The μ-calculus on trees and Rabin's complementation theorem

André Arnold
LaBRI
Université Bordeaux I

Abstract

In this paper we show that the function associated with any closed or non closed term of the μ-calculus on trees can be represented by a recognizable set of trees whose nodes are labeled by letters and by sets of variables. Rabin's complementation theorem is an immediate consequence of this result.

1 Introduction

There is no need to recall the importance, in logic and computer science, of Rabin's complementation theorem [10]: *for any tree automaton $\mathcal{A}$ there is a tree automaton $\mathcal{A}'$ such that $L(\mathcal{A}')$, the set of infinite trees accepted by $\mathcal{A}'$, is the complement of $L(\mathcal{A})$.*

The Rabin's proof of this theorem is usually considered as a hard one, and several other proofs have been offered [3, 5, 6, 2]. All these proofs obey the same pattern. They consider a class $\mathcal{O}$ of objects such that a set of trees $L(O)$ is associated with any object O in this class, and they show that $\mathcal{O}$ has the two following properties:

> *Equivalence property*: for any set T of trees, there is an automaton $\mathcal{A}$ such that $T = L(\mathcal{A})$ if and only if there is $O \in \mathcal{O}$ such that $T = L(O)$.
>
> *Complementation property*: for any O, there is O' in $\mathcal{O}$ such that $L(O')$ is the complement of $L(O)$.

Obviously, $\mathcal{O}$ is chosen so that the above Complementation property is simpler to prove than the complementation property for automata.

In the original Rabin's proof, $\mathcal{O}$ is the class of automata with the so-called Rabin acceptance condition (while the automaton Rabin considers are defined with the Muller acceptance condition).

For Gurevich and Harrington [3] and for Muchnik [5], an object of $\mathcal{O}$ is a two-players game G on trees, together with one player $p \in \{I, II\}$. The set $L(G, p)$ is the set of trees on which player p has a winning finite-state strategy for the game G. In this case the Equivalence property is

$$\begin{aligned} \forall\mathcal{A},\ \exists G &: L(\mathcal{A}) = L(G, I), \\ \forall G,\ \exists\mathcal{A} &: L(\mathcal{A}) = L(G, II), \end{aligned}$$

and the Complementation property amounts to saying that for any game on any tree, one of the two players has a winning finite-state strategy (see also the survey of Thomas [11]).

For Muller and Schupp [6], $\mathcal{O}$ is the class of alternating automata, which contains the usual automata, so that the Equivalence property is: for any alternating automaton $\mathcal{B}$ there is an automaton $\mathcal{A}$ such that $L(\mathcal{B}) = L(\mathcal{A})$. They do not prove this property but claim it can be easily obtained by Gurevich and Harrinton's method. The Complementation property is obtained as a consequence of the determinacy of certain games.

For Emerson and Jutla [2], $\mathcal{O}$ is the class of all closed terms of the μ-calculus on trees (cf. [9, 4, 7, 1]), for which the Complementation property is easy: every term τ can be syntactically transformed into its dual $\tilde{\tau}$ where $L(\tilde{\tau})$ is the complement of $L(\tau)$. On the other hand, Niwiński has proved [8] the Equivalence property for the μ-calculus without intersection, which applies also to non closed μ-terms. But even if τ does not contain intersection, its dual may contain intersection. Thus, Emerson and Jutla, still using a game-theoretical argument, prove that for every closed μ-term τ there is an alternating automaton $\mathcal{B}$ of a special form such that $L(\tau) = L(\mathcal{B})$ and the Equivalence property for the μ-calculus becomes a consequence of the Equivalence property for these alternating automata, that they prove by the same kind of arguments as those used for the Equivalence property for games.

In this paper we prove directly the Equivalence property for the μ-calculus, without any use of alternating automata or game-theoretical arguments. Indeed we prove a more general Equivalence property which applies also to non closed μ-terms and which generalizes Niwiński's Equivalence property. In [7], Niwiński considers the function $\tau^{\mathcal{T}}[x_1, x_2, \ldots, x_n] : \wp(T)^n \to \wp(\mathcal{T})$ associated with a μ-term τ whose free variables are $x_1, x_2, \ldots, x_n$ (so that when τ is closed, $\tau^{\mathcal{T}}[] = L(\tau)$), and shows that this function can be represented by a recognizable set of incomplete infinite trees, whose leaves are labeled by variables in $\{x_1, x_2, \ldots, x_n\}$.

For instance, the μ-term $\tau_1 = a(x, b(y, x))$ defines the function $\tau_1^{\mathcal{T}}[x, y]$ such that $\tau_1^{\mathcal{T}}[x, y](T_x, T_y) = \{a(t_1, b(t_2, t_3)) \mid t_1 \in T_x, t_2 \in T_y, t_3 \in T_x\}$, the μ-term $\tau_2 = a(y, x)$ defines the function $\tau_2^{\mathcal{T}}[x, y]$ such that $\tau_2^{\mathcal{T}}[x, y](T_x, T_y) = \{a(t_1, t_2) \mid t_1 \in T_y, t_2 \in T_x\}$. Then the function $\tau^{\mathcal{T}}[x, y]$ associated with $\tau = \tau_1 \wedge \tau_2$ is defined by

$$\tau^{\mathcal{T}}[x, y](T_x, T_y) = \{a(t_1, b(t_2, t_3)) \mid t_1 \in T_x, t_1 \in T_y, b(t_2, t_3) \in T_x, t_2 \in T_y, t_3 \in T_x\}.$$

We suggest to represent this function by something like $a(\{x, y\}, \langle b, \{x\}\rangle(y, x))$. In such incomplete trees, leaves are labeled by single variables, as usual, or by sets of variables, but also internal nodes may be labeled by variables, expressing some other constraints that have to be taken into account when substituting trees for variables.

This leads us to consider complete infinite trees whose nodes are labeled by a letter *and* a set of variables, called f-trees, because we use them as a concrete representation of *functions* over trees (such trees have been considered both by Muchnik [5], where variables are called "dead-ends", and by Rabin [10]). We give a proper algebraic meaning to sets of f-trees (called f-sets when they satisfy an additional technical condition) by associating with every f-set F over the variables $\{x_1, x_2, \ldots, x_n\}$, a mapping $F[x_1, x_2, \ldots, x_n] : \wp(\mathcal{T})^n \to \wp(\mathcal{T})$ and by defining an operation of composition of f-sets denoted by $F[x_1 := F_1, \ldots, x_n := F_n]$.

Then we prove our main result:

(1) for any μ-term τ there exists an f-set $\tau^{\mathcal{F}}$ such that the function associated with $\tau^{\mathcal{F}}$ is equal to the function associated with τ, and

(2) the f-set $\tau^{\mathcal{F}}$ is a recognizable set of trees.

This proof is done by structural induction on τ, the only difficult step being to prove the following result:

> if F is a recognizable f-set of f-trees, then $\nu x.F$ and $\mu x.F$, the greatest and least fixed points of the equation $G = F[x := G]$, are recognizable.

We claim that all the difficulty of Rabin's theorem is concentrated in this result.

Of course, all methods used so far to prove Rabin's theorem can be used to prove this single result. We choose in this paper to adopt the more direct one, which turns out to be the Rabin's method:
(1) we prove that $\nu x.F$ and $\mu x.F$ are the projection of some sets $F^{\nu x}$ and $F^{\mu x}$, defined in terms of F (for proving that a set T is recognizable, it is usually a good technique to prove first that it is the projection of some recognizable set T', because one can expect that an automaton for T' is easier to construct that an automaton for T).
(2) we prove that $F^{\nu x}$ and $F^{\mu x}$ are recognizable. Indeed this is almost exactly Lemmas 3.2 and 3.4 of Rabin [10].

Beyond their interest in the present proof of Rabin's complementation theorem, we believe that f-sets should be a useful tool to define recognizable functions over sets of trees in the way suggested by Thomas in [12] or to investigate the problem posed in [1] of whether or not the alternation-depth hierarchy of the μ-calculus with intersection is finite.

This paper is organized as follows. In Section 1 we recall some notions and notations about trees. In section 2 we define the notions of f-trees and f-sets, we define the composition of f-sets and prove some properties of this operation. In Section 3 we define the μ-calculus on trees and show how to associate an f-set with every μ-term. In Section 4 we characterize the f-sets $\nu x.F$ and $\mu x.F$ as projections of sets $F^{\nu x}$ and $F^{\mu x}$. In section 5, we prove the Equivalence property for the μ-calculus. In Section 6 we prove that $F^{\nu x}$ and $F^{\mu x}$ are recognizable.

2 Preliminary notions

2.1 Notations

We denote by $\emptyset$ the empty set as well as the unique function defined on the empty set.

For any set E, $\wp(E)$ is the set of subsets of E. If s and s' are two functions from E into $\wp(F)$, we write $s \subseteq s'$ iff $\forall e \in E, s(e) \subseteq s'(e)$. If $s : E \to \wp(F)$ then $s \cap X$ is the function defined by $(s \cap X)(e) = s(e) \cap X$; similarly, $s - X$ is defined by $(s - X)(e) = s(e) - X$.

For any infinite sequence $u \in E^\omega$ of elements of E, we denote by $u(n)$, for $n \geq 1$, the nth element of this sequence, so that $u = u(1)u(2)\cdots u(n)\cdots$, and by $u[n]$, for $n \geq 0$, the sequence $u(1)u(2)\cdots u(n)$ of length n. In particular, $u[0] = \varepsilon$.

In all what follows we consider a finite alphabet A.

2.2 Trees

Let E be any set. An *infinite full binary tree* (or, simply, a *tree*) over E is a mapping $t : \{l, r\}^* \to E$. We denote by $\mathcal{T}(E)$ the set of all trees over E. If E is the fixed set A, then $\mathcal{T}(A)$ will be written $\mathcal{T}$.

If t is a tree and if $u \in \{l, r\}^*$, then $t|_u$ is the tree defined by $\forall v \in \{l, r\}^*, t|_u(v) = t(uv)$. In particular, $t|_\varepsilon = t$.

If $a \in A$ anf if t_1, t_2 are trees, then $a(t_1, t_2)$ is the tree t defined by $t(\varepsilon) = a, t(lu) = t_1(u), t(ru) = t_2(u)$, so that $t|_l = t_1$ and $t|_r = t_2$. If T_1 and T_2 are two sets of trees then $a(T_1, T_2) = \{a(t_1, t_2) \mid t_i \in T_i, ; i = 1, 2\}$.

3 f-sets

3.1 Definitions

Let X be a finite set of variables. An *f-tree* over X is a pair $\langle t, s\rangle$ where $t \in \mathcal{T}$ is a tree over the alphabet A and $s \in \mathcal{T}(\wp(X))$ is a tree over the subsets of X.

An *f-set* over X is a set F of f-trees over X such that

$$\text{if } \langle t, s\rangle \in F \text{ and } s \subseteq s' \text{ then } \langle t, s'\rangle \in F \qquad \text{(F)}.$$

We denote by $\mathcal{F}(X)$ the set of all f-sets over X. In particular, $\mathcal{F}(\emptyset)$ can be identified with $\wp(\mathcal{T})$.

