PATRICIA'S BAD DISTRIBUTIONS

LOUIGI ADDARIO-BERRY, PAT MORIN, AND RALPH NEININGER

ABSTRACT. The height of a random PATRICIA tree built from independent, identically distributed infinite binary strings with arbitrary diffuse probability distribution μ on $\{0,1\}^{\mathbb{N}}$ is studied. We show that the expected height grows asymptotically sublinearly in the number of leaves for any such μ , but can be made to exceed any specific sublinear growth rate by choosing μ appropriately.

1. INTRODUCTION AND RESULTS

The PATRICIA tree is a space efficient data structure for strings which improves on the trie. For the purpose of this note it is sufficient to introduce these tree structures for binary strings: Label the nodes of the complete infinite rooted binary tree by the elements of $\bigcup_{k=0}^{\infty} \{0,1\}^k$, starting at the root with \emptyset and left and right child of a node labelled $v \in \{0,1\}^k$ with v0 and v1, respectively. Here, for $v \in \{0,1\}^k$ with $v = (v_1, \ldots, v_k)$ we abbreviate vas $v = v_1 \ldots v_k$ and denote $vi := v_1 v_2 \ldots v_k i$ for i = 0, 1.

The coming definitions are depicted in Figure 1. For distinct infinite binary strings $x_1, \ldots, x_n \in \{0, 1\}^{\mathbb{N}}$ a finite tree called a *trie* (or *radix search tree*) to represent the strings x_1, \ldots, x_n is constructed by first associating with each x_i the infinite path in $\bigcup_{k=0}^{\infty} \{0, 1\}^k$ consisting of the nodes whose labels are the prefixes of x_i . The node labelled with the shortest such prefix that is not a prefix of any x_j with $j \in \{1, \ldots, n\} \setminus \{i\}$ becomes a leaf in the trie representing string x_i for $i = 1, \ldots, n$. The resulting tree, which is a finite binary tree with n leaves, is the trie representing x_1, \ldots, x_n . Next, starting from the trie, all vertices with out-degree 1 (i.e. with exactly one child) are deleted edge. This results in the PATRICIA tree, which was introduced independently by Morrison [20] and Gwehenberger [11] and first systematically analysed by Knuth [16]. The PATRICIA tree contains all the information needed to retrieve the strings and to perform operations such as sorting, searching and selecting; for broad expositions, see [17, 19, 22].

PATRICIA trees have been analysed assuming various probabilistic models for the input strings; where usually the infinite strings are assumed to be independent and identically distributed over $\{0,1\}^{\mathbb{N}}$. Note that atoms of such a distribution result in identical strings with positive probability, and

Date: March 11, 2024.

²⁰¹⁰ Mathematics Subject Classification. 68P05, 60C05, 68R15, 68P10.

in this case the construction of the trie does not lead to a finite tree. Hence, the law of the strings is usually assumed to be diffuse (non-atomic). Special cases of such diffuse probability distributions have been considered in the analysis of algorithms on strings such as the Bernoulli models, Markov model, dynamical sources or the density model; see [3, 5, 10, 12–14, 18] and the references given in these papers.



Figure 1. On the left the trie for the strings 00000..., 00001..., 0100..., 0101..., 1100..., and 1101... is shown. It's leaves are the full black vertices, the indicated children of the full black vertices do not belong to the trie. Vertices with out-degree 1 within the trie are indicated by arrows. On the right the resulting PATRICIA tree by deleting corresponding edges is shown.

In the present note we focus on the height of a PATRICIA tree, which is the maximal (graph) distance of any leave from the root. The asymptotic behavior of the height of tries under the Bernoulli models is covered in Devroye [4, 6]. For general diffuse laws concentration of the height of PA-TRICIA trees is studied (assuming only independence of the infinite strings not necessarily identical distribution) in Devroye [7] based on results from [1]; see also [15] for concentration of the height of PATRICIA trees in the Bernoulli model.

