ISRAEL JOURNAL OF MATHEMATICS **156** (2006), 93–110

MINIMAL MODELS FOR NONINVERTIBLE AND NOT UNIQUELY ERGODIC SYSTEMS

BY

TOMASZ DOWNAROWICZ

Institute of Mathematics, Wrocław University of Technology
ul. Wybrzeze Wyspiańskiego 27, 50-370 Wrocław, Poland
e-mail: downar@pwr.wroc.pl

ABSTRACT

Let (Y, S) be a (not necessarilly invertible) topological dynamical system on a zero-dimensional metric space Y without periodic points. Then there exists a minimal system (X, T) with the same simplex of invariant measures as (Y, S). More precisely, there exists a Borel isomorphism between full sets in Y and X such that the adjoint map on measures is a homeomorphism between the corresponding sets of invariant measures in the weak* topology. As an application we construct a minimal system carrying isomorphic copies of all nonatomic invariant measures.

Introduction

Let (X, Σ, μ) be a standard Borel probability space and let T be a measurable measure-preserving transformation from X into itself, i.e., such that $\mu(A) = \mu(T^{-1}(A))$ for every $A \in \Sigma$. Then (X, Σ, μ, T) is called a measure-theoretic dynamical system or an endomorphism. An invertible T is often called an automorphism. A measure-theoretic dynamical system is called ergodic if all T-invariant sets (i.e., $A \in \Sigma$ satisfying $T(A) \subset A$) have either measure 1 or 0. Two measure-theoretic dynamical systems (X, Σ, μ, T) and (X', Σ', μ', T') are said to be isomorphic if there exists a bimeasurable bijection $\psi\colon X_0 \to X_0'$, where $X_0 \in \Sigma$, $X_0' \in \Sigma'$, $\mu(X_0) = \mu'(X_0') = 1$, which sends the measure μ to μ' (i.e., $\mu(A) = \mu'(A')$ whenever $A' = \psi(A)$, $A \in \Sigma$), and which is equivariant, i.e., $\psi \circ T = T' \circ \psi$ μ-almost everywhere. A system isomorphic to an ergodic one is ergodic.

Received September 30, 2004

By an **assignment** we will mean a function Ψ defined on an abstract metrizable Choquet simplex $\mathcal{P}$, whose "values" are measure-theoretic dynamical systems, i.e., for $p \in \mathcal{P}$, $\Psi(p)$ has the form $(X_p, \Sigma_p, \mu_p, T_p)$. Two assignments, Ψ on a simplex $\mathcal{P}$ and Ψ' on a simplex $\mathcal{P}'$, are said to be **equivalent** if there exists an affine homeomorphism of Choquet simplexes $\pi: \mathcal{P} \to \mathcal{P}'$ such that for every $p \in \mathcal{P}$ the systems $\Psi(p)$ and $\Psi'(p')$, where $p' = \pi(p)$, are isomorphic.

By a topological dynamical system we shall mean a pair (X, T), where X is a compact metric space and T is a continuous map of X into itself. A topological dynamical system (X, T) is minimal if for every $x \in X$ the orbit $\{T^n(x): n \in \mathbb{N}\}$ is dense in X. In the context of a topological dynamical system (X, T), by a "measure" we will always mean a probability measure on the Borel sigma-field $\mathcal{B}_X$. By $\mathcal{P}_T(X)$ we will denote the collection of all T-invariant measures on X, i.e., measures μ preserved by T, in other words such that $(X, \mathcal{B}_X, \mu, T)$ becomes a measure-theoretic dynamical system. It is well known that $\mathcal{P}_T(X)$ is a nonempty compact, for the weak* topology of measures, metrizable Choquet simplex whose extreme points are precisely the ergodic invariant measures. A topological dynamical system (X, T) determines a natural assignment on the simplex $\mathcal{P}_T(X)$ by the rule: $\mu \mapsto (X, \mathcal{B}_X, \mu, T)$.

This note contributes to the investigation of the following abstract problem: *Characterize the assignments equivalent to the natural assignments arising from minimal topological dynamical systems.* The renowned Jewett–Krieger theorem solves the problem for the trivial (one-point) simplex and automorphisms; every assignment of an ergodic automorphism can be equivalently realized by a minimal (strictly ergodic) invertible zero-dimensional topological system. A. Rosenthal [R] proved an analogous theorem also for ergodic endomorphisms. We will provide insight into the case of nontrivial simplexes.

As an application of our result we will create a noninvertible version of the "universal system" of Weiss. Recall that B. Weiss [W] constructs a minimal invertible system whose assignment's range contains (up to isomorphism) all possible invertible measure-preserving transformations.

Let us mention that I. Kornfeld and N. Ormes have recently obtained results overlapping with our main theorem. We will discuss the similarities and differences in the next section.

Refinements of the problem and formulation of the main result

For nontrivial simplexes there is no known characterization of the assignments realizable in minimal systems. Likewise, there is no characterization of the assignments realizable in topological (not necessarily minimal) systems. Though, there are some obvious restrictions on such assignments. Let us begin with the most general and obvious ones:

(R1) Ψ assigns ergodic systems to extreme points;

(R2) Ψ obeys the **ergodic decomposition rule**: if $p = \int e d\xi_p(e)$, where ξ_p is the unique probability measure with barycenter at p, supported by the extreme points e of $\mathcal{P}$, and $\Psi(p) = (X_p, \Sigma_p, \mu_p, T_p)$, then μ_p admits a decomposition $\mu_p = \int \mu_e d\xi_p(e)$ with each μ_e ergodic, preserved by the transformation T_p, and such that $(X_p, \Sigma_p, \mu_e, T_p)$ is isomorphic to $\Phi(e)$.

