MULTIPLE RECURRENCE AND ALGORITHMIC RANDOMNESS

RODNEY G. DOWNEY, SATYADEV NANDAKUMAR, AND ANDRÉ NIES

ABSTRACT. This work contributes to the programme of studying effective versions of "almost everywhere" theorems in analysis and ergodic theory via algorithmic randomness. We determine the level of randomness needed for a point in a Cantor space $\{0,1\}^{\mathbb{N}}$ with the uniform measure and the usual shift so that effective versions of the multiple recurrence theorem of Furstenberg holds for iterations starting at the point. We consider recurrence into closed sets that possess various degrees of effectiveness: clopen, Π_1^0 with computable measure, and Π_1^0. The notions of Kurtz, Schnorr, and Martin-Löf randomness, respectively, turn out to be sufficient. We obtain similar results for multiple recurrence with respect to the k commuting shift operators on $\{0,1\}^{\mathbb{N}^k}$.

1. INTRODUCTION

We begin with some background that will lead to our main definition. A measurable operator $T\colon X \to X$ on a probability space $(X, \mathcal{B}, \mu)$ is called *measure preserving* if $\mu T^{-1}(A) = \mu A$ for each $A \in \mathcal{B}$.

Theorem 1.1 (Furstenberg multiple recurrence theorem, see [4] Thm. 7.15)**.**
Let $(X, \mathcal{B}, \mu)$ be a probability space. Let $T_1, \ldots, T_k$ be commuting measure preserving operators on X. Let $A \in \mathcal{B}$ with $\mu A > 0$. There is $n > 0$ such that $0 < \mu \bigcap_i T_i^{-n}(A)$.

For this paper the following equivalent formulation of Thm. 1.1 will matter. We verify their equivalence in Subsection 1.

Corollary 1.2. *With the hypotheses of Thm. 1.1, for μ-a.e. $x \in A$, there is an $n > 0$ such that $x \in \bigcap_i T_i^{-n}(A)$.*

An important special case is that T_i is the power V^i of a measure-preserving operator V.

Corollary 1.3. *Let $(X, \mathcal{B}, \mu)$ be a probability space. Let V be a measure preserving operator. Let $A \in \mathcal{B}$ and $\mu A > 0$. For each k, for μ-a.e. $x \in A$ there is n such that $\forall i.1 \leq i \leq k\,[x \in V^{-ni}(A)]$.*

(A full "almost everywhere" version of Corollary 1.3 would assert that for μ-a.e. x there is n such that $\forall i.1 \leq i \leq k\,[x \in V^{-ni}(A)]$. Note that we can expect such a version only for ergodic operators, even if $k = 1$. For, if A is S-invariant, then an iteration starting from $x \notin A$ will never get into A.)

We will mostly assume that $(X, \mathcal{B}, \mu)$ is Cantor space $2^{\mathbb{N}}$ with the product measure λ. In the following X, Y, Z will denote elements of Cantor space. We will work with the shift $T\colon 2^{\mathbb{N}} \to 2^{\mathbb{N}}$ as the measure preserving operator. Thus, $T(Z)$ is obtained by deleting the first entry of the bit sequence Z. We note that this operator is (strongly) mixing, and hence strongly ergodic, namely, all of its powers are ergodic.

The following is our central definition.

Definition 1.4. Let $\mathcal{P} \subseteq 2^{\mathbb{N}}$ be measurable, and let $Z \in 2^{\mathbb{N}}$. We say that Z is *k-recurrent in $\mathcal{P}$* if there is $n \geq 1$ such that

$$(\diamond) \qquad Z \in \bigcap_{1 \leq i \leq k} T^{-ni}(\mathcal{P}).$$

We say that Z is *multiply recurrent* in $\mathcal{P}$ if Z is k-recurrent in $\mathcal{P}$ for each $k \geq 1$.

In other words, Z is k-recurrent in $\mathcal{P}$ if there is $n \geq 1$ such that removing $n, 2n, \ldots, kn$ bits from the beginning of Z takes us into $\mathcal{P}$.

We only consider multiple recurrence for closed sets. Note that for (multiple) recurrence in the sense of Thm. 1.1, this is not an essential restriction, because any set of positive measure contains a closed subset of positive measure.

