CENTRALIZERS IN THE GROUP OF INTERVAL EXCHANGE TRANSFORMATIONS

BY

DANIEL BERNAZZANI

Department of Mathematics, Rice University, Houston, TX 77005, USA
e-mail: dbernazz@gmail.com

ABSTRACT

We study the group of interval exchange transformations. Let T be an m-interval exchange transformation. By the rank of T we mean the dimension of the $\mathbb{Q}$-vector space spanned by the lengths of the exchanged subintervals. We prove that if T satisfies Keane's infinite distinct orbit condition and $\operatorname{rank}(T) > 1 + \lfloor m/2 \rfloor$, then the only interval exchange transformations which commute with T are its powers.

In the case that T is a minimal 3-interval exchange transformation, we prove a more precise result: T has a trivial centralizer in the group of interval exchange transformations if and only if T satisfies the infinite distinct orbit condition.

1. Introduction

An interval exchange transformation (IET) is a bijective map $T : [0, 1) \to [0, 1)$ defined by partitioning the unit interval $[0, 1)$ into finitely many subintervals and then rearranging these subintervals by translations. The formal definition appears below. The permutation group of the set $\{1, 2, \dots, m\}$ will be denoted by S_m.

Definition 1.1: Let $m \in \mathbb{N}$. Let $\pi \in S_m$ and let $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ be a vector in the simplex

$$\Delta_m = \left\{ (\lambda_1, \lambda_2, \ldots, \lambda_m) \in \mathbb{R}^m : \lambda_i > 0, \sum_i \lambda_i = 1 \right\}.$$

Let

$$\beta_0 = 0 \quad \text{and} \quad \beta_j = \sum_{i=1}^{j} \lambda_j \text{ for } 1 \leq j \leq m.$$

The set $\{\beta_0, \beta_1, \ldots, \beta_m\}$ partitions $[0,1)$ into m subintervals of the form $I_j = [\beta_{j-1}, \beta_j)$. Define $T_{(\pi, \lambda)} : [0,1) \to [0,1)$ by

$$T_{(\pi, \lambda)}(x) = x - \left(\sum_{i<j} \lambda_i \right) + \left(\sum_{\pi(i)<\pi(j)} \lambda_i \right), \text{ for } x \in I_j.$$

The map $T_{(\pi, \lambda)}$ rearranges the intervals I_j by translations according to the permutation π. We will refer to a map constructed in this manner as an m-IET. For convenience, we sometimes drop the reference to π and λ and simply denote an IET by a single letter, typically T.

The set of all IETs forms a group $\mathbb{G}$ under composition. Given $T \in \mathbb{G}$, let $C(T)$ denote the centralizer of T in $\mathbb{G}$ and let $\langle T \rangle$ denote the cyclic subgroup generated by T. Novak proved that the quotient $C(T)/\langle T \rangle$ is typically finite [6, Proposition 5.3]. However, there are examples of IETs which satisfy the hypotheses of Novak's theorem for which $C(T) \neq \langle T \rangle$ (some of these examples will be described in the next section of this paper). In fact, to the best of the author's knowledge, the only previously known examples where it could be proven that $C(T) = \langle T \rangle$ are due to del Junco [3], who constructed IETs which do not commute with any Lebesgue measure preserving transformations of $[0,1)$ except for their powers.

In this paper, we will show that for a large class of IETs, the groups $C(T)$ and $\langle T \rangle$ coincide. Recall that an IET T is said to be minimal if for each $x \in [0,1)$, the orbit

$$\mathcal{O}_T(x) = \{T^n(x) : n \in \mathbb{Z}\}$$

is dense in $[0,1)$. Recall also that an IET is said to be of rotation type if there exists $\alpha \in \mathbb{R}$ such that $T(x) = x + \alpha \pmod{1}$ for all $x \in [0,1)$.

THEOREM 1.2: *Let T be a minimal IET which is not of rotation type. Suppose that the lengths of the exchanged subintervals are linearly independent over $\mathbb{Q}$. Then*

$$C(T) = \langle T \rangle.$$

Theorem 1.2 is a simplified version of our main result, which will be stated in the next section of this paper (see Theorem 2.9).

We will denote n-fold compositions $T \circ \cdots \circ T$ by T^n. We will say that T has an n^{th} root in $\mathbb{G}$ if there exists $S \in \mathbb{G}$ such that $T = S^n$. If T is minimal, then the existence of an n^{th} root for some $n \geq 2$ implies that $C(T) \neq \langle T \rangle$. In a previous paper [1], the author proved that many IETs do not have any nontrivial roots in $\mathbb{G}$. In particular, it was proven that an IET which satisfies the hypotheses of Theorem 1.2 does not have an n^{th} root for any $n \geq 2$. Though the non-existence of roots is weaker than the statement that $C(T) = \langle T \rangle$, our proofs of the results in this paper depend on the results in [1].

In the case that T is a 3-IET, we will prove a more precise result. We recall the following definition.

Definition 1.3: Let T be an m-IET. Let $\beta_1, \beta_2, \ldots, \beta_{m-1}$ be as in Definition 1.1. We say that T satisfies the **infinite distinct orbit condition** (IDOC) if each of the orbits $\mathcal{O}_T(\beta_1), \mathcal{O}_T(\beta_2), \ldots, \mathcal{O}_T(\beta_{m-1})$ is infinite and

$$\mathcal{O}_T(\beta_i) \cap \mathcal{O}_T(\beta_j) = \emptyset$$

for $i \neq j$.

The IDOC was originally formulated by Keane, who showed that any IET which satisfies it and exchanges two or more intervals is minimal [4].

THEOREM 1.4: *Let T be a minimal 3-IET which is not of rotation type. Then $C(T) = \langle T \rangle$ if and only if T satisfies the infinite distinct orbit condition.*

Let T be a 3-IET which is not of rotation type, and let $\lambda_1, \lambda_2, \lambda_3$ be the lengths of the exchanged subintervals. As explained in [1, page 251], T is minimal if and only if $\frac{\lambda_1 - \lambda_3}{1 - \lambda_3} \notin \mathbb{Q}$. Moreover, assuming that T is minimal, T satisfies the IDOC if and only if

$$n(\lambda_1 - \lambda_3) \neq \lambda_1 + m(1 - \lambda_3)$$

for every pair of integers m, n. This enables one check whether or not $C(T) = \langle T \rangle$, provided one understands the rational dependencies among $\lambda_1, \lambda_2, \lambda_3$.

