Math. Struct. in Comp. Science (2001), *vol.* 11, *pp.* 273–297. Printed in the United Kingdom
© 2001 Cambridge University Press

Mu-depth 3 is more than 2: a game-theoretic proof

GIACOMO LENZI

LaBRI, Université Bordeaux I, 351, Cours de la Libération, F-33405 Talence cedex, France

and

Dip. Mat. Appl. 'U.Dini', Università di Pisa, Via Bonanno 25B, 56100 Pisa, Italia

Email: lenzi@labri.u-bordeaux.fr

Received 10 February 1999; revised 14 March 2000

We consider the *propositional modal mu-calculus*, a logic proposed by Kozen in 1983. In this logic two operators μ and ν are present, which express the least and greatest fixpoints of monotone operators on sets. Bradfield in 1998 proved for any n the existence of a mu-calculus formula that requires n alternations of μ and ν. In this paper we consider the particular case $n = 3$ and we exhibit a new formula requiring 3 alternations. Our proof is independent of the technique of Bradfield, and is based on a new kind of game on infinite trees.

1. Introduction

The subject of this paper is the *propositional modal mu-calculus* (or mu-calculus for short), a logic introduced in Kozen (1983), which nowadays represents a very active research area in both theoretical and practical Computer Science. The 'mu' means that the mu-calculus is a logic capable of expressing least (also called minimal) and greatest (also called maximal) solutions of fixpoint equations $X = f(X)$, where f is a monotone function mapping some complete ordered set into itself; the notation $\mu X.f(X)$ is used for the least fixpoint of the function f, and the notation $\nu X.f(X)$ is used for the greatest fixpoint of f. The presence of the two extremal fixpoint operators makes the mu-calculus very powerful, in spite of its simple syntax; in particular, the mu-calculus is capable of expressing various kinds of correctness properties for finite or infinite state systems, such as safety, liveness, fairness, *etc*., so it is widely applied in the area of Computer Science known as System Verification.

This paper addresses a specific expressiveness problem of the mu-calculus, the *fixpoint alternation depth (f.a.d.) hierarchy problem*. The mu-calculus formulas can be classified according to their fixpoint alternation depth, *i.e*., the number of non-trivial nestings of alternating least and greatest fixpoints. One alternation is sufficient (and necessary) for expressing safety and liveness, and two alternations are sufficient (and necessary) for

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

expressing fairness; but for a long time, no property was known that requires three or more alternations. Bradfield in 1996 proved that such properties exist, and that the f.a.d. hierarchy is infinite (see the final section for some more historical remarks on the problem).

In this paper we reprove a particular case of Bradfield's theorem, in the sense that we exhibit a new explicit formula that lies in the third level of the hierarchy. Our proof is independent of the work of Bradfield and is claimed to be more direct. In particular, we use directly some semantical properties of f.a.d. 2 formulas that allow us to find mu-calculus properties that are not expressible by any f.a.d. 2 formula. More precisely, we will study the complexity of expressing a particular f.a.d. 2 formula (the well-known 'infinitely often P' formula), and in this task we will find it useful to introduce a new kind of game on infinite trees, which is the main technical novelty of this paper.

The rest of the paper is organized as follows. Section 2 recalls the standard syntax and semantics of the mu-calculus, together with some related auxiliary concepts. Section 3 introduces the classical notions of Fischer–Ladner closure of a sentence, and of the unrolling operation on the subformulas of a sentence, and a new particular kind of unrolling of sentences, which we call 'collection'. Section 4 recalls the classical notion of fixpoint alternation depth, again with some related auxiliary material, and Section 5 introduces the related and natural notion of composition alternation depth of a f.a.d. n formula. Section 6 introduces a new kind of complexity class for the mu-calculus, which we call P^+-complexity classes, where P is an atomic proposition. Section 7 presents a seemingly new kind of game on infinite trees, and Section 8 deals with a composition lemma for the mu-calculus. The main result of the paper is given in Section 9, where a mu-calculus formula is presented that is hard for fixpoint alternation depth three, and a proof of its hardness is set up. The final section contains some concluding remarks.

2. The mu-calculus: syntax and semantics

In this section we begin by recalling the definition of the syntax of Kozen's mu-calculus L_μ (Kozen 1983), which is an extension of Modal Logic with a least fixpoint operator μ and a greatest fixpoint operator ν. Modalities are indexed by a set L of labels; so for each $l \in L$ we have the modalities $\langle l \rangle$ and $[l]$.

The mu-calculus is interpreted over transition systems; that is, sets equipped with one binary relation for each label. The formula $[l]\phi$ means that ϕ holds on every l-successor of the current node, and the formula $\langle l \rangle \phi$ means that ϕ holds on some l-successor of the current node. Moreover, the formula $\mu X.\phi(X)$ means that the current node is in the least set L equal to $\phi(L)$, and the formula $\nu X.\phi(X)$ means that the current node is in the greatest set G equal to $\phi(G)$.

We define the syntax of the L_μ formulas ϕ in the following way:

$$\phi ::= P \mid \neg P \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid [l]\phi \mid \langle l \rangle \phi \mid X \mid \mu X.\phi \mid \nu X.\phi, \tag{1}$$

where P ranges over a set At of atomic propositions, l ranges over a set L of labels and X ranges over a set Var of fixpoint variables.

We can regard μ and ν as variable binders, hence we use the obvious notion of free and

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

bound variable occurrence, and we call the formulas without free variable occurrences *closed formulas* or *sentences*.

Note that, in the syntax above, the formulas are already in the so-called *negation normal form*: negations appear only in front of the atomic propositions. In this way the usual mu-calculus condition of syntactical monotonicity (that every variable appears within the scope of an even number of negations) is trivially satisfied by all formulas. We will see that considering formulas in negation normal form is essential for the definition of (fixpoint) alternation depth, which will be very important later.

However, negation may be recovered as a defined connective, at least in the case of *sentences*: given a sentence ϕ, we define $\neg\phi$ to be the result of exchanging P and $\neg P$, $\wedge$ and $\vee$, $[l]$ and $\langle l\rangle$, μ and ν in ϕ.

Another useful notion is the dual ϕ^d of any formula ϕ (not necessarily a sentence), defined as the result of exchanging $\wedge$ and $\vee$, $[l]$ and $\langle l\rangle$, μ and ν in ϕ.

Finally, we say that an occurrence of an atom in a formula is *positive* if it is not preceded by a negation, and *negative* otherwise. Given an atom P, a formula is said to be *positive in* P if all its occurrences of P are positive, and, similarly, a formula is said to be *negative in* P if all its occurrences of P are negative.

The precise meaning of a formula of some logic can be given by specifying a *semantics* of the formulas of the logic. Here we recall the standard Kripke semantics for L_μ, which is taken from Kozen (1983).

We recall that a *transition system* is a structure $T = (U, \{\overset{l}{\rightarrow}\}_{l\in L})$, where U is a set (the universe of the system) and each $\overset{l}{\rightarrow}$ is a binary relation on U.

A valuation on T is a function V from $At \cup Var$ to the powerset $\wp(U)$ of U. A model is a pair $M = (U, V)$; U is also called the universe of M.

Given a model M and a mu-calculus formula ϕ, the *semantics* of ϕ in M is a subset $||\phi||M$ of the universe U of M, defined inductively by:

$$||P||M = V(P); \tag{2}$$

$$||\neg P||M = U - V(P); \tag{3}$$

$$||X||M = V(X); \tag{4}$$

$$||\phi_1 \wedge \phi_2||M = ||\phi_1||M \cap ||\phi_2||M; \tag{5}$$

$$||\phi_1 \vee \phi_2||M = ||\phi_1||M \cup ||\phi_2||M; \tag{6}$$

$$||\langle l\rangle\phi||M = \{s \in U | \; t \in ||\phi||M \text{ for some } t \text{ with } s \overset{l}{\rightarrow} t\}; \tag{7}$$

$$||[l]\phi||M = \{s \in U | \; t \in ||\phi||M \text{ for all } t \text{ with } s \overset{l}{\rightarrow} t\}; \tag{8}$$

$$||\mu X.\phi||M = \bigcap\{E \subseteq U : ||\phi||M[X\backslash E] \subseteq E\}; \tag{9}$$

$$||\nu X.\phi||M = \bigcup\{E \subseteq U : E \subseteq ||\phi||M[X\backslash E]\}, \tag{10}$$

where the model $M[X\backslash E]$ is the one that maps the variable X to the set E and is identical

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

to M elsewhere. Likewise, we will use the notation $M[X_1\backslash E_1,\ldots,X_m\backslash E_m]$ to denote the model that maps each X_i to E_i, and is identical to M elsewhere.

By Tarski's theorem (Tarski 1955), the last two clauses set $||\mu X.\phi||M$ equal to the least fixpoint of the function f mapping $E \subseteq U$ to $||\phi||M[X\backslash E]$, and set $||\nu X.\phi||M$ equal to the greatest fixpoint of f.

We say that two formulas ϕ and ϕ' are *semantically equivalent*, written $\phi \Leftrightarrow \phi'$, if they have the same semantics; that is, for every model M we have $||\phi||M = ||\phi'||M$. Likewise, we say that the formula ϕ *implies* the formula ϕ', written $\phi \Rightarrow \phi'$, when for every model M we have $||\phi||M \subseteq ||\phi'||M$. Finally, we say that a formula ϕ is true at the point t of the model M, written $M, t \models \phi$, when $t \in ||\phi||M$.

The example of a transition system that we will consider here is the n-ary tree T_n, whose universe is the set $\{l_1,\ldots,l_n\}^*$ of all the finite words over the alphabet $l_1,\ldots,l_n$, including the empty word Λ, and whose transition relations are $\xrightarrow{l_1},\ldots,\xrightarrow{l_n}$, defined by: $w \xrightarrow{l_i} w'$ if and only if $w' = wl_i$. Given a word w, we call wl_i the *i-th child* of w.

The points of the n-ary tree, that is the n-ary words, are ordered by the prefix relation, that is $w \leqslant w'$ if w' begins with w, or, putting it another way, there is a word z with $w' = wz$.