Goal of this paper. For background on randomness notions see [10, Ch. 3] or [1, Chapters 6 and 7]. Since limit laws in probability theory embody certain intuitive properties associated with randomness, it is important to establish that objects of study in algorithmic randomness satisfy such laws. Indeed, the Law of Large Numbers [11], the Law of Iterated Logarithm [12] Birkhoff's Ergodic theorem [13], [3] and the Shannon-McMillan-Breiman Theorem [6], [7] are satisfied by Martin-Löf random sequences. These results serve to justify the intuition that the set of Martin-Löf points is a canonical set of measure 1 which obeys all effective limit laws which are true almost everywhere.

In this vein, we analyse how weaker and weaker effectiveness conditions on a closed set $\mathcal{P}$ ensure multiple recurrence when starting from a sequence Z that satisfies a stronger and stronger randomness property for an algorithmic test notion. We begin with the strongest effectiveness condition, being clopen; in this case it is easily seen that weak (or Kurtz) randomness of Z suffices. As most general effectiveness condition we will consider being effectively closed (i.e. Π_1^0); Martin-Löf-randomness turns out to be the appropriate notion. The proof will import some method from the case of a clopen $\mathcal{P}$. Note that Π_1^0 subsets of Cantor space are often called Π_1^0 *classes*.

The major part of the present work establishes the multiple recurrence property of algorithmically random sequences under the left-shift transformation. In the question of "structure versus randomness" in measure-preserving dynamical systems, this represents a setting which has a high degree of randomness. In a final section, we also provide a proof of the principle, well-known in practice, that a very structured system, namely Kronecker systems also exhibit multiple recurrence for all points.

Related results. In the theory of algorithmic randomness, we are interested to know which of the well-studied notions of algorithmic randomness are characterized by effective versions of limit laws. Of particular interest are the computable versions of theorems in dynamical systems. An early result is the classic theorem of Kučera that a sequence X is Martin-Löf random if and only if for every Π_1^0 class P with positive measure, there is a suffix of X which is in P [8]. This result can be recast into ergodic-theoretic language using the left-shift transformation. Bienvenu, Day, Hoyrup, Mezhirov and Shen [2] have generalized this to arbitrary computable ergodic transformations. They show that for a computable probability space and a computable ergodic transformation T, a member x in the space is Martin-Löf random if and only if for every Π_1^0 set P of positive measure, there is an n such that $T^n x \in P$. These can be viewed as how Poincaré recurrence relates to algorithmic randomness notions.

Birkhoff's ergodic theorem that for every ergodic transformation T defined on a probability space $(X, \mathcal{F}, \mu)$ and for every function $f \in L^1(X)$, the average $\lim_{n \to \infty} 1/n (\sum_{i=0}^{n-1} f(T^i x))$ converges to $\int f d\mu$ for μ-almost every $x \in X$. The question of how Birkhoff's ergodic theorem relates to algorithmic randomness has also been studied. Gács, Hoyrup and Rojas [5] establish that a point x in a computable probability space under a computable ergodic transformation $T : X \to X$ is Schnorr random if and only if it obeys the Birkhoff ergodic theorem with respect to all Π_1^0 sets of computable positive measure — *i.e.* the asymptotic frequency with which the orbit of x visits such a Π_1^0 set P is exactly the probability of P. Franklin, Greenberg, Miller and Ng [3] show that under the same setting but without the assumption that the measure of the Π_1^0 set is computable, a point x is Martin-Löf random if and only if it obeys the Birkhoff ergodic theorem with respect to Π_1^0 sets of positive measure.

Proof that Thm. 1.1 and Cor. 1.2 are equivalent. Cor. 1.2 clearly yields Thm. 1.1 because it implies that $\bigcap_i T_i^{-n}(A)$ has positive measure for some n. Conversely, let us show that Thm.

1.1 yields Cor. 1.2. Let $R_n = \bigcap_i T_i^{-n}(A)$. We recursively define a sequence $\langle n_p \rangle_{p<N}$ of numbers and a descending sequence $\langle A_p \rangle_{p<N}$ of sets, where $0 < N \leq \omega$.

Let $n_0 = 0$, and $A_0 = A$. Suppose n_p and A_p have been defined. If $\mu A_p = 0$ let $N = p+1$ and finish. Otherwise, let n_{p+1} be the least $n > n_p$ such that $\mu(A_p \cap R_n) > 0$, and let $A_{p+1} = A_p - R_n$.

Let $A_N = \bigcap_{p<N} A_p$. Then $\mu A_N = 0$: This is clear if N is finite. If $N = \omega$ and $\mu A_N > 0$, by Thm. 1.1 there is n such that $\mu(R_n \cap A_N) > 0$. This contradicts the definition of A_{p+1} where $n_p < n \leq n_{p+1}$.