2. Centralizers in $\mathbb{G}$

In this section we will define the rank of an IET and use this to state our main result, a generalization of Theorem 1.2. We will also explain how our results are related to the work of Novak [6] and the previous work of the author [1].

The study of centralizers in $\mathbb{G}$ was initiated by Novak. Given $T \in \mathbb{G}$, let $d(T^n)$ denote the number of discontinuities of T^n. Novak showed that for a given $T \in \mathbb{G}$, there are only two possibilities for the growth rate of $d(T^n)$ [6, Theorem 1.1]. The first possibility is that the growth rate is bounded: there exists a constant M such that $d(T^n) \leq M$ for all $n \geq 1$. The second possibility is that $d(T^n)$ grows linearly: there exists a natural number c such that $\lim_{n \to \infty} d(T^n)/n = c$. In the second case, Novak showed that $C(T)$ cannot be too large.

THEOREM 2.1 (Novak [6, Proposition 5.3]): *Suppose that T is a minimal IET and that $d(T^n)$ grows linearly in n. Then $\langle T \rangle$ has finite index in $C(T)$.*

An examination of the proof of Proposition 2.3 of [6] makes it clear that any IET T satisfying the IDOC will exhibit linear discontinuity growth, provided that T is not of rotation type. Novak makes a similar observation on p. 381 of his paper.

It is now convenient to introduce a definition. If T is an m-IET, and $\beta_1, \ldots, \beta_{m-1}$ are as in Definition 1.1, then it is clear that the discontinuities of T must be among $\beta_1, \ldots, \beta_{m-1}$. However, it is possible that T is continuous at some of these points. Whether or not T is continuous at these points depends on the permutation π. More specifically, if $1 \leq i \leq m - 1$, then $T_{(\pi,\lambda)}$ is discontinuous at β_i if and only if $\pi(i + 1) \neq \pi(i) + 1$. This motivates the following definition.

Definition 2.2: We will say that $\pi \in S_m$ is **separating** if $\pi(i+1) \neq \pi(i)+1$ for $1 \leq i \leq m - 1$.

For example, the permutation $\tau = (4213) \in S_4$ is separating, while the permutation $\sigma = (4231) \in S_4$ is not. The following result, whose proof can be found in [1, Proposition 2.3], shows that there is no loss of generality in only considering IETs defined by separating permutations.

PROPOSITION 2.3: *Let T be an IET with precisely $m-1$ discontinuities. There exists a separating permutation $\pi \in S_m$ and a vector $\lambda \in \Delta_m$, both of which are unique, such that $T = T_{(\pi,\lambda)}$.*

We will make use of the following special case of Theorem 2.1.

PROPOSITION 2.4: *Let $m \geq 3$. Suppose that T is an m-IET defined by a separating permutation. If T satisfies the IDOC, then $\langle T \rangle$ has finite index in $C(T)$.*

Proof. Recall that any IET which satisfies the IDOC is minimal. Therefore, by Theorem 2.1 and the remarks immediately following it, we only need to show that T is not of rotation type. Since the permutation defining T is separating, T has $m - 1 \geq 2$ discontinuities, so T is not of rotation type. ∎

It is natural to wonder under what circumstances we have equality $C(T) = \langle T \rangle$. The hypotheses of Theorem 2.1 are not enough to guarantee this. For example, let S be a 3-IET with permutation (321) which satisfies the IDOC. Let $T = S^n$ for some $n \geq 2$. Then T is minimal and has linear discontinuity growth, but $C(T) \neq \langle T \rangle$, since $S \in C(T)$ and $S \notin \langle T \rangle$.

As the preceding example shows, one obstruction to having $C(T) = \langle T \rangle$ is the existence of an n^{th} root for some $n \geq 2$. In a previous paper, the author showed that most IETs do not have nontrivial roots. In order to state this result precisely, we recall the following definition.

Definition 2.5: Let T be an IET. Let $\gamma_1 < \gamma_2 < \cdots < \gamma_{m-1}$ be the points at which T is discontinuous. Let $\gamma_0 = 0$ and $\gamma_m = 1$. We will refer to the dimension of the $\mathbb{Q}$-vector space spanned by $\{\gamma_i - \gamma_{i-1} : 1 \leq i \leq m\}$ as the **rank** of T. This will be denoted by rank(T).

The term "rank" was originally used in this setting by Boshernitzan [2], who proved that minimal rank two IETs are uniquely ergodic and described an algorithm which tests a rank two IET for minimality and aperiodicity.

Using the notation of Definition 1.1, we observe that if $T_{(\pi,\lambda)}$ is an m-IET and π is separating, then the discontinuities of T are precisely the points $\beta_1, \beta_2, \ldots, \beta_{m-1}$, so rank($T$) is equal to the dimension of the $\mathbb{Q}$-vector space spanned by $\lambda_1, \lambda_2, \ldots, \lambda_m$. If π is not separating, then rank($T_{(\pi,\lambda)}$) may be smaller than the dimension of the $\mathbb{Q}$-vector space spanned by $\lambda_1, \lambda_2, \ldots, \lambda_m$.

Notice that if T is an m-IET defined by a separating permutation, then rank(T) can assume any value from 1 to m. The following result was established by the author [1, Theorem 2.4].

THEOREM 2.6: *Let T be a minimal m-IET defined by a separating permutation. Suppose that $\mathrm{rank}(T) > 1 + \lfloor m/2 \rfloor$. Then T does not have an n^{th} root in $\mathbb{G}$ for any $n \geq 2$.*

A second obstruction to having $C(T) = \langle T \rangle$ is the presence of torsion in $C(T)$. For convenience, we will give an example defined on $[0, 2)$ rather than $[0, 1)$. Let $S : [0, 1) \to [0, 1)$ be a 3-IET with permutation (321) which satisfies the IDOC. Define $U : [0, 2) \to [0, 2)$ by

$$ U(x) = \begin{cases} S(x), & x \in [0, 1), \\ x, & x \in [1, 2). \end{cases} $$

Let $P : [0, 2) \to [0, 2)$ be the 2-IET of order two which interchanges $[0, 1)$ and $[1, 2)$. Since the U fixes $[1, 2)$ and PUP fixes $[0, 1)$, U and PUP commute with one another. Consider the map $T = PUPUP = UPU$. It is not difficult to see that T is minimal and that $d(T^n) \approx 2n$, so T exhibits linear discontinuity growth. However, P commutes with T and $P \notin \langle T \rangle$.

