$ -direction, so the z -coordinates of the fixed points are indeed the vertices of the hyperbox $\Pi_{i=1}^n[-r_i, r_i]$ and the hyperbox itself is clearly the convex hull. The z -coordinates z_1, \dots, z_n and rotation angles $\theta_1, \dots, \theta_n$ are the action-angle coordinates on $\text{Arm}_3(n; \vec{r})$ and the fact that α is measure-preserving is an immediate consequence of Theorem 1. \square

Since we can sample $\Pi_{i=1}^n[-r_i, r_i] \times T^n$ directly, this gives a direct sampling algorithm for (a full-measure subset of) $\text{Arm}_3(n; \vec{r})$. Of course, direct sampling of fixed-edgelenh arms is straightforward even without symplectic geometry, but this description of arm space has additional implications for confinement problems: if we can describe a confinement model by additional linear constraints on the action variables, this automatically yields a toric symplectic structure on the space of confined arms. We give examples in the next two sections, then in Section 3.3 we use this machinery to provide a symplectic explanation for Rayleigh’s formula for the probability density function (pdf) of the distance between the endpoints of a random equilateral arm.

3.1. Slab-Confined arms

One system of linear constraints on the action variables of equilateral arms is the ‘slab’ confinement model:

Definition 3. Given a polygon p in \mathbb{R}^3 with vertices v_1, \dots, v_n , let $\text{zWidth}(p)$ be the maximum absolute value of the difference between z -coordinates of any two vertices. We define the subspace $\text{SlabArm}(n, h) \subset \text{Arm}_3(n; \vec{1})$ to be the space of equilateral (open) space n -gons up to translation which obey the constraint $\text{zWidth}(p) \leq h$.

This is a slab constraint model where the endpoints of the walk are free (one could also have a model where one or both endpoints are on the walls of the slab). We now rephrase this slab constraint in action-angle variables.

Proposition 4. A polygon p in $\text{Arm}_3(n; \vec{1})$ given in action-angle coordinates by $(z_1, \dots, z_n, \theta_1, \dots, \theta_n)$ lies in $\text{SlabArm}(n, h)$ if and only if the vector $\vec{z} = (z_1, \dots, z_n)$ of action variables lies in the parallelotope $P(n, h)$ given by the inequalities

$$-1 \leq z_i \leq 1, \quad -h \leq \sum_{k=i}^j z_k \leq h.$$

Hence, there is a measure-preserving map

$$\alpha : P(n, h) \times T^n \rightarrow \text{SlabArm}(n, h)$$

given by restricting the action-angle map of Proposition 2.

Proof. This follows directly from Definition 3: $\sum_{k=i}^j z_k$ is the difference in z -height between vertex i and j so this family of linear constraints encodes $\text{zWidth}(p) \leq h$. The other constraints just restate the condition that \vec{z} lies in the moment polytope $[-1, 1]^n$ for $\text{Arm}_3(n; \vec{1})$. \square

Corollary 5. The probability that a random polygon $p \in \text{Arm}_3(n; \vec{1})$ lies in $\text{SlabArm}(n, h)$ is given by $\text{Vol } P(n, h)/2^n$.

This probability function should be useful in computing the entropic force exerted by an ideal polymer on the walls of a confining slab. Figure 1 shows a collection of these moment polytopes for different slab widths, and the corresponding volumes.

3.2. Half-space confined arms

A similar problem is this: suppose we have a freely jointed chain which is attached at one end to a plane (which we assume for simplicity is the xy -plane), and must remain in the halfspace on




FIG. 1: This figure shows the moment polytopes corresponding to 3-edge arms contained in slabs of width h as subpolytopes of the cube with vertices $(\pm 1, \pm 1, \pm 1)$, which is the moment polytope for unconfined arms. In this case, we can compute the volume of these moment polytopes directly using `polymake` [27]. We conclude, for instance, that the probability that a random 3-edge arm is confined in a slab of width $1/2$ is $1/16$.

one side of the plane. This models a polymer where one end of the molecule is bound to a surface (at an unknown site). The moment polytope is

$$\mathcal{H}_n = \{\vec{z} \in [-1, 1]^n \mid z_1 \geq 0, z_1 + z_2 \geq 0, \dots, z_1 + \dots + z_n \geq 0, -1 \leq z_i \leq 1\} \quad (3)$$

and the analogue of Proposition 4 holds in this case.

We can understand this condition on arms in terms of a standard random walk problem: the z_i are i.i.d. steps in a random walk, each selected from the uniform distribution on $[-1, 1]$, and we are interested in conditioning on the event that all the partial sums are in $[0, \infty)$. A good deal is known about this problem: for instance, Caravenna gives an asymptotic pdf for the end of a random walk conditioned to stay positive, which is the height of the end of the chain above the plane [17]. If we could find an explicit form for this pdf, we could analyze the stretching experiment where the free end of the polymer is raised to a known height above the plane using magnetic or optical tweezers (cf. [65]).

We can directly compute the partition function for this problem; this is the volume of subpolytope (3) of the hypercube. This result is also stated in a paper of Bernardi, Duplantier, and Nadeau [7]. The proof is a pleasant combinatorial argument which is tangential to the rest of the paper, so we relegate it to Appendix B.

Proposition 6. *The volume of the polytope (3) is $\frac{1}{2^n} \binom{2n}{n} = \frac{(2n-1)!!}{n!}$.*

3.3. Distribution of failure to close lengths

We now apply the action-angle coordinates to give an alternate formula for the pdf of end-to-end distance in a random walk in \mathbb{R}^3 with fixed step lengths and show that it is equivalent to

Rayleigh's sinc integral formula [59]. This pdf is key to determining the Green's function for closed polygons, which in turn is fundamental to the Moore–Grosberg [52] and Diao–Ernst–Montemayor–Ziegler [22–24] sampling algorithms and to expected total curvature calculations [16, 29]. For mathematicians, we note that this pdf is required in order to estimate the entropic elastic force exerted by an ideal polymer whose ends are held at a fixed distance. Such experiments are actually done in biophysics – Wuite et al. [72] (cf. [13]) made one of the first measurements of the elasticity of DNA by stretching a strand of DNA between a bead held in a micropipette and a bead held in an optical trap.

We first establish some lemmas:

Lemma 7. *The pdf of a sum of independent uniform random variates in $[-r_1, r_1]$ to $[-r_n, r_n]$ is given by the pushforward of the Lebesgue measure on $\prod_{i=1}^n [-r_i, r_i]$ to $[-\sum r_i, \sum r_i]$ by the linear function $\sum x_i$. This is given by*

$$f_n(x) = \frac{1}{\prod_{i=1}^n 2r_i} \frac{1}{\sqrt{n}} \text{SA}(x, r_1, \dots, r_n), \quad (4)$$

where $\text{SA}(x, r_1, \dots, r_n)$ is the volume of the slice of the hypercube $\prod_{i=1}^n [-r_i, r_i]$ by the plane $\sum x_i = x$. The function f_n is everywhere $n - 2$ times differentiable for $n > 2$.

Proof. This is the coarea formula applied to the hyperbox $\prod_{i=1}^n [-r_i, r_i]$ and the function $\sum x_i$, normalized by the volume of $\prod_{i=1}^n [-r_i, r_i]$. To see that this function is $n - 2$ times differentiable, note that the pdf $f_2(x)$ has a corner at $x = 0$ (and is linear away from this corner) and that we can express $f_n(x)$ in terms of f_{n-1} by the convolution integral

$$f_n(x) = \int_{-r_n}^{r_n} f_{n-1}(x - y) \frac{1}{2r_n} dy.$$

where $1/2r_n$ is the pdf of the uniform random variate x_n on $[-r_n, r_n]$. □

We have

Proposition 8. *The pdf of the end-to-end distance $\ell \in [0, \sum r_i]$ over the space of polygonal arms $\text{Arm}_3(n; \vec{r})$ is given by*

$$\phi_n(\ell) = \frac{\ell}{2^{n-1} R \sqrt{n-1}} (\text{SA}(\ell - r_n, r_1, \dots, r_{n-1}) - \text{SA}(\ell + r_n, r_1, \dots, r_{n-1})).$$

where $R = \prod_{i=1}^n r_i$ is the product of the edgelengths and $\text{SA}(x, r_1, \dots, r_{n-1})$ is the volume of the slice of the hyperbox $\prod_{i=1}^{n-1} [-r_i, r_i]$ by the plane $\sum_{i=1}^{n-1} x_i = x$.

Proof. From our moment polytope picture, we can see immediately that the sum z of the z -coordinates of the edges of a random polygonal arm in $\text{Arm}_3(n; \vec{r})$ has the pdf of a sum of uniform

random variates in $[-r_1, r_1] \times \cdots \times [-r_n, r_n]$, or $f_n(z)$ in the notation of Lemma 7. Since this is a projection of the spherically symmetric distribution of end-to-end distance in \mathbb{R}^3 to the z -axis (\mathbb{R}^1), Equation 29 of [42] applies¹, and tells us that the pdf of ℓ is given by

$$\phi_n(\ell) = -2\ell f'_n(\ell).$$

To differentiate $f_n(\ell)$, we use the following observation (cf. Buonacore [12]):

$$f_n(x) = \int_{-r_n}^{r_n} f_{n-1}(x-y) \frac{1}{2r_n} dy = \frac{F_{n-1}(x+r_n) - F_{n-1}(x-r_n)}{2r_n}, \quad (5)$$

where $F_{n-1}(x)$ is the cdf of a sum of uniform random variates in $[-r_1, r_1] \times \cdots \times [-r_{n-1}, r_{n-1}]$. Differentiating and substituting in the results of Lemma 7 yields the formula above. \square

Since we will often be interested in equilateral polygons with edgelenhth 1, we observe

Corollary 9. *The pdf of the end-to-end distance $\ell \in [0, n]$ over the space of equilateral arms $\text{Arm}_3(n; \vec{1})$ is given by*

$$\phi_n(\ell) = \frac{\ell}{2^{n-1}\sqrt{n-1}} (\text{SA}(\ell-1, [-1, 1]^{n-1}) - \text{SA}(\ell+1, [-1, 1]^{n-1})). \quad (6)$$

where $\text{SA}(x, [-1, 1]^{n-1})$ is the volume of the slice of the standard hypercube $[-1, 1]^{n-1}$ by the plane $\sum_{i=1}^{n-1} x_i = x$.

The reader who is familiar with the theory of random walks may find the above corollary rather curious. As mentioned above, the standard formula for this pdf as an integral of sinc functions was given by Rayleigh in 1919 and it looks nothing like (6). The derivation given by Rayleigh of the sinc integral formula has no obvious connection to polyhedral volumes, but in fact by the time of Rayleigh's paper a connection between polyhedra and sinc integrals had already been given by George Pólya in his thesis [57, 58] in 1912. This formula has been rediscovered many times [10, 48]. First, we state the Rayleigh formula [23, 59] in our notation:

$$\phi_n(\ell) = \frac{2\ell}{\pi} \int_0^\infty y \sin \ell y \text{sinc}^n y dy, \quad (7)$$

where $\text{sinc } x = \sin x/x$ as usual. Now Pólya showed that the volume of the central slab of the hypercube $[-1, 1]^{n-1}$ given by $-a_0 \leq \sum x_i \leq a_0$ is given by

$$\text{Vol}(a_0) = \frac{2^n a_0}{\pi} \int_0^\infty \text{sinc } a_0 y \text{sinc}^{n-1} y dy. \quad (8)$$

¹ Lord's notation can be slightly confusing: in his formula for $p_3(r)$ in terms of $p_1(r)$, we have to remember that $p_3(r)$ is not itself a pdf on the line, it is a pdf on \mathbb{R}^3 . It only becomes a pdf on the line when multiplied by the correction factor $4\pi r^2$ giving the area of the sphere at radius r in \mathbb{R}^3 .

Our $\text{SA}(x, [-1, 1]^{n-1})$ is the $(n-1)$ -dimensional volume of a face of this slab; since it is this face (and its symmetric copy) which sweep out n -dimensional volume as a_0 increases, we can deduce that

$$\text{SA}(x, [-1, 1]^{n-1}) = \frac{\sqrt{n-1}}{2} \text{Vol}'(x),$$

and we can obtain a formula for $\text{SA}(x, [-1, 1]^{n-1})$ by differentiating (8). After some simplifications, we get

$$\text{SA}(x, [-1, 1]^{n-1}) = \frac{2^{n-1}\sqrt{n-1}}{\pi} \int_0^\infty \cos xy \operatorname{sinc}^{n-1} y \, dy.$$

Using the angle addition formula for $\cos(a+b)$, this implies that

$$\begin{aligned} \text{SA}(\ell-1, [-1, 1]^{n-1}) - \text{SA}(\ell+1, [-1, 1]^{n-1}) &= \frac{2^{n-1}\sqrt{n-1}}{\pi} \int_0^\infty 2 \sin y \sin \ell y \operatorname{sinc}^{n-1} y \, dy \\ &= \frac{2^n \sqrt{n-1}}{\pi} \int_0^\infty y \sin \ell y \operatorname{sinc}^n y \, dy. \end{aligned}$$

Multiplying by the coefficient $\frac{\ell}{2^{n-1}\sqrt{n-1}}$ from (6) recovers the Rayleigh form (7) of the pdf ϕ_n .

Given (6) and (7), the pdf of the failure-to-close vector $\vec{\ell} = \sum \vec{e}_i$ with length $|\vec{\ell}| = \ell$ can be written in the following forms:

$$\begin{aligned} \Phi_n(\vec{\ell}) &= \frac{1}{4\pi\ell^2} \phi_n(\ell) = \frac{1}{2^{n+1}\pi\ell\sqrt{n-1}} (\text{SA}(\ell-1, [-1, 1]^{n-1}) - \text{SA}(\ell+1, [-1, 1]^{n-1})) \\ &= \frac{1}{2\pi^2\ell} \int_0^\infty y \sin \ell y \operatorname{sinc}^n y \, dy \end{aligned} \quad (9)$$

The latter formula for the pdf appears in Grosberg and Moore [52] as equation (B5). Since Grosberg and Moore then actually evaluate the integral for the pdf as a finite sum, one immediately suspects that there is a similar sum form for the slice volume terms in (6). In fact, we have several options to choose from, including using Pólya's finite sum form to express (8) and then differentiating the sum formula with respect to the width of the slab. We instead rely on the following theorem, which we have translated to the current situation.

