Ergodic Theory in Combinatorics

Talk by Florian Richter

Notes by Sanjana Das

February 10, 2023

§1 Introduction

Definition 1.1. The *upper density* of a set $E \subseteq \mathbb{N}$ is defined as

$$\overline{d}(E) = \limsup_{n \to \infty} \frac{|E \cap \{1, \ldots, n\}|}{n}.$$

Example 1.2

We have $\overline{d}(\mathbb{N}) = 1$, $\overline{d}(2\mathbb{N}+1) = \frac{1}{2}$, $\overline{d}(\text{squarefree integers}) = \frac{6}{\pi^2}$, and $\overline{d}(\mathbb{P}) = 0$.

Think of sets with positive density as 'large.' In Ramsey theory, large sets contain a lot of structure; we'll look at what structure we can find in positive-density subsets of $\mathbb{N}$.

Theorem 1.3 (Roth 1953)

If $\overline{d}(E) > 0$, then E contains a 3-term arithmetic progression $\{a, a+d, a+2d\}$.

Roth's theorem was a special case of a more general conjecture. In Ramsey theory, van der Waerden's theorem states that if you partition $\mathbb{N}$ into finitely man classes, one class must contain arbitrarily long arithmetic progressions; but it does not state which one. Erdős and Turán conjectured that any set with positive density contains arbitrarily long arithmetic progressions. This was eventually proved by Szemerédi.

Theorem 1.4 (Szemeredi 1975)

If $\overline{d}(E) > 0$, then E contains a k-term arithmetic progression for all $k \geq 1$.

Question 1.5. What other arrangements do sets of positive density contain?

Theorem 1.6 (Furstenberg–Sárközy 1978)

If $\overline{d}(E) > 0$, then E contains two elements a and b such that $b - a$ is a perfect square.

Furstenberg proved this result, as well as Szemerédi's theorem, using ergodic theory.

§2 Intersective Sets

The underlying question here is the following:

Question 2.1. If $\overline{d}(E) > 0$, then for how many integers do we have $E \cap (E - n) \neq \emptyset$?

The Furstenberg–Sárközy theorem states that this is true for some perfect square n.

Definition 2.2. A set $R \subseteq \mathbb{N}$ is *intersective* if for all sets $E \subseteq \mathbb{N}$ with $\overline{d}(E) > 0$, there exists $n \in R$ with $E \cap (E - n) \neq \emptyset$.

Equivalently, R is intersective if $R \cap (E - E) \neq \emptyset$ for all sets E of positive density.

Example 2.3

Some examples and non-examples of intersective sets.

Intersective	Non-intersective
$\mathbb{N}$	Any finite set
$k\mathbb{N}$	$2\mathbb{N} + 1$
$\{n^2 \mid n \in \mathbb{N}\}$ (by the Furstenberg–Sárközy theorem)	$\{n^2 + 1 \mid n \in \mathbb{N}\}$
$\{p(n) \mid n \in \mathbb{N}\}$ for polynomials p with a root mod every prime	
$\mathbb{P} - 1$	$\mathbb{P}$
$D - D$ for any infinite $D \subseteq \mathbb{N}$	Any lacunary set

For example, $k\mathbb{N}$ is intersective because any infinite set E contains two numbers of the same residue class mod k.

For polynomials, if $\{p(n)\}$ is intersective then p must have a root mod every prime (otherwise we could take E to be an arithmetic progression); the converse is true as well (this is not obvious).

In particular, although lacunary sets (ones which grow exponentially when enumerated in order) are never intersective, there do exist very sparse sets that are intersective ($D - D$ is intersective; it is not lacunary, as it has bounded gaps).

Remark 2.4. We don't have a complete characterization of intersective sets — for example, we don't know whether $\lfloor e^{\sqrt{n}} \rfloor$ is intersective. This is true even for sets with polynomial growth, although we do know that $\lfloor n \log n \rfloor$, $\lfloor n^c \rfloor$, and $\lfloor n\beta \rfloor$ are all intersective.

Fact 2.5 — If R is intersective, then:

1. $R \cap b\mathbb{N} \neq \emptyset$ for all $b \in \mathbb{N}$;
2. If $R = R_1 \cup R_2$, then at least one R_i is intersective. (In other words, being intersective is partition-regular.)

