A map with topological minimal self-joinings in the sense of del Junco

Jonathan L. King*

University of Florida, Gainesville 32611-2082, squash@math.ufl.edu

Abstract. Andrés del Junco has proposed a definition of topological minimal self-joinings intended to parallel Dan Rudolph's measure-theoretic concept. By means of a rank-two "cutting and stacking", this article constructs the first example of a system (a subshift) satisfying his proposed definition of 2-fold topological minimal self-joinings.

The second part of the article shows that 2-fold topological minimal self-joinings does not imply 3-fold and that no map has 4-fold topological minimal self-joinings. This latter result follows from a generalization of a theorem of Schwartzman.

§0 Introduction

In 1979 Dan Rudolph showed that Don Ornstein's rank-1 mixing map, T, could be built to possess a measure-theoretic property which has come to be called ***minimal self-joinings*** of all orders. For order 2 (our tacit assumption, if no adjective is present) this says that whenever two copies of T sit simultaneously as factors of a third, ergodic, transformation then the two copies are either independent or are identified by some power of T. Equivalently, letting (X, μ) denote the space on which T acts, any $T \times T$-invariant ergodic measure on $X \times X$ projecting to μ on each coordinate is either product measure $\mu \times \mu$ or else is supported on the graph of some T^n. One can similarly define K-fold minimal self-joinings by considering the possible ergodic measures living on $X^{\times K}$, the cartesian K-th power of X. If T has minimal self-joinings of all orders then any automorphism of $T^{\times \mathbb{N}}$ is simply a cartesian product of powers of T composed with a permutation of the coordinates. Rudolph constructed such a map in [R] and used the automorphism property to fabricate a menagerie of counterexample transformations. Later investigations showed, [K] and [K,T], that the maps with minimal self-joinings play the role of elementary building blocks for a class of maps containing the finite-rank mixing maps.

What analogue does this notion have in the topological category of a homeomorphism T of a compact metric space X? Using minimal sets as the analogue of ergodic measures, Nelson Markley proposed in [M] a definition paralleling Rudolph's measure-theoretic definition. He and Joe Auslander proved in [A,M] that this property for order 2 implies the property for all orders and

*Partially supported by National Science Foundation Postdoctoral Research Fellowship. Affiliation when the article was submitted: *Dept. of Mathematics, UC Berkeley*

Ergodic Theory and Dynamical Systems, vol. 10, (1990), 745–761.

powers. They then went on to establish a structure theorem roughly analogous to the measure-theoretic one of [K]. However, an exact analogue of Rudolph's theory is not obtained because, for such transformations, the automorphism group of $T^{\times\mathbb{N}}$ seems difficult to pin down.

In [J] Andrés del Junco took orbit closures as the natural analogue of ergodic measures, giving rise to a different definition of topological minimal self-joinings (below) which is appealing to measure-theorists because it permits an analogue of product measure. However, since no examples fulfilling his definition were known, he worked with a complex alternative definition designed to be fulfilled by the topological Chacón's transformation and for it obtained an analogue of a large part of the Rudolph theory.

This note constructs a map satisfying Andrés' original "natural" definition, in the hope of resuscitating interest in it. The method by which the map is constructed seems also to be of independent interest.

Section 2 strengthens an old theorem of Schwartzman to show that 4-fold topological minimal self-joinings does not exist, thus answering negatively a question of [J]. By contrast, in the measure-theoretic setting, all known maps with 2-fold minimal self-joinings have msj of all orders and indeed "2-fold $\Longrightarrow$ N-fold" is a central open question related to the (in)famous problem of whether 2-fold mixing implies N-fold mixing. What is known and will appear elsewhere, is that 4-fold minimal self-joinings implies N-fold, in the measure-theoretic category.

I thank Joe Auslander and Nelson Markley for cheerfully suffering and critically examining early versions of this argument. This work was supported by a National Science Foundation Postdoctoral Research Fellowship.

Nomenclature. We use $[n\mathinner{..}m)$ to indicate a half-open "interval of integers" $\mathbb{Z}\cap[n,m)$. A similar convention holds for closed and open intervals of integers.

Having fixed a finite alphabet $\mathbf{A}$, a ***name*** denotes a doubly infinite string of letters from $\mathbf{A}$ ie., is a point $x\in\mathbf{A}^{\mathbb{Z}}$. Use $x[n]$ to denote the n-th letter in x and let $x[n\mathinner{..}m)$ denote the substring

$$x[n]\,x[n+1]\,x[n+2]\,\ldots\,x[m-1]$$

with the obvious meaning if $m=\infty$. A ***word*** means a finite string of letters from $\mathbf{A}$. Words are always indexed from zero ie., if word W is of length h then $W=W[0\mathinner{..}h)$. The general purpose term "string" may denote a finite, half-infinite, or bi-infinite string of letters.

Let $\#X$ denote the cardinality of the set X. Agree to use $\coloneqq$ to mean "is defined to be"; in $a\coloneqq b$ the expression b defines the symbol a.

Topological notions. This paper concerns a homeomorphism $T\colon X\to X$ of a compact metric space with $|\cdot,\cdot|$ denoting the metric. Use $\mathcal{O}_T(x)$ or just $\mathcal{O}(x)$ to represent the orbit $\{T^n x\}_{n\in\mathbb{Z}}$ of x and let $\overline{\mathcal{O}}(x)$ denote its orbit closure. An $x\in X$ is a ***transitive point*** if $\overline{\mathcal{O}}(x)=X$. Distinct points x and y are ***proximal*** if

$$\inf_{n\in\mathbb{Z}}\bigl|T^n x, T^n y\bigr|=0;$$

equivalently, there exist a point z and times $n(i)$ such that $T^{n(i)}x\to z$ and $T^{n(i)}y\to z$. Points x and y are future (past) ***asymptotic*** if $|T^n x, T^n y|$ goes to zero as $n\to+\infty$ ($n\to-\infty$).

Map T is ***topologically transitive*** if every invariant non-empty open set is dense; in the case of a metric space, this is equivalent to X possessing a transitive point. The map is topologically

weak-mixing if $T \times T$ is topologically transitive. If every point of T is transitive then T is said to be ***minimal***. Finally, a minimal map T is ***proximal orbit dense*** if for every pair of points x, y in distinct orbits there exists $n \neq 0$ such that x and $T^n y$ are proximal. This property is properly weaker than the following.

DEFINITION. A map $T\colon X \to X$ has ***topological minimal self-joinings*** in the sense of del Junco if two conditions hold.

(i) Every (non-zero) power of T is minimal.
(ii) For any pair of points $x, y \in X$ not in the same orbit, the pair $\langle x, y \rangle$ is a transitive point for $T \times T$.