These two restrictions apply to assignments arising from any (not necessarily minimal) topological dynamical systems and follow from the basics of ergodic theory. They allow us to focus on assignments defined only on the extreme points of simplexes, and associating ergodic measure-preserving transformations.

In minimal realizations another restriction is obvious:

(R3) The assigned measure-theoretic dynamical systems are nonatomic.

Indeed, the atomic part of an invariant measure is supported by finitely many periodic points, so that the only assignments involving atomic measures and realizable in minimal systems are those on trivial simplexes assigning a measure supported by a single periodic orbit. We exclude such trivial systems from our considerations (they are uniquely ergodic).

It is obvious that the above list of restrictions is incomplete. There must be some kind of "measurability" or even "semicontinuity" of the assignment involved, but due to lack of a natural topology or measurable structure in the "class of classes" of measure-theoretic dynamical systems modulo isomorphisms, they seem extremely difficult to capture. A manifestation of the existence of such type of restriction is seen in the following condition, valid without assuming minimality:

(R4) The entropy function $p \mapsto h(\Psi(p)) := h_{\mu_p}(T_p)$ must be a nondecreasing limit of upper-semicontinuous functions (see [D-S]).

In this note we exploit the following approach: an assignment determined by a non-minimal topological dynamical system should possess all the mysterious "measurability" or "semicontinuity" properties. Does minimality impose any further restrictions other than (R3)? In other words, if Ψ is an assignment determined by an arbitrary topological dynamical system (Y, S) having no periodic points (this is (R3) for such assignments), does there exist a minimal topological dynamical system (X, T) whose assignment is equivalent to Ψ?

We will answer this question affirmatively in the case of Y zero-dimensional:

THEOREM 1: *If Y is zero-dimensional and (Y, S) has no periodic points, then the assignment determined by (Y, S) is equivalent to an assignment determined by some minimal system (X, T).*

In particular, we will be able to easily create minimal systems with the simplex of invariant measures spanned by an arbitrarily preassigned finite collection of nonperiodic ergodic measures. Simply, we apply the above theorem to the disjoint union of the strictly ergodic realizations obtained by the Jewett–Krieger or Rosenthal theorems.

Independently, Kornfeld and Ormes ([K-O]) have recently investigated the assignments (under the name of "families of ergodic automorphisms") in almost the same spirit, and have proved a result, which can be formulated verbatim as Theorem 1* with three modifications:

- the system (Y, S) is assumed invertible,
- it is assumed to have at most countably many ergodic measures,
- the minimal system (X, T) is obtained within the topological orbit equivalence class of any *a priori* given minimal Cantor system whose simplex of invariant measures is affinely homeomorphic to that of (Y, S).

The realization within a prescribed orbit equivalence class is clearly an essential advantage of the quoted result over ours. Roughly interpreted, it says that the affine-topological form of the simplex of invariant measures is the only "recognizable in terms of ergodic theory" invariant of topological orbit equivalence. On the other hand, the invertibility and countability assumptions restrict the range of that theorem.

Let us remark that on simplexes with countably many extreme points, every nonnegative function h satisfying the decomposition integral formula is in fact a nondecreasing limit of upper-semicontinuous functions. Namely, by the monotone convergence theorem, it is the limit of the functions obtained as integral extensions of the functions $h \cdot \mathbf{1}_{\{e_1, e_2, \ldots, e_n\}}$, where (e_i) is some ordering of the extreme points. Thus, the restriction (R4) is in fact void for such simplexes. If it happens that the mysterious "measurability" or "semicontinuity" conditions are of a similar nature (i.e., some parameter functions are monotone limits of semicontinuous functions), then they are also void in the countable case, leading to the following, somewhat optimistic, yet appealing conjecture:

---

* *Added in proof*: While this paper was being refereed, Konrnfeld and Ormes improved their result making it more general (but less similar to our Theorem 1): they actually proved Conjecture 1 (see below) in the invertible case (maintaining the additional orbit equivalence class statement).

CONJECTURE 1: *Every nonatomic assignment on a simplex with countably many extreme points is realizable in a minimal system.*

Our Theorem 1 (as well as the theorem of [K-O]) automatically extends to systems (Y, S) with so-called small boundary property, i.e., existence of arbitrarily fine finite open covers by sets with boundaries having measure zero for all invariant measures. Every system with small boundary property has a zero-dimensional extension which yields exactly the same assignment. The construction is standard (see, e.g., [D] for a description). The small boundary property has been exploited in the works of E. Lindenstrauss (see [L] and reference therein). Lack of periodic orbits plus any of the properties listed below suffices for (Y, S) to have the small boundary property:

- Y is zero-dimensional (includes all subshifts over finite or countable alphabets);
- (Y, S) has finitely or countably many ergodic measures;
- S is invertible and Y finite-dimensional [K];
- (Y, S) is invertible, has finite topological entropy and a nonperiodic minimal topological factor [L].

Finally, as we shall see in the proof, zero-dimensionality is needed only for the existence of closed-and-open (clopen) markers, which immediately implies that

THEOREM 2: *Theorem 1 also holds in the case where Y is not zero-dimensional but (Y, S) admits a nonperiodic factor with the small boundary property.*

Technical lemmas, precise formulation of the result, and proofs

We will prove a statement stronger than equivalence of the assignments. We will prove that every zero-dimensional system without periodic points is conjugate to a minimal flow in a rather strong sense, which we specify in the following definition. Let (X, T) be a topological dynamical system. A Borel subset $X' \subset X$ is called a full set if it has measure 1 for every invariant measure $\mu \in \mathcal{P}_T(X)$.