One of the goals of this paper will be to establish that $C(T)$ is torsion-free under the following conditions.

THEOREM 2.7: *Let T be an m-IET defined by a separating permutation. Suppose that T satisfies the infinite distinct orbit condition and that $\mathrm{rank}(T) > 1 + \lfloor m/2 \rfloor$. Then the only element of $C(T)$ which has finite order is the identity.*

For $n \in \mathbb{N}$, let r_n denote the periodic IET defined by

$$ r_n(x) = x + \frac{1}{n} \pmod 1. $$

Suppose that T is minimal and that $S \in C(T)$ has finite order n. Then S is conjugate to r_n, and consequently T is conjugate to some element of $C(r_n)$. Novak gives a precise description of $C(r_n)$ [6, Proposition 5.5], and it is clear from Novak's description that all of the elements of $C(r_n)$ fail to satisfy the estimate on the rank given in Theorem 2.7. However, this argument does not prove Theorem 2.7 since conjugation need not preserve rank (see the Appendix for an example). Instead of trying to circumvent this difficulty, we will prove Theorem 2.7 using a different technique. Incidentally, our proof can be used to show that there actually is a conjugating map from S to r_n which does not alter $\mathrm{rank}(T)$.

The proof of Theorem 2.7 will be given in the fifth section of this paper. We will now show how to combine Proposition 2.4 and Theorems 2.6 and 2.7 to establish our main result.

LEMMA 2.8: *Let G be a group and let $x \in G$. Suppose that*

(1) $\langle x \rangle$ *has finite index in $C(x)$;*
(2) x *does not have an n^{th} root in G for any $n \geq 2$;*
(3) $C(x)$ *is torsion-free.*

Then $C(x) = \langle x \rangle$.

Proof. Let $y \in C(x)$. By (1), there exists $n \in \mathbb{N}$ such that $y^n \in \langle x \rangle$, say $y^n = x^m$. We will prove that $y \in \langle x \rangle$ by induction on the number of primes dividing n. If there are no such primes, then $n = 1$ and there is nothing to prove. Suppose that $n \geq 2$. Let p be a prime dividing n, say $n = up$. Write $m = qp + r$, where $0 \leq r < p$. We claim that $r = 0$. If not, then r and p are relatively prime, so there exist integers k and l such that $kr - lp = 1$. Then

$$y^{upk} = y^{nk} = x^{mk} = x^{kqp+kr} = x^{kqp+lp+1}.$$

Since x commutes with y, it follows that $(y^{uk} x^{-l-kq})^p = x$. This contradicts (2). So $r = 0$ and $m = qp$. We now have

$$y^{up} = y^n = x^m = x^{qp}.$$

After rearranging, we find that $(y^u x^{-q})^p = e$, where e denotes the identity in G. By (3), it must be that $y^u x^{-q} = e$. Therefore,

$$y^u = x^q.$$

Since $u = n/p$, it now follows by induction that $y \in \langle x \rangle$.   ∎

THEOREM 2.9: *Let T be an m-IET defined by a separating permutation. Suppose that T satisfies the infinite distinct orbit condition and that $\mathrm{rank}(T) > 1 + \lfloor m/2 \rfloor$. Then $C(T) = \langle T \rangle$.*

Proof. Recall that any IET which satisfies the IDOC is minimal. Notice also that we must have $m \geq 3$, since otherwise the inequality $\mathrm{rank}(T) > 1 + \lfloor m/2 \rfloor$ could not possibly be true. The hypotheses of Proposition 2.4 and Theorems 2.6 and 2.7 are all satisfied. The claim now follows from Lemma 2.8.   ∎

Recall that a permutation $\pi \in S_m$ is said to be **irreducible** if

$$\pi(\{1, 2, \ldots, k\}) \neq \{1, 2, \ldots, k\}$$

for any $k < m$. A result of Keane asserts that if $\pi \in S_m$ is irreducible and the coordinates of $\lambda \in \Delta_m$ are linearly independent over $\mathbb{Q}$, then the IET $T_{(\pi, \lambda)}$ satisfies the IDOC [4]. We will also have $\mathrm{rank}(T_{(\pi, \lambda)}) = m$ in this case, provided that π is separating. Combining these observations with Theorem 2.9 proves the following result, which justifies the statement that "most" IETs have trivial centralizers in $\mathbb{G}$.

COROLLARY 2.10: *Let $m \geq 3$. Let $\pi \in S_m$ be separating and irreducible. Let*

$$A = \{\lambda \in \Delta_m : C(T_{(\pi, \lambda)}) = \langle T_{(\pi, \lambda)} \rangle\}.$$

Then A is a residual subset of Δ_m of full Lebesgue measure.

For completeness, we explain why Theorem 1.2 follows from Theorem 2.9.

Proof of Theorem 1.2. Let T be a minimal IET which is not of rotation type. Suppose that the lengths of the exchanged subintervals are linearly independent over $\mathbb{Q}$. Let $m - 1$ be the number of discontinuities of T. The assumption that T is not of rotation type implies that T has at least two discontinuities. So $m \geq 3$. Choose $\pi \in S_m$ and $\lambda \in \Delta_m$ according to Proposition 2.3. The assumption that the lengths of the exchanged subintervals are linearly independent over $\mathbb{Q}$ implies that the dimension of the $\mathbb{Q}$-vector space spanned by $\lambda_1, \lambda_2, \ldots, \lambda_m$ is m. Since T is minimal, it is clear that π is irreducible. Therefore T satisfies the IDOC. Moreover, since π is separating, $\mathrm{rank}(T) = m$. Since $m \geq 3$, $\mathrm{rank}(T) > 1 + \lfloor m/2 \rfloor$, so Theorem 2.9 implies that $C(T) = \langle T \rangle$. ∎

We now turn our attention to Theorem 1.4. Our proof will follow the same strategy as our proof of Theorem 2.9. In place of Theorem 2.6, we will use following result, which was established by the author [1, Theorem 1.5].

