TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 349, Number 11, November 1997, Pages 4355–4365
S 0002-9947(97)02069-2

β-EXPANSIONS WITH DELETED DIGITS FOR PISOT NUMBERS β

STEVEN P. LALLEY

Abstract. An algorithm is given for computing the Hausdorff dimension of the set(s) $\Lambda = \Lambda(\beta, D)$ of real numbers with representations $x = \sum_{n=1}^{\infty} d_n \beta^{-n}$, where each $d_n \in D$, a finite set of "digits", and $\beta > 0$ is a Pisot number. The Hausdorff dimension is shown to be $\log \lambda / \log \beta$, where λ is the top eigenvalue of a finite 0-1 matrix A, and a simple algorithm for generating A from the data β, D is given.

1. Introduction

This paper concerns the set(s) $\Lambda = \Lambda(\beta, D)$ of real numbers with representations $x = \sum_{n=1}^{\infty} d_n \beta^{-n}$, where each $d_n \in D$, a finite set of "digits", and $\beta > 0$. These sets have been the subject of several recent studies. Keane, Smorodinsky, and Solomyak [3] considered the special case $D = \{0, 1, 3\}$ and $\beta \in (2.5, 3)$: they showed that although for almost every $\beta \in (2.5, 3)$ the Hausdorff dimension of Λ is 1, there is a sequence β_k of algebraic integers in (2.5, 3) such that the dimension of Λ is less than 1. Pollicott and Simon [5] showed, more generally, that under certain conditions the dimension of Λ equals $\log |D| / \log \beta$ for almost every β in a certain critical interval, but that the set of discontinuities is dense in this interval. These discontinuities are at algebraic numbers β. It is not known if there are discontinuities at transcendental values of β.

The main result of this paper is that the Hausdorff dimension of $\Lambda = \Lambda(\beta, D)$ is computable provided β is a Pisot number β (arbitrary) and D is a finite subset of $\mathbf{Z}[\beta]$. In particular, it will be shown that

$$\dim_H(\Lambda) = \frac{\log \lambda(\beta, D)}{\log \beta}, \tag{1}$$

where $\lambda = \lambda(\beta, D)$ is the largest eigenvalue of a certain 0-1 matrix $A = A(\beta, D)$ (and, therefore, λ is itself an algebraic integer). A simple algorithm for computing the matrix A from the data β, D will be provided. Thus, (1) permits (in theory!) the computation of $\dim_H(\Lambda)$ to any degree of accuracy. In the special case $\beta = 1 + \sqrt{3}$ (the larger root of $x^2 - 2x - 2 = 0$) and $D = \{0, 1, 3\}$, Eq. (1) may be used to obtain the approximation (and rigorous upper bound)

$$\dim_H(\Lambda) \approx 0.971847 \tag{2}$$

Received by the editors June 12, 1995.
1991 *Mathematics Subject Classification.* Primary 11K55, 28A78.
Key words and phrases. Hausdorff dimension, entropy, Pisot number.
Supported by NSF grant DMS-9307855.

Acknowledgment

The author would like to thank Professor Boris Solomyak for pointing out an error in the author's original computation of this number.

2. Finite Approximations to Λ

For the remainder of the paper $\beta > 1$ will be a fixed Pisot number with algebraic conjugates $\beta_2, \cdots, \beta_d$ and minimal polynomial $p(x) = x^d - \sum_{k=0}^{d-1} a_k x^k$, where each a_k is an integer and $d \geq 2$. (Recall that a *Pisot number* is an algebraic integer all of whose algebraic conjugates lie inside the unit disk.) The set D of *digits* will be an arbitrary finite subset of $\mathbf{Z}[\beta]$ (the set of all integer polynomial expressions in β) of cardinality at least 2; the maximum absolute value of an element of D will be denoted by h. Define

$$\Lambda = \{\sum_{n=1}^{\infty} d_n \beta^{-n} \ : \ d_n \in D\};$$

$$\Lambda_m = \{\sum_{n=1}^{m} d_n \beta^{-n} \ : \ d_n \in D\}, \qquad m = 1, 2, \ldots .$$

The finite sets Λ_m should be thought of as discrete approximations to Λ. (It can be shown that as $m \to \infty$, $\Lambda_m \to \Lambda$ in the Hausdorff metric, but this fact will not be needed.) The growth of the sets Λ_m determines the Hausdorff dimension of Λ, as the following propositions show.

