The Beta Expansion as a Dynamical System

Linas Vepštas

September 2020

linasvepstas@gmail.com

Abstract

A beta-expansion is a representation of a real number x written in the form $x = \sum_n k_n \beta^{-n}$ for a fixed $1 < \beta \leq 2$ and a string of binary digits k_n. Such expansions are not unique, as for almost all x and almost all β there are an uncountable number of equivalent expansions.

This text explores these classes of equivalent expansions. They can be organized into an infinite binary tree (whence uncountability). This tree induces a measure-preserving dynamical system, a subshift on the Cantor space $\{0,1\}^\omega$ that is an extension of the subshift arising from the iterated map $\beta x \mod 1$, commonly called the β-transformation. The resulting invariant measure (the Ruelle-Frobenius-Perron measure) is explored numerically and compared to the Renyi-Gelfond-Parry measure for the β-transformation. The distance between branch-points in the infinite binary tree can be expressed in terms of this measure.

1 Introduction

The the β-transform is the iterated map $\beta x \mod 1$ applied to the unit interval on the real number line. A very special case of this map is $\beta = 2$, which is named the "dyadic transform". The dyadic transform is, in a certain sense, "completely solvable", in that it has been heavily studied and is well-understood. It's dynamics is described by the Bernoulli process. It is ergodic, and corresponds to a shift operator acting on the Cantor space $\{0,1\}^\omega$: this is the shift that lops off one bit from an infinite bit-string. The shift can be expressed as a dynamical system; the function field on the dynamical system is described by the transfer operator (the Ruelle-Frobenius-Perron operator). The natural setting for the function field is on the Borel algebra (the sigma algebra) associated to the Cantor space. It splits naturally into two parts: an invariant measure (the Bernoulli measure), and a dissipative part. The spectrum of the transfer operator is discrete, when limited to the real-analytic functions; the spectrum is continuous when acting on square-integrable functions.

Not as much is known for the case of $\beta \neq 2$. The β-transform is still a shift, in that it can still be understood as a way of lopping off one bit from an infinite bit-string in $\{0,1\}^\omega$. However, not all possible bit-strings can occur; it is a subshift, in that there is a subspace $B \subset \{0,1\}^{\omega}$ that is invariant under the β-transform. The Bernoulli measure, when restricted to this subshift, and then mapped to the unit interval in the conventional way, becomes the invariant measure of the β-transform. The existence of an invariant measure was first demonstrated by Renyi,[1] and given explicit form by Gelfond[2] and Parry[3]. As an invariant measure, it is the eigenvalue-one (the Frobenius-Perron eigenvalue) eigenfunction of the transfer operator. The remaining spectrum of the transfer operator is poorly understood, but is explored by this author in a different text.[4] Understanding the dissipative spectrum is important, as it explicitly describes the rate of convergence of numerical explorations.

The Cantor space can be mapped to the reals in the unit interval, by interpreting a bit-string in $\{0,1\}^{\omega}$ as a binary expansion of a real number. Famously, the mapping is surjective, in that the dyadic rationals have two inequivalent binary expansions: a dyadic rational can be written either as $0.b_0 b_1 b_2 \cdots b_k 1000\cdots$ ending in an infinite string of zeros, or as $0.b_0 b_1 b_2 \cdots b_k 0111\cdots$ ending in an infinite string of one's. The dyadic rationals are in one-to-one correspondence with the "gaps" in the (conventional) Cantor set (the one with middle-thirds removed.) These gaps are also in one-to-one correspondence with the "gaps" between the left-right subtrees in an infinite tree; they are also in one-to-one correspondence with the cusps of modular forms (as the infinite binary tree can be mapped into the modular group). For the general β-shift, $\beta \neq 2$,the situation changes: rather than there being only two (in-)equivalent binary expansions, there are 2^{ω} (that is, uncountably many) of them! This text is an exploration of these expansions for the general case.

The remainder of the introduction describes the β-transform and the conventional mapping of the Cantor set to the reals. This is followed by a short review of dynamical systems, measure theory and function fields; it concludes with a literature review of prior results on the β-transform.

The main results follow in order: the "gaps" can be organized into an infinite binary tree; the locations of the branch points in the tree can be explicitly specified; the distance between branch-points can be explicitly stated. The specification of the branch-points induces a dynamical system; it has an invariant measure (which is an extension of the Gelfond-Parry measure). The distances between branch-points can be directly obtained from this measure. The text terminates with a short conclusion and some speculations.

1.1 Beta Shift

The beta shift is the iterated function

$$T_\beta(x) = \begin{cases} \beta x & \text{for } 0 \leq x < \frac{1}{2} \\ \beta\left(x - \frac{1}{2}\right) & \text{for } \frac{1}{2} \leq x \leq 1 \end{cases} \tag{1}$$

where, by convention, one takes $1 < \beta \leq 2$. Given a real number x in the unit interval $0 \leq x \leq 1$, the orbit $T_\beta^n(x)$ of the point x under the iterated shift generates a sequence of digits

$$k_n = k_n(x) = \begin{cases} 0 & \text{if } 0 \le T_\beta^n(x) < \frac{1}{2} \\ 1 & \text{if } \frac{1}{2} \le T_\beta^n(x) \le 1 \end{cases} \tag{2}$$

This bit-sequence is sometimes referred to as the "symbolic dynamics" of the orbit. The symbolic dynamics can be inverted to reconstruct the original value of x whenever $1 < \beta$. This is explicitly given by

$$x = \frac{k_0}{2} + \frac{1}{\beta}\left(\frac{k_1}{2} + \frac{1}{\beta}\left(\frac{k_2}{2} + \frac{1}{\beta}\left(\frac{k_3}{2} + \frac{1}{\beta}(\cdots)\right)\right)\right)$$

Written this way, the $T_\beta(x)$ clearly acts as a shift on this sequence:

$$T_\beta(x) = \frac{k_1}{2} + \frac{1}{\beta}\left(\frac{k_2}{2} + \frac{1}{\beta}\left(\frac{k_3}{2} + \frac{1}{\beta}\left(\frac{k_4}{2} + \frac{1}{\beta}(\cdots)\right)\right)\right)$$

That is, the shift T_β lops off the leading digit from the sequence $k_0, k_1, k_2, \cdots$. Multiplying out the above sequence, one obtains the so-called "β-expansion" of a real number x, namely the series

$$x = \frac{1}{2}\sum_{n=0}^{\infty}\frac{k_n}{\beta^n} \tag{3}$$

This map from (a subspace of) the Cantor space $\{0,1\}^\omega$ to the unit interval can be taken as a homomorphism of the β-shift. That is, we can take the β-shift as acting in the space of bit-strings, or, equivalently, acting on the unit interval.

Not every possible bit-sequence occurs in this system. Let $S_\beta \subset \{0,1\}^\omega$ denote the set of all bit-strings given by eqn 2. For $\beta \neq 2$ it is a proper subset, and by construction, the shift is an injection: $T_\beta S_\beta \subseteq S_\beta$. Setting $\beta = 2$ in eqn 2 gives the Bernoulli shift. This is a very special case; the Bernoulli shift is the full shift, in that $S_2 = \{0,1\}^\omega$ and $T_2 S_2 = S_2$. The homomorphism deserves explicit mention, as it plays a distinctive role as the canonical map from the Cantor space to the unit interval:

$$x = \frac{1}{2}\sum_{n=0}^{\infty}\frac{b_n}{2^n} \tag{4}$$

where the b_n are meant to be the bit-sequences obtained from iterating $b(x) = T_2(x)$.

1.2 Beta Transformation

After exactly one iteration of the beta shift, all initial points $\beta/2 \le x \le 1$ are swept up into the domain $0 \le x < \beta/2$, and never leave. Likewise, the range of the iterated beta-shift is $0 \le x < \beta/2$. Thus, an alternative representation of the beta shift, filling the entire unit square, can be obtained by dividing both the domain and range by $\beta/2$ to obtain the function

$$t_\beta(u) = \begin{cases} \beta u & \text{for } 0 \le u < \frac{1}{\beta} \\ \beta u - 1 & \text{for } \frac{1}{\beta} \le u \le 1 \end{cases} \tag{5}$$

This can be written more compactly as $t_\beta(x) = \beta x \mod 1$. In this form, the function is named "the beta-transform", written as the β-transformation, presenting a typesetting challenge to search engines when used in titles of papers. The orbit of a point x in the beta-shift is identical to the orbit of a point $u = 2x/\beta$ in the beta-transformation. Explicitly comparing to the beta-shift of eqn 1:

$$T_\beta^n(x) = \frac{\beta}{2} t_\beta^n\left(\frac{2x}{\beta}\right)$$

The beta-shift and the β-transformation are essentially "the same function"; this text works almost exclusively with the beta-shift, and is thus idiosyncratic, as it flouts the more common convention of working with the β-transformation. The primary reason for doing this is that the β-shift retains better contact the idea of being a subshift of the full shift.