Condition (ii) says that any pair which *could* be transitive under $T \times T$ *is* transitive: Evidently if x and y are in the same orbit, $y = T^3 x$ for instance, then the $T \times T$ orbit closure of $\langle x, y \rangle$ is the third ***off-diagonal***

$$D_3 \coloneqq \big\{ \langle z, T^3 z \rangle \bigm| z \in X \big\}$$

and thus certainly cannot be all of $X \times X$. ♦

Condition (ii) of topological minimal self-joinings is equivalent to the seemingly weaker condition that whenever points x and y are in different orbits, then they are proximal. For this along with the minimality of T yields

$$\overline{\mathcal{O}}_{T\times T}\big(x, y\big) \supset D_0 \,.$$

whenever x and y are points in different orbits. But for any k the points x and $T^{-k} y$ are in different orbits and so the orbit closure of $\langle x, T^{-k} y \rangle$ contains D_0; equivalently, $\overline{\mathcal{O}}_{T\times T}\big(x, y\big)$ contains D_k. Thus

$$\overline{\mathcal{O}}_{T\times T}\big(x, y\big) \supset \text{Closure}\big\{ \langle z, T^k z \rangle \bigm| z \in X \;\;\&\;\; k \in \mathbb{Z} \big\} \overset{\text{note}}{=\!=} X \times X.$$

This last equality follows by the minimality of T. ♦

The point(s) of difficulty. The topological Chacón's map is proximal orbit dense but does not have topological minimal self-joinings. The difficulty comes from the existence of a pair $\langle x, y \rangle$ of future asymptotic points. In the context of a symbolic space this means that there is some N for which $x[N \mathinner{..} \infty) = y[N \mathinner{..} \infty)$.

For a minimal T, any asymptotic pair x and y must be in distinct orbits. For if $y = T^p x$, say, then any limit point $z \coloneqq \lim_{i\to\infty} T^{n(i)} x$ is periodic—since $\lim_i T^{n(i)} y$ must equal both z and $T^p z$. But T is minimal hence X has no periodic points (a periodic minimal map has no asymptotic pairs). Thus p must have been zero.

Now topological minimal self-joinings says that for a pair $\langle x, y \rangle$ the orbit closure is either an off-diagonal or is all of $X \times X$. But suppose we have two points x and y which are future-asymptotic under some shift k (ie. x and $T^k y$ are future-asymptotic) and past-asymptotic under a different shift s. Then the $T \times T$ orbit closure of $\langle x, y \rangle$ consists of *two* off-diagonals, $D_k \cup D_s$. [Indeed, after shifting y so that neither k nor s is zero, the pair $\langle x, y \rangle$ is not proximal.]

The exceptional points of Chacón's map. This map fails to achieve topological minimal self-joinings by the slimmest of margins –there is one bad pair of orbits. There are names x and y for which $x(-\infty \mathinner{..} 0) = y(-\infty \mathinner{..} 0)$ and $x[0 \mathinner{..} \infty) = y[1 \mathinner{..} \infty)$. These names look like

$$\begin{aligned} x &= \overleftarrow{p}\,\overrightarrow{q} \\ y &= \overleftarrow{p}\,\mathsf{s}\,\overrightarrow{q} \end{aligned} \tag{0.1}$$

where $\overleftarrow{p}$ is a past string, $\overrightarrow{q}$ a future string, and $\mathtt{s}$ is a letter from the alphabet (see [J,K]).

What to avoid. Asymptoticity cannot be avoided—any expansive system $T\colon X \to X$ has future-asymptotic pairs and past-asymptotic pairs (this well-known fact follows from theorem 2.1). What must be prevented is that a pair $\langle x, y\rangle$ possesses, upto a shift, both future and past asymptoticity.

Say that a string is "valid" if it is a substring of some x in X. Suppose there existed pairs of valid words V_n and W_n, whose lengths went to infinity, such that the concatenations

$$V_n W_n \quad \text{and} \quad V_n \mathtt{s} W_n$$

were both valid strings. By taking a weak limit, the space X would be forced to have points x and y as in (0.1).

To prevent this, we will construct X so that when a word V is sufficiently long, it determines the letter which must follow it.

§1 The Construction

How might one go about building a shift-invariant symbolic space? Fix a finite alphabet $\mathbf{A}$. Suppose $\mathcal{V}$ is some collection of strings over $\mathbf{A}$: strings which are finite, left-infinite, right-infinite, or bi-infinite. Define $\mathrm{Cl}(\mathcal{V})$ to be the set of names $x \in \mathbf{A}^{\mathbb{Z}}$ for which:

Every finite substring of x is a substring of some $V \in \mathcal{V}$.

Since no mention of time is made in its definition, $\mathrm{Cl}(\mathcal{V})$ is a shift-invariant subset of $\mathbf{A}^{\mathbb{Z}}$. *A fortiori* $\mathrm{Cl}(\mathrm{Cl}(\mathcal{V})) = \overline{(}\mathcal{V})$ and consequently

$$\mathrm{Cl}(\mathcal{V}) \text{ is a closed shift-invariant subset of } \mathbf{A}^{\mathbb{Z}}.$$

The shift-left map on this forms a topological dynamical system.

Building the space X. Our alphabet will be $\{\mathtt{a}, \mathtt{b}, \mathtt{s}\}$; "$\mathtt{s}$" will be employed as a spacer between n-blocks. For $n = 0, 1, 2, \ldots$ we build two ***types*** of n-block

$$H_n^{\mathbf{0}} \quad \text{and} \quad H_n^{\mathbf{1}}$$

whose common length is denoted h_n. It will be convenient that the $\{h_n\}_n$ be even numbers; we leave verification to the reader that the operations done below do not prevent our arranging this.

Concatenation will be indicated by juxtaposition. However, we use $\bigotimes$ to write iterated concatenation: For example, if Letter(i) denotes the i-th Roman letter then $\bigotimes_{i=1}^{5}$ Letter(i) means $\mathtt{abcde}$. The expression $3\times$"$\mathtt{ab}$" indicates $\mathtt{ababab}$.

In the sequel, the symbols α and β take on the values $\mathbf{0}$ and $\mathbf{1}$. When not desiring to specify the type of an n-block we will write H_n.

Step A: *Set $H_0^{\mathbf{0}} :=$"$\mathtt{a}$" and $H_0^{\mathbf{1}} :=$"$\mathtt{b}$".*

Step B: *At stage n, with $H_n^{\mathbf{0}}$ and $H_n^{\mathbf{1}}$ known, pick a K much larger than h_n. For each α pick a sequence of **gap sizes***

$$\left\{g^{\alpha}(k)\right\}_{k=0}^{K-1}$$

with each $g^\alpha(k)$ either **0** *or* **1**. [We have suppressed the subscript n in K_n and in g^α_n.]

STEP C: *Define the two $(n+1)$-blocks as follows.*

$$H^\alpha_{n+1} \coloneqq \left(\bigotimes_{k=0}^{K-1} H_n^{g^\alpha(k)} \big[g^\alpha(k) \times \text{“}\mathtt{s}\text{”} \big] \right) H^\alpha_n .$$

Thus each n-block of type **1** is followed by a spacer; those of type **0** are followed by no spacer. The type of the rightmost n-block must agree with the type of its enclosing $(n+1)$-block.