Theorem 10 (Marichal and Mossinghoff [48]). *Suppose that $\vec{w} \in \mathbb{R}^n$ has all nonzero components and suppose z is any real number. Then the $(n-1)$ -dimensional volume of the intersection of the hyperplane $\langle \vec{w}, \vec{v} \rangle = z$ with the hypercube $[-1, 1]^n$ is given by*

$$\text{Vol} = \frac{|\vec{w}|_2}{(n-1)! \prod w_i} \sum_{A \subset \{1, \dots, n\}} (-1)^{|A|} \left(z + \sum_{i \notin A} w_i - \sum_{i \in A} w_i \right)_+^{n-1}, \quad (10)$$

where $|\vec{w}|_2$ is the usual (L^2) norm of the vector \vec{w} , and $x_+ = \max(x, 0)$, and we use the convention $0^0 = 0$ when considering the $n = 1$ case.

For our $\text{SA}(x, [-1, 1]^{n-1})$ function, the vector \vec{w} consists of all 1's, and using the fact that the number of subsets of $\{1, \dots, n\}$ with cardinality k is $\binom{n}{k}$ it follows that

Proposition 11. *The $n - 2$ dimensional volume $\text{SA}(x, [-1, 1]^{n-1})$ is given by the finite sum*

$$\text{SA}(x, [-1, 1]^{n-1}) = \frac{\sqrt{n-1}}{(n-2)!} \sum_{k=0}^{n-1} (-1)^k \binom{n-1}{k} (x+n-1-2k)_+^{n-2}. \quad (11)$$

We can combine this with (9) to obtain the explicit piecewise polynomial pdf for the failure-to-close vector (for $n \geq 2$):

$$\Phi_n(\vec{\ell}) = \frac{n-1}{2^{n+1}\pi\ell} \sum_{k=0}^{n-1} \frac{(-1)^k}{k!(n-k-1)!} ((n+\ell-2k-2)_+^{n-2} - (n+\ell-2k)_+^{n-2}) \quad (12)$$

When $n = 2$, recall that we use the convention $0^0 = 0$. When $n = 1$ the formula does not make sense, but we can easily compute $\Phi_1(\vec{\ell}) = \frac{1}{4\pi}\delta(1-\ell)$. This formula for $\Phi_n(\ell)$ is known classically, and given as (2.181) in Hughes [35]. The polynomials are precisely those given in (B13) of Grosberg and Moore [52].

3.4. The Expected Total Curvature of Equilateral Polygons

In Section 5.4 it will be useful to know exact values of the expected total curvature of equilateral polygons. Let $\text{Pol}_3(n; \vec{1}) \subset \text{Arm}_3(n; \vec{1})$ be the subspace of closed equilateral n -gons. Following the approach of [16, 29], we can use the pdf above to find an integral formula for the expected total curvature of an element of $\text{Pol}_3(n; \vec{1})$:

Theorem 12. *The expected total curvature of an equilateral n -gon is equal to*

$$E(\kappa; \text{Pol}_3(n; \vec{1})) = \frac{n}{2C_n} \int_0^2 \arccos\left(\frac{\ell^2-2}{2}\right) \Phi_{n-2}(\ell)\ell \, d\ell, \quad (13)$$

where C_n and $\Phi_{n-2}(\ell)$ are given explicitly in (15) and (12), respectively, and Table II shows exact values of the integral for small n .

This integral can be evaluated easily by computer algebra using the fact that $\Phi_{n-2}(\ell)$ is a piecewise polynomial in ℓ in combination with the identity $\int_0^2 \arccos\left(\frac{\ell^2-2}{2}\right) \ell^k \, d\ell = \frac{2^{2k+1}n \text{B}(k/2+1, k/2)}{(k+1)^2}$, where B is the Euler beta function. Of course, it would be very interesting to find a closed form.

Proof. The total curvature of a polygon is just the sum of the turning angles, so the expected total curvature of an n -gon is simply n times the expected value of the turning angle $\theta(\vec{e}_i, \vec{e}_{i+1})$ between any pair $(\vec{e}_i, \vec{e}_{i+1})$ of consecutive edges. In other words,

$$E(\kappa; \text{Pol}_3(n; \vec{1})) = nE(\theta; \text{Pol}_3(n; \vec{1})) = n \int \theta(\vec{e}_i, \vec{e}_{i+1}) P(\vec{e}_i, \vec{e}_{i+1}) \, d\text{Vol}_{\vec{e}_i} \, d\text{Vol}_{\vec{e}_{i+1}}, \quad (14)$$

where $P(\vec{e}_i, \vec{e}_{i+1}) \, d\text{Vol}_{\vec{e}_i} \, d\text{Vol}_{\vec{e}_{i+1}}$ is the joint distribution of the pair of edges.

The edges \vec{e}_i, \vec{e}_{i+1} are chosen uniformly from the unit sphere subject to the constraint that the remaining $n - 2$ edges must connect the head of \vec{e}_{i+1} to the tail of \vec{e}_i . In other words,

$$P(\vec{e}_i, \vec{e}_{i+1}) \, d\text{Vol}_{\vec{e}_i} \, d\text{Vol}_{\vec{e}_{i+1}} = \frac{1}{C_n} \Phi_1(\vec{e}_i) \Phi_1(\vec{e}_{i+1}) \Phi_{n-2}(-\vec{e}_i - \vec{e}_{i+1}) \, d\text{Vol}_{\vec{e}_i} \, d\text{Vol}_{\vec{e}_{i+1}},$$

where

$$C_n = \Phi_n(\vec{0}) = \frac{1}{2^{n+1} \pi (n-3)!} \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^{k+1} \binom{n}{k} (n-2k)^{n-3} \quad (15)$$

is the normalized $(2n - 3)$ -dimensional Hausdorff measure of the submanifold of closed n -gons. Notice that $\Phi_1(\vec{v}) = \frac{\delta(|\vec{v}|-1)}{4\pi}$ is the distribution of a point chosen uniformly on the unit sphere. In particular, we can re-write the integral (14) as

$$E(\kappa; \text{Pol}_3(n; \vec{1})) = \frac{n}{C_n} \int_{\vec{e}_i \in S^2} \int_{\vec{e}_{i+1} \in S^2} \theta(\vec{e}_i, \vec{e}_{i+1}) \frac{1}{16\pi^2} \Phi_{n-2}(-\vec{e}_i - \vec{e}_{i+1}) \, d\text{Vol}_{S^2} \, d\text{Vol}_{S^2}.$$

Moreover, at the cost of a constant factor 4π we can integrate out the \vec{e}_i coordinate and assume \vec{e}_i points in the direction of the north pole. Similarly, at the cost of an additional 2π factor we can integrate out the azimuth angle of \vec{e}_{i+1} and reduce the above integral to a single integral over the polar angle of \vec{e}_{i+1} , which is now exactly the angle $\theta(\vec{e}_i, \vec{e}_{i+1})$:

$$E(\kappa; \text{Pol}_3(n; \vec{1})) = \frac{n}{2C_n} \int_0^\pi \theta \Phi_{n-2}(\sqrt{2-2\cos\theta}) \sin\theta \, d\theta$$

since $\sqrt{2-2\cos\theta}$ is the length of the vector $\vec{\ell} = -\vec{e}_i - \vec{e}_{i+1}$. Changing coordinates to integrate with respect to $\ell = |\vec{\ell}| \in [0, 2]$ completes the proof. \square

4. THE (ALMOST) TORIC SYMPLECTIC STRUCTURE ON CLOSED POLYGONS

We are now ready to describe explicitly the toric symplectic structure on closed polygons of fixed edge lengths. We first need to fix a bit of notation. The space $\text{Pol}_3(n; \vec{r})$ of closed polygons

of fixed edge lengths $\vec{r} = (r_1, \dots, r_n)$, where polygons related by translation are considered equivalent, is a subspace of the Riemannian manifold $\text{Arm}_3(n; \vec{r})$ (with the product metric on spheres of varying radii). It has a corresponding subspace metric and measure, which we refer to as the *standard measure* on $\text{Pol}_3(n; \vec{r})$. There is a measure-preserving action of $\text{SO}(3)$ on $\text{Pol}_3(n; \vec{r})$, and a corresponding quotient space $\widehat{\text{Pol}}_3(n; \vec{r}) = \text{Pol}_3(n; \vec{r}) / \text{SO}(3)$. This quotient space inherits a pushforward measure from the standard measure on $\text{Pol}_3(n; \vec{r})$, and we call this the standard measure on $\widehat{\text{Pol}}_3(n; \vec{r})$, which we will shortly see (almost) has a toric symplectic structure.

There are many ways to triangulate a convex n -gon; each triangulation T joins $n - 3$ pairs of vertices of the n -gon with chords to create a total of $n - 2$ triangles. We call these $n - 3$ chords the *diagonals* of T . The edgelengths and diagonal lengths of T must obey a set of $3(n - 2)$ triangle inequalities, which we call the *triangulation inequalities*.

Finally, given a diagonal of a space polygon, we can perform what the random polygons community calls a *polygonal fold* or *crankshaft move* [1] and the symplectic geometry community calls a *bending flow* [37] by rotating half of the polygon rigidly with respect to the other half around the diagonal; the collection of such rotations around all of the $n - 3$ diagonals of a given triangulation will be our Hamiltonian torus action. We can now summarize the existing literature as follows:

Theorem 13 (Kapovich and Millson [37], Howard, Manon and Millson [34], Hitchin [33]). *The following facts are known:*

- $\widehat{\text{Pol}}_3(n; \vec{r})$ is a possibly singular $(2n - 6)$ -dimensional symplectic manifold. The symplectic volume is equal to the standard measure.
- To any triangulation T of the standard n -gon we can associate a Hamiltonian action of the torus T^{n-3} on $\widehat{\text{Pol}}_3(n; \vec{r})$, where the angle θ_i acts by folding the polygon around the i -th diagonal of the triangulation.
- The moment map $\mu : \widehat{\text{Pol}}_3(n; \vec{r}) \rightarrow \mathbb{R}^{n-3}$ for a triangulation T records the lengths d_i of the $n - 3$ diagonals of the triangulation.
- The moment polytope P for a triangulation T is defined by the triangulation inequalities.
- The action-angle map α for a triangulation T is given by constructing the triangles using the diagonal and edgelength data to recover their side lengths, and assembling them in space with (oriented) dihedral angles given by the θ_i .
- The inverse image $\mu^{-1}(\text{interior } P) \subset \widehat{\text{Pol}}_3(n; \vec{r})$ of the interior of the moment polytope P is an (open) toric symplectic manifold.

Here is a very brief summary of how these results work. Just as for Hamiltonian torus actions, in general there is a moment map associated to every Hamiltonian Lie group action on a symplectic manifold. In particular, Kapovich and Millson [37] pointed out that the symplectic manifold




FIG. 2: This figure shows how to construct an equilateral pentagon in $\widehat{\text{Pol}}(5; \vec{1})$ using the action-angle map. First, we pick a point in the moment polytope shown in Figure 3 at center. We have now specified diagonals d_1 and d_2 of the pentagon, so we may build the three triangles in the triangulation from their side lengths, as in the picture at left. We then choose dihedral angles θ_1 and θ_2 independently and uniformly, and join the triangles along the diagonals d_1 and d_2 , as in the middle picture. The right hand picture shows the final space polygon, which is the boundary of this triangulated surface.

$\text{Arm}_3(n; \vec{r})$ admits a Hamiltonian action by the Lie group $\text{SO}(3)$ given by rotating the polygonal arm in space (this is the diagonal $\text{SO}(3)$ action on the product of spheres) whose moment map μ gives the vector joining the ends of the polygon. The closed polygons $\text{Pol}_3(n; \vec{r})$ are the fiber $\mu^{-1}(\vec{0})$ of this map. While the group action does not generally preserve fibers of this moment map, it does preserve $\mu^{-1}(\vec{0}) = \text{Pol}_3(n; \vec{r})$ and in this situation, we can perform what is known as a *symplectic reduction* (or Marsden–Weinstein–Meyer reduction [49, 50]) to produce a symplectic structure on the quotient of the fiber $\mu^{-1}(\vec{0})$ by the group action. This yields a symplectic structure on the $(2n - 6)$ -dimensional moduli space $\widehat{\text{Pol}}_3(n; \vec{r})$. The symplectic measure induced by this symplectic structure is equal to the standard measure given by pushing forward the subspace measure on $\text{Pol}_3(n; \vec{r})$ to $\widehat{\text{Pol}}_3(n; \vec{r})$ because the “parent” symplectic manifold $\text{Arm}_3(n; \vec{r})$ is a Kähler manifold [33].

The polygon space $\widehat{\text{Pol}}_3(n; \vec{r})$ is singular if

$$\varepsilon_I(\vec{r}) := \sum_{i \in I} r_i - \sum_{j \notin I} r_j$$

is zero for some $I \subset \{1, \dots, n\}$. Geometrically, this means it is possible to construct a linear polygon with edgelengths given by \vec{r} . Since linear polygons are fixed by rotations around the axis on which they lie, the action of $\text{SO}(3)$ is not free in this case and the symplectic reduction develops singularities. Nonetheless, the reduction $\widehat{\text{Pol}}_3(n; \vec{r})$ is a complex analytic space with isolated singularities; in particular, the complement of the singularities is a symplectic (in fact Kähler) manifold to which Theorem 13 applies.

Both the volume and the cohomology ring of $\widehat{\text{Pol}}_3(n; \vec{r})$ are well-understood from this symplectic perspective [11, 32, 36, 38, 39, 46, 66]. For example:

Proposition 14 (Takakura [66], Khoi [38], Mandini [46]). *The volume of $\widehat{\text{Pol}}_3(n; \vec{r})$ is given by*

$$\text{Vol}(\widehat{\text{Pol}}_3(n; \vec{r})) = -\frac{(2\pi)^{n-3}}{2(n-3)!} \sum_I (-1)^{n-|I|} \varepsilon_I(\vec{r})^{n-3},$$

where the sum is over all $I \subset \{1, \dots, n\}$ such that $\varepsilon_I(\vec{r}) > 0$.

Corollary 15. *The volume of the space of equilateral n -gons is*

$$\text{Vol}(\widehat{\text{Pol}}_3(n; \vec{1})) = -\frac{(2\pi)^{n-3}}{2(n-3)!} \sum_{k=0}^{\lfloor n/2 \rfloor} (-1)^k \binom{n}{k} (n-2k)^{n-3}.$$

4.1. The knotting probability for equilateral hexagons

We immediately give an example application of this picture. In [16], we showed using the Fáry-Milnor theorem that at least $1/3$ of hexagons of total length 2 are unknotted by showing that their total curvature was too small to form a knot. We could repeat the calculation using our explicit formula for the expectation of the total curvature for equilateral hexagons above, but the results would be disappointing; only about 27% of the space is revealed to be unknotted by this method. On the other hand action-angle coordinates, coupled with results of Calvo, immediately yield a much better bound:

Proposition 16. *At least $3/4$ of the space $\widehat{\text{Pol}}_3(6; \vec{1})$ of equilateral hexagons consists of unknots.*

Proof. Consider the triangulation of the hexagon given by joining vertices 1, 3, and 5 by diagonals and its corresponding action-angle coordinates $\alpha : \mathcal{P} \times T^3 \rightarrow \widehat{\text{Pol}}_3(6; \vec{1})$. In [14], an impressively detailed analysis of hexagon space, Jorge Calvo defines a geometric² invariant of hexagons called the curl which is 0 for unknots and ± 1 for trefoils. In the proof of his Lemma 16, Calvo observes that any knotted equilateral hexagon with curl $+1$ has all three dihedral angles between 0 and π .