While such studies aim to show that the height behaves well with respect to applications from algorithms, Evans and Wakolbinger [8, 9] studied these random tree structures as tree-valued transient Markov chains from the perspective of Doob–Martin boundary theory. They asked (private communication) how high PATRICIA trees can grow for arbitrary diffuse probability distributions of the strings (see [9, Section 5] for specific examples). The subject of the present note is to answer this question by Theorems 1 and 2: The expected height grows always sublinearly, but can be made to exceed any fixed sublinear growth rate by the choice of an appropriate diffuse law. For a diffuse probability distribution μ on $\{0,1\}^{\mathbb{N}}$ and $(\Xi^{(j)})_{j\in\mathbb{N}}$ a sequence of independent and identically distributed random strings with law μ we denote by H_n^{μ} the height of the PATRICIA tree constructed from $\Xi^{(1)}, \ldots, \Xi^{(n)}$.

Theorem 1. For all diffuse probability distributions μ on $\{0,1\}^{\mathbb{N}}$ we have, as $n \to \infty$, that

$$\frac{\mathbb{E}[H_n^{\mu}]}{n} \to 0, \quad and \quad \frac{H_n^{\mu}}{n} \to 0 \ almost \ surely.$$

Theorem 2. For any sequence $\alpha = (\alpha_n)_{n \in \mathbb{N}}$ of positive numbers with $\alpha_n \to \infty$ as $n \to \infty$ there exists a diffuse probability distribution $\mu = \mu^{(\alpha)}$ on $\{0,1\}^{\mathbb{N}}$ such that for all n sufficiently large

$$\mathbb{E}[H_n^{\mu}] \ge \frac{n}{\alpha_n}.$$

We call a law μ on $\{0,1\}^{\mathbb{N}}$ causing large expected heights $\mathbb{E}[H_n^{\mu}]$ bad since such laws are undesirable from the point of view of the efficiency of algorithms based on PATRICIA trees. The remaining part of the present note contains proofs of these two theorems.

2. Proofs

2.1. **Proof of Theorem 1.** We start with a technical observation:

Lemma 3. Suppose μ is a diffuse probability distribution on $\{0, 1\}^{\mathbb{N}}$, and let $\Xi = (\xi_i)_{i \in \mathbb{N}}$ be random with law μ . Then for all ε there exists $k = k(\varepsilon) \in \mathbb{N}$ such that for any string $v = v_1 \dots v_k \in \{0, 1\}^k$, $\mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) < \varepsilon$.

Proof. Suppose for a contradiction that there exists $\varepsilon > 0$ such that for all $k \in \mathbb{N}$ there is a string $v_1 \ldots v_k \in \{0,1\}^k$ such that $\mathbb{P}(\xi_1 \ldots \xi_k = v_1 \ldots v_k) \ge \varepsilon$. Then by a compactness argument shown below there exists an infinite string $v = (v_i)_{i \in \mathbb{N}} \in \{0,1\}^\infty$ such that for all $k \in \mathbb{N}$, $\mathbb{P}(\xi_1 \ldots \xi_k = v_1 \ldots v_k) \ge \varepsilon$. The events $\{\xi_1 \ldots \xi_k = v_1 \ldots v_k\}$ are decreasing in k, so this implies that

$$\mathbb{P}(\Xi = v) = \lim_{k \to \infty} \mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) \ge \varepsilon,$$

which contradicts the assumption that μ is diffuse.

It remains to show the existence of the infinite string $v = (v_i)_{i \in \mathbb{N}} \in \{0, 1\}^{\mathbb{N}}$ such that for all $k \in \mathbb{N}$, $\mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) \geq \varepsilon$. Consider $\{0, 1\}$ as a topological space with the discrete topology (all subsets being open) and $\{0, 1\}^{\mathbb{N}}$ as the product space with the product topology. As a product of compact spaces $\{0, 1\}^{\mathbb{N}}$ is compact, and the projections $\Pi_k : \{0, 1\}^{\mathbb{N}} \to$ $\{0, 1\}^k$ given by

