

Stony Brook University



OFFICIAL COPY

The official electronic file of this thesis or dissertation is maintained by the University Libraries on behalf of The Graduate School at Stony Brook University.

© All Rights Reserved by Author.

THE EVENTUAL GAUSSIAN
DISTRIBUTION OF SELF-INTERSECTION
NUMBERS ON CLOSED SURFACES

A Dissertation Presented

by

Matthew Wroten

to

The Graduate School

in Partial Fulfillment of the
Requirements
for the Degree of

Doctor of Philosophy

in

Mathematics

Stony Brook University

August 2013

Matthew Murray Wroten

Dissertation
Doctorate of Philosophy

Moira Chas - Co-Advisor
Assistant Professor, Department of Mathematics

Dennis Sullivan - Co-Advisor
Distinguished Professor, Department of Mathematics

John Milnor - Chair of Committee
Distinguished Professor, Co-Director Institute for Mathematical Sciences

Martin Roček
Professor of Physics, C.N. Yang Institute for Theoretical Physics

Charles Taber
Interim Dean of the Graduate School

Abstract of Dissertation

**The Eventual Gaussian Distribution of Self-Intersection
Numbers on Closed Surfaces**

by
Matthew Murray Wroten
Stony Brook University
2013

Oriented loops on an orientable surface are, up to equivalence by free homotopy, in one-to-one correspondence with the conjugacy classes of the surface's fundamental group. These conjugacy classes can be expressed (not uniquely in the case of closed surfaces) as a cyclic word of minimal length in terms of the fundamental group's generators. The self-intersection number of a conjugacy class is the minimal number of transverse self-intersections of representatives of the class. By using Markov chains to encapsulate the exponential mixing of the geodesic flow and achieve sufficient independence, we can use a form of the central limit theorem to describe the statistical nature of the self-intersection number. For a class chosen at random among all classes of length n , the distribution of the self intersection number approaches a Gaussian when n is large. This theorem generalizes the result of Steven Lalley and Moira Chas to include the case of closed surfaces.

This dissertation is dedicated to my loving wife, Amelia Christine Wroten.

CONTENTS

Acknowledgements	vi
1. Introduction	1
2. The Markov Chain for Snakes	2
2.1. Snakes and the Fundamental Group	2
2.2. The Markov Chain	12
2.3. Infinite Snakes and Their Uniqueness Properties	13
2.4. The Existence of Infinite Snakes via Algorithm	18
2.5. Summary of Correspondence	23
3. The Self-Intersection Number	24
3.1. The Spine of a Snake	24
3.2. Calculating Linked Pairs	26
3.3. Properties of h_k on λ_n	27
4. Statistics	31
4.1. The Parry Measure	31
4.2. The U-Statistic	33
References	34

ACKNOWLEDGEMENTS

First I must give thanks to God for giving me the talent and life circumstances necessary to achieve this work.

I also thank my wife, Amelia, for all her support. She was a constant source of motivation and encouragement.

I thank my advisors, Dennis and Moira, for their guidance, direction, and mentorship.

I thank my parents for raising me and for always putting me in the best environment for my talents to thrive.

I thank my in-laws for providing me with the comfort and isolation of their home where I could write a large portion of this dissertation.

Lastly, I thank all my friends, family, and professors for their support and prayers.

1. INTRODUCTION

In 1987, Marshall Cohen and Martin Lustig studied the paths of closed geodesics on hyperbolic surfaces with boundary [3] [4]. They gave a precise combinatorial algorithm for computing the geometric self-intersection number of a closed geodesic. Closed geodesics are in one-to-one correspondence with conjugacy classes of fundamental groups. In 2009, Moira Chas ran a computer program to calculate the self-intersection number of all conjugacy classes of a given word-length for the doubly-punctured plane. She plotted the distribution of the self-intersection numbers of all (approximately) 4,000,000 curves and found it to be remarkably Gaussian in appearance. She consulted Steven Lalley and informed him of the discovery. At first he was doubtful that it could really be a Gaussian until she showed him the histogram. He said that it was the best Gaussian he'd ever seen and so they immediately got to work proving it. In August of 2011, they proved that, for a closed hyperbolic surface with boundary, the distribution of the self-intersection number (when suitably rescaled) approaches a Gaussian as the word-length of the curves approaches infinity [2].

From a combinatorial viewpoint, closed curves are simpler to analyze on a surface with boundary than on a closed surface. This is because for a surface with boundary, the fundamental group is just a free group. On a closed surface, however, there is a relationship among the generators. Because of this, there is not always a unique shortest way to represent a conjugacy class in terms of the generators. In his dissertation in 1987 [4], Martin Lustig described a unique way to codify these conjugacy classes by expressing them in terms of a larger alphabet of $16g$ symbols (as opposed to $4g$ which is the number of generators and inverses for a closed surface of genus g). He presented an algorithm to convert a conjugacy class into its unique representative and then calculate the self-intersection number.

It was natural and desirable to do a statistical analysis similar to the one done by Chas and Lalley [2] but for closed surfaces. The analogous result does hold here too: for a closed hyperbolic surface, the distribution of self-intersection numbers approaches a Gaussian as the word-length goes to infinity. First, the conjugacy classes were represented by a Markov process. Because of the complicated combinatorial nature of the problem, the Markov chain requires an astonishing $4g(6g - 1)$ states. For the simplest case, the surface of genus two, this corresponds to 88 states. The Parry measure on the Markov chains (the unique measure that maximizes entropy) describes, in a limiting sense, the uniform measure on periodic orbits (conjugacy classes of the fundamental group). Thus by analyzing a statistic such as self-intersection number on the Markov chain, we also know the distribution of the statistic over actual geodesics of a given word-length on the surface. Ergodicity and the exponential mixing of the geodesic flow, which is expressed in terms of the mixing properties of our Markov chain, allow us to achieve a sufficient amount of independence so that a central limit for Markov chains can be utilized. Finally, we have the limiting normal distribution of the intersection number.

More specifically, let S be a closed orientable surface of genus at least 2. Let $N_n(\alpha)$ denote the random variable obtained by evaluating N (the intersection number) on a uniformly chosen $\alpha \in F_n$ (the set of conjugacy classes $\pi_1(S)$ that have (minimal) combinatorial word length n). Then there exists constants κ and $\sigma^2 \geq 0$ (depending only on the genus of S) such that, as $n \rightarrow \infty$,

$$\frac{N_n(\alpha) - \kappa n^2}{n^{3/2}} \implies Normal(0, \sigma)$$

in the weak topology. This is our main theorem (Theorem 14).

2. THE MARKOV CHAIN FOR SNAKES

Throughout the paper (unless otherwise stated) we will explicitly argue the case where the genus is 2, but the arguments generalize to the cases of higher genus as well. In proving our main theorem, we will need a Markov chain whose periodic orbits correspond to conjugacy classes of the fundamental group (with standard alternating presentation) G of a given closed orientable hyperbolic surface S . Before defining the Markov chain, we will give motivation for its construction by presenting an intuitive classification of the elements of G . Then, after describing the Markov chain, we will make statements about the way it uniquely (Theorem 6) and completely (Theorem 8) represents conjugacy classes of G . Then we get the required correspondence (Theorem 9).

2.1. Snakes and the Fundamental Group. Our Markov chain will represent geometric limits of what we will call "snakes." Given two vertices in the Cayley graph of the fundamental group of our surface, we define a *snake* to be the union of all shortest paths from one vertex to the other. The length of a snake is defined as the number of edges in such a shortest path. We can give snakes a direction by declaring one vertex to be the start and the other vertex to be the finish. The n th step of a snake is defined as the union of the n th edges of each of the shortest directed paths from the start to the finish.

Although the the statement of the Main Theorem is purely topological, we will set up a specific geometry (WLOG) to help state and prove some useful theorems. For now, assume our surface has genus 2 and its fundamental group has the presentation $G = \langle a, b, c, d \mid abABcdCD = 1 \rangle$ (here and elsewhere, a capital letter denotes inverse), and that the surface is formed by identifying appropriate edges of a regular octagon in hyperbolic space with vertex angles $\pi/4$. We can allow such an octagon, F , to be the fundamental domain in the universal cover of the surface. Tessellating the universal cover with this octagon via the deck transformations of the group yields a graph, N (made of infinitely many infinite geodesics), with geodesic edges in hyperbolic space. We can embed the Cayley graph, Γ , of the group dually to N by connecting the centers of all adjacent octagons with geodesics. Let v_0 be the vertex of Γ at the center of F , corresponding to the identity element of G . There is a canonical bijection




FIGURE 1. This is G_2

between elements of G and vertices of Γ . Each vertex in Γ also corresponds to the end of a unique snake starting at v_0 . Thus the following theorem yields a complete classification of the elements of G .

Theorem 1. *The snakes in Γ of length n starting at v_0 are in bijective correspondence with the directed paths of length $n+1$ starting at St and finishing at Fi in the directed graph, G_2 (see Figure 1).*

Before proving this, we will need a lemma. Remember that the snake is made up of the union of shortest edgepaths in Γ so that knowing the structure of shortest edgepaths will tell us the structure of the snake. A shortest edgepath between two vertices of Γ is a path along the graph of Γ from one vertex to the other that traverses the minimal possible number of edges out of all paths in Γ between the vertices. We present a fundamental lemma about shortest edgepaths that will allow us to derive all of their characteristics and thus all of the characteristics of snakes.

Lemma 1 (Golden Lemma). *A shortest edgepath in Γ never crosses the same geodesic of N more than once. Further, a (directed) shortest edgepath can never cross a geodesic L of N into the half-plane bounded by L that doesn't contain its final vertex.*

Proof: Suppose E is a shortest edgepath between v_0 and v_n that crosses a geodesic, L , of N more than once. Let $v_0, v_1, \dots, v_{n-1}, v_n$ be, in order, the vertices traversed by E (we can immediately rule out the possibility that any vertex of Γ is traversed multiple times because the portion of the edgepath between the first and last such traversal could be truncated out to yield a shorter edgepath). Let v_i be the vertex of E immediately preceding the first crossing of L . Let v_j be the vertex of E immediately following the second crossing of L . The portion of E from v_{i+1} to v_{j-1} can be reflected across L to yield a path between v_i and v_j of length $j-i-2$, creating a contradiction [fig 4]. A (directed) shortest edgepath cannot cross a geodesic L of N into the half-plane bounded by L that doesn't contain its final vertex because it would eventually have to cross back, creating a contradiction ■

Given this lemma, we see that with every edge traversed along a shortest edgepath, we cross a geodesic in N , thus eliminating a half-plane from the allowable region that the rest of the edgepath can traverse. Using this idea, we can find precisely all of the shortest edgepaths from one vertex to another and thus our snake.

We will now describe the snake that starts at v_0 and ends at another vertex v_n of Γ . We will break this down into cases based on where v_n lies. Starting at v_0 , we will determine, edge by edge, where the shortest edgepath(s) can go based on the Golden Lemma. For a Markov chain of period n , we want the state of the i th step to correspond to the i th step in the snake which is a combination of the i th edge in each of the shortest edgepaths combined with information about where the next steps can be. We will build the snake step by step. At each step we will look at the final vertex or vertices of the edges in the previous step. We will then decide which edge or edges emanating from there could potentially lead via shortest edgepaths to v_n .

Case St: If $v_n = v_0$, we are done (Case Fi) so assume this is not the case. Then v_n lies outside the fundamental domain F . The 8 geodesics on the boundary of F divide the complement of F into 16 regions – 8 each of 2 types.