1.3 Dynamical Systems

A brief review of dynamical systems is in order, as it provides a coherent language with which to talk about and think about the beta-shift. The technical reason for this is that a subshift $S \subset \{0,1\}^\omega$ provides a more natural setting for the theory, and that a lot of the confusion about what happens on the unit interval is intimately entangled with the homomorphism 4 (or 3 as the case may be). Disentangling the subshift from the homomorphism provides a clearer insight into what phenomena are due to which component.

The review of dynamical systems here is more-or-less textbook-standard material; it is included here only to provide a firm grounding for later discussion.

The Cantor space $\{0,1\}^\omega$ can be given a topology, the product topology. The open sets of this topology are called "cylinder sets". These are the infinite strings in three symbols: a finite number of 0 and 1 symbols, and an infinite number of * symbols, the latter meaning "don't care". Set union is defined location-by-location, with $0 \cup * = 1 \cup * = *$ and set intersection as $0 \cap * = 0$ and $1 \cap * = 1$. Set complement exchanges 0 and 1 and leaves * alone: $\overline{0} = 1$, $\overline{1} = 0$ and $\overline{*} = *$. The topology is then the collection of all cylinder sets. Note that the intersection of any finite number of cylinder sets is still a cylinder set, as is the union of an infinite number of them. The product topology does *not* contain any "points": strings consisting solely of just 0 and 1 are *not* allowed in the topology. By definition, topologies only allow finite intersections, and thus don't provide any way of constructing "points". Of course, points can always be added "by hand", but doing so tends to generate a topology (the "box topology") that is "too fine"; in particular, the common-sense notions of a continuous function are ruined by fine topologies. The product topology is "coarse".

The Borel algebra, or sigma-algebra, takes the topology and also allows set complement. This effectively changes nothing, as the open sets are still the cylinder sets, although now they are "clopen", as they are both closed and open.

Denote the Borel algebra by $\mathscr{B}$. A shift is now a map $T : \mathscr{B} \to \mathscr{B}$ that lops off the leading symbol of a given cylinder set. This is provides strong theoretical advantages over working with "point dynamics": confusions about counting points and orbits and defining densities go away. This is done by recasting discussion in terms of functions $f:\mathscr{B}\to\mathbb{R}$ from Borel sets to the reals (or the complex numbers $\mathbb{C}$ or other fields, when this is interesting). An important class of such functions are the measures. These are functions $\mu:\mathscr{B}\to\mathbb{R}$ that are positive, and are "compatible" with the sigma algebra, in that $\mu(A\cup B)=\mu(A)+\mu(B)$ whenever $A\cap B=\varnothing$ and (for product-space measures) that $\mu(A\cap B)=\mu(A)\mu(B)$ for all $A,B\in\mathscr{B}$. The measure of the total space $\Omega=\{0,1\}^{\omega}$ is by convention unity: $\mu(\Omega)=1$.

The prototypical example of a measure is the Bernoulli measure, which assigns probability p to any string containing a single 0 and the rest all *'s. By complement, a string containing a single 1 and the rest all *'s has probability $1-p$. The rest follows from the sigma algebra: a cylinder set consisting of m zeros and n ones has measure $p^m(1-p)^n$. It is usually convenient to take $p=1/2$, the "fair coin"; the Bernoulli process is a sequence of coin tosses.

The map given in equation 4 is a homomorphism from the Cantor space to the unit interval. It extends naturally to a map from the Borel algebra $\mathscr{B}$ to the algebra of intervals on the unit interval. It is not an isomorphism: cylinder sets are both open and closed, whereas intervals on the real number line are either open, or closed (or half-open). It is convenient to take the map as a map to closed intervals, so that it's a surjection onto the reals, although usually, this detail does not matter. What does matter is if one takes $p=1/2$, then the Bernoulli measure is preserved: it is mapped onto the conventional measure on the real-number line. Thus, the cylinder set $0***\cdots$ is mapped to the interval $[0,1/2]$ and $1***\cdots$ is mapped to $[1/2,1]$ and both have a measure of $1/2$ and this extends likewise to all intersections and unions. Points have a measure of zero. That is, the homomorphism 4 preserves the fair-coin Bernoulli measure.

Although the map of eqn 4 is not an isomorphism, it is more-or-less safe for most practical purposes to treat it as such. For any closed interval, there is a unique corresponding cylinder set. The only trouble arises at a countable number of points: all dyadic rationals have two expansions, not one, and so there, the map is not invertable. However, the points all have a measure of zero in both cases, and there are only a countable number of such points, and the mapping always takes points to points. Thus, the inverse exists for almost all members of the algebra and is unique whenever an interval has positive measure.

Much of what is said above still holds for subshifts. Recall, a subshift S is a subspace $S\subset\{0,1\}^{\omega}$ that is invariant under the shift T, so that $TS=S$. The space S inherits a topology from $\{0,1\}^{\omega}$; this is the subspace topology. The Borel algebra $\mathscr{B}$ is similarly defined, as are measures. One can now (finally!) give a precise definition for an invariant measure: it is a measure μ such that $\mu\circ T^{-1}=\mu$, or more precisely, for which $\mu\left(T^{-1}(\sigma)\right)=\mu(\sigma)$ for almost all cylinder sets $\sigma\in S$. This is what shift invariance looks like. Note carefully that T^{-1} and not T is used in the definition. This is because T^{-1} is a surjection while T is not: every cylinder set σ in the subshift "came from somewhere"; we want to define invariance for all σ and not just for some of them.

The T^{-1} is technically called a "pushforward", and it defines a linear operator $\mathscr{L}_T$ on the space $\mathscr{F}$ of all functions $f:\mathscr{B}\to\mathbb{R}$. It is defined as $\mathscr{L}_T:\mathscr{F}\to\mathscr{F}$ by setting $\mathscr{L}_T:f\mapsto f\circ T^{-1}$. It is obviously linear, in that $\mathscr{L}_T(af+bg)=a\mathscr{L}_T(f)+$

$b\mathscr{L}_T(g)$. This pushforward is canonically called the “transfer operator” or the “Ruelle-Frobenius-Perron operator”. Like any linear operator, it has a spectrum. The precise spectrum depends on the space $\mathscr{F}$.

The canonical example is again the Bernoulli shift. For this, we invoke the inverse of the mapping of eqn 4 so that $f:[0,1]\to\mathbb{R}$ is a function defined on the unit interval, instead of $f:\mathscr{B}\to\mathbb{R}$. When $\mathscr{F}$ is the space of real-analytic functions on the unit interval, that is, the closure of the space of all polynomials in $x\in[0,1]$, then the spectrum of $\mathscr{L}_T$ is discrete. It consists of the Bernoulli polynomials $B_n(x)$ corresponding to an eigenvalue of 2^{-n}. That is,$\mathscr{L}_T B_n = 2^{-n}B_n$. Note that $B_0(x)=1$ is the invariant measure on the full shift. For the space of square-integrable functions $f:[0,1]\to\mathbb{R}$, the spectrum of $\mathscr{L}_T$ is continuous, and consists of the unit disk in the complex plane; the corresponding eigenfunctions are fractal. Even more interesting constructions are possible; the Minkowski question mark function provides an example of a measure on $\{0,1\}^\omega$ that is invariant under the shift defined by the Gauss map $h(x)\mapsto\frac{1}{x}-\left\lfloor\frac{1}{x}\right\rfloor$. That is, as a measure, it solves $\mathscr{L}_T ?' = ?'$ with ? the Minkowski question mark function, and $?'$ it’s derivative; note that the derivative is “continuous nowhere”. This rather confusing idea (of something being “continuous nowhere”) can be completely dispelled by observing that it is well-defined on all cylinder sets in $\mathscr{B}$ and is finite on all of them – not only finite, but less than one, as any good measure must obey.

These last examples are mentioned so as to reinforce the idea that working with $\mathscr{B}$ instead of the unit interval $[0,1]$ really does offer some strong conceptual advantages. They also reinforce the idea that the Bernoulli shift is not the only “full shift”. In the following text, we will be working with subshifts, primarily the beta-shift, but will draw on ideas from the above so as to make rigorous statements about measureability and invariance, without having to descend into either *ad hoc* hand-waving or provide painfully difficult (and confusing) reasoning about subsets of the real-number line.

