FREQUENCY-LOCKING ON THE BOUNDARY OF THE BARYCENTRE SET

OLIVER JENKINSON

Abstract. We consider the doubling map $T : z \mapsto z^2$ of the circle. For each T-invariant probability measure μ we define its barycentre $b(\mu) = \int_{S^1} z \, d\mu(z)$, which describes its average weight around the circle. We study the set Ω of all such barycentres, a compact convex set with non-empty interior. Its boundary $\partial\Omega$ has a countable dense set of points of non-differentiability, the worst possible regularity for the boundary of a convex set. We explain this behaviour in terms of the frequency-locking of rotation numbers for a certain class of invariant measures, each supported on the closure of a Sturmian orbit.

1. Introduction

A recurring theme in the study of chaotic dynamics is the occurrence of non-smooth phenomena (fractal attractors, irregular conjugacies, etc), even when the system itself is smooth. Such non-smooth behaviour has been observed even in the simplest of dynamical systems, one-dimensional discrete maps.

In this article we describe some experimental work leading to a new example of non-smooth behaviour, associated to arguably the simplest model of chaotic dynamics, the doubling map $z \mapsto z^2$ of the unit circle in the complex plane. This map has a complicated orbit structure, an abundance of invariant sets, and hence a highly non-trivial set $\mathcal{M}$ of invariant probability measures (a measure μ is T-invariant if $\mu(T^{-1}A) = \mu(A)$ for all Borel subsets A). Indeed the behaviour described in this article is further testimony to the complicated nature of $\mathcal{M}$.

Our problem is geometrical. We consider the most natural two-dimensional projection of $\mathcal{M}$, given by taking the barycentre (or, equivalently, the first moment) $b(\mu) = \int_{S^1} z \, d\mu(z)$ of each measure μ. Since $\mathcal{M}$ is weak* compact and convex (see [Walters 1982]), the barycentre set $\Omega = b(\mathcal{M})$ is a compact convex subset of the unit disc, and one easily checks that it has non-empty interior. The convexity of Ω means it is completely determined by its extremal points (those not expressible as a convex combination of other points). It is natural to ask questions about these extremal points, and about the boundary $\partial\Omega$.

Question 1. What *are* the points on the boundary? The only immediately obvious point is the fixed point 1, which supports a Dirac measure.
Question 2. What are the smoothness properties of $\partial\Omega$? Is it smooth? Piecewise smooth? Piecewise linear?
Question 3. Which invariant measures have barycentres on $\partial\Omega$? Do such *boundary measures* share a common structure?
Question 4. Is each point on $\partial\Omega$ the barycentre of a *unique* measure?

1991 Mathematics Subject Classification: 58F11, 58F15, 58F03.

Question 5. What are the ergodic properties of boundary measures? In particular, do such measures have positive entropy? Zero entropy?

It turns out that Ω is strictly convex, but that $\partial\Omega$ is highly non-smooth. In fact $\partial\Omega$ has a countable dense set of points of non-differentiability. This is the most pathological possible behaviour for the boundary of a convex planar set (see [Royden 1988]). Moreover, there is a fascinating parametrisation of $\partial\Omega$, which explains this non-differentiability in a way reminiscent of the frequency-locking of rotation numbers observed in parametrised families of degree-one circle maps.

The parametrisation of $\partial\Omega$ is by a remarkable one-parameter family of zero entropy measures known as *Sturmian*. A point lies on the boundary of Ω if and only if it is the barycentre of a Sturmian measure. There are many ways of defining Sturmian measures, the original (purely symbolic) definition going back to [Morse & Hedlund 1940]. For our purposes the most pertinent characterisation of Sturmian measures is the following (see [Bullett & Sentenac 1994]): a T-invariant measure is Sturmian if and only if its support is completely contained within some closed semi-circle.

In fact this support is always a rather thin set. It is either finite, or a Cantor set of zero Hausdorff dimension. Combinatorially, the dynamics on this support is a rotation (see [Bullett & Sentenac 1994]), so the measure can be assigned a *rotation number* between 0 and 1. This rotation number, which parametrises the family of Sturmian measures, is rational precisely when the support of the measure is finite. If the rotation number is irrational, then the support is contained in a *unique* closed semi-circle $[\lambda-\frac{1}{4},\lambda+\frac{1}{4}]$. If, however, the rotation number is rational, then there exists $[\lambda^-,\lambda^+]$ such that the support is contained in $[\lambda-\frac{1}{4},\lambda+\frac{1}{4}]$ for all $\lambda\in[\lambda^-,\lambda^+]$. This phenomenon, where the pre-image of a rational rotation number is a non-trivial parameter-space interval, is known as *frequency-locking*.