Some submodels of the n-ary tree are obtained by looking at the descendants of a given point. Given a point w of T_n, we use $w\{l_1,\ldots,l_n\}^*$ to denote the set of all the n-ary words beginning with w, and $T_n(w)$ to denote the transition system T_n restricted to the set $w\{l_1,\ldots,l_n\}^*$; and given a model M on T_n, we use $M(w)$ to denote the model on $T_n(w)$ such that $||P||M(w) = (||P||M) \cap w\{l_1,\ldots,l_n\}^*$ for every atom P, and, similarly, for the variables.

We use L_n to denote the mu-calculus over n labels $l_1,\ldots,l_n$, interpreted on the n-ary tree. In the following we are interested mostly in the case $n \leqslant 3$, hence we will adopt the notation a, b, c instead of l_1, l_2, l_3. Moreover, note that, on the n-ary tree, for every point w and every i from 1 to n there is exactly one point w' such that $w \xrightarrow{l_i} w'$ (that is, $w' = wl_i$), hence the formulas $\langle l_i\rangle\phi$ and $[l_i]\phi$ are equivalent, so we will use the notation $l_i\phi$ to denote either of the two.

We say that a model M over the n-ary tree satisfies (or is a model of) a formula $\phi \in L_n$ when $\Lambda \in ||\phi||M$. So to each mu-calculus formula ϕ we can associate the set of all the models of ϕ or, in terms of properties, the property of being a model of ϕ. Moreover, given a property (or set) $\mathscr{P}$ of models and a set Φ of formulas, we say that $\mathscr{P}$ is definable in Φ when there is $\phi \in \Phi$ such that $\mathscr{P}$ is the set of all the models of ϕ.

All the definitions and notation above are quite standard. However, we will also use some *ad hoc* notation. In particular, since we will frequently use tuples of formulas, we will take some notational liberty with them: we use overlined letters to denote tuples; we do not surround with parentheses the list of the component of a tuple; we use the set-theoretic notation $\alpha \in \overline{\phi}$ to denote that the formula α is a component of the tuple $\overline{\phi}$; and given a tuple of formulas, say $\overline{\phi} = \phi_1,\ldots,\phi_m$, and given a label l, we use $l\overline{\phi}$ to denote the tuple $l\phi_1,\ldots,l\phi_m$.

Some more *ad hoc* notions that we introduce are related to the Boolean structure of the set of models over a fixed transition system, say the n-ary tree. Let M, M' be two

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

models over the n-ary tree. We call the intersection of M and M' the model N such that $||P||N = ||P||M \cap ||P||M'$ for a y atom P, and similarly for the union. We say that M is included in M', written $M \subseteq M'$, when $||P||M \subseteq ||P||M'$ for every atom P. In a similar vein, given a subset S of the n-ary tree, we say that two models M, M' coincide on S when $(||P||M) \cap S = (||P||M') \cap S$ for any atom P.

Finally, a formula ϕ is said to be monotonic if $M \subseteq M'$ implies $||\phi||M \subseteq ||\phi||M'$.

3. FL-closure and related topics

In this section we first recall two classical notions in mu-calculus: the Fischer–Ladner closure of a sentence (see, for instance, Emerson and Streett (1989)); and the Kozen unrolling map associated to a sentence (introduced in Kozen (1983)). Moreover, we introduce a new operation on sentences, which we call *collection*, and which amounts to a particular kind of partial unrolling.

Given a mu-calculus sentence ψ, we will use $FL(\psi)$ to denote the *Fischer–Ladner closure* of ψ, that is, the smallest set of sentences containing ψ and closed under subsentence, negation and fixpoint unrolling, that is passing from $\sigma X.\phi(X)$ to $\phi(\sigma X.\phi(X))$, where σ is μ or ν.

The following properties of the Fischer–Ladner closure are classical.

Lemma 3.1.

1 The FL-closure of any sentence is a finite set.
2 (Boundary property) Let F be a finite, prefix closed subset of the n-ary tree. Let M, M' be two models on the n-ary tree that coincide on all the elements of F, and assume that they verify the same sentences from $FL(\psi)$ at the same points of the 'boundary' of F, that is the points of F having some child outside F. Then M and M' verify the same sentences from $FL(\psi)$ at the same points of F.

The second part of the lemma (what we call here the boundary property) means that, if we specify the validity of the atoms in a finite subset F of the n-ary tree and the validity of sentences in $FL(\psi)$ in the boundary of F, then the validity of these sentences is determined also in the entire F.

Now we recall from Kozen (1983) the definition of the Kozen unrolling map e_ψ associated to a sentence ψ. Let σ denote either μ or ν. If X is a bound variable of ψ, there is a unique subformula $\sigma X.p(X)$ in which X is the fixpoint variable. We denote this formula by σX. The variable X is called a μ-variable if $\sigma X = \mu X.p(X)$ and a ν-variable if $\sigma X = \nu X.p(X)$.

Define $p \leqslant q$ if q appears as a subformula of p, and $p < q$ if q appears as a proper subformula of p. (I use this notation instead of the reverse one to fit Kozen's notation in Kozen (1983)). For $\psi \leqslant p$, define $V_p = Y_1, \ldots, Y_n, n \geqslant 0$, to be the sequence of all variables Y such that $\sigma Y < p$, taken in the order

$$\sigma Y_1 < \ldots < \sigma Y_n < p. \tag{11}$$

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

For $\overline{X} = X_1, \ldots, X_k$ a subsequence of V_p and $\overline{q} = q_1, \ldots, q_k$ a sequence of formulas, define

$$p[\overline{X}\backslash\overline{q}] = p[X_k\backslash q_k][X_{k-1}\backslash q_{k-1}]\ldots[X_1\backslash q_1], \tag{12}$$

where $p[X\backslash q]$ denotes the formula p with all free occurrences of X replaced by q.

If $V_p = X_1, \ldots, X_n$, let σV_p denote the sequence $\sigma X_1, \ldots, \sigma X_n$.

Define the map e_ψ on subformulas of ψ by

$$e_\psi(p) = p[V_p\backslash\sigma V_p]. \tag{13}$$

We note, as in Kozen (1983), the equality

$$FL(\psi) = \{e_\psi(p)|\psi \leqslant p\}. \tag{14}$$

Finally, we note that, in order to prove our main result, we will need to study the syntactical and semantical properties of the sentences obtained from a given L_n sentence ψ by fixpoint unrolling. To this end, we introduce a special fixpoint unrolling procedure, called *collection* of an L_n sentence ψ at the label $l = l_i$, which consists of unrolling ψ in a suitably canonical way obtaining $\chi(l\overline{\theta})$, where $\chi(\overline{Q})$ is a sentence without l, $\overline{Q}$ is a tuple of atoms and $\overline{\theta}$ is a tuple of sentences, possibly containing l. More precisely, we have the following definition.

Definition 3.1. The collection of a L_n sentence ψ at the label $l = l_i$ is the sentence $Coll_l(\psi)$ obtained by the following procedure:

1 Let $l\overline{\omega}$ be the tuple of all the maximal subformulas of ψ beginning with l, say ordered from left to right.
2 We can write ψ as $\chi(l\overline{\omega})$, where $\chi(\overline{Q})$ is an l-free sentence and $\overline{Q}$ is a tuple of atoms occurring only positively in $\chi(\overline{Q})$ (note that this is not necessarily a composition, because if some component of $\overline{\omega}$ has some free variable X, then X is captured by σX in χ, where σ is μ or ν).
3 Replace in ψ each formula $\omega \in \overline{\omega}$ with its image $e_\psi(\omega)$ in the Kozen unrolling map e_ψ.
4 Eliminate vacuous fixpoints: that is, if a subformula in the resulting sentence has the form $\sigma X.\alpha$, where σ is μ or ν, and α does not contain X free, then replace $\sigma X.\alpha$ with α.

We illustrate the collection with a pair of examples.

Example 3.1. Let $n = 3$ and consider

$$\mu_1 = \mu X.c\mu Y.P \vee aX \vee bY. \tag{15}$$

Then we have

$$Coll_c(\mu_1) = c\mu Y.P \vee a\mu_1 \vee bY. \tag{16}$$

Here we have unrolled the variable X and eliminated the vacuous fixpoint μX.

Example 3.2. Consider the sentence

$$\mu_2 = \mu X.\nu Y.X \vee a(P \vee X \vee Y). \tag{17}$$

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

The collection of μ_2 at the label a is obtained by the following steps:

$$\begin{aligned} \mu_2 &= \mu X.\nu Y.X \vee a(P \vee X \vee Y) && (18)\\ &= \mu X.\nu Y.X \vee a(P \vee \mu_2 \vee Y) && (19)\\ &= \mu X.\nu Y.X \vee a(P \vee \mu_2 \vee \nu Y[X \backslash \mu_2]) && (20)\\ &= \mu X.X \vee a(P \vee \mu_2 \vee \nu Y[X \backslash \mu_2]), && (21) \end{aligned}$$

where the square brackets mean substitution, and we let

$$\nu Y = \nu Y.X \vee a(P \vee X \vee Y). \tag{22}$$

Hence in this case we obtain

$$Coll_a(\mu_2) = \chi(a\theta), \tag{23}$$

where

$$\begin{aligned} \chi &= \mu X.X \vee Q && (24)\\ \theta &= P \vee \mu_2 \vee \nu Y[X \backslash \mu_2]. && (25) \end{aligned}$$

We note that the collection of a sentence ψ at any label l is always semantically equivalent to ψ, because it is obtained from ψ by a series of fixpoint unrollings.

4. Fixpoint alternation depth

In this section we recall a classical notion of mu-calculus, that is, the *(fixpoint) alternation depth*. Intuitively, the fixpoint alternation depth of a mu-calculus formula measures the maximum number of nested, alternated fixpoints μ's and ν's appearing in the formula. Actually, in the literature we find several ways of formalizing this idea, see, for instance, Kaivola (1997) for a discussion, but here we stick to the most standard (and most powerful) definition of f.a.d., which was introduced by Niwiński in 1986, and which has now been adopted by the majority of mu-calculus researchers. The idea that makes the Niwiński's definition so powerful is the stability of Niwiński's sets under arbitrary composition of formulas.

The notion of f.a.d. can be formalized by defining a double sequence Σ_n and Π_n of sets of L_μ formulas as follows.