Case A: v_n is in one of the (open) regions with 3 boundary geodesic pieces, adjacent to F . In this case, any shortest edgepath must immediately cross into that region, otherwise it would cross a different geodesic of N and never be able to enter the region without recrossing it, violating the Golden Lemma (Figure 2). We will describe what follows Case A after discussing the other 8 regions.

In all of our diagrams, solid lines are part of Γ , dashed lines are part of N , the edge being discussed will be bold and marked with an arrow, and the region being discussed will be shaded.

Case B: v_n is in one of the (closed) regions with two boundary geodesic rays, adjacent to two of the regions of the Case A type. In this case, a shortest edgepath can cross first into either of the adjacent regions (but nowhere else) so there can (and




FIGURE 2. this is a Case A region

will) be multiple edgpaths in the snake (see Figure 3). From Case B, the next step in the snake is predetermined. By the Golden Lemma, the two edgpaths in the snake must continue around the octagon so that they can eventually end up in the Case B region. Thus we are immediately forced to a Case C.

Case C: v_n is still in the (closed) region described in Case B, but now the edgpaths of the snake need to progress one further step to opposite vertices of an octagon (see Figure 4). Again the edgpaths in the snake must continue around the octagon so we are forced into a Case C'.

Case C': v_n is still in the same (closed) region described as in Case C, but now the edgpaths have traversed one edge closer to (and one edge away from) the final vertex of the aforementioned octagon (by the Golden Lemma they must in order to eventually reach the region) (see Figure 5). Here there will be choices for the next step depending on where in the region of C' (or C or B) v_n lies. If v_n is in one of




FIGURE 3. this is a Case B region




FIGURE 4. this is a Case C region

the two wedge shaped regions on the left or right side of the C' region, then we must proceed to Case E. Otherwise we proceed to case D.

Case D: v_n is in the (open) region formed by removing two wedges from case E regions. Here, by the Golden Lemma, both branches of the snake must continue to the final edge of the octagon (see Figure 6). Before describing where the snake can go from here, we will continue describing the other cases.

Case E: v_n is now in a (closed) region shaped like the Case B region. If in the left (right) wedge from Case C', the edgpath on the right (left) branch of the snake




FIGURE 5. this is a Case C' region subdivided into its two subregions

must go to the final vertex of the aforementioned octagon. However, the left (right) branch can get to v_n either via that vertex, or the one to the left (right) (see figure 7). From here there will once again be two branches of the snake and we are forced into a Case C.

We return to the discussion of Case A. Once in case A, there are 7 edges that can be traversed during the next step (the eighth edge (the one we just traversed) is ruled out by the Golden Lemma). The edge(s) taken next is(are) determined by where in the region of Case A v_n lies. We consider the octagon of N containing the final vertex of the edge just taken. If we extend the geodesic segments that make up the boundary of this octagon, the region of Case A is divided into 14 pieces (see Figure 8). One is the current octagon. If v_n is there, then we are done (Case Fi). 6 of the regions correspond to entering Case B, and 5 of the regions correspond to




FIGURE 6. this is a Case D region

entering another Case A. The regions on the left and right extreme are almost like Case A, but a wedge has been removed from the far side (when in the previous step we assumed we were in Case A and not B). We will call such a region Case A'.

Case A': v_n is in a (open) region shaped like a Case A region but with a wedge removed from the left (right) (see Figure 9). By the Golden Lemma, the snake must immediately cross into this region. Again we extend the geodesics of the octagon of N surrounding the final vertex of the edge just taken. This divides the A' region again into 14 pieces. If v_n is in the octagon, we are done (Case Fi). 6 regions lead to Case B and 5 regions lead to case A. The region on the extreme right (left) leads to another Case A', but the region on the extreme left (right) will lead to a narrower A''.

Case A'': This is (open) shaped like Case A' but with the final edge on the left (right) removed. Again by the Golden Lemma, the snake must cross into this region immediately. Dividing the region as before yields only 12 regions (see Figure 10). If v_n is in the octagon, we're done (Case Fi). Apart from that, 5 regions (including the one on the far left (right)) lead to Case A, 5 regions lead to Case B, and the region on the far right (left) leads to another Case A'.

We conclude by finishing the discussion of Case D. After splitting up the region by extending geodesics segments of the boundary of the octagon of N containing the




FIGURE 7. this is a Case E region

vertex where the previous to branches of the snake merged, there are 12 pieces (see Figure 11). If v_n is in the octagon in the middle, we're done (Case Fi). Apart from that, 5 regions lead to Case B, 4 regions lead to Case A, and the regions on the extreme left and right lead to Case A'.

G_2 is constructed from the case by case scenario above as follows. Each Case is given a vertex of the graph with the same label. If there are x ways that Case Y leads directly to Case Z (corresponding to x Case Z shaped regions in a Case Y shaped region), then we draw x directed edges from vertex Y to vertex Z . We label each of these edges YZ_1, YZ_2, \dots, YZ_x . We say YZ_1 represents entering the Case Z shaped region furthest to the left relative to the direction the snake is progressing (Excepting with the StA and StB edges where an arbitrary starting point must be chosen). Then the rest of YZ_2, \dots, YZ_x represent in clockwise order the other Case Z shaped regions.

Proof of Theorem 1: First we show how to construct the snake corresponding to a directed path P in G_2 from St to Fi. Start at v_0 . Whenever we add an edge to P , this will correspond to the addition of a step to the snake and the entering of the




FIGURE 8. this is a Case A region subdivided into 14 pieces corresponding to the different possible subsequent regions




FIGURE 9. this is a Case A' region subdivided into 14 pieces corresponding to the different possible subsequent regions




FIGURE 10. this is a Case A'' region subdivided into 12 pieces corresponding to the different possible subsequent regions

corresponding region, unless the edge leads to F_i , in which case the snake stops there (we immediately have the correspondence between the length $n + 1$ of a path in G_2 and the length n of a snake). Each successive region is nonempty, and when F_i is reached, a specific vertex is chosen in Γ . Thus our correspondence is well-defined. To check that the correspondence is injective, one only needs to observe that in each Case, the region is split into non-overlapping sub-regions. All that remains is to check whether the correspondence is onto.

Choose a vertex v in Γ . Γ is connected so there exist finite edgepaths from v_0 to v . All of these edgepaths have finite, non-negative, integer lengths and since the non-negative integers are well ordered, there exists a (not necessarily unique) edgepath of minimal length. Call this length n . Proceed through the case by case procedure above, building a snake from v_0 toward v . This cannot end before the $n + 1$ step or we would have a path of length less than n . At step j , our procedure has constructed all edgepaths of length j that could possibly be a part of a shortest edgepath from v_0 to v by pruning only edges (and thus paths) that are ruled out by the Golden Lemma. Thus, at least one of the end-vertices of the n th step of the snake must be v . Suppose that there are two end-vertices of the n th step of the snake. Then the n th step would have to be a Case B, Case C, Case C', or Case E. However in each of these cases, it is assumed that the end-vertices of the step are not contained in the allowable region, but that v is, which is a contradiction. Thus the n th step must be a Case St, Case A, Case A', Case A'', or Case D, so the next step can be Case F_i (it




FIGURE 11. this is a Case D region subdivided into 12 pieces corresponding to the different possible subsequent regions

is worth noting here that this also leads to the converse of the Golden Lemma). This completes the proof ■

Following the above analysis, one can easily derive a generalized version of Theorem 1: The elements of the fundamental group of the closed surface of genus g with wordlength n according to the presentation

$$G = \langle x_1, x_2, \dots, x_{2g} \mid x_1 x_2 X_1 X_2 \dots x_{2g-1} x_{2g} X_{2g-1} X_{2g} = 1 \rangle$$

are in bijective correspondence to directed paths from St to Fi in the graph G_g .

2.2. The Markov Chain. Here we will describe in detail the Markov states for the surface of genus 2. We will use explicitly the presentation of the group $G = \langle a, b, c, d \mid abABcdCD = 1 \rangle$. Let Ω be the set of generators of G along with their inverses.

Each state in the Markov chain represents a step in a bi-infinite snake. This means that it represents the combination of the information of a region, and the directed edge or edges of that step.

Case A: This includes 8 different Markov states $A_k, k \in \Omega = a, b, c, d, A, B, C, D$. Here, k is simply the label of the edge traversed in this step.

Case B: This includes 8 different Markov states B_{jk} where j and $k \in \Omega$ are the edges traversed in this step with j on the left relative to the direction of traversal.

Case C: This includes 8 different Markov states C_{jk} where j and $k \in \Omega$ are the edges traversed in this step with j on the left relative to the direction of traversal.

Case C': This includes 8 different Markov states C'_{jk} where j and $k \in \Omega$ are the edges traversed in this step with j on the left relative to the direction of traversal.

Case D: This includes 8 different Markov states D_{jk} where j and $k \in \Omega$ are the edges traversed in this step with j on the left relative to the direction of traversal.

Case E: This includes 16 different Markov states E_{ijk}^L and E_{ijk}^R where $i, j,$ and $k \in \Omega$ are the edges traversed in this step with i on the left of j which is to the left of k relative to the direction of traversal. The L superscript is used when the j edge starts at the start of the i edge and ends at the end of the k edge, and the R superscript is used in the other case.

Case A': This includes 16 different Markov states $A'_k{}^L$ and $A'_k{}^R, k \in \Omega$. Here, k is the edge traversed in this step. The L superscript is used when a wedge has been cut from the left side (relative to the direction of the edge) of the Case A region to make this Case A' region and the R superscript is used in the other case.

Case A'': This includes 16 different Markov states $A''_k{}^L$ and $A''_k{}^R, k \in \Omega$. Here, k is the edge traversed in this step. The L superscript is used when a wedge has been cut from the left side (relative to the direction of the edge) of the Case A' region to make this Case A'' region and the R superscript is used in the other case.

Let \mathbb{A} , this set of 88 Markov states, be the alphabet for our Markov chain. Let Π be a map from \mathbb{A} to the vertices of G_2 that simply removes the subscripts and superscripts. The transition matrix for our Markov chain will now be described. Choose $J, K \in \mathbb{A}$. Then J can be followed by K if and only if there is a directed edge in G_2 from $\Pi(J)$ to $\Pi(K)$, and a choice of images in Γ of J and K can be made (edges and corresponding regions) such that the edge(s) of K start where the edge(s) of J end and the region of K is contained in that of J as one of the specific subregions corresponding to the choice of an edge in G_2 .

2.3. Infinite Snakes and Their Uniqueness Properties.

Definition 1 (Infinite Snake). *An Infinite Snake is the image of a Markov chain in the universal cover.*

Lemma 2. *An Infinite Snake never contains four or more consecutive edges of an octagon in Γ without containing all 8.*

Proof: Assume there are between 4 and 7 consecutive edges of a given octagon in Γ but not an 8th. The first 4 edges can be ordered 1 through 4 depending on their ordering of traversal. The octagon cannot be left and then returned to without violating the Golden Lemma. The second edge must be either a Case A' or a Case

B (if it were a Case A then the previous edge couldn't be in the octagon, if it were a Case A" then it must be at least the 3rd edge, and if it were a Case C, C', D, or E, then it would violate our assumption about the octagon being incomplete). Case B is ruled out now because then the 3rd edge cannot be in the same octagon. Thus the third edge must be a Case A". The corresponding region doesn't include the 4th edge so we are done ■

Theorem 2. *An Infinite Snake has an inverse which is also an Infinite Snake.*

Proof: We will construct the inverse of the Infinite Snake and show that it is indeed the image of a Markov chain. Let $\mathbb{J} = \dots J_{-2}J_{-1}J_0J_1J_2\dots$ be our Markov chain. We will make its inverse $\mathbb{K} = \dots K_{-2}K_{-1}K_0K_1K_2\dots$. Let $e = \{e_i\}_{-\infty}^{\infty}$ and $R = \{R_i\}_{-\infty}^{\infty}$ be edges and regions of an Infinite Snake image of \mathbb{J} where e_i is(are) the edge(s) corresponding to the i^{th} step in the Infinite Snake (the image of J_i), and R_i is the corresponding region.