The relation to the boundary of the barycentre set is the following. Let $w(\theta)$ denote the barycentre with maximal component in the $2\pi\theta$ direction (i.e. its projection to the line through the origin making angle $2\pi\theta$ with the positive real axis is maximal). This defines a parametrisation of $\partial\Omega$ by the angle of an outward-pointing normal. Now if $w\in\partial\Omega$ is a point of *non-differentiability*, then its pre-image under this parametrisation is a whole interval in θ-space. That is, the parametrisation *locks* at the value w.

Below we state more precisely the main result about $\partial\Omega$. This was first conjectured in [Jenkinson 1996]. A reduced form of the main result was also conjectured by [Hunt & Ott 1996], who gave a heuristic argument for the support of the boundary measures having zero Hausdorff dimension. Very recently the main result has been proved by [Bousch 1998], who at the same time introduced an arsenal of techniques applicable to similar ergodic optimisation problems. In this article, an hors d'œuvre for the main (fish!) dish of [Bousch 1998], we describe some of the experimental work leading to the main result, and draw out the parallels with frequency-locking.

Main Result. *Let $T:z\mapsto z^2$ be the doubling map of the circle, and let Ω be the corresponding barycentre set.*

Ω is strictly convex, and its boundary $\partial\Omega$ has a countable dense set of points of non-differentiability. The set of boundary measures (those whose barycentres lie on $\partial\Omega$) is precisely the set of Sturmian measures. $w\in\partial\Omega$ is a point of non-differentiability if and only if it is an atomic Sturmian measure.

In section 2 we describe our experimental approach, and the results it gave. In section 3 we interpret these results in terms of Farey fractions. In section 4 we take a closer look at Sturmian measures, their various characterisations, and their parametrisation by rotation number. In section 5 we give the frequency-locking analogy, and in section 6 describe how this causes the non-differentiability on the boundary $\partial\Omega$.

2. The Approach

As well as the multiplicative circle $S^1 \subset \mathbb{C}$, it will also be convenient to work with the additive circle $K = [0, 1)$, with addition defined modulo one, so that the doubling map becomes $T(x) = 2x \pmod 1$. We let $\mathcal{O}(x)$ denote the T-orbit of a point $x \in K$. However, our measures $\mu \in \mathcal{M}$ will always be given on $S^1 \subset \mathbb{C}$, so that barycentres $b(\mu)$ lie in the unit disc in $\mathbb{C}$.

We will also use the symbolic model of the doubling map, obtained by associating each $x \in K$ with its binary expansion. Periodic points for T correspond to periodic sequences in $\{0,1\}^{\mathbb{N}}$. For convenience we represent such sequences by the finite repeated block which defines them. For example the period-2 point $1/3$ has symbolic representation 010101010101.... , which we abbreviate to simply 01. In fact we will use a single block to represent a whole periodic *orbit* - we simply use the block corresponding to the smallest point in the orbit. So, for example, 01 represents the orbit $\{1/3, 2/3\}$. T has precisely $2^n - 1$ periodic points of period n, each of the form $\frac{j}{2^n-1}$ for some $0 \le j \le 2^n - 2$.

In this section we calculate the barycentres of those invariant probability measures concentrated on single periodic orbits. Each periodic orbit supports a unique such measure, so without ambiguity we will talk of the barycentre corresponding to a particular periodic orbit. For example $b(01)$ will denote the barycentre of the invariant measure supported on the orbit $\{1/3, 2/3\}$ (which has symbolic representation 01). Such atomic measures are weak*-dense in $\mathcal{M}$, so the corresponding barycentres are dense in Ω. By computing sufficiently many we expect to obtain a reasonable approximation to Ω. If x has period n under T, then the corresponding barycentre is the trigonometric sum

$$\frac{1}{n}\sum_{r=0}^{n-1} e^{2\pi i 2^r x} \quad .$$

Starting with period 1 we systematically increase the period and calculate the barycentres of all periodic orbits of a given (least) period n. We performed this for all $n \le 19$ (thereby considering around 2^{20} periodic points).

The symmetry of Ω about the real line (due to the fact that T commutes with complex conjugation) halves the task, meaning we can ignore all barycentres with negative imaginary part. Moreover, the convexity of Ω means that at any stage we are only interested in those barycentres which extend the convex approximation we already have.

We can compute the first few barycentres by hand. The fixed point 0 has barycentre 1 in the complex plane. The period-2 orbit $\{1/3,\ 2/3\}$ has barycentre $-1/2$.

The barycentre of the period-3 orbit $\{1/7,\ 2/7,\ 4/7\}$ is

$$\frac{1}{3}\left(e^{2\pi i/7} + e^{4\pi i/7} + e^{8\pi i/7}\right) = -\frac{1}{6} + \frac{\sqrt{7}}{6}\, i \simeq -0.1666666 + 0.4409585\, i\ .$$

(We ignore the conjugate orbit $\{3/7, 5/7, 6/7\}$ which just gives conjugate barycentre). So far each new barycentre has extended our convex approximation, but for period 4 we have the following:

Symbolic code	Barycentre
0001	$0.125 + 0.4841229\, i$
0011	-0.25

Although $b(0001)$ *does* extend our convex region, it is clear that $b(0011)$ does not (since it lies between $b(0)$ and $b(01)$). We therefore "throw away" the point $b(0011)$ - that is, we can freely ignore it when considering whether subsequent barycentres extend the region of convexity.