Proposition 1. *There exists a constant* $0 < \lambda \leq \beta$ *such that as* $m \to \infty$,

$$|\Lambda_m|^{\frac{1}{m}} \longrightarrow \lambda.$$

Proof. It is easily seen that $|\Lambda_{m+n}| \leq |\Lambda_m|\,|\Lambda_n|$. Consequently, the existence of the limit follows by the fundamental subadditivity lemma. That $\lambda \leq \beta$ follows from below, and that $\lambda > 0$ follows from Corollary 2 in section 5, which shows that λ is the top eigenvalue of a certain 0-1 matrix. □

Proposition 2. $\dim_H(\Lambda) = \log \lambda / \log \beta$.

Lemma 1. *The box dimension* $\dim_B(\Lambda)$ *of* Λ *satisfies*

$$\dim_B(\Lambda) \leq \frac{\log \lambda}{\log \beta}.$$

Proof. This is a direct consequence of Proposition 1. For each $m \geq 1$ let $\mathcal{U}_m$ be the collection of intervals $[x - \kappa\beta^{-m}, x + \kappa\beta^{-m}]$, where $x \in \Lambda_m$ and $\kappa = 2h/(\beta - 1)$. This is a covering of Λ by $|\Lambda_m|$ intervals of radius $\kappa\beta^{-m}$. The advertised inequality for $\dim_B(\Lambda)$ therefore follows from Proposition 1. □

Since a set's box dimension always dominates its Hausdorff dimension, the lemma implies that

$$\dim_H(\Lambda) \leq \frac{\log \lambda}{\log \beta}. \tag{3}$$

Thus, to complete the proof of Proposition 2, it suffices to establish the reverse inequality. This will be accomplished in section 6. It is in this direction that the hypothesis that β is a Pisot number will be used. (Note that neither the existence of the limit in Proposition 1 nor the inequality for the box dimension in Lemma 1 required this hypothesis.) The following lemma is the only part of the argument where the Pisot property is needed.

Lemma 2 (Separation Lemma)**.** *There exists a constant $C = C(\beta, D) > 0$ with the following property. For any $m \geq 1$ and any two* distinct *elements x, y of Λ_m,*

$$|x - y| \geq C\beta^{-m}.$$

Proof. It suffices to consider the case where $D \subset \mathbf{Z}$. This is because any polynomial expression $\sum_{n=1}^{m} d_n\beta^{-n}$ whose coefficients d_n lie in a finite subset D of $\mathbf{Z}[\beta]$ may be rewritten as a polynomial expression $\sum_{n=1}^{m+r} d'_n\beta^{-n}$ whose coefficients d'_n are all elements of a finite set D' of integers (r and D' depend only on D, not on m or the particular choice of $\sum_{n=1}^{m} d_n\beta^{-n}$). Let $x = \sum_{n=1}^{m} d_n\beta^{-n}$ and $y = \sum_{n=1}^{m} d'_n\beta^{-n}$ be distinct elements of Λ_m. Set

$$F(x) = \sum_{n=1}^{m} (d_n - d'_n)x^{m-n};$$

then $F(\beta) \neq 0$, and consequently $F(\beta_j) \neq 0$ for each of the algebraic conjugates β_j of β. Recall that the coefficients $d_n - d'_n$ are uniformly bounded in absolute value by a finite constant $2h$. Since β is a Pisot number, $|\beta_j| < 1$ for each j, and consequently $|F(\beta_j)| \leq 2h/(1-|\beta_j|)$. But the coefficients $d_n - d'_n$ of F are integers, so

$$F(\beta)\prod_{j=2}^{d} F(\beta_j) \in \mathbf{Z} - \{0\}.$$

It follows that $|F(\beta)| \geq (\prod_{j=2}^{d}(1 - |\beta_j|))/(2h)^{d-1}$. □

3. The Associated Digraph

The relation between the finite sets Λ_m and the infinite set Λ may be visualized with the aid of the directed graph $\mathcal{G}$ whose vertex set is $\bigcup_{m=0}^{\infty}(\Lambda_m \times \{m\})$ (with $\Lambda_0 = \{0\}$) and whose directed edges are

$$(\sum_{n=1}^{m} d_n\beta^{-n}, m) \longrightarrow (\sum_{n=1}^{m} d_n\beta^{-n} + d'\beta^{-m-1}, m+1), \quad d' \in D.$$