Define $f = \{f_i\}_{-\infty}^{\infty}$ such that $f_i = (e_{-i})^{-1}$ where inverse here simply means traversing the edge or edges in the opposite direction. Whenever $\Pi(J_i)$ does not equal A, A', or A", it is straightforward to deduce what K_{-i} need be, as there is only one state that has the required set of edges to project to f_{-i} . The following rules hold:

- If $J_i = B_{jk}$, then $K_{-i} = D_{k^{-1}j^{-1}}$
- If $J_i = D_{jk}$, then $K_{-i} = B_{k^{-1}j^{-1}}$
- If $J_i = C_{jk}$, then $K_{-i} = C'_{k^{-1}j^{-1}}$
- If $J_i = C'_{jk}$, then $K_{-i} = C_{k^{-1}j^{-1}}$
- If $J_i = E_{jkl}^{R(L)}$, then $K_{-i} = E_{l^{-1}k^{-1}j^{-1}}^{R(L)}$

Now suppose $\Pi(J_i)$ equals A, A', or A". If no octagon of Γ has both edge e_i and any edge of e_{i+1} as boundary components, then $K_{-i} = A_{j^{-1}}$ where $e_i = j$. If an octagon of Γ has both edge e_i and any edge of e_{i+1} as boundary components, but no edge of e_{i+2} , then $K_{-i} = A'_{j^{-1}}^{R(L)}$ where $e_i = j$ and the octagon is on the left(right) of the directed edge e_i . Lastly, if an octagon of Γ has edge e_i and any edge of e_{i+2} as boundary components, then $K_{-i} = A''_{j^{-1}}^{R(L)}$ where $e_i = j$ and the octagon is on the left(right) of the directed edge e_i .

We must now verify that \mathbb{K} as constructed satisfies the Markov chain transition conditions. In particular, we will check that K_{-i} can follow K_{-i-1} for arbitrary i . Clearly, as the edge or edges e_{i+1} continue those of e_i , so must the edges f_{-i} follow those of f_{-i-1} . There are only finitely many possibilities for J_i through J_{i+4} . By the above lemma, this sufficiently determines what K_{-i-1} and K_{-i} must be. In each of these cases, one can see there is a directed edge in G_2 from $\Pi(K_{-i-1})$ to $\Pi(K_{-i})$ with the region corresponding to f_{-i} being the appropriate subregion of f_{-i-1} ■

Theorem 3. *If $\mathbb{J} = \dots J_{-2}J_{-1}J_0J_1J_2\dots$ and $\mathbb{K} = \dots K_{-2}K_{-1}K_0K_1K_2\dots$ are two Markov chains that have Infinite Snake images with the same edges $e = \{e_i\}_{-\infty}^{\infty}$ where e_i corresponds to both steps J_i and K_i , then \mathbb{J} and \mathbb{K} are identical.*

Proof: When $\Pi(J_i) \in \{B, C, C', D, \text{ or } E\}$, $J_i = K_i$ trivially. If not, one can use the above lemma to see that $\#\{j < i \mid \text{an edge of } e_j \text{ is on the boundary of one of the same octagons as } e_i\} < 3$. From the definition of the transition matrix, one can deduce that this number must also be the number of "primes" following the A of J_i and J_k . The superscript of these states, if present, is determined by the which side of e_i the common octagon is on, and the subscript is determined by the label of e_i ■

The following theorem describe the sense in which Infinite Snakes are geometric limits of snakes.

Theorem 4. *Any finite segment of an Infinite Snake can be extended (by a bounded amount) to a snake between two points in Γ .*

Proof: Let $\mathbb{J} = \dots J_{-2} J_{-1} J_0 J_1 J_2 \dots$ be a Markov chain who has our Infinite Snake as an image. Let $e = \{e_i\}_{-\infty}^{\infty}$ and $R = \{R_i\}_{-\infty}^{\infty}$ be edges and regions of the Infinite Snake image of \mathbb{J} where e_i is(are) the edge(s) corresponding to the i^{th} step in the Infinite Snake (the image of J_i), and R_i is the corresponding region. Let $\{e_i \mid a \leq i \leq b\}$ be the segment of interest.

We must now find appropriate beginnings and ends of our snake. If $\Pi(J_a)$ is B or A, then e_a has one initial vertex which we can let be the start of the snake. If $\Pi(J_a)$ is A' or A'', we can prepend the snake with a Case A or a Case A and Case A' respectively and then the initial vertex of the Case A can be the start of the snake. If $\Pi(J_a)$ is C, C', D, or E, we can prepend as necessary until we prepend a Case C with a Case B, whose initial vertex can be the initial vertex of the snake.

Now to find an appropriate end vertex of our snake, if $\Pi(J_b)$ is D, A, A', or A'', then e_b has one final vertex which we can let be the end of our snake. If $\Pi(J_b)$ is B, C, C', or E, we can postpend the snake with steps until we reach a Case C' which can then be followed by a Case D. Its final vertex will be the end of our snake.

Now that we have a start vertex and an end vertex, we can verify that what we have is indeed a snake between them. We will do this by showing that we have a sequence of steps corresponding to a directed edgepath in G_2 from St to Fi. Start at St. The next vertex in G_2 will be A or B by construction. The edge from St to A or B is chosen appropriately according to the lables of the initial edge or edges in the snake. We continue making the directed edgepath by appropriately connecting the prepended Cases until we get to $\Pi(J_a)$. Then, for $a \leq i < b$ we can choose the edge in G_2 from $\Pi(J_i)$ to $\Pi(J_{i+1})$ because of the following: by definition of the transition matrix for our Markov chain, R_{i+1} is one of the m Case $\Pi(J_i + 1)$ regions in the Case $\Pi(J_i)$ region, $R(i)$. Each of these m regions correspond precisely to one of the m edges in G_2 from $\Pi(J_i)$ to $\Pi(J_{i+1})$ so we choose the appropriate one. Then if necessary, we choose the appropriate edges to any postpended Cases and then from the final Case A, A', A'', or D to Fi. We now have that our segment of the Infinite snake is a piece of the snake corresponding (up to deck tranformation of the initial vertex) to a directed path in G_2 from St to Fi ■

Corollary 1. *Given a geodesic L of N and an Infinite Snake, there is at most one direction in which edges of that Infinite Snake can cross L .*

Proof: Let $\mathbb{J} = \dots J_{-2}J_{-1}J_0J_1J_2\dots$ be a Markov chain who has our Infinite Snake as an image. Let $e = \{e_i\}_{-\infty}^{\infty}$ be edges of the Infinite Snake image of \mathbb{J} where e_i is(are) the edge(s) corresponding to the i^{th} step in the snake (the image of J_i). Suppose that, for some Geodesic L of N , an edge of e_a crosses L in one direction, but an edge of e_b crosses it in the other direction. WLOG assume $a \leq b$. Then by the previous theorem, there is a snake between some vertices v_0 and v_n containing the segment $\{e_i | a \leq i \leq b\}$ of the Infinite Snake. v_n has to be on one particular side of L . Thus, one of the two traversing edges must be in violation of the Golden Lemma ■

Theorem 5. *An Infinite Snake has exactly two limit points on the boundary of \mathbb{H}^2 .*

Before we prove this, we will need the following lemma.

Lemma 3. *Given any two distinct points on the boundary of \mathbb{H}^2 there exists a geodesic L of N that separates the two. The same holds for a point on the boundary and a point in the interior.*

Proof: Consider the geodesic l between the two points. l has infinite hyperbolic length so it cannot be contained in the interior of any of the finite-sized hyperbolic octagons bounded by geodesics of N . Therefore, it must cross one of these geodesics. l must then intersect some geodesic L of N transversely (if it intersected some geodesic of N non-transversely, then it would be a geodesic of N , each of which intersect infinitely many other geodesics of N transversely). Since l and L are geodesics that intersect transversely in hyperbolic space, they must intersect in exactly one point. Therefore the two limit points of l lie on opposite sides of L . The proof of the second statement follows similarly ■

Proof of Theorem 5: Let $\mathbb{J} = \dots J_{-2}J_{-1}J_0J_1J_2\dots$ be a Markov chain who has our Infinite Snake as an image. Let $e = \{e_i\}_{-\infty}^{\infty}$ be edges of the Infinite Snake image of \mathbb{J} where e_i is(are) the edge(s) corresponding to the i^{th} step in the snake (the image of J_i). For $0 \leq i < \infty$ let v_i be the the (left-most relative to the direction of the edges if there are multiple) final vertex of e_i . The sequence $e = \{e_i\}_0^{\infty}$ must have accumulation points in the closure of the Poincare disk. By the above theorem, since no snake ends two distinct steps on a given vertex (this would violate the shortest edgepath part of the definition), the vertices of e must be unique. Thus, since the only accumulation points of vertices of Γ are in the boundary of the disk, the same must hold for e .

Now suppose that b_1 and b_2 are distinct accumulation points of e . Then by the Lemma, there is a geodesic L of N separating them. Since points in the sequence e get arbitrarily close (in the Euclidian metric) to both b_1 and b_2 , there must be $i < j < k$ such that v_i and v_k are on one side of L but v_j is on the other. This leads to a violation of the above Corollary. Thus there is a unique accumulation point and it must be a limit point. We will call this the "forward limit point."

By constructing the inverse of the Infinite Snake via the previous theorem, one can find another limit point at the other end of the Infinite Snake. We will call this the "backward limit point." Suppose that they are the same point b_0 . Now via the above lemma construct L' separating v_0 and b_0 . By a similar argument to the above paragraph, one can again construct a contradiction to the above Corollary. This completes the proof ■

Definition 2 (Proper Infinite Snake). *A Proper Infinite Snake is an Infinite Snake such that all of its regions contain one of its limit points and the regions of its inverse all contain the other limit point.*

Theorem 6. *Given any two distinct points f and b on the boundary of \mathbb{H}^2 that do not bound a geodesic of N , if S and S' are Proper Infinite Snakes with f as forward and b as backward limit points, then S and S' are the same.*

Proof: Let $\mathbb{J} = \dots J_{-2}J_{-1}J_0J_1J_2\dots$ be a Markov chain with S as image. Let $e = \{e_i\}_{-\infty}^{\infty}$ be edges of the Infinite Snake image of \mathbb{J} where e_i is(are) the edge(s) corresponding to the i^{th} step in the snake (the image of J_i). We will divide this problem into three cases.