For period 5 we have:

Symbolic code	Barycentre
00001	$0.3083872 + 0.443599\, i$
00011	$-0.0786801 + 0.1745122\, i$
00101	$-0.329707 + 0.2876896\, i$

A short calculation shows that $b(00001)$ and $b(00101)$ *do* extend our convex region, but $b(00011)$ does not.

For period-6 we have:

Symbolic code	Barycentre
000001	$0.4285381 + 0.3919765\, i$
000011	$0.0833333 + 0.2204793\, i$
000101	$-0.124271 + 0.3709144\, i$
000111	0
001101	$-0.3042671 + 0.0210621\, i$

Of these barycentres, only $b(000001)$ extends our region of convexity.

Let us take stock. So far we have a polygon whose extremal points are the barycentres of those periodic orbits with symbolic codes 0, 01, 001, 0001, 00001, 00101, and 000001. All these barycentres have a *chance* of being bona fide extremal points of Ω, though we don't know for sure whether they are. It is possible (in theory) that some of these points will be "outflanked" - that we will calculate new barycentres which will contain some of the old extremal barycentres in their convex hull. So far we note that this has *not* happened - any barycentre which has been extremal at the n^{th} stage has remained extremal at later stages. In fact in our subsequent calculations (up to and including period 19) we continue to observe this "persistence" of extremal points. This is the first hint of some sort of order in our observations. We might start to believe that a barycentre which is extremal at the n^{th} stage is in fact a bona fide extremal point of Ω.

What about the symbolic codes giving extremal barycentres? So far we note that all codes of the form 0^k1 are extremal. In fact the only extremal codes *not* of this form are the fixed point 0 and the period-5 orbit 00101. No other noticeable patterns have yet emerged.

Lastly, we note that already the calculations are becoming somewhat laborious. The least of our problems is that the trigonometric sums are becoming longer. More pertinent is that the *number* of periodic-n orbits is growing (exponentially fast). Even more of a problem is that each new barycentre must be compared with

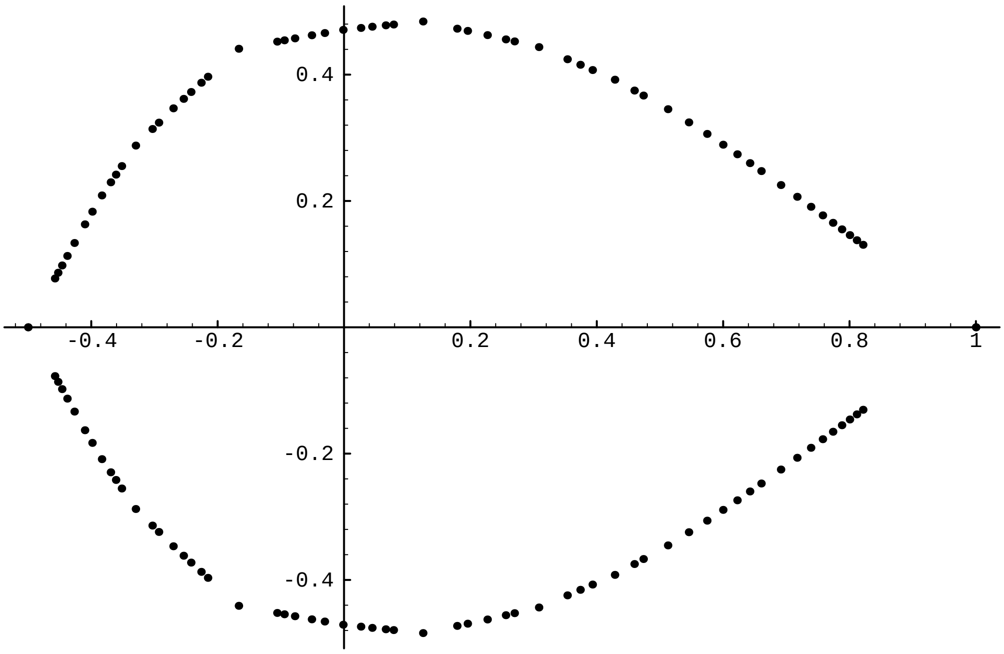

FIGURE 1. The 120 extremal points of the 19^{th} polygonal approximation to Ω

previous ones to see if it lies inside or outside the existing convex region - often this is a fairly delicate question, so we need to work with high precision. For all subsequent computations we used Mathematica.