Definition 1: By a **Borel* isomorphism** between two topological dynamical systems (X, T) and (Y, S) we shall understand an equivariant Borel-measurable bijection ϕ: $X' \to Y'$ between full invariant subsets $X' \subset X$ and $Y' \subset Y'$, such that the conjugated map ϕ^*: $\mathcal{P}_T(X) \to \mathcal{P}_S(Y)$ defined by the rule $\phi^*(\mu)(A) = \mu(\phi^{-1}(A))$ $(A \in \mathcal{B}_X)$ is a homeomorphism with respect to the weak* topology.

If ϕ is a Borel* isomorphism, then the pair ϕ and ϕ^* establishes an equivalence between the assignments determined by (X, T) and (Y, S); ϕ^* plays the role of an affine homeomorphism between the simplexes, while, for each pair of measures μ, $\nu = \phi^*(\mu)$, ϕ provides the isomorphism between $(X, \mathcal{B}_X, \mu, T)$ and $(Y, \mathcal{B}_Y, \nu, S)$. Notice that any equivariant Borel-measurable bijection ϕ between full invariant sets provides an affine bijection ϕ^* between the simplexes of invariant measures. By compactness of these simplexes, in order to verify ϕ^* as a homeomorphism (and thus a Borel* isomorphism) it suffices to check its weak* continuity.

We will exploit the notion of n-markers. The following lemma is a version of the so-called Krieger's marker lemma (see [B] for the invertible case) in the absence of periodic points.

Definition 2: A subset F of a topological dynamical system (Y, S) is called an n-**marker** ($n \in \mathbb{N}$) if

(1) the sets $T^{-i}(F)$ $(0 \leq i \leq n)$ are pairwise disjoint;
(2) the sets $T^{-i}(F)$ $(0 \leq i \leq m)$ cover Y for some $m \geq n$.

The system (X, T) is said to have the **marker property** if there exist clopen n-markers for all $n \in \mathbb{N}$.

LEMMA 1: *Every topological system admitting a zero-dimensional factor (Y, S) without periodic points has the marker property.*

Proof: It suffices to prove the lemma for (Y, S). Then the n-markers lift to n-markers in the larger system. For given n, every point $y \in Y$ belongs to a clopen set E_y such that $T^{-i}(E_y)$ $(0 \leq i \leq n)$ are pairwise disjoint. Choose $\mathcal{U} = \{U_j\}$, in a finite subcover by the sets E_y. The cover $\mathcal{U}' = \{U'_j\} = \{T^{-nm}(U_j)\}$, where m is the cardinality of $\mathcal{U}$, has the same property; each set has pairwise disjoint $n+1$ preimages; in addition, their nm forward images are also clopen. Define inductively

$$F_1 := U'_1,$$
$$F_{j+1} := F_j \cup \left(U'_{j+1} \setminus \bigcup_{-n \leq i \leq n} T^i(F_j) \right),$$

and set $F = F_m$. This is clearly a clopen set. The verification of (1) is straightforward, and it is also not hard to verify that $2n+1$ preimages of F_m cover Y. ∎

We are in a position to state (and prove) the main theorem of this note, which obviously covers Theorem 1 (and implies Theorem 2):

THEOREM 3: *If (Y, S) has the marker property, then it is Borel*-isomorphic to a minimal topological dynamical system (X, T). In particular, the assignment* *determined by (Y,S) is equivalent to the assignment determined by the minimal system (X,T).*

Proof: The mainframe of the proof relies on techniques developed in [D-L]. However, we needed to change most of the technical details. The proof is presented in a completely self-contained way. The construction of the Borel* isomorphism is by induction. As we want to reserve the letter n to denote the coordinates in the elements of our subshifts, we will use the letter t for the naturals enumerating the induction steps (and some other objects). It is an easy exercise to arrange a decreasing (with respect to inclusion) sequence of clopen n-markers. We will use a subsequence of such, denoted by F_t. By $p_t < q_t$ we will denote the upper and lower bounds on the gaps between visits in F_t, respectively. We require that the gap sizes grow so fast that the sequence

$$\epsilon_1 = \frac{1}{4}, \quad \epsilon_{t+1} := \frac{(r_t+4)q_t}{p_{t+1}} < \frac{1}{4}$$

is summable (the numbers r_t will be specified later). For easy reference we note the following equality:

(3) $p_{t+1}\epsilon_{t+1} = (r_t+4)q_t.$

Let $Y_0 = Y \cup \{0,1,2,\dots,\infty\}$ with such a metric d that the ϵ_t-ball around t is exactly $\{t, t+1, \dots, \infty\}$. Zeros will be considered "empty spaces", other integers are "markers" and $y \in Y$ are called "symbols".

To begin, we replace each y by the $(\mathbb{N}_0 \times \mathbb{N}_0)$-matrix $(y_{k,n})_{k\in\mathbb{N}_0, n\in\mathbb{N}_0}$ with the top row (indexed by $k=0$) filled with zeros except at positions n corresponding to the times of visits of $T^n y$ to the marker sets F_t, where we put the largest index t of a visited marker (the symbol ∞ will be used at most once in a representation of a point.) The row number 1 is filled with the orbit of y, i.e., it reads $(y_{1,0}, y_{1,1}, y_{1,2}, \dots) = (y, Ty, T^2y, \dots)$. All further rows are empty.

$$y \equiv (y_{k,n})_{k\in\mathbb{N}_0, n\in\mathbb{N}_0} = \begin{cases} t_o \quad t_1 \quad t_2 \quad \cdots & \leftarrow \text{ the top row, markers} \\ y_{1,0} \quad y_{1,1} \quad y_{1,2} \quad \cdots & \leftarrow \text{ first row, orbit} \\ 0 \quad 0 \quad 0 \quad \cdots & \leftarrow \text{ further rows, empty} \\ \vdots \quad \vdots \quad \vdots \quad \vdots & \end{cases}$$

Because the markers occur at visits to clopen sets, this procedure yields a topologically conjugate representation of (Y,S) as a subset of $Y_0^{\mathbb{N}_0\times\mathbb{N}_0}$ with the horizontal shift map S; $(S(y))_{k,n} = y_{k,n+1}$. From now on, by (Y,S) we shall mean this representation.

