NON-HOMOGENEOUS EXTENSIONS OF CANTOR MINIMAL SYSTEMS

ROBIN J. DEELEY, IAN F. PUTNAM, AND KAREN R. STRUNG

Abstract. Floyd gave an example of a minimal dynamical system which was an extension of an odometer and the fibres of the associated factor map were either singletons or intervals. Gjerde and Johansen showed that the odometer could be replaced by any Cantor minimal system. Here, we show further that the intervals can be generalized to cubes of arbitrary dimension and to attractors of certain iterated function systems. We discuss applications.

1. Introduction and statement of results

We consider dynamical systems consisting of a compact space, X, together with a homeomorphism, $\varphi : X \to X$. We say that such a system is *minimal* if the only closed sets $Y \subseteq X$ such that $\varphi(Y) = Y$ are $Y = X$ and $Y = \emptyset$. Equivalently, for every x in X, its orbit, $\{\varphi^n(x) \mid n \in \mathbb{Z}\}$, is dense in X.

There are a number of examples of such systems: rotation of the circle through an angle which is an irrational multiple of 2π, odometers and certain diffeomorphisms of spheres of odd dimension $d \geq 3$ constructed by Fathi and Herman [5].

All of these examples share one common feature: the spaces involved are homogeneous. There are several ways to make this more precise, but one simple way would be to observe that the group of homeomorphisms acts transitively on the points.

In [6], Floyd gave the first example of a minimal system where the space is not homogeneous in this (or an even stronger) sense. Floyd began with the 3^∞-odometer, (X, φ), which is a minimal system with X compact, metrizable, totally disconnected and without isolated points. Any two such spaces are homeomorphic and we refer to such a space as a *Cantor set*. Floyd then constructed another minimal system, $(\tilde{X}, \tilde{\varphi})$, together with a continuous surjection $\pi : \tilde{X} \to X$ satisfying $\pi \circ \tilde{\varphi} = \varphi \circ \pi$. In general, we refer to such a map as a *factor map*, we say that (X, φ) is a *factor* of $(\tilde{X}, \tilde{\varphi})$ and that $(\tilde{X}, \tilde{\varphi})$ is an *extension* of (X, φ). In Floyd's example, some points x in X have $\pi^{-1}\{x\}$ homeomorphic to the unit interval, $[0, 1]$, while for others, it is a single point. It is then quite easy to see, using the fact that X is totally disconnected, that the space $\tilde{X}$ has some connected components which are single points and some homeomorphic to the interval.

Date: November 20, 2020.

2010 *Mathematics Subject Classification.* 37B05, 46L35, 46L85, 19K99.

Key words and phrases. minimal dynamics.

RJD is currently funded by NSF Grant DMS 2000057 and was previously funded by Simons Foundation Collaboration Grant for Mathematicians number 638449. KRS is currently funded by GAČR project 20-17488Y and RVO: 67985840 and part of this work was carried out while funded by Sonata 9 NCN grant 2015/17/D/ST1/02529 and a Radboud Excellence Initiative Postdoctoral Fellowship. IFP is supported in part by an NSERC Discovery Grant.

This example has been generalized in several ways (for example [1, 2, 7, 8]). In Floyd's example, the points x with $\pi^{-1}\{x\}$ infinite all lie in a single orbit. Haddad and Johnson in [8] showed that the set of such x could be much larger and even have positive measure under the unique invariant measure for (X, φ). More importantly for our purposes, Gjerde and Johansen [7] showed that the 3^∞-odometer could be replaced with any minimal system, (X, φ), with X a Cantor set. Their principal tool was the Bratteli–Vershik model for such systems [12, Chapters 4 and 5]. We will describe this in more detail in Section 2.

Our aim here is to show that the interval, $[0, 1]$, appearing as $\pi^{-1}\{x\}$, can be replaced by more complicated spaces. We are particularly interested in the case of the n-dimensional cube (that is, $[0, 1]^n$), for any positive integer n.