$$(v_i)_{i\in\mathbb{N}} \xrightarrow{\Pi_k} v_1 \dots v_k$$

are continuous for all $k \in \mathbb{N}$. Hence, the set

$$V_k := \{ (v_i)_{i \in \mathbb{N}} \in \{0, 1\}^{\mathbb{N}} \mid \mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) \ge \varepsilon \}$$
$$= \bigcup_{\substack{v_1 \dots v_k \in \{0, 1\}^k \\ \mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) \ge \varepsilon}} \Pi_k^{-1}(\{v_1 \dots v_k\})$$

is closed and thus compact in $\{0,1\}^{\mathbb{N}}$. This implies that $(V_k)_{k\in\mathbb{N}}$ is a nested sequence of non-empty, compact sets. Hence, we have

$$\bigcap_{k=1}^{\infty} V_k \neq \emptyset.$$

Any element v of $\bigcap_{k=1}^{\infty} V_k$ has the desired property.

Proof of Theorem 1. Fix a diffuse probability distribution μ on $\{0,1\}^{\mathbb{N}}$. Let $\Xi^{(j)} = (\xi_i^{(j)})_{i \in \mathbb{N}}$ for $j \in \mathbb{N}$ be independent, identically distributed with law μ and denote by T_n the PATRICIA tree built from $\Xi^{(1)}, \ldots, \Xi^{(n)}$.

We first show that $H_n^{\mu}/n \to 0$ almost surely. Fix any $\varepsilon \in (0, 1/4)$. Let $k = k(\varepsilon)$ be as in Lemma 3, so that for any string $v = v_1 \dots v_k \in \{0, 1\}^k$, if $\Xi = (\xi_i)_{i \in \mathbb{N}}$ has law μ then $\mathbb{P}(\xi_1 \dots \xi_k = v_1 \dots v_k) < \varepsilon$. To prove $H_n^{\mu}/n \to 0$ almost surely we first show that

(1)
$$\mathbb{P}(\exists n_0 \,\forall n \ge n_0 : H_n^{\mu} \le k + 2\varepsilon n) = 1.$$

Note that if the event

$$E_{n,k} := \bigcup_{v_1 \dots v_k \in \{0,1\}^k} \{ |\{1 \le j \le n : \xi_1^{(j)} \dots \xi_k^{(j)} = v_1 \dots v_k\}| \ge 2\varepsilon n \}$$

does not occur then the subtrees of T_n rooted at nodes $v \in \{0, 1\}^k$ all have at most $2\varepsilon n$ leaves and so height less than $2\varepsilon n$; thus if $E_{n,k}$ does not occur then $H_n^{\mu} \leq k + 2\varepsilon n$. It follows that

$$\mathbb{P}(\exists n_0 \forall n \ge n_0 : H_n^{\mu} \le k + 2\varepsilon n)$$

$$\ge \mathbb{P}(E_{n,k} \text{ occurs for at most finitely many values } n)$$

$$= \mathbb{P}\left(\left(\limsup_{n \to \infty} E_{n,k}\right)^c\right) ,$$

so to prove (1) it suffices to show that the probability of $\limsup_{n\to\infty} E_{n,k}$ is 0. For this, simply note that

$$\mathbb{P}(E_{n,k}) \leq \sum_{v_1 \dots v_k \in \{0,1\}^k} \mathbb{P}(|\{1 \leq j \leq n : \xi_1^{(j)} \dots \xi_k^{(j)} = v_1 \dots v_k\}| \geq 2\varepsilon n)$$

$$\leq 2^k \mathbb{P}(Y_n \geq 2\varepsilon n),$$

where Y_n has the Binomial distribution $Bin(n, \varepsilon)$; the second inequality holds since the events that $\xi_1^{(j)} \dots \xi_k^{(j)} = v_1 \dots v_k$ are independent for distinct $1 \leq j \leq n$, and each has probability at most ε . A Chernoff bound then gives

$$\mathbb{P}(E_{n,k}) \le 2^k e^{-\varepsilon n/2}.$$

Since this is summable, it follows by the first Borel–Cantelli lemma that

$$\mathbb{P}\left(\limsup_{n\to\infty}E_{n,k}\right)=0,$$

hence we obtain (1). Now, note that for any $m_0 \in \mathbb{N}$,

$$\left\{\frac{H_n^{\mu}}{n} \to 0\right\} = \bigcap_{m=m_0}^{\infty} \bigcup_{n_0=1}^{\infty} \bigcap_{n=n_0}^{\infty} \left\{\frac{H_n^{\mu}}{n} \le \frac{3}{m}\right\} \,.$$