Since $\mu A_N = 0$, Cor. 1.2 follows.

Notation. For a set of strings $S \subseteq 2^{<\omega}$, by $[S]^{\prec}$ we denote the open set $\{Y \in 2^{\mathbb{N}} \colon \exists \sigma \in S \, [\sigma \prec Y]\}$. We write $\lambda[S]^{\prec}$ for the measure of this set, namely $\lambda([S]^{\prec})$.

Recall that we work with the shift $T \colon 2^{\mathbb{N}} \to 2^{\mathbb{N}}$ as the measure preserving operator. We will write Z_n for $T^n(Z)$, the tail of Z starting at bit position n. Thus, for any $\mathcal{C} \subseteq 2^{\mathbb{N}}$, $Z \in 2^{\mathbb{N}}$ and $k \in \mathbb{N}$, $Z \in T^{-k}(\mathcal{C}) \leftrightarrow Z_k \in \mathcal{C}$.

2. Multiple recurrence for weakly random sequences

Recall that Z is weakly (or Kurtz) random if Z is in no null Π_1^0 class.

Proposition 2.1. *Let $\mathcal{P} \subseteq 2^{\mathbb{N}}$ be a non-empty clopen set. Each weakly random bit sequence Z is multiply recurrent in $\mathcal{P}$.*

Proof. Suppose Z is not k-recurrent in $\mathcal{P}$ for some $k \geq 1$. We define a null Π_1^0 class $\mathcal{Q}$ containing Z. Let n_0 be least such that $\mathcal{P} = [F]^{\prec}$ for some set of strings of length n_0. Let $n_t = n_0(k+1)^t$ for $t \geq 1$. Let

$$\mathcal{Q} = \bigcap_{t \in \mathbb{N}} \{Y \colon \bigvee_{1 \leq i \leq k} Y_{in_t} \notin \mathcal{P}\}.$$

By definition of n_0 the conditions in the same disjunction are independent, so we have

$$\lambda(\bigvee_{1 \leq i \leq k} Y_{in_t} \notin \mathcal{P}) = 1 - (\lambda \mathcal{P})^k < 1.$$

By definition of the n_t for $t > 0$, the class $\mathcal{Q}$ is the independent intersection of such classes indexed by t. Therefore $\mathcal{Q}$ is null. Clearly $\mathcal{Q}$ is Π_1^0.

By hypothesis $Z \in \mathcal{Q}$. So Z is not weakly random. $\square$

3. Multiple recurrence for Schnorr random sequences

Theorem 3.1. *Let $\mathcal{P} \subseteq 2^{\mathbb{N}}$ be a Π_1^0 class such that $0 < p = \lambda \mathcal{P}$ and p is a computable real. Each Schnorr random Z is multiply recurrent in $\mathcal{P}$.*

We note that this also follows from a particular kind of effective version of Furstenberg multiple recurrence (Cor. 1.2), as explained in Remark 5.3 below. However, we prefer to give a direct proof avoiding Cor. 1.2.

Proof. We extend the previous proof, working with an effective approximation $\mathcal{B} = 2^{\mathbb{N}} - \mathcal{P} = \bigcup_s \mathcal{B}_s$ where the $\mathcal{B}_s$ are clopen. We may assume that $\mathcal{B}_s = [B_s]^{\prec}$ for some effectively given set B_s of strings of length s.

We fix an arbitrary $k \geq 1$ and show that Z is k-recurrent in $\mathcal{P}$. Given $v \in \mathbb{N}$ we will define a null Π_1^0 class $\mathcal{Q}_v \subseteq 2^{\mathbb{N}}$ which plays a role similar to the class $\mathcal{Q}$ before. We also define an "error class" $\mathcal{G}_v \subseteq 2^{\mathbb{N}}$ that is Σ_1^0 uniformly in v. Further, $\lambda \mathcal{G}_v$ is computable uniformly in v and $\lambda \mathcal{G}_v \leq 2^{-v}$, so that $\langle \mathcal{G}_v \rangle_{v \in \mathbb{N}}$ is a Schnorr test. If Z passes this Schnorr test then Z behaves essentially like a weakly random in the proof of Proposition 2.1, which shows that Z is k-recurrent for $\mathcal{P}$.