Definition 4.1. (The fixpoint alternation depth hierarchy)

1 $\Sigma_0 = \Pi_0$ is the set of formulas without fixpoints.
2 For all $n \in \mathbf{N}$, Σ_{n+1} is the least set of formulas such that:
   (a) $\Sigma_n \cup \Pi_n \subseteq \Sigma_{n+1}$.
   (b) If $\phi \in \Sigma_{n+1}$, then $\mu X.\phi \in \Sigma_{n+1}$.
   (c) If $\phi_1, \phi_2 \in \Sigma_{n+1}$, then $\phi_1 \wedge \phi_2, \phi_1 \vee \phi_2, \langle l \rangle \phi_1, [l]\phi_2 \in \Sigma_{n+1}$.
   (d) If $\phi, \phi_1, \ldots, \phi_m \in \Sigma_{n+1}$, $P_1, \ldots P_m$ are propositional letters occurring only positively in ϕ, and ϕ_i is free for P_i in ϕ for $1 \leqslant i \leqslant m$, then $\phi[P_1 \backslash \phi_1, \ldots, P_m \backslash \phi_m]$ is in Σ_{n+1}.
3 Π_{n+1} is defined dually to 2 by exchanging ν and μ.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Note that the condition above that ϕ_i is free for P_i in ϕ is necessary in order to have a semantically correct notion of formula composition, and, in particular, to prevent the well-known phenomenon of 'variable capturing'.

We note that the union of all the classes Σ_n and Π_n contains all the mu-calculus formulas. We use *fixpoint alternation depth* of the formula ϕ to mean the smallest n such that $\phi \in \Sigma_{n+1} \cap \Pi_{n+1}$.

To give some intuition of the f.a.d. classes, we note that Σ_1 is exactly the set of all the formulas whose fixpoints are all of type μ, and dually Π_1 is the set of all the formulas whose fixpoints are all of type ν. However, already the class Π_2 is quite large and complicated, due to the combined effect of the composition and of the ν operator. In fact Π_2 contains not only relatively simple formulas such as $\nu X.\mu Y.P \wedge aX \wedge bY$, but some more intricate formulas obtained by composition and ν. For example, inspired by Niwiński (1997), we note that the following formula is Π_2:

$$\phi_0 = \nu W.\mu X.\langle l\rangle(P \wedge X) \vee \nu Y.\mu Z.(\langle l\rangle(P \wedge (W \vee Z))) \vee \langle l\rangle(P \wedge Y). \tag{26}$$

To see this, note that the following formulas are Π_2:

$$\phi_1 = \mu X.\langle l\rangle(P \wedge X) \vee V \tag{27}$$

$$\phi_2 = \nu Y.\mu Z.(\langle l\rangle(P \wedge (W \vee Z))) \vee \langle l\rangle(P \wedge Y). \tag{28}$$

Hence, by composition, the following formula is also Π_2:

$$\phi_3 = \mu X.\langle l\rangle(P \wedge X) \vee \nu Y.\mu Z.(\langle l\rangle(P \wedge (W \vee Z))) \vee \langle l\rangle(P \wedge Y), \tag{29}$$

which after applying νW gives ϕ_0.

For each of the complexity classes introduced above we can define a notion of hardness: we stipulate that Σ_0-hard coincides with the set of all formulas, and for $n \geqslant 1$ we use Σ_n-hard to mean the complement of Π_n; and we operate similarly for Π_n-hard. Finally, we use f.a.d. $n+1$-hard to mean the complement of f.a.d. n.

The definition above implies that Σ_{n+1} is the closure of $\Sigma_n \cup \Pi_n$ under conjunction, disjunction, modalities, least fixpoints and composition. Actually, we note that conjunction, disjunction and modalities can be seen as particular cases of composition (for instance, $\phi_1 \wedge \phi_2$ is the composition of $P \wedge Q$ with ϕ_1 and ϕ_2, where P and Q are fresh atoms); so we can say that Σ_{n+1} is the closure of $\Sigma_n \cup \Pi_n$ under composition and least fixpoint; and, dually, Π_{n+1} is the closure of $\Sigma_n \cup \Pi_n$ under composition and greatest fixpoint.

Let us make some remarks on the f.a.d. classes. First, they are closed under subformulas and they are 'compositional' in the sense of the following lemma.

Lemma 4.1.

1 Let ϕ be a Σ_n formula: then every subformula of ϕ is Σ_n as well, and the same holds for Π_n.
2 (Compositionality) Let P be an atom and let $\alpha(P)$ be a formula. Let β, β' be formulas such that for any h, if $\beta \in \Sigma_h$, then also $\beta' \in \Sigma_h$, and that the same holds for Π_h. If $\alpha(\beta) \in \Sigma_k$ for some k, then $\alpha(\beta') \in \Sigma_k$ as well; and the same holds for $C_n\Pi_k$.

Proof. The first point can be proved by induction on n and, for $n > 0$, on the number of

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

composition and fixpoint operators necessary to obtain ϕ starting from some $\Sigma_{n-1} \cup \Pi_{n-1}$ formulas.

The second point can be proved by induction on α. □

Moreover, we have a neat characterization of the formulas $\Sigma_{n+1} \cap \Pi_{n+1}$, which is given by the last point of the following lemma.

Lemma 4.2.

1 Let ϕ be a sentence obtained by composing a formula α with a finite tuple $\beta_1, \ldots, \beta_k$ of formulas. Then all of $\alpha, \beta_1, \ldots, \beta_k$ are sentences.
2 A mu-calculus sentence is in $\Sigma_{n+1} \cap \Pi_{n+1}$ if and only if it is the composition of finitely many sentences in $\Sigma_n \cup \Pi_n$.
3 A mu-calculus formula is in $\Sigma_{n+1} \cap \Pi_{n+1}$ if and only if it is the composition of finitely many formulas in $\Sigma_n \cup \Pi_n$.

Proof. For the first point, α must be a sentence, otherwise the free variables of α would be free also in ϕ; and each β_i must also be a sentence, otherwise there would be variable capturing.

For the second point, let ϕ be a mu-calculus sentence belonging to $\Sigma_{n+1} \cap \Pi_{n+1}$. We prove that ϕ is a composition of sentences in $\Sigma_n \cup \Pi_n$ by induction on the length of ϕ. In fact, we can suppose that ϕ begins with a fixpoint operator, say μ (the case of ν is dual). Since $\phi \in \Pi_{n+1}$, we have ϕ is in the closure of $\Sigma_n \cup \Pi_n$ under composition and ν. So, since ϕ begins with μ, either $\phi \in \Sigma_n \cup \Pi_n$ and we are done, or ϕ is obtained by composing a formula α beginning with μ with a finite tuple $\beta_1, \ldots, \beta_k$ of formulas, which can be supposed to begin with fixpoints. Moreover, by the first point of the lemma, all of $\alpha, \beta_1, \ldots, \beta_k$ must be a sentence and they all belong to $\Sigma_{n+1} \cap \Pi_{n+1}$ (because $\beta_1, \ldots, \beta_k$ are subformulas of ϕ and α is obtained from ϕ by replacing some subformulas with fresh atoms, hence we can apply subformula closure and compositionality) and they are shorter than ϕ. Hence the inductive hypothesis applies, and $\alpha, \beta_1, \ldots, \beta_k$ are compositions of sentences in $\Sigma_n \cup \Pi_n$, hence ϕ is such a composition also. This concludes the proof of the first point.

The third point is a corollary of the proof of the second, obtained by dropping the assumption that the formulas in question are closed. □

Corollary 4.1. (Characterization of f.a.d.)

1 A mu-calculus sentence has f.a.d. at most n if and only if it is the composition of finitely many sentences in $\Sigma_n \cup \Pi_n$.
2 A mu-calculus formula has f.a.d. at most n if and only if it is the composition of finitely many formulas in $\Sigma_n \cup \Pi_n$.

We have said that to any set of formulas we can associate the set of properties definable in it. Applying this to the classes f.a.d. n gives a hierarchy of sets of properties. The problem of deciding whether the fixpoint alternation depth hierarchy on properties collapses has been open for more than a decade since the invention of mu-calculus. A complete solution has been found by Bradfield in 1996, as follows.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Theorem 4.1. (Bradfield) The fixpoint alternation depth hierarchy is infinite.

For more comments on the hierarchy problem, see the final section.

In the following we will often meet compositions of formulas or sentences, so it is useful to recall the classical semantical characterization of the composition operation, as given in the following lemma.

Lemma 4.3. (The semantics of composition) Let $\phi, \phi_1, \ldots, \phi_m$ be formulas, let $P_1, \ldots P_m$ be propositional letters occurring only positively in ϕ, such that ϕ_i is free for P_i in ϕ for $1 \leqslant i \leqslant m$. Then a model M satisfies the composition $\phi[P_1 \backslash \phi_1, \ldots, P_m \backslash \phi_m]$ if and only if there are sets $E_1, \ldots, E_m$ of points of M such that $M[P_1 \backslash E_1, \ldots, P_m \backslash E_m]$ satisfies $\phi(P_1, \ldots, P_m)$ and each element of E_i satisfies ϕ_i in M, for $1 \leqslant i \leqslant m$.

We conclude this section by introducing the *ad hoc* notion of block.

We use *block* to mean any sentence without any proper subsentence beginning with μ or ν.

Note that every mu-calculus sentence decomposes into a composition of finitely many blocks, which are unique up to renaming the atoms used for composition. So we can use 'block of ϕ' to mean any block occurring in any such decomposition of ϕ.

If β is a block of a sentence ϕ, then we can write ϕ as a double composition $\alpha(\beta(\overline{\gamma}))$, where α is a sentence and $\overline{\gamma}$ is a tuple of sentences. We say that the block β *surrounds* a subformula θ of ϕ if θ is a subformula of some component of $\overline{\gamma}$.

By compositionality, every Σ_n sentence is a composition of finitely many Σ_n blocks, and, similarly, for Π_n; and by the characterization of f.a.d., a block has f.a.d. at most n if and only if it is Σ_n or Π_n.

5. Composition alternation depth

In the previous section we defined the notion of fixpoint alternation depth, and we saw that every f.a.d. n formula is the composition of a finite number of formulas belonging to Σ_n or Π_n. Hence we can refine the fixpoint hierarchy by classifying the f.a.d. n formulas according to their 'composition alternation depth', that is the maximum number of nested alternating Σ_n and Π_n subformulas. So we introduce the 'composition classes' $C_n\Sigma_k$ and $C_n\Pi_k$, which are defined as follows.