In Figure 1 we plot the points obtained by considering orbits up to period 19, together with their complex conjugates. There are 120 such points, two of which lie on the real line.

Figure 1 strongly suggests that $\partial\Omega$ is non-differentiable. At the points 1 and $-\frac{1}{2}$ the non-differentiability seems particularly pronounced. Figure 1 also suggests that $\partial\Omega$ might in fact be piecewise linear. However the picture is misleading. Sufficiently precise calculations reveal that the "almost collinear" points on the boundary are not in fact collinear. Indeed it will turn out that the boundary does not contain any line segments.

3. The Farey Pyramid

The next step in our investigation was to examine the symbolic codes corresponding to those barycentres which we believe to be extremal. We observed that these codes seemed to be generated by a symbolic concatenation process, which we now describe. If we start with the codes 0 and 01 (the parents) we can concatenate to obtain 001 (the child), which we now place in between 0 and 01. We then repeat the process, concatenating each pair of adjacent codes and then placing the child in between each of its parents. So after the second stage we have the five codes 0, 0001, 001, 00101, 01. Continuing in this way we generate an infinite sequence of "levels", each level consisting of $2^{k-1}+1$ codes in (lexicographic) order.

We call the sequence of levels the *Farey pyramid*, as it resembles the well-known method of constructing Farey fractions (see [Hardy & Wright 1979], for example). The first four levels of the Farey pyramid are as follows.

Level									
1	0								01
2	0				001				01
3	0		0001		001		00101		01
4	0	00001	0001	0001001	001	00100101	00101	0010101	01

All the symbolic codes in the Farey pyramid are finite, and represent periodic orbits of the doubling map. Naturally this process also generates infinite non-periodic sequences if we concatenate infinitely often, thus giving an *extended Farey pyramid* (the infinite non-periodic sequences "at infinity" together with the finite codes in the Farey pyramid). The nature of this process means that all sequences in the extended Farey pyramid display a high level of self-similarity. Indeed the sequence of finite codes leading down to an infinite code at infinity is analogous to the sequence of continued fraction convergents to an irrational number. This analogy can be made precise if we replace each code by the frequency with which the symbol 1 appears in it.

In fact, the way the Farey codes are generated (by concatenation and interpolation) projects down to the barycentre set - a barycentre child has complex argument in between those of its barycentre parents.

Note that had we started with the symbolic codes 01 and 1 (this corresponds to coding the fixed point with a 1 rather than a 0) and repeated the concatenation procedure, we would have obtained a different (extended) Farey pyramid. We call this the *conjugate (extended) Farey pyramid*, since the codes obtained represent orbits conjugate to those in the (extended) Farey pyramid. In particular, the barycentres obtained are the complex conjugates of those arising from the Farey pyramid.

4. Sturmian Orbits and Sturmian Measures

Let us give some alternative characterisations of the sequences appearing in the extended Farey pyramid.

Definition 1. A sequence $\underline{x} = (x_1, x_2, \ldots) \in \{0,1\}^{\mathbb{N}}$ is said to be *Sturmian* if
(a) the number of 1's in any two sub-blocks of the same length differs by at most one, and
(b) it is recurrent (i.e. for all n, the length-n initial block of $\underline{x}$ occurs infinitely often in $\underline{x}$.)

If this sequence is the symbolic code of a point $x \in K$ then we say the orbit $\mathcal{O}(x)$ is a *Sturmian orbit.*

The simplest example of a non-Sturmian sequence is the periodic sequence defined by the finite block 0011 (recall from Section 2 that this was the first symbolic code which did *not* give an extremal barycentre).

Sturmian sequences were first studied by [Morse & Hedlund 1940]. Since then a variety of applications has been discovered, and an impressive body of literature has accumulated (see, for example, the extensive bibliographies in [Berstel] and [Brown 1993]). One application is in coding circle rotations. If $S : K \to K$ is rotation by angle α, then the points 0 and $1-\alpha$ define a partition of the circle into two intervals, which we label 0 and 1. The S-orbit of any point on the circle thus generates a sequence of 0's and 1's, and this sequence turns out to be always Sturmian. Another application is to coding trajectories in square billiard systems, which in turn is closely related to the notion of a cutting sequence of a line in the plane (see [Series 1985]). Sturmian sequences are also important in language theory. The non-periodic Sturmian sequences have sub-word complexity $n+1$ (i.e. there are precisely $n+1$ subwords of length n), the minimal possible complexity for a sequence which is not eventually periodic. Similarly, the periodic Sturmian sequences have sub-word complexity $n+1$ until this number reaches the least period of the sequence, after which the complexity function is constant.

In terms of the dynamics of the doubling map we have the following characterisations of Sturmian sequences (see [Bullett & Sentenac 1994]).

