FINITE RANK BRATTELI-VERSHIK DIAGRAMS ARE EXPANSIVE

Tomasz Downarowicz and Alejandro Maass

Abstract. The representation of Cantor minimal systems by Bratteli-Vershik diagrams has been extensively used to study particular aspects of their dynamics. A main role has been played by the symbolic factors induced by the way vertices of a fixed level of the diagram are visited by the dynamics. The main result of this article states that Cantor minimal systems that can be represented by Bratteli-Vershik diagrams with a uniformly bounded number of vertices at each level (called finite rank systems) are either expansive or topologically conjugate to an odometer. More precisely, when expansive, they are topologically conjugate to one of their symbolic factors.

0. Introduction

Since the construction by Herman, Putnam and Skau in [HPS] of the representation of Cantor minimal systems by means of Bratteli-Vershik diagrams, several aspects of their dynamics have followed from particular properties or invariants constructed from the diagrams. The complete invariant for orbit equivalence proposed in [GPS] is one of the most remarkable examples. Another interesting aspect of this theory has been the characterization of particular classes of minimal Cantor systems through Bratteli-Vershik diagrams. Among other results, stationary-expansive Bratteli-Vershik diagrams characterize substitution subshifts [DHS], expansive Bratteli-Vershik diagrams with constant number of incoming edges per level characterize Toeplitz systems [GJ] and expansive Bratteli-Vershik diagrams with a finite set of incidence matrices characterize linearly recurrent subshifts [CDHM].

There exists no universal checkable criterion for a Bratteli-Vershik diagram to be expansive. In the above examples expansiveness needed to be assumed or it is deduced from the specific form of the diagram (which implies finite rank). This problem motivates our article.

What does expansiveness mean in this context ?

Key words and phrases. minimal Cantor system, Bratteli-Vershik models, expansiveness.
Supported by Nucleus Millennium Information and Randomness P04-069-F.

Given a Bratteli-Vershik diagram, the vertices of a fixed level induce a natural partition of the associated state space and an associated symbolic factor (which we call a *level factor*). Moreover, since these partitions refine (generate the topology), the system appears as the inverse limit of the level factors. A big diversity can be observed in the behavior of these factors. If all of them are finite, then the system is an odometer and has only rational eigenvalues. In Bratteli-Vershik diagrams associated to Cantor minimal systems with only irrational eigenvalues (like Sturmian subshifts) all level factors are infinite and conjugate to the entire system. If the system is linearly recurrent then there is only a finite number of symbolic factors, so all (but finitely many) level factors are mutually conjugate [D].

In most of the results described in this introduction the main technical idea is to prove that the system is in fact conjugate to one of its level factors. This is also the idea we develop in this paper for the class of minimal Cantor systems which admit a Bratteli-Vershik representation such that the number of vertices per level is uniformly bounded, also called finite rank Bratteli-Vershik diagrams. Our main result states that every such system is conjugate either to one of its level factors or to an odometer. This result generalizes former expansiveness results for substitution subshifts, linearly recurrent systems and finite rank Toeplitz systems.

In Section 1 the main background about Bratteli-Vershik diagrams and its associated dynamical systems is given. The main result of the paper is stated and proved in Section 2.

1. Bratteli-Vershik representation

Consider a Cantor minimal system (X, T), i.e., a homeomorphism T on a compact metric zero-dimensional space X with no isolated points, such that the orbit $\{T^n x : n \in \mathbb{Z}\}$ of every point $x \in X$ is dense in X. We assume familiarity of the reader with the Bratteli-Vershik diagram representation of such systems, yet we recall it briefly in order to establish the notation. Fore more details see [HPS].

The vertices of the Bratteli-Vershik diagram are organized into countably many finite subsets $V_0, V_1, \ldots$ called *levels* (V_0 is a singleton $\{v_0\}$). Every edge e connects a vertex $s = s(e) \in V_{i+1}$ for some $i \geq 0$ with some vertex $t = t(e) \in V_i$. At least one edge goes upward and at least one goes downward from each vertex in V_{i+1}. Multiple arrows connecting the same vertices are admitted (see Figure 1). We assume that the diagram is *simple*, i.e., that there is a subsequence $(i_k)_{k \geq 0}$ such that from every vertex in $V_{i_{k+1}}$ there is an upward path (going upward at each level) to every vertex in V_{i_k}. For each vertex v (except v_0) the set of all edges going upward from v is ordered linearly. This induces a lexicographical order on all upward paths from v to v_0, and a partial order on all infinite upward paths arriving to v_0. We identify X with the set of all such infinite paths. We assume that this partial order has a unique minimal element x_m (i.e., such that all its edges are minimal for the local order) and a unique maximal one x_M (whose all edges are maximal for the local order). The map T sends every element x to its successor in the partial order and it sends x_M to x_m.