Although it is natural to generalize to more complicated spaces, let us explain briefly why we want such a result in the specific case of $[0, 1]^n$. The Elliott program aims to show that a broad class of C^*-algebras may be classified up to isomorphism by their K-theory [4]. One very useful way of constructing C^*-algebras is via groupoids [13] and it becomes a natural question: which C^*-algebras in the Elliott scheme can be realized via a groupoid construction? In view of the classification results themselves, this amounts to constructing groupoids whose associated C^*-algebras are classifiable and have some prescribed K-theory. If one begins with a minimal action of the integers on a Cantor set, it is known that the K_0-group is a simple acyclic dimension group and K_1 is the integers [9]. Moreover, any such K-theory can be realized from such a system [9].

In another direction, if one takes a minimal action, φ, of the integers on some space X and considers a closed, non-empty subset $Y \subseteq X$ such that Y meets each orbit at most once, one can construct the associated "orbit-breaking groupoid": the equivalence relation where the classes are either the original orbits of φ which do not meet Y or the half-orbits, split at Y. The change in K-theory passing from the crossed product C^*-algebra to the orbit-breaking subalgebra can be computed, essentially in terms of the K-theory of the space Y (see [11] for details).

Marrying these two ideas would seem to generate many interesting groupoids, except that the choices for $K^*(Y)$, where Y is a closed subset of the Cantor set, are very limited. Here, we would like to replace the dynamics (X, φ) with $(\tilde{X}, \tilde{\varphi})$, without changing the associated K-theory, but allowing us to find more interesting spaces Y inside of $[0, 1]^n \cong \pi^{-1}\{x\}$. These C^*-algebraic applications can be found in [3].

Our construction and proof follow those of Gjerde and Johansen in [7] quite closely and, in turn, their proof is quite similar to Floyd's original one [6]. One added feature here is that we use the framework of iterated function systems, as this allows us to replace the interval, $[0, 1]$, with the more complicated spaces.

Following usual conventions (see for example [10]) an *iterated function system* consists of a metric space, (C, d_C), and $\mathcal{F}$, a finite collection of maps $f : C \to C$ with the property that there is a constant $0 < \lambda < 1$ such that $d_C(f(x), f(y)) \leq \lambda d_C(x, y)$, for all x, y in C and f in $\mathcal{F}$. In particular, each map is continuous. We will require a few extra properties.

Definition 1.1. Let $(C, d_C, \mathcal{F})$ be an iterated function system. We say it is *compact* if the metric space (C, d_C) is compact. We also say it is *invertible* if

(1) each f in $\mathcal{F}$ is injective, and
(2) $\cup_{f\in\mathcal{F}} f(C) = C$.

The term "invertible" is meant to indicate that each map f in $\mathcal{F}$ has an inverse, $f^{-1} : f(C) \to C$. It is not ideal as it does not rule out the possibility that the images of the various f's overlap.

Of course, the restriction that each map is injective is quite important. On the other hand, it is well-known that any compact iterated function system has a fixed point set and the restriction of the maps to this set will satisfy the invariance condition [10, Section 3].

We list several simple examples of relevant iterated function systems. The first is the one originally used by Floyd [6] along with the subsequent examples [1, 2, 7, 8].

Example 1.2. Let $C = [0, 1]$, $f_i(x) = 2^{-1}(x + i)$ for x in $[0, 1]$ and $i = 0, 1$, and $\mathcal{F} = \{f_0, f_1\}$.

The next example is a fairly simple generalization of the last, but it is important as this is the example we need in our applications in [3].

Example 1.3. Let n be any positive integer, $C = [0, 1]^n$, $f_\delta(x) = 2^{-1}(x + \delta)$, for each x in $[0, 1]^n, \delta \in \{0, 1\}^n$ and $\mathcal{F} = \{f_\delta \mid \delta \in \{0, 1\}^n\}$.