1.4 Beta Transformation Literature Review and References

The β-transformation, in the form of $t_\beta(x)=\beta x \mod 1$ has been well-studied over the decades. The beta-expansion 2 was introduced by A. Renyi[1], who demonstrates the existence of the invariant measure. The ergodic properties of the transform were proven by W. Parry[3], who also shows that the system is weakly mixing.

An explicit expression for the invariant measure was independently obtained by A.O. Gel’fond[2] and by W. Parry[3], as a summation of step functions

$$\nu_\beta(y)=\frac{1}{F}\sum_{n=0}^{\infty}\frac{\varepsilon_n(y)}{\beta^n} \tag{6}$$

where ε_n is the digit sequence

$$\varepsilon_n(y)=\begin{cases}0 & \text{if } t_\beta^n(1)\le y\\ 1 & \text{otherwise}\end{cases}$$

and F is a normalization constant. By integrating $\varepsilon_n(y)$ under the sum, the normalization is given by

$$F=\sum_{n=0}^{\infty}\frac{t_{\beta}^{n}(1)}{\beta^{n}}$$

Similar to the way in which a dyadic rational $p/2^n$ has two different binary expansions, one ending in all-zeros, and a second ending in all-ones, so one may also ask if and when a real number x might have more than one β-expansion (for fixed β). In general, it can; N. Sidorov shows that almost every number has a continuum of such expansions![5] It is this result that this text primarily concerns itself with.

Conversely, the "univoke numbers" are those values of β for which there is only one, unique expansion for $x=1$. These are studied by De Vries.[6]

The β-transformation has been shown to have the same ergodicity properties as the Bernoulli shift.[7] The fact that the beta shift, and its subshifts are all ergodic is established by Climenhaga and Thompson.[8]

An alternative to the notion of ergodicity is the notion of universality: a β-expansion is universal if, for any given finite string of bits, that finite string occurs somewhere in the expansion. This variant of universality was introduced by Erdös and Komornik[9]. Its is shown by N. Sidorov that almost every β-expansion is universal.[10] Conversely, there are some values of β for which rational numbers have purely periodic β-expansions;[11] all such numbers are Pisot numbers.[12]

The symbolic dynamics of the beta-transformation was analyzed by F. Blanchard[13]. A characterization of the periodic points are given by Bruno Maia[14]. A discussion of various open problems with respect to the beta expansion is given by Akiyama.[15]

When the beta expansion is expanded to the entire real-number line, one effectively has a representation of reals in a non-integer base. One may ask about arithmetic properties, such as the behavior of addition and multiplication, in this base - for example, the sum or product of two β-integers may have a fractional part! Bounds on the lengths of these fractional parts, and related topics, are explored by multiple authors.[16, 17, 18]

Certain values of β - generally, the Pisot numbers, generate fractal tilings,[19, 20, 21, 11, 15] which are generalizations of the Rauzy fractal. An overview, with common terminology and definitions is provided by Akiyama.[22] The tilings, sometimes called (generalized) Rauzy fractals, an be thought of as living in a direct product of Euclidean and p-adic spaces.[23]

The set of finite beta-expansions constitutes a language, in the formal sense of model theory and computer science. This language is recursive (that is, decidable by a Turing machine), if and only if β is a computable real number.[24]

The zeta function, and a lap-counting function, are given by Lagarias[25]. The Hausdorff dimension, the topological entropy and general notions of topological pressure arising from conditional variational principles is given by Daniel Thompson[26]. A proper background on this topic is given by Barreira and Saussol[27].

Except for the results from Sidorov, none of the topics or results cited above are made use of, or further expanded on, or even touched on below.

2 Beta expansions

The Bernoulli process and it's relationship to the Cantor set provides the baseline to which the beta-shift can be compared. The beta-shift has a number of very striking differences, however, which manifest themselves in the structure of the beta-transformations. This section provides the main results being reported in this text.

More precisely, again consider the binary expansion of a real number

$$x = 0.b_0 b_1 b_2 \cdots$$

The dyadic rationals, rationals in the form $r = (2k+1)/2^n$ have two equivalent representations:

$$r = 0.b_0 b_1 b_2 \cdots b_{n-2} 1000 \cdots$$

and

$$r = 0.b_0 b_1 b_2 \cdots b_{n-2} 0111 \cdots$$

These are equivalent as real numbers; they are inequivalent as elements of the Cantor space 2^ω. The term "gap" flies to mind when one writes

$$c = \sum_{n=0}^{\infty} \frac{b_n}{3^{n+1}}$$

the collection of such real numbers c being the conventional Cantor set as it appears in the reals. It obviously has "gaps". These are the objects of study here. They have a natural interpretation in terms of the infinite binary tree: a gap occurs whenever one makes a left-right decision in traveling down the tree. The left sub-tree is distinct from the right, and the gap is what lies in between. The word "gap" is also interesting, as it suggests that the real numbers can be constructed from the Cantor set by "gluing together the gaps": a nicely intuitive idea.

This idea fits nicely with that of Lebesgue dimension. The Cantor space is zero-dimensional, in that every point in the Cantor space is covered by a single open set (a cylinder set), and every refinement of cylinder set is such that the point is contained in one set. This is because, in the product topology, all open sets are either disjoint, or are subsets of one-another. The real number line has Lebesgue covering dimension of one in the natural topology: refinements of open covers chain together. The double-membership, the double-expansion of the dyadic rationals can be understood as the "root source" of why the real numbers have to be one dimensional.

2.1 Gaps

It is tempting to assume that this general structure persists for the beta transform. In fact, nothing could be farther from the truth. This was demonstrated by N. Sidorov and other authors[28, 5]; we recap and expand upon those findings here.

The β-shift given in eqn 1 generates a sequence of bits k_n defined in 2 which can be summed as 3. The summation provides a real number x for $0 \le x \le 1$, the β-expansion of x. The question arises: is this expansion unique, or are there other sequences k_n that are equivalent? The surprising answer is that, for almost all x, there are in fact an *uncountable* number of equivalent expansions! This is in stark contrast to the binary case, where almost all x had a unique expansion, and a countable number of them (the dyadic rationals) had two expansions (but no more). As before, we will call a set of equivalent expansions a "gap".

The first important result is that the gap is an infinite binary tree. Starting with the sequence k_n, there is a countable sequence of locations $n_1, n_2, \cdots$ at which an alternate expansion can be made. Precisely, define

$$m = \left\lfloor \log_\beta \frac{\beta - 1}{2 - \beta} \right\rfloor + 1$$

where $\lfloor \cdot \rfloor$ denotes the floor (the integer part of a number). Let c_m be the bit-sequence $c_m = 100 \cdots 0$ starting with 1, followed by at least m zeros. Sidorov shows, by appeal to the Poincaré recurrence theorem, that this bit-sequence occurs in an infinite number of locations in k_n. Let $n_1, n_2, \cdots$ be the locations of the first bit of c_m (the location of the 1-bit). Picking one of these locations, say n_j, write a new sequence k'_n by setting

$$k'_n = \begin{cases} k_n & \text{for } n < n_j \\ 0 & \text{for } n = n_j \\ \kappa_n & \text{for } n_j < n \end{cases} \tag{7}$$

where κ_n is the bit-sequence obtained by resuming the expansion 2 where it has left off. That is,

$$\kappa_{n_j+1+p} = k_p \left(\beta T_\beta^{n_j}(x) \right)$$

for integer $0 \leq p$. This new sequence k'_n does not evade the theorem; it too has an infinite sequence of the cylinders c_m embedded in it. Thus one obtains an infinite binary tree: at location n_1, one gets to make a decision: should the expansion proceed with k_{n_1+1} being a zero, or a one? Should one go low, or go high? Go left, or go right? This is the binary decision tree; the gap is an infinite binary tree.

A visualization of the gaps is shown in figure 1.

2.2 Unique β-expansions

Glenndinning and Sidorov consider the conjugate problem: are there any points in the unit interval for which the β-expansion is unique? The answer is that there are none, when $1 < \beta \leq \varphi$ where $\varphi = \left(1 + \sqrt{5}\right)/2$ the Golden ratio; there are two when $\varphi < \beta \leq \beta_f$ where $\beta_f = 1.7548776\cdots$ is the root of $\beta^3 = 2\beta^2 - \beta + 1$, a countable number when $\beta_f < \beta < \beta_c$ and an uncountable number 2^ω (the cardinality of the Cantor set) when $\beta_c \leq \beta < 2$.[28] The last accumulate towards $\beta = 2$.

The number β_c is very interesting; lets call it the Komornik–Loreti number. It is a transcendental number built from the Thue–Morse sequence, having the approximate value of $\beta_c = 1.787231650\cdots$ Above this number, there is only one unique β-expansion of $x = 1$. Although it plays this important role in the theory of β-expansions, it does not seem to appear in any of the constructions in this text.