THEOREM 2.11: *Let T be a minimal 3-IET which is not of rotation type. Then T has an n^{th} root in $\mathbb{G}$ for some $n \geq 2$ if and only if T fails to satisfy the infinite distinct orbit condition.*

In place of Theorem 2.7, we will use the following result.

THEOREM 2.12: *Let T be a 3-IET which satisfies the infinite distinct orbit condition and which is not of rotation type. Then $C(T)$ is torsion-free.*

The proof of Theorem 2.12 will be given in the sixth section of this paper. We now proceed with the proof of Theorem 1.4.

Proof of Theorem 1.4. Let T be a minimal 3-IET which is not of rotation type.

Suppose that T satisfies the IDOC. The assumption that T is not of rotation type implies that T must be defined by the permutation (321), which is separating. Thus the hypotheses of Proposition 2.4 and Theorems 2.11 and 2.12 are all satisfied. An application of Lemma 2.8 shows that $C(T) = \langle T \rangle$.

Conversely, suppose that T does not satisfy the IDOC. By Theorem 2.11, there exists $S \in \mathbb{G}$ and $n \geq 2$ such that $T = S^n$. Since T has infinite order in $\mathbb{G}$, it is clear that $S \notin \langle T \rangle$. Therefore $C(T) \neq \langle T \rangle$. $\blacksquare$

We remark that if T is a minimal 3-IET which is not of rotation type and which does not satisfy the IDOC, then Theorems 1.7 and 1.8 of [1] can be used to show that $C(T)$ is actually uncountable and contains a subgroup isomorphic to the circle. We omit the details.

The main results of this paper have been established, but the proofs of Theorems 2.7 and 2.12 are still pending. The rest of this paper will be devoted to proving these two results.

3. The first return map

In this section we briefly review some facts about first return maps which will be used in our proofs of Theorems 2.7 and 2.12. Given an IET T, let $D(T)$ denote the set of points at which T is discontinuous.

Let $\mathcal{A}$ denote the set algebra consisting of finite unions of half-open intervals $[a, b)$ contained in $[0, 1)$. If $J \in \mathcal{A}$, let $E(J)$ denote the set of endpoints of the connected components of J. For example, if $J = [0, \frac{1}{3}) \cup [\frac{4}{5}, \frac{5}{6})$, then $E(J) = \{0, \frac{1}{3}, \frac{4}{5}, \frac{5}{6}\}$.

It is well-known that if T is an IET and $J \in \mathcal{A}$, then the first return map to J is essentially an IET. The following lemma is a precise formulation of this fact.

LEMMA 3.1: *Let T be an IET. Suppose that $J \in \mathcal{A}$ and that no point in $D(T)$ belongs to the interior of J. Let $P \subseteq J$ consist of those points x in the interior of J for which there exists an $n \geq 1$ such that $T^j(x) \notin J \cup D(T) \cup E(J)$ for $0 < j < n$ and $T^n(x) \in D(T) \cup E(J)$. The set P is finite. Therefore P*

partitions J into finitely many subintervals $J_1, J_2, \ldots, J_k$. There exist positive integers $m_1, m_2, \ldots, m_k$ such that

(1) for each i, the restriction of T^j to J_i is a translation for $1 \leq j \leq m_i$;
(2) for each i, $T^j(J_i) \cap J = \emptyset$ for $0 < j < m_i$;
(3) for each i, T^{m_i} translates J_i onto a subinterval of J;
(4) the intervals $T^{m_i}(J_i)$, $1 \leq i \leq k$, are pairwise disjoint.

For a proof of this result, see, e.g., [7, Lemma 4.2].

Definition 3.2: Let T be an IET. Suppose that $J \in \mathcal{A}$ satisfies the hypothesis of Lemma 3.1. For each $x \in J$, let

$$n_J(x) = \inf\{n \geq 1 : T^n(x) \in J\}.$$

According to the lemma, $n_J(x) = m_i$ for $x \in J_i$. The map $T_J : J \to J$ defined by

$$T_J(x) = T^{n_J(x)}(x)$$

is the **first return map** to J. The integers $m_1, m_2, \ldots, m_k$ are the **return times**.

For convenience, we will sometimes refer to first return maps as IETs, even though this is technically inconsistent with Definition 1.1, since the domain is not necessarily $[0, 1)$.

4. Fundamental discontinuities

In this section we use a result of Novak to derive some facts which will be used in our proofs of Theorems 2.7 and 2.12.

Let T be a minimal IET. Let $p \in D(T)$. Since $\mathcal{O}_T(p)$ is infinite, there exists a smallest positive integer N_p such that $T^n(p) \notin D(T) \cup \{0\}$ for $n \geq N_p$. If T^{N_p} is continuous at p, then all of the iterates T^n for $n \geq N_p$ are continuous at p. On the other hand, if T^{N_p} is discontinuous at p, then all of the iterates T^n for $n \geq N_p$ are discontinuous at p. This dichotomy played an important role in Novak's study of the discontinuity growth rate of IETs [6].

Definition 4.1: Let T be a minimal IET. We will say that $p \in D(T)$ is a **fundamental discontinuity** if $T^n(p) \notin D(T)$ for all negative n and T^n is discontinuous at p for all sufficiently large positive n. We will denote the set of fundamental discontinuities by $D_f(T)$.

We emphasize that Definition 4.1 implies that distinct fundamental discontinuities belong to different T-orbits. The following result is due to Novak. It is not stated as a separate lemma in [6], but it is established in the proof of Proposition 5.3 in [6].