The rest of the proof is elementary, but we give all the steps here as this is the first of many such arguments below. Formally, the knot probability is the expected value of the characteristic function

$$\chi_{\text{knot}}(p) = \begin{cases} 1 & \text{if } p \text{ is knotted,} \\ 0 & \text{if } p \text{ is unknotted.} \end{cases}$$

² Interestingly, curl is independent from topological invariant given by the handedness of the trefoil, so there are actually four different types of equilateral hexagonal trefoils.

By Calvo’s work, χ_{knot} is bounded above by the sum of characteristic functions $\chi_{\text{curl}=+1} + \chi_{\text{curl}=-1}$ and $\chi_{\text{curl}=+1}$ is bounded above by the simpler characteristic function

$$\chi(d_1, d_2, d_3, \theta_1, \theta_2, \theta_3) = \begin{cases} 1 & \text{if } \theta_i \in [0, \pi] \text{ for } i \in \{1, 2, 3\}, \\ 0 & \text{otherwise.} \end{cases}$$

Now Theorem 13 tells us that almost all of $\widehat{\text{Pol}}_3(6; \vec{1})$ is a toric symplectic manifold, so (2) of Theorem 1 holds for integrals over this polygon space. In particular, χ doesn’t depend on the d_i , so its expected value over $\widehat{\text{Pol}}_3(6; \vec{1})$ is equal to its expected value over the torus T^3 of θ_i . This expected value is clearly $1/8$, and a similar argument holds for $\chi_{\text{curl}=-1}$. This means that the knot probability is no more than $1/4$, as desired. \square

Of course, this bound is still a substantial underestimate of the fraction of unknots. Over a 12 hour run of the “PTSMCMC” Markov chain sampler of Section 5.5, we examined 1,318,001 equilateral hexagons and found 173 knots. Using the 95% confidence level Geyer IPS error estimators of Section 5.3, we estimate the knot probability among unconfined equilateral hexagons to be $1.3 \times 10^{-4} \pm 0.2 \times 10^{-4}$, or between 1.1 and 1.5 in 10,000.

4.2. The Fan Triangulation and Chordlengths

While Theorem 13 applies to any triangulation, we now work out the details for a particular triangulation. The “fan” triangulation is created by joining vertex v_1 of the polygon to vertices v_3, \dots, v_{n-1} . This gives rise to the following set of triangulation inequalities. As shown in Figure 3, we number the diagonals d_1, \dots, d_{n-3} so that the first triangle has edgelengths d_1, r_1, r_2 , the last triangle has edgelengths d_{n-3}, r_{n-1}, r_n , and all the triangles in between have edgelengths in the form d_i, d_{i+1}, r_{i+2} . The corresponding triangulation inequalities, which we call the “fan triangulation inequalities” are:

$$\begin{aligned} |r_1 - r_2| \leq d_1 \leq r_1 + r_2 & & r_{i+2} \leq d_i + d_{i+1} & & |r_n - r_{n-1}| \leq d_{n-3} \leq r_n + r_{n-1} \\ |d_i - d_{i+1}| \leq r_{i+2} & & & & \end{aligned} \quad (16)$$

Definition 17. The *fan triangulation polytope* $P_n(\vec{r}) \subset \mathbb{R}^{n-3}$ is the moment polytope for $\widehat{\text{Pol}}_3(n; \vec{r})$ corresponding to the fan triangulation and is determined by the fan triangulation inequalities (16). The fan triangulation polytopes $P_5(\vec{1})$ and $P_6(\vec{1})$ are shown in Figure 3.

This description of the moment polytope follows directly from Theorem 13.

Applying Theorem 1 to this situation gives necessary and sufficient conditions for uniform sampling on $\widehat{\text{Pol}}_3(n; \vec{r})$. These could be used to test proposed polygon sampling algorithms given statistical tests for uniformity on convex subsets of Euclidean space and on the $(n - 3)$ -torus.




FIG. 3: This figure shows the fan triangulation of a 7-gon on the left and the corresponding moment polytopes for equilateral space pentagons and equilateral space hexagons. For the pentagon moment polytope, we show the square with corners at $(0, 0)$ and $(2, 2)$ to help locate the figure, while for the hexagon moment polytope, we show the box with corners at $(0, 0, 0)$ and $(2, 3, 2)$ to help understand the geometry of the figure. The vertices of the polytopes correspond to polygons fixed by the torus action given by rotating around the diagonals. For instance, the $(2, 2)$ point in the pentagon’s moment polytope corresponds to the configuration given by an isocles triangle with sides 2, 2, and 1. The diagonals lie along the long sides; rotating around them is a rotation of the entire configuration in space, and is hence trivial because we are considering equivalence classes up to the action of $SO(3)$. The $(2, 3, 2)$ point in the hexagon’s moment polytope corresponds to a completely flat (or “lined”) configuration double-covering a line segment of length 3. Here all the diagonals lie along the same line and rotation around the diagonals does nothing.

Proposition 18. *A polygon in $\widehat{\text{Pol}}_3(n; \vec{r})$ is sampled according to the standard measure if and only if its diagonal lengths $d_1 = |v_1 - v_3|$, $d_2 = |v_1 - v_4|$, \dots , $d_{n-3} = |v_1 - v_{n-1}|$ are uniformly sampled from the fan polytope $P_n(\vec{r})$ and its dihedral angles around these diagonals are sampled independently and uniformly in $[0, 2\pi)$.*

The fan triangulation polytope also gives us a natural way to understand the probability distribution of chord lengths of a closed random walk. To fix definitions,

Definition 19. Let $\text{ChordLength}(k, n; \vec{r})$ be the length $|v_1 - v_{k+1}|$ of the chord skipping the first k edges in a polygon sampled according to the standard measure on $\widehat{\text{Pol}}_3(n; \vec{r})$. This is a random variable.

The expected values of squared chordlengths for equilateral polygons have been computed by a rearrangement technique, and turn out to be quite simple:

Proposition 20 (Cantarella, Deguchi, Shonkwiler [15] and Millett, Zirbel [73]). *The second moment of the random variable $\text{ChordLength}(k, n; \vec{1})$ is $\frac{k(n-k)}{n-1}$.*

It is obviously interesting to know the other moments of these random variables, but this problem seems considerably harder. In particular, the techniques used in the proofs of Proposition 20 don't apply to other moments of chordlength. Here is an alternate form for the chordlength problem which allows us to make some explicit calculations:

Theorem 21. *The expected value of the random variable $\text{ChordLength}(k, n; \vec{\Gamma})$ is coordinate d_{k-1} of the center of mass of the fan triangulation polytope $P_n(\vec{\Gamma})$. For n between 4 and 8, these expectations are given by the fractions*

$n \setminus k$	2	3	4	5	6
4	1				
5	17/15	17/15			
6	14/12	15/12	14/12		
7	461/385	506/385	506/385	461/385	
8	1,168/960	1,307/960	1,344/960	1,307/960	1,168/960

(17)

The p -th moment of $\text{ChordLength}(k, n; \vec{\Gamma})$ is coordinate d_{k-1} of the p -th center of mass of $P_n(\vec{\Gamma})$.

Proof. Since the measure on $\widehat{\text{Pol}}_3(n; \vec{\Gamma})$ is invariant under permutations of the edges, the pdf of chord length for any chord skipping k edges must be the same as the pdf for the length of the chord joining v_1 and v_{k+1} . But this chord is a diagonal of the fan triangulation, so its length is the coordinate d_{k-1} of the fan triangulation polytope $P_n(\vec{\Gamma})$. Since these chord lengths don't depend on dihedral angles, their expectations over polygon space are equal to their expectations over $P_n(\vec{\Gamma})$ by (2) of Theorem 1, which applies by Theorem 13. But the expectation of a power of a coordinate over a region is simply a coordinate of the corresponding center of mass. We obtained the results in the table by a direct computer calculation using `polymake` [27], which decomposes the polytopes into simplices and computes the center of mass as a weighted sum of simplex centers of mass. \square

It would be very interesting to get a general formula for these polytope centers of mass.

4.3. Closed polygons in (rooted) spherical confinement

Following Diao et. al. [22], we say that a polygon p is in rooted spherical confinement of radius R if every vertex of the polygon is contained in a sphere of radius R centered at the first vertex of the polygon. As a subspace of the space of closed polygons of fixed edgelengths, the space of confined closed polygons inherits a toric symplectic structure. In fact, the moment polytope for this structure is a very simple subpolytope of the fan triangulation polytope:

Definition 22. The *confined fan polytope* $P_{n,R}(\vec{r}) \subset P_n(\vec{r})$ is determined by the fan triangulation inequalities (16) and the additional linear inequalities $d_i \leq R$.

As before, we immediately have action-angle coordinates $P_{n,R}(\vec{r}) \times T^{n-3}$ on the space of rooted confined polygons. We note that the vertices of the confined fan triangulation polytope corresponding to a space of confined polygons are *not* all fixed points of the torus action since this is not the entire moment polytope; new vertices have been added by imposing the additional linear inequalities. As before, we get criteria for sampling confined polygons (directly analogous to Proposition 18 for unconfined polygons):

Proposition 23. *A polygon in $\widehat{\text{Pol}}_3(n; \vec{r})$ is sampled according to the standard measure on rooted sphere-confined closed space polygons in a sphere of radius R if and only if its diagonal lengths $d_1 = |v_1 - v_3|$, $d_2 = |v_1 - v_4|$, \dots , $d_{n-3} = |v_1 - v_{n-1}|$ are uniformly sampled from the confined fan polytope $P_{n,R}(\vec{r})$ and its dihedral angles around these diagonals are sampled independently and uniformly in $[0, 2\pi)$.*

We can also compute expected values for chordlengths for confined polygons following the lead of Theorem 21, but here our results are weaker because the pdf of chordlength is no longer simply a function of the number of edges skipped:

Theorem 24. *The expected length of the chord joining vertex v_1 to vertex v_{k+1} in a polygon sampled according to the standard measure on polygons in rooted spherical confinement of radius R is given by coordinate d_{k-1} of the center of mass of the confined fan triangulation polytope $P_{n,R}(\vec{r})$. For n between 4 and 10, $\vec{r} = \vec{1}$, and $R = 3/2$, these expectations are*

$n \setminus k$	2	3	4	5	6	7	8	(denominator)
4	3/4							
5	8/9	8/9						
6	293/336	316/336	293/336					
7	281/320	298/320	298/320	281/320				
8	23,237	24,752	24,402	24,752	23,237			26,496
9	46,723	49,718	49,225	49,225	49,718	46,723		53,256
10	1,145,123	1,218,844	1,205,645	1,210,696	1,205,645	1,218,844	1,145,123	1,305,344

where for $n = 8, 9$, and 10 we moved the common denominator of all fractions in the row to the right-hand column.

The proof is just the same as the proof of Theorem 21, and again we use `polymake` for the computations. The data show an interesting pattern: for 8, 9, and 10 edge polygons, the confinement is tight enough that the data reveals small parity effects in the expectations. For 10-gons, for instance, vertex v_5 is on average closer to vertex v_1 than vertex v_4 is! We also calculated the exact expectation of chordlength for equilateral 10-gons confined to spheres of other radii. The results are shown in Figure 4.




FIG. 4: Each line in this graph shows the expected length of the chord joining vertices v_1 and v_k in a random equilateral 10-gon. The 10-gons are sampled from the standard measure on polygons in rooted spherical confinement. From bottom to top, the confinement radii are 1.25, 1.5, 1.75, 2, 2.5, 3, 4, and 5. Polygons confined in a sphere of radius 5 are unconfined. Note the small parity effects which emerge in tighter confinement. These are exact expectations, not the result of sampling experiments.

5. MARKOV CHAIN MONTE CARLO FOR CLOSED AND CONFINED RANDOM WALKS

We have now constructed the action-angle coordinates on several spaces of random walks, including closed walks, closed walks in rooted spherical confinement, standard (open) random walks, and open random walks confined to half-spaces or slabs. In each case, the action-angle coordinates have allowed us to prove some theorems and make some interesting exact computations of probabilities on the spaces. To address more complicated (and physically interesting) questions, we will now turn to numerically sampling these spaces.

Numerical sampling of closed polygons underlies a substantial body of work on the geometry and topology of polymers and biopolymers (see the surveys of [55] and [6], which contain more than 200 references), which is a topic of interest in statistical physics. Many of the physics questions at issue in these investigations seem to be best addressed by computation. For instance, while our methods above gave us simple (though not very tight) theoretical bounds on the fraction of unknots among equilateral 6-gons, a useful theoretical bound on, say, the fraction of unknots among 1,273-gons seems entirely out of reach. On the other hand, it is entirely reasonable to work on developing well-founded algorithms with verified convergence and statistically defensible error bars for experimental work on such questions, and that is precisely our aim in this part of the paper.

5.1. Current Sampling Algorithms for Random Polygons

A wide variety of sampling algorithms for random polygons have been proposed. They fall into two main categories: Markov chain algorithms such as polygonal folds [51] or crankshaft

moves [40, 70] (cf. [1] for a discussion of these methods) and direct sampling methods such as the “triangle method” [53] or the “generalized hedgehog” method [69] and the methods of Moore and Grosberg [52] and Diao, Ernst, Montemayor, and Ziegler [22–24] which are both based on the “sinc integral formula” (7).

Each of these approaches has some defects. No existing Markov chain method has been proved to converge to the standard measure, though it is generally conjectured that they do. It is unclear what measure the generalized hedgehog method samples, while the triangle method clearly samples a submanifold³ of polygon space. The Moore–Grosberg algorithm is known to sample the correct distribution, but faces certain practical problems. It is based on computing successive piecewise-polynomial distributions for diagonal lengths of a closed polygon and directly sampling from these one-dimensional distributions. There is no problem with the convergence of this method, but the difficulty is that the polynomials are high degree with large coefficients and many almost-cancellations, leading to significant numerical problems with accurately evaluating them⁴. These problems are somewhat mitigated by the use of rational and multiple-precision arithmetic in [52], but the number of edges in polygons sampled with these methods is inherently limited. For instance, the text file giving the coefficients of the polynomials needed to sample a random closed 95-gon is over 25 megabytes in length. Diao et. al. avoid this problem by approximating these distributions by normals, but this approximation means that they are not quite⁵ sampling the standard measure on polygon space.