Thus, for $\varepsilon = \frac{1}{m}$ with fixed $m \ge m_0$ we can choose n sufficiently large so that $k(\varepsilon)/n \le \varepsilon$ and obtain

$$\left\{H_n^{\mu} \le k(\varepsilon) + 2\varepsilon n\right\} \subset \left\{\frac{H_n^{\mu}}{n} \le \frac{3}{m}\right\}$$

and see that (1) implies $H_n^{\mu}/n \to 0$ almost surely.

Finally, note that by construction of the PATRICIA tree we deterministically have $H_n^{\mu} \leq n-1$, thus $H_n^{\mu}/n \leq 1$. Hence, we obtain from $H_n^{\mu}/n \to 0$ almost surely and dominated convergence that $\mathbb{E}[H_n^{\mu}]/n \to 0$.

2.2. **Proof of Theorem 2.** As building blocks for our bad distributions we first define a set of auxiliary probability distributions $(\mu_N, N \in \mathbb{N})$, on $\{0,1\}^{\mathbb{N}}$ as follows. For fixed $N \in \mathbb{N}$ we choose T uniformly at random from $\{1,\ldots,N^2\}$. Independently of T, let $(B_i)_{i\in\mathbb{N}}$ be independent Bernoulli $(\frac{1}{2})$ distributed random variables. Then define a sequence $(\vartheta_i)_{i\in\mathbb{N}}$ by

(2)
$$\vartheta_i = \begin{cases} 0, & \text{if } i < T, \\ 1, & \text{if } i = T, \\ B_{i-T}, & \text{if } i > T. \end{cases}$$

Now, μ_N is defined as the law of the string $\Theta = (\vartheta_i)_{i \in \mathbb{N}}$. Note that by definition μ_N is diffuse for all $N \in \mathbb{N}$. We use the notation

$$\langle \Theta \rangle := \min\{i \in \mathbb{N} \mid \vartheta_i = 1\}$$

for the index of the first entry of Θ equal to 1.

Lemma 4. For any $n \in \{1, \ldots, N\}$ we have $\mathbb{E}[H_n^{\mu_N}] \ge n-2$.

Proof. Let $1 \leq n \leq N \in \mathbb{N}$ and $\Theta^{(1)}, \ldots, \Theta^{(n)}$ be i.i.d. with law μ_N . We consider the set $A := \{\langle \Theta^{(1)} \rangle, \ldots, \langle \Theta^{(n)} \rangle\} \subset \{1, \ldots, N^2\}$. By construction of the PATRICIA tree we have

(3)
$$H_n^{\mu_N} \ge |A| - 1,$$

where |A| denotes the cardinality of A, i.e., the number of distinct elements within $\{\langle \Theta^{(1)} \rangle, \ldots, \langle \Theta^{(n)} \rangle\}$. For all $1 \leq i < j \leq n$ we have $\mathbb{P}(\langle \Theta^{(i)} \rangle = \langle \Theta^{(j)} \rangle) = 1/N^2$. Hence, we obtain

(4)
$$\mathbb{E}[|A|] \ge n - \mathbb{E}\left[\sum_{1 \le i < j \le n} \mathbf{1}_{\{\langle \Theta^{(i)} \rangle = \langle \Theta^{(j)} \rangle\}}\right] \ge n - \frac{n^2}{2N^2} \ge n - 1,$$

since $n \leq N$. Now, (3) and (4) imply the assertion.