LEMMA 4.2 (Novak): *Suppose that T is a minimal IET and that $S \in C(T)$. Then S permutes the orbits of the fundamental discontinuities of T. More precisely, for each $x \in D_f(T)$, there exists a unique $y \in D_f(T)$ such that*

$$S(\mathcal{O}_T(x)) = \mathcal{O}_T(y).$$

COROLLARY 4.3: *Suppose that T is a minimal IET and that $S \in C(T)$ has finite order q. Then S induces a $\mathbb{Z}/q\mathbb{Z}$ action on the set $D_f(T)$. Specifically, the generator of $\mathbb{Z}/q\mathbb{Z}$ acts on $D_f(T)$ by $x \mapsto y$ where y is the unique element of $D_f(T)$ with the property that $S(\mathcal{O}_T(x)) = \mathcal{O}_T(y)$.*

It will be useful to know that each of the orbits of the $\mathbb{Z}/q\mathbb{Z}$ action described in Corollary 4.3 has size q. We will prove this, but first we need the following lemma.

LEMMA 4.4: *Let T be a minimal IET and let $S \in C(T)$. If S has a fixed point, then S is the identity map.*

Proof. Let F be the set of points in $[0, 1)$ which are fixed by S. The definition of an IET implies that F is a finite union of half-open intervals. Since T commutes with S, the set F must be T-invariant. Since T is minimal, $F = [0, 1)$. ∎

COROLLARY 4.5: *Let T be a minimal IET and let $S \in C(T)$. If S has finite order q, then each of the orbits of the $\mathbb{Z}/q\mathbb{Z}$ action described in Corollary 4.3 has size q.*

Proof. Let $x \in D_f(T)$ and suppose that the orbit of x under the $\mathbb{Z}/q\mathbb{Z}$ action has size u, where $1 \le u < q$. Then $S^u(\mathcal{O}_T(x)) = \mathcal{O}_T(x)$, so $S^u(x) = T^k(x)$ for some k. The map $S^u T^{-k}$ fixes x, so Lemma 4.4 implies that $S^u T^{-k} = I$, where I denotes the identity map. So $S^u = T^k$. Since T has infinite order and S^u has finite order, it must be that $k = 0$. But then $S^u = I$, contrary to the fact that S has order q. ∎

The $\mathbb{Z}/q\mathbb{Z}$ action described above will play an important role in our proofs of Theorems 2.7 and 2.12. We will also use the following fact.

LEMMA 4.6: *Suppose that T satisfies the IDOC. Then every discontinuity of T, with the possible exception of $T^{-1}(0)$, is a fundamental discontinuity.*

Proof. Let $p \in D(T)$ and suppose that $p \neq T^{-1}(0)$. Since T satisfies the IDOC, $T^n(p) \notin D(T)$ for all $n < 0$. Let $n \geq 1$. Since T satisfies the IDOC, T is continuous at all points in the forward orbit of p. Since $T(p)$ is contained in the interior of $[0, 1)$, it follows that T^{n-1} is a translation on some open interval containing $T(p)$. Combining this with the fact that T is discontinuous at p, it follows that T^n is discontinuous at p. Since this is true for all $n \geq 1$, p is a fundamental discontinuity of T. ∎

5. Proof of Theorem 2.7

In this section we will prove Theorem 2.7. It is convenient to introduce the following notation. Let T be an IET. Given $x \in [0, 1)$ and integers $n \leq m$, let

$$\mathcal{O}_T(x, n, m) = \{T^j(x) : n \leq j \leq m\}.$$

Each set $\mathcal{O}_T(x, n, m)$ is just a finite part of the T-orbit of x. Notice that if S is an IET which commutes with T, then $S(\mathcal{O}_T(x, n, m)) = \mathcal{O}_T(S(x), n, m)$.

THEOREM 5.1: *Suppose that T is a minimal IET. Let d_f denote the number of fundamental discontinuities. Let d_{nf} denote the number of discontinuities which are not fundamental. If $S \in C(T)$ has finite order q, then*

$$\mathrm{rank}(T) \leq \frac{d_f}{q} + d_{nf} + 1.$$

Proof. For clarity, suppose first that every discontinuity of T is fundamental. Consider the $\mathbb{Z}/q\mathbb{Z}$-action on $D_f(T)$ described in Corollary 4.3. According to Corollary 4.5, the orbits of this action have size q. Let $x_1, x_2, \ldots, x_k$ be representatives for the distinct $\mathbb{Z}/q\mathbb{Z}$-orbits in $D_f(T) = D(T)$.

Using Lemma 4.4, it is not hard to see that every S-orbit has size q. Therefore there exists a subinterval $K = [a, b) \subseteq [0, 1)$ such that the sets $S^i(K)$ for $0 \leq i \leq q - 1$ are pairwise disjoint and the restriction of S to each of these sets is a translation. Notice that the set

$$J = \bigcup_{i=0}^{q-1} S^i(K)$$

is S-invariant. We claim that if K is chosen appropriately, then J will not contain any of the discontinuities of T.

In order to see this, let N be a positive integer which will be specified later. By shrinking K if necessary, we can assume that none of the points in the set $\bigcup_{i=1}^{k} \mathcal{O}_T(x_i, -N, N)$ belong to J. Since J is S-invariant, it follows that none of the points in

$$\bigcup_{j=0}^{q-1} \bigcup_{i=1}^{k} S^j(\mathcal{O}_T(x_i, -N, N))$$

belong to J. Each of the sets

$$S^j(\mathcal{O}_T(x_i, -N, N)) = \mathcal{O}_T(S^j(x_i), -N, N)$$

is part of the T-orbit of one of the discontinuities of T. Moreover, every discontinuity of T belongs to one of the orbits $\mathcal{O}_T(S^j(x_i))$, where $1 \leq i \leq k$ and $0 \leq j \leq q-1$. It follows that if N is chosen large enough, then all of the discontinuities of T belong to $\bigcup_{j=0}^{q-1} \bigcup_{i=1}^{k} S^j(\mathcal{O}_T(x_i, -N, N))$ and therefore $D(T) \cap J = \emptyset$, as claimed.

We can also ensure that the endpoints of the components of J belong to the orbits of some of the points in $D(T)$. Indeed, let M_1 be the largest negative integer such that $T^{M_1}(x_1) \in J$. Let M_2 be the smallest positive integer such that $T^{M_2}(x_1) \in J$. By construction, $M_1 < -N$ and $M_2 > N$. By shrinking K, we can assume that the endpoints of the q connected components of J are precisely the points $S^i(T^{M_j}(x_1))$ for $j = 1, 2$ and $i = 0, 1, \ldots, q-1$. Notice that none of the points in the set $\bigcup_{j=0}^{q-1} S^j(\mathcal{O}_T(x_1, M_1 + 1, M_2 - 1))$ belong to J.