Figure 1: Gap Visualization

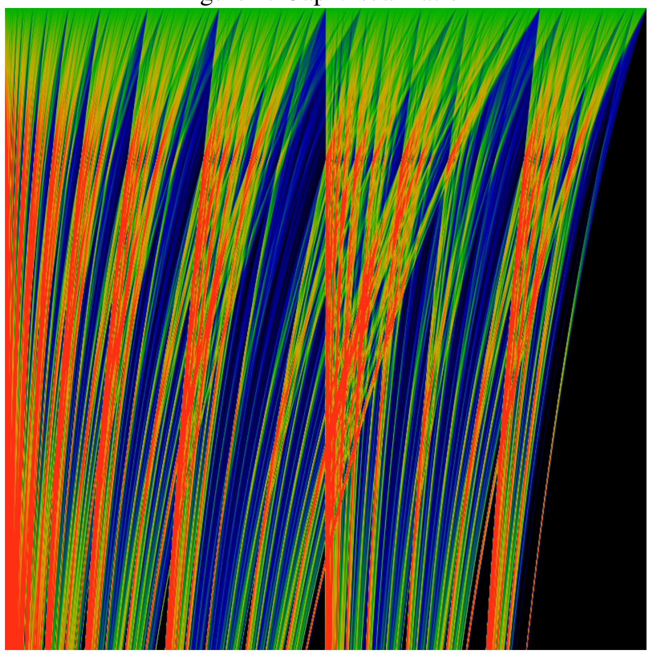

This visualizes the multitude of β-expansions available for $\beta = 1.5$ and $0 \leq x \leq 1$ from left to right. A parameter α runs from $\alpha = \beta/2$ at the top to $\alpha = 2$ at the bottom. For each value of x, an average of 512 distinct expansions k'_n were computed, and then for each such expansion, the sum $w = \frac{1}{2}\sum_{n=0}^{\infty} k'_n \alpha^{-n}$ was computed. The sum is normalized so that $0 \leq w \leq 1$ for all possible expansions. A density visualization is created by incrementing the pixel located at (w, α). The color coding is such that pixels with an average density of 1.0 are green, denser areas tend to yellow and read, while less dense areas tend to blue. Black areas have no hits. Thus, the top row is green, as for the top row, $\alpha = \beta$ and so all expansions give the same value $x = w$, and so the distribution is uniform. As α increases, a separation of the expansions can be seen. This can be understood as follows. Consider the expansions of $x = 0.5000000001$, the first few of which are

1000000000000000000000000000000000000000000000000
0101000000100100101000000000010000000100001000000 1010
0101000000100011100000001001000000100001001000100 1000
0101000000011100000000100100100100100000010010000 0
0101000000011010100100000000100100010000001001001 0000

The first expansion will clearly yeild $w = 1/2$, independent of α. This accounts for the obvious vertical feature at $x = 1/2$. The remaining tracks will bend to the left as α increases, top to bottom. The general slope of each track is dominated by the first non-zero bit in the expansion; this accounts for thier intersections. Figures for other values of β follow similar trends, with more extreme bunching for small β, and a more vertical and uniform arrangement for larger β.

2.3 The tree of expansions

The bit sequence k_n generated in 2 is unique; one may wonder, where are all those other expansions hiding? The answer becomes apparent after some thought: they are the result of failing to take a 'mod 1' at certain times during the iteration. That is, the sequence $T_\beta^n(x)$ obeys $0 \le T_\beta^n(x) < \beta/2$ for all $n > 0$. The 2 can generate the desired gaps only if the input sequence breaks out of this bound at the cylinder locations n_j. As there are an uncountable number of such locations for (almost) every x, then perhaps these accumulate to form a non-zero measure, extending beyond $0 \le x < \beta/2$?

This argument can be made more plausible if it can be shown that the spacing between the alternative expansion points tends to a constant. So, for example, if every fifth bit is a location where an alternative expansion can be made, then one might expect that one-fifth of the extended measure to fall outside the bound $0 \le x < \beta/2$. This intuitive idea needs to be firmed up, as the expansion points are not merely points in a linear sequence, but are points in a binary tree. That is, the length of the branches in the tree must be measured, and averaged.

Towards this end, let D denote the infinite binary tree, and let γ denote a specific location on the tree. In other words, γ is an index into the tree. The natural representation for γ is as a finite sequence of left-right moves. In this case, γ can be represented as a finite string of binary digits. Alternately, γ can be represented as a sequence of left-moves, followed by a mirror, and again some left moves, *etc.* so that $\gamma = g^{a_1} r g^{a_2} r \cdots r g^{a_m}$ or, more compactly, $\gamma = (a_1, a_2, \cdots, a_m)$ for integers a_j. These are obviously the binary and the continued-fraction representations for a location on the tree. It is convenient to define $\text{len} : D \to \mathbb{N}$ by setting $\text{len}\gamma$ to the number of left-right moves in γ. That is, len just counts the depth into the tree; equivalently, len γ is the length of the string γ (in the binary representation).

Let $p : D \to \mathbb{N}$ be a function that assigns a natural number to each node on the binary tree. Such a function can be said to be strictly increasing if (and only if) $p(\gamma') > p(\gamma)$ whenever $\text{len}\,\gamma' > \text{len}\,\gamma$. For the present situation, we wish to take $p(\gamma)$ to be the distance to the branch point at γ, i.e. to be the length of the bitstring k'_n where the branch is taken. The average distance between branch points is then

$$l(\beta) = \lim_{n\to\infty} \frac{1}{n+1} \sum_{\gamma:\text{len}\gamma\le n} \frac{p(\gamma L) + p(\gamma R) - 2p(\gamma)}{2^{1+\text{len}\gamma}}$$

where γL and γR are the left and right locations under γ and the sum is limited to trees of depth n. Regrouping terms, let

$$s_n(\beta) = \frac{1}{2^n} \sum_{\gamma:\text{len}\gamma=n} p(\gamma)$$

be the average distance (path-length) down to the n'th level in the tree. The numbering here is such that $n = 0$ corresponds to the first branch point, which is, in general, a few bits into the string. This implies that the average distance between branch-points can then be taken to be

$$\ell_n(\beta) = \frac{s_n(\beta)}{n+1}$$

so that

$$\ell(\beta) = \lim_{n\to\infty} \ell_n(\beta)$$

The simplest route to obtaining some insight into $\ell(\beta)$ is by numerical exploration; for the measure, numerical histogramming is a straight-forward way of getting a glimpse. Towards that end, the first 22 values of $\ell_n(\beta)$ were computed for several hundred distinct values of x at $\beta = 1.5$ and $\beta = 1.7$. In both cases, it seems that the value of the $\ell_n(\beta)$ does have an x dependence, but it appears to be small and bounded, so that it's more or less flat. The numerical data hint at some fractal behavior, but it is small and hard to discern. A sample distribution at $n = 12$ and $\beta = 1.66$ is shown in figure 2. The numerics hint that $\ell_n(\beta)$ converges uniformly to a relatively flat distribution (as a function of x). Thus, for the remainder, define

$$\hat{\ell}_n(\beta) = \int_0^1 \ell_n(\beta)\,dx$$

This appears to be well-defined, as the results from Sidorov state that there are distinct β-expansions for almost every $0 < x < 1$, and certainly the numerics show nothing amiss. This definition will soon be replaced by a more appropriate one, in the next section, which integrates over an appropriate invariant measure. However, as that measure has not yet been specified, the above will serve as a suitable interim definition. It's a "reasonable" definition of the average distance between branch points.

To understand the β and n dependence, the first 16 values of $\hat{\ell}_n(\beta)$ were computed for the entire interval $1 < \beta < 2$. These settle down very quickly: the first decimal place is determined by $n = 1$, the second decimal place by $n = 5$ and the third decimal place by $n = 10$. Explicit values are shown in figure 3. Due to the rapid convergence, this figure can be taken as the numeric estimate for $\hat{\ell}(\beta)$. The rapid convergence can be intuitively explained by the spectrum of transfer operator of the beta shift. The decaying eigenvalues λ are located near the circle $\lambda = 1/\beta$, they describe how quickly transients die away. That is, one expects that

$$\hat{\ell}_n(\beta) = \hat{\ell}(\beta) + \mathscr{O}\left(\frac{1}{\beta^n}\right)$$

describes the rate of convergence.

The figure 3 shows prominent discontinuities at larger β values. These occur at the locations where the length of the branching cylinder changes. The first one is at $\beta = \varphi = 1.618\cdots$ and the remaining ones occur at solutions to $\beta^m(2-\beta) = \beta - 1$. The approximate sizes of the jumps are shown in the table below. The numerics hint that each jump is twice as large as the jump before, in the limit of large m.