By a block of length l we shall mean any "vertical strip" of the form $B = (B_{k,n})_{k\in\mathbb{N}_0, 0\le n<l}$ with entries in Y_0. We define a distance D between blocks as $\mathrm{diam}(Y_0)$ for blocks of different lengths, and

$$D(B, B') = \sup\{\sum_{k=0}^{\infty} 2^{-k} d(B_{k,n}, B'_{k,n}) \colon 0 \le n < l\}$$

for blocks of the same length l.

Let n_1, n_2 be the coordinates of some two consecutive markers larger than or equal to t (i.e., $t_{n_1} \ge t$, $t_{n_2} \ge t$, and $t_n < t$ for $n_1 < n < n_2$) in the matrix representation of some $y \in Y$. The block $(y_{k,n})_{k\in\mathbb{N}_0, n_1<n\le n_2}$ will be called an **original** t-block. By the arrangement of the marker sets F_t, each original $(t+1)$-block is a concatenation of original t-blocks, all but the last one ending with the marker t, and the last one ending with some $m \ge t+1$. The terminal section of length $p_{t+1}\epsilon_{t+1}$ of a $(t+1)$-block will be called a **buffer**. By (3), the buffer contains at least $r_t + 4$ complete t-blocks.

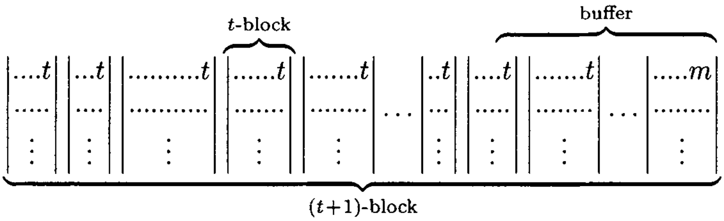

Each $y \in Y$ (in the matrix representation) decomposes, for each t, in a unique way to concatenation of original t-blocks, the first one typically incomplete (truncated on the left).

By compactness of Y_0, there is a finite family of original t-blocks ϵ_t-dense in the metric D among all original t-blocks. We define the number r_t appearing in (3) as the cardinality of this family (the numbers p_{t+1} and ϵ_{t+1} should be established after that).

We will now define a sequence of equivariant injective and continuous maps ϕ_t from Y into $Y_0^{\mathbb{N}_0\times\mathbb{N}_0}$. The map ϕ_1 is the identity. Assume that ϕ_t is defined as a length-preserving continuous code on t-blocks, replacing original t-blocks by their images called **regular** t-blocks. Assume also that the contents of the row number 1 of the original t-block B is memorized in the "bottom line" of $\phi_t(B)$, i.e.,

$$B_{1,n} = (\phi_t(B))_{k_n,n},$$

with $k_n = \sup\{k \colon (\phi_t(B))_{k,n} \ne 0\}$, and that $k_n \le t$ for each n, i.e., that all rows below t in the regular t-blocks are empty. Then we define the map ϕ_{t+1}

as a code on the original $(t+1)$-blocks B as follows: Let V be a concatenation of selected r_t original t-blocks which are ϵ_t-dense among all original t-blocks, with the terminal markers changed to t (this will not affect the density), and arranged in a certain (arbitrary) order. Let $W = CCV$, where C is a fixed original t-block of maximal length q_t. The length of W is at most $(r_t+2)q_t$ (the length of the buffer minus $2q_t$). The choice of W does not depend on B. Then proceed:

(A) First apply the code ϕ_t to all component t-blocks of B and of W; this creates $\phi_t(B)$ and $\phi_t(W)$.

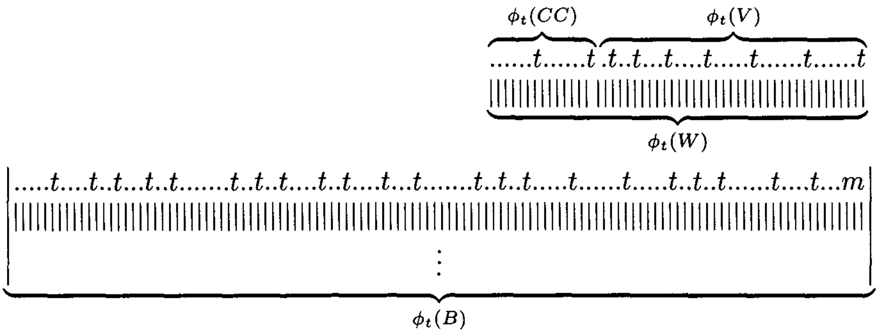

Figure 1. Step (A). Creating the block W, applying the code ϕ_t. The short vertical lines represent the t rows containing symbols (and also zeros), vertical dots mark completely empty rows (this is skipped in the presentation of $\phi_t(W)$).

(B) Determine a "cutting place" (the algorithm for that will be described later; it depends upon B). Let R denote the rectangle in rows 0 through t (including the marker row) from cutting place to the right end of $\phi_t(B)$, let Q denote the rectangle of the same size in $\phi_t(W)$, also ending at the right end.

