ANNALES

DE

L'INSTITUT FOURIER

Lluís ALSEDÀ, David JUHER & Pere MUMBRÚ

On the preservation of combinatorial types for maps on trees

Tome 55, nº 7 (2005), p. 2375-2398.

<http://aif.cedram.org/item?id=AIF_2005__55_7_2375_0>

Ann. Inst. Fourier, Grenoble
55, 7 (2005), 2375–2398

ON THE PRESERVATION OF COMBINATORIAL TYPES FOR MAPS ON TREES

by Lluís ALSEDÀ, David JUHER & Pere MUMBRÚ

1. Introduction and statement of the main result.

We deal with a classical problem in combinatorial dynamics: Given a topological space X, a continuous map $f : X \to X$ and a periodic orbit $A \subset X$ of f, what can be said about the dynamics (periodic orbits, topological entropy—see [1]—, etc.) of f in terms of $f_{|A}$?

A well known case is when X is a closed interval $I \subset \mathbb{R}$. Indeed, if $f : I \to I$ is a continuous map then intrinsic information can be obtained by considering the *"pattern"* of A which is characterized by the permutation π_A induced by $f_{|A}$ (see [7] or [9]). To each pattern π we may associate a map f_π which has a finite invariant set B such that the permutation induced by $f_{\pi|B}$ is π and f_π is monotone between consecutive points of B (a "connect-the-dots" map). Such a map is called a π-*monotone model* and its existence has some important consequences. During the 1980's and the early 1990's, monotone models for interval maps (and a trivial generalization to star maps) were used by several authors to tackle a wide variety of problems (see for instance [4], [5], [7], [8], [9] and [10]).

We are interested in continuous maps defined on trees (from now on, such a map will be called a *tree map*). In [2], the authors introduce a notion of *pattern* of a finite invariant set of a tree map, and prove that for any

(*) The authors have been partially supported by DGES grant number BFM2002-01344.
Keywords: Tree maps, minimal dynamics.
Math. classification: 37E25.

pattern π there exist π-monotone models which minimize the topological entropy. In this paper, which is a continuation and a natural extension of [2], we want to investigate whether these π-monotone models also minimize the dynamics (measured by the set of periodic orbits). Let us start by introducing some notation.

Given any subset X of a topological space, we will denote by $\mathrm{Int}(X)$ and $\mathrm{Cl}(X)$ the interior and the closure of X, respectively. The boundary of X will be denoted by $\mathrm{Bd}(X)$. For a finite set A we will denote its cardinality by $|A|$.

By an *interval* we mean any space homeomorphic to $[0,1] \subset \mathbb{R}$. A *tree* is a uniquely arcwise connected space that is either a point or a union of finitely many intervals. If T is a tree and $x \in T$, we define the *valence* of x to be the number of connected components of $T \setminus \{x\}$. Each point of valence 1 will be called an *endpoint* of T and the set of such points will be denoted by $\mathrm{En}(T)$. A point of valence different from 2 will be called a *vertex* of T, and the set of vertices of T will be denoted by $V(T)$. The closure of each connected component of $T \setminus V(T)$ will be called an *edge* of T. Any tree which is a union of $r > 1$ intervals whose intersection is a unique point x of valence r will be called an *r-star*, and x will be called its *central point*.

Given a tree T and $P \subset T$ we will define the *convex hull* of P, denoted by $\langle P \rangle_T$ or simply by $\langle P \rangle$, as the smallest closed connected subset of T containing P. When $P = \{x, y\}$, we will write $\langle x, y \rangle$ or $[x, y]$ to denote $\langle P \rangle$. The notations (x, y), $(x, y]$ and $[x, y)$ will be understood in the natural way.

Next we recall the definitions of *pattern* and *monotone model*, which are the central notions used in this paper. For further discussion on these ideas, we refer the reader to [2]. Let T be a tree and let $A \subset T$ be a finite subset of T. The pair (T, A) will be called a *pointed tree*. A set $Q \subset A$ is said to be a *discrete component* of (T, A) if either $|Q| > 1$ and there is a connected component C of $T \setminus A$ such that $Q = \mathrm{Cl}(C) \cap A$, or $|Q| = 1$ and $Q = A$. We say that two pointed trees (T, A) and (T', A') are *equivalent* if there exists a bijection $\phi\colon A \to A'$ which preserves discrete components. The equivalence class of a pointed tree (T, A) by this relation will be denoted by $[T, A]$.

Let (T, A) and (T', A') be equivalent pointed trees, and let $\theta\colon A \to A$ and $\theta'\colon A' \to A'$ be maps. We will say that θ and θ' are *equivalent* if $\theta' = \varphi \circ \theta \circ \varphi^{-1}$ for a bijection $\varphi\colon A \to A'$ which preserves discrete components. The equivalence class of θ by this relation will be denoted by $[\theta]$. If $[T, A]$ is an equivalence class of pointed trees and $[\theta]$ is an equivalence class of maps then the pair $([T, A], [\theta])$ will be called a *pattern.* If in addition θ is a cycle over A, then the pattern is said to be *periodic.*

A triplet (T, A, f) will be called a *model* if $f : T \to T$ is a tree map and $A \subset T$ is a finite f-invariant set. We say that a model (T, A, f) *exhibits* a pattern $(\mathcal{T}, \Theta)$, or that (T, A, f) *is a model of* the pattern $(\mathcal{T}, \Theta)$, if $\mathcal{T} = [T, A]$ and $\Theta = [f_{|A}]$. This pattern will be denoted by $[T, A, f]$.

Now we extend the notion of a "connect-the-dots" interval map to the setting of trees. Let $f : T \to T$ be a tree map. Given $x, y \in T$ we say that $f_{|[x,y]}$ is *monotone* if either $f([x, y])$ is a point or it is an interval and, given two homeomorphisms $\phi : [0, 1] \to [x, y]$ and $\varphi : f([x, y]) \to [0, 1]$, then $\varphi \circ f \circ \phi : [0, 1] \to [0, 1]$ is monotone (as a real function). If $P \subset T$ is a finite f-invariant set such that $\mathrm{En}(T) \subset P$, we say that f is P*-monotone* if $f_{|[x,y]}$ is monotone whenever $[x, y] \cap P = \{x, y\}$. In this case we will say that the model (T, P, f) is *monotone.*

Remark 1.1. — When (T, P, f) is a monotone model, it is shown in Proposition 4.2 of [2] that the image of each vertex z is uniquely determined and is either a vertex or belongs to P. In fact, if we take three different points $a, b, c \in P$ in such a way that $z \in [a, b] \cap [a, c] \cap [b, c]$ and $\langle\{a, b, c\}\rangle_T \setminus P$ is connected, then it can be easily seen that $f(z)$ is the only point contained in $f([a, b]) \cap f([a, c]) \cap f([b, c])$.

The next theorem, which is the main result of [2], states the existence and minimality of the topological entropy for monotone models of tree maps.

THEOREM 1.2 (Theorem A of [2]). — *Let $(\mathcal{T}, \Theta)$ be a pattern. Then the following statements hold.*

(a) *There exists a monotone model exhibiting the pattern $(\mathcal{T}, \Theta)$.*

(b) *The topological entropy of any monotone model is the minimum within the class of models which exhibit the pattern $(\mathcal{T}, \Theta)$.*

As we have said, the goal of this paper is to show that the dynamics of monotone models is also minimal from the point of view of the set of periodic orbits. To be more precise we recall some more notions and results from [2].

Let $f : T \to T$ be a tree map, and let $x, y \in T$ be fixed points of f^n for some $n \in \mathbb{N}$. We say that x and y are f*-monotone equivalent* if either $x = y$ or $f^n{}_{|\langle x,y\rangle}$ is monotone. It is not difficult to see that the periods of x and y

are either equal, or one is a multiple of the other. Observe that, since x and y are fixed points of f^n, the f-monotone relation is an equivalence relation.

Remark 1.3. — It is easy to see that if x and y are f-monotone equivalent then $f^i(x)$ and $f^i(y)$ are also f-monotone equivalent, for each $i \geq 0$.

Given a model (T, A, f), we say that a periodic point of f is *significant* if it is not f-monotone equivalent to any element of $A \cup V(T)$ and its period is minimal within its f-monotone equivalence class. The significant periodic points of a monotone model of a pattern $\mathcal{P}$ can be essentially put in a one-to-one correspondence with the loops of a certain combinatorial graph, which is uniquely associated to $\mathcal{P}$. Let us introduce these notions in detail.