Definition 5.1. (The C_n hierarchy)

1 $C_n\Sigma_1 = \Sigma_n$ and $C_n\Pi_1 = \Pi_n$.
2 For $k \geqslant 1$, $C_n\Sigma_{k+1}$ is the set of formulas obtained by composing a Σ_n formula with a tuple of formulas belonging to $C_n\Pi_k$.
3 Dually, $C_n\Pi_{k+1}$ is the set of formulas obtained by composing a Π_n formula with a tuple of formulas belonging to $C_n\Sigma_k$.

We note that, for any n, f.a.d. n is the union over k of $C_n\Sigma_k$ (or $C_n\Pi_k$). Hence the C_n-alternation depth may be defined as usual, that is, the C_n-a.d. of an f.a.d. n formula ψ is the smallest k such that $\psi \in C_n\Sigma_{k+1} \cap C_n\Pi_{k+1}$.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Also, in the case of the composition classes $C_n\Sigma_k$ and $C_n\Pi_k$ we have the notion of hardness: we define $C_n\Sigma_0$-hard to coincide with the set of all formulas, and for $k \geqslant 1$ we call $C_n\Sigma_k$-hard the complement of $C_n\Pi_k$; and we operate similarly for $C_n\Pi_k$-hard.

We note that later we will be interested in studying the composition complexity of *sentences* rather than of arbitrary formulas. Now, we have seen that if a sentence splits as a composition of $k+1$ formulas $\alpha(\beta_1, \ldots, \beta_k)$, then all these formulas are sentences as well. So, in order to evaluate the composition complexity of a sentence, it is sufficient to look at its subsentences rather than at all of its subformulas. For instance, we have the following sufficient condition for a sentence to be $C_n\Sigma_k$- hard or $C_n\Pi_k$-hard.

Lemma 5.1. Let ϕ be a sentence and let n, k be natural numbers. If ϕ has a Σ_n-hard block surrounding a $C_n\Pi_k$-hard sentence, then ϕ is $C_n\Sigma_{k+1}$-hard; and the dual condition holds for $C_n\Pi_{k+1}$- hardness.

Let us make some more remarks on the C_n classes. Like f.a.d. classes, they are closed under subformulas and they are 'compositional' in the sense of the following lemma.

Lemma 5.2. (Compositionality) Let P be an atom, let n be a natural number and let $\alpha(P)$ be a formula. Let β, β' be formulas such that for any h, if $\beta \in C_n\Sigma_h$, then also $\beta' \in C_n\Sigma_h$, and that the same holds for $C_n\Pi_h$. If $\alpha(\beta) \in C_n\Sigma_k$ for some k, then $\alpha(\beta') \in C_n\Sigma_k$ as well; and the same holds for $C_n\Pi_k$.

Proof. The proof is by induction on α. □

In the rest of this section we develop some theory about the behaviour of the composition complexity when rewriting a sentence by fixpoint unrollings; in particular, we will show that the composition complexity does not increase upon unrolling. More precisely, we have the following series of lemmas.

Lemma 5.3. If ψ is a $C_n\Sigma_k$ sentence beginning with a fixpoint and ψ' is obtained by unrolling the outermost variable of ψ, then ψ' is also $C_n\Sigma_k$. The same holds for $C_n\Pi_k$.

Proof. Assume that $\psi \in C_n\Sigma_k$ (the case $\psi \in C_n\Pi_k$ is dual). Then ψ has the form $\sigma X.\alpha(X, \beta_1, \ldots, \beta_h)$, where each β_i is a $C_n\Sigma_k$ sentence and $\sigma X.\alpha(X, Q_1, \ldots, Q_h)$ is Σ_n for any tuple of fresh atoms $Q_1, \ldots, Q_h$. Then, taking another fresh atom P, we have that $\alpha(P, \beta_1, \ldots, \beta_h)$ is also Σ_n. Now ψ' has the form $\alpha(\psi, \beta_1, \ldots, \beta_h)$, namely it is a composition of a Σ_n sentence with a tuple of $C_n\Sigma_k$ sentences, and hence it is also in $C_n\Sigma_k$. □

Lemma 5.4. If ψ is a $C_n\Sigma_k$ sentence beginning with a fixpoint and ψ' is obtained by replacing any variable Y of the outermost block of ψ with $e_\psi(Y)$, then ψ' is also $C_n\Sigma_k$. The same holds for $C_n\Pi_k$.

Proof. ψ' is obtained from ψ by iterating some finite number of times the operation of unrolling the outermost variable. □

Lemma 5.5. If ψ is a $C_n\Sigma_k$ sentence and ψ' is obtained by replacing any variable Y of ψ with $e_\psi(Y)$, then ψ' is also $C_n\Sigma_k$. The same holds for $C_n\Pi_k$.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Proof. Let β be the smallest subsentence of ψ containing Y. Let β' be obtained from β by replacing Y with $e_\psi(Y)$. Then ψ is a composition of the form $\alpha(\beta)$, and ψ' has the form $\alpha(\beta')$. Now, by a previous lemma, if $\beta \in C_n\Sigma_h$, then $\beta' \in C_n\Sigma_h$ as well, and the same holds for $C_n\Pi_h$. Hence, if $\psi = \alpha(\beta) \in C_n\Sigma_k$, then by compositionality, $\psi' = \alpha(\beta') \in C_n\Sigma_k$ as well, and the same holds for $C_n\Pi_k$. □

Lemma 5.6. If ψ is a $C_n\Sigma_k$ sentence containing a (possibly non-closed) subformula ω and ψ' is obtained by replacing ω with $e_\psi(\omega)$, then ψ' is also $C_n\Sigma_k$. The same holds for $C_n\Pi_k$.

Proof. $e_\psi(\omega)$ is obtained by replacing every free variable Y of ω with $e_\psi(Y)$, and hence the previous lemma applies. □

Lemma 5.7. If ψ is a $C_n\Sigma_k$ sentence and $\overline{\omega}$ is a set of (possibly non-closed) subformulas of ψ, and ψ' is obtained by replacing each $\omega \in \overline{\omega}$ with $e_\psi(\omega)$, then ψ' is also $C_n\Sigma_k$. The same holds for $C_n\Pi_k$.

Proof. The proof follows by iterating the previous lemma. □

Finally, the previous lemma implies the following corollary.

Corollary 5.1. If ψ is a $C_n\Sigma_k$ sentence and $\psi' = Coll_l(\psi)$ for some label l, then ψ' is $C_n\Sigma_k$ as well. The same holds for $C_n\Pi_k$.

6. P^+-complexity

In the following sections we will have to evaluate the complexity of nesting mu-calculus formulas. A useful tool to do this is the introduce the notion of P^+-complexity, where P is an atom.

Intuitively, the P^+-complexity of a formula is the complexity of the formula when all the subsentences negative in P are discarded. Formally, the definition is as follows.

Definition 6.1. Let P be an atom. A formula is said to be $\Sigma_n(P^+)$ if it is the composition of a Σ_n formula with finitely many sentences negative in P, and similarly for $\Pi_n(P^+)$.

We have the following characterization of $\Sigma_n(P^+)$ and $\Pi_n(P^+)$.

Lemma 6.1. A sentence is $\Sigma_n(P^+)$ if and only if every block surrounding some positive occurrence of P is Σ_n; and similarly for $\Pi_n(P^+)$.

Now we can define hardness classes in the usual way: we define $\Pi_0(P^+)$-hard to be the set of all formulas, and for $n \geqslant 1$ we define $\Pi_n(P^+)$-hard to be the complement of $\Sigma_n(P^+)$; and we operate similarly for $\Sigma_n(P^+)$-hard.

From the above we get a characterization of $\Sigma_n(P^+)$-hard and $\Pi_n(P^+)$-hard, as follows.

Lemma 6.2. A sentence is $\Sigma_n(P^+)$-hard if and only if it contains a Σ_n-hard block surrounding some positive occurrence of P; and similarly for $\Pi_n(P^+)$-hard.

See the final section for some comments on these definitions.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

7. Some games on infinite trees

In this section we introduce a new kind of game on infinite trees. However, before defining the games let us explain the motivations behind this definition.

As we will see, the main result of this paper is the f.a.d. 3-hardness of the property defined by the following formula:

$$\phi = \nu X.P \wedge I(cI^d(cX)), \tag{30}$$

where P is an atom, and for any fresh atom Q, $I(Q)$ is the formula saying that, on the binary tree $\{a,b\}^*$, there is a path with an infinity of Q's, $I^d(Q)$ is the dual of $I(Q)$, and c is a fresh label.

The intuition behind the proof is that ϕ satisfies the fixpoint equation

$$\phi \Leftrightarrow P \wedge I(cI^d(c\phi)), \tag{31}$$

and hence, for any number n, ϕ is equivalent to a nested composition of n formulas alternatingly $\Pi_2(Q^+)$-hard (that is, $P \wedge I(Q)$) and $\Sigma_2(Q^+)$-hard (that is, $I^d(Q)$), so the C_2-complexity of ϕ is infinite and ϕ cannot be expressed in f.a.d. 2.

In order to fill in the details of the proof according to the idea above, the first step is to prove that the property $P \wedge I(Q)$, or equivalently $I(Q)$, is $\Pi_2(Q^+)$-hard. That property $I(Q)$ is Π_2-hard is a long known result dating back at least to Emerson and Clarke (1980), but in the literature there seems to be no specific tool for studying hardness with positivity constraints such as $\Pi_2(Q^+)$-hardness. For this reason, in this section we introduce a seemingly new expressiveness technique based on games, and, in particular, a two player game G on infinite trees and its minor variant G^-, which will allow us to prove the $\Pi_2(Q^+)$-hardness of property $I(Q)$.

The game $G = G(n,\psi)$ depends on a number $n \in \mathbf{N}$ and on a sentence $\psi \in L_n$. The trees we consider are always given by the complete n-ary tree $T_n = \{l_1,\ldots,l_n\}^*$ equipped with some mu-calculus valuation. Later we will often consider n fixed and will omit it from the notation.