Δm	β	$\Delta\ell(\beta)$
3-2	$1.618\cdots$	1.1
4-3	$1.7548776663\cdots$	3.4
5-4	$1.8667603992\cdots$	9
6-5	$1.9331849819\cdots$	21
7-6	$1.9671682128\cdots$	50
8-7	$1.9838613961\cdots$	107

Figure 2: Track length distribution

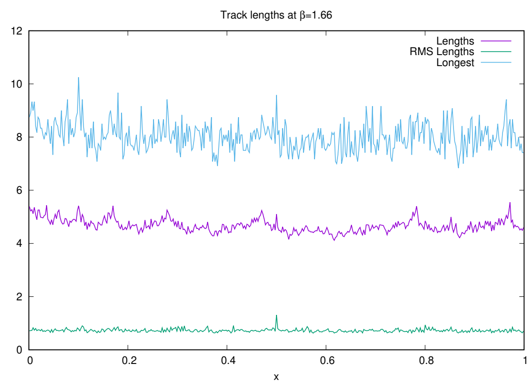

This shows the average distance between branch-points $\ell_n(\beta)$ at $\beta = 1.66$ and $n = 12$ as a function of the starting point x of the sequence. A total of three samples were taken for each of 403 locations on the interval $0 < x < 1$; thus $3 \times 403 = 1209$ distinct points were examined. Each of these points were expanded to $2^{12} = 4096$ distinct paths γ. Thus, each graphed datapoint is the average of these 4096 paths, times three. There does seem to be some kind of fractal-like variation of ℓ over the interval, but no such structure appears in the root-mean-square deviation of the lengths. The line marked "longest" was the longest observed distance, over the 4096 paths, times three. It seems bounded and featureless. The average over the entire interval was $\hat{\ell} = 4.673$, and rms $\hat{\ell} = 1.310$.

Figure 3: Average distance between branch points

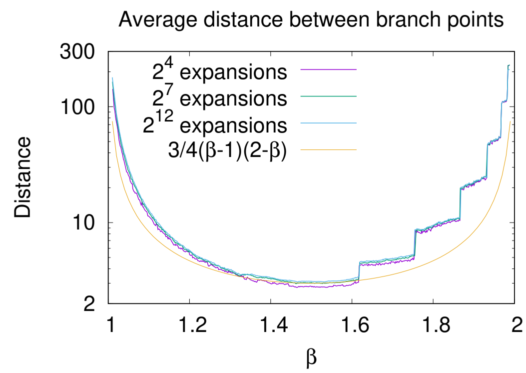

This shows the numerically estimated distance $\hat{\ell}_n(\beta)$ between branch points, as a function of β. Four curves are shown: one at $n=4$, one $n=7$, one at $n=12$ and, behaving as an apparent lower bound, the curve $3/4(\beta-1)(2-\beta)$; the lower bound broken only by the $n=4$ curve. The $\hat{\ell}_n(\beta)$ are estimated from the $\ell_n(\beta)$ by sampling 200 random points in the interval $0<x<1$. It would appear that $\ell_n(\beta)$ is fairly well-converged even at these low values of n. Considerations in a later section suggests a sharper lower bound of $\varphi(\beta-1)(2-\beta)/2$ holds in the limit, although it would be broken by the graphs here. Taking the difference $\hat{\ell}(\beta)-\varphi(\beta-1)(2-\beta)/2$ suggests that the steps on the right are flat (level), following a geometric progression of approximately $(2\varphi)^{m-1}$. The numerics hints that $\hat{\ell}$ does have a minute jittery, fractal shape; it seems that some of the jitter in the graph is inherent, and not numeric artifacts.

The distribution of the distances seems to be well-behaved. Define the variation as

$$\text{var } s_n(\beta) = \frac{1}{2^n} \sum_{\gamma:\text{len}\gamma=n} p^2(\gamma)$$

so that the root-mean-square variation becomes

$$\text{rms } \hat{\ell}(\beta) = \lim_{n\to\infty} \sqrt{\frac{\int_0^1 \text{var} s_n(\beta)\,dx - \hat{\ell}_n^2(\beta)}{n}}$$

The numeric evidence suggests that the limit exists, and that $\hat{\ell}(\beta) = \text{rms}\hat{\ell}(\beta)$ holds exactly. Likewise, define the longest distance between two branch-points at a given level as

$$\max s_n(\beta) = \max_{\gamma:\text{len}\gamma=n} p(\gamma)$$

and similarly, the same quantity averaged over all starting points $0 < x < 1$ in the unit interval:

$$\max \hat{\ell}(\beta) = \lim_{n\to\infty} \frac{\int_0^1 \max s_n dx}{n}$$

The numeric evidence suggests that this limit exists and that $\max \hat{\ell}(\beta) = ((2-\beta)/(1+\beta))\,\hat{\ell}(\beta)$ holds exactly.

2.4 Orbits, subshifts and measures

The branch-points can be thought of as being “dense in the tree” if $\ell(\beta) < \infty$. Intuitively, having dense branch-points suggests that the corresponding map should have a positive measure on the real number line. In this case, one then expects the out-of-bounds measure to be roughly proportional to $1/\ell(\beta)$. This can be made precise.

Each alternative expansion arises from an alternative orbit. Let γ specify that alternative; that is, γ is a specific path down the binary tree. It is a sequence of left-right moves down the tree. Denote these moves as $\gamma = (\gamma_0, \gamma_1, \gamma_2, \cdots)$ so that $\gamma_k = 0$ for a left-move and $\gamma_k = 1$ for a right move. Correspondingly, let $(p_0, p_1, p_2, \cdots)$ be locations of the alternative choices in the β-expansion along that particular path. Given a walk γ down the infinite tree, the alternative orbit for x is then given by

$$\tau_n(x;\gamma) = \begin{cases} T_\beta^n(x) & \text{for } n \le p_0 \\ (1-\gamma_0)\beta T_\beta^{p_0}(x) + \gamma_0 T_\beta^{p_0+1}(x) & \text{for } n = p_0+1 \\ T_\beta^{n-p_0-1}\left(\tau_{p_0+1}(x)\right) & \text{for } p_0+1 < n \le p_1 \\ \cdots & \cdots \\ (1-\gamma_k)\beta\tau_{n-1}(x) + \gamma_k T_\beta(\tau_{n-1}(x)) & \text{for } n = p_k+1 \\ T_\beta^{n-p_k-1}\left(\tau_{p_k+1}(x)\right) & \text{for } p_k+1 < n \le p_{k+1} \\ \cdots & \cdots \end{cases} \tag{8}$$

This orbit has a corresponding bit string expansion

$$\kappa_n(x;\gamma) = \begin{cases} 0 & \text{if } 0 \le \tau_n(x) < \frac{1}{2} \text{ and n} \neq p_k \\ 1 & \text{if } \frac{1}{2} \le \tau_n(x) \le 1 \text{ and n} \neq p_k \\ \gamma_k & \text{if } n = p_k \end{cases} \tag{9}$$

This bit-string is exactly the same as that of eqn 7, except that it uses a more precise notation for the path γ. By construction, one has that each of these bit-strings provide an inequivalent β-expansion for the same x. That is, one has, for all γ (and almost all x and β) that

$$x = \frac{1}{2}\sum_{n=0}^{\infty} \frac{\kappa_n(x;\gamma)}{\beta^n} \tag{10}$$

by construction.

Note that the orbits $\tau_n(x;\gamma)$ are no longer confined to the interval $[0,\beta/2]$, but can briefly venture outside of it. For any fixed path γ, there are a countable number of such locations. However, there are an uncountable number of paths γ down the tree, and so one might reasonably expect that these orbits accumulate on the real number line with a non-zero measure.

This argument can be made slightly more formal. Let S consists of all of the strings generated by the β-shift, that is, the set of all strings given by eqn 2. Let $\mathscr{S}$ be the corresponding subshift (notice that a calligraphic $\mathscr{S}$ was used for the latter, and an italic S for the former). Letting $\mathscr{B}$ denote the sigma algebra on the Cantor space $\{0,1\}^{\omega}$, the subshift $\mathscr{S} \subset \mathscr{B}$ inherits the subspace sigma algebra from $\mathscr{B}$. The elements of $\mathscr{B}$, and of $\mathscr{S}$, are the cylinder sets.