The set $\mathcal{F}(X)$ ordered by inclusion is a complete lattice whose minimal element is the empty set and the maximal one is $\mathcal{T} \times \mathcal{T}(\wp(X))$.

3.2 Composition of f-sets

3.2.1 Definition

Let $Y, X_1, \ldots, X_n$ be finite sets of variables, let $F, G_1, \ldots, G_n$ be f-sets, respectively over $Y, X_1, \ldots, X_n$, and let $x_1, \ldots, x_n$ be variables. We denote by $F[x_1 := G_1, \ldots, x_n := G_n]$ the f-set F' over $Z = (Y - \{x_1, \ldots, x_n\}) \cup X_1 \cup \cdots \cup X_n$ defined by
$\langle t, s'\rangle \in F'$ iff there exists $\langle t, s\rangle \in F$ such that

$$s - \{x_1, \ldots, x_n\} \subseteq s',$$

$$\forall i = 1, \ldots, n, \forall u \in \{l, r\}^*, x_i \in s(u) \Rightarrow \langle t|_u, s'|_u \cap X_i\rangle \in G_i.$$

The fact that $\langle t, s'\rangle \in F'$ and $s' \subseteq s'' \Rightarrow \langle t, s''\rangle \in F'$ is a straightforward consequence of the definition

In particular, if $F \in \mathcal{F}(\{x_1, \ldots, x_n\})$ and if $G_1, \ldots, G_n$ are in $\mathcal{F}(\emptyset) = \wp(\mathcal{T})$, then $F[x_1 := G_1, \ldots, x_n := G_n]$ is in $\mathcal{F}(\emptyset)$, so that F can be seen as a mapping from $\wp(\mathcal{T})^n$ into $\wp(\mathcal{T})$, when the order of variables is given. To be more precise, if we denote by $\vec{x}$ a sequence $x_1, x_2, \ldots, x_n$ of variables in a given order, containing all the variables of F, then $F[\vec{x}]$ is the mapping from $\wp(\mathcal{T})^n$ into $\wp(\mathcal{T})$ defined by

$$F[\vec{x}](T_1, \ldots, T_n) = F[x_1 := T_1, \ldots, x_n := T_n].$$

3.2.2 Examples

Identity For any variable x, let I_x be the set $\{\langle t, s\rangle \in \mathcal{T} \times \mathcal{T}(\wp(\{x\})) \mid x \in s(\varepsilon)\}$. It is obviously an f-set which has the following property:

Proposition 1 $\forall F \in \mathcal{F}(X), I_x[x := F] = F$.

Proof "$\subseteq$" If $\langle t, s'\rangle$ is in $I_x[x := F]$, there exists $\langle t, s\rangle \in I_x$ such that $s - \{x\} \subseteq s'$ and $x \in s(u) \Rightarrow \langle t|_u, s'|_u \cap X\rangle \in F$. Since $\langle t, s\rangle \in I_x$, $x \in s(\varepsilon)$, thus $\langle t, s' \cap X\rangle \in F$. But $\langle t, s'\rangle$ is an f-tree over $(\{x\} - \{x\}) \cup X = X$, so that $s' \cap X = s'$. Hence, $\langle t, s'\rangle \in F$.

"$\supseteq$" Conversely, let s be defined by

$$\forall u \in \{l, r\}^*, s(u) = \begin{cases} \{x\} & \text{if } u = \varepsilon \\ \emptyset & \text{otherwise} \end{cases}$$

and let $\langle t, s'\rangle \in F$. We have $\langle t, s\rangle \in I_x, s - \{x\} = \emptyset \subseteq s'$, and if $x \in s(u)$ then $u = \varepsilon$ and $\langle t|_u, s'|_u \cap X\rangle = \langle t, s' \cap X\rangle = \langle t, s'\rangle \in F$, so that $\langle t, s'\rangle \in I_x[x := F]$. $\square$

Union and intersection For two given variables x_1 and x_2, we define the f-sets V and Λ in $\mathcal{F}(\{x_1, x_2\})$ by

$$\begin{aligned} \langle t, s\rangle \in V &\Leftrightarrow x_1 \in s(\varepsilon) \text{ or } x_2 \in s(\varepsilon), \\ \langle t, s\rangle \in \Lambda &\Leftrightarrow x_1 \in s(\varepsilon) \text{ and } x_2 \in s(\varepsilon). \end{aligned}$$

It is easy to check the following result:

Proposition 2 *For $G_1 \in \mathcal{F}(X_1), G_2 \in \mathcal{F}(X_2)$, let*

$$G'_1 = \{\langle t, s\rangle \in \mathcal{T} \times \mathcal{T}(\wp(X_1 \cup X_2)) \mid \langle t, s \cap X_1\rangle \in G_1\} \in \mathcal{F}(X_1 \cup X_2)$$

and

$$G'_2 = \{\langle t, s\rangle \in \mathcal{T} \times \mathcal{T}(\wp(X_1 \cup X_2)) \mid \langle t, s \cap X_2\rangle \in G_2\} \in \mathcal{F}(X_1 \cup X_2).$$

Then

$$\begin{aligned} V[x_1 := G_1, x_2 := G_2] &= G'_1 \cup G'_2, \\ \Lambda[x_1 := G_1, x_2 := G_2] &= G'_1 \cap G'_2. \end{aligned}$$

In particular, for $T_1, T_2 \in \wp(\mathcal{T}) = \mathcal{F}(\emptyset)$, $V[x_1 := T_1, x_2 := T_2] = T_1 \cup T_2$ and $\Lambda[x_1 := T_1, x_2 := T_2] = T_1 \cap T_2$.

Rooting For any letter a, let R_a be the element of $\mathcal{F}(\{x_1, x_2\})$ defined by

$$\langle t, s\rangle \in R_a \Leftrightarrow t(\varepsilon) = a, x_1 \in s(l), x_2 \in s(r).$$

Proposition 3 *For $G_1 \in \mathcal{F}(X_1), G_2 \in \mathcal{F}(X_2)$, $R_a[x_1 := G_1, x_2 := G_2] =$*

$$\{\langle t, s\rangle \in \mathcal{T} \times \mathcal{T}(\wp(X_1 \cup X_2)) \mid t(\varepsilon) = a, \langle t|_l, s|_l \cap X_1\rangle \in G_1, \langle t|_r, s|_r \cap X_2\rangle \in G_2\}.$$

In particular, for $T_1, T_2 \in \mathcal{T} = \mathcal{F}(\emptyset)$, $R_a[x_1 := T_1, x_2 := T_2] = a(T_1, T_2)$.

3.3 Some properties of composition

Monotonicity and continuity It follows directly from its definition that composition is monotonic for inclusion with respect to any argument:

Proposition 4 *Let $F \subseteq F' \in \mathcal{F}(Y)$, $G_i \subseteq G'_i \in \mathcal{F}(X_i)$. Then*

$$F[x_1 := G_1, \ldots, x_n := G_n] \subseteq F'[x_1 := G'_1, \ldots, x_n := G'_n].$$

It follows that for any $F \in \mathcal{F}(X)$, the equation $G = F[x := G]$ has a least and a greatest fixed point, both in $\mathcal{F}(X - \{x\})$, denoted respectively by $\mu x.F$ and $\nu x.F$.

Moreover, the composition is both sup- and inf-continuous in its first argument.

Proposition 5 *Let F_i, for $i \in I$, a family of f-sets in $\mathcal{F}(Y)$, and let $G_i \in \mathcal{F}(X_i)$. Then*

$$\begin{aligned}(\bigcup_{i\in I} F_i)[x_1 := G_1, \ldots, x_n := G_n] &= \bigcup_{i\in I} F_i[x_1 := G_1, \ldots, x_n := G_n],\\ (\bigcap_{i\in I} F_i)[x_1 := G_1, \ldots, x_n := G_n] &= \bigcap_{i\in I} F_i[x_1 := G_1, \ldots, x_n := G_n].\end{aligned}$$

Proof The case of the union is straightforward from the definition. Let us consider the case of the intersection. Obviously,

$$(\bigcap_{i\in I} F_i)[x_1 := G_1, \ldots, x_n := G_n] \subseteq \bigcap_{i\in I} F_i[x_1 := G_1, \ldots, x_n := G_n].$$

Conversely, let $\langle t, s'\rangle$ be in $\bigcap_{i\in I} F_i[x_1 := G_1, \ldots, x_n := G_n]$. Then for any $i \in I$, there is $\langle t, s_i\rangle \in F_i$ such that $s_i - \{x_1, \ldots, x_n\} \subseteq s'$ and $x_j \in s_i(u) \Rightarrow \langle t|_u, s_i|_u \cap X_j\rangle \in G_j$. Let us consider $s = \bigcup_{i\in I} s_i$. Since $s_i \subseteq s$, by (F), $\langle t, s\rangle \in F_i$, thus $\langle t, s\rangle \in \bigcap_{i\in I} F_i$. If $x_j \in s(u)$, then $x_j \in s_i(u)$ for some i, thus $\langle t|_u, s_i|_u \cap X_j\rangle \in G_j$. But $s_i \subseteq s \Rightarrow s_i|_u \cap X_j \subseteq s|_u \cap X_j$, and, by (F), $\langle t|_u, s|_u \cap X_j\rangle \in G_j$. Therefore, $\langle t, s'\rangle \in (\bigcap_{i\in I} F_i)[x_1 := G_1, \ldots, x_n := G_n]$. $\square$

Associativity Now we prove that the composition is associative in the following sense:

Proposition 6 *Let $F \in \mathcal{F}(Z)$, let $G_i \in \mathcal{F}(X_i)$ for $i = 1, \ldots, n$, and $H_j \in \mathcal{F}(Y_j)$ for $j = 1, \ldots, k$. Let $x_1, \ldots, x_n, y_1, \ldots, y_k$ be $n + k$ distinct variables.*

Let $G'_i = G_i[y_1 := H_1, \ldots, y_k := H_k]$. Then

$$(F[x_1 := G_1, \ldots, x_n := G_n])[y_1 := H_1, \ldots, y_k := H_k] = F[x_1 := G'_1, \ldots, x_n := G'_n, y_1 := H_1, \ldots, y_k := H_k].$$

Proof To avoid heavy notations, we prove this result for the case $n = k = 1$. The proof for the general case is quite similar.

Let $F \in \mathcal{F}(Z), G \in \mathcal{F}(X), H \in \mathcal{F}(Y)$. Let us consider
$F' = F[x := G] \in \mathcal{F}((Z - \{x\}) \cup X)$,
$K' = F'[y := H] \in \mathcal{F}((Z - \{x, y\}) \cup (X - \{y\}) \cup Y)$,
$G' := G[y := H] \in \mathcal{F}((X - \{y\}) \cup Y)$ and

$K = F[x := G', y := H] \in \mathcal{F}((Z - \{x,y\}) \cup (X - \{y\}) \cup Y)$,
and let us prove that $K = K'$.