Remark 2.6. Intersective also generalizes to k-fold intersective. These are not equivalent — the set $\{n \mid n^2\alpha \mod 1 \in (1/2, 1/2+\varepsilon)\}$ is intersective, but is not double intersective because it is not centered. It also generalizes to *quadratic-intersective* — a set R is quadratic-intersective if for all $\overline{d}(E) > 0$ there exists $n \in R$ with $E \cap (E - n^2) \neq \emptyset$. There is an open conjecture that quadratic-intersective and 2-intersective are equivalent.

§3 Measure-Preserving Systems

Shifting variables plays well with density — if a set has density $\frac{1}{2}$ and we shift it by 1, then it still has density $\frac{1}{2}$. So we can think of $(\mathbb{N}, \overline{d})$ as a fake probability space, with the fake probability measure preserved under shifting.

Furstenberg's idea was to turn this into a measure-preserving system.

Definition 3.1. A *measure-preserving system* is a triple (X, μ, T) where:

- X is a compact metric space;
- μ is a Borel probability measure on X;
- T is a transformation $X \to X$ which preserves μ, meaning that $T\mu = \mu$.

($T\mu$ is the *pushforward* measure; the statement $T\mu = \mu$ means that $\mu(T^{-1}A) = \mu(A)$ for all A.)

Instead of just considering T, you can also look at its powers — $T^{n+m} = T^n \circ T^m$, so T gives an action of $(\mathbb{Z}, +)$. To understand the arithmetic structure of the integers, you can consider their actions; these dynamics can be understood through ergodic theory, which can be used to say more about arithmetics.

Example 3.2

One measure-preserving system is to take X to be the torus $\Pi = \mathbb{R}/\mathbb{Z}$, μ the Lebesgue measure, and $T(x) = x + \alpha \mod 1$ (i.e., a rotation).

Another is to again take X to be the torus and μ the Lebesgue measure, and $T(x) = 2x \mod 1$.

(Note that $x \mapsto 2x \mod 1$ is a measure-preserving transformation because $\mu(T^{-1}A) = \mu(A)$; however, it is *not* true that $\mu(TA) = \mu(A)$.)

These examples have very different behavior — a rotation is very deterministic, while $2x \mod 1$ is chaotic. The dichotomy between structure and randomness is often useful (for example, Szemerédi's regularity lemma). The idea in dynamical systems and ergodic theory is similar — we take a system and decompose it into things which look like the first example, and things which look like the second.

One of the most basic results in ergodic theory is the following.

Theorem 3.3 (Poincaré Recurrence)

If (X, μ, T) is a measure-preserbing system, then for every $A \subseteq X$ with $\mu(A) > 0$, there is an integer $n \in \mathbb{N}$ such that $\mu(A \cap T^{-n}A) > 0$.

This says that the orbit of any set must eventually come back to itself.

Question 3.4. What can we say about the sets of integers n that make this happen — how large or how structured are they? (In terms of dynamical systems, around what times can we come back to our initial position?)

Definition 3.5. A set $R \subseteq \mathbb{N}$ is a *set of recurrence* if for all (X, μ, T) and all subsets $A \subseteq X$ with $\mu(A) > 0$, there is some $n \in \mathbb{R}$ with $\mu(A \cap T^{-n}A) > 0$.

This definition feels very similar to the one of an intersective set, and in fact Furstenberg showed that these two notions are the same.

Theorem 3.6 (Furstenberg)

A set R is intersective if and only if R is a set of recurrence.

Then the Furstenberg–Sárközy theorem states that $\{n^2 \mid n \in \mathbb{N}\}$ is intersective; to prove this, it's enough to prove that it's a set of recurrence. This is how Furstenberg proved the theorem.

Theorem 3.7 (Furstenberg's Quadratic Recurrence Theorem 1978)

The set $\{n^2 \mid n \in \mathbb{N}\}$ is a set of recurrence.

Furstenberg also used similar ideas to prove Szemerédi's theorem (by establishing a higher-order version of Poincare's conjecture with multiple intersections).

Theorem 3.8 (Furstenberg's Multiple Recurrence Theorem 1978)

For any $A \subseteq X$ with $\mu(A) > 0$ and any $k \in \mathbb{N}$, there is some integer n such that

$$\mu(A \cap T^{-n}A \cap T^{-2n}A \cap \cdots \cap T^{-kn}A) > 0.$$

Using the same ideas, this can be used to show that then any set of positive integers with positive density contains a $(k+1)$-term arithmetic progression.