Let $\mathcal{P} = [T, A, f]$ be a pattern. Any (unordered) binary subset of a discrete component will be called a *basic path* of (T, A). The *$\mathcal{P}$-path graph* is the oriented graph whose vertices are in one-to-one correspondence with the basic paths π_i of (T, A) and there is an arrow from the vertex i to the vertex j if the corresponding basic paths satisfy $\pi_j \subset \langle f(\pi_i)\rangle$. Note that this definition is independent of the particular choice of the model (T, A, f). Let $\pi_0 \to \pi_1 \to \cdots \to \pi_{n-1} \to \pi_0$ be a loop α of length n in the $\mathcal{P}$-path graph. The length of such a loop α will be denoted by $|\alpha|$. The loop α and a point $x \in T$ are said to be *associated* if $f^i(x) \in \langle \pi_{i \bmod n}\rangle_T$ for each $i \geq 0$. We note that if in addition x is a periodic point, then the period of x is a divisor of $|\alpha|$. Recall that a loop is called *simple* if it is not an n-repetition (with $n \geq 2$) of any other loop. Now we are ready to state the two results of [2] which describe the sense in which the monotone models have minimal dynamics.

Theorem 1.4 (Theorem C of [2]). — *Let (T,A,f) be a monotone model exhibiting a pattern $(\mathcal{T},\Theta)$. Then the following statements hold.*

(a) *For each significant periodic point x of f of period n there exists a unique simple loop β of length n in the $(\mathcal{T},\Theta)$-path graph such that x and β are associated.*

(b) *Each simple loop β of length n in the $(\mathcal{T},\Theta)$-path graph is associated either to a significant periodic point of f of period n or to a periodic point which is f-monotone equivalent to a point of $A \cup V(T)$ and whose period is a divisor of n. In both cases, the point associated to β is unique up to f-monotone equivalence.*

THEOREM 1.5 (Theorem D of [2]). — *Let (T,A,f) be a model exhibiting a pattern $(\mathcal{T},\Theta)$ and let β be a simple loop of length n of the $(\mathcal{T},\Theta)$-path graph. Then there exists a fixed point x of f^{2n} such that β and x are associated.*

From Theorems 1.4 and 1.5 it follows that the set of periods of a monotone model is *essentially* (up to f-monotone equivalence and period-doubling) contained in the set of periods of each model of the same pattern.

Theorem A below is a stronger version of the above results for periodic patterns. It says that, when $(\mathcal{T},\Theta)$ is periodic, there is essentially a period-preserving injective map from the set of (almost all) significant periodic points of f into the set of periodic points of each model exhibiting $(\mathcal{T},\Theta)$. Moreover, the corresponding orbits are associated to the same loop in the path graph. Therefore, the relative position of the points inside the respective trees is essentially preserved.

Before stating in detail this result, we remark that, for simplicity, we will deal with a particular kind of monotone models, which will be called *A-linear models*. An A-linear model is essentially a special kind of "simplified" monotone model which does not exhibit unnecessary duplicated orbits. An A-linear model (T,A,f) behaves (with respect to some natural metric) like an affine map of the real line (with respect to the Lebesgue metric) on each connected component of $T \setminus (A \cup V(T))$. This condition is not restrictive: in Section 2 we prove that any pattern admits an A-linear model. We also prove that the monotone equivalence relation defined before Remark 1.3 is particulary simple in A-linear models.

Let (T,A,f) be a monotone model and let x be a significant n-periodic point of f. In particular, $x \notin V(T)$. The point x will be called *positive* if f^n preserves orientation in a neighborhood of x, and *negative* if f^n reverses orientation.

Finally, for any pointed tree (S,P) we define $M(S,P)$ as the number of basic paths π such that $\operatorname{Int}(\langle\pi\rangle_S) \cap V(S) \neq \emptyset$ multiplied by the maximum number of vertices of S contained in the interior of the convex hull of a basic path. It is easy to see that the number $M(S,P)$ can be upper bounded by $\binom{|\operatorname{En}(S)|}{2} \cdot |V(S)|^2$.

This paper is devoted to prove:

THEOREM A. — *Let $g\colon S \to S$ be a tree map and P be a periodic orbit of g. Let (T,A,f) be an A-linear model of the pattern $[S,P,g]$ and Λ_f be* *the set of significant periodic points x of f such that either x is positive or $|\operatorname{Orb}_f(x)| > M(S,P)$. Then:*

(a) *There exists a period-preserving injective map μ from Λ_f into the set of periodic points of g such that, for each $x \in \Lambda_f$, x and $\mu(x)$ are associated to the same loop in the $[S,P,g]$-path graph.*

(b) *The complement of Λ_f in the set of periodic points of f can be written as a disjoint union $\Omega \cup \Gamma \cup C$, where:*

- (b.1) *Ω is finite and coincides with the set of negative significant periodic points of f with period not greater than $M(S,P)$,*
- (b.2) *Γ is a union of periodic orbits contained in $V(T) \setminus A$. Thus, the period of each point in Γ is not greater than $|V(T) \setminus A|$,*
- (b.3) *$C = C_0 \cup C_1 \cup \dots C_{k-1}$ for a divisor k of $|A|$, where C_i intersects A, $f(\langle C_i \rangle_T) = \langle C_{i+1 \bmod k} \rangle_T$ and f maps bijectively C_i onto $C_{i+1 \bmod k}$ for $0 \le i \le k-1$. Moreover, $\langle C_i \rangle_T \cap \langle C_j \rangle_T = \emptyset$ whenever $i \neq j$ and $(\Omega \cup \Gamma) \cap \bigcup_i \langle C_i \rangle_T = \emptyset$. The period of each point in C is not greater than $8|P| - 4$.*

Consider a periodic pattern $\mathcal{P}$ and take an A-linear model (T, A, f) of $\mathcal{P}$. As it has been said, the significant periodic orbits of f can be easily obtained by studying the loops of the $\mathcal{P}$-path graph (see Theorem 1.4 and Proposition 3.4). Therefore, Theorem A (a) can be used to obtain relevant information about the periodic orbits of any particular model (S, P, g) which exhibits the pattern $\mathcal{P}$. For instance, from Theorem A it follows that, for any $n > M(S, P)$, the number of n-periodic points of g is lower bounded by the number of significant n-periodic points of the A-linear model. Moreover, the fact that the μ-corresponding orbits are associated to the same loop allows us localizing the orbits of g in S with a certain precision.

On the other hand, (b) of Theorem A tells us that the complement of Λ_f (that is, the set of periodic points of f which are not controlled by (a) of Theorem A) consists of a finite set $\Omega \cup \Gamma$ together with a set C contained in a cycle of k connected components, each of them intersecting A. The periods of the points in Ω, Γ and C are respectively bounded by $M(S, P)$, $|V(T) \setminus A|$ and $8|P| - 4$. It is easy to check that $|V(T) \setminus A| < M(S, P)$, so each point in $\Omega \cup \Gamma$ has period not greater than $M(S, P)$. The cardinality of Γ is at most $|V(T) \setminus A|$ while, on the other hand, it is not difficult to calculate an explicit upper bound for the cardinality of Ω in terms of $|P|$ (see Remark 3.5). For simplicity, this bound has not been made explicit in the statement of Theorem A.

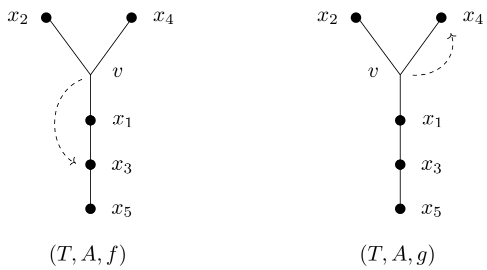

Figure 1. *On the left, an A-linear model of a pattern $\mathcal{P}$, with $A = \{x_1, \ldots, x_5\}$, $f(x_i) = x_{i+1}$ for $1 \le i < 5$ and $f(x_5) = x_1$. On the right, a model (T, A, g) of $\mathcal{P}$ such that $g(v) = x_4$ and g is monotone between consecutive points of $A \cup V(T)$.*

Example 1.6. — Consider the pattern $\mathcal{P}$ of the corresponding A-linear model (T, A, f) shown on the left side of Figure 1. Take $(S, P) = (T, A)$. Since there are only three convex hulls of basic paths whose interior contains one vertex of S and this vertex is unique, it follows that $M(S, P) = 3$. By checking the loops of the $\mathcal{P}$-path graph one easily gets that $\Gamma = \emptyset$, $C = A$ (see the proof of Theorem A (b) for explicit definitions of Γ, C and Ω) and the set of periods of the points in Λ_f is $\mathbb{N} \setminus \{1, 2, 3\}$. Moreover, Ω consists of three negative periodic points $\{y, w, w'\}$, where y is a fixed point associated to the loop $\{x_1, x_2\} \to \{x_1, x_2\}$ and $\{w, w'\}$ is a 2-periodic orbit associated to the loop $\{x_3, x_5\} \to \{x_1, x_4\} \to \{x_3, x_5\}$. Hence, from Theorem A it follows that the set of periods of each model (S, P, g) of the pattern $\mathcal{P}$ contains $\mathbb{N} \setminus \{1, 2, 3\}$.