Case 1: There exists i such that $\Pi(J_i)$ is B or D. By considering inverses, we can assume it is D WLOG. Let R_i be the corresponding region for this step. f is clearly in R_i . Consider the two geodesics that contain boundary rays of R_i . Let U be the open region between them. Since U contains the Case B region (call it T) corresponding to the inverse of step i , b must also be in U . By the Corollary, S' must then be contained in U . Let v be the final vertex of e_i and let V be the octagon of N containing it. Since V separates the parts of U containing f and b , S' must pass through V and thus contain v . The only step with edges in U that can have v as a final vertex, f contained in the corresponding region, and b contained in the inverse step's region is the Case D step of S . Thus S' contains this step. Since at every step, only one next step has a region that also contains f , the subsequent steps of S' must be identical to those of S . To determine what the preceding steps must be, one can use this same argument on the inverses of S and S' to first determine that the edges must be the same. Using the process for constructing inverses of Infinite Snakes, we then uniquely construct the inverse of the inverse of the remaining part of S' . The steps themselves (in particular whether an A, A', or A'' is present) are uniquely determined from the edges.

Case 2: There exists i such that $\Pi(J_i)$ is E. WLOG (by mirror symmetry) we can assume that the superscript of J_i is L . Let R_i be the region for this step, and let T be the inverse region for this step. Consider the region V made up of the two octagons of N that each intersect two edges of e_i . Let L be the geodesic of N on the common boundaries. Let U be the region bounded between the two geodesics of N which contain the arc boundaries of these two octagons that are farthest from each other. Since $f \in R_i \subset \text{Int}(U)$ and $b \in T \subset \text{Int}(U)$, S' must also be in U by

the corollary. As in Case 1, since V separates the portions of U containing b and f , S' must pass through V and thus must contain at least one of the vertices of Γ contained in V . By rotational symmetry and by considering inverses, we can assume WLOG that S' contains the vertex v_r in the center of the right (relative to direction of S) octagon of V . There are only two possible steps with edges in U that can have v_r as an initial vertex, f contained in the corresponding region, and b contained in an inverse step's region. These are the Case E step of S and a Case C step. However based on the regions corresponding to the Case C step and its Case C' inverse, the only place f and b could be (and still be consistent with S) are on opposite ends of L , a contradiction. Thus S' contains the Case E step of S . As in Case 1, we can then deduce further that S is identical to S' .

Case 3: For all i , $\Pi(J_i) \in \{A, A', A''\}$. This can be divided into two sub-cases.

Case 3a: There exists j such that $\Pi(J_j)$ is A' . WLOG assume the supercript of J_j is L , and let $\gamma \in N$ be the geodesic containing the left ray boundary of the corresponding region. Let v be the initial vertex of e_j and let V be the octagon of N containing it. e_{j-1} must be the edge that enters V immediately counterclockwise from where e_j leaves it. Now e_{j-2} cannot be the edge crossing γ and ending at the initial vertex of e_{j-1} , otherwise $\Pi(J_j)$ would be A'' . Thus the inverse edge of e_{j-1} must have its region on the V side of γ . b is in this region, and f , in the region for e_i , must be on the same side of γ . Let U be the region bounded by γ and the six geodesics that contain boundary arcs of V not crossed by e_j or e_{j-1} . f and b are in U , but are on opposite sides of V . Thus, by the corollary, S' must pass through V and contain v . The edges of S' can be uniquely constructed outwardly from v since at every vertex along the way, f and b are contained in the emanating regions of S and its inverse and there is only one edge leading from the vertex into each of those. The edges then determine the steps of S' uniquely.

Case 3b: There exists k such that $\Pi(J_k)$ is A . Let v be the initial vertex of e_k and let V be the octagon of N containing it. Notice that e_{k-1} and e_k do not intersect adjacent edges of V . Thus, the region U bounded by the 6 geodesics containing the other edges of V is disconnected when V is removed. The two disconnected regions must each contain either b or f since they contain the regions corresponding to e_k and the inverse of e_{k-1} . The argument proceeds as in Case 3a and we are done ■

2.4. The Existence of Infinite Snakes via Algorithm. We are trying to establish a correspondence between periodic elements of our Markov chain and conjugacy classes of G . The previous result gives us a form of uniqueness for such a Markov Chain. Now we need to show the existence of such a Markov chain given a conjugacy class. To this end, we will describe an algorithm, nearly identical to Lustig's [4] but adapted to our language, for creating an Infinite Snake (that is an image of a Markov Chain) which has the same limit points as a lift of a given conjugacy class of G . This will require an easier way to recognize Infinite Snakes.

First, we will describe topological or graphical structure of an Infinite Snake called a Topologic Snake.

Definition 3 (Topological Snake). *A Topological Snake is a (finite, infinite, or bi-infinite) sequence of arcs and/or bubble chains. A bubble chain is a sequence (finite, infinite, or bi-infinite) of bubbles (topologic circles) where each bubble shares a single proper topological interval with each neighboring bubble in the sequence, but no points with any other bubbles. In each bubble chain, we assign a marked point to the part of the first (last) bubble disjoint from "boundary intervals", and we call this the entry (exit) point. In case the bubble chain contains only one bubble, we require its entry point and exit point be distinct. The arcs are directed, meaning they also have an entry point and an exit point (the endpoints). The entry point of each bubble chain or arc is identified with the exit point of the preceding bubble chain or arc.*

Now we will describe sufficient conditions for an embedding of a Topological Snake in Γ to be an Infinite Snake. We will call such an embedding a Geometrical Snake.

Definition 4 (Geometrical Snake). *A Geometrical Snake is an embedding in Γ of a Topological Snake that satisfies the following 4 Properties:*

1. *Each bubble is mapped to an octagon of Γ .*
2. *Each bubble must have two "special points" chosen from entry points, exit points, or endpoints of boundary intervals. These two points must embed to opposite vertices of an octagon. If the bubble has an entry point and/or exit point (if it is the first and/or last in its bubble chain), such point(s) must be the special point(s).*
3. *Two bubbles in the same bubble chain cannot share a special point.*
4. *If the embedding contains four or more edges of an octagon in Γ , it contains all eight.*

Theorem 7. *A Geometrical Snake that isn't finite is an image of a Markov chain and must also be an Infinite Snake.*

Proof: We will show how to construct a Markov chain whose image is a given Geometrical Snake, S . First we partition the edges of S into steps. Each edge in the image of one of the arc sections of the Topological Snake will be in its own step of type A, A', or A". For edges in the images of bubbles, we consider the geodesic segment between the special points of that bubble. We group two edges of that bubble in the same step if they are mirror images of each other about that segment. Properties 1 and 2 allow this to be clearly defined. If no edge in one of these steps is a boundary interval, we have a unique choice of a B, C, C', or D step. If an edge is a boundary interval, then it is grouped with its reflection in each of the two adjacent octagons. Property 3 ensures that these three edges form an E step. As in the construction of inverse snakes, we can use Property 4 to distinguish between steps of type A, A', and A", and then check that steps of types D, B, A, A', and A" follow each other in ways consistent with the transition matrix. Steps of type B,

C, C', D, and E follow each other in ways consistent with the transition matrix by construction and by Properties 1, 2, and 3 ■

It is now time to present an algorithm that converts a conjugacy class into a Geometrical Snake. We start with a cyclic word in the generators of G . The algorithm first uses the relation to shorten the word if possible, and then it adds other edges to include other loops in the conjugacy class of the same length. The next theorem will prove that this algorithm always yields the properties characteristic of a Geometrical Snake.

Let $W = w_1w_2\dots w_l$ be our (nontrivial) cyclic word in the generators of G . The indices here are understood modulo l , the current wordlength (which can shrink as we use relations). The algorithm, which will be proved to be finite, will consist of five steps.

Step 1. Scan W for instances where w_i and w_{i+1} are inverses and remove them. Iterate this process until there are no more such instances. Call the resulting cyclic word W' .

Step 2. Scan W' for any "long cycles" and replace them with the corresponding equivalent "short cycle". A "long (short (half)) cycle" is a sequence of generators of G that comprise more than (less than (exactly)) half of a permutation of the relation or its inverse. Repeat Step 1 if necessary. Then iterate this process until there are no more long cycles. Call the resulting cyclic word $W'' = w''_1w''_2\dots w''_l$.

Step 3. Scan W'' for any half cycle $w''_i\dots w''_{i+3}$. Replace it with the equivalent half cycle $v_1v_2v_3v_4$ if and only if $v_4w''_{i+4}w''_{i+5}w''_{i+6}$ is again a half cycle. Repeat this procedure with the new half cycle and so on (the next theorem will prove that this process is finite). Any time a long cycle or neighboring inverses appear, eliminate them as in Step 2. Iterate this process until all half cycles have been checked and, when needed, replaced. Call the resulting cyclic word W''' .

Step 4. Let $E = \dots e_{-1}e_0e_1\dots$ be an edgepath in Γ that periodically reads off the cyclic word W''' .

Step 5. Any time there are 4 consecutive sides of an octagon in E , add the other 4 and do this for the corresponding octagon in each of the other periods, creating a new Topological snake. Iterate this process for each new Topological Snake until it cannot be done anymore. Call the final collection of edges GS .

Theorem 8. *The above process yields a Geometrical Snake that has the same limit points as a lift of W .*

Before proving this theorem, we will need a lemma about the minimality of E .

Lemma 4. *E is locally minimal in the sense that any finite piece of it $e_i e_{i+1} \dots e_{j-1} e_j$ is a shortest edgepath.*

Proof: Steps 1 and 2 are clearly finite since every replacement involves a shortening. In step 3, we need to insure that no replacement of a half cycle creates another replacement which creates another replacement and so on endlessly. Note that when

one replacement gives rise to another, both replacements must have the same orientation (both changing clockwise cycles to counter-clockwise or vice versa). This must be the case since Steps 1 and 2 keep the word reduced. After continuing such consecutive replacements once through the length of the word, we come across a cycle that has already been switched to the opposite orientation, so the procedure must halt. When we change one half cycle, at most two new half cycles can be created, one before and one after. If one is created before, it can be immediately confirmed that it does not need to be replaced. If a half cycle is replaced because it creates a new one after, we can check that it does not need to be replaced again. If the new succeeding half cycle does not need to be replaced, then we have a net decrease in the number that need to be checked. Otherwise we continue with the next one and the next one, but this can't go on forever, and when it stops, the net number that need to be checked decreases. Thus, the net number of half cycles that need to be checked can never increase, but must eventually decrease, so the process terminates.

We will now introduce some terminology of Birman and Series [1] so that we can use a theorem of theirs to prove the lemma. Two cycles c_1 and c_2 are called "consecutive" if they lie on the boundaries of different octagons in Γ sharing a common edge E , such that c_1 ends on the endpoint of E where c_2 begins, but neither cycle contains the other endpoint of E . A sequence of consecutive cycles is called a "chain". A "long chain" is a chain of $n > 1$ cycles, c_1, \dots, c_n , where c_1 and c_n are half cycles and all the rest have one less edge than these". Such a chain of $3n + 2$ can be replaced by the chain of length $3n$ running along the other side of the octagons, and this is exactly what Step 3 does.