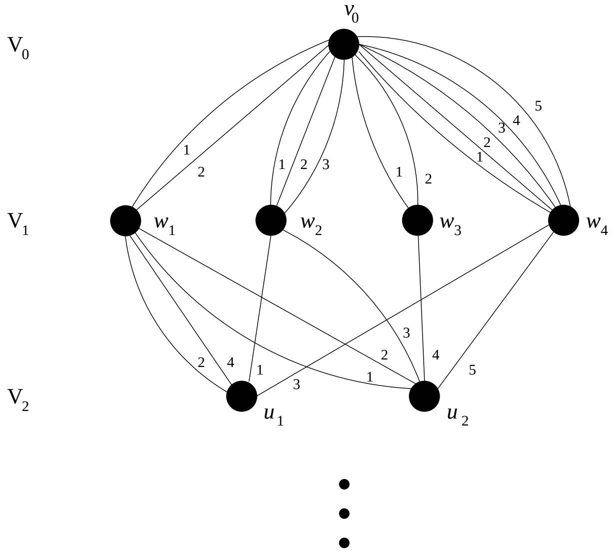

Figure 1. The first three levels of a Bratteli-Vershik diagram.

Now we provide a symbolic interpretation of the Bratteli-Vershik representation of (X, T). It is almost identical with the standard approach of considering Kakutani-Rohlin towers, but we find our symbolic approach a bit more convenient, as it allows to code individual points and play with block codes and other block manipulations.

With each vertex $v \in V_i$ we associate a rectangular matrix called an *i-symbol*, that is also denoted v. The i-symbol v has $i+1$ rows and some finite width. The last (bottom) row (number i) contains one symbol v (the name of the vertex in V_i or of the i-symbol itself) framed with a box extending over the full width of the matrix and height of one row. (Instead of a framed single symbol one can imagine a string of the form $v, v, v, \dots, v$.) The next row above it (number $i-1$) consists of several concatenated boxes labeled by the names of the vertices belonging to V_{i-1}. The number of these boxes equals the number of edges emerging upward from v and they carry labels corresponding to the target vertices of these edges ordered the same way as these edges are ordered in the diagram. In the next row (number $i-2$), aligned with each box of row $i-1$ (labeled, say, w) there is a concatenation of smaller boxes labeled by the names of the vertices in V_{i-2}, following the same rule as described for the row $i-1$ (now with reference to w in place of v). And so on, until row number 0, where all the boxes carry the same label v_0 and they all are one position wide. This determines the widths of all boxes in all rows of the described i-symbol (see Figure 2). The i-symbol associated to $v \in V_i$ is no more than its symbolic coding determined by the order of the paths in the top $i+1$ levels of the diagram.

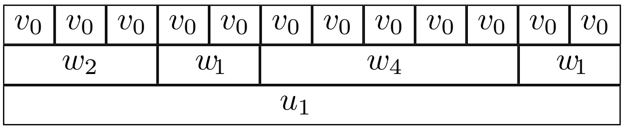

Figure 2. The 2-symbol corresponding to the vertex u_1 of Figure 1.

An element $x \in X$ can be represented as an infinite matrix with rows indexed from 0 to ∞ (each row extending from $-\infty$ to ∞), with the property that for every $i \geq 0$ the top $i+1$ rows of the matrix are an infinite concatenation of the i-symbols. The i-symbol crossing the coordinate 0 corresponds to the vertex v_i passed by the path x at the level i, and the position of the 0 coordinate in this i-symbol is determined by the position of the path connecting the vertex v_i with v_0 used by x in the lexicographical order defined on such paths. The map T corresponds to the usual horizontal shift on the set of such matrices. The finite alphabet system obtained by projecting onto the top $i+1$ rows will be denoted by X_i and we let π_i be the corresponding factoring map (projection). This is the mentioned before ith level factor of (X,T). Clearly the row i carries all the information about the system X_i, because it determines all rows above it. (In terms of the Bratteli-Vershik diagram it must be remembered, however, that the dynamics of X_i is determined by the higher-index levels of the diagram.) Also notice that the procedure known as *telescoping* the Bratteli-Vershik diagram corresponds to simply deleting some collection of rows (still leaving infinitely many of them), which leads to a topologically conjugate representation of the same system (X,T).

