Continuous fragment of the μ-calculus

Gaëlle Fontaine*

Institute for Logic, Language and Computation,
Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands
gfontain@science.uva.nl
http://staff.science.uva.nl/~gfontain

Abstract. In this paper we investigate the Scott continuous fragment of the modal μ-calculus. We discuss its relation with constructivity, where we call a formula constructive if its least fixpoint is always reached in at most ω steps. Our main result is a syntactic characterization of this continuous fragment. We also show that it is decidable whether a formula is continuous.

Key words: mu-calculus, automata, Scott continuity, constructive fixpoints, preservation results

1 Introduction

This paper is a study into the fragment of the modal μ-calculus that we call continuous. Roughly, given a proposition letter p, a formula φ is said to be continuous in p if it monotone in p and if in order to establish the truth of φ at a point, we only need finitely many points at which p is true. The continuous fragment of the μ-calculus is defined as the fragment of the μ-calculus in which $\mu x.\varphi$ is allowed only if φ is continuous in x.

We prove the following two results. First, Theorem 2 gives a natural syntactic characterization of the continuous formulas. Informally, continuity corresponds to the formulas built using the operators $\vee$, $\wedge$, $\Diamond$ and μ. Second, we show in Theorem 3 that it is decidable whether a formula is continuous in p.

We believe that this continuous fragment is of interest for a number of reasons. A first motivation concerns the relation between continuity and another property, constructivity. The constructive formulas are the formulas whose fixpoint is reached in at most ω steps. Locally, this means that a state satisfies a least fixpoint formula if it satisfies one of its *finite* approximations. It is folklore that if a formula is continuous, then it is constructive. The other implication does not strictly hold. However, interesting questions concerning the link between constructivity and continuity remain. In any case, given our Theorem 2, continuity can be considered as the most natural candidate to approximate constructivity syntactically.

* The research of the author has been made possible by VICI grant 639.073.501 of the Netherlands Organization for Scientific Research (NWO).

Next this fragment can be seen as a natural generalization of PDL in the following way. We define the completely additive formulas as the formulas built using the operators $\vee$, $\Diamond$ and μ. That is, the syntax is the same as for the continuous formulas, except that the conjunction is not allowed. Then it was observed by Yde Venema (personal communication) that PDL coincides with the fragment of the μ-calculus in which $\mu x.\varphi$ is allowed only if φ is completely additive. In this perspective, the continuous fragment appears as a natural extension of PDL.

Another reason for looking at this fragment (which also explains the name) is the link with Scott continuity. A formula is continuous in p iff it is continuous with respect to p in the Scott topology on the powerset algebra (with all other variables fixed). Scott continuity is of key importance in many areas of theoretical computer sciences where ordered structures play a role, such as domain theory (see, e.g., [1]). For many purposes, it is sufficient to check that a construction is Scott continuous in order to show that it is computationally feasible.

Finally our results fit in a model-theoretic tradition of so-called preservation results (see, e.g., [2]). Giovanna D'Agostino and Marco Hollenberg have proved some results of this kind in the case of the μ-calculus (see, e.g., [3] and [4]). Their proofs basically consist in identifying automata corresponding to the desired fragment and in showing that these automata give the announced characterization. The proof of our main result is similar as we also first start by translating our problem in terms of automata. We also mention that a version of our syntactic characterization in the case of first order logic has been obtained by Johan van Benthem in [5].

The paper is organized as follows. First we recall the syntax of the μ-calculus and some basic properties that will be used later on. Next we define the continuous fragment and we show how it is linked to Scott continuity, constructivity and PDL. Finally we prove our main result (Theorem 2) which is a syntactic characterization of the fragment and we show that it is decidable whether a formula is continuous (Theorem 3). We end the paper with questions for further research.

2 Preliminaries

We introduce the language and the Kripke semantic for the μ-calculus.