Proof of Theorem 2. Without loss of generality we may assume that $\alpha_n = o(n)$. There exists an $n_0 \in \mathbb{N}$ such that $\alpha_n \geq 8$ for all $n \geq n_0$. We define $\beta_n := \lfloor \log_2 \alpha_n \rfloor - 2$ and a sequence $(A(n))_{n \in \mathbb{N}}$ as a generalized inverse of $(\beta_n)_{n \in \mathbb{N}}$ by

(5)
$$A(n) := \max\{m \in \mathbb{N} \mid \beta_m \le n\}, \quad n \in \mathbb{N}.$$

The probability distribution $\mu^{(\alpha)}$ on $\{0,1\}^{\mathbb{N}}$ is obtained in two stages. First, let G be a random variable with geometric distribution with parameter $\frac{1}{2}$, i.e., with $\mathbb{P}(G = k) = (\frac{1}{2})^k$ for $k \in \mathbb{N}$. Then define a sequence $(\lambda_i)_{i \in \mathbb{N}}$ by

(6)
$$\lambda_i = \begin{cases} 0, & \text{if } i < G, \\ 1, & \text{if } i = G, \\ \vartheta_{i-G}, & \text{if } i > G, \end{cases}$$

where $\Theta = (\vartheta_i)_{i \in \mathbb{N}}$ is independent of G and has law $\mu_{A(k)}$ defined in (2) with $A(\cdot)$ defined in (5). We then define $\mu = \mu^{(\alpha)}$ as the law of $\Lambda = (\lambda_i)_{i \in \mathbb{N}}$. Since the $\mu_{A(k)}$ are diffuse, we obtain that μ is diffuse.

Now, let $\Lambda^{(j)} = (\lambda_i^{(j)})_{i \in \mathbb{N}}$ for $j \in \mathbb{N}$ be independent with law μ . For $n \geq n_0$, by construction,

$$X_n := \left| \left\{ 1 \le j \le n : \left(\lambda_1^{(j)}, \dots, \lambda_{\beta_n}^{(j)} \right) = (0, \dots, 0, 1) \right\} \right|$$

is $\operatorname{Bin}(n, 2^{-\beta_n})$ -distributed. To get rid of the floors in the definition of β_n denote by X'_n a $\operatorname{Bin}(n, 4/\alpha_n)$ -distributed random variable. Note that $2^{-\beta_n} \geq 4/\alpha_n$. By Okamoto's inequality, see [21] or [2, Exercise 2.12] we have

$$\mathbb{P}\left(X_n < \frac{2n}{\alpha_n}\right) \le \mathbb{P}\left(X'_n - \frac{4n}{\alpha_n} < -\frac{2n}{\alpha_n}\right)$$
$$\le \exp\left(-\frac{n(2/\alpha_n)^2}{2(4/\alpha_n)(1-4/\alpha_n)}\right) \le \exp\left(-\frac{n}{2\alpha_n}\right).$$

Hence, with high probability at least $\lceil 2n/\alpha_n \rceil$ of the *n* strings start with the prefix $(0, \ldots, 0, 1)$ of length β_n and thus have suffixes $(\lambda_{\beta_n+1}^{(j)}, \lambda_{\beta_n+2}^{(j)}, \ldots)$ drawn independently from $\mu_{A(\beta_n)}$ for the respective *j*. For all $n \ge n_0$ we have $\lceil 2n/\alpha_n \rceil \le n \le A(\beta_n)$. Hence, by Lemma 4, $\lceil 2n/\alpha_n \rceil$ such strings cause an expected height of at least $2n/\alpha_n - 2$. Together we obtain for all sufficiently large *n*, note also $\alpha_n = o(n)$, that

$$\mathbb{E}[H_n^{\mu}] \ge \mathbb{P}\left(X_n \ge \frac{2n}{\alpha_n}\right) \mathbb{E}\left[H_n^{\mu} \middle| X_n \ge \frac{2n}{\alpha_n}\right]$$
$$\ge \left(1 - \exp\left(-\frac{n}{2\alpha_n}\right)\right) \left(\frac{2n}{\alpha_n} - 2\right)$$
$$\ge \frac{n}{\alpha_n}.$$

Acknowledgement. This research was mainly done during the Sixteenth Annual Workshop on Probability and Combinatorics at McGill University's Bellairs Institute in Holetown, Barbados. We thank Bellairs Institute for its hospitality and support. We also thank Jasmin Straub for comments on a draft of this note.