In the game $G(\psi)$, the idea is that player *Two* chooses a sequence M_k of models of ψ, and *One* answers with finite subsets F_{k+1} of T_n, with the constraint that M_{k+1} must preserve the validity in M_k of each sentence in $FL(\psi)$ at each point of F_{k+1}. At the end, *One* wins if the limit of the M_k's verifies ψ, and *Two* wins otherwise.

More formally, we have the following definition.

Definition 7.1. (The game G) Let n be a positive integer and let ψ be a sentence in L_n. The game $G(n,\psi)$ is played by two players *One* and *Two*, and is defined by the following playing rules.

If ψ is unsatisfiable, then we stipulate that *One* wins.

Otherwise, *Two* picks a model M_0 of ψ, and then *One* and *Two* play alternately a sequence of finite sets $F_1, F_2, \ldots, F_k, \ldots$ and a sequence of models $M_1, M_2, \ldots, M_k \ldots$ of ψ according to the following conditions:

1 F_k is a finite, non-empty, prefix closed subset of T_n.
2 $F_k \subseteq F_{k+1}$ and $\bigcup_k F_k = T_n$.
3 $M_k, x \models s$ implies $M_{k+1}, x \models s$ for any $s \in FL(\psi)$ and $x \in F_{k+1}$.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

In this way the sequence of models M_k converges (in the usual topology of the sets of models) to a limit model M_∞: namely, for all $x \in T_n$ and any letters P, $M_\infty, x \models P$ iff $M_k, x \models P$ for all k except a finite number.

The winning rule is: *One* wins if $M_\infty \models \psi$, and *Two* wins if $M_\infty \models \neg\psi$.

In order to grasp some intuition of the game G, consider the following example (due to a referee).

Take $n = 1$ and $\psi = \nu X.\mu Y.(P \vee aY) \wedge aX$, namely ψ says that there is an infinity of P's. Suppose that *Two* adopts the following strategy: M_0 verifies P everywhere, and then M_{k+1} verifies P exactly on the positions in F_{k+1} that verify P in M_k, and on the positions outside F_{k+1} that are multiple of 2^k. If *One* plays badly, for example taking as F_k the first k nodes, he loses; but if he plays well, for instance taking F_k to be the first 2^{k+1} nodes, he wins.

We consider also a minor variant G^- of G, which looks easier for *One* to win.

Definition 7.2. (The game G^-) Let P be an atom, let n be a positive integer and let ψ be a sentence in L_n whose only atom is P. The game $G^-(n, P, \psi)$ has the same rules as G, with the additional constraint that the models M_k are decreasing.

Here we are mainly interested in those formulas ψ such that $G(\psi)$ or $G^-(\psi)$ are determined, namely one of player *One* or *Two* has a winning strategy, which works no matter how the other plays. Actually, the games do have a tendency to be determined, and one could even conjecture that $G(\psi)$ and $G^-(\psi)$ are determined for any mu-calculus sentence ψ, but we do not need such a strong result here. All we need is that *One* wins in some cases and *Two* wins in other cases. First let us consider some instances of G or G^- won by player *One*. As a first example we have the following.

Lemma 7.1. If ψ is negative in P and contains no atom but P, then player *One* has a winning strategy for $G^-(\psi)$.

Proof. The strategy is trivial, because in G^- the valuations of P in the M_k's are decreasing (for instance the set of all the words with length at most k will work as F_k). □

As a less trivial example, it turns out that *One* has a winning strategy for G in the previous infinity example, and the strategy consists in taking F_{k+1} large enough to contain at least k occurrences of P's in the model M_k, which is always possible because every M_k contains an infinity of P's.

More generally, we have the following determinacy result for Π_2 sentences.

Lemma 7.2. Let ψ be a Π_2 sentence. Then player *One* has a winning strategy for $G(\psi)$.

Proof. Let us assume first that ψ is a prenex Π_2 block, namely $\psi = \nu X.\mu Y.\chi(X, Y)$ where χ is a Π_0 formula. The idea is that ψ holds in a model M iff there is a set $X \subseteq T_n$ with $\Lambda \in X$ and satisfying $X \subseteq \mu Y.\chi(X, Y)$ in M. This yields a winning strategy for *One* for $G(\psi)$ which can be described as follows.

If ψ is unsatisfiable, then *One* wins and needs no strategy. Otherwise, assume that *One*

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

has played $F_1, \ldots, F_k$ and *Two* has chosen $M_0, M_1, \ldots, M_k$ for some k. Let us define what the strategy prescribes for the move $k+1$ of *One*, that is the choice of F_{k+1}.

Let X_k the set of the points verifying ψ in M_k. By definition of ψ we have

$$X_k \subseteq \mu Y.\chi(X_k, Y) \text{ in } M_k, \tag{32}$$

and, in particular,

$$X_k \cap F_k \subseteq \mu Y.\chi(X_k, Y) \text{ in } M_k, \tag{33}$$

where in the initial case $k = 0$ we let $F_0 = \varnothing$.

Since $X_k \cap F_k$ is finite, and since Σ_1 formulas have finite closure ordinals in finitely branching transition systems, there is a finite, prefix closed set $F \subseteq T_n$ including F_k containing all words of length at most $k+1$ and such that $X_k \cap F_k \subseteq \mu Y.\chi(X_k \cap F, Y)$ in every model coinciding with M_k in F. The strategy consists in letting F_{k+1} be one such set F.

The strategy is winning. In fact, let

$$X_\infty = \bigcup_k X_k \cap F_k : \tag{34}$$

first, by the last game rule we have

$$X_k \cap F_{k+1} \subseteq X_{k+1}. \tag{35}$$

Hence, on the one hand we have

$$X_\infty = \bigcup_k X_k \cap F_k = \bigcup_k X_k \cap F_{k+1}, \tag{36}$$

and on the other hand we have the inclusions

$$X_k \cap F_k \subseteq \mu Y.\chi(X_k \cap F_{k+1}, Y) \subseteq \tag{37}$$

$$\subseteq \mu Y.\chi(X_\infty, Y) \tag{38}$$

in every model coinciding with M_k in F_{k+1}, and hence also in M_∞.

Now we take the union over k:

$$X_\infty \subseteq \mu Y.\chi(X_\infty, Y) \text{ in } M_\infty. \tag{39}$$

Hence, by definition of greatest fixpoint,

$$X_\infty \subseteq \nu X.\mu Y.\chi(X, Y) = \psi \text{ in } M_\infty. \tag{40}$$

Now the root Λ of T_n is in $X_0 \cap F_1$, and hence is in X_∞ and verifies ψ in M_∞, so *One* wins.

Now let us assume that ψ is any, possibly non-prenex Π_2 sentence. Then the proof is the same as in the prenex case, modulo using vectorial notation. In fact we can use the following theorem, see Arnold (1988): for any Π_2 formula ψ there are h Σ_1 formulas $\mu_1(X_1, \ldots, X_h), \ldots, \mu_h(X_1, \ldots, X_h)$ such that ψ holds in an n-ary model M at the root Λ iff there are h subsets of the n-ary tree $E_1, \ldots, E_h$ such that

1 $\Lambda \in E_1$;
2 $E_i \subseteq \mu_i(E_1, \ldots, E_h)$ in M $(i = 1, \ldots, h)$.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

So, as in the prenex case, player *One* can construct a sequence

$$\overline{F}^k = (F_1^k, \ldots, F_h^k) \tag{41}$$

of h-tuples of finite sets such that

1 $\Lambda \in F_1^k$.
2 $F_i^k \subseteq \mu_i(\overline{F}^{k+1})$ in M_k $(i = 1, \ldots, h)$.
3 $F_i^k \subseteq F_i^{k+1}$ $(i = 1, \ldots, h)$.

In this way $\overline{F}^k$ converges to a solution $E_{1,\infty}, \ldots E_{h,\infty}$ of the system above in the limit model M_∞. Hence M_∞ verifies ψ. □

The next lemma shows that further winning examples may be obtained by composition.

Lemma 7.3. Let n be a positive integer. Let P, Q be different atoms. Let ψ_1 be an L_n sentence positive in Q whose only atoms are P and Q; and let ψ_2 be an L_n sentence whose only atom is P.

If *One* wins $G(\psi_1)$ and $G^-(\psi_2)$, then *One* wins $G^-(\psi_1[Q \backslash \psi_2])$ as well.

Proof. Suppose *One* has both a winning strategy W_1 for $G(\psi_1)$ and a winning strategy W_2 for $G^-(\psi_2)$. Then *One* can obtain a winning strategy W for $G^-(\psi_1(\psi_2))$ as follows.

If $\psi_1(\psi_2)$ has no models, then *One* wins by definition and hence needs no strategy. So suppose $\psi_1(\psi_2)$ has a model and let *Two* pick one, that is M_0.

For the induction step, suppose that *One* has chosen the finite sets $F_1, \ldots, F_k$, and that *Two* has chosen the models $M_0 \supseteq M_1 \supseteq \ldots \supseteq M_k$ of $\psi_1(\psi_2)$. By the semantics of composition, player *One* can choose a model M_k^1 of ψ_1 coinciding with M_k except possibly on Q, and such that $M_k(w)$ verifies ψ_2 for every Q-point w of M_k^1.

Let Q_k the set of the Q-points of M_k^1 lying in F_k. The idea is to preserve the modelhood of $\psi_1(Q)$ at the root, and the modelhood of ψ_2 at every point of Q_k. So *One* chooses the set F_{k+1} to be the prefix closure of the union of the set F_{k+1}^1, answered by the strategy W_1 to the sequence of moves $M_0^1, M_1^1, \ldots, M_k^1$, with all the sets $wF_{i,k+1,w}$, where $i = 0, 1, \ldots, k$, $w \in Q_i$ and $F_{i,k+1,w}$ is the answer of the strategy W_2 to the sequence of moves $M_i(w), M_{i+1}(w), \ldots, M_k(w)$. Note that $M_i(w), M_{i+1}(w), \ldots, M_k(w)$ are all models of ψ_2, and they are decreasing as required by G^-. Note also that F_{k+1} is a finite set.