Definition 1. *Let Prop be a finite set of proposition letters and let Var be a countable set of variables. The* formulas of the μ-calculus *are given by*

$$\varphi \ ::= \ \top \mid p \mid x \mid \varphi \vee \varphi \mid \neg\varphi \mid \Diamond\varphi \mid \mu x.\varphi,$$

where p ranges over the set $Prop$ and x ranges over the set Var. In $\mu x.\varphi$, we require that every occurrence of x is under an even number of negations in φ. The notion of closed μ-formula *or* μ-sentence *is defined in the natural way.*

As usual, we let $\varphi \wedge \psi$, $\Box\varphi$ and $\nu x.\varphi$ be abbreviations for $\neg(\neg\varphi \vee \neg\psi)$, $\neg\Diamond\neg\varphi$ and $\neg\mu x.\neg\varphi[\neg x/x]$. For a set of formulas Φ, we denote by $\bigvee \Phi$ the disjunction of formulas in Φ. Similarly, $\bigwedge \Phi$ denotes the conjunction of formulas in Φ.

Finally, we extend the syntax of the μ-calculus by allowing a new construct of the form $\nabla\Phi$, where Φ is a finite set of formulas. We will consider such a formula to be an abbreviation of $\bigwedge\{\Diamond\varphi : \varphi \in \Phi\} \wedge \Box \bigvee \Phi$. Remark that in [6], David Janin and Igor Walukiewicz use the notation $a \to \Phi$ for $\nabla_a \Phi$.

For reasons of a smooth presentation, we restrict to the unimodal fragment. All the results can be easily extended to the setting where we have more than one basic modality.

Definition 2. *A* Kripke frame *is a pair (M, R), where M is a set and R a binary relation on M. A* Kripke model *$\mathcal{M}$ is a triple (M, R, V) where (M, R) is a Kripke frame and $V : Prop \to \mathcal{P}(M)$ a valuation.*

If sRt, we say that t is a successor *of s and we write $R(s)$ to denote the set $\{t \in M : sRt\}$. A* path *is a (finite or infinite) sequence $s_0, s_1, \ldots$ such that $s_i R s_{i+1}$ (for all $i \in \mathbb{N}$).*

Definition 3. *Given a μ-formula φ, a model $\mathcal{M} = (M, R, V)$ and an assignment $\tau : Var \to \mathcal{P}(M)$, we define a subset $[\![\varphi]\!]_{\mathcal{M},\tau}$ of M that is interpreted as the set of points at which φ is true. This subset is defined by induction in the usual way. We only recall that*

$$[\![\mu x.\varphi]\!]_{\mathcal{M},\tau} = \bigcap\{U \subseteq M : [\![\varphi]\!]_{\mathcal{M},\tau[x:=U]} \subseteq U\},$$

where $\tau[x := U]$ is the assignment τ' such that $\tau'(x) = U$ and $\tau'(y) = \tau(y)$ for all $y \neq x$.

Observe that the set $[\![\mu x.\varphi]\!]_{\mathcal{M},\tau}$ is the least fixpoint of the map $\varphi_x : \mathcal{P}(M) \to \mathcal{P}(M)$ defined by $\varphi_x(U) := [\![\varphi]\!]_{\mathcal{M},\tau[x:=U]}$, for all $U \subseteq M$. Similarly, for a proposition letter p, we can define the map $\varphi_p : \mathcal{P}(M) \to \mathcal{P}(M)$ by $\varphi_p(U) := [\![\varphi]\!]_{\mathcal{M}[p:=U],\tau}$, where $\mathcal{M}[p := U]$ is the model (M, R, V') with $V'(p) = U$ and $V'(p') = V(p')$, for all $p' \neq p$.