Example 1.4. A minor variation on the last example would be to use instead $f_\delta(x) = 3^{-1}(x + \delta)$, for each x in $[0, 1]^n, \delta \in \{0, 1, 2\}^n$. On the other hand, if we instead let $\mathcal{F} = \{f_\delta \mid \delta \in \{0, 2\}\}$ when $n = 1$, or $\mathcal{F} = \{f_\delta \mid \delta \in \{0, 1, 2\}^2, \delta \neq (1, 1)\}$ for $n = 2$, this now fails the invariance condition of our definition. As mentioned above, standard results on iterated function systems show that C contains a unique closed set and restricting our maps to that set then satisfies all the desired conditions. Notice that when $n = 1$, the set in question is the Cantor ternary set, while for $n = 2$, it is the Sierpinski carpet [14].

Our main result is the following.

Theorem 1.5. *Let $(C, d_C, \mathcal{F})$ be a compact, invertible iterated function system and let (X, φ) be a minimal homeomorphism of the Cantor set. There exists a minimal extension, $(\tilde{X}, \tilde{\varphi})$ of (X, φ) with factor map $\pi : (\tilde{X}, \tilde{\varphi}) \to (X, \varphi)$ such that, for each x in X, $\pi^{-1}\{x\}$ is a single point or is homeomorphic to C. Moreover, both possibilities occur.*

Theorem 1.6. *Let $(C, d_C, \mathcal{F})$ be a compact, invertible iterated function system and let (X, φ) be a minimal homeomorphism of the Cantor set. If C is contractible, then the minimal extension $(\tilde{X}, \varphi)$ and the factor map $\pi : (\tilde{X}, \tilde{\varphi}) \to (X, \varphi)$ may be chosen so that*

$$\pi^* : K^*(X) \to K^*(\tilde{X})$$

is an isomorphism and so that π induces a bijection between the respective sets of invariant measures.

2. The construction and proofs

Just as for Gjerde and Johansen [7], we make critical use of the Bratteli–Vershik model for minimal systems on the Cantor set [12]. Briefly, the Bratteli-Vershik model takes some simple combinatorial data (an ordered Bratteli diagram) and produces a minimal homeomorphism of the Cantor set. In fact, every minimal homeomorphism of the Cantor set is produced in this way. A standard reference for Cantor minimal system is [12], in particular see Chapters 4 and 5 for more on the Bratteli–Vershik model.

We begin with a Bratteli diagram, (V, E), consisting of a vertex set V written as a disjoint union of finite, non-empty sets $V_n, n \geq 0$, with $V_0 = \{v_0\}$, and an edge set written as a disjoint union of finite, non-empty sets $E_n, n \geq 1$. Each edge e in E_n has a source, $s(e)$, in V_{n-1} and range, $r(e)$, in V_n. We may assume (see [12]) that our diagram has *full edge connections*, that is, every pair of vertices from adjacent levels is connected by at least one edge. We define the space X_E to consist of all infinite paths in the diagram, beginning at v_0. That is, a point $x = (x_1, x_2, \ldots)$, $x_n \in E_n, r(x_n) = s(x_{n+1})$. This space is endowed with the metric

$$d_E(x, y) = \inf\{2^{-n} \mid n \geq 0, x_i = y_i, 1 \leq i \leq n\}.$$

In addition, we may assume that the edge set E is endowed with an order such that two edges e, f are comparable if and only if $r(e) = r(f)$. The set of maximal edges and the set of minimal edges each form a tree and we assume that our diagram is *properly ordered*, meaning that each contains exactly one infinite path. Two finite paths from v_0 to V_n can be compared if they have the same range vertex by using a right-to-left lexicographic order. Infinite paths may be compared in a similar way: two paths are *cofinal* if they differ in only finitely many entries and can be compared using a right-to-left lexicographic order. The Bratteli–Vershik map, φ_E, takes an infinite path to its successor, and the unique infinite path with all edges maximal to the unique infinite path with all edges minimal. The system (X_E, φ_E) is a minimal Cantor system (provided X_E is infinite). Moreover, every minimal Cantor system is topologically conjugate to a Bratteli–Vershik system. In view of this, we may assume that $(X, \varphi) = (X_E, \varphi_E)$, for some properly ordered Bratteli diagram, (V, E).