Lemma 1. *Let $x \in K$, and let $T : K \to K$ be the doubling map. The following are equivalent.*
(a) $\mathcal{O}(x)$ is a Sturmian orbit under T.
(b) The symbolic code for x belongs to either the extended Farey pyramid or the conjugate extended Farey pyramid.
(c) The closure of the T-orbit $\mathcal{O}(x)$ is minimal (i.e. contains no proper closed T-invariant subsets) and is contained in some closed semi-circle $[\delta, \delta+\frac{1}{2}] \subset K$.
(d) The closure of the T-orbit $\mathcal{O}(x)$ is minimal, and the orbit is ordered. That is, if a, b, c are points on $\mathcal{O}(x)$, then their cyclic order around K is preserved by T.

We remark (see [Bullett & Sentenac 1994]) that if $\mathcal{O}(x)$ is a *non-periodic* Sturmian orbit contained in the semi-circle $[\delta, \delta+\frac{1}{2}]$, then both of the endpoints δ, $\delta+\frac{1}{2}$ belong to the orbit. We can then define the symbolic code of $\mathcal{O}(x)$ to be the symbolic code of the smaller (as elements of $[0,1)$) endpoint. We will be interested in invariant measures supported on the closure of a Sturmian orbit. We have the following result.

Lemma 2. *The closure of a Sturmian orbit has zero Hausdorff dimension. It supports a unique T-invariant probability measure. This measure, which we will call a Sturmian measure, is ergodic and has zero entropy. With the trivial exception of the Dirac measure concentrated on the fixed point, a Sturmian measure is not weak-mixing.*

Proof. Zero Hausdorff dimension is proved in [Bullett & Sentenac 1994]. Unique ergodicity and the other ergodic properties follow from the fact that the shift acting on the closure of a Sturmian orbit is conjugate (off a countable set) to a circle rotation (see [Morse & Hedlund 1940]). □

If $x \in K$ has symbolic code $(x_1, x_2, \ldots)$, we define its *rotation number*

$$\rho(x) = \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i \ ,$$

provided this limit exists. If $\mathcal{O}(x)$ is Sturmian then the *ordered* property (see Lemma 1) means that $\rho(x)$ is a rotation number in the usual sense of the word, since $\mathcal{O}(x)$ has the same combinatorial order as a circle rotation by angle $\rho(x)$.

Indeed we can associate a continuous degree-one circle map to $\mathcal{O}(x)$ in a natural way - see the discussion at the end of Section 5. By Lemma 2 we can unambiguously assign a rotation number to each Sturmian measure.

In fact the rotation number defines a one-to-one correspondence between Sturmian orbits and the interval $[0,1]$, with rationals corresponding precisely to the *periodic* Sturmian orbits. The only slight ambiguity concerns the fixed point, which has two possible codings, though this is unimportant. The interval $[0,\frac{1}{2}]$ corresponds to the extended Farey pyramid, while $[\frac{1}{2},1]$ corresponds to the conjugate extended Farey pyramid.

5. Frequency-locking of rotation numbers

Now we describe the frequency-locking of rotation numbers of Sturmian orbits as we vary the semi-circle to which they belong. For $\lambda \in K$, let $C_\lambda = [\lambda - \frac{1}{4}, \lambda + \frac{1}{4}] \subset K$ be the closed semi-circle centred around λ. [Bullett & Sentenac 1994], following earlier work of [Gambaudo, Lanford & Tresser 1984] and [Veerman 1986], [Veerman 1987], proved the following results.

Proposition 1. *Each semi-circle $C_\lambda \subset K$ contains a unique minimal closed T-invariant set A_λ. This set is the closure of some Sturmian orbit.*

Proposition 2. *Let A_λ denote the unique minimal closed T-invariant set contained in the semi-circle $C_\lambda = [\lambda - \frac{1}{4}, \lambda + \frac{1}{4}] \subset K$. There is a sequence of disjoint non-trivial intervals $[\lambda_i^-, \lambda_i^+] \subset K$ with the following properties.*

(a) $\cup_{i=1}^\infty (\lambda_i^-, \lambda_i^+)$ has full Lebesgue measure. Its complement is a Cantor set of zero Hausdorff dimension.

(b) There is a one-to-one correspondence between intervals in the sequence and periodic Sturmian orbits. If $[\lambda_i^-, \lambda_i^+]$ is an interval, and $\mathcal{O}$ is the corresponding periodic Sturmian orbit, then $A_\lambda = \mathcal{O}$ for all $\lambda \in [\lambda_i^-, \lambda_i^+]$.

(c) There is a one-to-one correspondence between points of $\left[\cup_{i=1}^\infty [\lambda_i^-, \lambda_i^+]\right]^c$ and non-periodic Sturmian orbits. If $\lambda \in \left[\cup_{i=1}^\infty [\lambda_i^-, \lambda_i^+]\right]^c$, and $\mathcal{O}$ is the corresponding non-periodic Sturmian orbit, then $A_\lambda = \overline{\mathcal{O}}$.