If $s \in [\![\varphi]\!]_{\mathcal{M},\tau}$, we write $\mathcal{M}, s \Vdash_\tau \varphi$ and we say that φ true *at $s \in \mathcal{M}$ under the assignment τ. If φ is a sentence, we simply write $\mathcal{M}, s \Vdash \varphi$.*

A formula φ is monotone *in a proposition letter p if for all models $\mathcal{M} = (M, R, V)$, all assignments τ and all sets $U, U' \subseteq M$ satisfying $U \subseteq U'$, we have $\varphi_p(U) \subseteq \varphi_p(U')$. The notion of monotonicity in a variable x is defined in an analogous way.*

Finally we use the notation $\varphi \models \psi$ if for all models $\mathcal{M}$ and all points $s \in \mathcal{M}$, we have $\mathcal{M}, s \Vdash \varphi$ implies $\mathcal{M}, s \Vdash \psi$.

When deciding whether a sentence is true at a point s, it only depends on the points accessible (in possibly many steps) from s. These points together with the relation and the valuation inherited from the original model form the submodel generated by s. We will use this notion later on and we briefly recall the definition.

Definition 4. *Let $\mathcal{M} = (M, R, V)$ be a model. A subset N of M is* downward closed *if for all s and t, sRt and $t \in N$ imply that $s \in N$. N is* upward closed *if for all s and t, sRt and $s \in N$ imply that $t \in N$.*

A model $\mathcal{N} = (N, S, U)$ is a generated submodel *of $\mathcal{M}$ if $N \subseteq M$, N is upward closed, $S = R \cap (N \times N)$ and $U(p) = V(p) \cap N$, for all $p \in Prop$. If N' is a subset of M, we say that $\mathcal{N} = (N, S, U)$ is the* submodel generated *by N' if $\mathcal{N}$ is a generated submodel and if N is the smallest upward closed set containing N'.*

In our proof, it will be often more convenient to work with a certain kind of Kripke models. That is, we will suppose that the models we are dealing with are trees such that each point (except the root) has infinitely many bisimilar siblings. We make this definition precise and we give the results needed to justify this assumption.

Definition 5. *A point s is a* root *of a model $\mathcal{M} = (M, R, V)$ if for every t distinct from s, there is a path from s to t. $\mathcal{M}$ is a* tree *if it has a root, every point distinct from the root has a unique predecessor and R is acyclic (that is, there is no non-empty path starting at a point t and ending in t).*

A model $\mathcal{M} = (M, R, V)$ is ω-expanded *if it is a tree such that for all $s \in M$ and all successors t of s, there are infinitely many distinct successors of s that are bisimilar to t.*

Proposition 1. *Let $\mathcal{M} = (M, R, V)$ be a model and let $s \in M$. There exists a tree $\mathcal{M}' = (M', R', V')$ that is ω-expanded such that s and the root s' of $\mathcal{M}'$ are bisimilar. In particular, for all μ-sentences φ, $\mathcal{M}, s \Vdash \varphi$ iff $\mathcal{M}', s' \Vdash \varphi$.*

Another way to look at formulas of the μ-calculus is to consider automata. In [6], David Janin and Igor Walukiewicz define a notion of automaton that operates on Kripke models and that corresponds exactly to the μ-calculus.

Definition 6. *A* μ-automaton $\mathbb{A}$ *over a finite alphabet Σ is a tuple (Q, q_0, δ, Ω) such that Q is a finite set of states, $q_0 \in Q$ is the initial state, $\delta : Q \times \Sigma \to \mathcal{P}\mathcal{P}(Q)$ is the transition map and $\Omega : Q \to \mathbb{N}$ is the parity function.*

Given a frame $\mathcal{M} = (M, R, V)$ with a labeling $L : M \to \Sigma$ and a point $s \in M$, an $\mathbb{A}$-game in $\mathcal{M}$ with starting position (s, q_0) is played between two players, the Duplicator and the Spoiler. The game is as follows: If we are in position (t, q) (where $t \in M$ and $q \in Q$), the Duplicator has to make a move. The Duplicator chooses a marking $m : Q \to \mathcal{P}\{u : tRu\}$ and then a description D in $\delta(q, L(t))$. If $u \in m(q)$, we say that u is marked *with q.*

The marking and the description have to satisfy the two following properties. First, if $q' \in D$, there exists a successor u of t that is marked with q'. Second, if u is a successor of t, there exists $q' \in D$ such that u is marked with q'. After the Duplicator has chosen a marking m, the Spoiler plays a position (u, q') such that $t \in m(q')$.