For the details, given $v \in \mathbb{N}$, we define a computable sequence $\langle n_t \rangle$. Let $n_0 = 1$. Let $n = n_t \geq (k+1)n_{t-1}$ be so large that

$$\lambda(\mathcal{B} - \mathcal{B}_n) \leq 2^{-t-v-k}.$$

3

As in the proof of Proposition 2.1, the class

$$\mathcal{Q}_v = \{Y \colon \forall t \bigvee_{1 \leq i \leq k} Y_{in_t} \in \mathcal{B}_{n_t}\}$$

is Π_1^0 and null. The "error class" for v at stage t is

$$\mathcal{G}_v^t = \{Y \colon \bigvee_{1 \leq i \leq k} Y_{in_t} \in \mathcal{B} - \mathcal{B}_{n_t}\}.$$

Notice that $\lambda \mathcal{G}_v^t \leq k2^{-t-v-k}$, and this measure is computable uniformly in v, t. Let $\mathcal{G}_v = \bigcup_t G_v^t$. Then $\lambda \mathcal{G}_v$ is also uniformly computable in v, and bounded above by 2^{-v}, as required.

If Z is Schnorr random, there is v such that $Z \notin \mathcal{G}_v$. Also, $Z \notin \mathcal{Q}_v$, so that for some t we have $Z_{in_t} \in \mathcal{P}$ for each i with $1 \leq i \leq k$, as required. $\qquad\square$

4. Multiple recurrence for ML-random sequences

For general Π_1^0 classes, the right level of randomness to obtain multiple recurrence is Martin-Löf randomness. We first remind the reader that even the case of 1-recurrence characterizes ML-randomness. This is a well-known result of Kučera [8].

Proposition 4.1. *Z is ML-random $\Leftrightarrow$ Z is 1-recurrent in each Π_1^0 class $\mathcal{P}$ with $0 < p = \lambda \mathcal{P}$.*

Proof. $\Rightarrow$: see e.g. [10, 3.2.24] or [1, 6.10].
$\Leftarrow$: ML-randomness of a sequence Z is preserved by adding bits at the beginning. By the Levin-Schnorr Theorem, the Π_1^0 class $\mathcal{P} = \{Y \colon \forall n K(Y \restriction_n) \geq n - 1\}$ consists entirely of ML-randoms. So, if Z is not ML-random, then no tail of Z is in the Π_1^0 class $\mathcal{P}$. Further, $\lambda \mathcal{P} \geq 1/2$. $\qquad\square$

Theorem 4.2. *Let $\mathcal{P} \subseteq 2^{\mathbb{N}}$ be a Π_1^0 class with $0 < p = \lambda \mathcal{P}$. Each Martin-Löf random Z is multiply recurrent in $\mathcal{P}$.*

Proof. As before we fix an arbitrary $k \geq 1$ in order to show that Z is k-recurrent in $\mathcal{P}$. First we prove the assertion under the additional assumption that $1 - 1/k < p$. This generalises Kučera's argument in '$\Rightarrow$' of the proposition above, where $k = 1$ and the additional assumption $0 < p$ is already satisfied.

Let $B \subseteq 2^{<\omega}$ be a prefix-free c.e. set such that $[B]^{\prec} = 2^{\mathbb{N}} - \mathcal{P}$. We may assume that $B_0 = \emptyset$ and for each $t > 0$, if $\sigma \in B_t - B_{t-1}$ then $|\sigma| = t$. We define a uniformly c.e. sequence $\langle C^r \rangle$ of prefix-free sets which also have the property that at stage t only strings of length t are enumerated.

For a string η and $u \leq |\eta|$, we write $(\eta)_u$ for the string η with the first u bits removed. Let C^0 only contain the empty string, which is enumerated at stage 0. Suppose $r > 0$ and C^{r-1} has been defined. Suppose σ is enumerated in C^{r-1} at stage s (so $|\sigma| = s$). For strings $\eta \succ \sigma$ we search for the failure of k-recurrence in $\mathcal{P}$ that would be obtained by taking s bits off η for k times. At stage $t > (k+1)s$, for each string η of length t such that $\eta \succ \sigma$ and

$$(*) \qquad\qquad \bigvee_{1 \leq i \leq k} (\eta)_{si} \in B_{t-si},$$

and no prefix of η is in C_{t-1}^r, put η into C^r at stage t.

Claim 4.3. *C^r is prefix-free for each r.*

This holds for $r = 0$. For $r > 0$ suppose that $\eta \preceq \eta'$ and both strings are in C^r. Let $t = |\eta|$. By inductive hypothesis the string η was enumerated into C^r via a unique $\sigma \prec \eta$, where $\sigma \in C^{r-1}$. Then $\eta = \eta'$ because we chose the string in C^r minimal under the prefix relation. This establishes the claim.