Le be the set of all strings given by 7, that is, 9, and construct the subshift $\mathscr{S}'$ likewise. That is a subshift follows in a relatively straight-forward fashion. Let T be the shift operator that simply lops one bit of of a bit-string. One then has that $TS' \subseteq S'$ because each such string is still a β-expansion, via eqn 10. That T is surjective should be fairly obvious: if there was some bit-string κ that could not be written as $T\kappa'$ for some other $\kappa' \in S'$, then the corresponding β-expansion would be neither βx nor βx mod 1, a contradiction. Thus $TS' = S'$. The operator T extends to an operator $\mathscr{T}$ acting on cylinder sets. By the same argument, $\mathscr{T}\mathscr{S}' = \mathscr{S}'$. That $\mathscr{T}$ is an injection follows essentially from the properties of a sigma-algebra: $\mathscr{T}$ preserves unions and intersections, so that if there were two different cylinder sets $\sigma, \sigma' \in \mathscr{S}'$ with $\mathscr{T}\sigma = \mathscr{T}\sigma'$, then $\sigma \backslash \sigma' = \varnothing$. That is, $\mathscr{T}^{-1} : \mathscr{S}' \to \mathscr{S}'$ is well-defined.

One may now consider some space of functions $\mathscr{F}$ with elements $f \in \mathscr{F}$ are maps $f : \mathscr{S}' \to \mathbb{R}$. Then $\mathscr{T}^{-1}$ acts as a pushforward, defining a linear operator $\mathscr{L}_T : \mathscr{F} \to \mathscr{F}$ by composition $\mathscr{L}_T f = f \circ \mathscr{T}^{-1}$. As a linear operator, this has a spectrum. As it corresponds to a shift, the Frobenius-Perron theorem applies. The largest eigenvalue is some invariant measure $\mu_\beta : \mathscr{S}' \to \mathbb{R}$ with $\mathscr{L}_T \mu_\beta = \mu_\beta$. This measure can be taken to be the restriction of the Bernoulli measure μ to the sub-algebra $\mathscr{S}'$. That is, μ_β does not actually depend on β, only the sub-algebra $\mathscr{S}' = \mathscr{S}'_\beta$ depends on β. This holds likewise for the Renyi–Gelfond–Parry measure: it is just the usual Bernoulli measure restricted to $\mathscr{S}$. It looks unusual when mapped to the unit interval, but that is entirely because the mapping of S to the unit interval is “unusual”. I believe that this is effectively an

Figure 4: Branching Measures

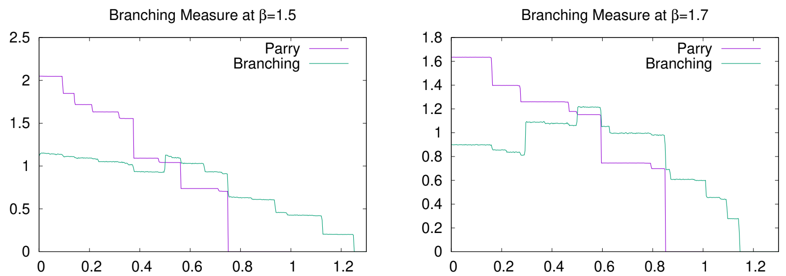

These figures show the Gelfond–Parry measure in purple, and the branching measure, obtained by bin-counting, in green, for two example choices of β. A total of $2^{10} = 1024$ distinct expansions were used to generate these images. Sampling noise was minimized by choosing 201500 distinct random initial values for x. The resulting trajectories $\tau_n(x)$ were histogrammed into 403 bins. The jitteriness in the graphs is purely numerical sampling noise; the actual distribution consists of nearly-flat segments. Note that $\tau_n(x)$ wanders into regions greater than one: without the "mod 1" constraint, it is free to do so. In order to obtain 1024 distinct expansions for each point, one must keep well more than $s_{10}(\beta)$ bits of precision. In this case, 85 bits of precision were used.

---

application of Ornstein theory to the β-transform. Ornstein theory basically just says that there is a large range of disparate phenomena, but that the underlying dynamics is that of the Bernoulli process; the phenomena look different only because they are obscured by different mappings from the Bernoulli process to the system in question.

All this is a round-about way of saying that we really want to do is to visualize the orbits of eqn 8, and to put this visualization on a firm footing. For every $0 \leq x \leq 1$ and path γ there is a bit-string κ given by eqn 9 and an orbit 8; these are in one-to-one correspondence between bitstrings and orbits. The observation that the underlying measure $\mu : \mathscr{S}' \to \mathbb{R}$ is just the Bernoulli measure implies that each and every point can be given an equal weighting. For numerical purposes, this states that the histogramming of the orbits of $\tau_n(x)$, giving each one equal weight as it is assigned to a bin, is an accurate technique to approximate to the true measure μ that arises from the one-to-one mapping of 9 to 8.

With that out of the way, visualizations of the measure can now be given. It's notationally convenient to re-insert the subscript β as the mapping itself does depend on β; thus write $\mu_\beta(x)$ for the mapped measure. This $\mu_\beta(x)$ is shown in figure 4 for a few select values of β. A visualization of the measure over the entire range $1 < \beta < 2$ is shown in figure 5.

2.5 Branching Transform and Measure

Rather than working with individual point orbits, some small amount of insight is gained by working with directly with densities on the real number line.

Figure 5: Branching diagram

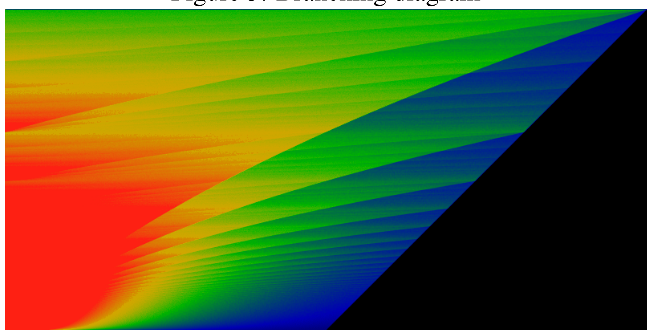

Above, the Parry-Gelfond measure $\nu_\beta(x)$; below the branching measure $\mu_\beta(x)$.

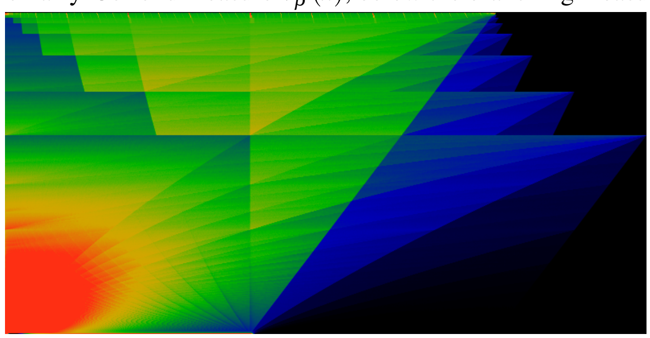

Both rectangles run $1 < \beta < 2$ bottom to top; the upper rectangle runs $0 \le x \le 1$ left to right; the lower runs $0 \le x \le (1+\varphi)/2$ left to right. Red indicated regions where $\mu_\beta(x) > 2$, green indicates regions where $\mu_\beta(x) \approx 1$ and blue shows regions with $\frac{1}{2} > \mu_\beta(x)$. The measure is vanishing in black regions. The top bar is green, as there, the measure is supported only on $0 \le x \le 1$ and is uniform on that interval. The visible speckling is due to the finite number of samples taken in generating this figure; increased sampling removes the speckles. The tiny red tick-marks at very top are numerical effects, due to $\ell(\beta)$ diverging as β approaches 2. The prominent plateaus are where size of the cylinder changes (the Δm jumps). The measures shown in the previous figure are a pair of horizontal slices through this figure. Perhaps less than obvious due to the color scheme, the rightmost-edge of the bottom-most plateau is a straight line, from $\beta = \varphi$, $x = (1+\varphi)/2$ at the plateau tip, to $\beta = 1$, $x = 1$ at the very bottom; the density fades to black making this hard to see. Most of the rays in both graphs are not actually straight lines; they are all curving arcs (that are sections of polynomials).

Starting with any number $0 \leq x \leq 1$, the first digit in it's k_n-expansion expansion will be a one if $\frac{1}{2} \leq x \leq 1$. A string of m zeros follow if $0 \leq T_\beta^j(x) < \frac{1}{2}$ for each $0 < j \leq m$. The second condition is satisfied if $0 \leq \beta^m \left(x - \frac{1}{2}\right) < \frac{1}{2}$; that is

$$\frac{1}{2} \leq x < \frac{1}{2}\left(1+\beta^{-m}\right)$$

Write the upper bound as $\alpha = \left(1+\beta^{-m}\right)/2$. The corresponding branching shift can be written as

$$\tau_\beta(x;G) = \begin{cases} (1-G)\,\beta x + G T_\beta(x) & \text{for } \frac{1}{2} \leq x < \alpha \\ T_\beta(x) & \text{otherwise} \end{cases}$$

where G is a random Bernoulli process. Iterating this generates the very same sequence as in eqn 8, where $\gamma \in G$ is understood to be one specific draw from the process. Rearranging terms, this simplifies to

$$\tau_\beta(x;G) = T_\beta(x) + (1-G)\,\frac{\beta}{2}\Theta\left(x-\frac{1}{2}\right)\Theta(\alpha - x)$$

where Θ is the Heaviside step function, indicating that the alternate term is valid only for $1/2 \leq x < \alpha$. Thus, for the greedy process $G = 1$, the original β-shift is regained.