Either player wins the game if the other player cannot make a move. An infinite match $(s, q_0), (s_1, q_1), \ldots$ is won by the Duplicator if the smallest element of $\{\Omega(q) : q$ appears infinitely often in $q_0, q_1, \ldots\}$ is even.

We say that $(\mathcal{M}, s)$ is accepted *by $\mathbb{A}$ if the Duplicator has a winning strategy in the $\mathbb{A}$-game in $\mathcal{M}$ with starting position (s, q_0).*

Remark that a model (M, R, V) can be seen as a frame (M, R) with a labeling $L : M \to \mathcal{P}(Prop)$ defined by $L(t) = \{p \in Prop : t \in V(p)\}$, for all $t \in M$.

Theorem 1. *[6] For every μ-automaton $\mathbb{A}$ (over the alphabet $\mathcal{P}(Prop)$), there is a sentence φ such that for all models $\mathcal{M}$ and all points $s \in \mathcal{M}$, $\mathbb{A}$ accepts $(\mathcal{M}, s)$ iff $\mathcal{M}, s \Vdash \varphi$. Conversely, for every sentence φ, there is a μ-automaton $\mathbb{A}$ such that for all models $\mathcal{M}$ and all points $s \in \mathcal{M}$, $\mathbb{A}$ accepts $(\mathcal{M}, s)$ iff $\mathcal{M}, s \Vdash \varphi$.*

3 Continuity

We define the notion of continuity for a formula and we show the connection with the Scott continuity. We also mention that these formulas are constructive and that there is a natural connection with PDL.

Definition 7. *Fix a proposition letter p. A sentence φ is* continuous in p *if for all models $\mathcal{M} = (M, R, V)$ and all $s \in M$, we have*

$$\mathcal{M}, s \Vdash \varphi \textit{ iff } \exists\, F \subseteq V(p) \textit{ s.t. } F \textit{ is finite and } \mathcal{M}[p := F], s \Vdash \varphi.$$

The notion of continuity in x *(where x is a variable) is defined similarly.*

That is, a formula φ is continuous in p iff it is monotone in p and whenever φ is true at a point in a model, we only need finitely many points where p is true in order to establish the truth of φ.

Continuity and Scott continuity

It does not seem very natural that a formula satisfying such a property should be called continuous. In fact, it is equivalent to require that the formula is Scott continuous with respect to p in the powerset algebra (with all other proposition letters fixed). In the next paragraph, we recall the definition of the Scott topology and we briefly show that the notion of Scott continuity and our last definition coincide.

Definition 8. *Let $\mathcal{M} = (M, R, V)$ be a model. A family $\mathcal{F}$ of subsets of M is* directed *if for all $U_1, U_2 \in \mathcal{F}$, there exists $U \in \mathcal{F}$ such that $U \supseteq U_1 \cup U_2$.*

A Scott open set *in the powerset algebra $\mathcal{P}(M)$ is a family $\mathcal{O}$ of subsets of M that is closed under upset (that is, if $U \in \mathcal{O}$ and $U' \supseteq U$, then $U' \in \mathcal{O}$) and such that for all directed family $\mathcal{F}$ satisfying $\bigcup \mathcal{F} \in \mathcal{O}$, the intersection $\mathcal{F} \cap \mathcal{O}$ is non-empty.*

As usual, a map $f : \mathcal{P}(M) \to \mathcal{P}(M)$ is Scott continuous *if for all Scott open sets $\mathcal{O}$, the set $f^{-1}[\mathcal{O}] = \{f^{-1}(U) : U \in \mathcal{O}\}$ is Scott open.*