By hypothesis $1 > q = k\lambda [B]^{\prec}$.

Claim 4.4. *For each $r \geq 0$ we have $\lambda [C^r]^{\prec} \leq q^r$.*

This holds for $r = 0$. Suppose now that $r > 0$. Let $\sigma \in C^{r-1}$. The local measure above σ of strings η, of a length t, such that $\bigvee_{1 \le i \le k} \eta_{si} \in B_{t-is}$ is at most q. The estimate follows by the prefix-freeness of C^r.

If Z is not k-recurrent in $\mathcal{P}$, then $Z \in [C^r]^{\prec}$ for each r, so Z is not ML-random.

We now remove the additional assumption that $1 - 1/k < p$. We define the sets C^r as before. Note that any string in C^r has length at least r. Everything will work except for Claim 4.4: if $\lambda[B]^{\prec} \ge 1/k$ then $\lambda[C^r]^{\prec}$ could be 1. To remedy this, we choose a finite set $D \subseteq B$ such that the set $\widetilde{B} = B - D$ satisfies $\lambda[\widetilde{B}]^{\prec} < 1/k$. Let $N = \max\{|\sigma|: \sigma \in D\}$. We modify the argument of Prop. 2.1, where the clopen set $\mathcal{P}$ there now becomes $2^{\mathbb{N}} - [D]^{\prec}$.

Let $C = \bigcup_r C^r$. Let G_m be the set of prefix-minimal strings η such that $\eta \in C$, and there exist m many $s > N$ as follows.

- $\eta \restriction_s \in C$, and
- for some i with $1 \le i \le k$, $\eta \restriction_{[si,s(i+1))}$ extends a string in D.

(Informally speaking, if there are arbitrarily long such strings along Z, then the attempted test $[C^r]^{\prec}$ might not work, because the relevant "block" $\eta \restriction_{[si,s(i+1))}$ may extend a string in D, rather than one in $\widetilde{B}$.)

The sets G_m are uniformly Σ_1^0. By choice of N and independence, as in the proof of Prop. 2.1 we have $\lambda[G_{m+1}]^{\prec} \le (1 - v^k)\lambda G_m$, where $v = \lambda(2^{\mathbb{N}} - [D]^{\prec})$. If Z is ML-random we can choose a least m^* such that $Z \notin [G_{m^*}]^{\prec}$.

Note that $m^* > 0$ since $G_0 = \{\emptyset\}$. So choose $\rho \prec Z$ such that $\rho \in G_{m^*-1}$. Then $\rho \in C^r$ for some r, and no τ with $\rho \preceq \tau \prec Z$ is in G_{m^*}.

We define a ML-test that succeeds on Z. Let $\widetilde{C}^r = C^r$. Suppose $u > r$ and $\widetilde{C}^{u-1}$ has been defined. For each $\sigma \in \widetilde{C}^{u-1}$, put into $\widetilde{C}^u$ all the strings $\eta \succ \sigma$ in C^u so that $(*)$ can be strengthened to $\bigvee_{1 \le i \le k}(\eta)_{si} \in \widetilde{B}_{t-is}$, where $s = |\sigma|$.

Let $q = k\lambda[\widetilde{B}]^{\prec}$. Note that $\lambda[\widetilde{C}^u]^{\prec} \le q^u$ as before. By the choice of m^* we have $Z \in \bigcap_{u \ge r}[\widetilde{C}^u]^{\prec}$, so since $q < 1$, an appropriate refinement of the sequence of open sets $\left\langle [\widetilde{C}^u]^{\prec} \right\rangle_{u \in \mathbb{N}}$ shows Z is not ML-random.

$\square$

5. Towards the general case

5.1. Recurrence for k shift operators.

The probability space under consideration is now $\mathcal{X} = \{0,1\}^{\mathbb{N}^k}$ with the product measure. For $1 \le i \le k$, the operator $T_i \colon \mathcal{X} \to \mathcal{X}$ takes one "face" of bits off in direction i. That is, for $Z \in \mathcal{X}$,

$$T_i(Z)(u_1, \ldots, u_k) = Z(u_1, \ldots, u_i + 1, \ldots, u_k).$$

Z is *multiply recurrent* in a class $\mathcal{P} \subseteq \mathcal{X}$ if $[Z \in \bigcap_{i \le k} T_i^{-n}(\mathcal{P})$ for some n.