The second coordinate m of a vertex will be called its *depth*, and the digit d' will be called the *color* of the directed edge (note that d' is independent of the representations chosen for the two vertices in question). The directed edges of $\mathcal{G}$ only connect successive depths $m, m+1$, so every path in $\mathcal{G}$ (a sequence of vertices, each successive pair connected by a directed edge) goes "down". In particular, if $\langle v_n\rangle_{n\geq 1} = \langle (x_n, n)\rangle_{n\geq 1}$ is an infinite path in $\mathcal{G}$ starting from the root vertex (0,0) then there is a sequence of digits d_n such that for each m,

$$x_m = \sum_{n=1}^{m} d_n\beta^{-n};$$

consequently, $\lim_{m\to\infty} x_m = x$ where $x = \sum_{n=1}^{\infty} d_n\beta^{-n}$. We will say that such a path *converges* to x. Thus, the *boundary* $\partial\mathcal{G}$ of $\mathcal{G}$ is identified with the set Λ. (Note, however, that a point $x \in \Lambda$ may have more than one such representation, so there will in general be many paths converging to the same limit point.)

For each vertex v, define the *descendants* of v to be those vertices w for which there is a finite path from v to w. Observe that, for any vertex v, the subgraph of $\mathcal{G}$ consisting of v and its descendants, together with the arrows emanating from them, is a replica of the entire graph G.

For any finite sequence $d_1 d_2 \dots d_m$ of digits (elements of D), let $v(d_1 d_2 \dots d_m)$ denote the vertex $(\sum_{n=1}^{m} d_n \beta^{-n}, m)$. Note that for any given vertex there may be several representations $v(d_1 d_2 \dots d_m)$. For any two vertices $x = v(d_1 d_2 \dots d_m)$ and $y = v(d'_1 d'_2 \dots d'_m)$ at the same depth m, define their $\mathcal{G}$-distance $\rho(x, y)$ by

$$\rho(x, y) = \beta^m \left| \sum_{n=1}^{m} (d_n - d'_n) \beta^{-n} \right|.$$

By the Separation Lemma, the distance ρ between distinct vertices of the same level is bounded below, and the bound is independent of the level. Define the constant

$$\kappa = 2h/(\beta - 1).$$

For any vertex x of $\mathcal{G}$ define its *neighborhood* $\mathcal{N}(x)$ to be the set of vertices y at the same depth such that $\rho(x, y) \le \kappa$.

Lemma 3. *Let $d_1 d_2 \dots$ and $d'_1 d'_2 \dots$ be arbitrary sequences of digits. If*

$$\lim_{m \to \infty} v(d_1 d_2 \dots d_m) = \lim_{m \to \infty} v(d'_1 d'_2 \dots d'_m)$$

then for every $m \ge 1$,

$$v(d'_1 d'_2 \dots d'_m) \in \mathcal{N}(v(d_1 d_2 \dots d_m)).$$

Proof. The hypothesis implies that $\sum_{n=1}^{\infty} d_n \beta^{-n} = \sum_{n=1}^{\infty} d'_n \beta^{-n}$. Since digits are bounded in absolute value by h, it follows that for each m,

$$\left| \sum_{n=1}^{m} d_n \beta^{-n} - \sum_{n=1}^{m} d'_n \beta^{-n} \right| \le 2h\beta^{-m}/(\beta - 1).$$

This is equivalent to the conclusion of the lemma, by the definition of κ and ρ. □

Note. If the elements of D are all nonnegative, then the bound $2h\beta^{-m}/(\beta - 1)$ in the preceding argument could be improved to $h\beta^{-m}/(\beta - 1)$. (See also the proof of Lemma 5 below.) In this case, we could use $\kappa = h/(\beta - 1)$ instead of $\kappa = 2h/(1 - \beta)$. This reduction can make a large difference in the size of the set of neighborhood types.

Say that two vertices x, y (not necessarily at the same depth) have the same *neighborhood type* if there is a bijective mapping between $\mathcal{N}(x)$ and $\mathcal{N}(y)$ that preserves the distance function ρ and identifies the vertices x and y. Let $\mathcal{T}$ be the set of all neighborhood types in $\mathcal{G}$.