References

- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart, A sharp concentration inequality with applications, Random Structures Algorithms 16 (2000), no. 3, 277–292. MR 1749290
- <u>Concentration inequalities</u>, Oxford University Press, Oxford, 2013, A nonasymptotic theory of independence, With a foreword by Michel Ledoux. MR 3185193
- Julien Clément, Philippe Flajolet, and Brigitte Vallée, Dynamical sources in information theory: a general analysis of trie structures, vol. 29, 2001, Average-case analysis of algorithms (Princeton, NJ, 1998), pp. 307–369. MR 1887308
- 4. Luc Devroye, A note on the probabilistic analysis of patricia trees, Random Structures & Algorithms **3** (1992), no. 2, 203–214.
- 5. ____, A study of trie-like structures under the density model, Ann. Appl. Probab. 2 (1992), no. 2, 402–434. MR 1161060
- 6. _____, Universal limit laws for depths in random trees, SIAM Journal on Computing **28** (1998), no. 2, 409–432.
- Universal asymptotics for random tries and PATRICIA trees, Algorithmica 42 (2005), no. 1, 11–29. MR 2131826
- 8. Steven N. Evans and Anton Wakolbinger, *Radix sort trees in the large*, Electronic Communications in Probability **22** (2017), no. none, 1 13.
- Steven N. Evans and Anton Wakolbinger, *PATRICIA bridges*, Genealogies of interacting particle systems, Lect. Notes Ser. Inst. Math. Sci. Natl. Univ. Singap., vol. 38, World Sci. Publ., Hackensack, NJ, [2020] ©2020, pp. 231–265. MR 4448159
- 10. Michael Fuchs, Hsien-Kuei Hwang, and Vytas Zacharovas, An analytic approach to the asymptotic variance of trie statistics and related structures, ArXiv **abs/1303.4244** (2013).
- Gernot Gwehenberger, Anwendung einer binären Verweiskettenmethode beim Aufbau von Listen, Elektronische Rechenanlagen 10 (1968), 223– 226.
- 12. Jasper Ischebeck, *Central limit theorems for fringe trees in patricia tries*, Preprint arXiv:2305.14900 [math.PR].
- 13. Svante Janson, *Renewal theory in the analysis of tries and strings*, Theor. Comput. Sci. **416** (2012), 33–54.
- 14. _____, Central limit theorems for additive functionals and fringe trees in tries, Electronic Journal of Probability **27** (2022), no. none, 1 63.

- Charles Knessl and Wojciech Szpankowski, *Limit laws for the height* in *PATRICIA tries*, vol. 44, 2002, Analysis of algorithms, pp. 63–97. MR 1932678
- Donald E. Knuth, The art of computer programming. Volume 3, Addison-Wesley Series in Computer Science and Information Processing, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1973, Sorting and searching. MR 445948
- The art of computer programming. Vol. 3, Addison-Wesley, Reading, MA, 1998, Sorting and searching, Second edition [of MR0445948]. MR 3077154
- 18. Kevin Leckey, Ralph Neininger, and Wojciech Szpankowski, Towards more realistic probabilistic models for data structures: The external path length in tries under the markov model, pp. 877–886.
- Hosam M. Mahmoud, Evolution of random search trees, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., New York, 1992, A Wiley-Interscience Publication. MR 1140708
- 20. Donald R. Morrison, *Patricia—practical algorithm to retrieve information coded in alphanumeric*, Journal of the ACM **15** (1968), 514–534.
- Masashi Okamoto, Some inequalities relating to the partial sum of binomial probabilities, Ann. Inst. Statist. Math. 10 (1958), 29–35. MR 99733
- 22. Wojciech Szpankowski, Average case analysis of algorithms on sequences, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2001, With a foreword by Philippe Flajolet. MR 1816272

DEPARTMENT OF MATHEMATICS AND STATISTICS, MCGILL UNIVERSITY, MONTRÉAL, CANADA

Email address: louigi.addario@mcgill.ca

School of Computer Science, Carleton University, Ottawa Ontario Canada Email address: morin@scs.carleton.ca

Institute of Mathematics, Goethe University Frankfurt, Frankfurt A.M., Germany

Email address: neininger@math.uni-frankfurt.de

8