We have $\langle t, \sigma\rangle \in K'$
iff $\exists \langle t, s'\rangle \in F'$ such that $s' - \{y\} \subseteq \sigma$ and $\forall u, y \in s'(u) \Rightarrow \langle t|_u, \sigma|_u \cap Y\rangle \in H$
iff $\exists \langle t, s\rangle \in F, s' \in \mathcal{T}(\wp((Z - \{x\}) \cup X))$ such that $s - \{x\} \subseteq s', s' - \{y\} \subseteq \sigma$ and
$\forall u, x \in s(u) \Rightarrow \langle t|_u, s'|_u \cap X\rangle \in G$,
$\forall u, y \in s'(u) \Rightarrow \langle t|_u, \sigma|_u \cap Y\rangle \in H$.

We have $\langle t, \sigma\rangle \in K$
iff $\exists \langle t, r\rangle \in F$ such that $r - \{x,y\} \subseteq \sigma$ and
$\forall u, x \in r(u) \Rightarrow \langle t|_u, \sigma|_u \cap ((X - \{y\}) \cup Y)\rangle \in G'$,
$\forall u, y \in r(u) \Rightarrow \langle t|_u, \sigma|_u \cap Y\rangle \in H$,
iff $\exists \langle t, r\rangle \in F$ such that $r - \{x,y\} \subseteq \sigma$ and
$\forall u, x \in r(u) \Rightarrow \exists \langle t|_u, r_u\rangle \in G$ such that $r_u - \{y\} \subseteq \sigma|_u \cap ((X - \{y\}) \cup Y)$ and
$\forall v, y \in r_u(v) \Rightarrow \langle t|_{uv}, \sigma|_{uv} \cap Y\rangle \in H$,
$\forall u, y \in r(u) \Rightarrow \langle t|_u, \sigma|_u \cap Y\rangle \in H$.

If $\langle t, \sigma\rangle \in K'$, we can prove that $\langle t, \sigma\rangle \in K$ by setting $r = s$ and $r_u = s'|_u \cap X$. We have $r - \{x,y\} = s - \{x,y\} \subseteq s' - \{y\} \subseteq \sigma$, $r_u - \{y\} = s'|_u \cap (X - \{y\}) \subseteq \sigma|_u \cap (X - \{y\})$, and
$x \in r(u) = s(u) \Rightarrow \langle t|_u, r_u\rangle = \langle t|_u, s'|_u \cap X\rangle \in G$,
$y \in r(u) = s(u) \Rightarrow y \in s'(u) \Rightarrow \langle t|_u, \sigma|_u \cap Y\rangle \in H$,
$y \in r_u(v) = s'(uv) \cap X \Rightarrow \langle t|_{uv}, \sigma|_{uv} \cap Y\rangle \in H$.

If $\langle t, \sigma\rangle \in K$, we can prove that $\langle t, \sigma\rangle \in K'$ by setting $s = r$, and s' defined by

$$\begin{array}{rcl} y \in s'(u) & \Leftrightarrow & y \in r(u) \text{ or } \exists v, w \in \{l, r\}^* : u = vw, x \in r(v), y \in r_v(w) \\ s'(u) - \{y\} & = & \sigma(u) \cap ((Z - \{x,y\}) \cup (X - \{y\})). \end{array}$$

Let us remark that $r_u \subseteq s'|_u \cap X$: if $z \in X - \{y\}$, $z \in r_u(v) \Rightarrow z \in \sigma|_u \cap (X - \{y\}) \Rightarrow z \in s'|_u$ and $y \in r_u(v) \Rightarrow y \in s'(uv) = s'|_u(v)$; thus, $r_u \subseteq s'|_u$ and we get the result, since $r_u \cap X = r_u$.

We have $s' - \{y\} \subseteq \sigma$ by definition of s'. We also have $s - \{x\} \subseteq s'$: if $z \in Z - \{x,y\}$ then $z \in s(u) = r(u) \Rightarrow z \in \sigma(u) \Rightarrow z \in s'(u)$; if $y \in s(u) = r(u)$ then $y \in s'(u)$.

If $x \in s(u) = r(u)$ then $\langle t|_u, r_u\rangle \in G$, and $r_u \subseteq s'|_u \cap X$ implies, by (F), $\langle t|_u, s'|_u \cap X\rangle \in G$.

If $y \in s'(u)$ then either $y \in r(u)$ thus $\langle t|_u, \sigma|_u \cap Y\rangle \in H$ or $u = vw$ with $y \in r_v(w)$ thus $\langle t|_{vw}, \sigma|_{vw} \cap Y\rangle \in H$ □

The composition is also compatible with fixed-point operators.

Proposition 7 *Let $F \in \mathcal{F}(Y)$ and $G_i \in \mathcal{F}(X_i)$, for $i = 1, \ldots, n$. Let $x, x_1, \ldots, x_n$ be distinct variables, such that $x \notin X_1 \cup \cdots \cup X_n$. Then $(\mu x.F)[x_1 := G_1, \ldots, x_n := G_n] = \mu x.(F[x_1 := G_1, \ldots, x_n := G_n])$ and $(\nu x.F)[x_1 := G_1, \ldots, x_n := G_n] = \nu x.(F[x_1 := G_1, \ldots, x_n := G_n])$.*

Proof Let us consider the case of the least fixed point.

Let F' be $F[x_1 := G_1, \ldots, x_n := G_n] \in \mathcal{F}(Z)$ with $Z = (Y - \{x_1, \ldots, x_n\}) \cup X_1 \cup \cdots \cup X_n$. Then $\mu x.F' \in \mathcal{F}(Z - \{x\})$ is the union of the sets F'_α, where α is any ordinal and F'_α is defined by: $F'_0 = \emptyset$, $F'_{\alpha+1} = F'[x := F'_\alpha]$, $F'_\beta = \bigcup_{\alpha<\beta} F'_\alpha$, for a limit ordinal β. Similarly, $\mu x.F \in \mathcal{F}(Y - \{x\})$ is the union of the sets F_α, defined by: $F_0 = \emptyset$, $F_{\alpha+1} = F[x := F_\alpha]$, $F_\beta = \bigcup_{\alpha<\beta} F_\alpha$, for a limit ordinal β.

Let us remark that $F'_{\alpha+1} = (F[x_1 := G_1, \ldots, x_n := G_n])[x := F'_\alpha]$, which is equal, by Proposition 6, to $F[x := F'_\alpha, x_1 := G_1[x := F'_\alpha], \ldots, x_n := G_n[x := F'_\alpha]]$. But $x \notin X_i \Rightarrow G_i[x := F'_\alpha] = G_i$. Thus $F'_{\alpha+1} = F[x := F'_\alpha, x_1 := G_1, \ldots, x_n := G_n]$.

Since $F'_0 = \emptyset = F_0[x_1 := G_1, \ldots, x_n := G_n]$, we can prove by induction, using Propositions 6 and 5, that $F_\alpha[x_1 := G_1, \ldots, x_n := G_n] = F'_\alpha$ and the result follows by Proposition 5.

The case of the greatest fixed point is quite similar, replacing union by intersection, and taking $F'_0 = \mathcal{T} \times \mathcal{T}(Z - \{x\})$, $F_0 = \mathcal{T} \times \mathcal{T}(Y - \{x\})$. Thus it is sufficient to prove that $F'_0 = F_0[x_1 := G_1, \ldots, x_n := G_n]$. In one direction,

$$F_0[x_1 := G_1, \ldots, x_n := G_n] \in \mathcal{F}((Y - \{x, x_1, \ldots, x_n\}) \cup X_1 \cup \cdots \cup X_n);$$

and, because of the hypothesis on x,

$$(Y - \{x, x_1, \ldots, x_n\}) \cup X_1 \cup \cdots \cup X_n = ((Y - \{x_1, \ldots, x_n\}) \cup X_1 \cup \cdots \cup X_n) - \{x\} = Z - \{x\}.$$

Thus, $F_0[x_1 := G_1, \ldots, x_n := G_n] \subseteq F'_0$.

Let us prove the converse inclusion. Let $\langle t, s' \rangle \in F'_0$. The f-tree $\langle t, s \rangle$, where s is defined by $s(u) = \emptyset$, for each $u \in \{l, r\}^*$, is in F_0. Then

- $s - \{x_1, \ldots, x_n\} \subseteq s'$,
- $\forall i = 1, \ldots, n, \forall u \in \{l, r\}^*, x_i \in s(u) \rightarrow \langle t|_u, s'|_u \cap X_i \rangle \in G_i$.

and $\langle t, s' \rangle \in F_0[x_1 := G_1, \ldots, x_n := G_n]$, as required. □

4 μ-calculus on trees

4.1 Syntax

Let us consider a set $X = \{x_1, x_2, \ldots\}$ of variables and a finite alphabet A. The set $\mathcal{M}$ of μ -terms τ is defined inductively, together with the sets $FV(\tau)$ of free variables of τ by:

- a variable x is a μ-term and $FV(x) = \{x\}$,
- if a is in A, and $\tau_1, \tau_2 \in \mathcal{M}$, then $\tau = a(\tau_1, \tau_2) \in \mathcal{M}$, and $FV(\tau) = FV(\tau_1) \cup FV(\tau_2)$,
- if τ_1 and τ_2 are in $\mathcal{M}$, then $\tau_1 \vee \tau_2$ and $\tau_1 \wedge \tau_2$ are in $\mathcal{M}$, and $FV(\tau_1 \vee \tau_2) = FV(\tau_1 \wedge \tau_2) = FV(\tau_1) \cup FV(\tau_2)$,
- if $\tau \in \mathcal{M}$, and if x is any variable, then $\tau' = \mu x.\tau$ and $\tau'' = \nu x.\tau$ are in $\mathcal{M}$, and $FV(\tau') = FV(\tau'') = FV(\tau) - \{x\}$.

A μ-term τ is *closed* if $FV(\tau) = \emptyset$.

4.2 Semantics

With each μ-term τ, and with any ordering $\vec{x} = \langle x_1, \dots, x_n \rangle$ of a set X of variables containing all the free variables of τ , we associate a monotonic mapping $\tau^{\mathcal{T}}[\vec{x}] : \wp(\mathcal{T})^n \to \wp(\mathcal{T})$ defined by induction on the construction of τ. Let $\vec{T} = T_1, T_2, \dots, T_n$. Then

- if $\tau = x_i$, then $\tau^{\mathcal{T}}[\vec{x}](\vec{T}) = T_i$,
- if $\tau = a(\tau_1, \tau_2)$, then $\tau^{\mathcal{T}}[\vec{x}](\vec{T}) = a(\tau_1^{\mathcal{T}}[\vec{x}](\vec{T}), \tau_2^{\mathcal{T}}[\vec{x}](\vec{T}))$,
- $(\tau_1 \vee \tau_2)^{\mathcal{T}}[\vec{x}](\vec{T}) = \tau_1^{\mathcal{T}}[\vec{x}](\vec{T}) \cup \tau_2^{\mathcal{T}}[\vec{x}](\vec{T})$, $(\tau_1 \wedge \tau_2)^{\mathcal{T}}[\vec{x}](\vec{T}) = \tau_1^{\mathcal{T}}[\vec{x}](\vec{T}) \cap \tau_2^{\mathcal{T}}[\vec{x}](\vec{T})$,
- $(\mu x.\tau)^{\mathcal{T}}[\vec{x}](\vec{T})$ and $(\nu x.\tau)^{\mathcal{T}}[\vec{x}](\vec{T})$ are, respectively, the least and the greatest solution of the equation

  $T = \tau^{\mathcal{T}}[x; \vec{x}](T, \vec{T})$, if x does not occur in $\vec{x}$,
  $T = \tau^{\mathcal{T}}[\vec{x}](T_1, \dots, T_{i-1}, T, T_{i+1}, \dots, T_n)$, if $x = x_i$ for some $i \in \{1, \dots, n\}$.