Lemma 4. *$\mathcal{T}$ is finite.*

Proof. Recall that for each vertex of $\mathcal{G}$ the set of all descendants constitutes a replica of $\mathcal{G}$. Hence, for any two digit sequences $d_1 d_2 \dots d_m$ and $d'_1 d'_2 \dots d'_k$ there is a distance-preserving inclusion

$$\mathcal{N}(v(d'_1 d'_2 \dots d'_k)) \longrightarrow \mathcal{N}(v(d_1 d_2 \dots d_m d'_1 d'_2 \dots d'_k)).$$

It follows that for any infinite sequence $d_1 d_2 \dots$ there is a chain of distance-preserving inclusions

$$\mathcal{N}(v(d_1)) \to \mathcal{N}(v(d_2 d_1)) \to \cdots \to \mathcal{N}(v(d_m \cdots d_2 d_1)) \to \cdots.$$

By the Separation Lemma, all such chains stabilize, because no neighborhood can contain more than $4h/C(\beta-1)+1$ distinct vertices. It follows by a routine argument that there are only finitely many neighborhood types. $\square$

Lemma 5. *Let $d_1d_2\dots d_{m+1}$ and $d'_1d'_2\dots d'_{m+1}$ be arbitrary sequences of digits of length $m+1$. If*

$$\rho(v(d_1d_2\dots d_{m+1}), v(d'_1d'_2\dots d'_{m+1})) \le \kappa$$

then

$$\rho(v(d_1d_2\dots d_m), v(d'_1d'_2\dots d'_m)) \le \kappa.$$

Proof. If the last inequality were not true then $|\sum_{n=1}^{m}(d_n - d'_n)\beta^{-n}| > \kappa\beta^{-m}$. Because the digits d_n, d'_n are bounded in modulus by h, it would then follow from the triangle inequality that

$$\left|\sum_{n=1}^{m+1}(d_n - d'_n)\beta^{-n}\right| > \kappa\beta^{-m} - 2h\beta^{-m-1} \ge \kappa\beta^{-m-1},$$

by definition of κ. This would contradict the hypothesis. $\square$

Corollary 1. *For any vertices v, v' of $\mathcal{G}$ such that $v \to v'$, the neighborhood type of v' is completely determined by that of v and the color of the directed edge $v \to v'$.*

For the sake of computation it will be necessary to have an algorithm for enumerating the set $\mathcal{T}$ of neighborhood types. For this purpose it is best to think of a neighborhood type as a finite set of real numbers contained in $[-\kappa, \kappa]$ with 0 as an element. (Thus, for a vertex (x, m) of $\mathcal{G}$, the neighborhood type is the set $\{\beta^m(x'-x) : x' \in \mathcal{N}((x,m))\}$. The neighborhood type of the root node (0,0) is the set $\{0\}$.) Define the *offspring* of a neighborhood type τ to be those neighborhood types τ' such that for some directed edge $v \to v'$ of $\mathcal{G}$, v has type τ and v' has type τ'. The offspring of τ may be enumerated without actually searching the graph for vertices of type τ: they are, for $d' \in D$, the finite sets

$$\{\beta x + (d'' - d') : x \in \tau, d'' \in D\} \cap [-\kappa, \kappa].$$

An algorithm for listing the members of $\mathcal{T}$ follows:

begin
  $\mathcal{T} := \{\{0\}\}$;
  $\mathcal{S} := \{\{0\}\}$;
  while $\mathcal{S} \neq \emptyset$ **do**
    begin
      $\mathcal{T} := \mathcal{T} \cup \mathcal{S}$;
      $\mathcal{S} := \{\text{offspring of } \mathcal{S}\}$;
      $\mathcal{S} := \mathcal{S} - \mathcal{T}$;
    end
  return $\mathcal{T}$
end

That this algorithm does in fact generate the entire set $\mathcal{T}$ of neighborhood types follows from Corollary 1 and the fact that for every neighborhood type τ there is a path in $\mathcal{G}$ from the root vertex (0,0) to a vertex of neighborhood type τ.

4. Admissible Paths

Each directed edge (arrow) of the digraph $\mathcal{G}$ may be labelled by triples (τ, τ', d'), where τ, τ' are the neighborhood types of the initial and terminal vertices, respectively, and d' is the color (digit) of the edge. The set $\mathcal{L}$ of labels is finite, since both $\mathcal{T}$ (the set of neighborhood types) and D (the set of colors) are finite. Say that a label $l = (\tau, \tau', d')$ is *admissible* if there are vertices v, v' of the digraph $\mathcal{G}$ such that

(1) $v \to v'$ is an edge of $\mathcal{G}$;

(2) the neighborhood types of v, v' are τ, τ'; and

(3) among all edges $v'' \to v'$ ending at v', the edge $v \to v'$ has the smallest color.