Finally, let T denote the shift-left map on X where: $X \coloneqq \mathrm{Cl}\big\{ H_n^{\mathbf{0}} \bigm| n \in \mathbb{N} \big\}$.

DEFINITION. Say that a word W is ***neatly n-blocked*** if it can be written in the form

$$W = \left(\bigotimes_{\ell=0}^{L-1} H_n^{\beta(\ell)} \big[\beta(\ell) \times \text{“}\mathtt{s}\text{”} \big] \right) H_n^{\beta(L)} \tag{1.1}$$

where each $\beta(\ell)$ is **0** or **1**.

Let $\mathrm{Spacer}_n(W)$ denote $\beta(L)$. This is a slight abuse of notation since $\beta(L)$ depends on the righthand side of (1.1) –which we have not bothered to show unique. However, this will cause no difficulty where we need it, in the following lemma. ♦

CONSISTENCY LEMMA. *For any pair $M \ge N$ and each α*

$$H^\alpha_M \text{ is neatly } N\text{-blocked.}$$

Thus for any $x \in X$, each substring $x[a \mathinner{..} b)$ can be extended to a larger enclosing substring $x[a' \mathinner{..} b')$ which is neatly N-blocked.

Remark. Consequently, given any position i on x there is a position j with

$$i - \frac{1}{2} h_N \le j \le i + \frac{1}{2} h_N$$

such that $x[j \mathinner{..} j + h_N) = H_N$.

PROOF. For any n and β the word H_n^β is neatly $(n-1)$-blocked, by definition. Evidently if words U and V are neatly $(n-1)$-blocked and $k \coloneqq \mathrm{Spacer}_{n-1}(U)$ then the concatenation

$$U \big[k \times \text{“}\mathtt{s}\text{”} \big] V$$

is neatly $(n-1)$-blocked. But STEP C implies $\mathrm{Spacer}_n(H_n^\beta) = \mathrm{Spacer}_{n-1}(H_n^\beta)$ and so for any word W:

$$W \text{ neatly } n\text{-blocked} \quad \Longrightarrow \quad W \text{ neatly } (n-1)\text{-blocked.}$$

By induction on n from M down to $N+1$, the word H^α_M is neatly N-blocked. ♦

Creating topological properties

The properties of T depend on our algorithm for choosing, at stage n, the parameters K and $\{g^\alpha(k)\}_{k=0}^{K-1}$. Below we give a first approximation of these gap sequences; this will be refined by later lemmas, whose proofs will require successive modifications to the sequences. These modifications will not affect properties previously obtained.

FIRST VERSION: *At stage n, set $K \coloneqq 4\kappa$ for some large κ. Let $\{g^\alpha(k)\}_{k=0}^{K-1}$ equal the following.*

$$\begin{array}{ll} \text{For } \alpha=\mathbf{0}: & \overbrace{\mathbf{000\cdots00}}^{\kappa\text{ many}}\overbrace{\mathbf{111\cdots11}}^{\kappa\text{ many}}\overbrace{\mathbf{000\cdots00}}^{\kappa\text{ many}}\overbrace{\mathbf{111\cdots11}}^{\kappa\text{ many}} \\ \text{For } \alpha=\mathbf{1}: & \underbrace{\mathbf{101010}\ \cdots\cdots\cdots\cdots\cdots\cdots\ \mathbf{1010}}_{\text{Length } K} \end{array} \tag{1.2a}$$

Each row has the same number of $\mathbf{1}$'s. Thus $H_{n+1}^{\mathbf{0}}$ and $H_{n+1}^{\mathbf{1}}$ indeed have the same length.

The next two lemmas need some notation. Given two words V and W of some length h and an integer s, let "$s+W$" indicate the word W shifted s units to the left. The phrase

"the intersection $V \cap [s+W]$ contains word U"

is to mean that there exists a position i such that

$$V[i \mathinner{..} i+u) = U \qquad \text{and} \qquad W[i+s \mathinner{..} i+s+u) = U$$

where u denotes the length of word U. (Thus i must satisfy $0 \le i$ and $i+u \le h$ as well as $0 \le i+s$ and $i+s+u \le h$.)

OPPOSITE TYPES LEMMA. *For any shift s satisfying $|s| \le \frac{1}{2}h_n$, the intersection*

$$H_n^{\mathbf{0}} \cap \left[s + H_n^{\mathbf{1}}\right]$$

contains a copy of $H_{n-1}^{\mathbf{0}}$.

PROOF. Let h denote h_{n-1}. Since the $\alpha = \mathbf{1}$ row of (1.2a) has an equal number of $\mathbf{1}$'s in its first half as in its second half, shifting $H_n^{\mathbf{1}}$ by at most $\frac{1}{2}h_n$ corresponds to a horizontal shift of the bottom row of (1.2a), by at most half its length, relative to the top row. The interval where the rows overlap will contain κ many $\mathbf{0}$'s sitting above the periodic pattern $\mathbf{0101}\cdots\mathbf{01}$. Thus the upstairs word $H_n^{\mathbf{0}}$ will contain a periodic pattern of period h sitting above a periodic pattern in $H_n^{\mathbf{1}}$ with period $2h+1$ (the length of the string $H_{n-1}^{\mathbf{0}}H_{n-1}^{\mathbf{1}}\mathsf{s}$ which arises from the $\mathbf{01}$ gap pattern).

These two numbers h and $2h+1$ are relatively prime. Consequently if we have chosen

$$\kappa \gg h \cdot (2h+1) \tag{1.2b}$$

then there will be some position, i, where the upstairs and downstairs periodic patterns start in synchronism. Starting there,

$$H_n^{\mathbf{0}}[i \mathinner{..} i+h) \;=\; H_{n-1}^{\mathbf{0}} \quad \text{and} \quad \left[s+H_n^{\mathbf{1}}\right][i \mathinner{..} i+2h+1) \;=\; H_{n-1}^{\mathbf{0}}H_{n-1}^{\mathbf{1}}\mathsf{s}$$

which gives the desired conclusion. ♦

The above lemma finds a common substring between two n-blocks; their relative shift is essentially arbitrary (just not too big) but their *types* had to be different. The lemma below obtains a similar conclusion where the restriction has been moved from the types to the shift.

ANY TYPES LEMMA. *Given any shift t for which $\frac{1}{2}h_n < |t| \leq \frac{1}{2}h_{n+1}$, given any types α and β, the intersection*

$$H^{\alpha}_{n+1} \cap \left[t + H^{\beta}_{n+1}\right]$$

contains the word $H^{\mathbf{0}}_{n-1}$.

PROOF. Without loss of generality t is positive. Suppress the subscript and let K and h denote K_n and h_n. Given a type $\gamma \in \{\mathbf{0}, \mathbf{1}\}$ let $i^{\gamma}(k)$ denote the position in H^{γ}_{n+1} commencing the k-th copy, traversing left to right, of an n-block. Thus

$$H^{\gamma}_{n+1}[i^{\gamma}(k) \mathinner{..} i^{\gamma}(k) + h) = H_n, \qquad \text{for } k = 0, 1, \ldots, K.$$

Note that $i^{\gamma}(0)$ equals zero.