In the previous paper [3], two of the authors and their coauthors proved a minimality result for *graph maps* which has the flavour of Theorem A (existence of a pattern-preserving injective map from the set of periodic orbits of a canonical model into that of a given graph map, with at most a finite number of exceptions). Since any tree is in particular a graph, one would expect that this result can be applied to tree maps. This idea does not work straightforwardly, mainly because the notion of *pattern* used in [3] for graph maps differs greatly from the corresponding notion used in [2] for tree maps. This difference arises from the distinct aims of both papers. However, by using some tools and results from [3], in a forthcoming paper we will prove that *there is a period-preserving injective map from the set of all negative significant periodic points of* *the A-linear model of a pattern $\mathcal{P}$ into the set of periodic points of any representative of $\mathcal{P}$.* It follows that the period-preserving injective map μ in Theorem A can be extended to the whole set of significant periodic points of the A-linear model. But observe that in this case we do not claim that the μ-corresponding orbits are associated to the same loop in the $\mathcal{P}$-path graph. Hence, this result fills the gap of negative significant periods between 2 and $M(S,P)$, but is weaker than Theorem A in the sense that it does not give any information about the localization of the corresponding orbits.

Thus one may wonder whether the hypotheses which define the set Λ_f in Theorem A (positivity or period greater than $M(S,P)$) are necessary in order to conclude that the μ-corresponding orbits are associated to the same loop in the path graph. Consider the pattern $\mathcal{P}$ and the fixed point y of the corresponding A-linear model (T,A,f) given in Example 1.6. The significant and negative fixed point y is associated to the loop $\{x_1,x_2\} \to \{x_1,x_2\}$. Now consider the model (T,A,g) of $\mathcal{P}$ shown in the right side of Figure 1. Then $g(v)=x_4$ and g is monotone between any pair of consecutive points of $A \cup V(T)$. The only fixed point of g located at the convex hull of the discrete component $\{x_1,x_2,x_4\}$ belongs to the interval $[v,x_4]$, and thus it is not associated to the loop $\{x_1,x_2\} \to \{x_1,x_2\}$. This example answers in the affirmative the above question. Thus Theorem A is optimum in this sense. Similar (and more complicated) examples can be build in which the period of the significant periodic point under consideration is greater than one.

This paper is organized as follows. In Section 2 we use the notion of a *canonical model*, first introduced in [2], to define a particular class of "simplified" monotone models called *linear models*. We also prove some minimality properties of the linear models. Finally in Section 3 we use these results to prove Theorem A.

Acknowledgements. — We thank an anonymous referee for his clever and detailed remarks, that have greatly improved the readability of this paper.

2. Linear models.

In this section we prove that any pattern admits a particular type of monotone model, which will be called an *A-linear model*. We also study the f-monotone equivalence classes of periodic points for A-linear models, which turn out to be specially simple.

We start by defining the notion of a *canonical model*, first introduced in [2]. Let (T, A, f) be a monotone model. We will say that $v_1, v_2 \in V(T) \setminus A$ are *f-identifiable* if either:

(i) $[f^i(v_1), f^i(v_2)] \cap A = \emptyset$ for all $i \geq 0$, or

(ii) if $[f^n(v_1), f^n(v_2)] \cap A \neq \emptyset$ for some $n \geq 0$, then $f^n(v_1) = f^n(v_2)$.

A monotone model (T, A, f) will be called a *canonical model* if it has no pairs of f-identifiable vertices. By Theorem B of [2], each pattern admits a canonical model. Moreover, it can be shown that this canonical model is unique (up to homeomorphisms of trees and conjugacy of maps). The procedure to obtain it from a generic monotone model (S, B, g) essentially consists of contracting the convex hull of each pair of f-identifiable vertices to a point. In doing it, besides keeping intact the B-monotonicity of the map and the pattern $[S, B, g]$, we eliminate some *non-significant* periodic orbits of vertices.

An *A-linear model* will be defined to be a canonical model which exhibits a certain piecewise linearity property. To introduce this condition, we must first consider an appropriate metric on trees. Given a tree T, a metric $d : T \times T \to [\mathbb{R}]$ such that $d(x, y) = d(x, z) + d(z, y)$ for each $x, y \in T$ and $z \in [x, y]$ will be called a *proper metric* (some authors use the term *taxicab metric*). Since a tree is a uniquely arcwise connected space, it admits a proper metric. For completeness, next we outline the basic ideas of the construction of a proper metric d on T. Let $I_1, I_2, \dots, I_n$ be the edges of T. For each of them choose a homeomorphism $h_i : I_i \to [0, 1]$ and let ϕ be the Lebesgue measure on $[0, 1]$. Take $x, y \in T$ and let $I_{i_1}, I_{i_2}, \dots, I_{i_k}$ be the set of edges of T whose interior intersects $[x, y]$. If $k = 1$, then $x, y \in I_{i_1}$ and we define $d(x, y) = \phi(h_{i_1}(x), h_{i_1}(y))$. Otherwise, assume that $x \in I_{i_1}$ and $y \in I_{i_k}$ and define $d(x, y) = k - 2 + d(x, a) + d(b, y)$, where a is the only point in $\mathrm{Bd}(I_{i_1}) \cap (x, y)$ and b is the only point in $\mathrm{Bd}(I_{i_k}) \cap (x, y)$. It is not difficult to show that d is a well defined proper metric on T.

Let T be a tree and let d be a proper metric on T. Let $I \subset T$ be a closed interval and let $f : I \to T$ be continuous. We say that f is *linear with respect to d* if $f(I)$ is either an interval or a point and there exists a real number $c \geq 0$ such that $d(f(x), f(y)) = c \cdot d(x, y)$ for all $x, y \in I$. The constant c will be called the *slope* of f. Observe that $f(I)$ reduces to a point if and only if $c = 0$.

Let $\mathcal{P}$ be a pattern. A canonical model (T, A, f) of $\mathcal{P}$ will be called *A-linear* if there exists a proper metric d on T such that f is linear with respect to d on each connected component of $T \setminus (A \cup V(T))$. The following result says that this condition is not restrictive.

Proposition 2.1. — *Each pattern $\mathcal{P}$ admits an A-linear model.*

Proof. — Let (T, A, g) be a canonical model of $\mathcal{P}$ and let d be a proper metric on T. The A-linear model is given by a map $f : T \to T$ such that $f_{|A \cup V(T)} = g_{|A \cup V(T)}$ and, for each $x \in T \setminus (A \cup V(T))$, $f(x)$ is defined to be the only point in $[f(v), f(w)]$ such that

$$d\big(f(v), f(x)\big) = \frac{d(f(v), f(w))}{d(v, w)} \cdot d(v, x)$$

where (v, w) is the only connected component of $T \setminus (A \cup V(T))$ which contains x. It is easy to check that f is well defined and satisfies the required properties. □

The following remark and lemma state the basic properties of the pre-images of the sets $A \cup V(T)$ and $V(T) \setminus A$ in an A-linear model (T, A, f).

Remark 2.2. — Let (T, A, f) be an A-linear model. For each $k \in \mathbb{N}$ set $V_k = f^{-k}(A \cup V(T))$. Since $A \cup V(T)$ is f-invariant, the sets V_k contain $A \cup V(T)$. Although V_k is not necessarily finite, the A-monotonicity of f implies that it has finitely many connected components, each of them being a point or a subtree on which f^k is constant. In addition, each connected component of $T \setminus V_k$ is an interval on which f^k is linear.

Lemma 2.3. — *If (T, A, f) is an A-linear model then $f^{-k}(V(T) \setminus A)$ is finite (or empty) for each $k \in \mathbb{N}$.*

Proof. — It is enough to prove that $f^{-k}(V(T) \setminus A) \cap K$ is finite or empty for the closure $K = [a, b]$ of any connected component of $T \setminus (A \cup V(T))$. Since $a, b \in A \cup V(T)$, by Remark 1.1 we have $f(a), f(b) \in A \cup V(T)$. If f is constant on K, then $\{a, b\} \cap A \neq \emptyset$ and $f(K) \subset A$, because (T, A, f) is canonical. If f is not constant on K, then f is linear with a positive slope on K, since (T, A, f) is A-linear. In both cases, it easily follows that $K \cap f^{-1}(X \setminus A)$ is finite or empty for each finite set $X \subset T$. In particular, $K \cap f^{-1}(V(T) \setminus A)$ is finite or empty. Since in addition this intersection is disjoint from A, the lemma easily follows from a simple inductive argument. □

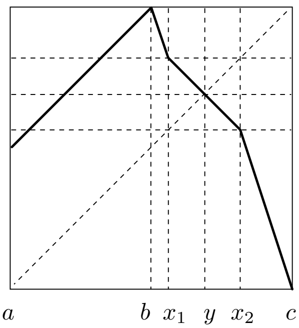

Figure 2. *A non-linear canonical model.*

The linearity condition in canonical models allows us removing most of the non-significant periodic orbits (not necessarily associated to orbits of vertices). For instance, Example 2.4 shows a non A-linear canonical model exhibiting a proper interval of non-significant 2-periodic points: these points will be removed when one constructs the corresponding A-linear model.