The strategy W so obtained is winning. In fact, let $Q_\infty = \bigcup_k Q_k$, let M_∞^1 the model M_∞ where a letter Q is added to every point of Q_∞, and let $M_\infty^2(w) = M_\infty(w)$ for each $w \in Q_\infty$: then M_∞^1 verifies ψ_1 since W_1 is winning, and $M_\infty^2(w)$ verifies ψ_2 since W_2 is winning. So M_∞ verifies $\psi_1(\psi_2)$. □

More generally we have the following corollary.

Corollary 7.1. Let h, n be positive integers and let $P, Q_1, \ldots, Q_h$ be different atoms. Let ψ be a sentence of L_n positive in Q_1 and … and Q_h, and whose only atoms are $P, Q_1, \ldots, Q_h$; and let $\chi_1, \ldots, \chi_h$ be sentences of L_n whose only atom is P.

If *One* wins $G(\psi)$ and $G^-(\chi_1)$ and … and $G^-(\chi_h)$, then *One* wins $G^-(\psi(\chi_1, \ldots, \chi_h))$ as well.

Proof. The proof is the same as above but using vectorial notation. □

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

By putting the above results together, we get the following corollary.

Corollary 7.2. Let P be an atom and let ψ be a sentence whose only atom is P, and is obtained by composing a Π_2 sentence with finitely many sentences negative in P. Then *One* wins $G^-(\psi)$.

We can also restate the corollary as follows.

Corollary 7.3. If P is an atom and ψ is a $\Pi_2(P^+)$ sentence containing no atom but P, then *One* wins $G^-(\psi)$.

So far we have examined cases where *One* wins. Let us consider a case where *Two* wins. Define the 'infinitely' formula

$$I(P) = \nu X.\mu Y.(P \wedge (aX \vee bX)) \vee aY \vee bY, \tag{42}$$

and the 'finitely' formula

$$J(P) = I^d(P) = \mu X.\nu Y.(P \vee (aX \wedge bX)) \wedge aY \wedge bY. \tag{43}$$

More precisely, formula I says that there is a path with an infinity of P; and, dually, formula J says that, for all paths, P holds on all points of the path except for a finite number of them.

For player *Two* we obtain the following lemma.

Lemma 7.4. If $\psi(P)$ is a sentence equivalent to $J(P)$ and containing no atom except P, then Player *Two* has a winning strategy for $G^-(\psi)$.

Proof. For this proof we will use P-models to mean the models M on the binary tree T_2 such that $||P'||M = \varnothing$ for any atom P' different from P.

The key to the proof is the definition of a suitable relation R between P-models. Precisely, given two P-models M, M', we say that MRM' holds iff M is strictly included in M' and $M(x)$ is strictly included in M' for infinitely many points x verifying $\neg P$ in M.

We note that the relation R is backwards directed on the P-models of J: that is, given a finite set $\mathscr{F}$ of P-models of J, there is a P-model N of J such that NRM for any $M \in \mathscr{F}$.

To prove this property of R, first let N_1 be the intersection of all the P-models belonging to $\mathscr{F}$; note that N_1 verifies J and is included in every element of $\mathscr{F}$. Now let N_2 be obtained by changing the valuation of N_1 at some point from P to $\neg P$. Then N_2 still verifies J and is strictly included in every element of $\mathscr{F}$.

Now let a^*b be the set of all the binary words of the form a^nb. Modify N_2 as follows: for any $w \in a^*b$, remove every w from the valuation of P and replace $N_2(w)$ with $N_2(w) \cap N_2$. Call the result N_3. Finally, for every point $w \in a^*b$, change the valuation of N_3 at some point above w from P to $\neg P$ and call the result N. Then N is in the relation R with every element of $\mathscr{F}$, as desired.

Now let us come to the proof of the lemma. Suppose that $\psi(P)$ is a sentence containing only the atom P and equivalent to $J(P)$. We describe a winning strategy for *Two* in $G^-(\psi)$.

Given a subset C of $FL(\psi)$, choose a P-model of $J \wedge \bigwedge C$ that is minimal with respect to R, if there is one, and call it $M(C)$. By the backward directedness of R on the P-models

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

of J, we can take a P-model $\tilde{M}$ of J with $\tilde{M}RM(C)$ for any C such that $M(C)$ exists (if $M(C)$ does not exist for any C, take as $\tilde{M}$ an arbitrary P-model of J).

Now *Two* plays a P-model M_0 such that $M_0R\tilde{M}$. Let *One* play the finite set $F_1 \subseteq T_2$. Now *Two* chooses a point $x_0 \in T_2$ such that:

1 The length of x_0 is greater than the length of every element of F_1 (hence $x_0 \notin F_1$ and $x_0 > \Lambda$).
2 $M_0(x_0)$ verifies $\neg P$.
3 $M_0(x_0)$ is strictly included in $\tilde{M}$.

Let C_0 be the set of sentences in $FL(\psi)$ that are true in $M_0(x_0)$. Then $M(C_0)$ does not exist, otherwise we would have $M_0(x_0) \subset \tilde{M}$ and $\tilde{M}RM(C_0)$, hence $M_0(x_0)RM(C_0)$ contrary to the minimality of $M(C_0)$. So there is no minimal P-model of $J \wedge \bigwedge C_0$, and, in particular, there is a P-model N_0 such that $N_0 \models J \wedge \bigwedge C_0$ and $N_0RM_0(x_0)$. Note that $N_0 \models \neg P$.

Let M_1 be the P-model obtained by M_0 by replacing $M_0(x_0)$ with N_0. So M_1 verifies at x_0 the same sentences of $FL(\psi)$ as M_0, and coincides with M_0 outside $T_2(x_0)$. In particular, M_1 verifies the same sentences as M_0 at every word of the same length as x_0, and, by the boundary property of $FL(\psi)$, M_1 verifies the same sentences as M_0 at every word of F_1, and hence M_1 fulfils the rules of the game G^-. So we can let *Two* play M_1. Note that $M_1(x_0) = N_0 \models J \wedge \neg P$.

At this point the reasoning can be moved from the entire tree T_2 to the subtree $T_2(x_0)$, and it can be iterated, with $M_1(x_0)$ in the place of M_0; and after a move F_2 of *One*, we obtain a point $x_1 > x_0$ and a P-model M_2 such that *Two* can play M_2 and $M_2(x_1) \models J \wedge \neg P$, and so on. At the limit we obtain an infinite chain of points $x_0 < x_1 < x_2 < \ldots$ such that each x_i verifies $\neg P$ in M_{i+1}, and hence the limit model verifies $I(\neg P)$, which is equivalent to $\neg J(P)$. □

Putting the last lemma and the last corollary together, we have the following corollary.

Corollary 7.4. Let Q be an atom. If a sentence $\psi \in L_2$ contains only the atom Q and is equivalent to $J(Q)$, then ψ is $\Sigma_2(Q^+)$-hard.

Corollary 7.5. The property defined by $J(Q)$ is $\Sigma_2(Q^+)$-hard.

Proof. If $\psi(Q, R)$ is a formula equivalent to $J(Q)$, where R is a tuple of variables or atoms different from Q, then $\psi(Q, \mathit{false})$ is a sentence equivalent to $J(Q)$ as well and its only atom is Q. By the previous corollary, $\psi(Q, \mathit{false})$ is $\Sigma_2(Q^+)$-hard, and hence $\psi(Q, R)$ is $\Sigma_2(Q^+)$-hard also. □

By dualizing we also have the following corollary.

Corollary 7.6. The property defined by $I(Q)$ is $\Pi_2(Q^+)$-hard.

In particular, we have that the properties defined by I and J are f.a.d. 2-hard, but this is well-known and follows by Emerson and Clarke (1980).

As a last, slight improvement we have the following corollary.

Corollary 7.7. The property defined by $P \wedge I(Q)$ is $\Pi_2(Q^+)$-hard.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Proof. Suppose for an absurdity that $\psi(P,Q)$ is equivalent to $P \wedge I(Q)$ and is $\Sigma_2(Q^+)$. Then $\psi(\neg P,Q)$ is equivalent to $(\neg P) \wedge I(Q)$ and is $\Sigma_2(Q^+)$ as well.

So the disjunction $\psi(P,Q) \vee \psi(\neg P,Q)$ is equivalent to $I(Q)$ and is $\Sigma_2(Q^+)$, contrary to the $\Pi_2(Q^+)$-hardness of the property defined by $I(Q)$. □

8. The composition lemma

In this section we continue the preparation required for the main result. Recall that in the previous section we showed that certain properties are $\Pi_2(Q^+)$-hard or $\Sigma_2(Q^+)$-hard. This means that expressing the properties in question requires a Π_2 or Σ_2 block, respectively, surrounding a positive occurrence of Q. In this section we take the next step, that is we find properties that require a nested alternating sequence of Σ_2 and Π_2 blocks of any fixed length. The idea for finding such properties is essentially to compose $\Sigma_2(Q^+)$-hard and $\Pi_2(Q^+)$-hard formulas. Our tool to realize this idea is a technical lemma, which we call the *composition lemma*.

Intuitively, the composition lemma says that in the mu-calculus, under certain hypotheses, expressing a composition requires nesting of what is required by the components.

Lemma 8.1. (Composition lemma) Let n be a positive integer. Let $\phi_1(Q)$ be an L_{n-1} sentence positive in Q, satisfiable, and such that for every model M of $\phi_1(Q)$ we have $||Q||M \neq \varnothing$ (which we call the non-emptyness property of ϕ_1). Let ϕ_2 be a sentence of L_n such that both ϕ_2 and $\neg\phi_2$ are satisfiable. Then for any natural number k:

1 If the property $\phi_1(Q)$ is $\Sigma_2(Q^+)$-hard and the property ϕ_2 is $C_2\Pi_k$-hard, then the property $\phi_1(l_n\phi_2)$ is $C_2\Sigma_{k+1}$-hard.
2 If the property $\phi_1(Q)$ is $\Sigma_2(Q^+)$-hard and the property ϕ_2 is $C_2\Sigma_k$-hard, then the property $\phi_1(l_n\phi_2)$ is $C_2\Pi_{k+1}$-hard.

See the final section for more comments on the composition lemma.

Proof. It is enough to prove the first part of the lemma because the second is dual. For this proof let $l = l_n$ and let v range over the set of words $\{l_1,\ldots,l_{n-1}\}^*$.