5.2. The Toric Symplectic Markov Chain Monte Carlo Algorithm

We introduce a Markov Chain Monte Carlo algorithm for sampling toric symplectic manifolds with an adjustable parameter $\beta \in (0, 1)$ explained below. We will call this the TORIC-SYMPLECTIC-MCMC(β) algorithm or TSMCMC(β) for convenience. Though we intend to apply this algorithm to our random walk spaces, it works on any toric symplectic manifold, so we state the results in this section and the next for an arbitrary toric symplectic manifold M^{2n} with moment map $\mu: M \rightarrow \mathbb{R}^n$, moment polytope P , and action-angle parametrization $\alpha: P \times T^n \rightarrow M$. The method is based on a classical Markov chain for sampling convex regions of \mathbb{R}^n called the “hit-and-run” algorithm: choose a direction at random and sample along the intersection of that ray with the region to find the next point in the chain. This method was introduced by Boneh and Golan [9]

³ It is hard to know whether this restriction is important in practice. The submanifold may be sufficiently “well-distributed” that most integrands of interest converge anyway. Or perhaps calculations performed with the triangle method are dramatically wrong for some integrands!

⁴ Hughes discusses these methods in Section 2.5.4 of his book on random walks [35], attributing the formula rederived by Moore and Grosberg [52] to a 1946 paper of Treloar [68]. The problems with evaluating these polynomials accurately were known by the 1970’s, when Barakat [5] derived an alternate expression for this probability density based on Fourier transform methods.

⁵ Again, it is unclear what difference this makes in practice.

and independently by Smith [62] as a means of generating random points in a high-dimensional polytope. There is a well-developed theory around this method which we will be able to make use of below.

The parameter β controls the relative weight given to the action and angle variables. At each step of $\text{TSMCMC}(\beta)$, with probability β we update the action variables by sampling the moment polytope P using hit-and-run and with probability $1 - \beta$ we update the angle variables by sampling the torus T^n uniformly. When $\beta = 1/2$ this is analogous to the random scan Metropolis-within-Gibbs samplers discussed by Roberts and Rosenthal [60] (see also [41]). The adjustable parameter β allows us to tune the speed and accuracy of Monte Carlo integration using $\text{TSMCMC}(\beta)$. As a trivial example, if we use $\text{TSMCMC}(\beta)$ to numerically integrate a functional which is almost constant in the angle variables, then we should let β be very close to 1.

```

TORIC-SYMPLECTIC-MCMC( $\vec{p}, \vec{\theta}, \beta$ )
  prob = UNIFORM-RANDOM-VARIATE(0, 1)
  if prob <  $\beta$ 
    then  $\triangleright$  Generate a new point in  $P$  using the hit-and-run algorithm.
          $\vec{v} = \text{RANDOM-DIRECTION}(n)$ 
          $(t_0, t_1) = \text{FIND-INTERSECTION-ENDPOINTS}(P, \vec{p}, \vec{v})$ 
          $t = \text{UNIFORM-RANDOM-VARIATE}(t_0, t_1)$ 
          $\vec{p} = \vec{p} + t\vec{v}$ 
    else  $\triangleright$  Generate a new point in  $T^n$  uniformly.
         for ind = 1 to  $n$ 
           do  $\theta_{ind} = \text{UNIFORM-RANDOM-VARIATE}(0, 2\pi)$ 
  return  $(\vec{p}, \vec{\theta})$ 

```

We now prove that the distribution of samples produced by this Markov chain converges geometrically to the distribution generated by the symplectic volume on M^{2n} . First, we show that the symplectic measure on M is invariant for TSMCMC.

To do so, recall that for any Markov chain Φ on a state space X , we can define the m -step transition probability $\mathcal{P}^m(x, A)$ to be the probability that an m -step run of the chain starting at x lands in the set A . This defines a measure $\mathcal{P}^m(x, \cdot)$ on X . The transition kernel $\mathcal{P} = \mathcal{P}^1$ is called *reversible* with respect to a probability distribution π if

$$\int_A \pi(dx) \mathcal{P}(x, B) = \int_B \pi(dx) \mathcal{P}(x, A) \text{ for all measurable } A, B \subset X. \quad (18)$$

In other words, the probability of moving from A to B is the same as the probability of moving from B to A . If \mathcal{P} is reversible with respect to π , then π is invariant for \mathcal{P} : letting $A = X$ in (18), we see that $\pi\mathcal{P} = \pi$.

In $\text{TSMCMC}(\beta)$, the transition kernel $\mathcal{P} = \beta\mathcal{P}_1 + (1 - \beta)\mathcal{P}_2$, where \mathcal{P}_1 is the hit-and-run kernel on the moment polytope and $\mathcal{P}_2(\vec{\theta}, \cdot) = \tau$, where τ is the uniform measure on T^n . Since hit-and-run is reversible on the moment polytope [63] and since \mathcal{P}_2 is obviously reversible with respect to τ , we have:

Proposition 25. *TSMCMC(β) is reversible with respect to the symplectic measure ν induced by symplectic volume on M . In particular, ν is invariant for TSMCMC(β).*

Recall that the total variation distance between two measures η_1, η_2 on a state space X is given by

$$|\eta_1 - \eta_2|_{\text{TV}} := \sup_{A \text{ any measurable set}} |\eta_1(A) - \eta_2(A)|.$$

We can now prove geometric convergence of the sample measure generated by $\text{TSMCMC}(\beta)$ to the symplectic measure in total variation distance.

Theorem 26. *Suppose that M^{2n} is a toric symplectic manifold with moment polytope P and action-angle coordinates $\alpha: P \times T^n \rightarrow M^{2n}$. Further, let $\mathcal{P}^m(\vec{p}, \vec{\theta}, \cdot)$ be the m -step transition probability of the Markov chain given by TORIC-SYMPLECTIC-MCMC(β) and let ν be the symplectic measure on M^{2n} .*

There are constants $R < \infty$ and $\rho < 1$ so that for any $(\vec{p}, \vec{\theta}) \in \text{int}(P) \times T^n$,

$$\left| \alpha_* \mathcal{P}^m(\vec{p}, \vec{\theta}, \cdot) - \nu \right|_{\text{TV}} < R\rho^m.$$

That is, for any choice of starting point, the pushforward by α of the probability measure generated by TORIC-SYMPLECTIC-MCMC(β) on $P \times T^n$ converges geometrically (in the number of steps taken in the chain) to the symplectic measure on M^{2n} .

Proof. Let λ be Lebesgue measure on the moment polytope P and, as above, let τ be uniform measure on the torus T^n . By Theorem 1, it suffices to show that

$$\left| \mathcal{P}^m(\vec{p}, \vec{\theta}, \cdot) - \lambda \times \tau \right|_{\text{TV}} < R\rho^m.$$

Since the transition kernels \mathcal{P}_1 and \mathcal{P}_2 commute, for any non-negative integers a and b and partitions i_1, \dots, i_k of a and j_1, \dots, j_ℓ of b we have

$$\left(\mathcal{P}_1^{i_1} \mathcal{P}_2^{j_1} \dots \mathcal{P}_1^{i_k} \mathcal{P}_2^{j_\ell} \right) (\vec{p}, \vec{\theta}, \cdot) = \left(\mathcal{P}_1^a \mathcal{P}_2^b \right) (\vec{p}, \vec{\theta}, \cdot) = \mathcal{P}_1^a(\vec{p}, \cdot) \times \mathcal{P}_2^b(\vec{\theta}, \cdot) = \mathcal{P}_1^a(\vec{p}, \cdot) \times \tau, \quad (19)$$

where the last equality follows from the fact that $\mathcal{P}_2(\vec{\theta}, \cdot) = \tau$ for any $\vec{\theta} \in T^n$.

The total variation distance between product measures is bounded above by the sum of the total variation distances of the factors (this goes back at least to Blum and Pathak [8]; see Sandler [61] for a proof), so we have that

$$\begin{aligned} \left| \mathcal{P}_1^a(\vec{p}, \cdot) \times \mathcal{P}_2^b(\vec{\theta}, \cdot) - \lambda \times \tau \right|_{\text{TV}} &= \left| \mathcal{P}_1^a(\vec{p}, \cdot) \times \tau - \lambda \times \tau \right|_{\text{TV}} \\ &\leq \left| \mathcal{P}_1^a(\vec{p}, \cdot) - \lambda \right|_{\text{TV}} + \left| \tau - \lambda \right|_{\text{TV}} \\ &= \left| \mathcal{P}_1^a(\vec{p}, \cdot) - \lambda \right|_{\text{TV}}. \end{aligned} \quad (20)$$

Using a result of Smith [63, Theorem 3], the right hand side is bounded above by $\left(1 - \frac{\xi}{n2^{n-1}}\right)^{a-1}$ where ξ is the ratio of the volume of P and the volume of the smallest round ball containing P . Let

$$\kappa := \left(1 - \frac{\xi}{n2^{n-1}}\right).$$

Then combining (19), (20), and the binomial theorem yields

$$\begin{aligned} \left| \mathcal{P}^m(\vec{p}, \vec{\theta}, \cdot) - \lambda \times \tau \right|_{\text{TV}} &= \left| (\beta \mathcal{P}_1 + (1 - \beta) \mathcal{P}_2)^m(\vec{p}, \vec{\theta}, \cdot) - \lambda \times \tau \right|_{\text{TV}} \\ &= \left| \sum_{i=0}^m \binom{m}{i} \beta^{m-i} (1 - \beta)^i (\mathcal{P}_1^{m-i}(\vec{p}, \cdot) \times \tau - \lambda \times \tau) \right|_{\text{TV}} \\ &\leq \sum_{i=0}^m \binom{m}{i} \beta^{m-i} (1 - \beta)^i \kappa^{m-i-1} \\ &= \frac{1}{\kappa} (1 + \beta(\kappa - 1))^m = \frac{1}{\kappa} \left(1 - \frac{\beta\xi}{n2^{n-1}}\right)^m. \end{aligned}$$

The ratio ξ of the volume of P and the volume of smallest round ball containing P is always a positive number with absolute value less than 1, so $0 < 1 - \beta\xi/n2^{n-1} < 1$. This completes the proof. \square

This proposition provides a comforting theoretical guarantee that TORIC-SYMPLECTIC-MCMC(β) will eventually work. The proof provides a way to estimate the constants R and ρ . However, in practice, these upper bounds are far too large to be useful. Further, the rate of convergence for any given run will depend on the shape and dimension of the moment polytope P and on the starting position x . There is quite a bit known about the performance of hit-and-run in general theoretical terms; we recommend the excellent survey article of Andersen and Diaconis [2]. To give one example, L6vász and Vempala have shown [44] (see also [43]) that the number of steps of hit-and-run required to reduce the total variation distance between $\mathcal{P}^m(x, \cdot)$ and Lebesgue measure by an order of magnitude is proportional⁶ to n^3 where n is the dimension of the polytope.

⁶ The constant of proportionality is large and depends on the geometry of the polytope, and the amount of time required to reduce the total variation distance to a fixed amount from the start depends on the distance from the starting point to the boundary of the polytope.

5.3. The Markov Chain CLT and Geyer's IPS Error Bounds for TSMCMC integration

We now know that the $\text{TSMCMC}(\beta)$ algorithm will eventually sample from the correct probability measure on any toric symplectic manifold, and in particular from the correct probability distributions on closed and confined random walks. We should pause to appreciate the significance of this result for a moment – while many Markov chain samplers have been proposed for closed polygons, none have been proved to converge to the correct measure. Further, there has never been a Markov chain sampler for closed polygons in rooted spherical confinement (or, as far as we know, for slab-confined or half-space confined arms).

However, the situation remains in some ways unsatisfactory. If we wish to compute the probability of an event in one of these probability spaces of polygons, we must do an integral over the space by collecting sample values from a Markov chain. But since we do not have any explicit bounds on the rate of convergence of our Markov chains, we do not know how long to run the sampler, or how far the resulting sample mean might be from the integral over the space. To answer these questions, we need two standard tools: the Markov Chain Central Limit Theorem and Geyer's Initial Positive Sequence (IPS) error estimators for MCMC integration [28]. For the convenience of readers unfamiliar with these methods, we summarize the construction here. Since this is basically standard material, many readers may wish to skip ahead to the next section.

Combining Proposition 26 with [67, Theorem 5] (which is based on [20, Corollary 4.2]) yields a central limit theorem for the $\text{TORIC-SYMPLECTIC-MCMC}(\beta)$ algorithm. To set notation, a run of the $\text{TSMCMC}(\beta)$ algorithm produces the sequence of points $((\vec{p}_0, \vec{\theta}_0), (\vec{p}_1, \vec{\theta}_1), \dots)$, where the initial point $(\vec{p}_0, \vec{\theta}_0)$ is drawn from some initial distribution (for example, a delta distribution). For any run R , let

$$\text{SMean}(f; R, m) := \frac{1}{m} \sum_{k=1}^m f(\vec{p}_k, \vec{\theta}_k)$$

be the sample mean of the values of a function $f: M^{2n} \rightarrow \mathbb{R}$ over the first m steps in R . We will use “ f ” interchangeably to refer to the original function $f: M^{2n} \rightarrow \mathbb{R}$ or its expression in action-angle coordinates $f \circ \alpha: P \times T^n \rightarrow \mathbb{R}$.

Let $E(f; \nu)$ be the expected value of f with respect to the symplectic measure ν on M . Then for each positive integer m the normalized sample error

$$\sqrt{m}(\text{SMean}(f; R, m) - E(f; \nu))$$

is a random variable (as it depends on the various random choices in the run R).

Proposition 27. *Suppose f is a square-integrable real-valued function on the toric symplectic manifold M^{2n} . Then regardless of the initial distribution, there exists a real number $\sigma(f)$ so that*

$$\sqrt{m}(\text{SMean}(f; R, m) - E(f; \nu)) \xrightarrow{w} \mathcal{N}(0, \sigma(f)^2), \quad (21)$$

where $\mathcal{N}(0, \sigma(f)^2)$ is the Gaussian distribution with mean 0 and standard deviation $\sigma(f)$ and the superscript w denotes weak convergence.

Given $\sigma(f)$ and a run R , the range $\text{SMean}(f; R, m) \pm 1.96\sigma(f)/\sqrt{m}$ is an approximate 95% confidence interval for the true expected value $E(f; \nu)$. Abstractly, we can find $\sigma(f)$ as follows.