Example 2.4. — Let $T = [a, c]$ be an interval. Take $A = \{a, b, c\}$ and $B = \{x_1, x_2\}$, where $a < b < x_1 < x_2 < c$. Now consider a map f such that $f(a) = b$, $f(b) = c$, $f(c) = a$, $f(x_1) = x_2$, $f(x_2) = x_1$ and is linear between consecutive points of $A \cup B$ (see Figure 2). Observe that A is f-invariant and f is A-monotone. Since in addition there are no f-identifiable vertices, (T, A, f) is a canonical model of the pattern $[T, A, f]$. On the other hand, f is not linear on $[b, c]$ since it has not constant slope, so that (T, A, f) is not an A-linear model. Let y be the unique fixed point of f in $[x_1, x_2]$. Since $f(x_1) = x_2$, $f(x_2) = x_1$ and f is linear in $[x_1, x_2]$, then f has slope -1 in $[x_1, x_2]$. Therefore, each point in $[x_1, x_2] \setminus \{y\}$ is 2-periodic and belongs to the f-monotone equivalence class of y. Moreover, the class of y contains no points of A (it is not difficult to see that f^{3k} is not monotone on $[x, w]$ for each $w \in A$ and each $k \in \mathbb{N}$). It follows that y is a significant point of f (because its period is minimal and its class contains no points of A) whose f-monotone equivalence class coincides with $[x_1, x_2]$.

As a consequence of the fact that the linearity condition eliminates most of the non-significant periodic orbits in canonical models, the f-monotone equivalence classes are particulary simple in an A-linear model. In particular, from the following technical lemma it follows that the significant periodic points of A-linear models are alone in their monotone equivalence classes.

Lemma 2.5. — *Let (T,A,f) be an A-linear model and let x be a periodic point of f which is not f-monotone equivalent to any point of A. Then the f-monotone equivalence class of x reduces to $\{x\}$.*

Proof. — During this proof, let us write $a \sim b$ to denote that a and b are f-monotone equivalent. First we claim that, for each $i > 0$, the f-monotone equivalence class of $f^i(x)$ contains no points of A. Assume now that there exist $y \in A$ and $i > 0$ such that $f^i(x) \sim y$. Let $n \in \mathbb{N}$ be such that $f^n(x) = x$. Then

$$x = f^n(x) = f^{n-(i \bmod n)}\big(f^{i \bmod n}(x)\big) = f^{n-(i \bmod n)}\big(f^i(x)\big).$$

Thus, by Remark 1.3, $x \sim f^{n-(i \bmod n)}(y) \in A$. This contradiction proves the claim.

Next we prove that

$$\text{if } a \sim x \sim b \text{ and } a \neq b \text{ then } [f^i(a), f^i(b)] \cap A = \emptyset \text{ for each } i \geq 0. \tag{2.1}$$

Let $i \geq 0$. By Remark 1.3, $f^i(a) \sim f^i(b)$. Hence, there exists $k \in \mathbb{N}$ such that $f^i(a)$ and $f^i(b)$ are fixed points of f^k and f^k is monotone on $[f^i(a), f^i(b)]$. Assume that there exists $w \in [f^i(a), f^i(b)] \cap A$. Let $\leq$ be the order in $[f^i(a), f^i(b)]$ such that $f^i(a) < f^i(b)$. Since $f^k|_{[f^i(a),f^i(b)]}$ is monotone and $f^i(a)$ and $f^i(b)$ are fixed points of f^k, it follows that $f^k([f^i(a), f^i(b)]) = [f^i(a), f^i(b)]$ and f^k is increasing (with respect to the ordering chosen above). Therefore, $\mathrm{Orb}_{f^k}(w) \subset [f^i(a), f^i(b)] \cap A$. Then, from the finiteness of A and the monotonicity of f^k it follows that there is a fixed point $w' \in [f^i(a), f^i(b)] \cap A$ of f^k. Therefore, $f^i(x) \sim f^i(a) \sim w' \in A$, in contradiction with the claim above. This proves (2.1).

Now assume that there exists some $z \neq x$ such that $z \sim x$. This will lead us to a contradiction. Let $m \in \mathbb{N}$ such that $f^m(x) = x$, $f^m(z) = z$ and f^m is monotone on $[x, z]$. Choose $e, e' \in \mathrm{En}(T)$ and an order $\leq$ in $[e, e']$ with $e < x < z < e'$.

By (2.1), $f^{-m}(A) \cap [x, z] = \emptyset$. Hence,

$$f^{-m}\big(A \cup V(T)\big) \cap [x, z] = f^{-m}\big(V(T) \setminus A\big) \cap [x, z],$$

which is a finite (or empty) set by Lemma 2.3. Moreover, f^m is linear between two consecutive points of $f^{-m}(V(T) \setminus A) \cap [x, z]$. Thus, f^m is strictly increasing on $[x, z]$ (with respect to the ordering chosen above) and there exist $\ell \geq 0$ and points $x = a_0 < a_1 < a_2 < \ldots < a_{\ell+1} = z$ such that $a_i \in f^{-m}(V(T) \setminus A)$ for $1 \leq i \leq \ell$ and f^m is linear on each interval of the form $[a_i, a_{i+1}]$. If for some $1 \le i \le \ell$ we had $f^m(a_i) \neq a_i$, then it easily follows that $(f^{mn}(a_i))_{n \ge 0}$ would be an infinite and strictly monotone sequence contained in $\langle a_i, b\rangle_T \cap V(T)$, where b is a fixed point of f^m in $[x, z]$. This contradicts the finiteness of $V(T)$. In consequence, $f^m(a_i) = a_i$ for $0 \le i \le \ell + 1$ and f^m is the identity map on $[x, z]$.

Observe that $\{x, z\} \not\subset V(T)$. Otherwise, by virtue of (2.1), x and z would be f-identifiable, contradicting the fact that (T, A, f) is a canonical model. Assume that $z \notin V(T)$ (the case $x \notin V(T)$ is symmetric). Since $e, e' \in \mathrm{En}(T) \subset A \subset f^{-m}(A \cup V(T))$ and $z \notin f^{-m}(A \cup V(T))$, from Remark 2.2 it follows that there exist $a, b \in [e, e']$ with $a < z < b$ and $[a, b] \cap f^{-m}(A \cup V(T)) = \{a, b\}$. Moreover, f^m is linear on $[a, b]$. Since, by the preceding paragraph, f^m is the identity on $[x, z]$, f^m is the identity on $[a, b]$. Thus $f^m(a) = a \in A \cup V(T)$, $f^m(b) = b \in A \cup V(T)$ and $a \sim z \sim x \sim b$. Since the f-monotone equivalence class of x does not contain any point of A, it follows that $a, b \in V(T)$. By (2.1), a and b are f-identifiable, in contradiction with the fact that (T, A, f) is a canonical model. $\square$

Remark 2.6. — Recall that a periodic point x of a monotone model (T, A, f) is significant if: (i) x is not f-monotone equivalent to any element of $A \cup V(T)$; and (ii) the period of x is minimal within its class. By Lemma 2.5, if (T, A, f) is A-linear then each point verifying (i) verifies also (ii). Therefore, *a periodic point of an A-linear model (T, A, f) is significant if and only if it is not f-monotone equivalent to any point of $A \cup V(T)$.*

The following remark will be used several times in the rest of this section. It easily follows from the definition of the f-monotone equivalence relation.

Remark 2.7. — Let (T, A, f) be a monotone model and let x and y be f-monotone equivalent periodic points of f. If z is a periodic point of f such that $z \in [x, y]$, then z is also f-monotone equivalent to x and y.

The rest of this section is devoted to state and prove Proposition 2.10, which describes the structure of the set of non-significant periodic points of A-linear models and will be used in the proof of statement (b) of Theorem A. To prove it, we need the next two technical lemmas.

LEMMA 2.8. — *Let (T, A, f) be a monotone model. Let $x, y, z \in T$ be three different points of a periodic orbit of f such that $y \in (x, z)$. Then x, y and z do not belong to the same f-monotone equivalence class.*

Proof. — Assume that x, y and z are f-monotone equivalent and this will lead us to a contradiction. From the definition of the f-monotone relation and the fact that $y \in (x,z)$ it easily follows that

$$f^i(y) \in \big(f^i(x), f^i(z)\big) \quad \text{for all } i \geq 0. \tag{2.2}$$

Set $X = \langle \operatorname{Orb}(x) \rangle_T$. Since $\operatorname{Orb}(x)$ is a periodic orbit containing y and $\operatorname{En}(X) \subset \operatorname{Orb}(x)$, there exists some $i \geq 0$ such that $f^i(y) \in \operatorname{En}(X)$, a contradiction with (2.2). □

Lemma 2.9. — *Let (T,A,f) be a monotone model of a periodic pattern and let x be a periodic point of f of period greater than $8|A|-4$. The following statements hold*:

(a) $\operatorname{Orb}(x) \cap (A \cup V(T)) = \emptyset$ *and there is a connected component of* $T \setminus (A \cup V(T))$ *containing five different points of* $\operatorname{Orb}(x)$.