(Recall that the set of edge colors is the digit set D, which is naturally ordered as a subset of $\mathbf{R}$. However, *any* order on the set of colors would work.) Note that this definition is independent of the choice of v, v' in the following sense: if there are vertices v, v' such that (1)-(3) hold, and if w, w' are vertices such that there is an edge $w \to w'$ with the same label as the edge $v \to v'$, then (1)-(3) hold for the pair w, w'. This follows from Lemma 5 and Corollary 1. Call a path γ in $\mathcal{G}$ *admissible* if every edge in γ has an admissible label. Denote by $\mathcal{L}$ the set of admissible labels.

Lemma 6. *For each vertex v of $\mathcal{G}$ there is a unique admissible path γ from the root vertex $(0,0)$ to v.*

Proof. This is certainly true for vertices v at depth 1, because for each such vertex there is only one edge terminating at v, namely $(0,0) \to v$, and this, by its uniqueness, has minimal color and therefore an admissible label. Suppose that the statement is true for all vertices at depth $m \geq 1$; we will show that it must then be true also for all vertices at depth $m+1$.

Let v' be any vertex at depth $m+1$. There is at least one directed edge $v \to v'$ with v a vertex at depth m; consequently, there is a unique edge $v \to v'$ with smallest color. By definition of admissibility, $v \to v'$ has an admissible label, and any other arrow $v'' \to v'$ terminating at v' has an inadmissible label. Thus, if there is an admissible path from the root vertex $(0,0)$ to v', then its final step must be $v \to v'$. But by the induction hypothesis there is a unique admissible path γ from the root vertex $(0,0)$ to v. The path obtained by adjoining the edge $v \to v'$ to γ is clearly admissible, and it is the only possible such path. □

Note. If the graph Γ is drawn in the natural way, with vertices at each level ordered in accordance with their usual order as real numbers, then the unique admissible path from the root to vertex v is simply the *rightmost* path in Γ from the root to v.

5. The Incidence Matrix

The incidence matrix A (the 0-1 matrix whose lead eigenvalue appears in equation (1)) has rows and columns indexed by the set $\mathcal{L}$ of admissible labels. For any two admissible labels $l = (\tau, \tau', d')$ and $l' = (\tau'', \tau''', d'')$ the l, l' entry of A is 1 if $\tau' = \tau''$ and 0 otherwise.

The matrix A determines a shift of finite type (Σ_A, σ). Here Σ_A is the sequence space consisting of all one-sided sequences $l_1 l_2 \dots$ with entries in $\mathcal{L}$ such that for every n, the (l_n, l_{n+1}) entry of A is 1. Observe that for every infinite admissible path in $\mathcal{G}$ the corresponding sequence of edge labels is an element of Σ_A. Conversely, for every element $l_1 l_2 \dots$ of Σ_A and every edge $v \to v'$ of $\mathcal{G}$ with label l_1 there is a unique admissible path with initial step $v \to v'$ for which the corresponding sequence of edge labels is $l_1 l_2 \ldots$.

Lemma 7. *Let $l, l' \in \mathcal{L}$ be admissible labels, and let $v \to v'$ be an edge of $\mathcal{G}$ with label l. Then the number of admissible paths of length $m+1$ with first step $v \to v'$ and final step labelled l' is $A^m_{l,l'}$.*

Proof. By induction on m. The case $m = 0$ is trivial. Suppose the result is true for some $m \geq 0$; then for each $l'' \in \mathcal{L}$ the number of length m admissible paths with first step $v \to v'$ and last step labelled l'' is $A^m_{l,l''}$. Now $A_{l'',l'} = 1$ iff the neighborhood type τ of the tail of l'' is the same as that of the head of l', and $A_{l'',l'} = 0$ otherwise. Thus, for any admissible path with last step $v'' \to v'''$ labelled l'' the number of edges with label l' emanating from v''' is $A_{l'',l'}$. Hence, by the induction hypothesis, the number of admissible paths of length $m+1 \geq 1$ with first step $v \to v'$ and final step labelled l' is

$$\sum_{l'' \in \mathcal{L}} A^m_{l,l''} A_{l'',l'} = A^{m+1}_{l,l'}.$$

□

For any $l \in \mathcal{L}$ denote by v_l the vector in $\mathbf{R}^{\mathcal{L}}$ with lth entry 1 and all other entries 0. Let τ_* be the neighborhood type of the root vertex (0,0), and let $u_* \in \mathbf{R}^{\mathcal{L}}$ be the vector with lth entry 1 if l is a label of the form (τ_*, τ', d') and 0 otherwise. (Any such label is necessarily admissible, since τ_* has only one vertex.) Let $\mathbf{1} \in \mathbf{R}^{\mathcal{L}}$ be the vector with all entries 1.