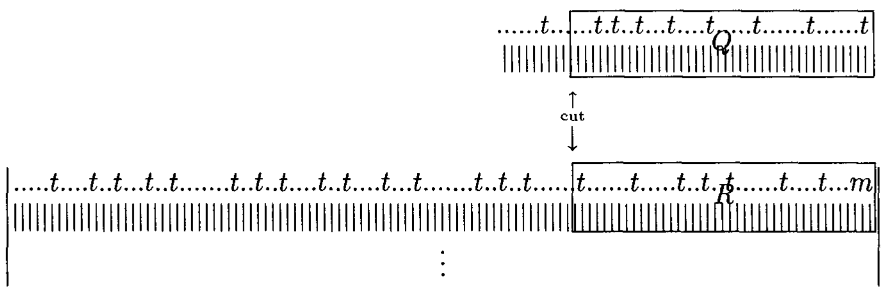

Figure 2. Step (B). Finding the cutting place and identifying the rectangles R and Q.

(C) Copy the contents of the bottom line of R into the row $t+1$. Then overwrite R by Q, except the terminal marker m which remains unchanged.

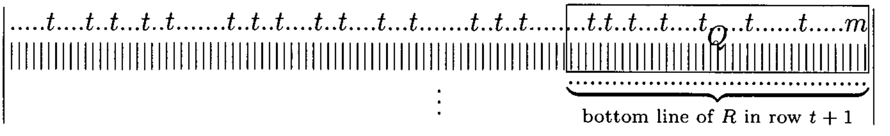

Figure 3. Step (C). Copying the bottom line of R into row $t+1$, replacing R by Q.

Notice the following properties of the code ϕ_{t+1}:

(4) $\phi_{t+1}(B)$ differs from $\phi_t(B)$ only within the buffer.

(5) The rows below $t+1$ in all regular $(t+1)$-blocks are empty, and the content of the original $(t+1)$-block is memorized in the bottom line, as required by the induction. Thus ϕ_{t+1} is a bijection.

(6) If the D-distance between two original $(t+1)$-blocks is smaller than ϵ_{t+1}, then they have the same markers (except the terminal one) and then ϕ_{t+1} does not increase the distance between them (in particular this implies continuity).

(7) Markers $t+1$ and higher have not been changed.

(8) Each regular $(t+1)$-block is a concatenation of regular t-blocks (possibly with symbols added in row $t+1$), and one exceptional block containing the cutting place. We call it an **irregular** t-block. Inductively, for each $s \leq t$, the regular $(t+1)$-block is a concatenation of regular s-blocks (possibly with symbols added below row s) and irregular s-blocks arising at all cutting places in codes $\phi_{s+1}, \ldots, \phi_{t+1}$.

We now describe the algorithm for finding the cutting place. It will be established in t steps enumerated decreasingly $t, t-1, t-2, \ldots, 1$. We will be comparing the markers in both $\phi_t(W)$ and in $\phi_t(B)$ (with their right ends aligned). Begin by placing a temporary mark $*$ under the marker t at the right end of the block $\phi_t(CC)$. In each step s find the nearest marker s in $\phi_t(B)$ left from (or at) the mark $*$, and then find the nearest marker s in $\phi_t(W)$ left from there. If the distance between these two markers s is at least $2p_s\epsilon_s$, then move the $*$ mark to the position $p_s\epsilon_s$ units left from the considered marker s in $\phi_t(B)$; otherwise, put it $p_s\epsilon_s$ units left from the considered marker s in $\phi_t(W)$ (see figure below).

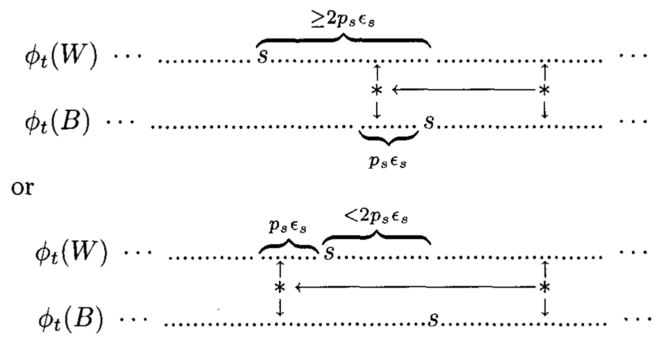

Then pass to step $s-1$. The position of the mark $*$ after step 1 is where we cut.

We now prove the following statement about the irregular blocks:

(9) An irregular s-block has length between $p_s(1-3\epsilon_s)$ and $2q_s$, and it contains a complete buffer of a regular s-block. Every such block appears within the buffer of a regular s'-block for some $s' > s$.

Assuming (9) for ϕ_t, we need to examine only the "new" irregular s-blocks created by the cuts in the code ϕ_{t+1}. Go back to the cutting algorithm for ϕ_{t+1}. In step t (the first one) of that algorithm we move the mark $*$ by no more than $q_t + 3p_t\epsilon_t$ positions. After that move, the distances from that mark $*$ to the nearest markers t in $\phi_t(W)$ on both sides and in $\phi_t(B)$ on the right are at least $p_t\epsilon_t$, and in $\phi_t(B)$ on the left it is at least $p_t(1-3\epsilon_t)$. In particular, the $*$ is now outside any buffers of the t-blocks, in an area where, by the inductive assumption, there are only regular $(t-1)$-blocks, with this area extending to the left far more than two such blocks. Thus, similarly, in the step $t-1$ we move the $*$ by not more than $q_{t-1} + 3p_{t-1}\epsilon_{t-1}$. Inductively, in each step $s \le t$ we move the $*$ by at most $q_s + 3p_s\epsilon_s$ positions and after that move the distances from it to the nearest markers s in $\phi_t(W)$ on both sides and in $\phi_t(B)$ on the right are at least $p_s\epsilon_s$, and in $\phi_t(B)$ on the left it is at least $p_s(1-3\epsilon_s)$. By the fast growth of the numbers q_s (see (3)), the moves in further steps amount to less than $p_s\epsilon_s$. Using this for $s = t$, because $q_t + 4p_t\epsilon_t < 2q_t$, we can see that the cutting place will not be moved beyond $\phi_t(W)$ or beyond the buffer in $\phi_t(B)$. For all $s \le t$ we conclude that eventually the cutting place falls at least $p_s\epsilon_s$ positions left from the nearest markers s in both $\phi_t(W)$ and $\phi_t(B)$ and at least $p_s(1-4\epsilon_s)$ right from the nearest marker s in $\phi_t(B)$. In particular, it falls outside the buffers of the s-blocks, thus the irregular $(s-1)$-block created by the code ϕ_{t+1} is concatenated from portions of two regular $(s-1)$-blocks, hence its length is at most $2q_{s-1}$. (Also, the irregular t-block created by this code is obviously created from regular t-blocks since there are no irregular t-blocks yet.) On the other hand, for each $s \leq t$ the irregular s-block inherits from $\phi_t(W)$ its right part of length at least $p_s\epsilon_s$, so it inherits the unchanged buffer, and from $\phi_t(W)$ it inherits its left part of length at least $p_s(1-4\epsilon_s)$, so the total length of the irregular s-block is at least $p_s(1-3\epsilon_s)$, as required. This concludes the proof of (9).