A special case of Theorem 2.8(a) of Birman and Series [1] says that any edgpath that doesn't contain a long chain or long cycle is minimal. Thus we only need to show that E contains no long chains or long cycles to prove the lemma. Suppose that E does contain a long cycle. Then W''' must be at least as long as this cycle because the relation doesn't repeat symbols. Thus the long cycle would be contained in a permutation of W''' , contradicting Step 2. Now suppose E contains a long chain. Since it is finite, it must be contained in an edgpath reading off $(W''')^n$. After cyclic permutation, we can assume WLOG that the chain starts at the beginning of $(W''')^n$. The chain must be longer than W''' or we'd have a contradiction of Step 3. Thus we can infer that W''' is one chain consisting of a half cycle followed cycles of one less length. Step 3 would eliminate this scenario too ■

Now we prove Theorem 8: $E = \dots e_{-1}e_0e_1\dots$ must be an imbedded arc by the previous theorem so it is a Topological Snake. It also vacuously satisfies Properties 1, 2, and 3 of being a Geometrical Snake. We must show that Step 5 is finite and that each addition of an octagon does not alter the facts that Properties 1, 2, and 3 are satisfied and that it is a Topological Snake. Also, on the way to satisfying Property 4, we will ensure the weaker condition, which we call Property 4', that at every step, there are never 5-7 consecutive edges of an octagon without all 8 being present. For ease of bookkeeping, when we complete octagons, we will group the new edges with

their mirror images as in the proof of the previous theorem. Let $F = \dots f_{-1} f_0 f_1 \dots$ where f_i contains e_i and any new edge grouped with it. It will follow inductively that for any new edge added to f_i , there will exist $m < \infty$ such that there is an edgepath of length $m + 1$ starting with the e_{i-m} and ending with that new edge of f_i and similarly there is an edgepath of length $m + 1$ starting with this new edge and ending with e_{i+m} . Call this Property 5. We will also maintain property called Property 6 that any time some f_a contains two edges of an octagon, F must contain all of them. We will now break down the completion of octagons in Step 5 into three types based on the number of adjacent (sharing an edge) octagons already completed.

Type 0: we complete an octagon that shares no edges with any previously completed octagon. To do this, we start with a half cycle $f_i \dots f_{i+3} = e_i \dots e_{i+3}$ that is the image of an arc or a piece of an arc of our Topological Snake. The four new (all four are new by Property 4') edges add a bubble chain of one bubble to our Topological Snake. They also are added simultaneously to ensure Property 6 remains satisfied. This bubble is mapped to an octagon, consistent with Property 1. The entry and exit point are opposite and can be chosen to be the special points consistent with Property 2. Property 3 is still vacuously true for this bubble chain. Property 5 will remain true trivially by simply using paths that go through the special points of this octagon. Now, to see that we still have a Topological Snake, we must check that none of the three new vertices already lie on different edges of F . If this were the case, we could create an edgepath using Property 5 that contradicts the minimality of E . Thus the new bubble is properly embedded. Similarly, Property 4' must still be satisfied. Otherwise we could use Property 5 to yield a contradiction of E being minimal.

Type 1: we complete an octagon that shares exactly one edge with a previously completed octagon. Start with a half cycle $d_j \dots d_{j+3}$ contained in $f_j \dots f_{j+3}$. By Property 5 we must have that any half cycle must be contained in four consecutive pieces of F to avoid contradicting the local minimality of E . By Property 6 we have that it must span all four. The edge bordering another completed octagon cannot be in f_{j+1} or f_{j+2} or we would have a contradiction of Property 2. Thus $d_a = e_a = f_a$ for $a \in \{j+1, j+2\}$. Thus, since by hypothesis only one edge of the half cycle borders a previously completed octagon, we have either $d_a = e_a = f_a$ for $a = j$ or for $a = j+3$. Assume the former is true (the other case is similar) so that d_{j+3} will border the previously completed octagon and its initial vertex, being the entry point of the bubble chain, must be a special point of this octagon. We now complete our current octagon. Properties 6 and 1 immediately remains satisfied. We have prepended the bubble chain containing the previously completed octagon with another bubble. Let the endpoints of our half cycle being completed be the new octagon's special points. This satisfies Property 2. It also satisfies Property 3 since special points of the neighboring octagons are on opposite sides of d_{j+3} . The other required Properties are checked to be satisfied as in Type 0 additions.

Type 2: we complete an octagon that shares exactly two edges with an previously completed octagon. As above, we start with a half cycle $d_k \dots d_{k+3}$ contained in $f_k \dots f_{k+3}$. Again we must have $d_a = e_a = f_a$ for $a \in \{k+1, k+2\}$. Thus d_k and d_{k+3} must be the edges bordering previously completed octagons. This means that they have special points on the endpoints of d_k and d_{k+3} nearest each other. Now we complete the octagon. We eliminate an arc between two bubble chains, joining them as one bubble chain with the addition of a new bubble. Let the endpoints of our half cycle being completed be the special points of the new octagon. All of the required properties can be checked to be satisfied as in the Type 1 additions.

In each of the above completion Types, the cardinality of $\{i | 1 \leq i \leq l''', e_i = f_i\}$ strictly decreases. Thus this process terminates. Since it cannot continue, F must finally satisfy Property 4.

Steps 1, 2, and 3 did not change the conjugacy class of W so there is a lift of W with the same limit points as E . Since $GS=F$ contains E , it too has the same limit points ■

2.5. Summary of Correspondence. Most of the time, when we create an Infinite Snake from a conjugacy class, it is a Proper Infinite Snake, and therefore unique. We will explore and classify the only instances where this is not the case. This discussion will be analogous to Theorem 3.5 in Lustig's paper [4].

Let W be a cyclic word in the generators of G . Let S be an Infinite Snake gotten from W via the above algorithm. Let $\mathbb{J} = \dots J_{-2} J_{-1} J_0 J_1 J_2 \dots$ be a Markov chain with S as image. Let R_i be the region corresponding to J_i for all i . Suppose that S is not a proper infinite snake. WLOG, suppose that its forward limit point, f is not contained in R_k for some k . Since f is a limit point of S the boundary of the Poincare disc, the edges of S get arbitrarily close to it in the euclidian metric. Now, if f weren't in $\overline{R_i}$ for some i , then it also wouldn't be in the closures of all subsequent regions. This cannot be the case since if f were a nonzero distance from the regions in the euclidian metric, then the edges of S closer to f would grow arbitrarily far away from their regions in the hyperbolic metric, which contradicts the definition of our regions. Thus, $f \in \overline{R_i}$ for all i . The periodicity of S and the nesting of the regions then dictate that f is not in any of the regions, so it must be in their boundaries and they all must be open. Type D regions cannot occur without Type C regions, so this implies that all regions are of Type A, A', or A". Let $\gamma \in N$ be the boundary of R_k containing f as a limit point. Then γ must bound all of the regions. This describes an infinite chain of cycles of length 3. Iterating backwards, one sees that γ 's other limit point is the other limit of S . We will call such an Infinite Snake "exceptional". Using this we can prove the following key theorem.

Theorem 9. *For $n > 6$, there is an n -to-1 correspondence between Markov chains of period n and primitive conjugacy classes of G of wordlength n .*

Proof: With our group presentation, all primitive conjugacy classes that yield exceptional Infinite Snakes via our algorithm have wordlength 6. For primitive conjugacy classes of wordlength greater than 6, the algorithm then yields a Proper Infinite Snake by the argument in the preceding paragraph. Choose a primitive cyclic word W of wordlength $n > 6$. Since there is an exceptional Infinite Snake with limit points corresponding to every geodesic of N , W 's Proper Infinite Snake, S must not have both its limit points on the same geodesic of N . Thus, up to deck transformation, it is the unique Infinite Snake for the conjugacy class of W by Theorem 6. There are then exactly n Markov chains of period n corresponding to the different shifts that read off S . Conversely, given a Markov chain of period n , one can read off the left-most edges of its Infinite Snake to obtain a corresponding cyclic word of length n ■

Let \mathbb{J}_n be the set of all elements of our Markov chain that have period n . Let F_n be the set of conjugacy classes of G of wordlength n . Let p_n be the pointwise map from \mathbb{J}_n to F_n described in the above proof. Let λ_n and μ_n be the uniform measures on \mathbb{J}_n and F_n respectively.

Theorem 10. *The pushforward under p_n of λ_n approaches μ_n in total variation norm (TV) as n approaches infinity.*

Proof: The number of elements of \mathbb{J}_n grows exponentially with n , and a counting argument reveals that the number of these that aren't primitive is on the order of the square root of them. Thus the fraction of non primitive elements of \mathbb{J}_n approaches 0. Since these then have negligible effect on measure λ_n , the previous theorem yields the desired result ■

3. THE SELF-INTERSECTION NUMBER

For $c \in F_n$ let $N(c)$ be the minimal number of self-intersections of a closed curve in the free homotopy class representing c . This number can be computed as Lustig did it [4].

Using the map p_n , N can also be defined on elements of \mathbb{J}_n . As in the paper by Chas and Lalley [2], we can express this value as a sum. Choose $\mathbf{X} \in \mathbb{J}_n$. Let $\sigma^i \mathbf{X}$ denote the i th cyclic permutation of \mathbf{X} . Construct via Lustig's algorithm [4] $H : (\mathbb{J}_n \times \mathbb{J}_n) \rightarrow \{0, 1\}$ that gives 1 if the inputs determine the first instance of a linked pair and 0 otherwise. Then we can write

$$N(\mathbf{X}) = \sum_{1 \leq i < j \leq n} H(\sigma^i \mathbf{X}, \sigma^j \mathbf{X}).$$

As in Chas and Lalley, H can be written as $H(\mathbf{X}, \mathbf{Y}) = \sum_{k=1}^n h_k(\mathbf{X}, \mathbf{Y})$ where h_k satisfies H0, H1, H2, and H3 of Chas and Lalley [2], described below.

3.1. The Spine of a Snake. Before we describe in detail how to calculate h_k , we introduce the *spine* of a snake which is equivalent to the "Normal Edgpaths" in

used in Lustig's paper [4]. The *spine* of a snake (or Infinite Snake, or Proper Infinite Snake) is a topological arc that runs along the center of the snake. In arc-like regions of a snake, the spine is coincident with the arc. When the spine reaches the entry point of a bubble chain, it travels along the geodesic from the entry point to the center of the octagon of the first bubble in the bubble chain. If it is the only bubble in the bubble chain, the spine continues along the geodesic from the center of the octagon to the exit point. When there are subsequent bubbles in the same chain, the spine travels along the geodesic from the center of the first octagon, and then to the next, until it reaches the center of the last bubble in the bubble chain. Then it exits along the geodesic from the center of the last octagon to the exit point as in the case of a bubble chain with a single bubble.

We will construct a step-by-step map from a snake to its spine. We will lose information about the Markov structure of the snake and the regions for each step, but we will keep exactly enough information to calculate self-intersection number via Lustig's algorithm [4]. The spine, as described above, is an embedded path in the graph Δ which we describe now.

Definition 5 (Δ). *The graph Δ , embedded in the hyperbolic plane, is the union of the graphs Γ and N , together with geodesic edges between each vertex of Γ and its 8 closest vertices of N .*

Now we construct the map ϕ from the alphabet \mathbb{A} to a smaller alphabet Φ that translates graphically into a map from a snake in Γ to its spine in Δ : let x, y , and z be arbitrary elements in Ω , the set of generators of G along with their inverses. Then we set

$$\begin{aligned}\phi(A_x) &= \phi(A_x^L) = \phi(A_x^R) = \phi(A_x^{L'}) = \phi(A_x^{R'}) = \alpha_x \\ \phi(B_{xy}) &= \beta_{xy}, \\ \phi(C_{xy}) &= \phi(C_{xy}') = \gamma, \\ \phi(D_{xy}) &= \delta_{xy}, \\ \phi(E_{xyz}^L) &= \epsilon_y, \text{ and} \\ \phi(E_{xyz}^R) &= \epsilon_{\bar{y}}.\end{aligned}$$

Here we will describe how ϕ translates graphically into a map from an Infinite Snake in Γ to its spine in Δ . Suppose $A_x(x \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the edge $x_0 \in \Gamma$ with label x . Then the spine of this step is the same edge in Δ and we label this edge (and any edge equivalent under deck transformation) α_x , consistent with ϕ . The same holds for steps of the form A_x^L , A_x^R , $A_x^{L'}$, and $A_x^{R'}$.