Corollary 2. $|\Lambda_m| = u_*^t A^m \mathbf{1}$.

Proof. This is an immediate consequence of Lemmas 7 and 6. □

Corollary 3. *If $\lambda > 0$ is the spectral radius of A, then*

$$\lim_{m\to\infty} |\Lambda_m|^{1/m} = \lambda. \tag{4}$$

Note. This gives another proof of Proposition 1, and shows that the limit λ is strictly positive.

Proof. The preceding corollary and the spectral radius formula imply that

$$\limsup_{m\to\infty} |\Lambda_m|^{1/m} \leq \lambda.$$

Thus, it must be shown that

$$\liminf_{m\to\infty} |\Lambda_m|^{1/m} \geq \lambda.$$

By the Perron-Frobenius theorem, λ is an eigenvector of A, with a nonnegative eigenvector. Consequently, there exist indices $l_1, l_2, \ldots, l_r$ and a constant $\varepsilon > 0$ such that for all sufficiently large n,

$$(\sum_{i=1}^r v_{l_i})^t A^n \mathbf{1}/\lambda^n \geq \varepsilon.$$

But each label l_i occurs somewhere in Γ, say at level k_i, and there is an admissible path from the root to l_i. It follows that for all large n,

$$\max_{1\leq i\leq r} u_*^t A^{n+k_i} \mathbf{1} \geq \varepsilon\lambda^n / r.$$

The preceding corollary and the fact that $|D|^{-1} \leq |\Lambda_n|/|\Lambda_{n-1}| \leq |D|$ now imply that for all sufficiently large n,

$$|\Lambda_n| \geq \varepsilon \lambda^{n-k_*}/r|D|^{k_*},$$

where k_* is the maximum of $k_1, k_2, \dots, k_r$. □

Note. In fact, the latter part of this argument will prove unnecessary for the main result below. The easy inequality $\limsup_{m\to\infty} |\Lambda_m|^{1/m} \leq \lambda$ implies that the box dimension of Λ is no larger than $\log\lambda/\log\beta$, by Proposition 1. In the next section we will give an independent proof that the Hausdorff dimension of Λ is no smaller than $\log\lambda/\log\beta$. Since the box dimension is no smaller than the Hausdorff dimension, it will then follow that $\liminf_{m\to\infty} |\Lambda_m|^{1/m} \geq \lambda$.

6. Maximum Entropy Measures

It is well known that the shift (Σ_A, σ) has an invariant probability measure of maximal entropy (see Parry [4] and/or Bowen [1]), and that under such a measure the coordinate process is a Markov chain. We will show that the maximal entropy measure (more precisely, a probability measure that is absolutely continuous with respect to the maximal entropy measure) induces a measure on Λ whose Hausdorff dimension is $\log\lambda/\log\beta$. This will imply that the Hausdorff dimension of Λ is at least $\log\lambda/\log\beta$, and by Proposition 2 it will then follow that the Hausdorff dimension of Λ *equals* $\log\lambda/\log\beta$.

Recall that λ is the spectral radius of the incidence matrix A. By the Perron-Frobenius theorem, there is a nonnegative vector $\mathbf{h}$ indexed by $\mathcal{L}$ such that

$$A\mathbf{h} = \lambda\mathbf{h}.$$

This eigenvector may be used to construct a transition probability matrix $\mathbf{P}$ on the set $\mathcal{L}^*$ of admissible labels such that $h(l) > 0$:

$$\mathbf{P}(l, l') = \frac{A(l, l')\mathbf{h}(l')}{\lambda\mathbf{h}(l)}.$$

Lemma 8. *There exist constants $0 < C_1 < C_2 < \infty$ with the following properties: (A) For every sequence $l_1, l_2, \dots$ in $\mathcal{L}^*$ and every integer $n \geq 1$,*

$$\prod_{i=1}^{n} \mathbf{P}(l_i, l_{i+1}) \leq C_2\lambda^{-n}.$$

(B) For every sequence $l_1, l_2, \dots$ in $\mathcal{L}^*$ such that $A(l_i, l_{i+1}) > 0$ for all i, and for every integer $n \geq 1$,

$$\prod_{i=1}^{n} \mathbf{P}(l_i, l_{i+1}) \geq C_1\lambda^{-n}.$$

Proof. This is an immediate consequence of the definition of $\mathbf{P}$, since the entries of $\mathbf{h}$ are strictly positive on $\mathcal{L}^*$ and the entries of A are zeros and ones. □