Another feature of the code ϕ_{t+1} is very important: it works on $(t+1)$-blocks B truncated on the left. For such blocks, while performing the algorithm for finding the cutting place, often we do not see the nearest marker s in $\phi_t(B)$ on the left of the $*$ mark. Nonetheless, in such cases, we do know that the cutting place is left from the truncation point (which suffices; in such case we simply do not cut). If this marker does appear in the visible part of the block $\phi_t(B)$, then, since we have full knowledge of the distance to the next marker s in $\phi_t(W)$ (even if that one appears left from the truncation point), we can complete shifting the $*$ (if it goes left from the truncation point, then we can stop the algorithm). Thus the map ϕ_{t+1} is applicable and shift-equivariant on all matrices representing the elements $y \in Y$.

Denote

$$Y_t = \{y \in Y\colon \text{the coordinate 0 falls in the buffer of an original } t\text{-block}\},$$

$$\begin{aligned} Y'' &= Y \setminus \bigcap_{t=1}^{\infty} \bigcup_{m=t}^{\infty} Y_t \\ &= \{y \in Y\colon \text{the coordinate 0 falls in the buffer of an} \\ &\qquad \text{original } t\text{-block for at most finitely many } t\}, \end{aligned}$$

$$\begin{aligned} Y' &= \bigcap_{n=0}^{\infty} S^{-n}(Y'') \\ &= \{y \in Y\colon \text{every coordinate } n \text{ falls in the buffer of an} \\ &\qquad \text{original } t\text{-block for at most finitely many } t\}. \end{aligned}$$

Notice that for fixed t and any $y \in Y$, the frequency of visiting the set Y_t by the orbit of y is at most ϵ_t. Thus ϵ_t estimates $\mu(Y_t)$ from above for all shift-invariant measures μ. By summability of the epsilons, Y'' is hence a full subset of Y. Therefore, Y' is a full invariant subset in (Y,S).

Now, for $y \in Y'$ the code images $\phi_t(y)$ converge coordinatewise (at each coordinate only finitely many codes intervene). Thus on Y' the limit map $\phi = \lim_t \phi_t$ is well defined (and, of course, measurable). Now let $X = \overline{\phi(Y')}$. This is a closed subset of $Y_0^{\mathbb{N}_0 \times \mathbb{N}_0}$, invariant under the horizontal shift T. For each t, every element $\phi(y)$ ($y \in Y'$) is a concatenation of regular and irregular t-blocks. This property obviously passes to all elements of X. Although in X we may not be able to determine whether the coordinate 0 is in a (truncated) regular or irregular t-block, still we may determine when it is certainly not in its buffer (when the first marker $m \geq t$ is far enough), and this happens with frequency at least $1 - \epsilon_t/(1 - 3\epsilon_t)$ (in the shortest irregular t-blocks). Thus, by the same argument as for Y' (now using summability of the numbers $\epsilon_t/(1 - 3\epsilon_t)$), the set

$$X' = \{x \in X\colon \text{every coordinate } n \text{ falls in a buffer of a regular or irregular } t\text{-block for at most finitely many } t\}$$

is a full invariant subset in (X, T).

We will now prove minimality of (X, T). For $y \in Y'$, every initial rectangle in $\phi(y)$ is part of a regular t-block (for t sufficiently large). Every regular t-block has an ϵ_t approximate in Q, which is later introduced (with symbols added in lower rows) in the buffers of all regular and irregular $(t+1)$-blocks, hence appears syndetically (with gaps bounded by $2q_{t+1}$) in each $\phi(y)$ ($y \in Y'$). By a standard argument, this already implies minimality of the closure of $\phi(Y')$.

Finally, we need to show that (Y, S) and (X, T) are Borel* conjugate. At first we will show that ϕ is injective on Y' and that the image contains X'. Injectivity is almost immediate: for $y \in Y'$, the number of nonempty rows at each coordinate of $\phi(y)$ is finite, so we can read the original y from the bottom line. Now let $x \in X'$. Because the row t is nonempty only in the buffers of the t-blocks, there are only finitely many symbols in the zero column of x. Let y be the lowest of them. Let y also denote the matrix representation of the symbol y. We need to show that $y \in Y'$ and that $\phi(y) = x$. Consider an initial rectangle in x. For sufficiently large t this rectangle is part of a truncated regular t-block which continues far to the right. This regular t-block is an image by ϕ_t of an original t-block, in which the symbol y appears at an appropriate place. This means that $\phi_t(y)$ coincides with x on that rectangle. Moreover, the nearest marker $m \geq t$ in y is at the same place as in x. Because this is true for all sufficiently large t, we obtain that $y \in Y'$ and $\phi(y) = x$.