Picture the word H^{α}_{n+1} written horizontally above the word H^{β}_{n+1} which has been shifted left by t positions. Let $\widehat{k}$ denote the rightmost n-block in word $t + H^{\beta}_{n+1}$ which overlaps, by at least $\frac{1}{2}h$, the leftmost n-block in H^{α}_{n+1}. In other words, pick $\widehat{k}$ largest such that $|s| \leq \frac{1}{2}h$ where

$$s \coloneqq i^{\beta}(\widehat{k}) - i^{\alpha}(0) - t.$$

The remark after the *Consistency lemma* (and that h is even) shows that such a $\widehat{k}$ exists. Moreover $\widehat{k} \geq 1$, since the given shift t strictly exceeds $\frac{1}{2}h$.

Suppose, for the sake of contradiction, that our lemma fails. The substrings

$$H^{\alpha}_{n+1}\left[i^{\alpha}(0) \mathinner{..} i^{\alpha}(0) + h\right) \quad \text{and} \quad H^{\beta}_{n+1}\left[i^{\beta}(\widehat{k}) \mathinner{..} i^{\beta}(\widehat{k}) + h\right)$$

are n-blocks, upstairs and downstairs respectively, with a relative shift of s. Were their types different, our lemma would be proved by an application of the *Opposite types lemma* to these n-blocks. Their types are consequently the same and so

$$g^{\alpha}_n(0) = g^{\beta}_n(\widehat{k}).$$

But this implies that $i^{\beta}(\widehat{k}+1) - i^{\alpha}(1) - t$, the relative shift between the succeeding n-blocks upstairs and downstairs, also equals s. Again we must be unable to apply the *Opposite types lemma* and so their following gaps must also be equal. By inductively stepping across the upstairs and downstairs $(n+1)$-blocks in this fashion we conclude that $g^{\alpha}_n(j)$ equals $g^{\beta}_n(\widehat{k}+j)$ for $j = 0, 1, \ldots, K - \widehat{k} - 1$.

Periodicity in the gap sequence. An upper bound on $\widehat{k}$ arises from its definition in that $i^{\beta}(\widehat{k})$ must be dominated by $t + \frac{1}{2}h$ and hence by $\frac{1}{2}h_{n+1} + \frac{1}{2}h_n$. Now each row of (1.2a) has about the same number of **1**'s in its first half as in its second half. Hence whether β is **0** or **1** we conclude that $\widehat{k}$ is essentially dominated by $\frac{1}{2}K$. Certainly by magnanimously only asserting that

$$\widehat{k} < \frac{2}{3}K \tag{1.2c}$$

we have absorbed any errors.

Evidently α must equal β: Otherwise we could have applied the *Opposite types lemma* to the given $(n+1)$-blocks and concluded that their intersection contains an $H_n^{\mathbf{0}}$, hence *a fortiori* contains an $H_{n-1}^{\mathbf{0}}$.

The preceding two paragraphs rewrite the conclusion of the "inductively stepping" paragraph as

$$\text{For } 0 \le j \le \tfrac{1}{3}K_n\text{:} \quad g_n^\alpha(j) = g_n^\alpha(\widehat{k}+j). \tag{1.3}$$

This is a periodicity condition on a gap sequence. Thus our proof will be completed by altering the definition of a gap sequence so as to make (1.3) impossible. This simply requires inserting a few distinct "marker" words

$$U \coloneqq \mathbf{1001} \quad V \coloneqq \mathbf{10001} \quad W \coloneqq \mathbf{100001}$$

which appear nowhere else in a gap sequence.

Second Version: *At stage n, with $H_n^{\mathbf{0}}$ and $H_n^{\mathbf{1}}$ known: Pick an even integer $\kappa \gg h_n \cdot (2h_n+1)$. Define $\{g^\alpha(k)\}_{k=0}^{K-1}$ as follows.*

$$\begin{array}{ll}
\textit{For } \alpha = \mathbf{0}: & \overbrace{\mathbf{0000}\cdots\mathbf{00}}^{\kappa \textit{ many}}U\overbrace{\mathbf{1111}\cdots\mathbf{11}}^{\kappa \textit{ many}}V\overbrace{\mathbf{0000}\cdots\mathbf{00}}^{\kappa \textit{ many}}W\overbrace{\mathbf{1111}\cdots\mathbf{11}}^{\kappa \textit{ many}} \\
\textit{For } \alpha = \mathbf{1}: & \underbrace{\mathbf{1010}\cdots\mathbf{10}\,U\,\mathbf{1010}\cdots\mathbf{10}\,V\,\mathbf{1010}\cdots\mathbf{10}\,W\,\mathbf{1010}\cdots\mathbf{10}}_{\textit{Length } \kappa}
\end{array}$$

Any subsequence of $\{g^\alpha(k)\}_{k=0}^{K-1}$ taking up at least a third of it, must contain one of the three marker words. But a particular marker word appears nowhere else in the gap sequence. This prohibits (1.3), since $\widehat{k}$ is not zero. ♦

Proving proximality of points in different orbits. Here starts the proof that T has topological minimal self-joinings. Choose points $x, y \in X$ which are *not* proximal. In several steps, the lemmas above will imply that x and y must be in the same orbit.

Obtaining the shift progression $\{s_n\}_{n=0}^\infty$. By the remark following the *Consistency lemma* there exists a large interval containing, say, $[-10h_n \mathinner{..} 10h_n]$ such that

$$x[a \mathinner{..} b] \text{ is neatly } n\text{-blocked} \tag{1.4}$$

where $a \le -10h_n < 10h_n \le b$. Thus there exists an index

$$i \in \left[-\tfrac{3}{2}h_n \mathinner{..} -\tfrac{1}{2}h_n\right] \tag{1.5}$$

such that $x[i \mathinner{..} i+h_n)$ is an n-block. Similarly, by applying the lemma to y we can find an index $j \in [i - \frac{1}{2}h_n \mathinner{..} i + \frac{1}{2}h_n]$ such that $y[j \mathinner{..} j+h_n) = H_n$. Let

$$s_n \coloneqq i - j$$

denote the relative shift between these two n-blocks.

The limit supremum of $|s_n|$ is finite. We show that there does not exist a sequence of $n \in \mathbb{N}$ along which s_n gets arbitrarily large. Fix n and set $t \coloneqq s_n = i - j$. By its definition, $|t|$ is dominated by $\frac{1}{2}h_n$. We would like to apply the *Any types lemma* to the two n-blocks at i and j. Indeed we could, if we knew that $|t|$ exceeded $\frac{1}{2}h_{n-1}$. Since, however, it might not, simply pass to a smaller value of the subscript. For note that for any $m < n$

$$x[i \mathinner{..} i+h_{m+1}) = x\big[i \mathinner{..} i+h_n\big)[0 \mathinner{..} h_{m+1}) = H_n[0 \mathinner{..} h_{m+1}) = H_{m+1}.$$

Similarly, $y[j \mathinner{..} j+h_{m+1})$ is an $(m+1)$-block of some type. Pick the value m less than n such that

$$\frac{1}{2}h_m < |t| \le \frac{1}{2}h_{m+1}.$$

By the *Any types lemma* there now exists a position p such that

$$x[p \mathinner{..} p+h_{m-1}) = y[p \mathinner{..} p+h_{m-1})$$

because they both equal $H^{\mathbf{0}}_{m-1}$.