Algorithmic randomness notions for points in $\mathcal{X}$ can be defined via the effective measure preserving isomorphism $\mathcal{X} \to 2^{\mathbb{N}}$ given by a computable bijection $\mathbb{N}^k \to \mathbb{N}$. Modifying the methods above, we show the following.

Theorem 5.1. *Let $\mathcal{P} \subseteq \mathcal{X}$ be a Π_1^0 class with $0 < p = \lambda\mathcal{P}$. Let $Z \in \mathcal{X}$.*

If Z is (a) Kurtz (b) Schnorr (c) ML-random, then Z is multiply recurrent in $\mathcal{P}$

in case (a) $\mathcal{P}$ is clopen (b) $\lambda\mathcal{P}$ is computable (c) for any $\mathcal{P}$.

Proof. For the duration of this proof, by an *array* we mean a map $\sigma \colon \{0, \ldots, n-1\}^k \to \{0,1\}$. We call n the *size* of σ and write $n = |\sigma|$. The letters σ, τ, ρ, η now denote arrays. If σ is an array of size n, for $s \le n$ and $i \le k$ let $(\sigma)_{i,s}$ be the array τ of size $n - s$ such that

$$\tau(u_1, \ldots, u_k) = \sigma(u_1, \ldots, u_i + s, \ldots, u_k)$$

for $u_1, \ldots, u_k \le n - s$. This operation removes s faces in direction i, and then in the remaining directions cuts the faces from the opposite side in order to obtain an array. For a set S of arrays we define $[S]^{\prec} = \{Y \in \mathcal{X} \colon \exists \sigma \in S\, [\sigma \prec Y]\}$ where the "prefix" relation $\prec$ is defined as expected.

Suppose that Z is not k-recurrent in $\mathcal{P}$ for some $k \geq 1$.

(a). As in Prop. 2.1 we define a null Π^0_1 class $\mathcal{Q} \subseteq \mathcal{X}$ containing Z. Let n_1 be least such that $\mathcal{P} = [F]^\prec$ for some set of arrays of that all have size n_1. Let

$$\mathcal{Q} = \bigcap_{r \geq 1}\{Y \colon \bigvee_{1 \leq i \leq k} T_i^{rn}(Y) \notin \mathcal{P}\}.$$

By the choice of n_1 the conditions in the same disjunction are independent, so we have

$$\lambda(\bigvee_{1 \leq i \leq k} T_i^{rn}(Y) \notin \mathcal{P}) = 1 - p^k < 1.$$

The Π^0_1 class $\mathcal{Q}$ is the independent intersection of such classes indexed by r. Therefore $\mathcal{Q}$ is null. By hypothesis $Z \in \mathcal{Q}$. So Z is not weakly random.

(b). We could modify the previous argument. However, this also follows by the general fact in Remark 5.3 below.

(c). The argument is similar to the proof of Theorem 4.2 above. The definition of the c.e. set B and its enumeration are as before, except that each string of length n is now an array of size n. In particular, an array enumerated at a stage s has size s.

Let C^0 only contain the empty array, which is enumerated at stage 0. Suppose $r > 0$ and C^{r-1} has been defined. Suppose σ is enumerated in C^{r-1} at stage s (so $|\sigma| = s$).

At a stage $t > 2s$, for each array η of size t such that $\eta \succ \sigma$ and

$$(*) \qquad \qquad \bigvee_{1 \leq i \leq k} (\eta)_{i,s} \in B_{t-s},$$

and no array that is a prefix of η is in C^r_{t-1}, put η into C^r at stage t. As before one checks that C^r is prefix-free for each r.

Choose a finite set of arrays $D \subseteq B$ such that the set $\widetilde{B} = B - D$ satisfies $\lambda[\widetilde{B}]^\prec < 1/k$. Let $N = \max\{|\sigma| \colon \sigma \in D\}$. Let $C = \bigcup_r C^r$. Let G_m be the set of prefix-minimal arrays η such that $\eta \in C$, and there exist m many $s > N$ as follows.

- $\eta \restriction_{\{0,\dots,s-1\}^k} \in C$, and
- for some i with $1 \leq i \leq k$, $(\eta)_{i,s}$ extends an array in D.