Consider a pair $\langle x,x'\rangle$ of distinct points in X such that $\pi_i(x)=\pi_i(x')$ for some $i \geq 1$. We call such pair *i-compatible.* Because $x \neq x'$, there exists some $j > i$ such that $\pi_j(x) \neq \pi_j(x')$. We then say that the pair is *j-separated.* The largest index i_0 for which the pair $\langle x,x'\rangle$ is i_0-compatible (and then it is (i_0+1)-separated) will be called *depth of compatibility* for this pair (or just *depth* for short). Equal elements have depth ∞. Observe that if $\langle x,x'\rangle$ is a pair of depth i and $\langle x,x''\rangle$ is a pair of depth $j > i$ then $\langle x',x''\rangle$ is a pair of depth i (hence never equal).

A j-separated pair $\langle x,x'\rangle$ is said *to have a common j-cut* if there is a coordinate $n \in \mathbb{Z}$ such that a j-symbol starts at n in both x and x' (see Figure 3). Notice that if a pair has a common j-cut then it also has a common j'-cut for each integer $j' \leq j$ (at the same coordinate n).

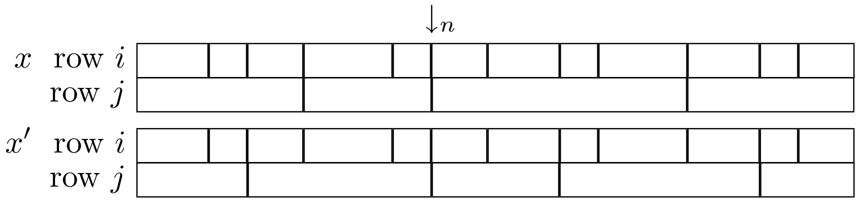

Figure 3. A common j-cut in a pair $\langle x,x'\rangle$ of depth i.

2. The main result

A Cantor minimal system (X,T) is said to have *topological finite rank* K if it admits a Bratteli-Vershik diagram representation with $\#V_i \le K$ for every $i \ge 0$ (in other words, there is a symbolic representation with no more than K i-symbols at each level i), and K is the smallest such integer. It is well known that a system has topological rank 1 if and only if it is an *odometer*, i.e., an inverse limit of a sequence of periodic systems. Such systems are also characterized as the only equicontinuous minimal Cantor systems.

Recall that (X,T) is expansive if there is an $\epsilon > 0$ such that the trajectory of any two distinct points is separated by ϵ at some moment. In the context of systems on the Cantor set, due to a classical result by Hedlund (1960), expansiveness is equivalent to topological conjugacy with a two-sided subshift. In particular, such system is topologically conjugate with one of its level factors X_i, $i \ge 0$. Once a system is conjugate to its level factor X_i, all further level factors X_j $(j \ge i)$ are all mutually conjugate.

The main result of this paper is the following theorem concerning systems with finite rank higher than 1.

Theorem 1. *Every Cantor minimal system of finite rank $K > 1$ is expansive.*

Proof. Suppose a Bratteli-Vershik diagram represents a minimal dynamical system (X,T) which is not expansive. Then for infinitely many levels i there exist pairs of points $\langle x_i, x_i' \rangle$ with depth of compatibility i. By telescoping the diagram we may assume that every $i \ge 1$ appears in this role for some pair. There are now two possibilities:

(1) There exists i_0 such that for all $i \ge i_0$ and every $j > i$ there exists a pair of depth i with a common j-cut.
(2) For infinitely many i, and then every sufficiently large $j > i$, any pair of depth i has no common j-cuts.

The proof is different for the above two cases. Recall that we are assuming that there is a $K \in \mathbb{N}$ such that $\#V_i \le K$ for each i.

Proof in case (1). In fact, we will prove that such case never occurs, even for $K = 1$.