Similarly, suppose $B_{xy}(x, y \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the pair of edges x_0 and $y_0 \in \Gamma$ with labels x and y . Then the spine of this step is the edge in Δ from the vertex at the start of x_0 and y_0 to the vertex of N at the center of the octagon of Γ containing them. We label this edge (and any edge equivalent under deck transformation) β_{xy} , consistent with ϕ .

If $C_{xy}(x, y \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the pair of edges x_0 and $y_0 \in \Gamma$ with labels x and y , then the spine of this step is just the vertex in N at the center of the octagon of Γ containing the two edges. We label this vertex (and any vertex equivalent under deck transformation) γ , consistent with ϕ . We do the same for any step of the form C'_{xy} .

Instead, if $D_{xy}(x, y \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the pair of edges x_0 and $y_0 \in \Gamma$ with labels x and y , then the spine of this step is the edge in Δ from the vertex of N at the center of the octagon of Γ containing x_0 and y_0 to the vertex of Γ where they meet. We label this edge (and any edge equivalent under deck transformation) δ_{xy} , consistent with ϕ .

Now suppose $E_{xyz}^L(x, y, z \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the set of edges $x_0, y_0,$ and $z_0 \in \Gamma$ with labels $x, y,$ and z . Then the spine for this step is the edge in $N \subset \Delta$ from the vertex of N at the center of the octagon of Γ containing y_0 and z_0 to the vertex of N at the center of the octagon of Γ containing x_0 and y_0 . We label this edge (and any edge equivalent under deck transformation) ϵ_y , consistent with ϕ .

Finally suppose $E'_{xyz}^L(x, y, z \in \Omega)$ is a step in a Markov chain and its Infinite Snake image is the set of edges $x_0, y_0,$ and $z_0 \in \Gamma$ with labels $x, y,$ and z . Then the spine for this step is the edge in $N \subset \Delta$ from the vertex of N at the center of the octagon of Γ containing x_0 and y_0 to the vertex of N at the center of the octagon of Γ containing y_0 and z_0 . We label this edge (and any edge equivalent under deck transformation) $\bar{\epsilon}_y$, consistent with ϕ .

We can define inverses of the elements of Φ by pushing forward via ϕ the inverses described in Theorem 2. This will give a well-defined map from Φ to itself, even though the inverse on single elements of \mathbb{A} wasn't always well-defined. In particular, we have

$$\overline{\alpha_x} = \alpha_{\bar{x}}, \overline{\beta_{xy}} = \delta_{\bar{y}\bar{x}}, \overline{\gamma} = \gamma, \overline{\delta_{xy}} = \beta_{\bar{y}\bar{x}}, \text{ and } \overline{\epsilon_y} = \bar{\epsilon}_y.$$

Notice that if an edge in Δ is labeled $\omega \in \Phi$, then the same edge, when traversed in the opposite direction, will be labeled $\bar{\omega}$.

3.2. Calculating Linked Pairs. Now we can discuss the cyclic ordering of edges around the vertices of Δ . This will be used to identify "linked pairs" of Cohen and Lustig [3] [4], corresponding to self-intersections. Here and elsewhere we assume that Γ is embedded the hyperbolic plane with the orientation that makes a counterclockwise traversal of the edges in an octagon read off the relation $abABcdCD$ of G in order (up to cyclic permutation). Then, the edges of Δ emanating from a vertex $v \in \Gamma \subset \Delta$ are, when read in counterclockwise order (up to cyclic permutation),

$$O_1 = \alpha_a \beta_{da} \alpha_d \beta_{Cd} \alpha_C \beta_{DC} \alpha_D \beta_{cD} \alpha_c \beta_{bc} \alpha_b \beta_{Ab} \alpha_A \beta_{BA} \alpha_B \beta_{aB}.$$

Also, the edges of Δ emanating from a vertex $\gamma \in N \subset \Delta$ are, when read in counterclockwise order (up to cyclic permutation)

$$O_2 = \delta_{AD} \epsilon_A \delta_{Ba} \epsilon_B \delta_{ab} \epsilon_a \delta_{bA} \epsilon_b \delta_{CB} \epsilon_C \delta_{Dc} \epsilon_D \delta_{cd} \epsilon_c \delta_{Dc} \epsilon_d.$$

With these orderings, we can create functions that assign an orientation to any cyclic word in the elements of Φ . Let ξ be a cyclic word in Φ . For $i \in \{1, 2\}$, set $o_i(\xi)$ equal to 1 if the letters of ξ are all in O_i and occur in cyclic order in O_i , -1 if they are all in O_i and occur in reverse cyclic order in O_i , and 0 otherwise.

With these orientations, we can construct functions u_k and v_k to detect intersections on pieces of spines of length k . Consider two (finite or infinite) strings $\zeta = \zeta_1\zeta_2\dots$ and $\xi = \xi_1\xi_2\dots$ of letters in Φ . We will now define u_k on such a pair of strings for each $k \geq 1$. First, unless the following conditions hold, we set $u_k(\zeta, \xi) = 0$:

- (a) each argument of u_k has length at least k , and
- (b) $\zeta_1 \neq \xi_1, \zeta_k \neq \xi_k$, and $\zeta_i = \xi_i$ for all $1 < i < k$.

Given that (a) and (b) hold for a particular ζ, ξ , and k , we define

$$u_k(\zeta, \xi) = \begin{cases} 1 & \text{if } k = 1, \text{ and } \zeta_1 \text{ and } \xi_1 \text{ are different elements of the set } \{\alpha_x, \epsilon_x\} \\ & \text{for the same } x \in \Omega; \\ 1 & \text{if } k = 2, \text{ and } o_1(\overline{\zeta_1} \overline{\xi_1} \zeta_2 \xi_2) = 1; \\ 1 & \text{if } k = 3, \text{ and } o_1(\overline{\zeta_1} \overline{\xi_1} \zeta_2) = o_1(\overline{\zeta_2} \zeta_3 \xi_3) \neq 0; \\ 1 & \text{if } k = 4, \text{ and either } o_1(\overline{\zeta_1} \overline{\xi_1} \zeta_2) = o_1(\overline{\zeta_3} \zeta_4 \xi_4) \neq 0, \text{ or} \\ & o_2(\overline{\zeta_1} \overline{\xi_1} \zeta_4 \xi_4) = 1; \\ 1 & \text{if } k \geq 5, \text{ and } o_1(\overline{\zeta_1} \overline{\xi_1} \zeta_2) + o_2(\overline{\zeta_1} \overline{\xi_1} \zeta_4) = \\ & o_1(\overline{\zeta_{k-1}} \zeta_k \xi_k) + o_2(\overline{\zeta_{k-3}} \zeta_k \xi_k) \neq 0; \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Now, for each k , we define $v_k(\zeta, \xi) = 0$ unless each of its arguments has length at least k , in which case we set

$$v_k(\zeta, \xi) = u_k(\zeta_1 \zeta_2 \dots \zeta_k, \overline{\xi_k} \overline{\xi_{k-1}} \dots \overline{\xi_1}).$$

u_k and v_k are defined for arbitrary finite or infinite strings of letters in Φ . We extend these definitions to *doubly* infinite sequences $\zeta = \dots\zeta_{-1}\zeta_0\zeta_1\dots$, and $\xi = \dots\xi_{-1}\xi_0\xi_1\dots$ by adopting the convention that

$$u_k(\zeta, \xi) = u_k(\zeta_1 \zeta_2 \dots \zeta_k, \xi_1 \xi_2 \dots \xi_k), \text{ and} \\ v_k(\zeta, \xi) = v_k(\zeta_1 \zeta_2 \dots \zeta_k, \xi_1 \xi_2 \dots \xi_k).$$

Finally, we can define h_k . If \mathbf{X} and \mathbf{Y} are (finite, infinite, or doubly infinite) strings of letters in \mathbb{A} satisfying the Markov transition rules, we set $h_k(\mathbf{X}, \mathbf{Y}) = u_k(\phi(\mathbf{X}), \phi(\mathbf{Y})) + v_k(\phi(\mathbf{X}), \phi(\mathbf{Y}))$. Here when we write ϕ acting on a string, it is to be interpreted as acting on all of the letters in the string.

3.3. Properties of h_k on λ_n . We will need to show that h_k satisfies 4 important properties of Chas and Lalley chas-lalleyin order to make use of some of their results. They are:

- (H0) Each function h_k is symmetric.
- (H1) There exists $C < \infty$ such that $|h_k| \leq C$ for all $k \geq 1$.

(H2) For each $k \geq 1$ the function h_k only depends on the first k entries of its arguments.

(H3) There exist constants $C \leq \infty$ and $0 < \theta < 1$ such that for all $n \geq k \geq 1$ and $0 \leq i < j < n$,

$$E_{\lambda_n} |h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})| \leq C\theta^k.$$

The first three of these are clearly satisfied for our h_k defined in Section 3.2. Before we prove (H3), we will need some preliminary results. We start with some properties of the transition matrix of our Markov chain. Let $\mathbf{A} = \{a_{ij} | 1 \leq i, j \leq 88\}$ be the transition matrix of our topological Markov chain defined in Section 2.2 on the alphabet \mathbb{A} of 88 Markov states. We have the following simple theorem.

Theorem 11. *The matrix \mathbf{A} of 0's and 1's is irreducible and aperiodic. Further, \mathbf{A}^n is strictly positive for $n \geq 11$.*

Proof: We will show that we can get from any state to A_a and from A_a to any state. This will imply irreducibility. From any state $J_1 \in \mathbb{A}$ such that $\Pi(J_1) \in \{B, C, C', D, E\}$ we can reach a state $J_2 \in \mathbb{A}$ such that $\Pi(J_2) = D$ in 3 steps. From any state $J_3 \in \mathbb{A}$ such that $\Pi(J_3) \in \{A, A', A'', D\}$ we can reach either A_a or A_c in one step and from A_c we can reach A_a in one step. From A_a we can get to A_c in one step and from these we can reach any J_4 such that $\Pi(J_4) = A$. From these we can reach any J_5 such that $\Pi(J_5) \in \{A', B\}$. From these we can reach any J_6 such that $\Pi(J_6) \in \{A'', C, C', D, E\}$ in at most 3 steps. This proves irreducibility. The fact that we can get from A_a to itself in one step then implies aperiodicity. The specific statement for positivity follows from the fact that we can go from any state to A_a in five or less steps, stay there for as many steps as we wish, and then go to any other state in 6 or less additional steps ■

Next we discuss a refinement of a general result from Section 8.5 of Walters [6] about the growth rate of periodic points of a Markov Chain. Let $\tau : \mathbf{Y}_M \rightarrow \mathbf{Y}_M$ be a two sided topological Markov chain with irreducible, aperiodic s by s transition matrix \mathbf{M} . Let P_n denote the cardinality of the set $\mathbb{J}_n = \{\mathbf{y} \in \mathbf{Y}_M | \tau^n(\mathbf{y}) = \mathbf{y}\}$. Let $\lambda_i (1 \leq i \leq s)$ be the eigenvalues of \mathbf{M} and let Λ be the eigenvalue that is positive and uniquely has the largest absolute value. We have

Theorem 12. *$\log(P_n)/n$ approaches $\log \Lambda$ as n approaches infinity. This convergence is exponential in the following sense: there exists $C < \infty$ and $0 < \beta < 1$ such that $|P_n/\Lambda^n - 1| \leq C\beta^n$ for all n .*

Proof: A simple counting argument shows that P_n equals the trace of \mathbf{M}^n or $\sum_{i=1}^s \lambda_i^n$. Let $\Lambda' < \Lambda$ be a number strictly greater in magnitude than all other eigenvalues of \mathbf{M} (it exists by Theorem 1.1 of Seneta [5]). Combining this gives

$$\left| \frac{\sum_{i=1}^s \lambda_i^n}{\Lambda^n} - 1 \right| \leq (s-1)\Lambda'/\Lambda^n.$$

Setting $C = s - 1$ and $\beta = \Lambda'/\Lambda$ gives the desired exponential convergence and the result follows ■

We will also use the following fact about how well a spine determines a snake.