We now consider the first return map T_J. We claim that the set P described in Lemma 3.1 consists of exactly d_f points. By definition, the points belonging to P are precisely those points in the interior of J whose forward T-orbit encounters one of the discontinuities of T or one of the endpoints of J before returning to J. By construction, a point has this property if and only if its forward T-orbit encounters one of the sets $S^j(\mathcal{O}_T(x_1, M_1 + 1, M_2 - 1))$, where $0 \leq j \leq q-1$, or one of the sets $S^j(\mathcal{O}_T(x_i, -N, N))$, where $2 \leq i \leq k$ and $0 \leq j \leq q-1$, before returning to J. Clearly, the backward T-orbit of each discontinuity contains precisely one point with this property, so $|P| = d_f$ as claimed.

The discussion in the preceding paragraph makes it clear that the set P is S-invariant. Indeed, since S commutes with T, the forward T-orbit of a point $p \in J$ encounters one of the sets $S^j(\mathcal{O}_T(x_i, -N, N))$ or returns to J at precisely the same time that the forward T-orbit of $S(p)$ encounters $S^{j+1}(\mathcal{O}_T(x_i, -N, N))$ or returns to J, respectively.

Since $|P| = d_f$ and J has q connected components, P partitions J into $q + d_f$ subintervals $J_1, J_2, \ldots, J_{q+d_f}$. Let $l_1, l_2, \ldots, l_{q+d_f}$ denote the lengths of these intervals, respectively. Let $m_1, m_2, \ldots, m_{q+d_f}$ be the return times for the map T_J. The set

$$A = \bigcup_{i=1}^{q+d_f} \bigcup_{j=0}^{m_i-1} T^j(J_i)$$

is T-invariant. Since T is minimal, it must be that $A = [0, 1)$. Moreover, each of the intervals $T^j(J_i)$, for $1 \le i \le q + d_f$ and $0 \le j \le m_i - 1$, must be contained in one of the intervals on which T is continuous. For if this were not the case, then some point in the interior of J would have to contain a discontinuity in its forward T-orbit. This point would then have to belong to P, a contradiction. It follows that each of the intervals on which are exchanged by T must be a disjoint union of some of the intervals $T^j(J_i)$, where $1 \le i \le q + d_f$ and $0 \le j \le m_i - 1$. These intervals are themselves just translates of the intervals J_i, for $1 \le i \le q + d_f$. Therefore rank(T) is at most equal to the number of distinct lengths appearing among $l_1, l_2, \ldots, l_{q+d_f}$.

Since the partition P which determines the intervals $J_1, J_2, \ldots, J_{q+d_f}$ is S-invariant, and S acts on the q connected components of J by translation, it follows that there are at most

$$\frac{q + d_f}{q} = \frac{d_f}{q} + 1$$

different lengths appearing among $l_1, l_2, \ldots, l_{q+d_f}$. It follows that

$$\text{rank}(T) \le \frac{d_f}{q} + 1,$$

which is exactly what we wanted to show.

In the case when some of the discontinuities are not fundamental, we can proceed as we did above. The only essential difference is that for each non-fundamental discontinuity, there might be an additional point in the set P. This point divides one of the intervals J_t into two intervals, and increases the rank of T by at most one. Our previous estimate on the rank is therefore increased by at most d_{nf}, the number of non-fundamental discontinuities. Thus

$$\text{rank}(T) \le \frac{d_f}{q} + 1 + d_{nf}.$$

This completes the proof.   ∎

Proof of Theorem 2.7. Let T be an m-IET defined by a separating permutation. Suppose that T satisfies the IDOC. Assume that $S \in C(T)$ has finite order $q \geq 2$. We have to prove that $\mathrm{rank}(T) \leq 1 + \lfloor m/2 \rfloor$.

Since T is defined by a separating permutation, $D(T) = m - 1$. Since T satisfies the IDOC, Lemma 4.6 tells us that there are two possibilities: either $d_f = m - 1$ and $d_{nf} = 0$ or else $d_f = m - 2$ and $d_{nf} = 1$. In the first case, Theorem 5.1 tells us that

$$\mathrm{rank}(T) \leq 1 + \frac{m-1}{q} \leq 1 + \frac{m}{2}.$$

Since $\mathrm{rank}(T)$ is an integer, it follows that $\mathrm{rank}(T) \leq 1 + \lfloor m/2 \rfloor$. In the second case, Theorem 5.1 tells us that

$$\mathrm{rank}(T) \leq 2 + \frac{m-2}{q} \leq 2 + \frac{m-2}{2} = 1 + \frac{m}{2}.$$

Once again, we conclude that $\mathrm{rank}(T) \leq 1 + \lfloor m/2 \rfloor$. ∎

6. Proof of Theorem 2.12

Unlike our proof of Theorem 2.7, which was combinatorial in nature, our proof of Theorem 2.12 is based on the notion of topological weak-mixing. The following definition of topological weak-mixing will be sufficient for our purposes.

Definition 6.1: Let λ denote Lebesgue measure on $[0, 1)$. If T is an IET, then the map $f \mapsto f \circ T$ defines a unitary operator on $L^2([0, 1), \lambda)$. We say that T is **topologically weak-mixing** if this operator has no non-constant piecewise continuous eigenfunctions.

To be clear, the term "piecewise continuous" here means that $f : [0, 1) \to \mathbb{C}$ has only finitely many discontinuities, that the left- and right-hand limits $\lim_{x \to a^-} f(x)$ and $\lim_{x \to a^+} f(x)$ exist for every $a \in (0, 1)$, and that the limits $\lim_{x \to 0^+} f(x)$ and $\lim_{x \to 1^-} f(x)$ both exist.

The next result follows from [5, Theorem 4.1], but we include a proof for completeness.