(b) *If (T,A,f) is A-linear then x is significant.*

Proof. — Let us prove (a). Set $Q = \operatorname{Orb}(x)$, which does not coincide with A since $|Q| > |A|$. By Remark 1.1, $A \cup V(T)$ is f-invariant. Since Q and A are periodic orbits, it follows that either $Q \subset V(T) \setminus A$ or $Q \cap (A \cup V(T)) = \emptyset$. It is easy to check that any tree S satisfies $|V(S) \setminus \operatorname{En}(S)| \leq |\operatorname{En}(S)|$. Since $\operatorname{En}(T) \subset A$, it follows that

$$\big|V(T) \setminus A\big| \leq \big|V(T) \setminus \operatorname{En}(T)\big| \leq \big|\operatorname{En}(T)\big| \leq |A|. \tag{2.3}$$

From (2.3) and the fact that $|A| < |Q|$ we have that $Q \not\subset V(T) \setminus A$. In consequence, $Q \cap (A \cup V(T)) = \emptyset$.

We prove that there exists a connected component of $T \setminus (A \cup V(T))$ containing at least five different points of Q. On the contrary, $|Q| \leq 4k$ where k is the number of connected components of $T \setminus (A \cup V(T))$. It is well known that the number of connected components of $S \setminus X$ equals $|X|-1$ for any tree S and any finite set $X \subset S$ such that $V(S) \subset X$. Hence, we have

$$|Q| \leq 4k = 4\big(|A \cup V(T)| - 1\big). \tag{2.4}$$

On the other hand, $|A \cup V(T)| = |A| + |V(T) \setminus A| \leq 2|A|$ by (2.3). Thus from (2.4) we get $|Q| \leq 8|A| - 4$, a contradiction.

Next we prove (b). Since $x \notin V(T)$, by Remark 2.6 and Lemma 2.5 it is enough to show that the f-monotone equivalence class of x does not contain any point of A. Assume, on the contrary, that x is f-monotone equivalent to $w \in A$.

From (a) it follows that there are five different points $a, b, c, d, e \in \mathrm{Orb}(x)$ and an order $\leq$ on $[a, e]$ such that $a < b < c < d < e$ and $[a, e] \cap (A \cup V(T)) = \emptyset$. Since x is f-monotone equivalent to $w \in A$, by Remark 1.3, c is f-monotone equivalent to some $w' \in A$. Thus there exists $m \in \mathbb{N}$ such that c and w' are fixed points of f^m and f^m is monotone on $[c, w']$. Since $[a, e] \cap (A \cup V(T)) = \emptyset$, either $a \in (w', c)$ or $e \in (c, w')$. Assume without loss of generality that $e \in (c, w')$. Since d and e are periodic points, by Remark 2.7 they are f-monotone equivalent to c, in contradiction with Lemma 2.8. □

Let (T, A, f) be an A-linear model. By Lemma 2.5, all the f-monotone equivalence classes which do not intersect A reduce to a single point. On the other hand, there are obviously finitely many classes intersecting A. Proposition 2.10 states that the set of these classes has a cyclic structure.

Proposition 2.10. — *Let (T,A,f) be an A-linear model of a periodic pattern. There exists a divisor k of $|A|$ and a labelling $C_0, C_1, \dots, C_{k-1}$ of the f-monotone equivalence classes intersecting A such that $f(\langle C_i \rangle) = \langle C_{i+1 \bmod k} \rangle$ and f maps bijectively C_i onto $C_{i+1 \bmod k}$ for $0 \leq i \leq k-1$. Moreover, $\langle C_i \rangle \cap \langle C_j \rangle = \emptyset$ whenever $i \neq j$.*

Proof. — By definition of the f-monotone equivalence relation, all the points in a class are periodic. This proves that f is injective on each class. Let us call any f-monotone equivalence class intersecting A an *A-class*. From Remark 1.3, the f-image of a class is contained in a class. Moreover, the f-image of an A-class is contained in an A-class because A is f-invariant. Since all the points in a class are periodic and A is a single periodic orbit, it easily follows that there exists $k \in \mathbb{N}$ and a labelling $C_0, C_1, \dots, C_{k-1}$ of the A-classes such that $f(C_i) \subset C_{i+1 \bmod k}$ for $0 \leq i \leq k-1$. Let x be an n-periodic point which belongs to some C_i. Then $f^{n-1}(x)$ is a pre-image of x which must belong to C_{i-1}. Therefore, f is onto and thus bijective on each C_i. Consequently, $f(C_i) = C_{i+1 \bmod k}$. Observe that k is a divisor of the period of each periodic point contained in an A-class. In particular, since A is a periodic orbit, k divides $|A|$.

We claim that each A-class is a closed set. For each $n \in \mathbb{N}$, let $\mathcal{P}_n$ denote the set of fixed points of f^i for $1 \leq i \leq n$. It is clear that $\mathcal{P}_n$ is a finite union of closed sets and thus is closed. Moreover, since (T, A, f) is an A-linear model, it easily follows that $\mathcal{P}_n$ has finitely many connected components. Let C be an A-class and set $n = 8|A| - 4$. By Lemma 2.9, the period of each point in C is not greater than n. Hence, $C \subset \mathcal{P}_n$. To prove the claim it is enough to show that $C \cap K$ is closed for each connected component K of $\mathcal{P}_n$ such that $C \cap K \neq \emptyset$. By Remark 2.7, the set $C \cap K$ is convex. From the definition of the f-monotone equivalence relation, if two periodic points x and y belong to an A-class then there exists some $m \in \mathbb{N}$ such that $f^m(x) = x$, $f^m(y) = y$ and f^m is monotone on $[x, y]$. Recall that the periods of x and y are either equal, or one is a multiple of the other. Therefore, all the points in $C \cap K$ are pairwise f-monotone equivalent and fixed points of f^m for some $m \leq n$. Thus f^m is the identity map on $C \cap K$. In consequence, each accumulating point of a sequence of points in $C \cap K$ belongs to $C \cap K$, and the claim follows.

Let x and y be different points belonging to an A-class, and let $m \in \mathbb{N}$ be such that $f^m(x) = x$, $f^m(y) = y$ and f^m is monotone on $[x, y]$. In particular, $f^m([x, y]) = [x, y]$. Assume that $(x, y) \cap A = \emptyset$. Since A is f-invariant and $f^m((x, y)) = (x, y)$, it follows that $f^i(x, y) \cap A = \emptyset$ for each $i \geq 0$. Moreover, since f is A-monotone, a simple inductive argument shows that $f^i((x, y)) = (f^i(x), f^i(y))$ for each $i \geq 0$. Consequently, since $f(C_i) = C_{i+1 \bmod k}$, we obtain that $f(\langle C_i \rangle) = \langle C_{i+1 \bmod k} \rangle$ for $0 \leq i \leq k-1$.

Finally we prove that $\langle C_i \rangle \cap \langle C_j \rangle = \emptyset$ whenever $i \neq j$. This is obvious when each A-class reduces to one point. Thus from now on assume that some A-class contains at least two points. Hence, each A-class contains at least two points, because f is a bijection over the A-classes. Assume that there are $0 \leq i < j \leq k - 1$ such that $X := \langle C_i \rangle \cap \langle C_j \rangle \neq \emptyset$. This will lead us to a contradiction. Observe that X is a subtree of T. Moreover, by the preceding paragraph $f(X) \subset \langle C_{i+1} \rangle \cap \langle C_{j+1} \rangle$, and iterating k times this argument we get $f^k(X) \subset X$. Thus, $f^k|_X$ is a tree map and so it has fixed points. Hence, there exists a periodic point $z \in X$ of f. By the claim above, each A-class is closed and thus $\operatorname{En}(\langle C_i \rangle) \subset C_i$. Since $z \in \langle C_i \rangle$ and each A-class contains at least two points, there exist $x, y \in \operatorname{En}(\langle C_i \rangle) \subset C_i$ such that $z \in [x, y]$. Analogously, there exist $x', y' \in \operatorname{En}(\langle C_j \rangle) \subset C_j$ such that $z \in [x', y']$. Since z is periodic, from Remark 2.7 it follows that $z \in C_i \cap C_j$, a contradiction since $C_i \cap C_j = \emptyset$. □

3. Proof of Theorem A.

We start this section with four preliminary results which will be used in the proof of Theorem A. They are three technical lemmas and

Proposition 3.4, which is a refinement of Theorem 1.4 and plays a central role in the proof of Theorem A.

LEMMA 3.1. — *Let $\mathcal{P}$ be a pattern. For $1 \le j \le s$, let $\pi_0 \to \pi_1^j \to \dots \to \pi_{n-1}^j \to \pi_0$ be s different loops of length n in the $\mathcal{P}$-path graph, and let $g\colon S \to S$ be a tree map exhibiting $\mathcal{P}$. Then there exist subsets $J^1, J^2, \dots, J^s$ of $\langle \pi_0 \rangle$ such that, for each $j = 1,2,\dots,s$, we have:*

(a) $J^j \subset \langle \pi_0 \rangle$, $g^i(J^j) \subset \langle \pi_i^j \rangle$ *for* $1 \le i < n$ *and* $g^n(J^j) = \langle \pi_0 \rangle$.