Lemma 5. *A string of letters in Φ can have at most 16 preimages under ϕ that are portions of Infinite Snakes.*

Proof: Given the string $\zeta_1\zeta_2\dots\zeta_n$, there are at most 16 preimages of ζ_1 (this maximum occurs when $\zeta_1 = \gamma$. Suppose that J_1, J_2 , and J_3 are in \mathbb{A} such that J_2 and J_3 can each immediately follow (according to the transition matrix \mathbf{A}) J_1 , but $J_2 \neq J_3$. Then it can be checked that $\phi(J_2) \neq \phi(J_3)$. It inductively follows that a choice of preimage for ζ_1 uniquely determines the rest of the preimage of the spine. Thus there are at most 16 such preimages ■

The next lemma describes how a spine of a snake is "straight" and cannot turn back on itself or be it's own inverse.

Lemma 6. *Suppose $\zeta_1\zeta_2\dots\zeta_i$ be a portion of a spine such that $\zeta_j = \overline{\zeta_{i+1-j}}$ for all $j \leq i$. Then $i \leq 2$.*

Proof: If $i = 3$, then $\zeta_2 = \overline{\zeta_2}$, so ζ_2 must be γ . Exactly one of ζ_1 or ζ_3 must then also be γ , but this is impossible since they must be inverses of each other (this is checked by considering all possibilities for their preimages under ϕ). If $i = 4$, then ζ_2 and ζ_3 must be inverses of each other. The only way for this to be a portion of a spine is for them both to be γ . If $x_1x_2x_3x_4$ is part of a snake such that $\phi(x_i) = \zeta_i$ where $\zeta_2 = \zeta_3 = \gamma$, then ζ_1 and ζ_4 cannot be inverses – a contradiction. If $i > 5$ we would have one of the previous two cases (depending on whether i were even or odd), so this too is impossible ■

Lastly, we will need a crude bound on the growth rate of strings of symbols in \mathbb{A} satisfying the transition matrix \mathbf{A} . Let Q_n be the number of such strings of length n .

Lemma 7. $Q_n \leq 88^{11}P_n$.

Proof: By Theorem 11, any string of length $n - 11$ can be made by removing the last 11 letters from some element of \mathbb{J}_n . Thus there are at most P_n strings of length $n - 11$. The crude bound in the lemma comes from considering adding any 11 elements of \mathbb{A} to such a string ■

We are now in position to prove that (H3) applies in our case.

Theorem 13. *There exist constants $C \leq \infty$ and $0 < \theta < 1$ such that for all $n \geq k \geq 1$ and $0 \leq i < j < n$,*

$$E_{\lambda_n} |h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})| \leq C\theta^k.$$

Proof: Since the measure λ_n is invariant under cyclic shifts, and h_k is symmetric in its inputs, it suffices to prove this in the case where $i = 0$ and $j \leq n/2$. By property (H1), we simply need to show that the fraction of elements \mathbf{x} of \mathbb{J}_n such

that $h_k(\mathbf{x}, \tau^j \mathbf{x}) \neq 0$ falls off exponentially in k . Let P_n be the cardinality of \mathbb{J}_n . Let Λ be the maximal eigenvalue of \mathbf{A} . By Theorems 12 and 11, there exists $C_0 < \infty$ and $0 < \beta < 1$ such that for all n ,

$$(1) \quad |P_n/\Lambda^n - 1| \leq C_0\beta^n.$$

Choose n_0 sufficiently large such that $C_0\beta^{n_0} < 1$. For all $n > n_0$, rearranging (1) gives

$$(2) \quad 0 < \Lambda^n(1 - C_0\beta^{n_0}) < \Lambda^n(1 - C_0\beta^n) < P_n.$$

We need to compute an upper bound on the number of elements \mathbf{x} of \mathbb{J}_n such that

$$(3) \quad h_k(\mathbf{x}, \tau^j \mathbf{x}) \neq 0.$$

Suppose $\mathbf{x} = \dots x_1 x_2 \dots x_n \dots$ satisfies (3). Then it is necessary that either

$$(4) \quad \phi(x_2 \dots x_{\lfloor (k-1)/2 \rfloor}) = \phi(x_{j+2} \dots x_{j+\lfloor (k-1)/2 \rfloor}),$$

or

$$(5) \quad \phi(x_2 \dots x_{\lfloor (k-1)/2 \rfloor}) = \overline{\phi(x_{j+k-1})} \overline{\phi(x_{j+k-2})} \dots \overline{\phi(x_{j+k+1-\lfloor (k-1)/2 \rfloor})}.$$

Assume \mathbf{x} satisfies (4). Then $\phi(x_2 \dots x_{\lfloor (k-1)/2 \rfloor})$ is determined uniquely by the $n - \lfloor (k-1)/2 \rfloor$ symbols $\phi(x_{\lfloor (k-1)/2 \rfloor + 1}) \dots \phi(x_n)$ which are determined by the string $x_{\lfloor (k-1)/2 \rfloor + 1} \dots x_n$ of length $n - \lfloor (k-1)/2 \rfloor < n - k/2 + 1$. By Lemma 7 and (1), there are at most $88^{11} \Lambda^{n-k/2+1} (1 + C_0\beta^{n-k/2+1}) < C_1 \Lambda^{n-k/2+1}$ such strings, where $C_1 = 88^{11}(C_0 + 1)$. Each of these yield at most 16 values for $x_2 \dots x_{\lfloor (k-1)/2 \rfloor}$ according to Lemma 5. If we include the 88 values x_1 can have, then there are maximum of $16(88)C_1 \Lambda^{n-k/2+1}$ values of $\mathbf{x} \in \mathbb{J}_n$ satisfying (4).

Now assume that \mathbf{x} satisfies (5). By Lemma 6, $j+k \leq n+4$. Then $\phi(x_5 \dots x_{\lfloor (k-1)/2 \rfloor})$ is determined uniquely by the $n - \lfloor (k-1)/2 \rfloor$ symbols $\phi(x_{\lfloor (k-1)/2 \rfloor + 1}) \dots \phi(x_n)$ which are determined by the string $x_{\lfloor (k-1)/2 \rfloor + 1} \dots x_n$ of length $n - \lfloor (k-1)/2 \rfloor < n - k/2 + 1$. As in the previous paragraph, there are at most $C_1 \Lambda^{n-k/2+1}$ such strings. Each yields at most 16 values for $x_5 \dots x_{\lfloor (k-1)/2 \rfloor}$ according to Lemma 5. If we include the 88 values each of x_1, x_2, x_3 , and x_4 can take, then there are a maximum of $16(88^4)C_1 \Lambda^{n-k/2+1}$ values of $\mathbf{x} \in \mathbb{J}_n$ satisfying (5).

Combining the results of the previous two paragraphs, there are $16(88^4 + 88)C_1 \Lambda^{n-k/2+1}$ values of $\mathbf{x} \in \mathbb{J}_n$ satisfying (3). Setting $C_2 = 16(88^4 + 88)C_1 \Lambda / (1 - C_0\beta^{n_0})$ and using (2) we have the upper bound

$$(6) \quad E_{\lambda_n} |h_k(\mathbf{x}, \sigma^j \mathbf{x})| \leq \frac{C_2 \Lambda^{n-k/2}}{\Lambda^n} = C_2 \Lambda^{-k/2}$$

when $n \geq n_0$. When $n < n_0$ we still have $E_{\lambda_n} |h_k(\mathbf{x}, \sigma^j \mathbf{x})| \leq 1$. Thus, setting $C = \max(C_2, \Lambda^{n_0 k/2})$ and $\theta = \Lambda^{-1/2}$ gives the desired result \blacksquare

4. STATISTICS

In this section, we state and prove the main theorem.

Theorem 14 (Main Theorem). *Let N_n denote the random variable obtained by evaluating N on a uniformly chosen $\alpha \in F_n$. Then there exists κ and $\sigma^2 \geq 0$ such that, as $n \rightarrow \infty$,*

$$\frac{N_n - \kappa n^2}{n^{3/2}} \implies \text{Normal}(0, \sigma)$$

in the weak topology.

Theorem 10 implies that the distribution of

$$f_n = \frac{\sum_{i=1}^n (\sum_{j=i+1}^n (\sum_{k=\infty}^{\infty} (h_k(\sigma^i \mathbf{X}, \sigma^j \mathbf{X})))) - \kappa n^2}{n^{3/2}}$$

on μ_n and λ_n approach each other. Thus, to prove the Main Theorem, we only need to prove that the distribution of f_n on λ_n approaches a Gaussian. To this end, we will prove (Theorem 17) that it is weakly equivalent to a normalized U-statistic, which according to Theorem 5.1 of Chas and Lalley [2] is Gaussian. This will prove the Main Theorem.

4.1. The Parry Measure. Before we define the U-statistic, we need to define the Parry measure on our topological Markov chain. By Theorem 11 and Theorem 1.1 of Seneta [5], \mathbf{A} has a positive eigenvalue (of uniquely maximal absolute value) λ with positive right and left eigenvectors. Let $u = (u_1, \dots, u_{88})$ be a strictly positive left eigenvector and $v = (v_1, \dots, v_{88})$ be a strictly positive right eigenvector of \mathbf{A} such that $\sum_{i=1}^{88} u_i v_i = 1$. Then let $\pi_i = u_i v_i$ and $q_{ij} = a_{ij} v_j / (\lambda v_i)$. Then, as in Walters [6], the stochastic matrix $Q = \{q_{ij}\}$ and the stationary probability vector $\pi = \{\pi_i\}$ give our Markov chain the Parry measure ν .

The Parry measure has some important properties. It is the unique measure with maximal entropy on our Markov chain $\mathbf{X}_{\mathbf{A}}$ according to Theorem 8.10 of Walters [6]. This implies according to Theorem 8.17 of Walters [6] that it "describes the distribution of the periodic points" of τ . This means that for a measurable set $S \subset \mathbf{X}_{\mathbf{A}}$

$$(7) \quad \lim_{n \rightarrow \infty} \frac{\#\{\mathbf{x} \in \mathbb{J}_n \mid \mathbf{x} \in S\}}{P_n} = \lim_{n \rightarrow \infty} (\lambda_n(\mathbb{J}_n \cap S)) = \nu(S).$$

We will use this result to produce an analogue of Theorem 13 for the Parry measure ν .