The variance of the left hand side of (21) is

$$m \text{Var}(\text{SMean}(f; R, m)) = \frac{1}{m} \sum_{i=1}^m \text{Var}(f(\vec{p}_i, \vec{\theta}_i)) + \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^m \text{Cov}(f(\vec{p}_i, \vec{\theta}_i), f(\vec{p}_j, \vec{\theta}_j)).$$

Since the convergence in Proposition 27 is independent of the initial distribution, $\sigma(f)$ will be the limit of this quantity for *any* initial distribution. Following Chan and Geyer [18], suppose the initial distribution is the stationary distribution. In that case, the quantities

$$\gamma_0(f) := \text{Var}(f(\vec{p}_i, \vec{\theta}_i))$$

and

$$\gamma_k(f) := \text{Cov}(f(\vec{p}_i, \vec{\theta}_i), f(\vec{p}_{i+k}, \vec{\theta}_{i+k}))$$

(the stationary variance and lag k autocovariance, respectively) are independent of i . Then

$$\sigma(f)^2 = \lim_{m \rightarrow \infty} \left(\gamma_0(f) + 2 \sum_{k=1}^{m-1} \frac{m-k}{m} \gamma_k(f) \right) = \gamma_0(f) + 2 \sum_{k=1}^{\infty} \gamma_k(f)$$

provided the sum on the right hand side converges.

In what follows it will be convenient to write the above as

$$\sigma(f)^2 = \gamma_0(f) + 2\gamma_1(f) + 2 \sum_{k=1}^{\infty} \Gamma_k(f), \quad (22)$$

where $\Gamma_k(f) := \gamma_{2k}(f) + \gamma_{2k+1}(f)$. Again, we emphasize that the quantities $\gamma_0(f)$, $\gamma_k(f)$, $\Gamma_k(f)$ are associated to the *stationary* Markov chain.

In practice, of course, these quantities and hence this expression for $\sigma(f)$ are not computable. After all, if we could sample directly from the symplectic measure on M^{2n} there would be no need for TSMCMC. However, as pointed out by Geyer [28], $\sigma(f)$ can be estimated from the sample data that produced $\text{SMean}(f; R, m)$. Specifically, we will estimate the stationary lagged autocovariance $\gamma_k(f)$ by the following quantity:

$$\bar{\gamma}_k(f) = \frac{1}{m} \sum_{i=1}^{m-k} [f(\vec{p}_i, \vec{\theta}_i) - \text{SMean}(f; R, m)][f(\vec{p}_{i+k}, \vec{\theta}_{i+k}) - \text{SMean}(f; R, m)] \quad (23)$$

Note that multiplication by $1/m$ rather than $1/(m-k)$ is not a typographical error (cf. [28], Sec 3.1). Let $\bar{\Gamma}_k(f) = \bar{\gamma}_{2k}(f) + \bar{\gamma}_{2k+1}(f)$. Then for any $N > 0$

$$\bar{\sigma}_{m,N}(f)^2 := \bar{\gamma}_0(f) + 2\bar{\gamma}_1(f) + 2 \sum_{k=1}^N \bar{\Gamma}_k(f) \quad (24)$$

is an estimator for $\sigma(f)^2$. We expect the $\bar{\Gamma}_k$ to decrease to zero as $k \rightarrow \infty$ since very distant points in the run of the Markov chain should become statistically uncorrelated. Indeed, since TSMCMC is reversible, Geyer shows this is true for the stationary chain:

Theorem 28 (Geyer [28], Theorem 3.1). Γ_k is strictly positive, strictly decreasing, and strictly convex as a function of k .

We expect, then, that any non-positivity, non-monotonicity, or non-convexity of the $\bar{\Gamma}_k$ should be due to k being sufficiently large that $\bar{\Gamma}_k$ is dominated by noise. In particular, this suggests that a reasonable choice for N in (24) is the first N such that $\bar{\Gamma}_N \leq 0$, since the terms past this point will be dominated by noise and hence tend to cancel each other.

Definition 29. Given a function f and a length- m run of the TSMCMC algorithm as above, let N be the largest integer so that $\bar{\Gamma}_1(f), \dots, \bar{\Gamma}_N(f)$ are all strictly positive. Then the *initial positive sequence estimator* for $\sigma(f)$ is

$$\bar{\sigma}_m(f) := \bar{\sigma}_{m,N}(f) = \bar{\gamma}_0(f) + 2\bar{\gamma}_1(f) + 2 \sum_{k=1}^N \bar{\Gamma}_k(f).$$

Slightly more refined initial sequence estimators which take into account the monotonicity and convexity from Proposition 28 are also possible; see [28] for details.

The pleasant result of all this is that $\bar{\sigma}_m$ is a statistically consistent overestimate of the actual variance.

Theorem 30 (Geyer [28], Theorem 3.2). For almost all sample paths of TSMCMC,

$$\liminf_{m \rightarrow \infty} \bar{\sigma}_m(f)^2 \geq \sigma(f)^2.$$

Therefore, we propose the following procedure for *Toric Symplectic Markov Chain Monte Carlo integration* which yields statistically consistent error bars on the estimate of the true value of the integral:

Toric Symplectic Markov Chain Monte Carlo Integration. Let f be a square-integrable function on a toric symplectic manifold M^{2n} with moment map $\mu : M \rightarrow \mathbb{R}^n$.

1. Find the fixed points of the Hamiltonian torus action. The moment polytope P is the convex hull of the images of these fixed points under μ .
2. Convert this vertex description of P to a halfspace description. In other words, realize P as the subset of points in \mathbb{R}^n satisfying a collection of linear inequalities.⁷

⁷ For small problems, this can be done algorithmically [4, 19, 27]. Generally, this will require an analysis of the moment polytope, such as the one performed above for the moment polytopes of polygon spaces.

3. Pick the parameter $\beta \in (0, 1)$. We recommend repeating the entire procedure for several short runs with various β values to decide on the best β for a given application. The final error estimate is a good measure of how well the chain has converged after a given amount of runtime.
4. Pick a point $(\vec{p}_0, \vec{\theta}_0) \in P \times T^n$. This will be the starting point of the Markov chain. Ideally, \vec{p}_0 should be as far as possible from the boundary of P .
5. Using $(\vec{p}_0, \vec{\theta}_0)$ as the initial input, iterate the TSMCMC(β) algorithm for m steps ($m \gg 1$). This produces a finite sequence $((\vec{p}_1, \vec{\theta}_1), \dots, (\vec{p}_m, \vec{\theta}_m))$ of points in $P \times T^n$.
6. Let $\text{SMean}(f; m) = \frac{1}{m} \sum_{i=1}^m f(\vec{p}_i, \vec{\theta}_i)$ be the average value of f over the run of points produced in the previous step.
7. Compute the initial positive sequence estimator $\bar{\sigma}_m(f)^2$.
8. $\text{SMean}(f; m) \pm 1.96\bar{\sigma}_m(f)/\sqrt{m}$ is an approximate 95% confidence interval for the true expected value of the function f .

5.4. Tuning the TSMCMC Algorithm for Closed and Confined Polygons

The TSMCMC(β) algorithm for polygons has several adjustable parameters. We must always choose a starting polygon. For unconfined polygons, we may choose any triangulation of the n -gon and get a corresponding moment polytope. Finally, we must make an appropriate choice of β . In this section, we report experimental results which address these questions. In our experiments, we always integrated total curvature and used equilateral closed polygons. At least for unconfined polygons, we know the exact value of the expectation from Theorem 12. To measure convergence, we used the Geyer IPS error estimate as a measure of quality (lower is better). Since different step types take very different amounts of time to run, we ran different variations of the algorithm for a consistent amount of CPU time, even though this led to very different step counts.

We discovered in our experiments that the rate of convergence of hit-and-run depends strongly on the start point. Our original choice of start point – the regular planar equilateral n -gon – turned out to be a very poor performer. While it seems like a natural choice mathematically, the regular n -gon is tucked away in a corner of the moment polytope and it takes hit-and-run quite a while to escape this trap. After a number of experiments, the most consistently desirable start point was obtained by folding the regular n -gon along the diagonals of the given triangulation and then randomly permuting the resulting edge set. We used this as a starting configuration in all of our unconfined experiments.

We also discovered that hit-and-run can converge relatively slowly when sampling high-dimensional polytopes, leading to very long-range autocorrelations in the resulting Markov chain.

Following a suggestion of Soteris [64], after considerable experimentation we settled on the convention that a single “moment polytope” step in our implementation of $\text{TSMCMC}(\beta)$ would represent ten iterations of hit-and-run on the moment polytope. This reduced autocorrelations greatly and led to better convergence overall. We used this convention for all our numerical experiments below.

The $\text{TORIC-SYMPLECTIC-MCMC}(\beta)$ algorithm depends on a choice of triangulation T for the n -gon to determine the moment polytope P . There is considerable freedom in this choice, since the number of triangulations of an n -gon is the Catalan number $C_{n-2} \sim 4^{n-2}/(n-2)^{3/2}\sqrt{\pi}$. We have proved above that the $\text{TORIC-SYMPLECTIC-MCMC}(\beta)$ algorithm will converge for any of these triangulations, but the rate of convergence is expected to depend on the triangulation, which determines the geometry of the moment polytope. This geometry directly affects the rate of convergence of hit-and-run; “long and skinny” polytopes are harder to sample than “round” ones (see Lovasz [43]).

To get a sense of the effect of the triangulation on the performance of $\text{TSMCMC}(\beta)$ we set $\beta = 0.5$ and $n = 23$ and ran the algorithm from 20 start points for 20,000 steps. We then took the average IPS error bar for expected total curvature over these 20 runs as a measure of convergence. We repeated this analysis for 300 random triangulations and 300 repeats of three triangulations that we called the “fan”, “teeth” and “spiral” triangulations. The results are shown in Figure 5. The definition of the fan and teeth triangulations will be obvious from that figure; the spiral triangulation is generated by traversing the n -gon in order repeatedly, joining every other vertex along the traversal until the triangulation is complete. Our experiments showed that this spiral triangulation was the best performing triangulation among our candidates, so we standardized on that triangulation for further numerical experiments.

We then considered the effect of varying the parameter β for the $\text{TSMCMC}(\beta)$ algorithm using the spiral triangulation. We ran a series of trials computing expected total curvature for 64-gons over 10 minute runs, while varying β from 0.05 (almost all dihedral steps) to 0.95 (almost all moment polytope steps) over 10 minute runs. We repeated each run 50 times to get a sense of the variability in the Geyer IPS error estimators for different runs. Since dihedral steps are considerably faster than moment polytope steps, the step counts varied from about 1 to 9 million. The resulting Geyer IPS error estimators are shown in Figure 6. Our recommendation from these experiments is to use the spiral triangulation and $\beta = 0.5$ for experiments with unconfined polygons. From the 50 runs using the recommended $\beta = 0.5$, the run with the median IPS error estimate produced an expected total curvature estimate of 101.724 ± 0.142 using about 4.6 million samples; recall that we computed in Table II that the expected value of total curvature for equilateral, unconfined 64-gons is a complicated fraction close to 101.7278.




FIG. 5: We tested the average IPS 95% confidence error estimate for the expected value of total curvature over random equilateral 23-gons over 20 runs of the TSMCMC(0.5) algorithm. Each run had a starting point generated by folding and permuting a regular n -gon as described above, and ran for 20,000 steps. We tried 300 random n -gons and 300 repetitions of the same procedure for the “spiral”, “fan”, and “teeth” triangulations shown above. Below each triangulation is shown the range of average error bars observed over 300 repetitions of the 20-start-point trials; for the random triangulation we report the best average error bar over a single 20-start-point-trial observed for any of the 300 random triangulations we computed. We can see that the algorithm based on the spiral triangulation generally outperforms algorithms based on even the best of the 300 random triangulations, while algorithms based on the fan and teeth triangulations converged more slowly.




FIG. 6: The figure above shows a box-and-whisker plot for the IPS error estimators observed in computing expected total curvature over 50 runs of the TSMCMC(β) algorithm for various values of β . The boxes show the $1/4$ to $3/4$ quantiles of the data, while the whiskers extend from the 0.05 quantile to the 0.95 quantile. While the whiskers show that there is plenty of variability in the data, the general trend is that the performance of the algorithm improves rapidly as β varies from 0.05 to 0.25, modestly as β varies from 0.25 to 0.5 and is basically constant for β from 0.5 to 0.95.

5.5. Crankshafts, folds and permutation steps for unconfined equilateral polygons

If we are dealing with unconfined equilateral polygons, we can define a transformation from $\widehat{\text{Pol}}_3(n; \vec{1})$ to itself by permuting the edge set of the polygon. Since the symplectic measure on the space $\widehat{\text{Pol}}_3(n; \vec{1}) = \text{Pol}_3(n; \vec{1})/\text{SO}(3)$ of unconfined equilateral poly-

gons agrees with the standard measure, which is the pushforward of the subspace measure on $\text{Pol}_3(n; \vec{1}) \subset \text{Arm}_3(n; \vec{1}) = S^2(1) \times \dots \times S^2(1)$, the action of the symmetric group S_n which permutes edges is measure-preserving on $\widehat{\text{Pol}}_3(n; \vec{1})$. As a consequence, we will see that we can mix permutation steps with standard TSMCMC steps without losing geometric convergence or the applicability of the Central Limit Theorem. Such a Markov chain is a mixture of dihedral angle steps, moment polytope steps, and permutation steps in some proportion. It is interesting to note that we can recover algorithms very similar to the standard “crankshaft” and “fold” Markov chains by allowing no moment polytope steps in the chain.

Since previous authors have observed that adding permutation steps can significantly speed up convergence in polygon samplers [1], we now experiment to see whether our algorithm, too, can be improved by mixing in some permutations. More precisely, we can define a new Markov chain POLYGON-PERMUTATION on $\widehat{\text{Pol}}_3(n; \vec{1})$ by permuting edges at each step:

```
POLYGON-PERMUTATION(pol)
   $\sigma = \text{UNIFORM-PERMUTATION}(n)$ 
   $\text{pol} = \text{PERMUTE-EDGES}(\text{pol}, \sigma)$ 
  return pol
```

Since the symplectic measure on $\widehat{\text{Pol}}_3(n; \vec{1})$ is permutation-invariant, the symplectic measure is stationary for POLYGON-PERMUTATION.