LEMMA 6.2: *Suppose that T satisfies the IDOC (and exchanges two or more intervals). Then any piecewise continuous eigenfunction of T must be continuous everywhere on $[0, 1)$.*

Proof. Let f be a piecewise continuous eigenfunction of T and let $\alpha \in \mathbb{C}$ be the eigenvalue associated to f.

Let $a \in (0,1)$ and let $\epsilon > 0$. Let $y \in (0,1)$ be any point at which f is continuous. Then there exists an $\eta > 0$ such that $|f(p) - f(q)| < \epsilon$ whenever $p, q \in (y - \eta, y + \eta)$. Consider the sets $\{T^k(a) : k \geq 0\}$ and $\{T^k(a) : k < 0\}$. Since T is minimal, both of these sets are dense in $[0,1)$. Since T satisfies the IDOC, at least one of these sets does not contain any of the discontinuities of T. Suppose that $\{T^k(a) : k \geq 0\}$ has this property. Then T^k is continuous at a for all $k > 0$. Choose $k \geq 0$ such that $T^k(a) \in (y - \eta, y + \eta)$. Since T^k is continuous at a, there exists $\delta > 0$ such that $T^k(a - \delta, a + \delta) \subseteq (y - \eta, y + \eta)$. If $x \in (a - \delta, a + \delta)$, then both $T^k(x)$ and $T^k(a)$ belong to $(y - \eta, y + \eta)$. Using the fact that $|\alpha| = 1$, we see that

$$|f(x) - f(a)| = |\alpha^k f(x) - \alpha^k f(a)| = |f(T^k(x)) - f(T^k(a))| < \epsilon.$$

This shows that f is continuous at a. An analogous argument can be given when the set $\{T^k(a) : k < 0\}$ contains no discontinuities of T. Continuity of f at 0 can be proven in a similar way. ∎

The following result is folklore. The author learned about it from Michael Boshernitzan.

PROPOSITION 6.3: *Let T be a 3-IET which satisfies the IDOC and which is not of rotation type. Then T is topologically weak-mixing.*

Proof. Let $f : [0,1) \to \mathbb{C}$ be a piecewise continuous eigenfunction of T. Let $\alpha \in \mathbb{C}$ be the associated eigenvalue. By Lemma 6.2, f is continuous everywhere on $[0,1)$.

If f vanishes at some point $x \in [0,1)$, then f also vanishes along the T-orbit of x, which is dense in $[0,1)$. Since f is continuous, this implies that f is identically zero. So we may assume that f vanishes nowhere. In particular, after multiplying f by a suitable constant, we may assume that $f(0) = 1$. Since f is continuous at 0, $\lim_{x \to 0^+} f(x) = 1$.

Let β_1 and β_2 be the two points at which T is discontinuous. Since T is not of rotation type, the permutation defining T must be (321). Therefore $T([\beta_2, 1))$ lies to the far left of the unit interval. It follows that

$$\lim_{x \to \beta_2^+} f(x) = \lim_{x \to \beta_2^+} \frac{f(T(x))}{\alpha} = \lim_{x \to 0^+} \frac{f(x)}{\alpha} = \frac{1}{\alpha}.$$

We also know that the interval $T([\beta_1, \beta_2))$ lies immediately to the left of $T([0, \beta_1))$. Since f is continuous at $T(0)$, it follows that

$$\lim_{x \to \beta_2^-} f(x) = \lim_{x \to \beta_2^-} \frac{f(T(x))}{\alpha} = \lim_{x \to 0^+} \frac{f(T(x))}{\alpha} = \lim_{x \to 0^+} f(x) = 1.$$

Since f is continuous at β_2, it must be that

$$1 = \lim_{x \to \beta_2^-} f(x) = \lim_{x \to \beta_2^+} f(x) = \frac{1}{\alpha}.$$

So $\alpha = 1$ and $f \circ T = f$. Since T is minimal, f must be a constant function. ∎

The next few results involve the concept of a tower over an IET, which we now explain. Let T be an m-IET. Here we allow the domain of T to be any finite interval $[a, b)$, possibly different from $[0, 1)$. Let $I_1, I_2, \ldots, I_m$ be the intervals which are exchanged by T and suppose that $f : [a, b) \to \mathbb{N}$ is constant on each of these intervals, say $f(x) = n_j$ for $x \in I_j$. We can define a new map T_f as follows. The domain will consist of those points of the form (x, i) where $x \in [a, b)$ and i is an integer satisfying $1 \le i \le f(x)$. The map T_f is defined by

$$T_f(x, i) = \begin{cases} (x, i+1) & \text{if } i+1 \le f(x), \\ (T(x), 1) & \text{otherwise.} \end{cases}$$

The domain may be visualized as a tower over $[a, b)$. The map T_f transports a point up to the next level of the tower, unless the point is already at the top, in which case it is transported back to the first level according to the original map T.

Definition 6.4: Let T and f be as in the preceding paragraph. We can view the map T_f as an IET by laying the levels of the tower end to end, and then rescaling so that the total length of the resulting interval is one. We will refer to an IET constructed in this manner as a tower of type $(n_1, n_2, \ldots, n_m)$ over T. If all of the n_i are equal, we will refer to T_f as a tower of constant height $n_1 = n_2 = \cdots = n_m$.

LEMMA 6.5: *Let T be a tower over a minimal 2-IET. Then T is not topologically weak-mixing.*

Proof. By [1, Theorem 1.8], T is conjugate in the group $\mathbb{G}$ to either a minimal 2-IET or a tower of constant height $d > 1$ over a minimal 2-IET. Let g denote the conjugating IET in either case.

Suppose first that gTg^{-1} is a 2-IET, say $x \mapsto x+\alpha \pmod 1$. Let $f(x) = e^{2\pi i x}$. Then the function $f \circ g$ is a non-constant piecewise continuous eigenfuncion of T, with eigenvalue $e^{2\pi i \alpha}$.