4.3 Duality

With each μ-term τ we associate its dual $\tilde{\tau}$ defined by induction on the construction of τ:

- if $\tau = x_i$, then $\tilde{\tau} = x_i$,
- if $\tau = a(\tau_1, \tau_2)$, then $\tilde{\tau} = a(\tilde{\tau}_1, \mathbf{1}) \vee a(\mathbf{1}, \tilde{\tau}_2) \vee \bigvee_{b \neq a} b(\mathbf{1}, \mathbf{1})$, where $\mathbf{1}$ is the closed μ-term $\nu x. \bigvee_{a \in A} a(x, x)$, (so that $\mathbf{1}^{\mathcal{T}} = \mathcal{T}$),
- if $\tau = \tau_1 \vee \tau_2$ then $\tilde{\tau} = \tilde{\tau}_1 \wedge \tilde{\tau}_2$,
- if $\tau = \tau_1 \wedge \tau_2$ then $\tilde{\tau} = \tilde{\tau}_1 \vee \tilde{\tau}_2$,
- $\widetilde{\mu x.\tau} = \nu x.\tilde{\tau}$, $\widetilde{\nu x.\tau} = \mu x.\tilde{\tau}$.

It is clear that $FV(\tau) = FV(\tilde{\tau})$. The term $\tilde{\tau}$ is said to be the dual of τ because of the following result, easily proved by induction on the construction of τ, where $\overline{T}$ is the complement in $\mathcal{T}$ of the set T.

Proposition 8 $\tilde{\tau}^{\mathcal{T}}[\vec{x}](T_1, \dots, T_n) = \overline{\tau^{\mathcal{T}}[\vec{x}](\overline{T_1}, \dots, \overline{T_n})}$. *If τ is closed,* $\tilde{\tau}^{\mathcal{T}} = \overline{\tau^{\mathcal{T}}}$.

4.4 μ-terms and f-sets: another semantics

Now, with any μ-term τ, we associate an f-set $\tau^{\mathcal{F}} \in \mathcal{F}(X)$, where $X = FV(\tau)$, by induction on the construction of τ.

- if $\tau = x$, then $\tau^{\mathcal{F}} = I_x$,
- if $\tau = a(\tau_1, \tau_2)$, then $\tau^{\mathcal{F}} = R_a[x_1 := \tau_1^{\mathcal{F}}, x_2 := \tau_2^{\mathcal{F}}]$,
- $(\tau_1 \vee \tau_2)^{\mathcal{F}} = V[x_1 := \tau_1^{\mathcal{F}}, x_2 := \tau_2^{\mathcal{F}}]$, $(\tau_1 \wedge \tau_2)^{\mathcal{F}} = \Lambda[x_1 := \tau_1^{\mathcal{F}}, x_2 := \tau_2^{\mathcal{F}}]$,
- $(\mu x.\tau)^{\mathcal{F}}$ is the least solution $\mu x.\tau^{\mathcal{F}}$ of the equation $G = \tau^{\mathcal{F}}[x := G]$, and $(\nu x.\tau)^{\mathcal{F}}$ is the greatest solution $\nu x.\tau^{\mathcal{F}}$ of this equation.

It is easy to check that $\tau^{\mathcal{F}} \in \mathcal{F}(FV(\tau))$. This f-set $\tau^{\mathcal{F}}$ is a "concrete" representation of the semantics of τ in the following sense.

Proposition 9 *For any ordering $\vec{x}$ of a set X containing $FV(\tau)$,* $\tau^{\mathcal{T}}[\vec{x}] = \tau^{\mathcal{F}}[\vec{x}]$.

Proof The proof is by induction on the construction of τ.

If $\tau = x_i$ then $\tau^{\mathcal{F}} = I_{x_i}$, and by extending the proof of Proposition 1, $I_{x_i}[\vec{x}](\vec{T}) = T_i = \tau^{\mathcal{T}}[\vec{x}](\vec{T})$.

For $\tau = \tau_1 \vee \tau_2$, or $\tau = \tau_1 \wedge \tau_2$, or $\tau = a(\tau_1, \tau_2)$, this follows from the induction hypothesis and the Propositions 6, 2, and 3.

Using the induction hypothesis, it remains to prove that $(\mu x.\tau^{\mathcal{F}})[\vec{x}](\vec{T})$, (resp. $(\nu x.\tau^{\mathcal{F}})[\vec{x}](\vec{T})$), is the least (resp. the greatest) solution of the equation

$$T = \tau^{\mathcal{F}}[x, \vec{x}](T, \vec{T}), \text{ if } x \text{ does not occur in } \vec{x},$$

$$T = \tau^{\mathcal{F}}[\vec{x}](T_1, \ldots, T_{i-1}, T, T_{i+1}, \ldots, T_n), \text{ if } x = x_i \text{ for some } i \in \{1, \ldots, n\},$$

which is a consequence of Proposition 7. $\square$

In [1], the authors introduced a composition operation on μ-terms: $\tau[x \leftarrow \tau']$ is the μ-term obtained by replacing each free occurrence of x in τ by τ', provided no such occurrence is in a subterm $\mu y.\tau''$ or $\nu y.\tau''$ of τ where y is a free variable of τ' (such a situation can always be reached by renaming the bound variables of τ).

It is easy to check that $(\tau[x \leftarrow \tau'])^{\mathcal{F}} = \tau^{\mathcal{F}}[x := \tau'^{\mathcal{F}}]$.

5 A characterization of fixed points

In this section we give a characterization of the sets $\mu x.F$ and $\nu x.F$ in terms of F. This characterization plays the role of Lemmas 3.10 and 3.11 in [10].

5.1 The greatest fixed points

Let $F \in \mathcal{F}(X)$. Let us define the set $F^{\nu x}$, included in $\mathcal{T} \times \mathcal{T}(\wp(X))$, by $\langle t, s\rangle \in F^{\nu x}$ iff $x \in s(\varepsilon)$ and $\forall u, x \in s(u) \Rightarrow \langle t|_u, s|_u\rangle \in F$. $F^{\nu x}$ is not necessarily an f-set, but it has the following properties, whose proofs are straightforward:

Proposition 10
(i) $F^{\nu x} \subseteq F$.
(ii) If $s \subseteq s'$ *and* $s \cap \{x\} = s' \cap \{x\}$, *then* $\langle t, s\rangle \in F^{\nu x} \Rightarrow \langle t, s'\rangle \in F^{\nu x}$.
(iii) If $\langle t, s\rangle \in F^{\nu x}$ *and* $x \in s(u)$ *then* $\langle t|_u, s|_u\rangle \in F^{\nu x}$.

Proposition 11 $\langle t, s'\rangle \in \nu x.F$ *iff there is* $\langle t, s\rangle \in F^{\nu x}$ *such that* $s' = s - \{x\}$.

Proof Let $G = \{\langle t, s - \{x\}\rangle \mid \langle t, s\rangle \in F^{\nu x}\}$. G is indeed an f-set: if $s - \{x\} \subseteq s' = s' \cap (X - \{x\})$ then $\langle t, s \cup s'\rangle \in F^{\nu x}$ by Proposition 10(ii), and $\langle t, (s \cup s') - \{x\}\rangle = \langle t, s'\rangle \in G$. Let us prove that $G = \nu x.F$.

To prove that $G \subseteq \nu x.F$, we prove that $G \subseteq F[x := G]$. If $\langle t, s'\rangle \in G$ there is $\langle t, s\rangle \in F^{\nu x} \subseteq F$ such that $s' = s - \{x\}$; if $x \in s(u)$, then $\langle t|_u, s|_u\rangle \in F^{\nu x}$ by Proposition 10(iii), and $\langle t|_u, s'|_u\rangle = \langle t|_u, s|_u - \{x\}\rangle \in G$. Thus, $\langle t, s'\rangle \in F[x := G]$.

Let $\langle t, s'\rangle$ be in $\nu x.F$. Let us define $\xi \in \mathcal{T}(\wp(\{x\})$ by $x \in \xi(u)$ iff $\langle t|_u, s'|_u\rangle \in \nu x.F$, and let $s = s' \cup \xi$. We have $s' = s - \{x\}$, so let us prove that $\langle t, s\rangle \in F^{\nu x}$. Obviously, $x \in \xi(\varepsilon)$, thus $x \in s(\varepsilon)$. If $x \in s(u)$, then $\langle t|_u, s'|_u\rangle \in \nu x.F$ and there exists $\langle t|_u, r_u\rangle \in F$ such that $r_u - \{x\} \subseteq s'|_u$ and $x \in r_u(v) \Rightarrow \langle t|_{uv}, s'|_{uv}\rangle \in \nu x.F$. Thus, $x \in r_u(v) \Rightarrow x \in \xi(uv)$; it follows that $r_u \subseteq s|_u$ and $\langle t|_u, s|_u\rangle \in F$ by (F). $\square$

5.2 The least fixed points

As for the greatest fixed point in Proposition 11, we characterize $\mu x.F$ by mean of a set $F^{\mu x}$, included in $\mathcal{T} \times \mathcal{T}(\wp(X))$, whose definition is slightly more complicated than that of $F^{\nu x}$.

Let s be in $\mathcal{T}(\wp(X))$. Let $U = \{u \in \{l,r\}^* \mid x \in s(u)\}$. An ordinal function for s is a mapping O that associates a successor ordinal with each element of U. (Indeed, O depends only on $s \cap \{x\}$).

Let us define, for an ordinal α, the mapping s^α by

$$s^\alpha(u) = \begin{cases} s(u) - \{x\} & \text{if } x \in s(u) \text{ and } O(u) \geq \alpha \\ s(u) & \text{otherwise.} \end{cases}$$

We define $F^{\mu x}_\alpha$ as the set of all $\langle t, s\rangle$ such that

$x \in s(\varepsilon)$,

s has an ordinal function O such that

$O(\varepsilon) \leq \alpha$,

$\forall u \in \{l,r\}^*, u \neq \varepsilon,\ x \in s(u) \Rightarrow O(u) < O(\varepsilon)$,

$\forall u \in \{l,r\}^*,\ x \in s(u) \Rightarrow \langle t|_u, s|_u^{O(u)}\rangle \in F$.

This set has the following properties.