The corresponding bit-sequence can then be written as

$$\begin{aligned} \kappa_n(x;G) &= \Theta\left(\tau^n(x;G) - \frac{1}{2}\right)\left(1 - \Delta\left(\tau^n(x;G)\right)\right) + G\Delta\left(\tau^n(x;G)\right) \\ &= \Theta\left(\tau^n(x;G) - \frac{1}{2}\right) + (G-1)\,\Delta\left(\tau^n(x;G)\right) \end{aligned}$$

where $\Delta(x) = \Theta(x-1/2)\,\Theta(\alpha - x)$ is the cylinder. This gives exactly the same bit-sequence as eqn 9.

For the push-forward, write $\mathscr{L}f = f \circ T^{-1}$ and $\mathscr{M}f = f \circ \tau^{-1}$. The first transfer operator is straight-forward; it is

$$[\mathscr{L}f](y) = \frac{1}{\beta}\left[f\left(\frac{y}{\beta}\right) + f\left(\frac{y}{\beta}+\frac{1}{2}\right)\right]\Theta\left(\frac{\beta}{2} - y\right)\Theta(y)$$

with Θ the Heaviside step function, indicating that the expression is valid only for $0 \leq y \leq \beta/2$ and is zero otherwise.

The branching shift has three pre-images, defined on the intervals $[0,1/2]$, $[1/2,\alpha]$ and $[\alpha,\infty]$. The corresponding pushfoward on these three consists of the maps

$$\mathscr{M}_{[0,1/2]}f : y \mapsto \begin{cases} \frac{1}{\beta} f\left(\frac{y}{\beta}\right) & \text{for } 0 \leq y \leq \beta/2 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathscr{M}_{[1/2,\alpha]}f : y \mapsto \begin{cases} \frac{1}{\beta}(1-G)\,f\left(\frac{y}{\beta}\right) & \text{for } \beta/2 \leq y \leq \beta\alpha \\ \frac{1}{\beta} G f\left(\frac{y}{\beta}+\frac{1}{2}\right) & \text{for } 0 \leq y \leq 1/2\beta^{m-1} \end{cases}$$

and

$$\mathscr{M}_{[\alpha,1]} f : y \mapsto \begin{cases} \frac{1}{\beta} f\left(\frac{y}{\beta} + \frac{1}{2}\right) & \text{for } 1/2\beta^{m-1} \leq y \\ 0 & \text{otherwise} \end{cases}$$

Combining these together, $\mathscr{M} = \mathscr{M}_{[0,1/2]} + \mathscr{M}_{[1/2,\alpha]} + \mathscr{M}_{[\alpha,1]}$, which, after re-arrangement, can be written in the form

$$[\mathscr{M} f](y) = [\mathscr{L} f](y) + \frac{1-G}{\beta} \Theta\left(y - \frac{\beta}{2}\right) \left[f\left(\frac{y}{\beta}\right) \Theta(\beta\alpha - y) - f\left(\frac{y}{\beta} + \frac{1}{2}\right)\right]$$

Thus, the usual β-transfer operator is obtained when the binary process is greedy, i.e. when $G = 1$.

There does not appear to be any easy way of writing the invariant measure. It appears to be a linear combination of three terms:

$$A(y) = \frac{\beta}{2} \sum_{n=0}^{\infty} \frac{\Theta(\tau^n(\beta/2) - y)}{\beta^n}$$

and

$$B(y) = 2(\beta-1)(1-\alpha) \sum_{n=0}^{\infty} \frac{\Theta(\tau^n(\alpha\beta) - y)}{\beta^n}$$

and

$$C(y) = -(\beta-1) \sum_{n=0}^{\infty} \frac{\Theta(\tau^n(\alpha\beta) - \tau^n(\beta/2) - y)}{\beta^n}$$

but the precise coefficients are hard to grasp. Numerical work indicates the measure is approximately $\mu \approx A + B + C$ (up to overall normalization) to within about 20% over the entire range of x and β. Assuming that these are indeed the only components contributing to the measure, that is, assuming $\mu = aA + bB + cC$, then it appears that a, b and c are discontinuous and fractal functions of β. This is illustrated in figure 6.

A reasonably interesting task would be to characterize the self-similarities evident in this figure – or even the simpler case: to characterize the self-similarities for the plain Gelfond-Parry normalization factor.

2.6 Invariant distance

The distance ℓ between branching points is directly related to likelihood of encountering a one followed by m zero's in a given string. In a previous section, an average value of ℓ was defined, but this was not quite correct, as the average was taken over a measure that was not shift-invariant. The branching measure μ_β is the appropriate measure for this average. This follows, as the set S' contained only the β-expansion strings, and the algebra of cylinders $\mathscr{S}'$ inherited a shift-invariant weighting from the Bernoulli measure. Thus, each string in S' can now be counted uniformly, which allows averages of properties of strings in S' to be properly taken. This argument applies to any quantity that may be assigned to individual points in the set S': the measure μ_β provides the correct manner in which to weight and count the contribution to an average.

Figure 6: Step Function Sums

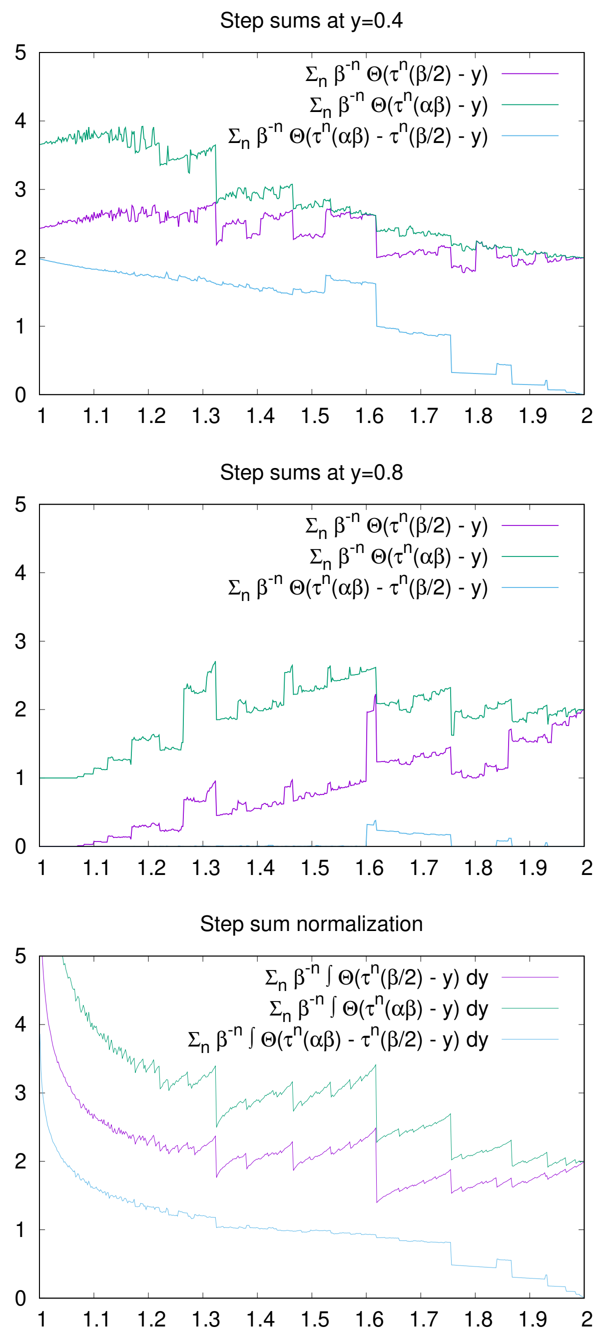

These figures show the β-sums over the iterated midpoint and upper-bound, and the normalization integral. All sums are as labeled; β runs along the horizontal axis.

The argument can also be rephrased in the language of physics. The distance $\ell(\beta) = \lim_{n\to\infty} \ell_n(\beta)$ has an a priori x-dependence, and so, when considering the distance between branch points in a thermalized orbit, one must use the corresponding thermalized measure for sampling. The μ_β provides that thermalized measure.