Theorem 15. *There exist constants $C \leq \infty$ and $0 < \theta < 1$ such that for all $k \geq 1$ and $0 \leq i < j$,*

$$E_{\nu} |h_k(\tau^i \mathbf{x}, \tau^j \mathbf{x})| \leq C \theta^k.$$

Proof: Fix i, j , and k . Let C and θ be as in Theorem 13. Let $S = \{\mathbf{x} \in \mathbf{X} | h_k(\tau^i \mathbf{x}, \tau^j \mathbf{x}) = 1\}$. Note that S is measurable since it is a finitely determined cylinder set. Since the range of h_k is $\{0, 1\}$,

$$(8) \quad E_\nu |h_k(\tau^i \mathbf{x}, \tau^j \mathbf{x})| = \nu(S), \text{ and}$$

$$(9) \quad E_{\lambda_n} |h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})| = \lambda_n \{\mathbf{x} \in \mathbb{J}_n | h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x}) = 1\}$$

For $\mathbf{x} \in \mathbb{J}_n \subset \mathbf{X}_A$, $\sigma \mathbf{x} = \tau \mathbf{x}$, so (9) becomes

$$(10) \quad E_{\lambda_n} |h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})| = \lambda_n (\mathbb{J}_n \cap S).$$

By (10) and Theorem 13, we have bound (independent of n)

$$(11) \quad \lambda_n (\mathbb{J}_n \cap S) \leq C\theta^k.$$

Finally, taking the limit on both sides of (11), and applying (7) explicitly gives $\nu(S) \leq C\theta^k$, so (8) gives the desired result \blacksquare

We will also need a theorem relating a substring measure to the Parry measure. Fix $0 < \delta < 1/4$ and set $m = m(n) = \lfloor n^\delta \rfloor$. Let $\lambda_{n,m}$ denote the distribution of the random substring $x_1 x_2 \dots x_{n-m}$ of $x_1 x_2 \dots x_n \in \mathbb{J}_n$. Also let $\nu_{n,m}$ denote the distribution of the random substring $x_1 x_2 \dots x_{n-m}$ of $\dots x_1 x_2 \dots x_{n-m} \dots \in \mathbf{X}_A$. Then

Theorem 16. *For n sufficiently large, $\lambda_{n,m}$ is absolutely continuous with respect to $\nu_{n,m}$, and the Radon-Nikodym derivative is sandwiched between bounds that tend to 1 so that the two measures agree in total variation norm.*

Proof: When $n \geq 11^{1/\delta}$, Theorem 11 guarantees absolute continuity because for any string $\dots x_1 x_2 \dots x_{n-m} \dots \in \mathbf{X}_A$, there exists at least one element $x_1 x_2 \dots x_n$ in \mathbb{J}_n since m will be at least 11. Now, at the substring $x_1 x_2 \dots x_{n-m}$,

$$\begin{aligned} \frac{d\lambda_{n,m}}{d\nu_{n,m}} &= \frac{\sum_{x_{n-m+1}, x_{n-m+2}, \dots, x_n \in \mathbb{A}} a_{x_1, x_2} a_{x_2, x_3} \dots a_{x_{n-m-1}, x_{n-m}} a_{x_{n-m}, x_{n-m+1}} \dots a_{x_n, x_1} / P_n}{u_{x_1} v_{x_1} \prod_{i=1}^{n-m-1} (a_{x_i, x_{i+1}} v_{x_{i+1}} / (\Lambda u_{x_i}))} \\ &= \frac{a_{x_1, x_2} a_{x_2, x_3} \dots a_{x_{n-m-1}, x_{n-m}} (\mathbf{A}^{m+1})_{x_{n-m}, x_1} / P_n}{a_{x_1, x_2} a_{x_2, x_3} \dots a_{x_{n-m-1}, x_{n-m}} u_{x_1} v_{x_{n-m}} / \Lambda^{n-m-1}} = \frac{(\mathbf{A}^{m+1})_{x_{n-m}, x_1} / P_n}{u_{x_1} v_{x_{n-m}} / \Lambda^{n-m-1}} \end{aligned}$$

According to Theorem 1.2 of Seneta [5], $(\mathbf{A}^r)_{x_i, x_j} \rightarrow \Lambda^r u_{x_j} v_{x_i}$ as n goes to infinity. We can use this and the bounds from Theorem 12 to force the Radon-Nikodym derivative bounds toward 1 as follows.

$$\frac{d\lambda_{n,m}}{d\nu_{n,m}} = \frac{(\mathbf{A}^{m+1})_{x_{n-m}, x_1} / P_n}{u_{x_1} v_{x_{n-m}} / \Lambda^{n-m-1}} \rightarrow \frac{\Lambda^{m+1} u_{x_1} v_{x_{n-m}} / \Lambda^n}{u_{x_1} v_{x_{n-m}} / \Lambda^{n-m-1}} = 1$$

as $n \rightarrow \infty$ independently of the choice of x_i 's \blacksquare

4.2. The U-Statistic. We will describe an example of a second order U-statistic as in Chas and Lalley [2]. When normalized, it will approach a Gaussian by Theorem 5.1 of their paper and we will show it also weakly approaches f_n . Call our Markov chain space $X_{\mathbf{A}}$. Assign it the Parry measure as above and let τ be the shift map on the sequence space $\mathbb{A}^{\mathbb{N}}$. For \mathbf{X} chosen from this stationary, aperiodic, irreducible Markov chain with measure ν , define the following second order U-statistic

$$N'_n(\mathbf{X}) = \sum_{i=1}^n \left(\sum_{j=i+1}^n \left(\sum_{k=1}^{\infty} (h_k(\tau^i \mathbf{X}, \tau^j \mathbf{X})) \right) \right).$$

Let f_n be the distribution of $(N(c) - n^2\kappa)/n^{3/2}$ on μ_n (the uniform measure on the set of conjugacy classes F_n of length n). Let g_n be the distribution of $(N'_n(\mathbf{x}) - \kappa n^2)/n^{3/2}$ on the Parry measure ν . This is our "normalized" version of the U-statistic. Before we relate this U-statistic to the intersection number, we will need an elementary result about distributions.

Lemma 8. *Suppose X_n and Y_n are two infinite sequences of random variables that are all defined on a common probability space. Let c_n be an infinite sequence of scalars. Choose $r > 0$. Let F_n and G_n denote the distributions of X_n and Y_n , respectively. Then,*

$$\|F_n - G_n\|_{TV} \rightarrow 0$$

implies that

$$(X_n - Y_n)/n^r \implies 0,$$

which itself implies that

$$\frac{Y_n - a_n}{n^r} \implies F \text{ if and only if } \frac{X_n - a_n}{n^r} \implies F.$$

The following theorem and proof closely follow Proposition 4.7 of Chas and Lalley [2].

Theorem 17. *For any metric ϱ which induces the topology of weak convergence on probability measures,*

$$(12) \quad \lim_{n \rightarrow \infty} \varrho(f_n, g_n) = 0,$$

so $f_n \implies Normal(0, \sigma)$ if and only if $g_n \implies Normal(0, \sigma)$.

Proof: Let f'_n be the distribution of $(N(\mathbf{x}) - n^2\kappa)/n^{3/2}$ under the uniform measure λ_n on \mathbb{J}_n . By Theorem 10, Lemma 8, and the fact that total variation norm is never increased by a mapping, f_n and f'_n approach each other in total variation distance. Thus, it suffices to prove (12) with f'_n replacing f_n . We will do this by truncating the formulas for the intersection number and U-statistic,

$$(13) \quad N(\mathbf{x}) = \sum_{i=1}^n \left(\sum_{j=i+1}^n \left(\sum_{k=1}^n (h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})) \right) \right) \text{ and}$$

$$(14) \quad N'_n(\mathbf{X}) = \sum_{i=1}^n \left(\sum_{j=i+1}^n \left(\sum_{k=1}^{\infty} (h_k(\tau^i \mathbf{X}, \tau^j \mathbf{X})) \right) \right),$$

twice each and then comparing the results using machinery similar to that in Walter's proof of (7).

Fix $0 < \delta < 1/4$ and set $m = m(n) = \lfloor n^\delta \rfloor$. By Theorems 13 and 15,

$$E_{\lambda_n} \left(\sum_{i=1}^n \left(\sum_{j=i+1}^n \left(\sum_{k=m(n)}^n |h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})| \right) \right) \right) \leq C n^2 \theta^{m(n)}, \text{ and}$$

$$E_{\nu} \left(\sum_{i=1}^n \left(\sum_{j=i+1}^n \left(\sum_{k=m(n)}^{\infty} |h_k(\tau^i \mathbf{x}, \tau^j \mathbf{x})| \right) \right) \right) \leq C n^2 \theta^{m(n)}.$$

These upper bounds decrease to 0 as $n \rightarrow \infty$. Thus, by Markov's inequality, both of these sums converge weakly to 0. According to Lemma 8, this implies that

$$\lim_{n \rightarrow \infty} \varrho(f'_n, f_n^A) = 0, \text{ and } \lim_{n \rightarrow \infty} \varrho(g_n, g_n^A) = 0.$$

Where f_n^A and g_n^A are gotten from f'_n and g_n respectively by replacing the bounds on the inner sums of (13) and (14) by $1 \leq k < m(n)$. Therefore, it suffices to prove

$$\lim_{n \rightarrow \infty} \varrho(f_n^A, g_n^A) = 0.$$

Notice that $|\{i, j, k | n - 2m(n) < j < n, 1 \leq i < j, 1 \leq k < m(n)\}| \leq 2(n)(m(n))^2 = 2n^{1+2\delta}$. Since this grows much slower than $n^{3/2}$, we can eliminate this index set from the summations describing f_n^A and g_n^A with negligible effect (each term is at most 1). Specifically, let $f_n^{A,B}$ and g_n^B be the distributions of the second truncations of f'_n and g'_n gotten by replacing the limits of summation in (13) and (14) by $1 \leq i < j < n - 2m(n)$ and $1 \leq k < m(n)$. Then by Lemma 8 it suffices to prove

$$\lim_{n \rightarrow \infty} \varrho(f_n^{A,B}, g_n^B) = 0.$$

Here, when $i < j < n - 2m(n)$ and $k < m(n)$, $h_k(\sigma^i \mathbf{x}, \sigma^j \mathbf{x})$ and $h_k(\tau^i \mathbf{X}, \tau^j \mathbf{X})$ only depend on the first $n - m(n)$ entries of \mathbf{x} . Thus $f_n^{A,B}$ and g_n^B are the same function on the two different measures $\lambda_{n,m}$ and $\nu_{n,m}$. Finally, by Theorem 16 and since total variation distance is never increased under a mapping, we have the desired result \blacksquare

REFERENCES

- [1] Joan S. Birman and Caroline Series. Dehn's algorithm revisited. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 451–478. Princeton Univ. Press, Princeton, NJ, 1987.
- [2] Chas, Moira, and Steven Lalley. "Self-intersections in combinatorial topology: statistical structure." *Inventiones Mathematicae*. 188. (2012): 429-463. Print.
- [3] Marshall Cohen and Martin Lustig. Paths of geodesics and geometric intersection numbers. I. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 479–500. Princeton Univ. Press, Princeton, NJ, 1987.

- [4] Martin Lustig. Paths of geodesics and geometric intersection numbers. II. In *Combinatorial group theory and topology (Alta, Utah, 1984)*, volume 111 of *Ann. of Math. Stud.*, pages 501–543. Princeton Univ. Press, Princeton, NJ, 1987.
- [5] E. Seneta. *Non-negative Matrices and Markov Chains*. New York: Springer-Verlag, 1981.
- [6] P. Walters. *Introduction to Ergodic Theory*. New York: Springer-Verlag, 1982.