Proposition 12
(i) If $\langle t, s\rangle \in F^{\mu x}_\alpha$ then $\langle t, s^\alpha\rangle \in F$.
(ii) If $s \subseteq s'$ and $s \cap \{x\} = s' \cap \{x\}$, then $\langle t, s\rangle \in F^{\mu x}_\alpha \Rightarrow \langle t, s'\rangle \in F^{\mu x}_\alpha$.
(iii) Let $\langle t, s\rangle \in F^{\mu x}_\alpha$ and let O be its ordinal function. Let u be any element of $\{l,r\}^*$ such that $x \in s(u)$; let r and O' be defined by

$r(\varepsilon) = \{x\} \cup s|_u^{O(u)}(\varepsilon)$,

$r(v) = s|_u^{O(u)}(v)$, *for* $v \neq \varepsilon$;

$O'(\varepsilon) = O(u)$, *and*

$O'(v) = O(uv)$, *for* $v \neq \varepsilon$ *such that* $x \in s(uv)$ *and* $O(uv) < O(u)$.

Then $\langle t|_u, r\rangle \in F^{\mu x}_{O(u)}$, and its ordinal function is O'.

Let $F^{\mu x}$ be the union of all $F^{\mu x}_\alpha$.

Proposition 13 *$\langle t, s'\rangle \in \mu x.F$ iff there is $\langle t, s\rangle \in F^{\mu x}$ such that $s' = s - \{x\}$.*

Proof Let G be the f-set $\{\langle t, s - \{x\}\rangle \mid \langle t, s\rangle \in F^{\mu x}\}$, which is the union of the f-sets $G_\alpha = \{\langle t, s - \{x\}\rangle \mid \langle t, s\rangle \in F^{\mu x}_\alpha\}$, and let $\mu x.F$ be the union of the f-sets F_α defined by $F_0 = \emptyset$, $F_{\alpha+1} = F[x := F_\alpha]$, $F_\beta = \bigcup_{\alpha<\beta} F_\alpha$. We prove by induction that $G_\alpha = F_\alpha$.

Since an ordinal function never has value 0, $F_0^{\mu x} = \emptyset$ and $G_0 = \emptyset = F_0$.

Let us prove that $G_{\alpha+1} = F[x := G_\alpha]$.

Let $\langle t, s'\rangle \in F[x := G_\alpha]$. There exist $\langle t, s\rangle \in F$ such that $s - \{x\} \subseteq s'$ and $x \in s(u) \Rightarrow \langle t|_u, s'|_u\rangle \in G_\alpha$. Applying the definition of G_α we get:
there exists $\langle t, s\rangle \in F$ such that $s - \{x\} \subseteq s'$ and $x \in s(u) \Rightarrow \exists s_u$ such that $s'|_u = s_u - \{x\}$ and $\langle t|_u, s_u\rangle \in F^{\mu x}_\alpha$.

Let $U_\varepsilon = \{u \in \{l,r\}^* \mid x \in s(u)\}$ and $s_\varepsilon = s$. For $u \in U_\varepsilon$, let $U_u = \{v \in \{l,r\}^* \mid x \in s_u(v)\}$ and let O_u be the ordinal function of s_u, that satisfies $O_u(\varepsilon) \leq \alpha$.

Let us define U by $u \in U$ iff $u = \varepsilon$ or $u = vw$ with $v \in U_\varepsilon$ and $w \in U_v$. For $u \in U$, we define

$$O(u) = \begin{cases} \alpha + 1 & \text{if } u = \varepsilon \text{ and } x \notin s(\varepsilon), \\ \inf\{O_v(w) \mid u = vw, v \in U_\varepsilon, w \in U_v\} & \text{otherwise.} \end{cases}$$

Let us check that $\forall u \in U, u \neq \varepsilon \Rightarrow O(u) \leq \alpha$ and $O(u) < O(\varepsilon)$. Let $u = vw \neq \varepsilon$ with $v \in U_\varepsilon$ and $w \in U_v$, such that $O(u) = O_v(w)$. If $v \neq \varepsilon$, then $O_v(w) \leq O_v(\varepsilon) \leq \alpha$; if $v = \varepsilon$ and $u = w \in U_\varepsilon$, then $O(u) \leq O_u(\varepsilon) \leq \alpha$. If $O(\varepsilon) = \alpha + 1$ then the second property is true; if $O(\varepsilon) \neq \alpha + 1$ then $x \in s(\varepsilon)$ and $O(\varepsilon) = O_\varepsilon(\varepsilon) > O_\varepsilon(u) \geq O(u)$.

Let us define $\xi \in \mathcal{T}(\wp(\{x\}))$ by $x \in \xi(u) \Leftrightarrow u \in U$ and let $\sigma = s' \cup \xi$. Since $s' = \sigma - \{x\}$, $\langle t, s'\rangle$ is in $G_{\alpha+1}$ if $\langle t, \sigma\rangle$ is in $F^{\mu x}_{\alpha+1}$. Let us prove this membership. Firstly, $x \in \sigma(\varepsilon)$ and O is an ordinal function for σ such that $O(\varepsilon) \leq \alpha+1$ and $O(v) \leq \alpha$ if $v \neq \varepsilon$. Secondly, let us prove that $\langle t|_u, \sigma|_u^{O(u)}\rangle \in F$ for all $u \in U$. If $O(\varepsilon) = \alpha + 1$, then $\varepsilon \notin U_\varepsilon$, thus $x \in s_\varepsilon(u) \Rightarrow O(u) \leq \alpha \Rightarrow x \in \xi^{O(\varepsilon)}(u)$. It follows that $s_\varepsilon \subseteq \sigma^{O(u)}$ and, since $\langle t, s_\varepsilon\rangle \in F$, then, by (F), $\langle t, \sigma^{O(\varepsilon)}\rangle \in F$. If $O(\varepsilon) = \beta + 1 \neq \alpha + 1$ then $O(u) = O_v(w)$ with $v \in U_\varepsilon$ and $w \in U_v$, and we know that $\langle t|_u, s_v|_w^{O_v(w)}\rangle = \langle t|_u, s_v|_w^{O(u)}\rangle \in F$; if $x \in s_v|_w^{O_v(w)}(w')$ and $O_v(ww') \leq O(u)$ then $x \in \xi(vww')$ and $O(vww') \leq O_v(ww') \leq O(u)$, hence, $x \in \xi|_u^{O(u)}(w')$ and $s_v|_w^{O(u)} \subseteq \sigma|_{vw}^{O(u)}$; thus $\langle t|_u, \sigma|_u^{O(u)}\rangle \in F$ by (F).

Now, let us prove that $G_{\alpha+1} \subseteq F[x := G_\alpha]$. If $\langle t, s'\rangle \in G_{\alpha+1}$ there is $\langle t, s\rangle \in F^{\mu x}_{\alpha+1}$ such that $s' = s - \{x\}$. Then $\langle t, s^{O(\varepsilon)}\rangle \in F$, by Proposition 12(i). If $x \in s^{O(\varepsilon)}(u)$ then $x \in s(u)$ and, by Proposition 12(iii), there is r such that $r - \{x\} = s|_u^{O(u)} - \{x\} = s'|_u$ and $\langle t|_u, r\rangle \in F^{\mu x}_{O(u)}$. Since $x \in s^{O(\varepsilon)}(u)$, $O(u) < O(\varepsilon) \leq \alpha$ and $\langle t|_u, r\rangle \in F^{\mu x}_\alpha$; thus, $\langle t|_u, s'|_u\rangle \in G_\alpha$. It follows that $\langle t, s'\rangle \in F[x := G_\alpha]$.

Finally, let us prove that for any limit ordinal β, $G_\beta = \bigcup_{\alpha<\beta} G_\alpha$, which is equivalent to $F^{\mu x}_\beta = \bigcup_{\alpha<\beta} F^{\mu x}_\alpha$. Since an ordinal function can take only successor ordinals as values, $\langle t, s\rangle \in F^{\mu x}_\beta$ iff $\langle t, s\rangle \in F^{\mu x}_{\alpha+1}$ with $\alpha + 1 < \beta$. □

6 f-sets and recognizable sets

6.1 Recognizable sets of trees

A tree automaton over the alphabet A is a tuple $\mathcal{A} = \langle Q, \delta, Q_I, \mathcal{F}\rangle$ where Q is a set of states, $\delta \subseteq Q \times A \times Q \times Q$ is a set of transitions, $Q_I \subseteq Q$ is a set of initial states and $\mathcal{F}$ is a subset of $\wp(Q)$.

A *run* of $\mathcal{A}$ on a tree $t \in \mathcal{T}(A)$ is a mapping $\rho : \{l,r\}^* \to Q$ such that $\rho(\varepsilon) \in Q_I$ and $\forall u \in \{l,r\}^*, \langle \rho(u), t(u), \rho(ul), \rho(ur)\rangle \in \delta$.

A run ρ is *accepting* if for any branch u (seen as an element of $\{l,r\}^\omega$), the set of states occuring infinitely often in the infinite sequence $\rho(u[0])\rho(u[1])\cdots\rho(u[n])\cdots$ belongs to $\mathcal{F}$.

A tree t is *accepted* by $\mathcal{A}$ if there is an accepting run of $\mathcal{A}$ on t. The set of trees accepted by $\mathcal{A}$ is denoted by $L(\mathcal{A})$. A set T of trees is *recognizable* if there is an automaton A such that $T = L(\mathcal{A})$, and we say that $\mathcal{A}$ *recognizes* T.

In the sequel, we will use a stronger criterion for a run to be accepting: for any branch u, the sequence $\rho(u[0])\rho(u[1])\cdots\rho(u[n])\cdots$ belongs to some recognizable subset K of Q^ω. This does not increase the class of accepted sets of trees (cf. Theorem 3.1 of [10]). Let us also remark that the set of infinite sequences of states such that the set of states occuring infinitely often is in $\mathcal{F}$ is a recognizable language K which has the following property:

(S) If $u \in Q^*$ and $w \in Q^\omega$, then $uw \in K \Leftrightarrow w \in K$.

Without loss of generality, we can consider a more restrictive condition on initial states:

(I) the set Q_I is a singleton $\{q_*\}$ such that q_* never occurs in the lefthand part of a transition (i.e., $\delta \subseteq Q \times A \times (Q - \{q_*\}) \times (Q - \{q_*\})$).
To get that condition, add a new state q_* and the transition $\langle q_*, a, q_1, q_2\rangle$ for any transition $\langle q, a, q_1, q_2\rangle \in \delta$ where q is an initial state.

It is well known (cf. [10]), and easy to prove, that the family $\mathrm{Rec}(A)$ of recognizable subsets of $\mathcal{T}(A)$ is closed under union and intersection.

Let us consider a second alphabet A' and a mapping $\pi : A \to A'$. The *projection* of a tree $t \in \mathcal{T}(A)$ is the tree $\pi(t) \in \mathcal{T}(A')$ defined by $\forall u \in \{l,r\}^*, \pi(t)(u) = \pi(t(u))$, and the projection of a set T is $\pi(T) = \{\pi(t) \mid t \in T\}$. The *inverse projection* of $T' \subseteq \mathcal{T}(A')$ is $\pi^{-1}(T') = \{t \in \mathcal{T}(A) \mid \pi(t) \in T'\}$. The family of all recognizable sets of trees is closed under projection and inverse projection.