Now, we can mix $\text{TSMCMC}(\beta)$ with POLYGON-PERMUTATION to get the following PERMUTATION-TORIC-SYMPLECTIC-MCMC(β, δ) algorithm, where $\delta \in [0, 1)$ gives the probability of doing a permutation step rather than a $\text{TSMCMC}(\beta)$ step. Recall that $\alpha: P \times T^{n-3} \rightarrow \widehat{\text{Pol}}_3(n; \vec{1})$ is the action-angle parametrization, where P is the moment polytope induced by the chosen triangulation.

```
PERMUTATION-TORIC-SYMPLECTIC-MCMC( $\vec{p}, \vec{\theta}, \beta, \delta$ )
   $\text{prob} = \text{UNIFORM-RANDOM-VARIATE}(0, 1)$ 
  if  $\text{prob} < \delta$ 
    then  $(\vec{p}, \vec{\theta}) = \alpha^{-1}(\text{POLYGON-PERMUTATION}(\alpha(\vec{p}, \vec{\theta})))$ 
    else  $(\vec{p}, \vec{\theta}) = \text{TORIC-SYMPLECTIC-MCMC}(\vec{p}, \vec{\theta}, \beta)$ 
  return  $(\vec{p}, \vec{\theta})$ 
```

Although POLYGON-PERMUTATION is *not* ergodic, the fact that it is stationary with respect to the symplectic measure is, by combining Proposition 26 and [67, Proposition 3], enough to imply that PERMUTATION-TORIC-SYMPLECTIC-MCMC(β, δ) is (strongly) uniformly ergodic.

Proposition 31. *Let $\widehat{\mathcal{P}}$ be the transition kernel for PTSMCMC(β, δ) with $0 < \beta < 1$ and $\delta < 1$ and let ν be the symplectic measure on $\widehat{\text{Pol}}_3(n; \vec{1})$. Then there exist constants $R < \infty$ and $\rho < 1$ so that for any $(\vec{p}, \vec{\theta}) \in \text{int}(P) \times T^{n-3}$,*

$$\left| \alpha_* \widehat{\mathcal{P}}^m(\vec{p}, \vec{\theta}, \cdot) - \nu \right|_{\text{TV}} < R\rho^m.$$

Just as in Proposition 27, since PTSMCMC(β, δ) is uniformly ergodic and reversible with respect to symplectic measure, it satisfies a central limit theorem:

Proposition 32. *Suppose $f: \widehat{\text{Pol}}_3(n; \vec{1}) \rightarrow \mathbb{R}$ is square-integrable. For any run R of PTSMCMC(β, δ), let $\text{SMean}(f; R, m)$ be the sample mean of the value of f over the first m steps of R . Then there exists a real number $\sigma(f)$ so that*

$$\sqrt{m}(\text{SMean}(f; R, m) - E(f; \nu)) \xrightarrow{w} \mathcal{N}(0, \sigma(f)^2).$$

The rest of the machinery of Section 5.3, including the initial positive sequence estimator for $\sigma(f)^2$, also applies. As a consequence, we get a modified **Toric Symplectic Markov Chain Monte Carlo Integration** procedure adapted to unconfined, equilateral polygons. Note that the full symmetric group S_n does not act on $\widehat{\text{Pol}}_3(n; \vec{r})$ when not all r_i are equal, so PTSMCMC(β, δ) cannot be used to sample non-equilateral polygons. However when many edgelengths are equal, a subgroup of the symmetric group which permutes only those edges certainly acts on $\widehat{\text{Pol}}_3(n; \vec{r})$. We recommend making use of this smaller set of permutations when possible. Permuting edges never preserves spherical confinement, so PTSMCMC(β, δ) is inapplicable to confined polygon sampling.

Having defined PTSMCMC(β, δ) and settled on a canonical starting point (the folded, permuted regular n -gon) and triangulation (the spiral), it remains to decide on the best values of β and δ . The question is complicated by the fact that the three different types of steps – permutations, folding steps, and moment-polytope hit-and-run steps – take different amounts of CPU time. To attempt to evaluate the various possibilities fairly, we ran experiments computing the expected total curvature for 64-gons where each experiment ran for 10 minutes of CPU time, completing between 2 million and 15 million steps depending on the mixture of step types. We measured the 95% confidence IPS error bars for each run, producing the data in Figure 7, and used the size of this error bar as a measure of convergence.

The data in Figure 7 show that the fraction δ of permutation steps is the most important factor in determining the rate of convergence in the PTSMCMC(β, δ) algorithm. This shows that the extra complication in defining PTSMCMC(β, δ) for unconfined equilateral polygons is worth it: the error bars produced by PTSMCMC(β, δ) to compute the expected total curvature of unconfined equilateral 64-gons are anywhere from 3 to 30 times smaller than the error bars for TSMCMC(β).

Larger values of β produce smaller error bars when $\delta = 0$, meaning that a large fraction of moment polytope steps are needed to produce mixing when there are no permutation steps. On the

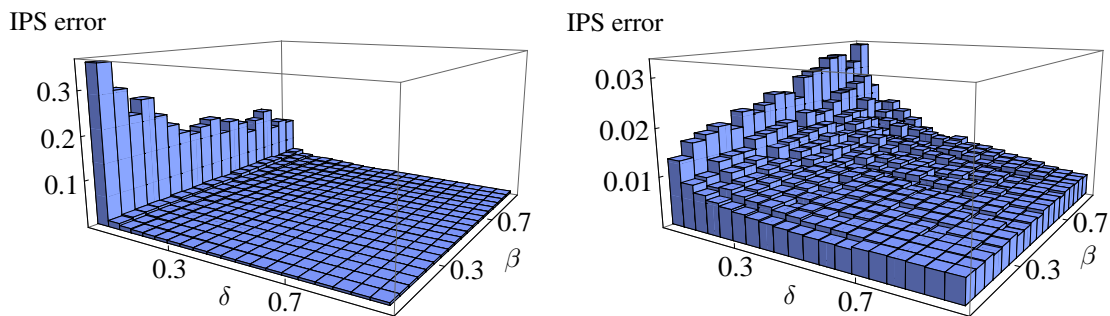


FIG. 7: This plot shows the IPS error estimator for the average total curvature of unconfined equilateral 64-gons. The IPS error was computed for 10 minute runs of the PTSMCMC(β, δ) Markov chain algorithm. The values of δ (the fraction of permutations among all steps) ranged from 0 to 0.95 in steps of 0.05, while the values of β (the fraction of moment polytope steps among non-permutation steps) ranged from 0 to 0.95 in steps of 0.05. When $\delta = 0$, this is just the TSMCMC(β) chain; removing those runs yields the plot on the right. We observed that convergence was very sensitive to δ , with error bars improving dramatically as soon as the fraction of permutation steps becomes positive: even the worst PTSMCMC(β, δ) run with $\delta > 0$ had error bars 3 times smaller than the error bars of the best TSMCMC(β) run. From the view at right, we can see that the error bars continue to improve more modestly as δ increases. Varying β has little effect on the error estimate when δ is large.

other hand, as we can see in Figure 8, even when $\delta = 0.05$ the permutation steps provide enough mixing that β has virtually no effect on the IPS standard deviation estimator. In this case, the effect of β on the size of the error bars is due to the fact that dihedral steps are faster than moment polytope steps, so runs with small β produce more samples and hence smaller error bars.




FIG. 8: These plots show the IPS error estimator and the IPS standard deviation estimator for the average total curvature of unconfined equilateral 64-gons using 10-minute runs of PTSMCMC($\beta, 0.05$). Although the IPS error estimate decreases as β decreases, the plot on the right demonstrates that the IPS standard deviation estimator is essentially constant – and presumably close to the true standard deviation of the total curvature function – across the different values of β . Since the IPS error estimate is proportional to the standard deviation estimate divided by the square root of the number of samples, we can see that the variation in IPS error bars for these runs is almost entirely due to the difference in the number of samples.

Once δ is large, varying β seems to have little effect on the convergence rate. In fact, though

our theory above no longer proves convergence, we seem to get a very competitive algorithm by removing moment polytope steps altogether ($\beta = 0$) and performing only permutations and dihedral steps. This algorithm corresponds to the “fold or crankshaft with permutations mixed in” method.

In practice, we make a preliminary recommendation of $\delta = 0.9$ and $\beta = 0.5$ for experimental work. These parameters guarantee convergence (by our work above) while optimizing the convergence rate. Using these recommended parameters, a 10 minute run of PTSMCMC(0.5,0.9) for unconfined, equilateral 64-gons produced just under 7 million samples and an expected total curvature of 101.7276 ± 0.0044 , which compares quite favorably to the actual expected total curvature of 101.7278.

Taking out the $\beta = 0$ runs, we observed that the absolute error in our computations of expected total curvature was less than our error estimate in 361 of 380 runs (95%), which is exactly what we would expect from a 95% confidence value estimator. We take this as solid evidence that the Markov chain is converging and the error estimators are working as expected.

5.6. Calculations on confined polygons

Recall from Definition 22 that a polygon is in spherical confinement in a sphere of radius R centered at vertex v_1 of the polygon if the vector \vec{d} of fan diagonals of the polygon lies in the confined fan polytope $P_{n,R}(\vec{r})$. This means that we can sample such polygons uniformly by restricting the hit-and-run steps in TSMCMC(β) to the confined fan polytope $P_{n,R}(\vec{r})$.

We again only explored the situation for equilateral polygons of edge length one. After some experimentation, we settled on the “folded triangle” as a start point. This polygon is constructed by setting each diagonal length d_i to one and choosing dihedrals randomly. This polygon is contained in spherical confinement for every $R \geq 1$, so we could use it for all of our experiments. We investigated 23-gons confined to spheres of radius 2, 4, 6, 8, 10 and 12, measuring the Geyer IPS error estimate for values of β selected from 0.05 (almost all dihedral steps) to 0.95 (almost all moment polytope steps) over 10 minute runs. Again, since dihedral steps are faster to run than moment polytope steps, the step counts varied over the course of the experiments. For instance, in the radius 2 experiments, we observed step counts as high as 35 million and as low as 7 million over runs with various β values. Our integrand was again total curvature. Since we do not have an exact solution for the expected total curvature of a confined n -gon, we were unable to check whether the error bars predicted actual errors. However, it was comforting to note that the answers we got from runs with various parameters were very consistent. We ran each experiment 50 times to get a sense of the repeatability of the Geyer IPS error bar; the results are shown in Figure 9.

We observed first that there is a clear trend in the error bar data. For the tightly confined runs, there was a noticeable preference for $\beta \sim 0.5$, while in less tight confinement the results generally continued to improve modestly as β increased. Still, we think the data supports a general




FIG. 9: These box-and-whisker plots show the results of computing the expected total curvature for confined equilateral 23-gons with edgelenhth 1. The confinement model here is “rooted spherical confinement”, meaning that each vertex is within a specified distance of the first vertex. For each β value, we repeated 10 minute experiments 50 times, computing 50 values for the Geyer IPS estimator. The boxes show the second and third quartiles of these 50 data points, while the whiskers show the 0.05 to 0.95 quantiles of the IPS estimators observed.

recommendation of $\beta = 0.5$ for future confined experiments, with a possible decrease to $\beta = 0.4$ in very tight confinement, and this is our recommendation to future investigators.

A very striking observation from Figure 9 is that the error bars for the tightly confined 23-gons in a sphere of radius 2 are about 10 times smaller than the error bars for the very loosely confined 23-gons in a sphere of radius 10. That is, our algorithm works better when the polygon is in tighter confinement! In some sense, this is to be expected, since the space being sampled is smaller. However, it flies in the face of the natural intuition that confined sampling should be numerically more difficult than unconfined sampling.

Using TSMCMC(0.5), we computed the expected total curvature of tightly confined equilateral 50- and 90-gons. Those expectations are shown in Table I. We can compare these data directly by looking at expected turning angles as in Figure 10. In this very tight confinement regime, the effect of confinement radius on expected turning angle dominates the effect of the number of edges.

6. COMPARISON WITH EXISTING WORK, CONCLUSION AND FUTURE DIRECTIONS

Now that we have laid out the symplectic theory of random walks and a few of its consequences, it is time to look back and see how we can reconcile this point of view with the existing understanding of closed random walks. In the methods of Moore and Grosberg [52] and Diao et. al. [22], closed random walks are generated incrementally, using distributions derived from the pdf $\Phi_n(\vec{\ell})$ given in (12) for the end-to-end distance of a random walk of n steps. To review, the key idea is that if we have taken $m - 1$ steps of an n -step closed walk and arrived at the m -th vertex \vec{v}_m ,

Confinement Radius	50-gons	90-gons
1.1	103.1120 ± 0.0093	185.701 ± 0.028
1.2	100.1900 ± 0.0089	180.261 ± 0.028
1.3	97.8369 ± 0.0088	175.947 ± 0.028
1.4	95.8891 ± 0.0090	172.346 ± 0.027
1.5	94.1979 ± 0.0091	169.271 ± 0.028
1.6	92.7501 ± 0.0094	166.660 ± 0.029
∞	79.74197470	142.5630093

TABLE I: This table shows the expected total curvature of equilateral 50- and 90-gons in rooted spherical confinement. We sampled equilateral 50- and 90-gons in confinement radii from 1.1 to 1.6 using 20-minute runs of TSMCMC(0.5) and computed the average total curvature and IPS error bars for each run. Each 50-gon run yielded about 14.5 million samples, while each 90-gon run yielded about 8 million samples. The bottom line shows the exact expectation of total curvature for unconfined polygons given by Theorem 12. More extensive information on expectations of confined total curvatures has been computed by Diao, Ernst, Montemayor, and Ziegler [26].




FIG. 10: The plot on the left shows the expected turning angles of equilateral 50-gons (blue) and equilateral 90-gons (red) in rooted spherical confinement of radii from 1.1 to 1.6. The horizontal lines show the expected turning angles for unconfined 50- and 90-gons computed using Theorem 12, which are $\simeq 1.59484$ and $\simeq 1.58403$, respectively. The plot on the right shows the differences between the expected turning angles of equilateral 50-gons and the expected turning angles of equilateral 90-gons. The blue dots show this difference for various confinement radii, while the red line shows the corresponding difference for unconfined polygons. Without confinement, we expect polygons with more edges to have smaller expected turning angle, since each individual edge feels less pressure to get back to the starting point. These data provide evidence this effect dissipates and even reverses in extremely tight confinement.

the pdf of the next vertex \vec{v}_{m+1} (conditioned on the steps we have already taken) is given by

$$P(\vec{v}_{m+1}|\vec{v}_1, \dots, \vec{v}_m) = \frac{\Phi_1(\vec{v}_{m+1} - \vec{v}_m)\Phi_{n-m-1}(\vec{v}_{m+1} - \vec{v}_1)}{\Phi_{n-m}(\vec{v}_m - \vec{v}_1)},$$

which is some complicated product of piecewise-polynomial $\Phi_k(\vec{\ell})$ functions. We can sample \vec{v}_{m+1} from this distribution and hence generate the rest of the walk iteratively.