Fix some $j > i_0 + K$ and for an integer $i \in [j-K, j-1]$ let $\langle x_i, x_i' \rangle$ be a pair of depth i with a common j-cut. For $i = j-1$ we have a $(j$–$1)$-compatible j-separated pair $\langle x_{j-1}, x_{j-1}' \rangle$ with a common j-cut. We assume without lost of generality that x_{j-1} and x_{j-1}' are j-separated to the right of this common j-cut.

If the first j-symbols right from this cut are the same in x_{j-1} and x_{j-1}' then the following j-cut is also common. So, there exists a common cut followed, in x_{j-1} and x_{j-1}', by different j-symbols, say u and v, respectively. If their lengths agree, since $\langle x_{j-1}, x_{j-1}' \rangle$ is a $(j-1)$-compatible pair, then u and v are $(j-1)$-compatible (their first $j-1$ rows coincide). We note this fact as it will play an important role in the argument. Moreover, then the following j-cut is also common, so we can continue checking the following pair of symbols to the right.

Otherwise suppose that u is longer than v. Then we modify the collection of j-symbols; namely we replace the last row in u by two boxes, the first is the same as the last row of v (and carries the symbol v), and the second has complementary length and carries a new symbol u' (see Figure 4).

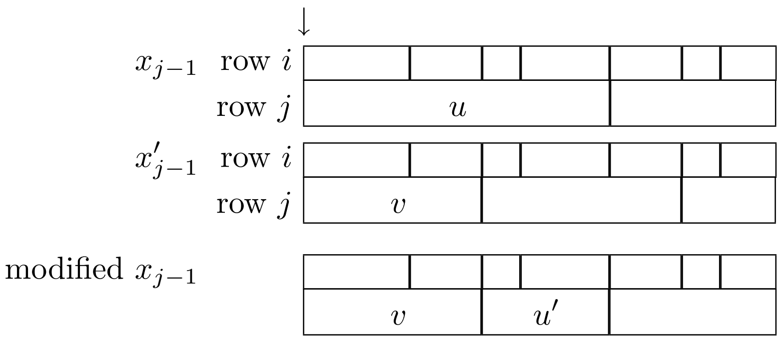

Figure 4. Change of a j-symbol. The arrow shows the common cut.

We replace every occurrence of the j-symbol u in every element of X_j by the concatenation vu'.

This procedure is reversible and preserves the number of j-symbols (it is a topological conjugacy of X_j) and it introduces more j-cuts (never removes them). If after this modification, all j-cuts in x_{j-1} and x'_{j-1} are common, this implies that u' is $(j-1)$-compatible with some j-symbol existing earlier. Again, we note this fact as important. If there are still some non-common cuts, we can repeat the previous step and replace some other j-symbol by a concatenation of one existing and one new j-symbol.

This procedure and the analogous one applied to the left of the starting j-cut must stop after finitely many such steps. So eventually we will obtain a new (but of the same cardinality) set of j-symbols (or a new equivalent Bratteli-Vershik diagram) for which the representations of x_{j-1} and x'_{j-1} we will have all j-cuts common, and there will necessarily be at least one pair of $(j-1)$-compatible but different (new) j-symbols (otherwise the pair would not be j-separated).

At this point we produce a topological factor X'_j of X_j by treating each pair of $(j-1)$-compatible j-symbols as one j-symbol. That is, in the language of Bratteli-Vershik diagrams, we amalgamate vertices of level j that are $(j-1)$-compatible (this replaces X_j by its topological factor, possibly even conjugate to X_{j-1}). It is important that in any case we STRICTLY REDUCE the number of j-symbols, so that in X'_j there is at most $K-1$ of them. Clearly, we have lost the j-separation of the pair $\langle x_{j-1}, x'_{j-1}\rangle$, but all the other pairs $\langle x_i, x'_i\rangle$ for $i \in [j-K, j-2]$ (more precisely, their images in the factor) remain j-separated, because they were separated already in rows with indices smaller than j, and these rows pass unchanged to the factor.

Moreover, our modification of the jth row only added more j-cuts, so the above pairs remain to have a common j-cut.