Fix a proposition letter p. A sentence φ is Scott continuous in p *if for all models $\mathcal{M} = (M, R, V)$, the map $\varphi_p : \mathcal{P}(M) \to \mathcal{P}(M)$ is Scott continuous.*

Remark that the Scott topology can be defined in an arbitrary partial order (see, e.g., [7]). It is a fairly standard result that a map f is Scott continuous iff it preserves directed joins. That is, for all directed family $\mathcal{F}$, we have $f(\bigcup \mathcal{F}) = \bigcup f[\mathcal{F}]$ (where $f[\mathcal{F}] = \{f(U) : U \in \mathcal{F}\}$). Now we check that our notion of continuity defined in a Kripke semantic framework is equivalent to the standard definition of Scott continuity.

Proposition 2. *A sentence is continuous in p iff it is Scott continuous in p.*

Proof. For the direction from left to right, let φ be a continuous sentence in p. Fix a model $\mathcal{M} = (M, R, V)$. We show that the map $\varphi_p : \mathcal{P}(M) \to \mathcal{P}(M)$ preserves directed joins.

Let $\mathcal{F}$ be a directed family. It follows from the monotonicity of φ that the set $\bigcup \varphi_p[\mathcal{F}]$ is a subset of $\varphi_p(\bigcup \mathcal{F})$. Thus, it remains to show that $\varphi_p(\bigcup \mathcal{F}) \subseteq \bigcup \varphi_p[\mathcal{F}]$. Take s in $\varphi_p(\bigcup \mathcal{F})$. That is, the formula φ is true at s in the model $\mathcal{M}[p := \bigcup \mathcal{F}]$. As φ is continuous in p, there is a finite subset F of $\bigcup \mathcal{F}$ such that φ is true at s in $\mathcal{M}[p := F]$. Now, since F is a finite subset of $\bigcup \mathcal{F}$ and since $\mathcal{F}$ is directed, there exists a set U in $\mathcal{F}$ such that F is a subset of U. Moreover, as φ is monotone, $\mathcal{M}[p := F], s \Vdash \varphi$ implies $\mathcal{M}[p := U], s \Vdash \varphi$. Therefore, s belongs to $\varphi_p(U)$ and in particular, s belongs to $\bigcup \varphi_p[\mathcal{F}]$. This finishes to show that $\varphi_p(\bigcup \mathcal{F}) \subseteq \bigcup \varphi_p[\mathcal{F}]$.

For the direction from right to left, let φ be a Scott continuous sentence in p. First we show that φ is monotone in p. Let $\mathcal{M} = (M, R, V)$ be a model. We check that $\varphi_p(U) \subseteq \varphi_p(U')$, in case $U \subseteq U'$. Suppose $U \subseteq U'$ and let $\mathcal{F}$ be the set $\{U, U'\}$. The family $\mathcal{F}$ is clearly directed and satisfies $\bigcup \mathcal{F} = U'$. Using the fact that φ_p preserves directed joins, we get that $\varphi_p(U') = \varphi_p(\bigcup \mathcal{F}) = \bigcup \varphi_p[\mathcal{F}]$. By definition of $\mathcal{F}$, we have $\bigcup \varphi_p[\mathcal{F}] = \varphi_p(U) \cup \varphi_p(U')$. Putting everything together, we obtain that $\varphi_p(U') = \varphi_p(U) \cup \varphi_p(U')$. Thus, $\varphi_p(U) \subseteq \varphi_p(U')$.