From the moment polytope point of view, the situation is considerably simpler. First, we observe that everything in the equation above can be expressed in terms of diagonal lengths in the fan triangulation polytope, since the length of the vector $\vec{\ell}$ is the only thing that matters in the formula for $\Phi_k(\vec{\ell})$. If we let $\vec{v}_1 = \vec{0}$ by convention, then conditioning on $\vec{v}_1, \dots, \vec{v}_m$ is simply restricting our attention to the slice of the moment polytope given by setting the diagonal lengths $d_1 = |\vec{v}_3|, d_2 = |\vec{v}_4|, \dots, d_{m-2} = |\vec{v}_m|$. The pdf $P(\vec{v}_{m+1}|\vec{v}_1, \dots, \vec{v}_m)$ is then the projection of the measure on this slice of the moment polytope to the coordinate d_{m-1} . This distribution is piecewise-polynomial precisely because it is the projection of Lebesgue measure on a convex polytope with a finite number of faces.

Of course, projecting volume measures of successive slices to successive coordinates is a perfectly legitimate way to sample a convex polytope, which is another explanation for why these methods work; they are basically sampling successive marginals of the coordinate distributions on a succession of smaller convex polytopes. By contrast, our method generates the entire vector of diagonal lengths d_1, \dots, d_{n-3} simultaneously according to their joint distribution by sampling the moment polytope directly. More importantly it offers a geometric insight into what this joint distribution *is* which seems like it would be very hard to develop by analyzing (12).

In conclusion, the moment polytope picture offers a clarifying and useful perspective on closed and confined random walks. It is clear that we have only scratched the surface of this topic in this paper, and that many fascinating questions remain to be explored both theoretically and computationally. In the interest of continuing the conversation, we provide an unordered list of open questions suggested by the work above.

- Previous studies of the relative efficiency of polygon sampling algorithms have focused on minimizing pairwise correlations between edges as a measure of performance. Proposition 18 suggests a more subtle statistical approach to evaluating sample quality: measure the uniformity of the distribution of diagonal lengths over the moment polytope and of dihedral angles over the torus (cf. [47]).
- It remains open to try to extend these methods to prove that a chain consisting only of permutation and dihedral steps is still strongly geometrically convergent on unconfined equilateral polygon space. This would lead directly to a proof of convergence for the crankshaft and fold algorithms, and hence place many years of sampling experiments using these methods on a solid theoretical foundation.
- Can we use the moment polytope pictures above for confined polygons to prove theorems about polygons in confinement? For instance, it would be very interesting to show that the expectation of total curvature is monotonic in the radius of confinement.
- What is the corresponding picture for random planar polygons? Of course, we can see the planar polygons as a special slice of the action-angle coordinates where the angles are all

zero or π . But is it true that sampling this slice according to Hausdorff measure in action-angle space corresponds to sampling planar polygons according to their Hausdorff measure inside space polygons⁸? If not, can we correct the measure somehow? Or is there another picture for planar polygons entirely?

- Can we understand the triangulation polytopes better? Can we compute their centers of mass explicitly, for example? It is well known that finding the center of mass of a high-dimensional polytope is algorithmically difficult, so we can't hope for a purely mechanical solution to the problem. But a deeper understanding of these polytopes seems likely to result in interesting probability theorems.
- Why are permutation steps so effective in the PTSMCMC Markov chain? It seems easy to compute that the number of points in the permutation group orbit of an n -edge polygon is growing much faster than the volume of equilateral polygon space computed by [38, 46, 66] and given above as Corollary 15. Can we prove that the points in this orbit are usually well-distributed over polygon space? This would give an appealing proof of the effectiveness of Grosberg's triangle method for polygon sampling [45, 53, 54].
- There is a large theory of "low-discrepancy" or "quasi-random" sequences on the torus which can provide better results in numerical integration than uniform random sampling. Would it be helpful to choose our dihedrals from such a sequence in the integration method above?
- Now that we can sample confined polygons quickly, with solid error bars on our calculations, what frontiers does this open in the numerical study of confined polymers? We take our cues from the pioneering work of Diao, Ernst, Montemayor, and Ziegler [21–24], but are eager to explore this new experimental domain. For instance, sampling tightly confined n -gons might be a useful form of "enriched sampling" in the hunt for complicated knots of low equilateral stick number, since very entangled polygons are likely to be geometrically compact as well.

We introduced a related probabilistic theory of non-fixed edge length closed polygons in a previous paper [15] by relating closed polygons with given total length to Grassmann manifolds. It remains to explain the connection between that picture and this one, and we will take up that question shortly.

⁸ These questions are less obvious than they may appear at first glance: the cylindrical coordinates θ and z are action-angle coordinates on the sphere, but it is not the case that the arclength measure on a curve in the θ - z cylinder pushes forward to the arclength measure on the image of the curve on the sphere, even though the area measure on the θ - z cylinder does push forward to the standard area measure on the sphere.

7. ACKNOWLEDGEMENTS

We are grateful to many more friends and colleagues for important discussions related to this project than we can possibly remember to name here. But to give it our best shot, Michael Usher taught us a great deal of symplectic geometry, Malcolm Adams introduced us to the Duistermaat-Heckman theorem, Margaret Symington provided valuable insight on moment maps, and Alexander Y Grosberg and Tetsuo Deguchi have been constant sources of insight and questions on polygon spaces in statistical physics. Yuanan Diao, Claus Ernst, and Uta Ziegler introduced us to the Rayleigh sinc integral form for the pdf of arm length (and to a great deal more). We were inspired by their insightful work on confined sampling to look at confinement models above. Ken Millett and Eric Rawdon have graciously endured our various doubts about the convergence of the crankshaft and fold algorithms for many years, and were the source of many pivotal conversations. Chris Soteris provided much appreciated expert guidance on Markov chain sampling. And we are especially indebted to Alessia Mandini, Chris Manon, Angela Gibney, and Danny Krashen for explaining to us some of the elements of the algebraic geometry of polygon spaces.

Shonkwiler was supported by the Simons Foundation and the UGA VIGRE II grant (DMS-07-38586), and Cantarella was supported by the Simons Foundation. We were supported by the Georgia Topology Conference grant (DMS-11-05699), which helped us organize a conference on polygon spaces in the summer of 2013. We are grateful to the Issac Newton Institute for the Mathematical Sciences, Cambridge, for support and hospitality during the program “Topological Dynamics in the Physical and Biological Sciences” in Fall 2012, when much of this work was completed.

Appendix A: Expected Total Curvature of Equilateral Closed Polygons for Small n

In Section 3.4, we found an exact integral formula for the expectation of total curvature for equilateral n -gons following the approach of Grosberg [29]. Grosberg analyzed the asymptotics of this formula for large numbers of edges, showing that the expected total curvature approaches the asymptotic value $n\frac{\pi}{2} + \frac{3\pi}{8}$. We are interested in evaluating the formula exactly for small n in order to provide a check on our numerical methods. We used *Mathematica* to evaluate the formula, obtaining the fractional expressions shown in Table II. Grosberg’s asymptotic value is shown in the rightmost column.

Though for space reasons it had to be truncated in the table, the exact value for the expected total curvature of equilateral, unconfined 64-gons is

$$\frac{4522188530226656504649836292227453294126904427946053625769754177967556412769571113455}{139655807027685559939231323004419270090691937881733899567960159577537880384373522432} \pi + \frac{288230376151711744}{491901992474628194486464288049342660789103293530486293575717158971541638355891307}.$$

n	Expected total curvature	Decimal	Asymptotic
3	2π	6.28319	5.89049
4	8	8	7.46128
5	$-2\pi + 9\sqrt{3}$	9.30527	9.03208
6	$6\pi - 8$	10.8496	10.6029
7	$\frac{316}{33}\pi - \frac{225}{22}\sqrt{3}$	12.369	12.1737
8	$\frac{15}{4}\pi + \frac{32}{15}$	13.9143	13.7445
9	$\frac{766}{289}\pi + \frac{11907}{2890}\sqrt{3}$	15.463	15.3153
10	$\frac{11}{2}\pi - \frac{64}{245}$	17.0175	16.8861
11	$\frac{90712}{14219}\pi - \frac{1686177}{1990660}\sqrt{3}$	18.5751	18.4569
12	$\frac{331545}{51776}\pi + \frac{512}{28315}$	20.1351	20.0277
13	$\frac{23336570}{3407523}\pi + \frac{2381643}{22716820}\sqrt{3}$	21.6969	21.5984
14	$\frac{877129}{118464}\pi - \frac{1024}{1282743}$	23.2601	23.1692
15	$\frac{3189814022}{403436289}\pi - \frac{1786291299}{207097295020}\sqrt{3}$	24.8244	24.74
16	$\frac{241091487}{28701184}\pi + \frac{4096}{168339171}$	26.3896	26.3108
17	$\frac{197198281266}{22161558721}\pi + \frac{44753178051}{88734881118884}\sqrt{3}$	27.9554	27.8816
18	$\frac{42415625107}{4513689728}\pi - \frac{8192}{15127913229}$	29.5219	29.4524
19	$\frac{240270145231776}{24279795663511}\pi - \frac{4277229018201}{194432603673396088}\sqrt{3}$	31.0888	31.0232
20	$\frac{111226176353241}{10700200165376}\pi + \frac{131072}{14288920862931}$	32.6561	32.594
32	$\frac{262929167708231675164189486733}{16044875932324628104050900992}\pi + \frac{134217728}{46358282926117706045930790075}$	51.4816	51.4436
64	$\simeq \frac{4.52218853 \times 10^{84}}{1.39655807 \times 10^{83}}\pi + \frac{288230376151711744}{4.91901992 \times 10^{80}}$	101.7278	101.7091

TABLE II: The expected total curvature of equilateral n -gons computed by evaluating (13) in *Mathematica* for $4 \leq n \leq 20$ and $n = 32, 64$ (the integral becomes singular when $n = 3$, but all triangles have total curvature 2π), together with Grosberg's asymptotic approximation. We see that for 64-gons we need 5 significant digits to distinguish the exact value from the asymptotic approximation.

Appendix B: Proof of Proposition 6

In this section we prove Proposition 6, which we restate here:

Proposition 33. *The polytope*

$$\mathcal{H}_n = \{\vec{z} \in [-1, 1]^n \mid z_1 \geq 0, z_1 + z_2 \geq 0, \dots, z_1 + \dots + z_n \geq 0, -1 \leq z_i \leq 1\}$$

has volume $\frac{1}{2^n} \binom{2n}{n} = \frac{(2n-1)!!}{n!}$.

Our proof is a modification of an argument originally suggested on MathOverflow by Johan Wästlund [71]; Bernardi, Duplantier, and Nadeau [7] seem to have had something similar in mind.

Proof. Suppose that $s_k(\vec{z}) = z_1 + \dots + z_k$ is the k -th partial sum of the coordinates of \vec{z} , and by convention we set $s_0(\vec{z}) = 0$. The polytope \mathcal{H}_n can be defined as the subset of the hypercube where all $s_k(\vec{z}) \geq 0$. In the remainder of the hypercube, the subset of \vec{z} where all the $s_k(\vec{z})$ are different has full measure: we now partition this set into a collection of n polytopes $\mathcal{S}_0, \dots, \mathcal{S}_n$ defined by

$$\mathcal{S}_k := \{ \vec{z} \in [-1, 1]^n - \mathcal{H}_n \mid \text{the smallest } s_i(\vec{z}) \text{ is } s_k(\vec{z}) \}.$$

We claim that $\text{Vol } \mathcal{S}_k = \text{Vol } \mathcal{H}_k \cdot \text{Vol } \mathcal{H}_{n-k}$ for $k = 1, \dots, n-1$ and $\text{Vol } \mathcal{S}_n = \text{Vol } \mathcal{H}_n$. Consider the linear map

$$L_k: \mathcal{S}_k \subset \mathbb{R}^n \rightarrow \mathbb{R}^k \times \mathbb{R}^{n-k}, \quad L_k(z_1, \dots, z_n) = ((-z_k, -z_{k-1}, \dots, -z_1), (z_{k+1}, \dots, z_n)).$$

It is clear that this map preserves unsigned volume. We claim the image is exactly $\mathcal{H}_k \times \mathcal{H}_{n-k}$. Consider the partial sums of $(-z_k, \dots, -z_1)$. The i -th partial sum is given by

$$s_i(-z_k, \dots, -z_1) = -z_k - z_{k-1} - \dots - z_{k-i+1} = s_{k-i}(z_1, \dots, z_n) - s_k(z_1, \dots, z_n).$$

The point $(-z_k, \dots, -z_1)$ is in \mathcal{H}_k if and only if this partial sum is positive for all $i \in \{1, \dots, k\}$. But that happens exactly when $s_k(\vec{z})$ is negative⁹ and the smallest partial sum among $s_1(\vec{z}), \dots, s_k(\vec{z})$. On the other hand, if we consider the partial sums of (z_{k+1}, \dots, z_n) , we get

$$s_i(z_{k+1}, \dots, z_n) = z_{k+1} + \dots + z_{k+i} = s_{k+i}(z_1, \dots, z_n) - s_k(z_1, \dots, z_n).$$

The point (z_{k+1}, \dots, z_n) is in \mathcal{H}_{n-k} if and only if this partial sum is positive for all $i \in \{1, \dots, n-k\}$. But that happens exactly when $s_k(\vec{z})$ is the smallest partial sum among $s_k(\vec{z}), \dots, s_n(\vec{z})$, proving the claim. When $k = n$, \mathcal{S}_n is just a reversed and negated copy of \mathcal{H}_n itself.

We now have the relation

$$\begin{aligned} \text{Vol}[-1, 1]^n &= 2^n = \text{Vol } \mathcal{H}_n + \sum \text{Vol } \mathcal{S}_k \\ &= 2 \text{Vol } \mathcal{H}_n + \sum_{k=1}^{n-1} \text{Vol } \mathcal{H}_k \text{Vol } \mathcal{H}_{n-k}. \end{aligned} \quad (\text{B1})$$

and we can prove the formula by induction on n .

When $n = 1$ the polytope $\mathcal{H}_1 = [0, 1]$ and so the formula holds. For the inductive step, assume that $\text{Vol } \mathcal{H}_k = \frac{1}{2^k} \binom{2k}{k}$ for all $k < n$. Then solving (B1) for $\text{Vol } \mathcal{H}_n$ yields

$$\text{Vol } \mathcal{H}_n = 2^{n-1} - \frac{1}{2^{n+1}} \sum_{k=1}^{n-1} \binom{2k}{k} \binom{2(n-k)}{n-k}. \quad (\text{B2})$$

⁹ Remember our convention that $s_0(\vec{z}) = 0$, which is applied when $i = k$.