Today we will see two proofs — of the equivalence between being intersective and being a set of recurrence, and if time permits, a sketch of the proof that $\{n^2 \mid n \in \mathbb{N}\}$ is a set of recurrence.

§4 Furstenberg's Correspondence

Theorem 4.1

For any set $E \subseteq \mathbb{N}$ there is a dynamical system (X, μ, T) and a subset $A \subseteq X$ with $\mu(A) = \overline{d}(E)$, such that for all $n_1, \ldots, n_k \in \mathbb{N}$,

$$\mu(A \cap T^{-n_1}A \cap \cdots \cap T^{-n_k}A) \leq \overline{d}(E \cap (E - n_1) \cap \cdots \cap (E - n_k)).$$

(In particular, if E has positive density, then so does A.) This means we can model the outer correlations of E by the outer correlations of the dynamical system. So in some sense, this is a way to turn the integers into an actual probability space.

Proof. We will take X to be $\{0,1\}^{\mathbb{N}\cup\{0\}}$ (this is a compact metric space). Identify E with its indicator function $1_E \in \{0,1\}^{\mathbb{N}\cup\{0\}}$.

We will take T to be the shift operator $(x_n)_{n\in\mathbb{N}} \to (x_{n+1})_{n\in\mathbb{N}}$ (i.e., we erase the first letter and shift the rest of the sequence to the left).

Our set A will be the set of all sequences $(x_n)_{n\in\mathbb{N}}$ with $x_0 = 1$; this is a clopen subset of the compact metric space.

Finally, we need to construct a measure. We first define the *finite* measures

$$\mu_N = \frac{1}{N}\sum_{n=1}^{N} \delta_{T^{-n}1_E}.$$

(We're placing a point mass on each of the points in the orbit of 1_E — the points 1_E, 1_{E-1}, 1_{E-2}, and so on — and averaging them.)

Choose a sequence (N_k) such that

$$\overline{d}(E) = \lim_{k\to\infty} \frac{1}{N_k}\sum_{n=1}^{N_k} 1_E(n) = \lim_{k\to\infty} \frac{1}{N_k}\sum_{n=1}^{N} 1_{E-n}(0).$$

(Such a sequence exists because $\overline{d}(E)$ is defined as the lim sup of this quantity over all N, and there must be a subsequence approaching the lim sup.)

Endowed with the weak-$*$ topology, this is a compact metric space, so every sequence has a convergent subsequence. Let μ be any weak-$*$ accumulation point of the sequence of measures μ_{N_k}. (A sequence of measures ν_n cconverges to ν in the weak-$*$ topology if and only if for all continuous functions f we have $\int f\, d\nu_n \to \int f\, d\nu$.)

We now have our measure μ. The fact that $\mu(A) = \overline{d}(E)$ follows from the above relation and the definition of A (since $\delta_{1_{E-n}}(A)$ is 1 if 1_{E-n} begins with a 1, or equivalently if $E - n$ contains 0, and is 0 otherwise). The same reasoning can be used to show $\mu(A\cap T^{-n_1}A\cap\cdots\cap T^{-n_k}A) \le \overline{d}(E\cap(E-n_1)\cap\cdots\cap(E-n_k))$. □

This proves one direction of Furstenberg's theorem — that if R is a set of recurrence, then it is intersective. (This is the direction that we need; we will not prove the other direction.)

§5 The Furstenberg–Sárközy Theorem

We have the Hilbert space $L^2(X,\mu)$ and a measure-preserving transformation T, so we can define an operator $U_T\colon L^2 \to L^2$ as $U_T f = f\circ T$.

Since T is measure-preserving, we have $\mu(T^{-1}A) = \mu(A)$. This means $\int U_T f\, d\mu = \int f\, d\mu$, so $\langle U_T f, U_T g\rangle = \langle f, g\rangle$. So U_T is a unitary operator. In particular, since unitary operators are isometric, we have $\|U_T f\|_{L^2} = \|f\|_{L^2}$.

The main theorem in ergodic theory is the following.

Theorem 5.1 (von Neumann's Ergodic Theorem)
If (X,μ,T) is a measure-preserving system and $f\in L^2$ is any function, then

$$\frac{1}{N}\sum_{n=1}^{N} U_T^n f \to Pf,$$

where P is an orthogonal projection onto $\mathcal{H}_{\text{inv}} = \{f\in L^2 \mid U^T f = f\}$.