We can now delete (or skip) the row $j-1$ and repeat the same algorithm using the common j-cuts in the pair $\langle x_{j-2}, x'_{j-2}\rangle$. This will lead to a new factor and again reduce the number of j-symbols by at least one and preserve the j-separation with a common cut of the pairs $\langle x_i, x'_i\rangle$ for $i \in [j-K, j-3]$. Having applied the above algorithm no more than $K-1$ times we obtain a factor in which the pair $\langle x_{j-K}, x'_{j-K}\rangle$ remains to be j-separated with a common j-cut while the alphabet of j-symbols has only one element. This is clearly impossible.

Proof in case (2). After telescoping, the required behavior described in (2) is observed at all levels: For each $i \geq 1$ every pair of depth i has no common $(i{+}1)$-cuts (hence no common j-cuts for any $j > i$). Recall the earlier assumption that for each $i \geq 1$ there exists at least one pair of depth i. In this case we will prove that for each $i_0 \geq 1$ X_{i_0} is periodic, which yields that (X, T) is an odometer and hence has topological rank 1.

Fix some $i_0 \geq 1$ and let $j = i_0 + K^2$. For $i \in [i_0, j-1]$ fix a pair $\langle x_i, x'_i\rangle$ of depth i. Because the $(i{+}1)$-separation without common cuts of a pair (and hence the depth for such pair) passes via any element τ of the enveloping semigroup (pointwise limit of a subsequence T^{n_k}; see [A]), we can replace the pairs $\langle x_i, x'_i\rangle$ by $\langle \tau_i(x_i), \tau_i(x'_i)\rangle$ with τ_i so chosen, that $\tau_i(x_i) = y_0$ is the same for all considered i (we use minimality of X). Then, for any $i', i \in [i_0, j-1]$, $i' > i$, the pair $\langle y_i, y_{i'}\rangle = \langle \tau_i(x'_i), \tau_{i'}(x'_{i'})\rangle$ has depth i and, by our assumption (2), has no common j-cuts. In this way we have found K^2+1 i_0-compatible and pairwise j-separated elements $y_0, y_{i_0}, y_{i_1} \dots y_{j-1}$ with no common j-cuts.

The reminder of the proof we call an *infection lemma*; it shows that a "periodicity law" spreads along the row number i_0 from one j-symbol to another, like an infection.

Lemma. (Infection lemma). *If there exist at least K^2+1 i-compatible points y_k, ($k \in [1, K^2+1]$) which, for some $j > i$, are pairwise j-separated with no common j-cuts, then X_i is periodic.*

Proof. Denote by $\hat{y}$ the common image of the points y_k in X_i. Analogously, by $\hat{v}$ we will denote the restriction of a j-symbol v to its top $i+1$ rows. Draw a diagram showing $\hat{y}$ and all elements y_k (it suffices to draw the jth row of each of them) one above another with aligned zero coordinate. Any fixed coordinate is covered by j-symbols differently in the points y_k. However, it may (and often will) happen that in some two of them we see overlapping copies of the same j-symbol v. Because there are no common j-cuts, such copies are shifted by some positive integer $l < |v|$. The j-symbol $v = v[1, |v|]$ then satisfies the "l-periodicity law" (concerning the projection to X_i): $\hat{v}(n) = \hat{v}(n+l)$ for every $n \in [1, |v|-l]$.

We make a general observation concerning such case: let v be a j-symbol and let l_v denote the smallest shift at which two copies of v overlap anywhere in the diagram. Now observe that if some position, say 0, is covered by two copies of v (no matter how shifted), then the l_v-periodicity law holds at this coordinate in $\hat{y}$ (see Figure 5).

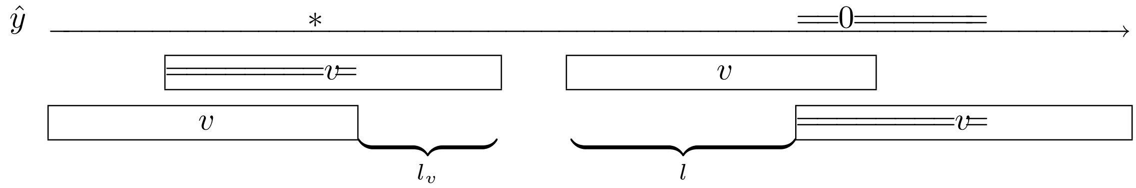

Figure 5. The periodicity law of the smallest shift l_v (marked by the long equality sign) is valid at any coordinate covered by two copies of v. (For the larger shift l the l-periodicity law does not have this property; it need not hold at the position marked by *)