Using the Chu-Vandermonde identity

$$\sum_{k=0}^n \binom{x}{k} \binom{y}{n-k} = \binom{x+y}{n}$$

with $x = y = -1/2$ and recalling that

$$\binom{-1/2}{m} = (-1)^m \binom{2m}{m} \frac{1}{2^{2m}} \quad \text{and} \quad \binom{-1}{p} = (-1)^p$$

for any positive integers m and p , we see that

$$\sum_{k=1}^{n-1} \binom{2k}{k} \binom{2(n-k)}{n-k} = \sum_{k=0}^n \binom{2k}{k} \binom{2(n-k)}{n-k} - 2 \binom{2n}{n} = 2^{2n} - 2 \binom{2n}{n}.$$

Therefore, equation (B2) simplifies to

$$\text{Vol } \mathcal{H}_n = \frac{1}{2^n} \binom{2n}{n},$$

as desired. □

-
- [1] Sotero Alvarado, Jorge Alberto Calvo, and Kenneth C Millett. The generation of random equilateral polygons. *Journal of Statistical Physics*, 143(1):102–138, 2011.
 - [2] Hans C Andersen and Persi W Diaconis. Hit and run as a unifying device. *Journal de la Société Française de Statistique*, 148(4):5–28, 2007.
 - [3] Michael F Atiyah. Convexity and commuting Hamiltonians. *Bulletin of the London Mathematical Society*, 14(1):1–15, 1982.
 - [4] David Avis and Komei Fukuda. A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra. *Discrete & Computational Geometry*, 8(3):295–313, 1992.
 - [5] Richard Barakat. Isotropic random flights. *Journal of Physics A: Mathematical, Nuclear and General*, 6(6):796, 1973.
 - [6] Craig J Benham and Steven P Mielke. DNA Mechanics. *Annual Review of Biomedical Engineering*, 7(1):21–53, 2005.
 - [7] Olivier Bernardi, Bertrand Duplantier, and Philippe Nadeau. A bijection between well-labelled positive paths and matchings. *Séminaire Lotharingien de Combinatoire*, 63:Art. B63e–13, 2010.
 - [8] Julius R Blum and Pramod K Pathak. A Note on the Zero-one Law. *Annals of Mathematical Statistics*, 43(3):1008–1009, 1972.
 - [9] Arnon Boneh and A Golan. Constraints redundancy and feasible region boundedness by random feasible point generator (RGPG) . In *Third European Congress on Operations Research - EURO III*, Amsterdam, 1979.

- [10] David Borwein and Jonathan M Borwein. Some remarkable properties of sinc and related integrals. *Ramanujan Journal. An International Journal Devoted to the Areas of Mathematics Influenced by Ramanujan*, 5(1):73–89, 2001.
- [11] Michel Brion. Cohomologie équivariante des points semi-stables. *Journal für die Reine und Angewandte Mathematik*, 421:125–140, 1991.
- [12] Aniello Buonocore, Enrica Pirozzi, and Luigia Caputo. A note on the sum of uniform random variables. *Statistics and Probability Letters*, 79(19):2092–2097, 2009.
- [13] Carlos Bustamante, Zev Bryant, and Steven B Smith. Ten years of tension: single-molecule DNA mechanics. *Nature*, 421:423–426, 2003.
- [14] Jorge Alberto Calvo. The embedding space of hexagonal knots. *Topology Appl.*, 112(2):137–174, 2001.
- [15] Jason Cantarella, Tetsuo Deguchi, and Clayton Shonkwiler. Probability Theory of Random Polygons from the Quaternionic Viewpoint. *Communications on Pure and Applied Mathematics*, 2013.
- [16] Jason Cantarella, Alexander Y Grosberg, Robert B Kusner, and Clayton Shonkwiler. The expected total curvature of random polygons. Preprint, [arXiv:1210.6537 \[math.DG\]](https://arxiv.org/abs/1210.6537), 2012.
- [17] Francesco Caravenna. A local limit theorem for random walks conditioned to stay positive. *Probability Theory and Related Fields*, 133(4):508–530, 2005.
- [18] Kung Sik Chan and Charles J Geyer. Discussion: Markov Chains for Exploring Posterior Distributions. *The Annals of Statistics*, 22(4):1747–1758, 1994.
- [19] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10(4):377–409, 1993.
- [20] Robert Cogburn. The central limit theorem for Markov processes. *Proc Sixth Berkeley Symp Math Statist Probab*, 1972.
- [21] Yuanan Diao, Claus Ernst, Anthony Montemayor, Eric J Rawdon, and Uta Ziegler. The knot spectrum of confined random equilateral polygons. University of North Carolina, Charlotte Technical Report Series Preprint 2013-15, 2013.
- [22] Yuanan Diao, Claus Ernst, Anthony Montemayor, and Uta Ziegler. Generating equilateral random polygons in confinement. *Journal of Physics A: Mathematical and Theoretical*, 44(40):405202, 2011.
- [23] Yuanan Diao, Claus Ernst, Anthony Montemayor, and Uta Ziegler. Generating equilateral random polygons in confinement II. *Journal of Physics A: Mathematical and Theoretical*, 45(27):275203, 2012.
- [24] Yuanan Diao, Claus Ernst, Anthony Montemayor, and Uta Ziegler. Generating equilateral random polygons in confinement III. *Journal of Physics A: Mathematical and Theoretical*, 45(46):465003, 2012.
- [25] Johannes J Duistermaat and Gerrit J Heckman. On the variation in the cohomology of the symplectic form of the reduced phase space. *Inventiones Mathematicae*, 69(2):259–268, 1982.
- [26] Claus Ernst and Uta Ziegler. Personal Communication.
- [27] Ewgenij Gawrilow and Michael Joswig. polymake: a framework for analyzing convex polytopes. In *Polytopes—combinatorics and computation (Oberwolfach, 1997)*, pages 43–73. Birkhäuser, Basel, 2000.
- [28] Charles J Geyer. Practical Markov Chain Monte Carlo. *Statistical Science*, 7(4):473–483, 1992.
- [29] Alexander Y Grosberg. Total curvature and total torsion of a freely jointed circular polymer with $n \gg 1$ segments. *Macromolecules*, 41(12):4524–4527, 2008.
- [30] Victor Guillemin and Shlomo Sternberg. Convexity properties of the moment mapping. *Inventiones Mathematicae*, 67(3):491–513, 1982.
- [31] Jean-Claude Hausmann and Allen Knutson. Polygon spaces and Grassmannians. *L’Enseignement*

- Mathématique. Revue Internationale. 2e Série*, 43(1-2):173–198, 1997.
- [32] Jean-Claude Hausmann and Allen Knutson. The cohomology ring of polygon spaces. *Université de Grenoble. Annales de l'Institut Fourier*, 48(1):281–321, 1998.
- [33] Nigel J Hitchin, Anders Karlhede, Ulf Lindström, and Martin Roček. Hyperkähler metrics and supersymmetry. *Communications in Mathematical Physics*, 108(4):535–589, 1987.
- [34] Benjamin Howard, Christopher Manon, and John J Millson. The toric geometry of triangulated polygons in Euclidean space. *Canadian Journal of Mathematics. Journal Canadien de Mathématiques*, 63(4):878–937, 2011.
- [35] Barry D Hughes. *Random Walks and Random Environments: Volume 1: Random Walks*. Clarendon Press, 1995.
- [36] Yasuhiko Kamiyama and Michishige Tezuka. Symplectic volume of the moduli space of spatial polygons. *Journal of Mathematics of Kyoto University*, 39(3):557–575, 1999.
- [37] Michael Kapovich and John J Millson. The symplectic geometry of polygons in Euclidean space. *Journal of Differential Geometry*, 44(3):479–513, 1996.
- [38] Vu The Khoi. On the symplectic volume of the moduli space of spherical and Euclidean polygons. *Kodai Mathematical Journal*, 28(1):199–208, 2005.
- [39] Frances Clare Kirwan. The cohomology rings of moduli spaces of bundles over Riemann surfaces. *Journal of the American Mathematical Society*, 5(4):853–906, 1992.
- [40] Konstantin V Klenin, Alexander V Vologodskii, Vadim V Anshelevich, Alexander M Dykhne, and Maxim D Frank-Kamenetskii. Effect of excluded volume on topological properties of circular DNA. *Journal of Biomolecular Structure and Dynamics*, 5(6):1173–1185, 1988.
- [41] Krzysztof Łatuszyński, Gareth O Roberts, and Jeffrey S Rosenthal. Adaptive Gibbs samplers and related MCMC methods. *The Annals of Applied Probability*, 23(1):66–98, 2013.
- [42] Reginald Douglas Lord. The use of the Hankel transform in statistics I. General theory and examples. *Biometrika*, 41(1/2):44–55, 1954.
- [43] László Lovász. Hit-and-run mixes fast. *Mathematical Programming. A Publication of the Mathematical Programming Society*, 86(3, Ser. A):443–461, 1999.
- [44] László Lovász and Santosh Vempala. Hit-and-Run from a Corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.
- [45] Rhonald C Lua, Nathan T Moore, and Alexander Y Grosberg. Under-knotted and over-knotted polymers: 2. Compact self-avoiding loops. In Jorge Alberto Calvo, Kenneth C Millett, Eric J Rawdon, and Andrzej Stasiak, editors, *Physical and Numerical Models in Knot Theory: Including Applications to the Life Sciences*, pages 385–398. World Scientific Publishing Co., Singapore, 2004.
- [46] Alessia Mandini. The Duistermaat-Heckman formula and the cohomology of moduli spaces of polygons. *Journal of Symplectic Geometry*, 2013.
- [47] Kanti V Mardia and Peter E Jupp. *Directional Statistics*. Wiley Series in Probability and Statistics. John Wiley & Sons Ltd., Chichester, 2000.
- [48] Jean-Luc Marichal and Michael J Mossinghoff. Slices, slabs, and sections of the unit hypercube. *Online Journal of Analytic Combinatorics*, 3:1–11, 2008.
- [49] Jerrold Marsden and Alan Weinstein. Reduction of symplectic manifolds with symmetry. *Reports on Mathematical Physics*, 5(1):121–130, 1974.
- [50] Kenneth R Meyer. Symmetries and integrals in mechanics. In *Dynamical systems (Proc. Sympos., Univ. Bahia, Salvador, 1971)*, pages 259–272. Academic Press, New York, 1973.
- [51] Kenneth C Millett. Knotting of regular polygons in 3-space. *Journal of Knot Theory and its Ramifications*, 3(3):263–278, 1994.
- [52] Nathan T Moore and Alexander Y Grosberg. Limits of analogy between self-avoidance and topology-

- driven swelling of polymer loops. *Physical Review E. Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 72(6):061803, 2005.
- [53] Nathan T Moore, Rhonald C Lua, and Alexander Y Grosberg. Topologically driven swelling of a polymer loop. *Proceedings of the National Academy of Sciences*, 101(37):13431–13435, 2004.
- [54] Nathan T Moore, Rhonald C Lua, and Alexander Y Grosberg. Under-knotted and over-knotted polymers: 1. Unrestricted loops. In Jorge Alberto Calvo, Kenneth C Millett, Eric J Rawdon, and Andrzej Stasiak, editors, *Physical and Numerical Models in Knot Theory: Including Applications to the Life Sciences*, pages 363–384. World Scientific Publishing Co., Singapore, 2004.
- [55] Enzo Orlandini and Stuart G Whittington. Statistical topology of closed curves: Some applications in polymer physics. *Reviews of Modern Physics*, 79(2):611–642, 2007.
- [56] Xavier Pennec. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, 2006.
- [57] George Pólya. *On a few questions in probability theory and some definite integrals related to them*. PhD thesis, Eötvös Loránd University, 1912.
- [58] George Pólya. Berechnung eines bestimmten integrals (calculating some definite integrals). *Mathematisches Annalen*, 74:204–212, 1913.
- [59] Lord Rayleigh. On the problem of random vibrations, and of random flights in one, two, or three dimensions. *Philosophical Magazine Series 5*, 37(220):321–347, 1919.
- [60] Gareth O Roberts and Jeffrey S Rosenthal. Geometric ergodicity and hybrid Markov chains. *Electronic Communications in Probability*, 2:13–25, 1997.
- [61] Wolfgang Sendlar. A note on the proof of the zero-one law of Blum and Pathak. *The Annals of Probability*, 3(6):1055–1058, 1975.
- [62] Robert L Smith. Monte Carlo procedures for generating random feasible solutions to mathematical programs. In *A Bulletin of the ORSA/TIMS Joint National Meeting*, Washington, DC, 1980. University of Pittsburgh.
- [63] Robert L Smith. Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research*, 32(6):1296–1308, 1984.
- [64] Chris Soteros. Personal Communication.
- [65] Terence R Strick, Vincent Croquette, and David Bensimon. Single-molecule analysis of DNA uncoiling by a type II topoisomerase. *Nature*, 404(6780):901–904, 2000.
- [66] Tatsuru Takakura. Intersection theory on symplectic quotients of products of spheres. *International Journal of Mathematics*, 12(1):97–111, 2001.
- [67] Luke Tierney. Markov chains for exploring posterior distributions. *Annals of Statistics*, pages 1701–1728, 1994.
- [68] Leslie Ronald George Treloar. The statistical length of long-chain molecules. *Transactions of the Faraday Society*, 42:77–82, 1946.
- [69] Rocco Varela, Kenneth Hinson, Javier Arsuaga, and Yuanan Diao. A fast ergodic algorithm for generating ensembles of equilateral random polygons. *Journal of Physics A: Mathematical and Theoretical*, 42(9):095204, 2009.
- [70] Alexander V Vologodskii, Vadim V Anshelevich, Alexander V Lukashin, and Maxim D Frank-Kamenetskii. Statistical mechanics of supercoils and the torsional stiffness of the DNA double helix. *Nature*, 280(5720):294–298, 1979.
- [71] Johan Wästlund (mathoverflow.net/users/14302). A random walk with uniformly distributed steps. MathOverflow. <http://mathoverflow.net/questions/94262> (version: 2012-04-17).
- [72] Gijs JL Wuite, Steven B Smith, Mark Young, David Keller, and Carlos Bustamante. Single-molecule studies of the effect of template tension on T7 DNA polymerase activity. *Nature*, 404(6773):103–106,

- 2000.
- [73] Laura Zirbel and Kenneth C Millett. Characteristics of shape and knotting in ideal rings. *Journal of Physics A: Mathematical and Theoretical*, 45(22):225001, 2012.