By the proved properties of ϕ, the map $\phi^*\colon \mathcal{P}_S(Y) \to \mathcal{P}_T(X)$ is an affine bijection. We need to verify its weak*-continuity. By the elementary properties of the weak* topology, it is seen that $\phi^* = \lim \phi_t^*$. We shall prove that this convergence is uniform. Because all the maps ϕ_t^* are continuous, this will suffice for continuity of ϕ^*. In the space $Y_0^{\mathbb{N}_0 \times \mathbb{N}_0}$ the weak* distance between measures depends on values these measures assume on open ϵ-cylinders over rectangular regions containing the coordinate $(0,0)$. Thus the uniform distance between ϕ_t^* and ϕ_m^* $(m > t)$ is estimated by the sum of change of frequencies of visits of $\phi_t(y)$ and $\phi_m(y)$ in a certain family of such ϵ-cylinders over sufficiently large rectangles and for sufficiently small ϵ. But these changes are small for large t, regardless of y, because the maps introduce modifications on rare sets of coordinates, with frequencies summable over t. The sequence of maps ϕ_t^* is thus uniformly Cauchy, and hence converges uniformly. ∎

Remark 1: Notice that the map ϕ restricted to Y' preserves the number of preimages by the shift, i.e., for $y \in Y'$, $\#S^{-1}(y) = \#T^{-1}(\phi(y))$. This follows easily from the fact that the orbit of each point is stored in the bottom line and all other entries are determined by this orbit. Such preservation does not follow automatically from Borel* isomorphism. There are easy examples of Borel-isomorphic pairs of strictly ergodic (i.e., minimal, uniquely ergodic), hence Borel* isomorphic systems, where all points from a set of positive measure in one system (which may even be invertible) receive, in the other system, additional preimages from a null set.

The universal system

In order to create a minimal universal system it now suffices to build a nonminimal zero-dimensional one having no periodic points.

THEOREM 4: *There exists a minimal zero-dimensional system (X_u, T_u) such that for any measure-preserving transformation (Y, Σ, ν, S) with ν nonatomic there exists an invariant measure μ on X_u such that $(X_u, \mathcal{B}_{X_u}, \mu, T_u)$ is isomorphic to (Y, Σ, ν, S).*

Proof: It is a standard fact that every system (Y, Σ, ν, S) has an isomorphic realization as an invariant measure on $(\mathfrak{C}^{\mathbb{N}_0}, \sigma)$, where σ denotes the left shift on the countable product of Cantor sets. Such a (non-minimal) universal system has, however, many periodic points, so we cannot apply our Theorem 1 yet. We need to "blow up" the periodic orbits into nonperiodic subsystems without affecting too much other invariant measures.

First, we practice the "blowing up" technique on shifts over finite alphabets. Let (Y, S) denote the one-sided shift over, say, l symbols. We represent (Y, S) as the shift on $(\mathbb{N}_0 \times \mathbb{N}_0)$-matrices with an additional three symbols $0, a, b$, where 0 is treated as an empty cell. Initially, each y is replaced by the matrix with y placed in row number 0 and all other rows empty. Choose and fix a member

z of some nonperiodic minimal two-sided subshift over two symbols a, b (for example, the classical Morse sequence). We now modify the elements of Y as follows:

For each $n \geq 1$ we scan every not eventually periodic point y for periodic patterns of period n and length larger than n. For every such pattern we find its right end (because y is not eventually periodic, the right end of such a periodic pattern is always well determined). If $n+k$ is the total length of such a pattern, we place $z[-k+1, 0]$ in row number n precisely under the leftmost k positions of that periodic pattern. In the figure below $y = 3333312112112112121333\ldots$ is a member of the shift over three symbols. Notice that the procedure works in a shift-equivariant way even when the repetitions reach the coordinate 0 (so we cannot predict its left end).

$$\begin{array}{cccccccccccccccccccccc}
3 & 3 & 3 & 3 & 3 & 1 & 2 & 1 & 1 & 2 & 1 & 1 & 2 & 1 & 1 & 2 & 1 & 2 & 1 & 3 & 3 & 3\ldots \\
a & b & b & a & 0 & 0 & 0 & a & 0 & 0 & a & 0 & 0 & a & 0 & 0 & 0 & 0 & 0 & b & a & 0\ldots \\
b & b & a & 0 & 0 & a & 0 & 0 & a & 0 & 0 & a & 0 & 0 & b & b & a & 0 & 0 & a & 0 & 0\ldots \\
b & a & 0 & 0 & 0 & b & b & a & a & b & a & b & b & a & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0\ldots \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots\ldots
\end{array}$$

Let Y' denote the closure of all matrices obtained in this manner from not eventually periodic sequences y. Each periodic point of period n admits in row n some sequences from the one-sided factor Z of the system generated by z. In other rows the filling is uniquely determined. All not eventually periodic points continue to admit a unique filling in all rows. In fact, the modification produces an extension (Y', S') of the full shift (Y, S). The factor map (projection on the first row) is 1-1 on not eventually periodic points. Each periodic orbit lifts to a joining of that orbit with finitely many copies of Z (which contains no periodic points). The set of all eventually periodic but not periodic points of (Y, S) carries no invariant measures and so does its lift in (Y', S'), thus we can completely ignore this set.