We note that there is an (essentially) unique ordered Bratteli diagram with $\#V_n = 1$ and $\#E_n = 3$, for all $n \geq 1$, and the associated Bratteli–Vershik map is the 3^∞-odometer considered by Floyd. More generally, an odometer is any system with $\#V_n = 1$, for all $n \geq 1$ (see for example [12, Chapter 11 Section 8]).

Recall that, in addition to the Cantor minimal system, we also have $(C, d_C, \mathcal{F})$, which is a compact, invertible iterated function system. Our final ingredient involves this system. To each edge e in E, we assign a function, denoted f_e, in $\mathcal{F} \cup \{\mathrm{id}_C\}$, where id_C is the identity function on C. We assume that this assignment satisfies the following three conditions:

(1) if e is either maximal or minimal, then $f_e \neq \mathrm{id}_C$,
(2) for every v in V, we have $\cup_{s(e)=v, f_e \neq \mathrm{id}_C} f_e(C) = C$,
(3) the set $\{e \in E \mid f_e = \mathrm{id}_C\}$ contains an infinite path.

Let us first mention that the following weaker third condition will suffice: there exists an infinite path $(e_1, e_2, \ldots)$ such that $f_{e_n} = \mathrm{id}_C$, for infinitely many n. Secondly, the fact that we can find a properly ordered Bratteli diagram satisfying these can be seen as follows. First, take any minimal Bratteli–Vershik (X_E, φ_E) with X_E infinite. It can be telescoped until every pair of vertices at adjacent levels have at least $\#\mathcal{F} + 1$ edges (see [7, page 94] or [12, page 22–23]). Select some infinite path, $(e_1, e_2, \ldots)$, which avoids all maximal and minimal edges and set $f_{e_n} = \mathrm{id}_C$ for all $n \geq 1$. Then f may be chosen so that it maps the set $s^{-1}\{v\} - \{e_1, e_2, \ldots\}$ surjectively to $\mathcal{F}$, for every v in V.

In the construction of Gjerde and Johansen [7], the edges e with $f_e = \mathrm{id}_C$ form a single infinite path. On the other hand, allowing more general subdiagrams, we capture examples such as those given in [8].

We now construct the system $(\tilde{X}, \tilde{\varphi})$. We endow $X_E \times C$ with the metric

$$d\left((x,c),(y,d)\right) = \max\{d_E(x,y), d_C(c,d)\},$$

for x, y in X_E and c, d in C. For each $n \geq 1$, define

$$\tilde{X}_n = \{(x,c) \in X_E \times C \mid c \in f_{x_1} \circ \cdots \circ f_{x_n}(C)\}.$$

It is immediate that each $\tilde{X}_n$ is closed and non-empty and that $\tilde{X}_n \supseteq \tilde{X}_{n+1}$. We let $\tilde{X} = \cap_{n \geq 1} \tilde{X}_n$, which is also closed and non-empty.

The quotient map $\pi : \tilde{X} \to X_E$ is defined by $\pi(x,c) = x$.

Lemma 2.1. *Suppose that $x \in X_E$. Then exactly one of the following hold:*

(1) Type 1: for infinitely many $n \geq 1$, $f_{x_n} \neq \mathrm{id}_C$,
(2) Type 2: there exists $n \geq 1$ such that $f_{x_i} = \mathrm{id}_C$, for all $i \geq n$.

Moreover, in the Type 1 case, $\pi^{-1}\{x\}$ consists of a single point, which we denote by (x, c_x). In the Type 2 case, we have

$$\pi^{-1}\{x\} = \{x\} \times f_{x_1} \circ \cdots \circ f_{x_n}(C)$$

which is homeomorphic to C, since each f_e is continuous and injective.

Proof. The set of n such that $f_{x_n} = \mathrm{id}_C$ is either finite or infinite, hence the two conditions are mutually exclusive and the only possibilities.