The sets G_m are uniformly Σ^0_1. By choice of N and independence $\lambda[G_{m+1}]^\prec \leq (1 - t^k)\lambda G_m$, where $t = \lambda(2^{\mathbb{N}} - [D]^\prec)$. If Z is ML-random we can choose a least m^* be such that $Z \notin [G_{m^*}]^\prec$, and $m^* > 0$ since $G_0 = \{\emptyset\}$. So choose $\eta \prec Z$ such that $\eta \in G_{m^*-1}$. Then $\eta \in C^r$ for some r, and no τ with $\eta \preceq \tau \prec Z$ is in G_{m^*}.

Let $\widetilde{C}^r = C^r$. Suppose $u > r$ and $\widetilde{C}^{u-1}$ has been defined. For each $\sigma \in \widetilde{C}^{u-1}$, put into $\widetilde{C}^u$ all the arrays $\eta \succ \sigma$ in C^u so that $(*)$ can be strengthened to $\bigvee_{1 \leq i \leq k}(\eta)_{i,s} \in \widetilde{B}_{t-s}$, where $s = |\sigma|$ and $t = |\eta|$.

Let $q = k\lambda[\widetilde{B}]^\prec$. Then $\lambda[\widetilde{C}^u]^\prec \leq q^u$ as before. By the choice of m^* we have $Z \in \bigcap_{u \geq r}[\widetilde{C}^u]^\prec$, so since $q < 1$, Z is not ML-random. □

5.2. The putative full result.

It is likely that a multiple recurrence theorem holds in greater generality. For background on computable probability spaces and how to define randomness notions for points in them, see e.g. [5].

Conjecture 5.2. Let (X, μ) be a computable probability space. Let $T_1, \dots, T_k$ be computable measure preserving transformations that commute pairwise. Let $\mathcal{P}$ be a Π^0_1 class with $\mu P > 0$.

If $z \in \mathcal{P}$ is ML-random then $\exists n[z \in \bigcap_{i \leq k} T_i^{-n}(\mathcal{P})]$.

Remark 5.3. Let U_n be the open set $\{x \colon x \notin \bigcap_{i \leq k} T_i^{-n}(\mathcal{P})\}$. Then $\mu(\mathcal{P} \cap \bigcap_n U_n) = 0$ by the classic multiple recurrence theorem in the version of Cor. 1.2. Since $\mathcal{P} \cap \bigcap_n U_n$ is Π^0_2, weak 2-randomness of z suffices for the k-recurrence.

Jason Rute has pointed out that if X is Cantor space and $\mu \mathcal{P}$ is computable, then $\exists n[z \in \bigcap_{i \leq k} T_i^{-n}(\mathcal{P})]$ for every Schnorr random $z \in \mathcal{P}$. For in this case $\mu \widehat{U}_n$ is uniformly computable

where $\widehat{U}_n = \bigcap_{i \leq n} U_i$. Let $\mathcal{P} = \bigcap_n \mathcal{P}_n$ where the $\mathcal{P}_n$ are clopen sets computed uniformly in n. Let $G_n = \mathcal{P}_n \cap \widehat{U}_n$. Then G_n is uniformly Σ_1^0 and $\mu(G_n)$ is uniformly computable. Refining the sequence $\langle G_n \rangle$ we obtain a Schnorr test capturing z.

6. Effective Multiple Recurrence for Irrational Rotations

In the foregoing sections, we have seen examples of mixing systems that display effective versions of Furstenberg multiple recurrence. At the other end of the spectrum, highly structured systems also exhibit multiple recurrence. Such systems are called compact. An example is irrational rotations of the unit circle. rotations. We discuss a fact from [4] which implies that *every* point in such a system is multiply recurrent with respect to every of its neighbourhoods, rather than merely the Kurtz random points. To establish this result, we briefly look at topological dynamical systems. Such a system consists of a compact space X and a continuous operator $T: X \to X$.

Definition 6.1 ([4], Def. 1.1). A point x in a topological dynamical system (X, T) is *k-recurrent* if the condition of Definition 1.4 holds for each neighbourhood $\mathcal{P}$ of x.

Definition 6.2. Given a compact group G and $a \in G$, let $T_a(x) = a \cdot x$. One calls (G, T_a) a *Kronecker system*.

For instance, for $\alpha \in \mathbb{R}$, the system $(\mathbb{R}/\mathbb{Z}, T)$ where $T(x) = \alpha + x \mod 1$ is a Kronecker system. There is a unique invariant probability measure, called the Haar measure, on a Kronecker system. Hence such a system can also be viewed as a measure-preserving system, which turns out to be compact in the measure theoretic sense. It is known that every compact *ergodic* system is equivalent to a Kronecker system in the sense that the two systems are isomorphic when viewed on the σ-algebra of measurable sets mod null sets.