The transition probability matrix $\mathbf{P}$ may be used to define a probability measure μ on Λ as follows. Select any edge $v \to v'$ of $\mathcal{G}$ whose label l is an element of $\mathcal{L}^*$. Let $Y_0, Y_1, Y_2, \dots$ be a Markov chain with initial state $Y_0 = l$ and transition probability matrix $\mathbf{P}$; note that with probability 1 all entries of the sequence $Y_0Y_1Y_2\dots$ are elements of $\mathcal{L}^*$. The sequence $Y_0Y_1Y_2\dots$ determines a random path $X_0X_1X_2\dots$ in $\mathcal{G}$ (each X_n is a vertex of $\mathcal{G}$) starting at $X_0 = v$ and whose first step $X_0 \to X_1$ is

$v \to v'$. The random path $X_0X_1X_2\dots$ converges to a random point Z of Λ (recall that the boundary of the graph $\mathcal{G}$ may be naturally identified with Λ). The random variable Z induces a probability measure μ on Λ (the "distribution" of Z) by

$$\mu(A) = P\{Z \in A\}.$$

(*Note*: There is a different measure μ for each choice of initial edge $v \to v'$. For our purposes, any such choice will suffice.)

Lemma 9. *For μ-a.e. x,*

$$\lim_{\varepsilon \to 0} \frac{\log \mu([x-\varepsilon, x+\varepsilon])}{\log \varepsilon} = \frac{\log \lambda}{\log \beta}. \tag{5}$$

Proof. It suffices to show that

$$\lim_{n\to\infty} \frac{1}{n} \log \mu([x-\kappa\beta^{-n}, x+\kappa\beta^{-n}]) = -\log \lambda$$

for μ-a.e. x. By definition of μ, it suffices to consider only those $x \in \Lambda = \partial\mathcal{G}$ for which there is an (infinite) admissible path $x_0x_1x_2\dots$ in $\mathcal{G}$ (necessarily unique, by Lemma 6) converging to x, such that (i) $x_0 = v$ and $x_1 = v'$; (ii) the labels l_n of the arrows $x_n \to x_{n+1}$ are elements of $\mathcal{L}^*$; and (iii) successive labels l_n, l_{n+1} satisfy $A(l_n, l_{n+1}) = 1$.

Consider any infinite admissible path $x'_0x'_1x'_2\dots$ in $\mathcal{G}$ starting at $x'_0 = v, x'_1 = v'$ and converging to a point $x' \in \Lambda$ such that $|x' - x| \le \kappa\beta^{-m-r}$, where r is the depth of v. Then, by Corollary 3, the depth $(m+r)$ approximants x_m, x'_m to x, x', respectively, must satisfy $\rho(x_m, x'_m) \le 2\kappa$. By the Separation Lemma, there is a constant $K = K(2\kappa) < \infty$ such that any ρ-neighborhood of radius 2κ can contain no more than K distinct vertices; hence, there are at most K possibilities for x'_m. For each such possibility x'_m there is at most one admissible path from v to x'_m, by Lemma 6, and by Lemma 8 the μ-probability of this path is no larger than $C_2\lambda^{-m}$. Thus,

$$\mu([x-\kappa\beta^{-m-r}, x+\kappa\beta^{-m-r}]) \le KC_2\lambda^{-m}.$$

On the other hand, any admissible path that begins $x_0x_1x_2\dots x_m$ must converge to a point of $[x-\kappa\beta^{-m-r}, x+\kappa\beta^{-m-r}]$, and by Lemma 8 the probability of the set of such paths is at least $C_1\lambda^{-m}$, so

$$\mu([x-\kappa\beta^{-m-r}, x+\kappa\beta^{-m-r}]) \ge C_1\lambda^{-m}.$$

□

Theorem 1. $\dim_H(\Lambda) = \log\lambda / \log\beta$.

Proof. The Hausdorff dimension is never larger than the box dimension, so by Lemma 1 it suffices to show that $\dim_H(\Lambda)$ is at least as large as $\log\lambda/\log\beta$. The preceding lemma shows that Λ supports a probability measure satisfying Eq. (5). But Frostman's Lemma (see [2], Ch. 1, ex.) implies that no measure satisfying Eq. (5) can be supported by a Borel set of Hausdorff dimension smaller than $\log\lambda/\log\beta$. □

7. The Hausdorff Measure on Λ

Define $\delta = \log\lambda/\log\beta$, and consider the δ-dimensional Hausdorff measure H_δ restricted to the set Λ.