6.2 The main result

Because of the natural isomorphism $\mathcal{T}(A) \times \mathcal{T}(\wp(X)) \simeq \mathcal{T}(A \times \wp(X))$, any f-set is a set of trees over the alphabet $A \times \wp(X)$.

Then it is clear that, for two given sets X and Y of variables, the mapping $\langle t, s\rangle \rightsquigarrow \langle t, s - Y\rangle$ (resp. $\langle t, s \cap Y\rangle$) is a projection from $\mathcal{T}(A \times \wp(X))$ to $\mathcal{T}(A \times \wp(X - Y))$ (resp. $\mathcal{T}(A \times \wp(X \cap Y))$.

We can now state the main theorem of this paper.

Theorem 1 *Let τ be a μ-term and let $X = FV(\tau)$ be its set of free variables. The f-set $\tau^{\mathcal{F}} \in \mathcal{F}(X)$ is recognizable as a subset of $\mathcal{T}(A \times \wp(X))$.*

The proof is by induction on the construction of τ. The set I_x is obviously recognizable.

If $\tau = \tau_1 \vee \tau_2$ (resp. $\tau = \tau_1 \wedge \tau_2$) and if $G_1 = \tau_1^{\mathcal{F}}$ and $G_2 = \tau_2^{\mathcal{F}}$, then $\tau^{\mathcal{F}} = U[x_1 := G_1, x_2 := G_2]$ (resp.$\tau^{\mathcal{F}} = V[x_1 := G_1, x_2 := G_2]$), and, by Proposition 2, $\tau^{\mathcal{F}}$ is the union (resp. the intersection) of the inverse projections G_1' and G_2' of G_1 and G_2. Thus, if G_1 and G_2 are recognizable, so is $\tau^{\mathcal{F}}$.

If $\tau = a(\tau_1, \tau_2)$, then, by Proposition 3, $\tau^{\mathcal{F}}$ is the set of all trees $\langle a, Y\rangle(t_1, t_2)$ with $Y \subseteq X$, such that t_i belongs to some inverse projection of $\tau_i^{\mathcal{F}}$. Thus if $\tau_1^{\mathcal{F}}$ and $\tau_2^{\mathcal{F}}$ are recognizable, so is $\tau^{\mathcal{F}}$.

Since $(\mu x.\tau)^{\mathcal{F}} = \mu x.\tau^{\mathcal{F}}$ and $(\nu x.\tau)^{\mathcal{F}} = \nu x.\tau^{\mathcal{F}}$, it remains to prove that if an f-set F is recognizable, so are the f-sets $\mu x.F$ and $\nu x.F$. The Propositions 11 and 13 show that $\nu x.F$ and $\mu x.F$ are the projections of $F^{\nu x}$ and $F^{\mu x}$, thus it remains to show the following result.

Lemma 1 *If F is recognizable, then $F^{\nu x}$ and $F^{\mu x}$ are recognizable.*

This lemma is nothing but Lemmas 3.2 and 3.4 of [10]. However, for completeness of this paper, we prove it in the last section of this paper.

6.3 The complementation theorem

The complementation theorem is now easy. By Niwiński's theorem [8], for any recognizable set K of trees there is a closed μ-term τ, not using the intersection operator, such that $K = \tau^{\mathcal{T}}$. By Proposition 8, $\overline{K} = \tilde{\tau}^{\mathcal{T}}$, by Proposition 9, $\overline{K} = \tilde{\tau}^{\mathcal{F}}$, and by Theorem 1, $\overline{K}$ is recognizable.

For completeness we give a sketch of the proof of the Niwiński's result, in a slightly more general form, borrowing much of the Muchnik's ideas.

Proposition 14 *For any recognizable f-set $F \in \mathcal{F}(X)$, there is a μ-term τ, with $FV(\tau) \subseteq X$, such that $F = \tau^{\mathcal{F}}$.*

Proof Let $\mathcal{A} = \langle Q, \delta, q_*, \mathcal{F}\rangle$ be an automaton, with a single initial state, over the alphabet $A' = A \times \wp(X)$.

We say that $\mathcal{A}$ has the property (F') if for all $q, q', q'' \in Q$, and for all $\langle a, Y\rangle, \langle a, Y'\rangle \in A'$ such that $Y \subseteq Y'$,

$$\langle q, \langle a, Y\rangle, q', q''\rangle \in \delta \Rightarrow \langle q, \langle a, Y'\rangle, q', q''\rangle \in \delta.$$

For $q \in Q$, let $L(q, \mathcal{A})$ be the set recognized by the automaton $\mathcal{A}_q = \langle Q, \delta, q, \mathcal{F}\rangle$. If $\mathcal{A}$ has the property (F'), then, for any $q \in Q$, $L(q, \mathcal{A})$ is an f-set in $\mathcal{F}(X)$.

Moreover, we may assume that a recognizable f-set is always recognized by an automaton having the property (F'), by adding to δ all transitions $\langle q, \langle a, Y'\rangle, q', q''\rangle$ such that $\langle q, \langle a, Y\rangle, q', q''\rangle \in \delta$ and $Y \subseteq Y'$. Thus, from now on, all automata we are considering will have this property.

We say that a state q is a *sink state* of $\mathcal{A}$ if $\{q\} \in \mathcal{F}$ and, for all $\langle a, Y\rangle$,

$$\langle q, \langle a, Y\rangle, q', q''\rangle \in \delta \Rightarrow q = q' = q''.$$

We say that q is *x-dead*, for $x \in X$, if $\langle q, \langle a, Y\rangle, q', q''\rangle \in \delta$ iff $x \in Y$ and $q' = q''$ is a sink state. A state q is *non trivial* if it is neither a sink state nor x-dead, for any x.

We prove by induction on the number of non trivial states of $\mathcal{A}$ that for any q, there is a μ-term τ_q, with $FV(\tau_q) \subseteq X$, such that $\tau_q^{\mathcal{F}} = L(q, \mathcal{A})$.

If $\mathcal{A}$ has no non trivial state, then q is a sink state or a x-dead state.

If q is a sink state, then $L(q, \mathcal{A}) = \mathcal{T} \times \mathcal{T}(X)$. We can take $\tau_q = \mathbf{1}_X$ where $\mathbf{1}_X = \nu z. \bigvee_{a \in A} \bigvee_{Y \subseteq X} (a(z, z) \wedge \bigwedge_{y \in Y} y)$, where z is a variable not in X. One can check that $\mathcal{T} \times \mathcal{T}(X) \subseteq (\bigvee_{a \in A} \bigvee_{Y \subseteq X} (a(z, z) \wedge \bigwedge_{y \in Y} y))^{\mathcal{F}}[x := \mathcal{T} \times \mathcal{T}(X)]$, thus $\mathcal{T} \times \mathcal{T}(X) = \mathbf{1}_X^{\mathcal{F}}$.

If q is an X-dead state, then $L(q,\mathcal{A}) = \{\langle t,s\rangle \in \mathcal{T}\times\mathcal{T}(X) \mid x \in s(\varepsilon)\}$. Thus $\tau_q = \mathbf{1}_X \wedge x$.

Let us assume that $\mathcal{A}$ has $n+1$ non trivial states, and let q be in Q. If q is trivial, then we define τ_q as above. If q is non trivial, then we choose an arbitary numbering $q_0,\ldots,q_n$ of the non trivial states of Q, with $q = q_0$, and for any $i \in \{0,\ldots,n\}$ we define the automaton $\mathcal{B}_i = \langle Q\cup\{\overline{q}\}, \delta_i, q_i, \mathcal{F}_i\rangle$ over the alphabet $A'' = A\times\wp(X\cup\{x\})$, where x is a variable not in X and $\overline{q}$ is a state not in Q, by

$$\mathcal{F}_i = \{P - \{q_{i\oplus 1}\} \mid P\in\mathcal{F}\}\cup\{\{\overline{q}\}\},$$

where $i\oplus 1 = i+1 \bmod (n+1)$, and

$$\begin{array}{rcl}\delta_i &=& \{\langle q,\langle a,Y\rangle,q',q''\rangle \mid \langle a,Y\rangle\in A'', q\neq q_{i\oplus 1}, \langle q,\langle a, Y-\{x\}\rangle,q',q''\rangle\in\delta\}\\ && \cup\{\langle q_{i\oplus 1},\langle a,Y\rangle,\overline{q},\overline{q}\rangle \mid \langle a,Y\rangle\in A'', x\in Y\}\\ && \cup\{\langle \overline{q},\langle a,Y\rangle,\overline{q},\overline{q}\rangle \mid \langle a,Y\rangle\in A''\}.\end{array}$$

By construction, this automaton has the property (F'), so that $L(q_i,\mathcal{B}_i)$ is an f-set in $\mathcal{F}(X\cup\{x\})$. It is useful to notice that $L(q_i,\mathcal{A}) = L(q_i,\mathcal{B}_i)[x := L(q_{i\oplus 1},\mathcal{A})]$.

Clearly, $\overline{q}$ is a sink state, $q_{i\oplus 1}$ is x-dead, thus, $\mathcal{B}_i$ has n non trivial states and we can apply the induction hypothesis: for each $i\in\{0,\ldots,n\}$ there exists a μ-term τ_i such that $\tau_i^{\mathcal{F}} = L(q_i,\mathcal{B}_i)$.

Now, let $\tau = \tau_0[x \leftarrow \tau_1[x \leftarrow \tau_2[\cdots\tau_{n-1}[x\leftarrow\tau_n]\cdots]]]$, ($\tau[x\leftarrow\tau']$ is defined in Section 4.4).

Thus, the μ-term τ_q ($q=q_0$) such that $\tau_q^{\mathcal{F}} = L(q,\mathcal{A})$ is defined by

$$\tau_q = \begin{cases}\nu x.\tau & \text{if } \{q_0,q_1,\ldots,q_n\}\in\mathcal{F},\\ \mu x.\tau & \text{if } \{q_0,q_1,\ldots,q_n\}\notin\mathcal{F}.\end{cases}$$ □

Note that the μ-terms we have constructed contain intersection, while Niwiński's result asserts that intersection is not necessary. The reason is that the term τ associated with F has the property $\tau^{\mathcal{F}} = F$. Indeed, to prove that any recognizable subset of $\mathcal{T}$ is the interpretation of a closed μ-term it is enough to prove the property $\tau^{\mathcal{T}}[\vec{x}] = F[\vec{x}]$, which is weaker than the previous one by Proposition 9. If we want only this weaker property, we can adapt the previous proof to get rid of intersection: just replace the μ-term $\mathbf{1}_X$ by $\mathbf{1}$ and the μ-term $\mathbf{1}_X\wedge x$ by x.

7 Proof of Lemma 1

Let us consider F as a recognizable subset of $\mathcal{T}(A')$ where A' is the alphabet $A\times\wp(X)$ (it makes sense to say that a letter $a'\in A'$ contains x and to use the notation $a'-\{x\}$).

Let $\mathcal{A} = \langle Q,\delta,q_*,K\rangle$ be a tree automaton having the property (I), where K has the property (S), such that $F = L(\mathcal{A})$.