In the first case, we have

$$\mathrm{diam}(f_{x_1} \circ \cdots \circ f_{x_n}(C)) \leq \lambda^m \mathrm{diam}(C),$$

where m is the number of $1 \leq l \leq n$ with $f_{x_l} \neq \mathrm{id}_C$. The conclusion follows.

In the second case, it is clear that

$$(\{x\} \times C) \cap \tilde{X}_i = \{x\} \times f_{x_1} \circ \cdots \circ f_{x_n}(C),$$

for any $i \geq n$ and now, taking the intersection over all i, we obtain the desired conclusion. □

We are now prepared to define our self-map of $\tilde{X}$ and show it is a minimal homeomorphism. Let (x,c) be in $\tilde{X}$. We consider three cases separately. First, we assume that x is Type 1 and that $x \neq x^{\max}$. Then x contains a non-maximal edge and we let n be the first such edge. Thus, we have $\varphi_E(x) = (y_1, \ldots, y_n, x_{n+1}, \ldots)$, for some path $y_1, \ldots, y_n$ with $r(y_n) = r(x_n)$. The fact that x is Type 1 means that x uniquely determines $c = c_x$ and we define $\tilde{\varphi}(x,c) = (\varphi_E(x), c_{\varphi_E(x)})$. To see this is well-defined, it suffices to note that, as x is Type 1 and $\varphi_E(x)$ differs from x in only finitely many entries, $\varphi_E(x)$ is also Type 1 and $c_{\varphi_E(x)}$ is well-defined.

Secondly, suppose that $x = x^{\max}$. It follows that, for all n, x_n is maximal and hence $f_{x_n} \neq \mathrm{id}_C$. So $x^{\max}$ is also Type 1 and c is once again uniquely determined as $c = c_x$. The same argument shows that $x^{\min}$ is Type 1 as well. We define $\tilde{\varphi}(x^{\max}, c_{x^{\max}}) = (x^{\min}, c_{x^{\min}})$.

Finally, we consider the case that x is Type 2. In particular, this implies that x is not equal to $x^{\max}$. Let n be as in the definition of Type 2. (Notice that such an n is not unique: we deal with this issue shortly.) As mentioned, x_n is not maximal. Hence $\varphi_E(x) = (y_1, \ldots, y_n, x_{n+1}, \ldots)$, for some path $y_1, \ldots, y_n$ with $r(y_n) = r(x_n)$. (It is worth noting that it is possible that $x_n = y_n$.) Since c is in $f_{x_1} \circ \cdots \circ f_{x_n}(C)$ we can define

$$\tilde{\varphi}(x,c) = (\varphi_E(x), f_{y_1} \circ \cdots \circ f_{y_n} \circ f_{x_n}^{-1} \circ \cdots \circ f_{x_1}^{-1}(c)).$$

As mentioned in the previous paragraph the choice of n is not unique and hence we need to check that the above definition is independent of the choice of n. This follows from the observation that if n' is another such integer, we may assume that $n' > n$, and then all the maps $f_{x_{n+1}}, \ldots, f_{x_{n'}}$ are all equal to id_C.

The proof that $\tilde{\varphi}$ is bijective is as follows: by simply reversing the order on the edge set, the construction will yield another map which is easily seen to be the inverse of $\tilde{\varphi}$.

We claim that $\tilde{\varphi}$ is continuous. To show this, it will be convenient to define, for any path $p = (p_1, \ldots, p_n)$ in (V, E) from v_0 to V_n, the sets

$$\tilde{X}_n(p) = \{x \in X_E \mid (x_1, \ldots, x_n) = p\} \times f_{p_1} \circ \cdots \circ f_{p_n}(C).$$

It is an easy matter to check that $\tilde{X}_n(p)$ is a closed subset of $\tilde{X}_n$, that for $p \neq q$ of length n, $\tilde{X}_n(p)$ and $\tilde{X}_n(q)$ are disjoint, and that the union of all such sets over paths of length p is exactly $\tilde{X}_n$. From this, it follows that each such set is also an open subset $\tilde{X}_n$. It also follows that $\tilde{X} \cap \tilde{X}_n(p)$ is clopen in $\tilde{X}$.