If X is a finite set, this corresponds to convergence to the 1-eigenspace — there, the point is that any eigenvalue that is not 1 will average out to 0. Here something similar happens — if f is invariant under U_T then this average converges to f itself, while if f is orthogonal to the invariant space then it converges to 0.

Now we will use this to prove that $\{n^2 \mid n\in\mathbb{N}\}$ is a set of recurrence. We will first prove this for the specific system corresponding to rotation on a circle (as in the initial example), before seeing the proof in the abstract case.

Theorem 5.2
The set $\{n^2\alpha \mod 1\}$ (for fixed irrational α) is dense in Π.

This is an analog of the Furstenberg–Sárközy theorem for this specific system. This problem was open for decades (if α is irrational then it's known $\{n\alpha \mod 1\}$ is dense; it's natural to ask if the same is true for n^2), and was eventually solved by Hardy. A couple of years later it was proven by Weyl, who came up with the notion of a *uniform distribution* — he showed that it's not just dense, but uniformly distributed. Uniform distribution can be quantified in the following ways.

Definition 5.3. A sequence $(x_n)_{n\in\mathbb{N}}$ is *uniformly distributed* in Π if one of the following equivalent conditions holds:

1. For all intervals $[a, b] \subseteq \Pi$, we have $\frac{1}{N}\#\{1 \le n \le N \mid x_n \in [a, b]\} \to b - a$.
2. For all continuous functions f on Π, $\frac{1}{N}\sum_{n=1}^{N} f(x_n) \to \int f\,dx$.
3. The above is true for all *exponential* functions, i.e., $\frac{1}{N}\sum_{n=1}^{N} e(hx_n) \to 0$ for all nonzero integers h (where $e(x) = e^{2\pi i x}$).

Theorem 5.4

The set $\{n^2\alpha \mod 1\}$ is uniformly distributed.

Proof. We will check the third condition — we want to show that

$$\frac{1}{N}\sum_{n=1}^{N} e(hn^2\alpha) \to 0$$

for all nonzero h. We can assume $h = 1$, by absorbing it into α otherwise.

Note that the above average is shift-invariant (as $N \to \infty$), so we can consider $\frac{1}{N}\sum_{n=1}^{N} e((n+1)^2\alpha)$ instead, and more generally we can shift by any fixed finite number h. We can then average over these values of h, so it suffices to show that

$$\frac{1}{N}\sum_{n=1}^{N}\frac{1}{H}\sum_{h=1}^{H} e((n+h)^2\alpha) \to 0.$$

(Think of N as much larger than H — we are taking $N \to \infty$ first.)

We will actually show that the *square* of the above expression goes to 0. By Cauchy–Schwarz, this square is at most

$$\frac{1}{N}\sum_{n=1}^{n}\left|\frac{1}{H}\sum_{h=1}^{H} e((n+h)^2\alpha)\right|^2.$$

Now using $|z|^2 = z\overline{z}$, we can expand this out as

$$\frac{1}{H^2}\sum_{h_1,h_2}\frac{1}{N}\sum_{n} e((n+h_1)^2\alpha - (n+h_2)^2\alpha).$$

The expression inside the exponent is $2n(h_1 - h_2)\alpha$ plus a constant, which is linear; this means we have a geometric series, which must go to 0, so we are done. □

Van der Corput realized you can do this with things other than squares — if you have a sequence (x_n) of complex numbers with $\frac{1}{N}\sum_{n=1}^{N} x_{n+h}\overline{x_n} \to 0$ for all $h \in \mathbb{N}$, then $\frac{1}{N}\sum_{n=1}^{N} x_n \to 0$ as well. So if all outer correlations of a sequence go to 0, the sequence does as well.

In this proof, we used basic properties of averages and Cauchy–Schwarz. So it's reasonable to think that the generalization of Van der Corput holds not just for $\mathbb{C}$, but in any place where we have Cauchy–Schwarz. In particular, the same thing works for any Hilbert space.

Theorem 5.5

If $(x_n) \subseteq L^2$ and $\frac{1}{N}\sum_{n=1}^N \langle x_{n+h}, x_n\rangle \to 0$ for all $h \in \mathbb{N}$, then $\left\|\frac{1}{N}\sum_{n=1}^N x_n\right\|_{L^2} \to 0$.