Let S be the set of tuples $\langle d, q_1, i_1, j_1, q_2, i_2, j_2, \ldots, q_k, i_k, j_k\rangle$ such that

$d\in\{l,r\}$,

$1\le k\le |Q|$,

$\forall m\in\{1,\ldots,k\},\ q_m\in Q$,

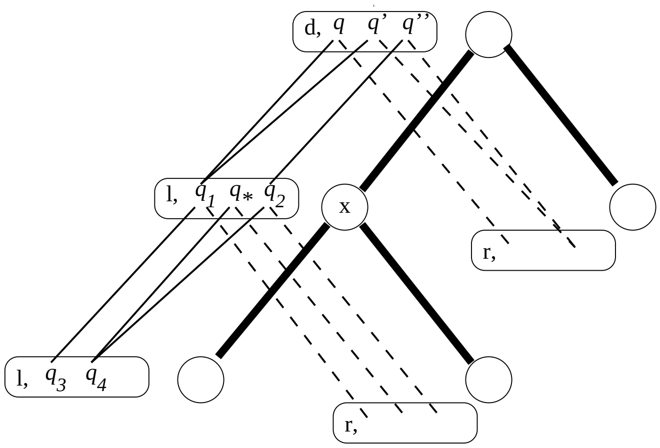

Figure 1: A (part of a) run

$$\forall m \in \{1,\ldots,k\},\ 1 \le i_m \le |Q|, \text{ and } 1 \le j_m \le |Q|,$$

$$\forall m, m' \in \{1,\ldots,k\},\ m \ne m' \Rightarrow q_m \ne q_{m'}.$$

Obviously S is finite.

The automaton $\mathcal{A}_\nu$ Now we can define the automaton $\mathcal{A}_\nu$ recognizing $F^{\nu x}$. Its set of states is S, its initial states are all the elements of S in the form $\langle d, q_*, i, j\rangle$, and its acceptance criterion is a set K_ν to be defined later on. Its transitions are defined by the following conditions: Let $s = \langle d, q_1, i_1, j_1, \ldots, q_k, i_k, j_k\rangle$, $s' = \langle d', q'_1, i'_1, j'_1, \ldots, q'_{k'}, i'_{k'}, j'_{k'}\rangle$, and $s'' = \langle d'', q''_1, i''_1, j''_1, \ldots, q''_{k''}, i''_{k''}, j''_{k''}\rangle$. Let b a letter of A'. Then $\langle s, b, s', s''\rangle$ is a transition of $\mathcal{A}_\nu$ iff

(i) $d' = l$ and $d'' = r$,

(ii) $\forall m : 1 \le m \le k,\ 1 \le i_m \le k',\ 1 \le j_m \le k''$,

(iii.a) $\forall m : 1 \le m' \le k',\ q'_{m'} \ne q_* \Rightarrow \exists m : 1 \le m \le k, i_m = m'$,

(iii.b) $\forall m : 1 \le m'' \le k'',\ q''_{m''} \ne q_* \Rightarrow \exists m : 1 \le m \le k, j_m = m''$,

(iv) $\exists n : 1 \le n \le k, q_n = q_*$ iff b contains x,

(v) $\forall m : 1 \le m \le k, \langle q_m, b, q'_{i_m}, q''_{j_m}\rangle \in \delta$.

It should be noted that the condition (I) on q_* together with (v) implies

(vi) $\forall m : 1 \le m \le k,\ q'_{i_m} \ne q_*, q''_{j_m} \ne q_*$.

Let us consider a run ρ of this automaton on a tree t. A part of such a run is shown in Figure 1.

For $u \in \{l,r\}^*$, let $\rho(u) = \langle d(u), q_1(u), i_1(u), j_1(u), \ldots, q_{k(u)}(u), i_{k(u)}(u), j_{k(u)}(u)\rangle$. Let us remark that because of (i) $d(ul) = l$ and $d(ur) = r$. Let $u \in \{l,r\}^*$ such that $t(u)$ contains x. By (iv), $x \in t(u)$ iff there is n, $1 \leq n \leq k(u)$ such that $q_n(u) = q_*$, and by definition of the initial states, $x \in t(\varepsilon)$. If $x \in t(u)$, we can define a run of $\mathcal{A}$ on $t|_u$ in the following way. For any $w \in \{l,r\}^*$ let $m(w)$ be the integer defined inductively by

$$m(\varepsilon) = n,$$

$$m(wd) = \begin{cases} i_{m(w)} & \text{if } d = l, \\ j_{m(w)} & \text{if } d = r. \end{cases}$$

Then the mapping $w \rightsquigarrow q_{m(w)}(w)$ is a run of $\mathcal{A}$ because of (v). This run accepts $t|_u$ if for any $w \in \{l,r\}^\omega$, the sequence $q_{m(w[0])}(w[0])q_{m(w[1])}(w[1])\ldots$ is in K. This leads us to the following definitions.

A *position* in ρ is a pair $\langle u, m\rangle$ with $u \in \{l,r\}^*$ and $1 \leq m \leq k(u)$.
Let $w \in \{l,r\}^\omega$ be a branch. A *track* of ρ along w is an infinite sequence of positions $\theta = \langle w[p], m_0\rangle\langle w[p+1], m_1\rangle \ldots \langle w[p+n], m_n\rangle \ldots$, for $p \geq 0$, such that
$\forall n \geq 0,\ m_{n+1} = \begin{cases} i_{m_n}(w[p+n]) & \text{if } d(w[p+n+1]) = l, \\ j_{m_n}(w[p+n]) & \text{if } d(w[p+n+1]) = r. \end{cases}$

If $\theta = \langle w[p], m_0\rangle\langle w[p+1], m_1\rangle \ldots, \langle w[p+n], m_n\rangle \ldots$ is a track, its *trace* is the sequence $q_{m_0}(w[p])q_{m_1}(w[p+1]) \ldots q_{m_n}(w[p+n]) \ldots \in Q^\omega$.

Let $\sigma = s_0 s_1 \ldots s_n \ldots$, with

$$s_n = \langle d(n), q_1(n), i_1(n), j_1(n), \ldots, q_{k(n)}(n), i_{k(n)}(n), j_{k(n)}(n)\rangle$$

be an infinite sequence of S^ω. This sequence is said to have property (B) if
$\forall n \geq 0, \forall m : 1 \leq m \leq k(n)$,

$$\begin{aligned} 1 \leq i_m(n) \leq k(n+1) \quad &\text{if} \quad d(n+1) = l, \\ 1 \leq j_m(n) \leq k(n+1) \quad &\text{if} \quad d(n+1) = r. \end{aligned}$$

Let us remark that if ρ is a run of $\mathcal{A}_\nu$ and $w \in \{l,r\}^\omega$, then, because of (ii) any sequence $\rho(w[n])\rho(w[n+1])\ldots$ has property (B), and that the notion of a track can be also defined for a sequence σ having property (B).

Let B be the set of sequences of S^ω having the property B. Let K_ν be the set of sequences in B defined by $\sigma \in K_\nu$ iff for any trace $\chi = q_0 q_1 \ldots$ of a track of σ, if $q_0 = q_*$ then $\chi \in K$.

The following result is almost obvious.

Proposition 15 *The set B is a recognizable subset of S^ω. The set K_ν is recognizable provided K is a recognizable subset of Q^ω.*

Now it is clear that if t is accepted by $\mathcal{A}_\nu$ then every $t|_u$ such that x occurs in $t(u)$ is accepted by $\mathcal{A}$. Hence, $t \in F^{\nu x}$.

Conversely let t be in $F^{\nu x}$. For any $u \in \{l,r\}^*$ such that x occurs in $t(u)$, let ρ_u be an accepting run of $\mathcal{A}$ on $t|_u$ so that we can associate with a node u the set of pairs $s''_u = \{\langle v, \rho_v(w)\rangle \mid u = vw, x \in t(u)\}$. Since x occurs in $t(\varepsilon)$, $s''_\varepsilon = \{\langle \varepsilon, q_*\rangle\}$. We get

s'_u by withdrawing from s''_u all pairs $\langle v, \rho_v(w)\rangle$ such that $v = v_1v_2$ with $v \neq v_1$ and $\rho_v(w) = \rho_{v_1}(v_2w)$. Then we add to each element $\langle v, \rho_v(w)\rangle$ of s'_u two "pointers" i_v and j_v on s'_{ul} and s'_{ur} defined by i_v (resp. j_v) is the unique $\langle v', \rho_{v'}(w')\rangle$ of s'_{ul} (resp. s'_{ur}) such that $\rho_{v'}(w') = \rho_v(wl)$ (resp. $\rho_v(wr)$). It is not difficult to see that we have constructed a run of $\mathcal{A}_\nu$ on t. Moreover, for any branch w, for any n, for any $\langle v, q\rangle \in s'_{u[n]}$, the sequence $v_0 = v; v_{p+1} = d_p$ with $d_p = i_{v_p}$ if $w(n+p+1) = l$, j_{v_p} if $w(n+p+1) = r$, is decreasing for the prefix ordering on $\{l, r\}^*$ thus ultimately stationary. Therefore, since K has property (S), every branch of this run is in K_ν.

The automaton $\mathcal{A}_\mu$ The automaton $\mathcal{A}_\mu$ recognizing $F^{\mu x}$ has the same states and the same initial states as $\mathcal{A}_\nu$. Its transitions are defined by the same conditions as for $\mathcal{A}_\nu$, except for the last two ones which become

(iv.a) $\exists n : 1 \leq n \leq k, q_n = q_*$ iff b contains x,

(iv.b) $\exists n : 1 \leq n \leq k, q_n = q_*$ and $k > 1$ implies $n < k$

(v.a) if b does not contain x then $\forall m : 1 \leq m \leq k, \langle q_m, b, q'_{i_m}, q''_{j_m}\rangle \in \delta$,

(v.b) if b contain x, let n such that $q_n = q_*$, then

$$\forall m : 1 \leq m \leq n, \langle q_m, b - \{x\}, q'_{i_m}, q''_{j_m}\rangle \in \delta,$$
$$\forall m : n < m \leq k, \langle q_m, b, q'_{i_m}, q''_{j_m}\rangle \in \delta.$$

Before defining its acceptance condition, let us explain how to construct a run of this automaton on a tree t' of $F^{\mu x}_{\alpha+1}$. First we write this tree as $\langle t, \xi\rangle$ where $t(u) = t'(u) - \{x\}$ and $\xi(u) = t'(u) \cap \{x\}$. Let O be the ordinal function associated with ξ. For any u such that $x \in \xi(u)$ let ρ_u be a run of $\mathcal{A}$ accepting $\langle t|_u, \xi|_u^{O(u)}\rangle$. We associate with a node u the set of triples $s''_u = \{\langle O(v), v, \rho_v(w)\rangle \mid u = vw, x \in \xi(v)\}$. Since x occurs in $\xi(\varepsilon)$, $s''_\varepsilon = \{\langle O(\varepsilon), \varepsilon, q_*\rangle\}$. Let us partially order the triples $\langle \beta, u, q\rangle$ where β is an ordinal, $u \in \{l, r\}^*$, and q is a state, by $\langle \beta, u, q\rangle \leq \langle \beta', u', q'\rangle$ iff $\beta < \beta'$ or $\beta = \beta'$ and u is a prefix of u'. Let us assume that we are at a node u containing x, such that we are just starting a new run ρ_u of $\mathcal{A}$. The ordering of the triples in s''_u is such that $\langle O(u), u, \rho_u(\varepsilon) = q_*\rangle < \langle \beta', u', q'\rangle$ iff $O(u) < \beta'$ (this is a consequence of the following remark: if $\langle \beta', u', q'\rangle \neq \langle O(u), u, \rho_u(\varepsilon) = q_*\rangle$ is in s''_u, then u' is a strict prefix of u). Thus, all runs ρ_v such that $\langle O(v), v, \rho_v(w)\rangle$ is before $\langle O(u), u, q_*\rangle$ are recognizing $\langle t|_v, \xi|_v^{O(v)}\rangle$ with $O(v) \leq O(u)$ and do not see the x at u.