This m depends on t and so we write it as $m[t]$. If, along a subsequence, the shift progression $t(\ell) \coloneqq s_{n(\ell)}$ goes to infinity (in absolute value) as $\ell \to \infty$, then $\lim_{\ell\to\infty} m[t(\ell)]$ equals infinity. Thus x and y have equal, aligned, substrings of arbitrarily large length—they would be proximal, contrary to our standing assumption. ♦

The point y is in the orbit of x. The preceding says that there is an infinite subset $F \subset \mathbb{N}$ such that $n \mapsto s_n$ is constant, say, 17, on F. Fix an $n \in F$ and let i and $j \overset{\text{note}}{=\!=} i - 17$ be as in the *Obtaining the shift progression* paragraph. Thus

$$x[i \mathinner{..} i+h_n) = H^{\alpha}_n \qquad \text{and} \qquad y[j \mathinner{..} j+h_n) = H^{\beta}_n$$

for some α and β. We may have taken n large enough that x and y have no equal, aligned, substring of length equal to h_{n-1}. Consequently, the *Opposite types lemma* allows us to conclude that $\alpha = \beta$.

Since the type of a block determines the number of spacers following it,

$$x[i \mathinner{..} i') = y[j \mathinner{..} j')$$

where $i' \coloneqq i + h_n + \alpha$ and $j' \coloneqq j + h_n + \beta$. By (1.4) position i' commences an n-block on x and analogously for j' on y; say $x[i' \mathinner{..} i'+h_n) = H^{\alpha'}_n$ and $y[j' \mathinner{..} j'+h_n) = H^{\beta'}_n$. As above, the *Opposite types lemma* must fail to be applicable to this pair of blocks and so α' equals β'.

Putting this all together yields

$$x[i \mathinner{..} i+L) = y[i-17 \mathinner{..} i-17+L)$$

where L denotes $h_n + \alpha + h_n$. Make explicit their dependence on n and write $i(n)$ and $L(n)$. From (1.5) we have that

$$i(n) \leq -\frac{1}{2}h_n \qquad \text{and} \qquad i(n) + L(n) \geq -\frac{3}{2}h_n + 2h_n = \frac{1}{2}h_n.$$

Thus $i(n) \to -\infty$ and $i(n) + L(n) \to +\infty$ as $n \to \infty$ inside of F. Consequently, $y(-\infty \mathinner{..} \infty)$ equals $x(-\infty \mathinner{..} \infty)$ shifted left by 17.

Since y has been shown to be in the orbit of x, our map T fulfills the proximality condition for topological minimal self-joinings. ♦

Minimality of T. It suffices to show that given any valid word W and any $x \in X$, there exists a position i for which $x\big[i \mathinner{..} i+\text{len}(W)\big)$ equals W. But X is the closure of $\{H_n^{\mathbf{0}}\}_{n=1}^{\infty}$ and so without loss of generality $W = H_n^{\mathbf{0}}$ for some n. By the *Consistency lemma*, the name x contains an $(n+1)$-block of some type. And *both* types of $(n+1)$-block contain $H_n^{\mathbf{0}}$.

Total minimality (T^n is minimal for all non-zero n) can be seen directly by a similar argument—but it also follows on general principles. The proximality condition of tmsj implies that T is (topologically) weak-mixing. And it is well-known, [Ke], that a weak-mixing minimal map is totally minimal. ♦

Remark. It is not hard to see that the foregoing T is uniquely ergodic. We remark without proof that the gap sequences can be modified so that the map still has topological minimal self-joinings but now supports two ergodic measures.

An application

The cartesian square $T^{\times 2}$ is not conjugate to $T^{\times 3}$. First note that a point $\langle x, y, z\rangle$ is in a minimal set for $T^{\times 3}$ if and only if $\mathcal{O}(x) = \mathcal{O}(y) = \mathcal{O}(z)$. For the closure $\overline{\mathcal{O}}_{T\times T}(x,y)$ must be a minimal set for $T^{\times 2}$ (since the projection map is a homomorphism and the homomorphic image of a minimal set is minimal) and so $y \in \mathcal{O}(x)$. Similarly $z \in \mathcal{O}(x)$.

Discussions with Joe Auslander produced this proof that $T^{\times 3}$ is not a factor of $T^{\times 2}$. For suppose

$$\varphi\colon X \times X \to X \times X \times X$$

were a homomorphism of $T^{\times 2}$ onto $T^{\times 3}$. Set $\mathbf{a} \coloneqq \langle x, x, y\rangle$ where $x, y \in X$ with $y \notin \mathcal{O}(x)$. Set

$$\mathbf{a}' = \langle w, z\rangle \coloneqq \varphi^{-1}(\mathbf{a}).$$

Evidently $\mathcal{O}_{T\times T\times T}(\mathbf{a})$ is not dense, since the first two coordinates of $\mathbf{a}$ are equal. Thus $\mathcal{O}_{T\times T}(\mathbf{a}')$ is not dense. Since T has topological minimal self-joinings this implies that w and z are in the same T-orbit; thus the orbit closure $\overline{\mathcal{O}}_{T\times T}(\mathbf{a}')$ is a minimal set. Hence the φ-image of this set,

$$\overline{\mathcal{O}}_{T\times T\times T}(\mathbf{a}),$$

is a minimal subset of $X \times X \times X$. But by the paragraph above, this would imply that y is in the T-orbit of x. ♦

§2 Higher order topological minimal self-joinings

The N-fold generalization of topological minimal self-joinings is that for any N points $x_1, \ldots, x_N$ inhabiting N different T-orbits, the tuple $\langle x_1, \ldots, x_N \rangle$ is a transitive point for $T^{\times N}$. As before, this is equivalent to asking that every such N points be proximal under T.

Two-fold topological minimal self-joinings does not imply three-fold. Our example above fails to have 3-fold topological minimal self-joinings, which can be seen as follows. Notice, for the gap sequences chosen, that the sequence of type α begins with an α. Thus the $(n+1)$-block of type $\mathbf{0}$ starts and ends with $H_n^{\mathbf{0}}$ and the $(n+1)$-block of type $\mathbf{1}$ begins and ends with an n-block of its same type. This will imply the existence of three distinct points $x, y, z \in X$ such that $\langle x, y \rangle$ are future-asymptotic and $\langle y, z \rangle$ are past-asymptotic. As argued earlier, minimality ensures that neither x nor z is in $\mathcal{O}(y)$. Nor could $z \in \mathcal{O}(x)$ since this implies that $\overline{\mathcal{O}}_{T\times T}(x, y)$ consists of two off-diagonals and thus is not dense. So the points $\{x, Ty, z\}$ are in three distinct orbits and yet ... they cannot be triply proximal.