(b) $J^j = \cup_{i=1}^m [a_i, b_i]$, *with:*

(b.1) $a \le a_1 < b_1 \le a_2 < b_2 \le \dots \le a_m < b_m \le b$ *(where* $[a,b] = \langle \pi_0 \rangle$ *and* $\le$ *is an order on* $[a,b]$*).*

(b.2) $g^n|_{\{a_1, b_1, \dots, a_m, b_m\}}$ *is monotone with respect to* $\le$.

(b.3) $g^n(b_i) = g^n(a_{i+1})$ *for* $i = 1,2,\dots,m-1$.

(b.4) $g^n([a_i,b_i]) \subset [g^n(a_i), g^n(b_i)]$.

Furthermore, $\operatorname{Int}(\langle J^j \rangle) \cap \operatorname{Int}(\langle J^k \rangle) = \emptyset$ *when* $j \ne k$.

Proof. — Statement (a) is simply a particular instance of Lemma 3.2 of [2], and statements (b.1)–(b.4) follow immediately from its proof. □

Lemma 3.2 states that from a closed sequence of coverings of sets contained in discrete components we can derive the existence of a periodic orbit of vertices.

LEMMA 3.2. — *Let (T,A,f) be a monotone model and let $n \in \mathbb{N}$. Assume that for each $0 \le i < n$ there is a set $D_i \subset A$ satisfying:*

(a) *D_i is a subset of a discrete component of (T,A) with $|D_i| \ge 3$.*

(b) *If π is a basic path contained in D_i, then there is a basic path σ contained in $D_{i+1 \bmod n}$ such that $\langle f(\pi) \rangle \supset \langle \sigma \rangle$.*

Then there exists a periodic point w of f such that, for each $i \ge 0$, $f^i(w) \in \operatorname{Int}(\langle D_{i \bmod n} \rangle) \cap V(\langle D_{i \bmod n} \rangle)$.

Proof. — Set $T_i = \langle D_i \rangle$ for $0 \le i < n$. By (a), $A \cap T_i = D_i = \operatorname{En}(T_i)$ and $|\operatorname{En}(T_i)| \ge 3$. Hence, $V(T_i) \cap \operatorname{Int}(T_i) \ne \emptyset$ for $0 \le i < n$. We claim that if $v \in V(T_i) \cap \operatorname{Int}(T_i)$ then $f(v) \in V(T_{i+1 \bmod n}) \cap \operatorname{Int}(T_{i+1 \bmod n})$. Observe that if the claim holds then, since $V(T)$ is finite, the lemma follows.

Let us prove the claim. Choose three different points $a, b, c \in A \cap T_i$ such that $\langle \{a, b, c\} \rangle$ is a 3-star with central point v. From (b) it follows that $f(a), f(b), f(c)$ are 3 pairwise different points. Hence, since f is A-monotone, $\langle\{f(a), f(b), f(c)\}\rangle$ is either an interval or a 3-star. If it is an interval, then there is a point in $\{a, b, c\}$ (assume without loss of generality that it is b) such that $f(b) \in (f(a), f(c))$. Thus, $f(a)$ and $f(c)$ belong to two different discrete components, in contradiction with (b). Therefore, $\langle\{f(a), f(b), f(c)\}\rangle$ is a 3-star. By Remark 1.1, $f(v)$ is the central point of $\langle\{f(a), f(b), f(c)\}\rangle$. From (b) it easily follows that there exist points a', b', c' in $D_{i+1 \bmod n}$ such that $a' \in (f(v), f(a)]$, $b' \in (f(v), f(b)]$ and $c' \in (f(v), f(c)]$. Hence, $f(v)$ is also the central point of the 3-star $\langle\{a', b', c'\}\rangle$, which is contained in $T_{i+1 \bmod n}$, and the claim follows. □

Theorem A (a) will essentially be a corollary of the next lemma. To state this result, we need to introduce a few more notions.

To each loop in a $(\mathcal{T}, [\theta])$-path graph we can associate a *sign* as follows. First we endow each basic path π of $(\mathcal{T}, [\theta])$ with an ordering $\leq_\pi$. Clearly, each $\leq_\pi$ induces a linear ordering on $\langle\pi\rangle_T$ and $\langle\theta(\pi)\rangle_T$. Let $\preceq_\pi$ be the linear ordering induced by $\leq_\pi$ on $\langle\theta(\pi)\rangle_T$ and let $\pi \to \sigma$ be an arrow in the $(\mathcal{T}, [\theta])$-path graph. Set $\pi = \{a, b\}$ and $\sigma = \{c, d\}$ and assume that $a <_\pi b$ and $c <_\sigma d$. Then the sign of the arrow $\pi \to \sigma$ is $+1$ if $\theta(a) \preceq_\pi c \prec_\pi d \preceq_\pi \theta(b)$ and -1 if $\theta(a) \preceq_\pi d \prec_\pi c \preceq_\pi \theta(b)$. A loop $\pi_0 \to \pi_1 \to \cdots \to \pi_{n-1} \to \pi_0$ is said to be *positive* if the product of the signs of the arrows $\pi_0 \to \pi_1$, $\pi_1 \to \pi_2$, ... , $\pi_{n-1} \to \pi_0$ is $+1$ and *negative* if it is -1. It is easy to see that the sign of a loop is independent from the particular choice of the set of orderings $\leq_\pi$. If (T, A, f) is a model of $(\mathcal{T}, [\theta])$ and there is a periodic point x of f associated to a positive (negative) loop, then we also say that x is *positive* (respectively, *negative*).

Now we are ready to state and prove Lemma 3.3. By Theorem 1.4, each significant periodic point of a monotone model is associated to a unique simple loop in the path graph. Therefore, the loop β in the statement of Lemma 3.3 is unique. Recall also that the constant $M(S, P)$ has been defined in just before Theorem A.

Lemma 3.3. — *Let (S,P,g) be a model of a periodic pattern $(\mathcal{T},\Theta)$ and let (T,A,f) be a monotone model of $(\mathcal{T},\Theta)$. Let $x \in T$ be a significant n-periodic point of f and let β be the simple loop of length n in the $(\mathcal{T},\Theta)$-path graph such that x and β are associated. The following statements hold*:

(a) *If either β is positive or $n > M(S,P)$ then there exists a point $z \in S \setminus P$ associated to β such that $g^n(z) = z$.*

(b) *Each fixed point of g^n in $S \setminus P$ associated to β is an n-periodic point of g.*

Consequently, if either β is positive or $n > M(S,P)$ then there is an n-periodic point of g in $S \setminus P$ associated to β.

Proof. — We start by claiming that if a fixed point z of g^n is associated to β then $z \notin P$. Indeed: if $z \in P$ then $\mathrm{Orb}_g(z) = P$ and, given a bijection $\theta : P \to A$ which preserves discrete components, $\theta(z)$ is a point of A associated to β. Since β is simple and x is also associated to β, from Theorem 1.4(b) it follows that x and $\theta(z)$ are f-monotone equivalent, a contradiction with the fact that x is significant. This proves the claim.

Let $\pi_0 \to \pi_1 \to \cdots \to \pi_{n-1} \to \pi_0$ be the loop β. By the claim above, to prove (a) it is enough to show that there exists a fixed point of g^n associated to β. We will consider two cases.

- *Case 1*: *β is positive.*

By Lemma 3.1 (with $s = 1$), there exists a finite union $J = \bigcup_{i=1}^m [a_i, b_i] \subset S$ of intervals with pairwise disjoint interiors such that, if $\langle \pi_0 \rangle_S = [a,b]$ and $\le$ is an ordering on $[a,b]$, then $a \le a_1 < b_1 \le a_2 < b_2 \le \cdots \le a_m < b_m \le b$ and:

(i) $g^i(J) \subset \langle \pi_i \rangle$ for $0 \le i < n$ and $g^n(J) = [a,b]$.

(ii) $g^n|_{\{a_1, b_1, \dots, a_m, b_m\}}$ is monotone with respect to $\le$.

(iii) $g^n(b_i) = g^n(a_{i+1})$ for $i = 1, 2, \dots, m-1$.

Since β is positive, $g^n|_{\{a_1, b_1, \dots, a_m, b_m\}}$ is increasing. An easy argument, analogous to the one used in Lemma 3 of [6], shows that there exists $z \in J$ such that $g^n(z) = z$. By (i), z and β are associated. This ends the proof of (a) in this case.

- *Case 2*: *β is negative* (then $n > M(S,P)$ by hypothesis).