We get s'_u by withdrawing from s''_u all $\langle O(v), v, \rho_v(w)\rangle$ which are not minimal among the elements $\langle O(v'), v', \rho_{v'}(w')\rangle$ such that $\rho_v(w) = \rho_{v'}(w')$, and we order s'_u by the previous order on the triples. Still, we have $\langle O(u), u, \rho_u(\varepsilon) = q_*\rangle < \langle \beta', u', q'\rangle$ iff $O(u) < \beta'$ and any run ρ_v such that $\langle O(v), v, \rho_v(w)\rangle$ is before $\langle O(u), u, q_*\rangle$ does not see x at u. This is exactly what says the condition (v.b). Moreover, if $u \neq \varepsilon$ then $O(u) < O(\varepsilon)$, thus $\langle O(u), u, \rho_u(\varepsilon)\rangle < \langle O(\varepsilon), \varepsilon, \rho_\varepsilon(u)\rangle$, which is expressed by condition (iv.b).

As for $\mathcal{A}_\nu$ we associate pointers with the elements of s'_u. Any trace χ of a track of this run starting in q_* is in K (because K has property (S), because if $\rho_u(v)$ and $\rho_{u'}(v')$ occur in this order in the trace then $\geq \langle O(u), u, \rho_u(v)\rangle\langle O(u'), u', \rho_{u'}(v')\rangle$, and because this order we have defined on triples is well founded), thus we can take K_ν

as acceptance criterion for $\mathcal{A}_\mu$, but it is not restrictive enough. Along a track, the ordinal associated with each position is nonincreasing, but it decreases when we go from $\langle \beta', u', q'\rangle$ to $\langle O(u), u, \rho_u(\varepsilon) = q_*\rangle$ with $\langle O(u), u, \rho_u(\varepsilon) = q_*\rangle < \langle \beta', u', q'\rangle$. Thus, it must be impossible to see q_* infinitely often, traveling along pointers or going from a triple to a preceding one in the same node (let us remark that a track can contain q_* at most in its first position because of (vi)).

Formally, let ρ be a run $\mathcal{A}_\mu$ on $\langle t, \xi\rangle$, with

$$\rho(u) = \langle d(u), q_1(u), i_1(u), j_1(u), \dots, q_{k(u)}(u), i_{k(u)}(u), j_{k(u)}(u)\rangle.$$

Let $w \in \{l, r\}^\omega$ be a branch. A *walk* in ρ along w is a finite or infinite sequence positions $\varpi = p\langle w[p_0], m_0\rangle\langle w[p_1], m_1\rangle \dots, \langle w[p_n], m_n\rangle \dots$, with $p_0 \geq 0$, such that $\forall n \geq 0$,

$$p_{n+1} = p_n + 1, \text{ or } p_{n+1} = p_n$$

$$m_{n+1} = \begin{cases} m_n - 1 & \text{if } p_{n+1} = p_n \\ i_{m_n}(w[p_n]) & \text{if } p_{n+1} = p_n + 1 \text{ and } d(w[p_{n+1}]) = l, \\ j_{m_n}(w[p_n]) & \text{if } p_{n+1} = p_n + 1 \text{ and } d(w[p_{n+1}]) = r. \end{cases}$$

The trace of this walk is the sequence $q_{m_0}(w[p_0])q_{m_1}(w[p_1]) \dots q_{m_n}(w[p_n]) \dots$

We can also define a walk, and its trace, for any sequence σ having property B, and we define M as the set of sequences in B such that $\sigma \in M$ iff the trace of any infinite walk along σ contains only finitely many occurrences of q_*.

Proposition 16 *The set M is recognizable subset of S^ω.*

Obviously, any branch of the run constructed above is in M. So, we take the recognizable set $K_\mu = K_\nu \cap M$ as acceptance criterion for $\mathcal{A}_\mu$ and we have just shown that any tree of $F^{\mu x}$ is accepted by $\mathcal{A}_\mu$. Let us prove the converse.

Let ρ be an accepting run of $\mathcal{A}_\mu$ on $\langle t, \xi\rangle$. Let P be the set of all positions in ρ and let us define the mapping $\mathcal{P} : \wp(P) \to \wp(P)$ by $\forall p \in P, R \subseteq P,\ p \in \mathcal{P}(R)$ iff there is a finite walk $\varpi = p_0 p_1 \dots p_n$, with $p = p_0, p_n \in R$, such that the trace of $p_1 \dots p_n$ contains at least one occurrence of q_*.

Obviously $\mathcal{P}$ is monotonic and has a greatest fixed point $R_\nu = \bigcap_\alpha R_\alpha$ where $R_0 = P, R_{\alpha+1} = \mathcal{P}(R_\alpha)$, and $R_\beta = \bigcap_{\alpha<\beta} R_\alpha$. It is easy to see that R_ν is empty: if $p_0 \in R_\nu$ there is a walk $p_0 p_1 \dots p_{n_1}$ with $p_{n_1} \in R_\nu$ and the trace of $p_1 \dots p_{n_1}$ contains q_*, there is also a walk $p_{n_1} p_{n_1+1} \dots p_{n_2}$ with $p_{n_2} \in R_\nu$ and the trace of $p_{n_1+1} \dots p_{n_2}$ contains q_*, and so on; thus the trace infinite walk $p_1 \dots p_{n_1} p_{n_1+1} \dots p_{n_2} p_{n_2+1} \dots p_{n_m} p_{n_m+1} \dots$ contains infinitely many occurrences of q_*, and the branch that this walk is along is not in M, a contradiction.

Thus, with each position p in P, we can associate a successor ordinal $\Omega(p)$ defined by $\Omega(p)$ is the least ordinal α such that $p \notin R_\alpha$. This ordinal exists, because R is empty, and is strictly greater than 0, since $p \in R_0$. Indeed, $\Omega(p)$ is a successor ordinal: if it is a limit ordinal β , then $p \notin R_\beta$ and $p \in R_\alpha$ for all $\alpha < \beta$, a contradiction with the definition of R_β. This mapping Ω has the following properties.

Proposition 17
(i) For $p = \langle w, m\rangle, p' = \langle w', m'\rangle \in P$ such that $w = w'l$ and the pointer $i_{m'}$ in state $\rho(w')$ is equal to m, or $w = w'r$ and the pointer $j_{m'}$ in state $\rho('w)$ is equal to m,

$\Omega(p) \leq \Omega(p')$.

(ii) For $p = \langle w, m\rangle, p' = \langle w, m'\rangle \in P$ *with* $m < m'$, $\Omega(p) \leq \Omega(p')$.

(iii) For $p = \langle w, m\rangle, p' = \langle w, m'\rangle \in P$ *with* $m < m'$, *if there is* n, $m \leq n < m'$ *such that the* n*th state in* $\rho(w)$ *is* q_*, *then* $\Omega(p) < \Omega(p')$.

Proof It follows from the definition of $\mathcal{P}$, that $\forall R \subseteq P,\ p \in \mathcal{P}(R) \Rightarrow p' \in \mathcal{P}(R)$ in case (i) and (ii), and $\forall R \subseteq P,\ p \in R \Rightarrow p' \in \mathcal{P}(R)$ in case (iii). □

Now, we define an ordinal function O on $\langle t, \xi\rangle$ in the following way. If $x \in \xi(u)$, then, by (iv.a) and (I), there is a unique n such that the nth state of $\rho(u)$ is q_* and we set $O(u) = \Omega(\langle u, n\rangle)$.

It is easy to see that this ordinal function satisfies $\forall u \in \{l, r\}^*, x \in s(u),\ u \neq \varepsilon \Rightarrow O(u) < O(\varepsilon)$: consider the "inverse walk" obtained by going from $\langle u, n\rangle$ as above to the "last" position at u, which is not $\langle u, n\rangle$ itself by (iv.b), then climbing back to the root of the tree by (iii.a) and (iii.b). Finally, it can be proved, using (v.a) and (v.b), that if $x \in \xi(u)$ the run of $\mathcal{A}$ constructed by the track originating in the position $\langle u, n\rangle$ accepts the tree $\langle t|_u, \xi|_u^{O(u)}\rangle$. Thus $\langle t, \xi\rangle \in F^{\mu x}$.

Acknowledgements I thank B. Courcelle and D. Niwiński for carefully reading this text and for their numerous comments, remarks and suggestions.

References

[1] A. Arnold, D. Niwiński. Fixed point characterization of weak monadic logic definable sets of trees. In M. Nivat, A. Podelski, editors, *Tree Automata and Languages*, pages 159–188. Elsevier, 1992.

[2] E. A. Emerson, C. S. Jutla. Tree automata, μ-calculus and determinacy. In *Proc. 32nd IEEE Symp. on Foundations of Comput. Sci.*, pages 368–377, 1991.

[3] Y. Gurevich, L. Harrington. Trees, automata, and games. In *14th ACM Symp. on Theor. of Comput.*, pages 60–65, 1982.

[4] D. Kozen. Results on the propositional μ-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.

[5] A. A. Muchnik. Games on infinite trees and automata with dead-ends: a new proof for the decidability of the monadic second order theory of two successors. In *Bull. EATCS*, 48:220–267, 1992, (translation of a Russian paper of 1984).

[6] D. E. Muller, P. E. Schupp. Alternating automata on infinite trees. *Theor. Comput. Sci.*, 54:267–276, 1987.

[7] D. Niwiński. On fixed point clones. In L. Kott, editor, *Proc. 13th ICALP*, pages 464–473. Springer–Verlag, Lect. Notes Comput. Sci. 226, 1986.

[8] D. Niwiński. Fixed points vs. infinite generation. In *Proc. 3rd IEEE Symp. on Logic in Comput. Sci.*, pages 402–409, 1988.

[9] V. Pratt A decidable *mu*-calculus. In *Proc. 22nd IEEE Symp. on Foundations of Comput. Sci.*, pages 421–427, 1981.

[10] M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Soc.*, 141:1–35, 1969.

[11] W. Thomas. Automata on infinite objects. In J. Van Leeuwen, editor, *Handbook of Theoretical Computer Science (vol. B)*, pages 133–191. 1990.

[12] W. Thomas. Infinite trees and automaton definable relations over ω-words. *Theor. Comput. Sci.*, 103:143–159, 1992.