Let ψ be a sentence equivalent to $\phi_1(l\phi_2)$. Collect ψ at l and use $\chi(l\overline{\theta})$), where $\overline{\theta} = \theta_1,\ldots,\theta_h$, to denote the result.

First we are going to split ϕ_2 and every element of $\overline{\theta}$ into finitely many disjoint pieces. To do this, let

$$\overline{V} = V_1,\ldots,V_m \tag{44}$$

be the sequence of all the maximal non-trivial conjunctions of elements of $\phi_2,\theta_1,\ldots,\theta_h$ and negated elements of $\phi_2,\theta_1,\ldots,\theta_h$, ordered in some way; so $m \leqslant 2^{h+1}$, and the sentences V_i are satisfiable, mutually exclusive and exhaustive. Every sentence $\theta \in \overline{\theta}$ is equivalent to a disjunction $D(\theta)$ of sentences from $\overline{V}$, and also ϕ_2 is a disjunction D_2 of sentences from $\overline{V}$; moreover, D_2 is different from $\bigvee\overline{V}$, otherwise ϕ_2 would be trivial.

Now we take fresh atoms $R_1,\ldots,R_m$ and we define the sentences

$$\phi_1'(R_1,\ldots,R_m) = \phi_1(\bigvee\{R_i|V_i \in D_2\}) \tag{45}$$

$$\chi'(R_1,\ldots,R_m) = \chi(\bigvee\{R_i|V_i \in D(\theta_1)\},\ldots,\bigvee\{R_i|V_i \in D(\theta_h\})), \tag{46}$$

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

so we obtain the following chain of equivalences:

$$\phi_1(l\phi_2) \iff \phi_1'(lV_1,\ldots,lV_m) \iff \chi'(lV_1,\ldots,lV_m) \iff \chi(l\overline{\theta}). \tag{47}$$

At this point we have to distinguish two cases.

Case 1: If there is $V \in D_2$ such that $V \in D(\theta)$ implies $\theta \in C_2\Pi_k$-hard for any $\theta \in \overline{\theta}$, take $W \in \overline{V} - D_2$ and take a fresh atom P.

Now we claim the following equivalence:

$$\chi'([lV\backslash P],[lW\backslash\neg P],\mathit{false}) \iff \phi_1'([lV\backslash P],\mathit{false}), \tag{48}$$

where the square brackets mean substitution, and where the *false* replaces every lV' with V' in $\overline{V}$ different from V and W.

To prove the claim above, let us distinguish the two implications.

For the implication from left to right, let M be a model of $\chi'(P,\neg P,\mathit{false})$. Define a model N in the following way. Fix a model M_V of V and a model M_W of W. For any l-free word v, let the evaluation of every atom in N at v coincide with that of M. Moreover, for any word v verifying P in M, take $N(vl) = M_V$, and for any word v verifying $\neg P$ in M, take $N(vl) = M_W$. Then by the semantics of composition, N verifies $\chi'(lV,lW,\mathit{false})$, and hence *a fortiori* it also verifies $\chi'(lV_1,\ldots,lV_m)$. By the previous equivalence, N verifies $\phi_1'(lV_1,\ldots,lV_m)$, and by the semantics of composition, this means that there are some sets $X_1,\ldots,X_m$ of l-free words such that N verifies $\phi_1'(X_1,\ldots,X_m)$ and N verifies lV_i at every element of X_i. Now, by definition of N, for $V_i = V$ we must have $X_i \subseteq ||P||M$, otherwise some element of X_i should verify in N both lV and lW, which are incompatible formulas; and, similarly, for $V_j = W$ we must have $X_j \subseteq ||\neg P||M$, and for $V_s \neq V, W$ we must have $X_s = \varnothing$. Moreover, since M and N agree on the l-free words, M also verifies $\phi_1'(X_1,\ldots,X_m)$. So, by the semantics of composition, M verifies $\phi_1'([lV\backslash P],\mathit{false})$.

Conversely, let M be a model of $\phi_1'([lV\backslash P],\mathit{false})$. Build N as in the previous implication. Then, again by the semantics of composition, N verifies $\phi_1'(lV,\mathit{false})$, and hence *a fortiori* it also verifies $\phi_1'(V_1,\ldots,V_m)$. By the previous equivalence, N verifies $\chi'(lV_1,\ldots,lV_m)$, and by the semantics of composition, this means that there are some sets $X_1,\ldots,X_m$ of l-free words such that N verifies $\chi'(X_1,\ldots,X_m)$ and N verifies lV_i at every element of X_i. Now, again by definition of N, for $V_i = V$ we must have $X_i \subseteq ||P||M$, for $V_j = W$ we must have $X_j \subseteq ||\neg P||M$, and for $V_s \neq V, W$ we must have $X_s = \varnothing$. And again, since M and N agree on the l-free words, M also verifies $\chi'(X_1,\ldots,X_m)$. So, by the semantics of composition, M verifies $\chi'([lV\backslash P],[lW\backslash\neg P],\mathit{false})$. This concludes the proof of the claim.

In particular, since $\phi_1'([lV\backslash P],\mathit{false})$ is equivalent to $\phi_1(P)$, we obtain

$$\chi'([lV\backslash P],[lW\backslash\neg P],\mathit{false}) \iff \phi_1(P). \tag{49}$$

Now, since $\phi_1(P)$ is $\Sigma_2(P^+)$-hard, we have that some positive P is within a Σ_2-hard block in $\chi'([lV\backslash P],[lW\backslash\neg P],\mathit{false})$, hence some θ with $V \in D(\theta)$ is within some Σ_2-hard block in $\chi(l\overline{\theta})$. Hence, by our assumption on V, some $\theta \in C_2\Pi_k$-hard is surrounded by some Σ_2-hard block in $\chi(l\overline{\theta})$, that is $\chi(l\overline{\theta})$ is $C_2\Sigma_{k+1}$-hard. So, since collection does not increase the complexity, ψ is $C_2\Sigma_{k+1}$-hard as well, and we are done.

Case 2: If the opposite of Case 1 holds, namely for every $V \in D_2$ there is $\theta \notin C_2\Pi_k$-hard

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

with $V \in D(\theta)$, then for each $V \in D_2$ define a subset of $\overline{V}$ by

$$U(V) = \bigcap \{D(\theta) | \theta \notin C_2\Pi_k - \textit{hard and } V \in D(\theta)\}. \tag{50}$$

Note that for some $V \in D_2$ we must have $U(V)$ not included in D_2, otherwise ϕ_2 would be equivalent to a finite union of mu-calculus sentences not in $C_2\Pi_k$-hard, contrary to the $C_2\Pi_k$-hardness of the property ϕ_2.

So we can pick $V \in D_2$ and $W \in U(V) - D_2$. By taking a fresh atom P and by reasoning as in Case 1, we again obtain

$$\chi'([lV \backslash P], [lW \backslash \neg P], \textit{false}) \iff \phi_1(P). \tag{51}$$

Once again, since $\phi_1(P)$ is $\Sigma_2(P^+)$-hard, we have that some positive P is within some Σ_2-hard block in ψ, and hence some θ with $V \in D(\theta)$ is within Γ in $\chi(l\overline{\theta})$. Calling an element of $\overline{\theta}$ *irrelevant* if it is replaced by $P \vee \neg P = \textit{true}$ in the above equivalence, and *relevant* otherwise, we note that all the θ's with both $V \in D(\theta)$ and $W \in D(\theta)$ are irrelevant. Hence every relevant θ with $V \in D(\theta)$ verifies $W \notin D(\theta)$, and hence $\theta \in C_2\Pi_k$-hard by choice of W. Now a relevant θ with $V \in D(\theta)$ must exist, otherwise $\phi_1(P)$ would be equivalent to a formula negative in P, and this is impossible because ϕ_1 is satisfiable and has the 'non-emptyness property' defined in the statement.

Hence, picking one such θ, we have that θ is in $C_2\Pi_k$-hard and is surrounded by some Σ_2-hard block in $\chi(l\overline{\theta})$, and hence $\chi(l\overline{\theta})$ is $C_2\Sigma_{k+1}$-hard and so is ψ, and we are done again. □

9. Proof of the main result

This section is devoted to stating and proving the main result of this paper.

Theorem 9.1. The following mu-calculus sentence defines an f.a.d. 3-hard property on the ternary tree $\{a, b, c\}^*$:

$$\phi = \nu X.P \wedge I(cJ(cX)), \tag{52}$$

where we recall that

$$I(P) = \nu X.\mu Y.(P \wedge (aX \vee bX)) \vee aY \vee bY \tag{53}$$

$$J(P) = \mu X.\nu Y.(P \vee (aX \wedge bX)) \wedge aY \wedge bY. \tag{54}$$

Proof. For any formula α mentioned in the proof, let 'α' stand for 'the property defined by α'.

It is enough to show that ϕ is $C_2\Pi_{2k}$-hard for any k. This can be proved inductively. For $k = 0$ we observe that we can apply the composition lemma to $J(c\phi)$; now $J(Q)$ is $\Sigma_2(Q^+)$-hard, and ϕ is $C_2\Pi_0$-hard just because this last set is trivial. Hence, by the composition lemma, we have that $J(c\phi)$ is Σ_2-hard, that is, $C_2\Sigma_1$-hard. Now we consider $P \wedge I(cJ(c\phi))$, and, since $P \wedge I(Q)$ is $\Pi_2(Q^+)$-hard and as we have just proved that $J(c\phi)$ is $C_2\Sigma_1$-hard, by the composition lemma, we conclude that $P \wedge I(cJ(c\phi))$ is $C_2\Pi_2$-hard. But this last formula is equivalent to ϕ, and hence ϕ is $C_2\Pi_2$-hard.

For the inductive step the schema is the same: assuming that ϕ is $C_2\Pi_{2k}$-hard, by applying the composition lemma, we get that $J(c\phi)$ is $C_2\Sigma_{2k+1}$-hard, and that $P \wedge I(cJ(c\phi))$

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

is $C_2\Pi_{2k+2}$-hard. But the last sentence is equivalent to ϕ, and hence ϕ is $C_2\Pi_{2k+2}$-hard. This concludes the proof of the theorem. □

As a byproduct of the proof we obtain the following corollary.