From the definition of $M(S,P)$ and the fact that $n > M(S,P)$, it follows that there is a basic path π in the loop β satisfying the following property: if s is the number of occurrences of π in the loop β, and $r = |\,\mathrm{Int}(\langle \pi \rangle) \cap V(S)|$, then $s > r$. Assume without loss of generality that $\pi = \pi_0$. By considering the s shifts of β starting at π_0, we have s loops $\pi_0 \to \pi_1^j \to \cdots \to \pi_{n-1}^j \to \pi_0$ for $1 \le j \le s$. Since β is simple, it is not difficult to see that these loops are pairwise different. Then, by Lemma 3.1, there exist subsets $J^1, J^2, \dots, J^s$ of $\langle \pi_0 \rangle$ which consist of finite unions of closed intervals such that for each $j = 1, 2, \dots, s$ we have

$J^j \subset \langle \pi_0 \rangle$, $g^i(J^j) \subset \langle \pi_i^j \rangle$ for $1 \le i < n$, and $g^n(J^j) = \langle \pi_0 \rangle$. Moreover, $\operatorname{Int}(\langle J^k \rangle) \cap \operatorname{Int}(\langle J^j \rangle) = \emptyset$ for $j, k \in \{1, 2, \dots, s\}$ with $j \neq k$.

Choose an ordering $\le$ for $\langle \pi_0 \rangle$. Without loss of generality, assume that the sets J^j for $1 \le j < s$ are labelled in such a way that $x \le y$ for each pair of points x, y such that $x \in \langle J^j \rangle$ and $y \in \langle J^{j+1} \rangle$. Observe that, given $j \in \{1, 2, \dots, s\}$, if there are no fixed points of g^n in J^j then there exist $w \in \langle J^j \rangle \cap V(S)$ and $a, b \in J^j$ such that $a < w < b$, $g^n(a) = g^n(b) = w$ and $g^n((a, b)) \cap \langle \pi_0 \rangle = \emptyset$. In particular, $w \in \operatorname{Int}(\langle J^j \rangle)$ (see Figure 3). Since there are r vertices in $\operatorname{Int}(\langle \pi_0 \rangle)$ and $s > r$, necessarily there exists $k \in \{1, 2, \dots, s\}$ such that $\operatorname{Int}(\langle J^k \rangle) \cap V(S) = \emptyset$. Then there is a fixed point of g^n in J^k. This ends the proof of (a).

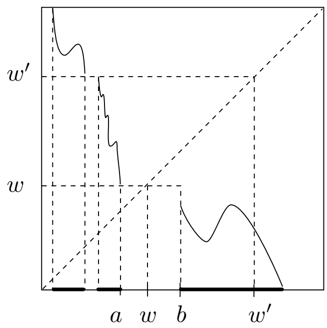

Figure 3. The graph of g^n on some J^i. The points w and w' are vertices and the shadowed segments correspond to the intervals that form J^i.

Now let us prove (b). Let z be a fixed point of g^n in $S \setminus P$ associated to β. We must prove that z is an n-periodic point of g.

We claim that given any path $\sigma_0 \to \sigma_1 \to \cdots$ in the $[S, P, g]$-path graph such that $\sigma_i = \sigma_j$ for some $i, j \ge 0$, it follows that:

$$(3.1) \qquad \begin{cases} \text{If } \operatorname{Int}(\langle \sigma_{i+1} \rangle_T) \cap \operatorname{Int}(\langle \sigma_{j+1} \rangle_T) \neq \emptyset \text{ or, equivalently,} \\ \operatorname{Int}(\langle \sigma_{i+1} \rangle_S) \cap \operatorname{Int}(\langle \sigma_{j+1} \rangle_S) \neq \emptyset, \text{ then } \sigma_{i+1} = \sigma_{j+1}. \end{cases}$$

Let us prove this claim. By the definition of the path graph, $\langle f(\sigma_i) \rangle_T \supset \langle \sigma_{i+1} \rangle_T$ and $\langle f(\sigma_i) \rangle_T \supset \langle \sigma_{j+1} \rangle_T$. Moreover, if $\operatorname{Int}(\langle \sigma_{i+1} \rangle_T) \cap \operatorname{Int}(\langle \sigma_{j+1} \rangle_T)$ is not empty then σ_{i+1} and σ_{j+1} belong to the same discrete component. Then (3.1) follows immediately from the fact that $\langle f(\sigma_i) \rangle_T$ is an interval (since f is A-monotone), and the claim is proved.

Now assume that $n = lk$ for some $l > 1$ and that $|\operatorname{Orb}_g(z)| = k$. This will lead us to a contradiction. Until the end of the proof, the integer subindexes will be considered modulo n. The fact that $z \notin P$ implies that $g^i(z) \in \operatorname{Int}(\langle \pi_i \rangle_S)$ for $i \geq 0$. Therefore, since $g^i(z) = g^{i+k}(z)$ for each $i \geq 0$,

$$\begin{cases} \operatorname{Int}(\langle \pi_i \rangle_S) \cap \operatorname{Int}(\langle \pi_{i+k} \rangle_S) \neq \emptyset \text{ for each } i \geq 0 \text{ or, equivalently,} \\ \operatorname{Int}(\langle \pi_i \rangle_T) \cap \operatorname{Int}(\langle \pi_{i+k} \rangle_T) \neq \emptyset \text{ for each } i \geq 0. \end{cases} \tag{3.2}$$

For $s = 0, 1, \dots, k-1$ set $\Pi_s = \{\pi_s, \pi_{s+k}, \dots, \pi_{s+(\ell-1)k}\}$. Now we prove that $|\Pi_s| = \ell$ for each $0 \leq s < k$. By considering (if necessary) a shift of β instead of β itself, we may assume that $s = 0$. Thus we must prove that if $i, j \in \{0, 1, \dots, \ell-1\}$ and $i \neq j$ then $\pi_{ik} \neq \pi_{jk}$. By considering again a shift of β, we can assume that $i = 0$ and $1 \leq j < \ell$. We proceed by induction on j. For $j = 1$, we must see that $\pi_0 \neq \pi_k$. If $\pi_0 = \pi_k$ then from (3.1) and (3.2) it follows that $\pi_1 = \pi_{k+1}$, $\pi_2 = \pi_{k+2}, \dots, \pi_{(\ell-1)k} = \pi_n$. Thus $\beta = \alpha^\ell$ with $\alpha = \pi_0 \to \pi_1 \to \cdots \to \pi_k$, contrary to the fact that β is simple. Now assume that we have proved that no basic path $\pi_k, \pi_{2k}, \dots, \pi_{(j-1)k}$ is π_0 for some $j < \ell - 1$. If $\pi_0 = \pi_{jk}$, then we write $\ell = qj + r$ for some $q \in \mathbb{N}$ and $0 \leq r < j$. As above, by using $j(q-1)k$ times (3.1) and (3.2) we get $\beta = \alpha^q \gamma$, where α is the loop $\pi_0 \to \pi_1 \to \ldots \to \pi_{jk} = \pi_0$ and γ is $\pi_{qjk} \to \pi_{qjk+1} \to \pi_{qjk+2} \to \ldots \to \pi_{qjk+rk} = \pi_n = \pi_0$. If $r = 0$ then $\beta = \alpha^q$, a contradiction with the fact that β is simple. If $0 < r < j$, since $\pi_0 = \pi_{qjk}$ we can use again rk times (3.1) and (3.2) in order to obtain $\pi_1 = \pi_{qjk+1}, \pi_2 = \pi_{qjk+2}, \dots, \pi_{rk} = \pi_{jqk+rk} = \pi_n = \pi_0$, contrary to the induction hypotheses. Therefore, $\pi_0 \neq \pi_{jk}$. Consequently, $|\Pi_s| = \ell$ for each $0 \leq s < k$ as we claimed.

For $s \in \{0, 1, \dots, k-1\}$, we define $D_s = \bigcup_{\pi \in \Pi_s} \pi$. In view of (3.2), each D_s is contained in a discrete component. Therefore, (i) of Lemma 3.2 is satisfied because $l \geq 2$, while (ii) is obviously verified. Hence, by Lemma 3.2 there is a periodic point w of f such that $f^s(w) \in \operatorname{Int}(\langle D_{s \bmod k} \rangle_T) \cap V(T)$ for each $s \geq 0$. In particular, there exists $0 \leq i < l$ such that $w \in \langle \pi_{ik} \rangle_T$. Note that, for each $j \geq 0$, $\pi_{ik+j+1 \bmod n}$ is the only basic path in $D_{j+1 \bmod k}$ contained in $f(\langle \pi_{ik+j \bmod n} \rangle)$. Consequently, w is associated to a shift of the loop β starting with π_{ik}. Therefore, there is a point $w' \in \operatorname{Orb}_f(w)$ associated to β and, by Theorem 1.4, x and w' are f-monotone equivalent. Since $w' \in V(T)$, this is a contradiction with the fact that x is significant.

□

Proposition 3.4. — *Let $g\colon S \to$ be a tree map, let P be a periodic orbit of g and let (T,A,f) be an A-linear model of the pattern $[S,P,g]$.*

There exists an injective map β from the set of significant periodic points x of f into the set of simple loops in the $[S,P,g]$-path graph such that x is associated to $\beta(x)$ and the length of $\beta(x)$ equals the period of x.