Lemma 6.3. *(see [4], Chapter 1) Every point in a Kronecker system is 1-recurrent.*

This is established by showing that there is some recurrent point $x_0 \in G$, by first considering a minimal subsystem consisting of points with dense orbits, and by applying the Zorn's lemma. Since G is a group, if there is any recurrent point in the system, then every point must be recurrent.

We can use the lemma itself to strengthen it.

Lemma 6.4. *Every point in a Kronecker system (G, T_a) is multiply recurrent.*

Proof. Given an integer $k \geq 2$, consider the tuple $t = (a, a^2, \dots, a^k)$ in the compact group $H = G^k$. The system (H, t) is also Kronecker. Every point is 1-recurrent in this system. In particular, for every $x \in G$, the point $y = (x, \dots, x) \in H$ is 1-recurrent.

If V is an open nbhd of x, then the cartesian power ${}^k V$ is an open nbhd of y. So $y \cdot t^n \in {}^k V$ for some n. This means that $x a^{in} \in V$ for each $i \leq k$, as required. $\qquad\square$

Let α be a computable irrational. Brown Westrick (see [9]) has proved that every ML-random point in $[0, 1]$ is multiply recurrent in each Π_1^0 class of positive measure for at least one of the operators $x \to (x + \alpha) \mod 1$ or $x \to (x - \alpha) \mod 1$. Note that this lends some further evidence to Conjecture 5.2. However, even for Kronecker systems with computable group structure and computable Haar measure, the conjecture is open.

References

[1] R. Downey and D. Hirschfeldt. *Algorithmic randomness and complexity.* Springer-Verlag, Berlin, 2010. 855 pages.
[2] L. Bienvenu, A. Day, M. Hoyrup, I. Mezhirov, and A. Shen. A constructive version of Birkhoff's ergodic theorem for Martin-Löf random points. *Information and Computation*, 210, 21-30, 2012.
[3] J. Franklin, N. Greenberg, J. Miller, and Keng Meng Ng. Martin-löf random points satisfy Birkhoff's ergodic theorem for effectively closed sets. *Proceedings of the American Mathematical Society*, 140:3623–3628, 2012.
[4] H. Furstenberg. *Recurrence in ergodic theory and combinatorial number theory.* Princeton University Press, 2014.

[5] P. Gács, M. Hoyrup, and C. Rojas. Randomness on computable probability spaces - a dynamical point of view. *Theory Comput. Syst.*, 48(3):465–485, 2011.

[6] M. Hochman. Upcrossing inequalities for stationary sequences and applications. *Annals of Probability*, 37(6):2135–2149, 2009.

[7] M. Hoyrup. The dimension of ergodic random sequences. In *Symposium on Theoretical Aspects of Computer Science*, pages 567–576, 2012.

[8] A. Kučera. Measure, Π_1^0-classes and complete extensions of PA. In *Recursion theory week (Oberwolfach, 1984)*, volume 1141 of *Lecture Notes in Math.*, pages 245–259. Springer, Berlin, 1985.

[9] A. Nies (editor). Logic Blog 2016. Available at `cs.auckland.ac.nz/~nies`, 2016.

[10] A. Nies. *Computability and randomness*, volume 51 of *Oxford Logic Guides*. Oxford University Press, Oxford, 2009. 444 pages. Paperback version 2011.

[11] M. van Lambalgen. *Random Sequences*. Academish Proefschri't, Amsterdam, 1987.

[12] V. G. Vovk. The law of the iterated logarithm for random Kolmogorov or stochastic sequences. *SIAM Theory of Probability and Applications*, 32(3):413–425, 1987.

[13] V. V. V'yugin. Effective convergence in probability and an ergodic theorem for individual random sequences. *Theory of Probability and Its Applications*, 42(1):39–50, 1997.

DEPARTMENT OF MATHEMATICS, STATISTICS AND OPERATIONS RESEARCH, VICTORIA UNIVERSITY OF WELLINGTON, WELLINGTON, NEW ZEALAND

E-mail address: `rod.downey@msor.vuw.ac.nz`

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING, INDIAN INSTITUTE OF TECHNOLOGY KANPUR, KANPUR, UTTAR PRADESH, INDIA.

E-mail address: `satyadev@cse.iitk.ac.in`

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF AUCKLAND, AUCKLAND, NEW ZEALAND

E-mail address: `andre@cs.auckland.ac.nz`