Theorem 2. $0 < H_\delta(\Lambda) < \infty$.

Proof. Consider again the covering $\mathcal{U}_m$ of Λ introduced in the proof of Lemma 1. The cardinality of the covering is $u_*^t A^{m-1}\mathbf{1}$, by Corollary 2, and each interval in the covering has radius $\kappa\beta^{-m}$. It follows from the definition of the outer Hausdorff measure that

$$H_\delta(\Lambda) \leq \lim_{m\to\infty}(u_*^t A^{m-1}\mathbf{1})\kappa\beta^{-m\delta} = \lim_{m\to\infty}(u_*^t A^{m-1}\mathbf{1})\kappa\lambda^{-m} < \infty,$$

the last inequality because λ is the spectral radius of A.

The inequality $H_\delta(\Lambda) > 0$ is a consequence of the existence of a probability measure μ on Λ with the property (5) above. The proof of Lemma 9 shows that for a suitable constant $\gamma > 0$, $\mu(J) \leq \gamma|J|^\delta$ for every interval J centered at a support point of μ. Since $\sum\mu(J) = 1$ for every covering of Λ, it follows that for every covering,

$$\sum |J|^\delta \geq 1/\gamma > 0.$$

$\square$

8. A Numerical Result

Consider the special case where $\beta = \beta_2 = 1+\sqrt{3}$ is the larger root of $x^2-2x-2 = 0$ and $D = \{0, 1, 3\}$. This is the most difficult case considered in [3]; there it is shown that for these parameters Λ has Lebesgue measure 0. Using the algorithm described in section 3, we have found that there are 43 distinct neighborhood types. For each pair τ, τ' of neighborhood types such that τ' is an offspring of τ, there is at most one color d such that (τ, τ', d) is an admissible label. Consequently, the matrix A may be "collapsed" to an equivalent matrix B with rows and columns indexed by the neighborhood types and with 0-1 entries indicating the absence or presence of admissible labels (τ, τ', d) of some color; the top eigenvalue λ of B is the same as the top eigenvalue of A. The eigenvalue λ may be approximated from above using the spectral radius formula

$$\lambda = \lim_{n\to\infty} \downarrow \|B^n\|^{\frac{1}{n}};$$

using $n = 1024$ gives the upper bounds

$$\lambda \leq 2.65584 \text{ and } \delta \leq .971847$$

A lower bound for λ is $b_n^{1/n}$, where b_n is the minimum entry of B^n; using $n = 1024$ gives the lower bounds

$$\lambda \geq 2.63855 \text{ and } \delta \geq .965351$$

A *Mathematica* notebook containing the code implementing the algorithms for computing the set of neighborhood types and the matrix B is available from the author. The *Mathematica* computations use exact integers, so the only numerical errors are in extraction of the 1024^{th} roots.

Although the algorithm described in this paper applies in principle to any Pisot number β and any digit set, in practice it is useable in very few cases. Even in the next simplest case, $D = \{0, 1, 3\}$ and $\beta = \beta_3$, the largest root of $x^3 - 2x^2 - 2x - 2 = 0$, computation is impractical, as there are 4017 distinct neighborhood types.

References

[1] Bowen, R. (1975) Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms. *Springer Lecture Notes in Math.* **470**. MR **56:**1364

[2] Falconer, K. (1985) *The Geometry of Fractal Sets.* Cambridge Univ. Press. MR **88d:**28001

[3] Keane, M., Smorodinsky, M., and Solomyak, B. (1995) On the morphology of $\gamma-$expansions with deleted digits. *Trans. Amer. Math. Soc.* **347**, 955-966. MR **95h:**11079

[4] Parry, W. (1964) Intrinsic Markov chains. *Trans. Amer. Math. Soc.* **112**, 55-64. MR **28:**4579

[5] Pollicott, M. and Simon, K. (1995) The Hausdorff dimension of $\lambda-$expansions with deleted digits. *Trans. Amer. Math. Soc.* **347**, 967-983. MR **95h:**11080

Department of Statistics, Mathematical Sciences Bldg., Purdue University, West Lafayette, Indiana 47907

E-mail address: `lalley@stat.purdue.edu`