Proof. — We start by defining the map β. For each significant periodic point x of f, by Theorem 1.4(a) there is a unique simple loop in the $[S,P,g]$-path graph, which we define to be $\beta(x)$, of length $|\operatorname{Orb}_f(x)|$ and associated to x. Let us see that the map β is injective. Take a significant periodic point $x' \neq x$ and assume that $\beta(x') = \beta(x)$. By Theorem 1.4 (b), x and x' are f-monotone equivalent, in contradiction with Lemma 2.5. This ends the proof of the proposition. □

Remark 3.5. — For a periodic pattern $\mathcal{P}$ of period n, the number of vertices of the $\mathcal{P}$-path graph coincides with the number of basic paths of $\mathcal{P}$, which is not greater than $\binom{n}{2}$. On the other hand, any finite combinatorial graph with r vertices has at most r^k loops of length k. Therefore, by Proposition 3.4, an A-linear model f of $\mathcal{P}$ has at most $\binom{n}{2}^k$ significant periodic points of period k. Consequently, if (S,P,g) is any model of $\mathcal{P}$ then the number of significant periodic points of f with period not greater than any positive integer (in particular, $M(S,P)$) is finite.

Finally we are ready to prove Theorem A.

Proof of Theorem A. — We start by proving (a). We have to define the map μ on Λ_f. Let $x \in \Lambda_f$ and set $n = |\operatorname{Orb}_f(x)|$. Let β be the injective map from the set of significant periodic points of f into the set of simple loops in the $[S,P,g]$-path graph given by Proposition 3.4, so that x and $\beta(x)$ are associated and the length of $\beta(x)$ is n. Since $x \in \Lambda_f$, either $\beta(x)$ is positive or $n > M(S,P)$. Hence, by Lemma 3.3, we can choose an n-periodic point of g in $S \setminus P$, denoted by $\mu(x)$, associated to $\beta(x)$.

To end the proof of (a) we have to show that μ is injective. Take $x, x' \in \Lambda_f$ with $x \neq x'$ and assume that $\mu(x) = \mu(x')$. This will lead us to a contradiction. Set $y = \mu(x) \in S \setminus P$. Let us write $\beta(x)$ and $\beta(x')$ as $\pi_0 \to \pi_1 \to \cdots \to \pi_{n-1} \to \pi_0$ and $\pi'_0 \to \pi'_1 \to \cdots \to \pi'_{n-1} \to \pi'_0$ respectively. Since $y \in \operatorname{Int}(\langle \pi_0 \rangle_S) \cap \operatorname{Int}(\langle \pi'_0 \rangle_S)$, it easily follows that $g^i(y) \in \operatorname{Int}(\langle \pi_i \rangle_S) \cap \operatorname{Int}(\langle \pi'_i \rangle_S)$ for each $0 \le i < n$. Hence,

$$\begin{cases} \operatorname{Int}(\langle \pi_i \rangle_S) \cap \operatorname{Int}(\langle \pi'_i \rangle_S) \neq \emptyset \text{ for } 0 \le i < n \text{ or, equivalently,} \\ \operatorname{Int}(\langle \pi_i \rangle_T) \cap \operatorname{Int}(\langle \pi'_i \rangle_T) \neq \emptyset \text{ for } 0 \le i < n. \end{cases} \tag{3.3}$$

Since β is injective, $\beta(x) \neq \beta(x')$. This means that $\pi_j \neq \pi'_j$ for some $0 \le j < n$. Since $\operatorname{Int}(\langle \pi_j \rangle_S) \cap \operatorname{Int}(\langle \pi'_j \rangle_S) \neq \emptyset$ by (3.3), the same argument which has been used to prove (3.1) shows that $\pi_{j-1} \neq \pi'_{j-1}$. Repeating backwards the same argument, we obtain that $\pi_i \neq \pi'_i$ for each $0 \leq i < n$.

For each $0 \leq i < n$, we set $D_i = \pi_i \cup \pi'_i$. Observe that each D_i obviously satisfies the hypothesis (ii) of Lemma 3.2, while (i) is satisfied as a consequence of (3.3) and the fact that $\pi_i \neq \pi'_i$. Hence, by Lemma 3.2 there is a periodic point w of f such that $f^i(w) \in \operatorname{Int}(\langle D_{i \bmod n}\rangle_T) \cap V(\langle D_{i \bmod n}\rangle_T)$ for each $i \geq 0$. In particular, $f^i(w) \in \operatorname{Int}(\langle \pi_{i \bmod n}\rangle_T)$ for each $0 \leq i < n$. Therefore, w is associated to $\beta(x)$. By Theorem 1.4, x and w belong to the same f-monotone equivalence class. Since $w \in V(T)$, this is a contradiction with the fact that x is significant. This ends the proof of (a).

Next we prove (b). Write the complement of Λ_f as $\Omega \cup \Gamma \cup C$, where Ω is the set of significant negative periodic points of f whose period is not greater than $M(S, P)$, Γ is the set of non-significant periodic points of f which are not f-monotone equivalent to any element of A, and C is the set of periodic points of f being f-monotone equivalent to some point in A. It is clear that Ω, Γ and C are pairwise disjoint.

The period of each point in Ω is not greater than $M(S, P)$. Hence, Ω is finite by Remark 3.5. This proves (b.1).

From the definition of Γ, Remark 2.6 and Lemma 2.5 it immediately follows that each point in Γ belongs to $V(T) \setminus A$. Moreover, since $f(V(T)) \subset A \cup V(T)$ and A is a periodic orbit, it follows that if $x \in \Gamma$ then $\operatorname{Orb}_f(x) \subset V(T) \setminus A$. This proves (b.2).

Finally we prove (b.3). By Proposition 2.10, $C = C_0 \cup C_1 \cup \ldots C_{k-1}$ for a divisor k of $|A|$, where C_i intersects A, $f(\langle C_i\rangle_T) = \langle C_{i+1 \bmod k}\rangle_T$ and f maps bijectively C_i onto $C_{i+1 \bmod k}$ for $0 \leq i \leq k-1$. Moreover, $\langle C_i\rangle_T \cap \langle C_j\rangle_T = \emptyset$ whenever $i \neq j$. By Remark 2.7, if a point is periodic and belongs to $\langle C_i\rangle_T$ then it belongs to C_i. Therefore, from the definition of the sets Ω, Γ and C it follows that $(\Omega \cup \Gamma) \cap \bigcup_i \langle C_i\rangle_T = \emptyset$. Finally, since each point in C is non-significant, Proposition 2.9 implies that each point in C has period not greater than $8|P| - 4$. This ends the proof of the theorem. $\square$

BIBLIOGRAPHY

[1] R. Adler, A. Konheim, M. McAndrew, Topological entropy, Trans. Am. Math. Soc., 114 (1965), 309–319.

[2] Ll. Alsedà, J. Guaschi, J. Los, F. Mañosas, P. Mumbrú, Canonical representatives for patterns of tree maps, Topology, 36 (1997), 1123–1153.

[3] Ll. Alsedà, F. Gautero, J. Guaschi, J. Los, F. Mañosas, P. Mumbrú, Patterns and minimal dynamics for graph maps, Prepublicacions UAB, 2002.

[4] Ll. Alsedà, J. Llibre, M. Misiurewicz, Periodic orbits of maps of Y, Trans. Amer. Math. Soc., 313 (1989), 475–538.

[5] Ll. Alsedà, J. Llibre, M. Misiurewicz, Combinatorial dynamics and entropy in dimension one, in Advanced Series in Nonlinear Dynamics, vol. 5, World Scientific, second edition, 2002.

[6] Ll. Alsedà, J. Llibre, M. Misiurewicz, C. Simó, Twist periodic orbits and topological entropy for continuous maps of the circle of degree one which have a fixed point, Ergod. Th. & Dynam. Sys., 5 (1985), 501–517.

[7] S. Baldwin, Generalizations of a theorem of Sharkovskii on orbits on continuous real-valued functions, Discrete Math., 67 (1987), 111–127.

[8] S. Baldwin, An extension of Sharkovskii's Theorem to the n-od, Ergod. Th. & Dynam. Sys., 11 (1991), 249–271.

[9] M. Misiurewicz, Z. Nitecki, Combinatorial patterns for maps of the interval, Mem. Amer. Math. Soc., 94 (1991), no. 456.

[10] Y. Takahashi, A formula for the topological entropy of one–dimensional dynamics, Sci. Papers College Gen. Ed. Univ. Tokyo, 30 (1980), 11–22.

Texte reçu le 5 janvier 2005,
révisé le 9 mai 2005,
accepté le 13 juin 2005.

Lluís ALSEDÀ,
Universitat Autònoma de Barcelona
Departament de Matemàtiques, Edifici Cc
08913 Cerdanyola del Vallès, Barcelona (Spain)
alseda@mat.uab.es

David JUHER,
Universitat de Girona
Departament d'Informàtica i Matemàtica Aplicada
Lluís Santaló s/n, 17071 Girona (Spain)
juher@ima.udg.es

Pere MUMBRÚ,
Universitat de Barcelona
Departament de Matemàtica Aplicada i Anàlisi
Gran Via 585, 08071 Barcelona (Spain)
pmumbru@ub.edu