(d) The rotation number map $\lambda \mapsto \rho(A_\lambda)$ is a weakly increasing continuous surjection of $[-\frac{1}{4}, \frac{3}{4}]$ onto $[0,1]$. It is constant on each of the intervals $[\lambda_i^-, \lambda_i^+]$.

The map $\lambda \mapsto \rho(A_\lambda)$, which by Proposition 2 is continuous, weakly monotonic, locally constant on a set of full measure but not globally constant, is called a *devil's staircase*. The phenomenon of *frequency-locking* of rotation numbers at rational values is well-known in the theory of parametrised families of degree-one circle maps. If S_λ is such a parametrised family then certain mild conditions ensure that the map associating rotation numbers (in the usual sense of the term) to parameter values is a devil's staircase, and that the inverse image of each rational value is a non-trivial interval. Further details can be found on page 392 of [Katok & Hasselblatt 1995], or in [Newhouse, Palis & Takens 1983]. The connection to our situation arises because each A_λ is an orbit closure lying in a semi-circle C_λ. The corresponding T-orbit is therefore also an orbit of a certain continuous degree-one map S_λ. We simply define S_λ to be the restriction of the doubling map on C_λ and constant otherwise. For further details see [Boyland 1986] or [Veerman 1986].

6. A Parametrisation of the boundary

We will parametrise the boundary of the barycentre set Ω by $\theta \in K$, which we think of as indexing a normal to $\partial\Omega$ in the $2\pi\theta$ direction. We let $q(\theta) = \sup_{w\in\Omega} < w, e^{2\pi i\theta} >$ denote the maximal component in the $2\pi\theta$ direction, and $w(\theta) = \{w \in \Omega : < w, e^{2\pi i\theta} > \ = q(\theta)\}$ the set of barycentres achieving this maximum. Here $< \cdot, \cdot >$ is the usual inner product.

Clearly any $w(\theta)$ must lie in $\partial\Omega$, and is either a point or a line segment, and conversely any point of $\partial\Omega$ must belong to some $w(\theta)$.

We remark that $q(\theta) = \sup_{\mu\in\mathcal{M}} \int f_\theta d\mu$, where $f_\theta(x) = \cos 2\pi(x-\theta) =< e^{2\pi i x}, e^{2\pi i\theta} >$. This variational characterisation of $q(\theta)$ motivates the following definition.

Definition 2. For a given continuous function $f : K \to \mathbb{R}$ we say the measure $m \in \mathcal{M}$ is *f-maximal* if $\int f dm = \sup_{\mu\in\mathcal{M}} \int f d\mu$.

The weak* compactness of $\mathcal{M}$ ensures that a maximal measure always exists, though in general it need not be unique. From the above discussion we clearly have the following.

Lemma 3. *The barycentre $b(\mu)$ lies on the boundary $\partial\Omega$ if and only if μ is f_θ-maximal for some $\theta \in K$.*

We now consider the smoothness properties of $\partial\Omega$, and relate them to the behaviour of the parametrisation $\theta \mapsto w(\theta)$.

We will consider separately the two symmetric halves of $\partial\Omega$, thinking of them as the graphs of two real-valued functions. Let $\partial\Omega = \partial\Omega^+ \cup \partial\Omega^-$, where

$$\partial\Omega^+ = \{w \in \partial\Omega : \mathrm{Im}(w) \geq 0\} \qquad \text{and} \qquad \partial\Omega^- = \{w \in \partial\Omega : \mathrm{Im}(w) \leq 0\} \ .$$

Let $g^+ : \Omega \cap \mathbb{R} \to \mathbb{R}$ be the function whose graph is $\partial\Omega^+$ (thinking of $\partial\Omega^+$ as lying in $\mathbb{R}^2$). The convexity of Ω means that g^+ is a concave function, so by [Royden 1988], page 113, its left and right derivatives exist at every point, with the left derivative always greater than or equal to the right derivative. Moreover, the left and right derivatives are equal to each other except on a countable set. Analogous differentiability properties hold for the convex function $g^- = -g^+$ whose graph is $\partial\Omega^-$. In particular, g^+ and g^- both have at most countably many points of non-differentiability.