Distances are accessible via some simple considerations. The corresponding cylinders are to be sampled out of the measure μ_β. That is, the probability $p(\beta)$ of encountering such an x corresponds to

$$p(\beta) = \int_{1/2}^{\alpha} d\mu_\beta$$

This gives the expectation value for the distance between such cylinders:

$$\langle \ell \rangle = \int \ell(\beta)\, d\mu_\beta = \frac{1}{p(\beta)}$$

where the integral is understood to run over the entire range. These arguments can be readily verified numerically, and one finds (surprisingly!?) that $\langle \ell \rangle = \hat{\ell}$ to within numerical uncertainty. That is, the β-dependence is exactly as that shown in figure 3. The numerics do suggest a slightly tighter bound it appears that $\varphi(\beta-1)(2-\beta)/2 \leq \langle \ell \rangle$ for all β. This bound is the most strict near $\beta \approx 1.5$; everywhere else it is easily met.

3 Conclusion

Beta expansions of real numbers offer a counterpoint to more conventional p-adic analysis and algebraic number theory. Compared to the vast amount of literature devoted to these topics, there is only a thimble-full of accumulated knowledge about β-expansions. Picking almost any topic in algebraic number theory, there are no known comparable statements for β-expansions. This is not because (this does not seem to be due to) the sterility of the β-expansion, where no interesting statements can be made, or that it reduces to some trivial situation, where all statements are obvious. On the contrary, there seems to be a severe shortage of useful algebraic tools and general theorems to be applied; for the most part, insight requires numerical exploration to confirm even the most basic ideas.

To put it differently: taking a real number β and multiplying it by another real number x seems about as basic an operation as one could imagine. It feels like the mod 1 aspect should be incidental; this is about understanding multiplication. From the point of view of an iterated dynamical system, the β-transform is oodles simpler than the famous logistic map or the Mandelbrot set: it is, after all, linear. Despite being simpler, it seems equally daunting. Getting a better understanding of it seems like a pre-requisite: the more complex dynamical-systems maps can be thought of as being kind-of-like algebraic varieties built from the β-expansion. Yet, no cornerstone has been laid for an algebraic geometry of beta expansions. The aim of this text was to move a few pebbles around.

To recap: different beta expansions for a given real number are organized into an infinite binary tree. Binary trees are inherently interesting, due to their intimate relationship with elliptic curves (the infinite binary tree embeds into the modular group $SL(2,\mathbb{Z})$ and thence into the hyperbolic plane). The real numbers already have a natural binary-tree structure via the 2-adic binary expansion (the Cantor set). Imagining two trees, one for x and one for β, the β-expansion provides a third tree that sits askew these two. It seems reasonable to expect a rich structure to emerge from this entanglement.

References

[1] A. Rényi, "Representations for real numbers and their ergodic properties", *Acta Math Acad Sci Hungary*, 8, 1957, pp. 477–493.

[2] A.O. Gel'fond, "A common property of number systems", *Izvestiya Akad Nauk SSSR Seriya Matematicheskaya*, 23, 1959, pp. 809–814, URL http://www.mathnet.ru/php/archive.phtml?wshow=paper&jrnid=im&paperid=3814&option_lang=eng.

[3] W. Parry, "On the β-expansion of real numbers", *Acta Math Acad Sci Hungary*, 11, 1960, pp. 401–416.

[4] Linas Vepstas, "Spectrum of the Beta Transformation", , 2020, URL https://www.linas.org/math/beta-spectrum.pdf, pre-print on personal website; submitted for publication.

[5] Nikita Sidorov, "Almost every number has a continuum of β-expansions.", *The American Mathematical Monthly*, 110, 2003, pp. 838–842, URL http://www.maths.manchester.ac.uk/nikita/amm.pdf.

[6] Martijn de Vries and Vilmos Komornik, "Unique Expansions of Real Numbers", *ArXiv*, arXiv:math/0609708, 2006, URL https://www.esi.ac.at/static/esiprpr/esi1810.pdf.

[7] Karma Dajani, et al., "The natural extension of the beta-transformation", *Acta Math Hungary*, 73, 1996, pp. 97–109, URL https://www.researchgate.net/publication/2257842.

[8] Vaughn Climenhaga and Daniel J. Thompson, "Intrinsic ergodicity beyond specification: beta-shifts, S-gap shifts, and their factors", *Israel Journal of Mathematics*, 2010, URL https://arxiv.org/abs/1011.2780.

[9] P. Erdös and V. Komornik, "Developments in non-integer bases", *Acta Math Hungar*, 79, 1998, pp. 57–83.

[10] Nikita Sidorov, "Universal β-expansions", *Arxiv*, 2002, URL https://arxiv.org/abs/math/0209247v1.

[11] Boris Adamczewski, et al., "Rational numbers with purely periodic β-expansion", *Bull London Math Soc*, 42, 2010, pp. 538–552, URL http://adamczewski.perso.math.cnrs.fr/AFSS_BLMS.pdf.

[12] K. Schmidt, “On periodic expansions of Pisot numbers and Salem numbers”, *Bull London Math Soc*, 12, 1980, pp. 269–278.

[13] F. Blanchard, “Beta-expansions and Symbolic Dynamics”, *Tehoretical Comp Sci*, 65, 1989, pp. 131–141.

[14] Bruno Henrique Prazeres de Melo e Maia, “An equivalent system for studying periodic points of the beta-transformation for a Pisot or a Salem number”, , 2007, URL http://repositorio.ual.pt/bitstream/11144/471/1/bmaia_thesis.pdf.

[15] Shigeki Akiyama, “Finiteness and Periodicity of Beta Expansions - Number Theoretical and Dynamical Open Problems”, , 2009, URL http://math.tsukuba.ac.jp/~akiyama/papers/proc/BetaFiniteCIRM.pdf.

[16] Louis-Sébastien Guimond, et al., “Arithmetics on beta-expansions”, *Acta Arithmetica*, 112, 2001, pp. 23–40, URL https://www.researchgate.net/profile/Edita_Pelantova/publication/259299735_Arithmetics_of_beta-expansions/links/5434e32a0cf294006f736e7c/Arithmetics-of-beta-expansions.pdf.

[17] Bernat Julien, “Arithmetics in β-numeration”, *Discrete Mathematics and Theoretical Computer Science*, 9, 2006, URL http://www.iecl.univ-lorraine.fr/~Julien.Bernat/arithbetanum.pdf.

[18] M. Hbaib and Y. Laabidi, “Arithmetics in the set of beta-polynomials”, *Int J Open Problems Compt Math*, 6, 2013, URL http://www.i-csrs.org/Volumes/ijopcm/vol.6/vol.6.3.1.pdf.

[19] W. P. Thurston, “Groups, tilings and finite state automata”, *AMS Colloquium Lectures*, 1989.

[20] Valérie Berthé and Anne Siegel, “Tilings associated with beta-numeration and substitutions”, *Integers: Electronc Journal of Combinatorial Number Theory*, 5, 2005.

[21] Sh. Ito and H. Rao, “Purely periodic β-expansions with Pisot unit base”, *Proc Amer Math Soc*, 133, 2005, pp. 953–964.

[22] Shigeki Akiyama, “Beta expansion and self-similar tilings”, , 2017, URL http://cloud.crm2.uhp-nancy.fr/pdf/Manila2017/Akiyama.pdf.

[23] Valérie Berthé and Anne Siegel, “Purely periodic β-expansions in the Pisot non-unit case”, *ArXiv*, arXiv:math/0407282, 2004, URL https://hal.archives-ouvertes.fr/hal-00002208/document.

[24] Jakob Grue Simonsen, “Beta-Shifts, their Languages, and Computability”, *Theory of Computing Systems*, 48, 2011, pp. 297–318, URL http://www.diku.dk/~simonsen/papers/j12.pdf.

[25] Leopold Flatto, et al., "The Zeta Function of the Beta Transformation", *Ergodic Theory and Dynamical Systems*, 14, 1994, pp. 237–266.

[26] Daniel J. Thompson, "Irregular sets and conditional variational principles in dynamical systems", , 2010, URL https://people.math.osu.edu/thompson.2455/thesis_thompson.pdf.

[27] L. Barreira and B. Saussol, "Variational Principles and Mised Multifractal Spectra", *Transactions of the American Mathematical Society*, 353, 2001, pp. 3919–3944, URL http://www.math.univ-brest.fr/perso/benoit.saussol/art/mixed.pdf.

[28] Paul Glendinning and Nikita Sidorov, "Unique Representations of Real Numbers in Non-Integer Bases.", *Mathematical Research Letters*, 8, 2001, p. 535–543, URL https://www.researchgate.net/publication/243081115_Unique_Representations_of_Real_Numbers_in_Non-Integer_Bases.