Corollary 9.1. Define a sequence of formulas ψ_k by:

1 $\psi_0(P) = P$.
2 $\psi_{k+1}(P) = I(cJ(c\psi_k))$.

Then for any k, property ψ_k has f.a.d. 2 and is $C_2\Pi_{2k}$-hard. Hence the C_2-hierarchy is infinite.

Proof. The proof is again by induction on k using the composition lemma. □

10. Conclusion

In this section we discuss briefly the pros and cons of the results and of the methods of the present paper.

This paper is part of a longstanding investigation on the mu-calculus hierarchy problem. Let us first make some historical remarks on the problem (see also Lenzi (1997)).

The first steps towards the solution were made by Emerson and Clarke (Emerson and Clarke 1980), and Niwiński (Niwiński 1986). Emerson and Clarke showed that the infinity formula lies in the second level of the hierarchy; and Niwiński proved that the hierarchy is infinite if one suitably restricts the use of conjunction of formulas. Then at the beginning of the nineties, Emerson and Jutla proposed a sequence W_n of candidate hard formulas (Emerson and Jutla 1991). These formulas express the existence of a winning strategy for one of the two players of a suitable game on graphs.

The infinity of the whole hierarchy has been first proved by Bradfield (Bradfield 1996), with an ingenious indirect proof, but without giving explicit examples of hard formulas. Then a sequence ϕ_n of hard examples was proposed by the author in Lenzi (1996) and, with an expanded proof, in Lenzi (1997), but only for the positive fragment of the mu-calculus. Later, Bradfield gave another sequence B_n of examples and a proof of their hardness (Bradfield 1998).

Recently, Arnold proved that all the three sequences above are hard using a new and elegant technique (Arnold 1999). In particular, Arnold proposed a notion of reduction between mu-calculus formulas, and shows that: under the reduction, for every n, both W_n and B_n are complete for Σ_n and, for $n \geqslant 2$, ϕ_n is complete for Σ_{n-2} (it remains open whether ϕ_n is complete for Σ_n); and no Σ_n-complete property can lie in Π_n, by a nice diagonal argument, using the Banach fixpoint theorem for contractions in complete metric spaces.

By the historical remarks above, it is clear that by now we know quite a lot of hard formulas for any level of fixpoint alternation depth. In this paper we have considered a formula ϕ and shown that it is hard for f.a.d. 3. So the potential interest of this paper is not in the result in itself, but rather in the proof method. In fact, let us recall the idea of the proof.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

The idea is that ϕ satisfies the fixpoint equation

$$\phi \Leftrightarrow P \wedge I(cI^d(c\phi)), \tag{55}$$

and hence, for any number n, ϕ is equivalent to a nested composition of n formulas alternatingly $\Pi_2(Q^+)$-hard (that is, $P \wedge I(Q)$) and $\Sigma_2(Q^+)$-hard (that is, $I^d(Q)$), so the C_2-complexity of ϕ is infinite and ϕ cannot be expressed in f.a.d. 2.

In order to fill in the proof, the crux is to show that:

1 $I(Q)$ is $\Pi_2(Q^+)$-hard (and the dual statement).
2 The result of replacing Q in a $\Pi_2(Q^+)$-hard property with a $C_2\Sigma_k$-hard one is $C_2\Pi_{k+1}$-hard (and the dual statement).

The second point is what we have called the 'composition lemma'. There is quite a subtle point about the statement of this lemma. In fact, the reader may wonder whether we can replace this strange definition of, say, $\Pi_2(Q^+)$-hardness with more standard properties, such as semantical monotonicity in Q, Π_2-hardness and absence of atoms different from Q. But, to my knowledge, these three properties imply $\Pi_2(Q^+)$-hardness only assuming the still conjectural Lyndon theorem for $\Sigma_2(Q^+)$, that is every monotonic $\Sigma_2(Q^+)$ formula is equivalent to a positive $\Sigma_2(Q^+)$ formula (note that the Lyndon theorem is true for the full mu-calculus, see D'Agostino (2000)).

From the above it follows that, on the positive side, proving the Lyndon theorem for $\Sigma_2(Q^+)$ would greatly simplify the proof of the main result, and, in particular, there would be no need for the games G and G^- introduced in this paper.

Now, in the literature there seems to be no tool for finding, say, $\Pi_n(Q^+)$-hard properties, even though several Π_n-hard properties are known for any n (once again unless the Lyndon theorem is true for $\Pi_n(Q^+)$, in which case the problem becomes trivial). So in order to prove the existence of $\Pi_2(Q^+)$-hard properties, I needed to figure out one such tool, and what I found is the games G and G^- above. The point of these games is that every sentence ψ such that player *Two* has a winning strategy for $G^-(\psi)$ must define a $\Pi_2(Q^+)$-hard property. So we have a sufficient condition for $\Pi_2(Q^+)$-hardness. Unfortunately, the method relies upon the particular semantical properties of the class Π_2, and hence it does not generalize easily to $\Pi_n(Q^+)$-hardness for $n > 2$; and this is why we cannot use this method as it is to find f.a.d. n hard formulas for $n > 3$.

Moreover, applying to sentences the procedure that we call 'collection' seems necessary to prove the composition lemma: in fact we said that, if we 'split' a sentence ψ as $\chi(l\overline{\omega})$, where $\overline{\omega}$ is a tuple of formulas some of which have free variables, then this splitting is not a legal composition because there is variable capturing, and hence we cannot apply the semantics of composition, which is crucial to prove the lemma.

Note also that the composition lemma generalizes from the C_2-complexity case to the C_n-complexity case, for every natural number n. So, following the same pattern as for $n = 2$, we can show that for any n, if a $\Pi_n(Q^+)$-hard formula exists, then the C_n hierarchy is infinite.

It must be said that both the main result and the proof method of this paper are already essentially in Lenzi (1997), except for a different choice of the definition of fixpoint alternation depth. In fact, in Lenzi (1997) the Emerson–Lei definition of f.a.d. is

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

used, which was proposed in Emerson and Lei (1986) and which is more restrictive than the Niwiński definition (the f.a.d. classes are closed only under composition of *sentences*, rather than of arbitrary formulas). So in Lenzi (1997) it is shown that ϕ is hard for Emerson–Lei f.a.d. 3, and here we extend the result to Niwiński f.a.d. In order to get the extension it is sufficient to note that:

1 The characterization of all the formulas of f.a.d. $\leqslant n$ (or equivalently of the formulas in $\Sigma_{n+1} \cap \Pi_{n+1}$) as compositions of Σ_n and Π_n formulas is valid both for the Emerson–Lei definition and for the one of Niwiński.
2 Likewise, the characterization of Π_2 by systems of inclusions taken from Arnold (1988), which allows us to prove that player *One* has a winning strategy for $G(\psi)$ for any sentence $\psi \in \Pi_2$, is also valid for both definitions.

We note that the semantical counterpart of the above characterization of $\Sigma_{n+1} \cap \Pi_{n+1}$, that is, any property that is both Π_{n+1} and Σ_{n+1} is a composition of Σ_n properties and Π_n properties, is true for $n = 0, 1$ (Arnold and Niwiński 1992), and is an interesting open problem for $n \geqslant 2$.

In spite of the achievements mentioned above, in the author's opinion, there is still a lot to be done in order to understand the structure of the mu-calculus and its f.a.d. levels. However, further investigations on this kind of problem are deferred to future work.

Acknowledgements

This paper is an elaboration of Chapter 4 of my Ph. D. thesis (Lenzi 1997). My thanks go to the Ph. D. thesis committee, and, in particular, to Allen Emerson and Furio Honsell for their encouragement.

References

Arnold, A. (1988) Logical definability of fixed points. *Theoretical Computer Science* **61** 289–297.

Arnold, A. (1999) The μ-calculus alternation-depth hierarchy is strict on binary trees. *Theoret. Informatics Appl.* **33** 329–339.

Arnold, A. and Niwiński, D. (1992) Fixed point characterization of weak monadic definable sets of trees. In: Nivat, M. and Podelski, A. (eds.) *Tree, automata and languages*, Elsevier 159–188.

Bradfield, J. (1996) The modal mu-calculus alternation hierarchy is strict. Proceedings of CONCUR '96. *Springer-Verlag Lecture Notes in Computer Science* **1119** 233–246.

Bradfield, J. (1998) Simplifying the modal mu-calculus alternation hierarchy. Proceedings of STACS'98. *Springer-Verlag Lecture Notes in Computer Science* **1373** 39–49.

D'Agostino, G. (2000) *Journal of Symbolic Logic* (to appear).

Emerson, E. and Clarke, E. (1980) Characterizing correctness properties of parallel programs using fixpoints. In: Proceedings of the 7th International Colloquium on Automata, Languages and Programming. *Springer-Verlag Lecture Notes in Computer Science* **85** 169–181.

Emerson, E. and Jutla, C. (1991) Tree automata, μ-calculus and determinacy. In: *Proceedings of FOCS '91* 368–377.

Emerson, E. and Lei, C. (1986) Efficient model checking in fragments of the propositional Mu-calculus. *Proceedings of LICS '86* 267–278.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press

Emerson, E. and Streett, R. (1989) An automata-theoretic decision procedure for the propositional Mu-calculus. *Information and Computation* **81** 249–264.

Kaivola, R. (1997) *Using automata to characterize fixed point temporal logics*, Ph. D. Thesis, University of Edinburgh.

Kozen, D. (1983) Results on the propositional mu-calculus. *Theoretical Computer Science* **27** 333–354.

Lenzi, G. (1996) A hierarchy theorem for the mu-calculus. Proceedings of ICALP'96. *Springer-Verlag Lecture Notes in Computer Science* **1099** 87–97.

Lenzi, G. (1997) *The mu-calculus and the hierarchy problem*, Ph. D. Thesis, Scuola Normale Superiore, Pisa.

Niwiński, D. (1986) On fixed point clones. Proceedings of the 13th International Colloquium on Automata, Languages and Programming. *Springer-Verlag Lecture Notes in Computer Science* **226** 464–473.

Niwiński, D. (1997) Fixed point characterization of infinite behavior of finite-state systems. *Theoretical Computer Science* **189** 1–69.

Tarski, A. (1955) A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* **55** 285–309.

https://doi.org/10.1017/S0960129500003194 Published online by Cambridge University Press