We say that $w = c + ig^+(c)$ (resp. $w = c + ig^-(c)$) is a *point of differentiability* of $\partial\Omega$ if $c \in \Omega \cap \mathbb{R}$ is a point of differentiability of g^+ (resp. g^-). Otherwise we say that w is a *point of non-differentiability*. The above analysis omits a discussion of the possible differentiability at the points $1, -1/2$, where $\partial\Omega$ intersects the real line. This could easily be done by considering $\partial\Omega = \partial\Omega^r \cup \partial\Omega^l$, say, where

$$\partial\Omega^r = \{w \in \partial\Omega : \mathrm{Re}(w) \geq 0\} \qquad \text{and} \qquad \partial\Omega^l = \{w \in \partial\Omega : \mathrm{Re}(w) \leq 0\} \ ,$$

and introducing functions $g^r, g^l : \Omega \cap i\mathbb{R} \to \mathbb{R}$ whose graphs are $\partial\Omega^r, \partial\Omega^l$ respectively. In fact it turns out that both 1 and $-1/2$ are points of *non-differentiability* of $\partial\Omega$.

With the above definitions, we see that $\partial\Omega$ can have at most countably many points of non-differentiability. If $w \in \partial\Omega$ is a point of non-differentiability then the left and right gradients to $\partial\Omega$ at w do not agree, so that $\partial\Omega$ does not have a unique tangent at w. Rather, there is a non-trivial interval $[\theta^-, \theta^+]$ such that, for any $\theta \in [\theta^-, \theta^+]$, the line through w of gradient $\tan(2\pi\theta + \pi/2)$ is a tangent to $\partial\Omega$. In other words, $\theta \mapsto w(\theta)$ is constant on the interval $[\theta^-, \theta^+]$. So bad behaviour (non-differentiability) of $\partial\Omega$ corresponds to good behaviour (local constancy) of $\theta \mapsto w(\theta)$. In terms of the parametrised family f_θ, a sufficient (but a priori not necessary) condition for the local constancy of $\theta \mapsto w(\theta)$ on $[\theta^-, \theta^+]$ is the existence of some measure μ which is the unique f_θ-maximal measure for all $\theta \in [\theta^-, \theta^+]$.

This discussion, and the strong analogy with Proposition 2, led us [Jenkinson 1996] to conjecture the following theorem, which in particular implies the main result described in the introduction. This theorem was proved by [Bousch 1998]. As notation, if $\mathcal{O}$ is a Sturmian orbit, then $\mu_{\mathcal{O}}$ will denote the corresponding Sturmian measure (the unique T-invariant probability measure supported on $\overline{\mathcal{O}}$).

Theorem. *Let $T : K \to K$ be the doubling map of the circle, and let $f_\theta : K \to \mathbb{R}$ be the family of functions defined by $f_\theta(x) = \cos 2\pi(x - \theta)$. There is a sequence of disjoint non-trivial intervals $[\theta_i^-, \theta_i^+] \subset K$ with the following properties.*

(a) $\cup_{i=1}^\infty(\theta_i^-, \theta_i^+)$ has full Lebesgue measure. Its complement in K is a Cantor set of zero Hausdorff dimension.

(b) There is a one-to-one correspondence between intervals in the sequence and periodic Sturmian orbits. If $[\theta_i^-, \theta_i^+]$ is an interval, and $\mathcal{O}$ is the corresponding periodic Sturmian orbit, then the Sturmian measure $\mu_{\mathcal{O}}$ is the unique f_θ-maximal measure for all $\theta \in [\theta_i^-, \theta_i^+]$.

(c) There is a one-to-one correspondence between points of $\left[\cup_{i=1}^\infty[\theta_i^-, \theta_i^+]\right]^c$ and non-periodic Sturmian orbits. If $\theta \in \left[\cup_{i=1}^\infty[\theta_i^-, \theta_i^+]\right]^c$, and $\mathcal{O}$ is the corresponding non-periodic Sturmian orbit, then the Sturmian measure $\mu_{\mathcal{O}}$ is the unique f_θ-maximal measure.

(d) Let $\mathcal{O}_\theta$ be the Sturmian orbit corresponding to parameter value θ. There is a length one ordered interval $I \subset K$ such that the map $\theta \mapsto \rho(\mathcal{O}_\theta)$ is a weakly increasing continuous surjection of I onto $[0,1]$. It is constant on each of the intervals $[\theta_i^-, \theta_i^+]$.

Note that if in the above Theorem we replace the family of functions f_θ by the family χ_{C_λ} of characteristic functions of the semi-circles $C_\lambda = [\lambda - \frac{1}{4}, \lambda + \frac{1}{4}]$, then we have precisely Proposition 2, where the ordered interval I is $[-\frac{1}{4}, \frac{3}{4}]$. One might imagine that the intervals of constancy of $\theta \mapsto \rho(\mathcal{O}_\theta)$ are precisely the same as those of $\lambda \mapsto \rho(A_\lambda)$, so that the maximal measure for any f_θ is the unique one whose support lies in the semi-circle where f_θ is positive. However, this is clearly *not* the case. For example, the fixed point $0 \in K$ is the unique Sturmian orbit contained in the semi-circle C_λ for all $\lambda \in [-\frac{1}{4}, \frac{1}{4}]$. Thus the map $\lambda \mapsto \rho(A_\lambda)$ is constant on the interval $[-\frac{1}{4}, \frac{1}{4}]$. In contrast, the map $\theta \mapsto \rho(\mathcal{O}_\theta)$ is clearly *not* constant on the whole of $[-\frac{1}{4}, \frac{1}{4}]$. To see this just note that the Dirac measure concentrated on the fixed point $x = 0$ is maximal for the function $f_0(x) = \cos 2\pi x$, while it is certainly not maximal for $f_{\frac{1}{4}}(x) = \sin 2\pi x$.