Now suppose that gTg^{-1} is a tower of constant height $d > 1$ over a 2-IET. Let $I_1, I_2, \ldots, I_d$ be the levels of the tower and let $\chi_1, \chi_2, \ldots, \chi_d$ be their characteristic functions, respectively. Let ζ_d be a primitive d^{th} root of unity. Let

$$f(x) = \sum_{k=1}^{d} \zeta_d^k \chi_k(x).$$

Then $f \circ g$ is a non-constant piecewise continuous eigenfunction of T, with eigenvalue ζ_d. ∎

PROPOSITION 6.6: *Let T be a 3-IET which satisfies the IDOC and which is not of rotation type. If $S \in C(T)$ has order two, then T is not topologically weak-mixing.*

Proof. Let β_1 and β_2 be the discontinuities of T. According to Lemma 4.6, at least one of β_1, β_2 is a fundamental discontinuity. Since S has order two, Corollary 4.5 implies that both β_1 and β_2 are fundamental discontinuities, and that S permutes their orbits: $S(\mathcal{O}_T(\beta_1)) = \mathcal{O}_T(\beta_2)$. As in the proof of Theorem 5.1, there exists some interval $K \subseteq [0, 1)$ such that

(1) K and $S(K)$ are disjoint;
(2) the set $J = K \cup S(K)$ satisfies the hypothesis of Lemma 3.1;
(3) the set P described in Lemma 3.1 is S-invariant and contains precisely two points.

It follows that the first return map T_J exchanges four intervals, say $J_1, J_2 \subseteq K$ and $J_3 = S(J_1), J_4 = S(J_2) \subseteq S(K)$. Let m_1, m_2, m_3, m_4 be the return times. Since the set J is S-invariant and S commutes with T, the return time function $n_J(x) = \inf\{n \geq 1 : T^n(x) \in J\}$ must be S-invariant. Therefore, $m_1 = m_3$ and $m_2 = m_4$. This implies that the restriction of S to J commutes with T_J.

Let $\pi \in S_4$ be the permutation which describes how T_J rearranges the intervals J_1, J_2, J_3, J_4. Notice that S acts on J according to the permutation $(13)(24)$. Since S commutes with T_J, π must commute with $(13)(24)$. So π must be one of the following: the identity, (13), (24), $(13)(24)$, $(12)(34)$, $(14)(23)$, (1234), or (1432). We can rule out the first five possibilities, since these would result in T_J, and hence T, not being minimal.

Let P be a tower of type (m_1, m_2) over a rotation. Let K_1 and K_2 be the intervals which are exchanged by the underlying rotation. Notice that the lengths of K_1 and K_2 are twice those of J_1 and J_2, respectively. Let $g : [0, 1) \to [0, 1)$ be the two-to-one map which is defined as follows. For $0 \le j \le m_1 - 1$, g scales each of the intervals $T^j(J_1)$ and $T^j(J_3)$ by a factor of two, and then maps both onto $P^j(K_1)$. For $0 \le j \le m_2 - 1$, g scales each of the intervals $T^j(J_2)$ and $T^j(J_4)$ by a factor of two, and then maps both onto $P^j(K_2)$. Using the fact that π is one of (1234), (1432), (14)(23), one can verify that $g \circ T = P \circ g$.

By Lemma 6.5, there exists a non-constant piecewise continuous eigenfunction, say f, of P. Let $\alpha \in \mathbb{C}$ be the associated eigenvalue. Then, for any $x \in [0, 1)$, we have

$$f(g(T(x))) = f(P(g(x))) = \alpha f(g(x)).$$

Since g is locally an affine map, it is clear that $f \circ g$ is piecewise continuous and non-constant. This proves that T is not topologically weak mixing.   ∎

Proof of Theorem 2.12. Let T be a 3-IET which satisfies the IDOC and which is not of rotation type. We have to prove that $C(T)$ is torsion-free.

Suppose that $S \in C(T)$ has finite order $q > 1$. By Lemma 4.6, T has at least one fundamental discontinuity. Since T has at most two fundamental discontinuities, Corollary 4.5 implies $q \le 2$, so it must be that $q = 2$. Proposition 6.6 now implies that T is not topologically weak-mixing. This contradicts Proposition 6.3.   ∎

7. Appendix

In the second section of this paper, we stated that conjugation in the group $\mathbb{G}$ need not preserve the rank of an IET. For the benefit of the reader, we provide here a simple example to illustrate this phenomenon. For convenience, our example will be defined on $[0, 2)$ rather than $[0, 1)$.

Given a real number s, let $R_s : [0, 1) \to [0, 1)$ be the rotation (2-IET) defined by $R_s(x) = x + s \pmod 1$. Let α, β be two irrational numbers lying in the interval $(0, 1)$. Assume furthermore that $\beta \notin \mathbb{Q}(\alpha)$. Define a map $S : [0, 2) \to [0, 2)$ as follows:

$$S(x) = \begin{cases} R_\alpha(x) & \text{if } x \in [0, 1), \\ R_\beta(x - 1) + 1 & \text{if } x \in [1, 2). \end{cases}$$

Let $P : [0, 2) \to [0, 2)$ be the 2-IET which interchanges $[0, 1)$ and $[1, 2)$. Finally, let $T = PS$. The reader can verify that

(1) every point in $[0, 1)$ returns to $[0, 1)$ after exactly two iterations of T;
(2) the first return map induced by T on $[0, 1)$ is $R_{\alpha + \beta}$.

These two properties together imply that T is conjugate to a tower of constant height two over $R_{\alpha + \beta}$. Such a map has rank two. However, since $\beta \notin \mathbb{Q}(\alpha)$, T has rank three.

References

[1] D. Bernazzani, *Most interval exchanges have no roots*, Journal of Modern Dynamics **11** (2017), 249–262.

[2] M. Boshernitzan, *Rank two interval exchange transformations*, Ergodic Theory and Dynamical Systems **8** (1988), 379–394.

[3] A. del Junco, *A family of counterexamples in ergodic theory*, Israel Journal of Mathematics **44** (1983), 160–188.

[4] M. Keane, *Interval exchange transformations*, Mathematische Zeitschrift **141** (1975), 25–31.

[5] A. Nogueira and D. Rudolph, *Topological weak-mixing of interval exchange maps*, Ergodic Theory and Dynamical Systems **17** (1997), 1183–1209.

[6] C. Novak, *Discontinuity growth of interval exchange maps*, Journal of Modern Dynamics **3** (2009), 379–405.

[7] M. Viana, *Ergodic theory of interval exchange maps*, Revista Matemática Complutense **19** (2006), 7–100.