Let p be any path in E from v_0 to V_n which is not maximal. Let q be its successor among such paths. Define a map $\psi : \tilde{X}_n(p) \to \tilde{X}_n(q)$ by

$$\psi(z, c) = (\varphi_E(z), f_{q_1} \circ \cdots \circ f_{q_n} \circ f_{p_n}^{-1} \circ \cdots \circ f_{p_1}^{-1}(c)),$$

for (z, c) in $\tilde{X}_n(p)$. Observe that this is clearly a homeomorphism.

We will show $\psi|_{\tilde{X} \cap \tilde{X}_n(p)} = \tilde{\varphi}|_{\tilde{X} \cap \tilde{X}_n(p)}$. From this and the fact that $\tilde{X} \cap \tilde{X}_n(p)$ is clopen in $\tilde{X}$, it follows that $\tilde{\varphi}$ is continuous on $\tilde{X} \cap \tilde{X}_n(p)$.

First, suppose that z is a Type 2 point in $\tilde{X} \cap \tilde{X}_n(p)$. Choose m such that $f_{z_i} = \mathrm{id}_C$, for all $i \geq m$, Without loss of generality, we may assume that $m > n$. We know that $\varphi_E(z) = (q_1, \ldots, q_n, z_{n+1}, \ldots)$ and hence, for any c in C,

$$\begin{aligned} \tilde{\varphi}(z, c) &= ((q_1, \ldots, q_n, z_{n+1}, \ldots), f_{q_1} \circ \cdots \circ f_{q_n} \circ f_{z_{n+1}} \circ \cdots \circ f_{z_m} \\ &\quad \circ f_{z_m}^{-1} \circ \cdots f_{z_{n+1}}^{-1} \circ f_{p_n}^{-1} \circ \cdots \circ f_{p_1}^{-1}(c)) \\ &= ((q_1, \ldots, q_n, z_{n+1}, \ldots), f_{q_1} \circ \cdots \circ f_{q_n} \circ f_{p_n}^{-1} \circ \cdots \circ f_{p_1}^{-1}(c)) \\ &= \psi(z, c). \end{aligned}$$

Now, we consider a point z of Type 1 in $\tilde{X} \cap \tilde{X}_n(p)$. The same argument as above shows that, for any $m > n$, we have

$$\psi(\tilde{X} \cap \tilde{X}_n(p, z_{n+1}, \ldots, z_m)) = \tilde{\varphi}(\tilde{X} \cap \tilde{X}_n(p, z_{n+1}, \ldots, z_m)).$$

The point $\psi(z, c_z)$ is the unique point which lies in the left-hand side for every $m > n$, while $\tilde{\varphi}(z, c_z)$ is the unique point that lies in the right-hand side, for every $m > n$. Hence, we conclude they are equal.

We have now shown that $\tilde{\varphi}$ is continuous on every set $\tilde{X}_n(p)$, where p is a finite path which is not maximal. But such sets, allowing both n and p to vary, contain every point of $\tilde{X}$, except $(x^{\max}, c_{x^{\max}})$. It follows from general topological arguments using the facts that $\tilde{\varphi}$ is a bijection and that $\tilde{X}$ is compact, that $\tilde{\varphi}$ is continuous everywhere.

Finally, we need to show that $(\tilde{X}, \tilde{\varphi})$ is minimal. Let (x, c) and (y, d) be in $\tilde{X}$ and let $\epsilon > 0$. It suffices for us to find a point (z, e) in the orbit of (x, c) within distance ϵ of (y, d).

We first consider the case that y is of Type 1. We choose n sufficiently large so that $2^{-n} < \epsilon$ and so that, if we let m be the number of $1 \leq l \leq n$ such that

$f_{y_l} \neq \mathrm{id}_C$, then $\lambda^m \mathrm{diam}(C) < \epsilon$. It follows that $\tilde{X} \cap \tilde{X}_n(y_1, \ldots, y_n)$ is a clopen set containing (y, d) of diameter less than ϵ.