Because we consider K^2+1 elements and there are only K j-symbols, the coordinate 0 is covered by at least $K+1$ (≥ 2) overlapping copies of the same j-symbol, say v. Let l_v be the minimal shift for v appearing in the diagram. Then the l_v-periodicity law holds at 0. Let I be the largest interval containing 0 where the l_v-periodicity law holds. If I is not bounded on the right then $\hat{y}$ is eventually l_v-periodic and, by minimality, X_i is a periodic system.

So suppose I has a right end. Let m be the first coordinate right from I. Restrict the diagram to some $K+1$ elements y_k in which the coordinate 0 is covered by v. In such diagram the coordinate m is still covered by at least two copies of the same j-symbol, say w. If $w=v$ then the l_v-periodicity law holds at m, which is ruled out. The coordinate m is aligned with some position $r>1$ of the extreme left copy of w. Because the coordinate zero is covered by v in all considered elements y_k, and $w \neq v$, this copy of w cannot extend beyond I on the left, hence has its left part $w[1,r-1]$ inside the interval I, where the l_v-periodicity law holds. Now take the other copy of w. It is shifted to the right by a positive shift, so the coordinate m is aligned with a position $n \in [1,r-1]$ in this copy. This implies that the l_v-periodicity law holds at m and thus $m \in I$, a contradiction (see Figure 6).

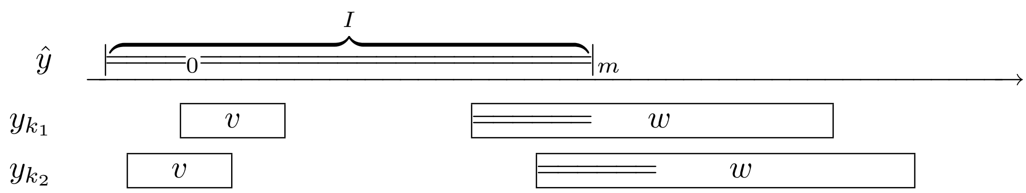

Figure 6. The second copy of w "infects" the coordinate m with the l_v-periodicity law.

We have proved that I has no right end. This concludes the proof of the infection lemma, and hence also of the main theorem. □ □

The Infection Lemma seems to have its own interest and maybe can be applied to study other combinatorial problems concerning Bratteli-Vershik diagrams. So it is interesting to ask,

Question 1. Is it possible to replace the expression $K^2 + 1$ in the assuptions of the lemma by a smaller one?

References

[A] *J. Auslander*, Minimal flows and their extensions, vol. 153, North-Holland Mathematics Studies, North-Holland, Amsterdam, 1988.

[CDHM] *M. Cortez, F. Durand, B. Host, A. Maass*, Continuous and measurable eigenfunctions of linearly recurrent dynamical Cantor systems, J. London Math. Soc. **67** (2003), 790–804.

[D] *F. Durand*, Linearly recurrent subshifts have a finite number of non-periodic subshift factors, Ergodic Th. & Dynam. Sys. **20** (2000), 1061–1078..

[DHS] *F. Durand, B. Host, C. Skau*, Substitutional dynamical systems, Bratteli diagrams and dimension groups, Ergodic Th. & Dynam. Sys. **19** (1999), 953–993.

[GPS] *T. Giordano, I. Putnam, C. Skau*, Topological orbit equivalence and C^*-crossed products, J. Reine Angew. Math. **469** (1995), 51–111.

[GJ] *R. Gjerde, O. Johansen*, Bratteli-Vershik models for Cantor minimal systems: applications to Toeplitz flows , Ergodic Th. & Dynam. Sys. **20** (2000), 1687–1710.

[HPS] *R. Herman, I. Putnam, C. Skau*, Ordered Bratteli diagrams, dimension groups and topological dynamics, Internat. J. Math. **3** (1992), 827–864.

Department of Mathematical Engineering and Center of Mathematical Modeling, University of Chile, Av. Blanco Encalada 2120, 5to piso, Santiago, Chile.

E-mail address: downar@im.pwr.wroc.pl, amaass@dim.uchile.cl