To make these names it will be convenient to let $\overrightarrow{\mathbf{0}}$ and $\overrightarrow{\mathbf{1}}$ denote certain right-infinite strings, that is, points in $\mathbf{A}^{\mathbb{N}}$. Define them inductively by

$$\overrightarrow{\mathbf{0}}[0 \mathinner{..} h_n) \coloneqq H_n^{\mathbf{0}} \qquad \text{and} \qquad \overrightarrow{\mathbf{1}}[0 \mathinner{..} h_n) \coloneqq H_n^{\mathbf{1}}, \qquad \text{for } n = 0, 1, 2, \ldots.$$

This produces valid strings because H_{n+1}^{α} commences with an H_n^{α}. Similarly, since H_{n+1}^{α} ends with H_n^{α}, the left-infinite pasts $\overleftarrow{\mathbf{0}}$ and $\overleftarrow{\mathbf{1}}$ specified inductively as

$$\overleftarrow{\mathbf{0}}[-h_n \mathinner{..} 0) \coloneqq H_n^{\mathbf{0}} \qquad \text{and} \qquad \overleftarrow{\mathbf{1}}[-h_n \mathinner{..} 0) \coloneqq H_n^{\mathbf{1}},$$

are also well-defined. Concatenate the futures to the pasts as follows:

$$\begin{aligned} x &\coloneqq \overleftarrow{\mathbf{1}}\,\mathbf{s}\,\overrightarrow{\mathbf{0}} && \text{since } H_n^{\mathbf{1}}\mathbf{s}H_n^{\mathbf{0}} \text{ is valid;} \\ y &\coloneqq \overleftarrow{\mathbf{0}}\overrightarrow{\mathbf{0}} && \text{since } H_n^{\mathbf{0}}H_n^{\mathbf{0}} \text{ is valid;} \\ z &\coloneqq \overleftarrow{\mathbf{0}}\overrightarrow{\mathbf{1}} && \text{since } H_n^{\mathbf{0}}H_n^{\mathbf{1}} \text{ is valid.} \end{aligned}$$

These names x, y and z are points in X because the corresponding concatenations of blocks, listed at right, are valid words. ♦

Four-fold minimal self-joinings cannot exist

In the case of symbolic systems, four-fold topological minimal self-joinings does not happen because of the existence of asymptotic points. Any expansive system $T\colon X \to X$ has a pair of distinct points $x, y \in X$ which are future-asymptotic and a pair of points $p, q \in X$ which are past-asymptotic. The $T^{\times 4}$ orbit of $\langle p, q, x, y \rangle$ is not dense.

To handle general not-necessarily-expansive maps, we need the following result, theorem 10.30 in [G,H], due to S. Schwartzman. For completeness we include a demonstration: the neat proof below is slight variation of one due to Mike Boyle, Will Geller and Jim Propp. Say that two distinct points x and y are ***future*** ***ε-bounded*** if $|T^n x, T^n y| \le \varepsilon$ for all $n \ge 1$ and define "past ε-bounded" analogously.

THEOREM 2.1. *Suppose X is infinite. For any positive ε there exists a future ε-bounded pair.*

PROOF. Suppose, for the sake of contradiction, that there is no future ε-bounded pair. Let M be the supremum of those natural numbers N for which there exists a pair x, y with

$$|x, y| \ge \varepsilon \qquad \text{and} \qquad \forall n \in [1 \,..\, N)\colon \ \big|T^n x, T^n y\big| \le \varepsilon. \tag{$*$}$$

If M is infinite then there exists a pair x_N, y_N fulfilling $(*)$ and without loss of generality

$$x := \lim_{N\to\infty} x_N \quad \text{and} \quad y := \lim_{N\to\infty} y_N$$

exist. Evidently $|x, y| \ge \varepsilon$. Moreover, for each $k \in \mathbb{Z}_+$

$$\big|T^k x, T^k y\big| = \lim_{N\to\infty} \big|T^k x_N, T^k y_N\big| \le \lim_{N\to\infty} \varepsilon = \varepsilon$$

since T^k is continuous. Thus one is forced to conclude that M is finite after all.

Any power of T is uniformly continuous, X being compact. So there exists δ such that

$$|x, y| < \delta \qquad \Longrightarrow \qquad \big|T^m x, T^m y\big| \le \varepsilon \text{ for every } m \in [0 \,..\, M)$$

and so the maximality of M for $(*)$ implies that $\big|T^{-1} x, T^{-1} y\big| < \varepsilon$. Applying this iteratively yields that

$$\big|T^{-i} x, T^{-i} y\big| < \varepsilon \qquad \text{for } i = 1, 2, 3, \ldots$$

(In other words the family $\big\{T^{-i}\big\}_{i=1}^{\infty}$ is equicontinuous.) Now cover X with finitely many, L, open balls of diameter δ. Let $F \subset X$ be any collection of $L+1$ points. For each integer N the set $T^N(F)$ has $L+1$ points and so by the pigeon-hole principle there exists a pair $x, y \in F$ of distinct points such that $T^N x$ and $T^N y$ are in the same δ-ball. Hence

$$\forall m < N\colon \qquad \big|T^m x, T^m y\big| < \varepsilon. \tag{2.2}$$

Renaming x and y to x_N and y_N we can drop to a subsequence $N(i) \nearrow \infty$ for which

$$x_{N(1)} = x_{N(2)} = \ldots \quad \text{and} \quad y_{N(1)} = y_{N(2)} = \ldots.$$

Call these two points x and y. Now (2.2) holds for N replaced by any $N(i)$; hence for $N = \infty$. *A fortiori* x and y are future ε-bounded. ♦

The next several paragraphs are devoted to strengthening this result to produce an ε-bounded pair x, y in *distinct* orbits ie., with $y \notin \mathcal{O}(x)$.

The semigroup of continuous maps. Consider a compact metric space $\big(X, |\cdot,\cdot|\big)$. Then $\big(\mathbf{C}, \|\cdot,\cdot\|\big)$ is a complete metric space, where $\mathbf{C}$ denotes the set of continuous maps $f\colon X \to X$ under the supremum norm

$$\|f, g\| := \sup_{x\in X} |f(x), g(x)|.$$

Any $f \in \mathbf{C}$ is uniformly continuous and so there is a function $\varepsilon_f(\cdot)$, with $\varepsilon_f(r) \searrow 0$ as $r \searrow 0$, such that

$$\big|f(x), f(y)\big| \leq \varepsilon_f\big(|x, y|\big)$$

for all $x, y \in X$. Letting fg denote composition $f \circ g$, the triangle inequality yields

$$\|f'g', fg\| \leq \|f', f\| + \varepsilon_f\big(\|g', g\|\big)$$

for any $f', g', f, g \in \mathbf{C}$. Thus $(f, g) \mapsto fg$ is a continuous map from $\mathbf{C} \times \mathbf{C} \to \mathbf{C}$.