We define z in X_E as follows: $z_i = y_i$ for all $1 \leq i \leq n$, z_{n+1} is any edge in E_{n+1} with $s(z_{n+1}) = r(y_n)$, and $r(z_{n+1}) = s(x_{n+2})$ and $z_i = x_i$, for any $i \geq n+2$. It follows at once that z and x are cofinal. Hence, there is an integer k such that $\varphi_E^k(x) = z$. It is then clear that $\tilde{\varphi}^k(x, c)$ is in $\tilde{X} \cap \tilde{X}_n(y_1, \ldots, y_n)$ and hence within ϵ of (y, d).

Now we consider the case that y is of Type 2. First, choose m sufficiently large so that $2^{-m} < \epsilon$ and so that $f_{y_n} = \mathrm{id}_C$, for all $n > m$. Define $z_i = y_i$, for all $1 \leq i \leq m$. Observe that d is in $f_{y_1} \circ \cdots \circ f_{y_m}(C)$ and we let $d_m = f_{y_1}^{-1} \circ \cdots \circ f_{y_m}^{-1}(d)$. Choose $n > m$ such that $\lambda^{n-m} \mathrm{diam}(C) < \epsilon$. For each $m < l \leq n$, we define z_l and d_l inductively, using our second hypothesis on the assignment $e \to f_e$, so that $s(z_l) = r(z_{l-1})$, $f_{z_l} \neq \mathrm{id}_C$ and $f_{z_l}(d_l) = d_{l-1}$. This obviously implies that $f_{z_{m+1}} \circ \cdots \circ f_{z_n}(d_n) = d_m$. We define z_{n+1} to be any edge with $s(z_{n+1}) = r(z_n)$ and $r(z_{n+1}) = s(x_{n+2})$ and then $z_l = x_l$, for $l \geq n+2$. This means that z is a path in X_E which is cofinal with x. Hence, there is an integer k such that $\varphi^k(x) = z$.

We claim that $\tilde{\varphi}^k(x, c)$, which we denote (z, e), is within ϵ of (y, d). First, as $\varphi_E^k(x)_i = z_i = y_i$, for all $1 \leq i \leq n$, we have $d(\varphi_E^k(x), y) < 2^{-n} \leq 2^{-m} < \epsilon$. This also means that e is in $f_{z_1} \circ \cdots \circ f_{z_n}(C)$. On the other hand, we know that

$$\begin{aligned} d &= f_{y_1} \circ \cdots \circ f_{y_m}(d_m) \\ &= f_{z_1} \circ \cdots \circ f_{z_m}(d_m) \\ &= f_{z_1} \circ \cdots \circ f_{z_n}(d_n) \\ &\in f_{z_1} \circ \cdots \circ f_{z_n}(C). \end{aligned}$$

We also know that, since $f_{z_l} \neq \mathrm{id}_C$, for $m < l \leq n$, we have

$$\begin{aligned} \mathrm{diam}(f_{z_1} \circ \cdots \circ f_{z_n}(C)) &\leq \mathrm{diam}(f_{z_{m+1}} \circ \cdots \circ f_{z_n}(C)) \\ &\leq \lambda^{n-m} \mathrm{diam}(C) \\ &< \epsilon. \end{aligned}$$

This completes the proof of Theorem 1.5.

We now turn to the proof of Theorem 1.6. Let $\pi_n : \tilde{X}_n \to X_E$ be the obvious extension of π: simply projecting onto the first factor. For any path p, from v_0 to V_n: π_n maps $\tilde{X}_n(p)$ to $X(p) = \{x \in X_E \mid x_i = p_i, 1 \leq i \leq n\}$. As we assume C is contractible, so is $f_{p_1} \circ \cdots \circ f_{p_n}(C)$ and this map induces an isomorphism on K-theory. Taking the union over all paths p of length n, we see that π_n induces an isomorphism from $K^*(\tilde{X}_n) \cong \oplus_p K^*(\tilde{X}_n(p))$ to $K^*(X_E) \cong \oplus_p K^*(X(p))$.