[Bullett & Sentenac 1994] show that if an ordered periodic orbit has (rational) rotation number p/q (in lowest terms), then the corresponding interval of constancy has length $\frac{1}{2(2^q-1)}$. An open problem is to describe the scaling behaviour of the lengths of the intervals of constancy of $\theta \mapsto \rho(\mathcal{O}_\theta)$. Below we give the intervals of constancy corresponding to the symbolic codes on the first three levels of the Farey pyramid. Note that the intervals corresponding to the conjugate codes (i.e. obtained by swapping 0's and 1's) in the conjugate Farey pyramid are obtained by reflecting in the mid-point $\frac{1}{2}$ of our parameter circle K. The codes 0 and 01, whose corresponding orbits are symmetric in the circle, have intervals of constancy which are also symmetric.

Code	Interval of Constancy	Length
0	$[-0.149550,\ 0.149550]$	0.2991
01	$[0.420148,\ -0.420148]$	0.159704
001	$[0.279199,\ 0.367215]$	0.088016
0001	$[0.216946,\ 0.266213]$	0.049267
00101	$[0.374417,\ 0.404815]$	0.030398

References

[Berstel] J. Berstel, Recent results in Sturmian words, *Developments in Language Theory (Magdeburg 1995)*, World Scientific, 1996, p. 13–24.

[Bousch 1998] T. Bousch, Le poisson n'a pas d'arêtes, *preprint*, Laboratoire de Topologie et Dynamique, Université d'Orsay, 1998.

[Boyland 1986] P. L. Boyland, Bifurcations of circle maps: Arnol'd tongues, bistability and rotation intervals, *Comm. Math. Phys.*, **106** (1986), 353–381.

[Brown 1993] T. C. Brown, Descriptions of the characteristic sequence of an irrational, *Canad. Math. Bull.*, **36** (1993), 15–21.

[Bullett & Sentenac 1994] S. Bullett and P. Sentenac, Ordered orbits of the shift, square roots, and the devil's staircase, *Math. Proc. Camb. Phil. Soc.*, **115** (1994), 451–481.

[Gambaudo, Lanford & Tresser 1984] J. M. Gambaudo, O. Lanford and C. Tresser, Dynamique symbolique des rotations, *C. R. Acad. Sci. Paris*, **1299** (1984), 823–825.

[Hardy & Wright 1979] G. H. Hardy and E. M. Wright, *An Introduction to the Theory of Numbers* (5th edition), Oxford, 1979.

[Hunt & Ott 1996] B. R. Hunt and E. Ott, Optimal periodic orbits of chaotic systems occur at low period, *Phys. Rev. E*, **54** (1996), 328–337.

[Jenkinson 1996] O. Jenkinson, Conjugacy rigidity, cohomological triviality, and barycentres of invariant measures, *Ph.D. Thesis*, Warwick University, (1996), available at `http://iml.univ-mrs.fr/~omj/`

[Katok & Hasselblatt 1995] A. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems*, Cambridge University Press, 1995.

[Morse & Hedlund 1940] M. Morse and G. A. Hedlund, Symbolic Dynamics II. Sturmian Trajectories, *Am. J. Math.*, **62** (1940), 1–42.

[Newhouse, Palis & Takens 1983] S. Newhouse, J. Palis and F. Takens, Bifurcations and stability of families of diffeomorphisms, *I.H.E.S. Publ. Math.*, **57** (1983), 5–72.

[Royden 1988] H. L. Royden, *Real Analysis* (3rd Edition), Macmillan, 1988.

[Series 1985] C. Series, The geometry of Markoff numbers, Math. Intelligencer, **7**, no. 3 (1985), 20–29.

[Veerman 1986] P. Veerman, Symbolic dynamics and rotation numbers, *Physica*, **134A** (1986), 543–576.

[Veerman 1987] P. Veerman, Symbolic dynamics of order-preserving orbits, *Physica*, **29D** (1987), 191–201.

[Walters 1982] P. Walters, *An Introduction to Ergodic Theory*, Springer-Verlag, Berlin, Heidelberg, New York, 1982.

UPR 9016 CNRS, Institut de Mathématiques de Luminy, 163 avenue de Luminy, case 907, 13288 Marseille, cedex 9, France

E-mail address: `omj@iml.univ-mrs.fr` , `omj@maths.warwick.ac.uk`