Let $\mathbf{H} \subset \mathbf{C}$ denote the subcollection of invertible maps, necessarily homeomorphisms. It is possible for a sequence of $f_n \in \mathbf{H}$ to have its limit in the complement $\mathbf{C} \setminus \mathbf{H}$. However

> *If $g \coloneqq \lim_{n\to\infty} f_n$ with each $f_n \in \mathbf{H}$ and $f_n f_k = f_k f_n$ for all n and k, then $g \in \mathbf{H}$.*

For whenever two homeomorphisms f and F commute then

$$\big\|f^{-1}, F^{-1}\big\| = \big\|f^{-1} \circ (fF), F^{-1} \circ (fF)\big\| = \big\|F, f\big\|.$$

Hence $\|{f_n}^{-1}, {f_k}^{-1}\| = \|f_n, f_k\|$ and so $\big\{{f_n}^{-1}\big\}_1^\infty$ is Cauchy and $G \coloneqq \lim_n {f_n}^{-1}$ exists. Letting I denote the identity map,

$$Gg = \lim_{n\to\infty} f_n^{-1} f_n = I = \lim_{n\to\infty} f_n f_n^{-1} = gG$$

and one concludes that g is invertible.

Define the norm $\|f\|$ to be $\|f, I\|$ and note that

$$\|fg\| \leq \|f, I\| + \varepsilon_I\big(\|g, I\|\big) = \|f\| + \|g\|$$

for all $f, g \in \mathbf{C}$. If $g \in \mathbf{H}$ then $\|g^{-1}\| = \|g^{-1} \circ g, I \circ g\| = \|g\|$. Thus

$$\|fg\| \geq \|f\| - \|g\|.$$

Fix henceforth a $T \in \mathbf{H}$ and let $\mathbf{A}(T)$ denote its automorphism group consisting of those $S \in \mathbf{H}$ such that $ST = TS$. Here is a standard type of argument for normed groups.

Uncountability Lemma. *Suppose there exist mutually commuting maps $\big\{P_n\big\}_1^\infty \subset \mathbf{A}(T)$, with $P_n \neq I$, such that $\|P_n\| \to 0$. Then for any positive ε the ball*

$$\big\{S \in \mathbf{A}(T) \;\big|\; \|S\| \leq \varepsilon\big\}$$

is uncountable.

Proof. By dropping to a subsequence of the $\{P_n\}$ we can arrange that $\sum_1^\infty \|P_n\| \leq \varepsilon$ and, for each K,

$$\big\|P_K\big\| > \sum_{k=K+1}^{\infty} \big\|P_k\big\|.$$

For a bit string $\vec{b} \in \{0,1\}^{\mathbb{N}}$ define $R_K \in \mathbf{A}(T)$ by

$$R_K \coloneqq P_1^{b(1)} P_2^{b(2)} \cdots P_K^{b(K)}.$$

Thus $\|R_N, R_{N+m}\| \leq \sum_{k=N+1}^{N+m} \|P_k\|$ which is less than $\|P_N\|$; hence $\{R_K\}_{K=1}^{\infty}$ is Cauchy. So $S_{\vec{b}} \coloneqq \lim_{K\to\infty} R_K$ exists and $\|S_{\vec{b}}\| \leq \varepsilon$. Moreover, $S_{\vec{b}}$ commutes with T because each R_K commutes with T and composition is continuous.

Finally, if bit strings $\vec{b}$ and $\vec{c}$ are distinct, then $S_{\vec{b}} \neq S_{\vec{c}}$. For let K be smallest for which $b(K) \neq c(K)$; say $b(K) = 1$ and $c(K) = 0$. Then

$$\begin{aligned}\|S_{\vec{b}} S_{\vec{c}}^{-1}\| &= \Big\|P_K \prod_{k=K+1}^{\infty} P_k^{b(k)} P_k^{-c(k)}\Big\| \\ &\geq \|P_K\| - \sum_{k=K+1}^{\infty} \|P_k^{b(k)-c(k)}\| \\ &\geq \|P_K\| - \sum_{k=K+1}^{\infty} \|P_k\|,\end{aligned}$$

which is positive. ♦

Proposition 2.3a. *Suppose $a \in X$ is a fixed point for T and there exists a sequence of points $z_n \to a$ and negative times $\{k(n)\}_1^{\infty}$ such that*

$$\inf_n |T^{k(n)} z_n, a| > 0.$$

Then for all $\delta > 0$ there exists a point $y \neq a$ for which $\langle y, a\rangle$ is future δ-bounded.

Proof. Reduce δ to smaller than the above infimum and discard those z_n which are not within δ of a. By moving each $k(n)$ closer to zero, if need be, we may now assume that

$$\forall \ell \in (k(n) \mathinner{..} 0]: \qquad |T^{\ell} z_n, a| \leq \delta. \tag{2.4}$$

Without loss of generality $y \coloneqq \lim_n T^{k(n)} z_n$ exists. Evidently $|y, a| \geq \delta$. If $\lim_n k(n) \neq -\infty$ then we could drop to a subsequence of $\{k(n)\}$ which was constant, say, always -7. But then

$$|y, a| \;=\; \lim_{n\to\infty} |T^{k(n)} z_n, a| \;=\; \lim_{n\to\infty} |T^{-7} z_n, T^{-7} a| \;=\; 0,$$

which is a contradiction. Thus $k(n) \to -\infty$. So given any $m > 0$, eventually $m + k(n)$ is in $(k(n) \mathinner{..} 0]$ for all large n. Consequently,

$$|T^m y, T^m a| = |T^m y, a| = \lim_{n\to\infty} |T^{m+k(n)} z_n, a| \leq \delta$$

by inequality (2.4). Thus y and a are future δ-bounded. ♦

PROPOSITION 2.3b. *Suppose $\#X = \infty$ and every minimal set is finite. Then for all ε there exists a future ε-bounded pair in distinct orbits.*

PROOF. *First suppose X consists only of periodic points.* Then we can pick a convergent sequence $\{z_n\}_1^\infty$ of points in *distinct* orbits, with

$$\forall n : \quad \mathcal{O}_T(z_n) \neq \mathcal{O}_T(a)$$

where $a := \lim_n z_n$. Thus a is a fixed point of $S := T^p$, where p denotes the least period of a. Choose δ sufficiently small that

$$\text{For all } y \in X \text{ and } m \in [0 \,..\, p), \qquad |y, a| \leq \delta \implies \bigl|T^m y, T^m a\bigr| \leq \varepsilon.$$

Consequently, having taken δ smaller than the distance between any two of the p points of $\mathcal{O}_T(a)$, it will suffice to show that

There exists a point $y \neq a$ such that the pair $\langle y, a\rangle$ is future δ-bounded with respect to S. (2.5)

We may as well assume that no z_n can play the role of y and so there exist integers $\bigl\{k(n)\bigr\}_n$ such that

$$\bigl|S^{k(n)} z_n, a\bigr| \;\geq\; \delta\,.$$

Since each z_n is a periodic point, we may take each $k(n)$ to be negative. The preceding proposition may now be applied to S to produce a point y which fulfills (2.5).