Now using von Neumann's Ergodic Theorem, we can split the space of functions into invariant functions, and functions whose averages go to 0, as

$$L^2(X,\mu) = \{f \in L^2 \mid U_T f = f\} \oplus \left\{f \in L^2 \mid \left|\frac{1}{N}\sum_{n=1}^N U_T^n f\right| \to 0\right\}.$$

But if T is measure-preserving, so is T^2, and T^3, and so on. So we get an infinite number of such decompositions; taking the direct sum of the sets on the left and the intersection of the sets on the right gives the decomposition

$$L^2(X,\mu) = \{f \in L^2 \mid U_T^i f = f \text{ for some } i \in \mathbb{N}\} \oplus \left\{f \in L^2 \mid \left|\frac{1}{N}\sum_{n=1}^N U_T^{in} f\right| \to 0 \text{ for all } i \in \mathbb{N}\right\}.$$

This means we can write any function f as $f_1 + f_2$, where f_1 is in the first set and f_2 in the second. This corresponds to splitting the dynamical system into things that kind of look like progressions, and things that look the opposite.

Proof of Furstenberg–Sárközy. We want to show that $\mu(A \cap T^{-n^2}A) > 0$ for some n. Instead of showing this for an explicit n (which may be hard to find), we will take an average and show that the average is positive — so we instead consider

$$\frac{1}{N}\sum_{n=1}^N \mu(A \cap T^{-n^2}A) = \frac{1}{N}\sum_{n=1}^N \langle U_T^{n^2} 1_A, 1_A\rangle.$$

Now splitting $1_A = f_1 + f_2$, this sum becomes

$$\frac{1}{N}\sum_{n=1}^N \langle U_T^{n^2} f_1, f_1\rangle + \frac{1}{N}\sum_{n=1}^N \langle U_T^{n^2} f_2, f_2\rangle.$$

In reality there are four terms, not 2; but f_1 and f_2 are from orthogonal subspaces which are both invariant under U_T, so the cross-terms must be 0 by orthogonality. So something really nice has happened, which is the idea behind Furstenberg's approach — we've isolated two completely different behaviors (completely rational and completely irrational), which we can think of as a structured and a random component. We will show that the random component is 0 and the structured component is positive, which suffices.

First, for the random component, we want to show $\frac{1}{N}\sum \langle U_T^{n^2} f_2, f_2\rangle \to 0$. To do this, we apply Van der Corput — taking $x_n = U_T^{-n^2} f_2$, it is enough to show that $\frac{1}{N}\sum_{n=1}^N \langle x_{n+h}, x_n\rangle \to 0$. We have

$$\langle x_{n+h}, x_n\rangle = \langle U_T^{(n+h)^2} f_2, U_T^{n^2} f_2\rangle,$$

and since U_T is unitary, we can move all copies of it to one side and cancel to get that this equals

$$\langle U_T^{(2nh+h^2)} f_2, f_2\rangle.$$

Since f_2 belongs to the space with $\frac{1}{N}\sum_n U_T^{nh} f \to 0$, this term (averaged over n) must go to 0, as desired.

This means $\left\|\frac{1}{N}\sum_{n=1}^N U_T^{n^2} f_2\right\| \to 0$, so our expression also goes to 0 (by Cauchy–Schwarz).

Next, we want to show that the term from f_1 is positive. First, all terms in the sum are nonnegative by properties of orthogonal projections, as f_1 is the projection of 1_A onto this space: if f_1 is invariant under U_T^i, then the function $\tilde{f}_1 = \max\{f_1, 0\}$ where we cut away all negative parts is also invariant. But the orthogonal projection is the function in the relevant space closest to the original function. If f_1 is not nonnegative, then

$\tilde{f}_1$ is strictly closer to 1_A than f_1 was, which is a contradiction; so we must have $\tilde{f}_1 = f_1$, and f_1 must be nonnegative.

This means we can ignore all terms except for the multiples of i; so our expression is at least

$$\frac{1}{N}\sum_{n=1}^{N/i}\langle U_T^{(in)^2} f_1, f_1\rangle \geq \frac{1}{i}\left\lVert f_1\right\rVert_{L^2}^2.$$

Finally, the fact that $\lVert f_1\rVert_{L^2} > 0$ follows from the assumption that $\mu(A) > 0$ — we have

$$0 < \mu(A) = \langle 1_A, 1\rangle,$$

and since 1 (the constant function) is in the invariant space, we can write this as $\langle 1_A, P1\rangle$, where P is the orthogonal projection. But by a property of orthogonal projections, this is then equal to $\langle P1_A, P1\rangle = \langle f_1, 1\rangle$. So f_1 cannot be the zero function, and must hvae nonzero norm. $\square$