We can now deal with the full shift over the Cantor alphabet. By representing the elements of $\mathfrak{C}$ as one-sided 0-1 sequences written as vertical columns (directed downward), we can picture $(\mathfrak{C}^{\mathbb{N}_0}, \sigma)$ as the horizontal shift on 0-1 $(\mathbb{N}_0 \times \mathbb{N}_0)$-matrices. The 0-1 full shift (Y, S) visible in the top row is a factor of $(\mathfrak{C}^{\mathbb{N}_0}, \sigma)$. The "blowing up" technique applied to(Y, S) produces (Y', S') with no periodic points. Now consider the fiber product (X, T) of (Y', S') with $(\mathfrak{C}^{\mathbb{N}_0}, \sigma)$ over (Y, S). This is a zero-dimensional system without periodic points which is an extension of $(\mathfrak{C}^{\mathbb{N}_0}, \sigma)$, injective on points which are not eventually periodic in the top row. However, every point which is eventually periodic in the top row has now obtained multiple lifts. In this manner we have affected some nonperiodic ergodic measures (by joining them with something). But this does not matter. Every such measure obviously admits another representation in the full shift over the Cantor alphabet, supported by sequences not periodic in the top row (and this copy will not be affected by our "blowing up" technique). Namely, if a measure admits a representation with some row almost surely nonperiodic, then we can simply switch rows and put this row on the top. If a nonperiodic ergodic measure admits a representation, where with probability one every row is periodic, then the measure is isomorphic to an adding machine. An adding machine has a nonperiodic one-row representation in the form of a regular 0-1 Toeplitz flow, which we can use as the first row and make all other rows constant. The application of Theorem 1 is now available for (X,T) and concludes the construction of the minimal universal system (X_u, T_u). ∎

Appendix

There are several reasons why the construction of ϕ in the proof of Theorem 3 is so complicated. We will try to convince the reader that it is not unnatural and that, in fact, we do not have much of a choice. For minimality we must introduce certain syndetically repeating rectangles in the top rows. Since we want the codes to be injective, we must "memorize" the original contents somewhere. For the sake of shift equivariance on one-sided sequences, we must do it in the same column. Otherwise, code images of two elements y, y' differing only at coordinate 0 might either differ further to the right (then σy and $\sigma y'$ would have to have, by shift equivariance, differing code images, while in fact $\sigma y = \sigma y'$) or they might not differ at all, contradicting injectivity. Further, we need to decide whether to introduce the new rectangles left or right from the markers. Suppose we choose the right hand side. Then consider two elements y, y' differing only at coordinate zero in such a way that $y_{0,0} = t \neq y'_{0,0}$. The code ϕ_t would impose modifications of y at several initial coordinates, while it would not on y'. Then again, σy and $\sigma y'$ would have to have differing code images. Thus the changes must be made only left from the markers. This is why the t-blocks have their dividing markers included at their right ends, and we change only the terminal "buffers". Further complications arise from avoiding accumulations of the cutting places. For example, if we attempted to insert rectangles of a fixed length not depending on B (for instance, to insert each time the same rectangle Q), then, because the $(t+1)$-blocks B being coded have varying structure of markers $s \leq t$, somewhere it could happen that the cutting place of step $t+1$ falls very close right from a cutting place of some previous step s. This would leave in the image by ϕ_{t+1} only few (the leftmost) entries of the rectangle Q' plugged in step s. If this happened for several t, involving the same s, syndetic occurrences of Q' could be destroyed. Thus the cutting length must be variable. In invertible systems, based on the knowledge of the length of each component t-block of each $(t+1)$-block being coded, we may assign the cutting precisely at the contact place of two such components and then compose the block Q as (ϕ_t applied to) a concatenation of original t-blocks of the form $C_1C_2\cdots C_kV$, with $C_1C_2\cdots C_k$ individually selected for the desired length. Then no "irregular" t-blocks would be created, making further arguments much simpler. Such a technique has been in fact used in a preliminary version of this paper concerning invertible maps only. But in noninvertible systems we must cope with the situation where the initial block being coded is truncated (hence the structure of components unpredictable), so that the cutting length cannot be determined, yet we must be able to determine the "visible part" of the plugged block. This rules out variable concatenations of the form $C_1C_2\cdots C_k$ and has lead us to using a fixed "extended plugged block" (in the proof that is $\phi_t(CCV)$) and only play with a variable cutting place. This unavoidably creates the irregular t-blocks, hence we need a cutting algorithm to control their length.

References

[B] M. Boyle, *Lower entropy factors of sofic systems*, Ergodic Theory and Dynamical Systems **3** (1983), 541–557.

[D] T. Downarowicz, *Entropy structure*, Journal d'Analyse Mathématique **96** (2005), 57–116.

[D-L] T. Downarowicz and Y. Lacroix, *Almost 1-1 extensions of Furstenberg–Weiss type*, Studia Mathematica **130** (1998), 149–170.

[D-S] T. Downarowicz and J. Serafin, *Possible entropy functions*, Israel Journal of Mathematics **172** (2002), 217–247.

[K-O] I. Kornfeld and N. Ormes, *Topological realizations of families of ergodic automorphisms, multitowers and orbit equivalence*, Israel Journal of Mathematics **155** (2006), 335–357.

[K] J. Kulesza, *Zero-dimensional covers of finite-dimensional dynamical systems*, Ergodic Theory and Dynamical Systems **15** (1995), 939–950.

[L] E. Lindenstrauss, *Mean dimension, small entropy factors and an imbedding theorem*, Publications Mathématiques de l'Institut des Hautes Études Scientifiques **89** (1999), 227–262.

[R] A. Rosenthal, *Strictly ergodic models for noninvertible transformations*, Israel Journal of Mathematics **64** (1988), 57–72.

[W] B. Weiss, *Countable generators in dynamics–universal minimal models*, Contemporary Mathematics, Vol. 94, American Mathematical Society, Providence, RI, 1989, pp. 321–326.