There exists an $x \in X$ with infinite T-orbit, is the other possibility. Pick $j(n) \nearrow \infty$ such that $a := T^{j(n)}x$ exists and is a point in a minimal set. Defining p, S and δ as above condition (2.5), it suffices to establish that condition.

Dropping to a subsequence of $\bigl\{j(n)\bigr\}_n$ we can assume that all the $j(n)$ are congruent modulo p; say, to r. By replacing x by $T^{-r}x$ we can write

$$S^{-k(n)}x \xrightarrow[n]{} a$$

where $k(n)$ is the negative number $-[j(n) - r]/p$. Setting $z_n := S^{-k(n)}x$ we have that

$$\inf_n \bigl|S^{k(n)} z_n, a\bigr| \;=\; |x, a| \;>\; 0$$

and so we may again apply to S the preceding proposition. ♦

We now can obtain the desired strengthening of Schwartzman's Theorem.

BOUNDEDNESS THEOREM, 2.6. *Suppose T is a homeomorphism of an infinite compact metric space X. Then for every positive ε there exist a future ε-bounded pair $x, y \in X$ with $\mathcal{O}(x) \neq \mathcal{O}(y)$.*

PROOF. By the foregoing proposition we may assume that T has an infinite minimal set and so we may take T to be minimal. Suppose, for the sake of contradiction, that the conclusion fails for ε. By Schwartzman's theorem there exist distinct y and z which are future ε-bounded. Hence

$z = T^\ell y$ for some integer ℓ, and $T^\ell \neq I$. By minimality, for any point x there exists $N(i) \nearrow \infty$ such that $T^{N(i)} y \to x$ and thus $T^{N(i)} z \to T^\ell x$. Hence $|x, T^\ell x| \leq \varepsilon$ for all $x \in X$. In other words,

$$\left\|T^\ell\right\| \leq \varepsilon.$$

Choosing a sequence $\varepsilon_n \searrow 0$ we obtain integers $\ell(n)$ such that

$$\left\|T^{\ell(n)}\right\| \longrightarrow 0\,, \qquad \text{as } n \to \infty,$$

and $T^{\ell(n)} \neq I$. Thus we are entitled to the conclusion of the *Uncountability lemma*.

Fix some $x \in X$. For any S in the ε-ball in $\mathbf{A}(T)$ we have that

$$\left|T^n x, T^n(Sx)\right| \;=\; \left|T^n x, S(T^n x)\right| \;\leq\; \|S\| \;\leq\; \varepsilon$$

for all $n \in \mathbb{Z}$. So the conclusion of our purported theorem can fail only if the point $S(x)$ lies in the T-orbit of x, a countable set. Since there are uncountably many S in the ε-ball, there must exist distinct $S_1, S_2 \in \mathbf{A}(T)$ such that $S_1(x)$ equals $S_2(x)$. But then $S_1(T^n x) = S_2(T^n x)$ for every integer n and so S_1 and S_2 agree on the T-orbit of x, a dense set. Thus $S_1 = S_2$. ♦

No map has 4-fold topological minimal self-joinings. Suppose $T\colon X \to X$ is minimal and $\#X = \infty$. Fixing a small ε, the *Boundedness theorem* gives us future and past bounded pairs $f, f', p, p' \in X$ such that

$$\left|T^n f, T^n f'\right| \leq \varepsilon \;\text{ and }\; \left|T^{-n} p, T^{-n} p'\right| \leq \varepsilon \qquad \text{for } n = 0, 1, 2, \ldots$$

with $\mathcal{O}(f) \neq \mathcal{O}(f')$ and $\mathcal{O}(p) \neq \mathcal{O}(p')$.

First suppose that the four points manage to inhabit only two orbits; say $p = T^i(f)$ and $p' = T^{i+s}(f')$ for some integers i and s. Any point $\langle x, y\rangle$ of $\overline{\mathcal{O}}_{T\times T}(f, f')$ must satisfy $|x, y| \leq \varepsilon$ or $|x, T^s y| \leq \varepsilon$, depending on whether it is in the future or past orbit closure of $\langle f, f'\rangle$. If T has even just 2-fold tmsj then $\overline{\mathcal{O}}_{T\times T}(f, f')$ equals $X \times X$ and so, fixing any point y, this implies that

$$\mathbf{B}(y) \cup \mathbf{B}(T^s y) = X\,,$$

where $\mathbf{B}(y)$ is the closed radius-ε ball centered at y. But we could have taken ε to be so small that no two radius-ε balls can cover X.

Suppose instead that $\{f, f', p, p'\}$ inhabit three orbits; say $T^s(p') = f'$ and $p \notin \mathcal{O}(f)$. Then the triple $\langle f, f', T^s p\rangle$ belies 3-fold tmsj since its $T^{\times 3}$ orbit closure contains no triple of the form $\langle z, z', z\rangle$ with $|z, z'| > \varepsilon$.

Similarly, if $\{f, f', p, p'\}$ inhabit four orbits then

$$\langle z, z', z, z'\rangle \notin \overline{\mathcal{O}}_{T^{\times 4}}(f, f', p, p')\,,$$

which prohibits 4-fold topological minimal self-joinings.

References

[G,H] W. Gottschalk, G. Hedlund, *Topological Dynamics*, vol. 36, AMS Colloquium Publications, Providence, RI, 1955.

[A,M] J. Auslander, N. Markley, *Minimal flows of finite almost periodic rank*, Ergodic Theory and Dynamical Systems **8** (1988), 155-172.

[J] A. del Junco, *On minimal self-joinings in topological dynamics*, Ergodic Theory and Dynamical Systems **7**, part 2 (1987), 211–227.

[J,K] A. del Junco, M. Keane, *On generic points in the cartesian square of Chacón's transformation*, Ergodic Theory and Dynamical Systems (to appear).

[Ke] H. Keynes, *The structure of weakly mixing minimal transformation groups*, Illinois J. Math **15** (1971), 475-489.

[K] J.L. King, *Joining-rank and the structure of finite rank mixing transformations*, J. d'Analyse Math. **51** (1988), 182–227.

[K,T] J.L. King, J-P Thouvenot, *A canonical structure theorem for finite joining-rank maps*, J. d'Analyse Math. **56** (1991), 211–230.

[M] N. Markley, *Topological minimal self-joinings*, Ergodic Theory and Dynamical Systems **3** (1983), 579–599.

[R] D. Rudolph, *An example of a measure-preserving map with minimal self-joinings, and applications*, J. d'Analyse Math. **35** (1979), 97-122.

Filename: `Article/10Tmsj/tmsj.ams.tex`
As of: *Mon Aug 18, 1997* Typeset: *6September2015*