As $\tilde{X} = \cap_n \tilde{X}_n$, it is also the inverse limit of

$$\tilde{X}_1 \leftarrow \tilde{X}_2 \leftarrow \tilde{X}_2 \leftarrow \cdots$$

where the maps are the inclusions. As K-theory is continuous, the conclusion follows.

If we choose our assignment $e \to f_e$ so that the edges which are assigned id_C form a single infinite path, then π is a bijection, except on a single orbit of φ_E. This orbit has measure zero under any invariant probability measure on X, so it lifts to a unique $\tilde{\varphi}$-invariant measure on $\tilde{X}$.

Acknowledgements. The authors thank the Banff International Research Station and the organisers of the workshop Future Targets in the Classification Program

for Amenable C*-Algebras where this project was initiated. Thanks also to the Department of Mathematics and Statistics at the University of Victoria and the Department of Mathematics of the University of Colorado Boulder for research visits facilitating this collaboration. Work on the project was also facilitated by the Lorentz Center where the first and third authors attended a conference on Cuntz–Pimsner algebras in June 2018. The authors are also grateful to the referee for a thorough reading of an earlier version of the paper and many helpful suggestions.

References

[1] Joseph Auslander. Mean-L-stable systems. *Illinois J. Math.*, 3(4):566–579, 12 1959.
[2] Joseph Auslander. *Minimal flows and their extensions*, volume 153 of *North-Holland Mathematics Studies*. North-Holland Publishing Co., Amsterdam, 1988. Notas de Matemática [Mathematical Notes], 122.
[3] Robin J. Deeley, Ian F. Putnam, and Karen R. Strung. Constructions in minimal amenable dynamics and applications to the classification of C*-algebras. Preprint, 2019.
[4] George A. Elliott. The classification problem for amenable C^*-algebras. In *Proceedings of the International Congress of Mathematicians, Zürich, 1994*, volume 1,2, pages 922–932. Birkhäuser, Basel, 1995.
[5] Albert Fathi and Michael R. Herman. Existence de difféomorphismes minimaux. pages 37–59. Astérisque, No. 49, 1977.
[6] Edwin E. Floyd. A nonhomogeneous minimal set. *Bull. Amer. Math. Soc.*, 55(10):957–960, 10 1949.
[7] Richard Gjerde and Ørjan Johansen. C*-algebras associated to non-homogeneous minimal systems and their K-theory. *Math. Scand.*, 85(1):87–104, 1999.
[8] Kamel N. Haddad and Aimee S. A. Johnson. Auslander systems. *Proc. Amer. Math. Soc.*, 125(7):2161–2170, 1997.
[9] Richard H. Herman, Ian F. Putnam, and Christian F. Skau. Ordered Bratteli diagrams, dimension groups and topological dynamics. *Internat. J. Math.*, 3(6):827–864, 1992.
[10] John E. Hutchinson. Fractals and self-similarity. *Indiana Univ. Math. J.*, 30(5):713–747, 1981.
[11] Ian F. Putnam. An excision theorem for the K-theory of C^*-algebras. *J. Operator Theory*, 38(1):151–171, 1997.
[12] Ian F. Putnam. *Cantor minimal systems*, volume 70 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2018.
[13] Jean Renault. *A groupoid approach to C^*-algebras*, volume 793 of *Lecture Notes in Mathematics*. Springer, Berlin, 1980.
[14] Wacław Sierpiński. Sur une courbe *cantor*ienne qui contient une image biunivoque et continue de toute courbe donnée. *C. R. Acad. Sci., Paris*, 162:629–632, 1916.

Department of Mathematics, University of Colorado Boulder Campus Box 395, Boulder, CO 80309-0395, USA
Email address: robin.deeley@gmail.com

Department of Mathematics and Statistics, University of Victoria, Victoria, B.C., Canada V8W 3R4
Email address: ifputnam@uvic.ca

Institute of Mathematics, Czech Academy of Sciences, Žitná 25, 115 67 Prague, Czech Republic
Email address: strung@math.cas.cz