To show that φ is continuous in p, it remains to show that if $\mathcal{M}, s \Vdash \varphi$, then there exists a finite subset F of $V(p)$ such that $\mathcal{M}[p := F], s \Vdash \varphi$. Suppose that the formula φ is true at s in $\mathcal{M}$. That is, s belongs to $\varphi_s(V(p))$. Now let $\mathcal{F}$ be the family $\{F \subseteq V(p) : F \text{ finite}\}$. It is not hard to see that $\mathcal{F}$ is a directed family satisfying $\bigcup \mathcal{F} = V(p)$. Since φ_p preserves directed joins, we obtain $\varphi_p(V(p)) = \varphi_p(\bigcup \mathcal{F}) = \bigcup \varphi_p[\mathcal{F}]$. From $s \in \varphi_p(V(p))$, it then follows that $s \in \bigcup \varphi_p[\mathcal{F}]$. Therefore, there exists $F \in \mathcal{F}$ such that $s \in \varphi_p(F)$. That is, F is a finite subset of $V(p)$ such that $\mathcal{M}[p := F], s \Vdash \varphi$.

Continuity and constructivity

A formula is constructive if its fixpoint is reached in at most ω steps. Formally, we have the following definition.

Definition 9. *Fix a proposition letter p. A sentence φ is* constructive *in p if for all models $\mathcal{M} = (M, R, V)$, the least fixpoint of the map $\varphi_p : \mathcal{P}(M) \to \mathcal{P}(M)$ is equal to $\bigcup \{\varphi_p^i(\emptyset) : i \in \mathbb{N}\}$ (where φ_p^i is defined by induction by $\varphi_p^0 = \varphi_p$ and $\varphi_p^{i+1} = \varphi_p \circ \varphi_p^i$).*

Locally, this means that given a formula φ constructive in p and a point s in a model at which $\mu p.\varphi$ is true, there is some natural number n such that s belongs to the finite approximation $\varphi_p^n(\emptyset)$. We observe that a continuous formula is constructive.

Proposition 3. *A sentence φ continuous in p is constructive in p.*

Proof. Let φ be a sentence continuous in p and let $\mathcal{M} = (M, R, V)$ be a model. We show that the least fixpoint of φ_p is $\bigcup\{\varphi_p^i(\emptyset) : i \in \mathbb{N}\}$.

Let $\mathcal{F}$ be the family $\{\varphi_p^i(\emptyset) : i \in \mathbb{N}\}$. It is enough to check that $\varphi_p(\bigcup \mathcal{F}) = \bigcup \mathcal{F}$. First remark that $\mathcal{F}$ is directed. Therefore, $\varphi_p(\bigcup \mathcal{F}) = \bigcup \varphi_p[\mathcal{F}]$. It is also easy to prove that $\bigcup \varphi_p[\mathcal{F}] = \bigcup \mathcal{F}$. Putting everything together, we obtain that $\varphi_p(\bigcup \mathcal{F}) = \bigcup \mathcal{F}$ and this finishes the proof.

Remark that a constructive sentence might not be continuous.

Example 1. Let φ be the formula $\Box p \wedge \Box\Box\bot$. Basically, φ is true at a point s in a model if the depth of s is less or equal to 2 (that is, there are no t and t' satisfying $sRtRt'$) and all successors of s satisfy p. It is not hard to see that φ is not continuous in p. However, we have that for all models $\mathcal{M} = (M, R, V)$, $\varphi_p^2(\emptyset) = \varphi_p^3(\emptyset)$. In particular, φ is constructive in p.

Example 2. Let ψ be the formula $\nu x.p \wedge \Diamond x$. The formula ψ is true at a point s if there is a infinite path starting from s and at each point of this path, p is true. This sentence is not continuous in p. However, it is constructive, since for all models $\mathcal{M} = (M, R, V)$, we have $\psi_p(\emptyset) = \emptyset$.

Observe that in the previous examples we have $\mu p.\varphi \equiv \mu p.\Box\Box\bot$ and $\mu p.\psi \equiv \mu p.\bot$. Thus, there is a continuous sentence (namely $\Box\Box\bot$) that is equivalent to φ, modulo the least fixpoint operation. Similarly, there is a continuous sentence (the formula $\bot$) that is equivalent to ψ, modulo the least fixpoint operation. This suggests the following question.

Question 1 (Yde Venema). Given a constructive formula φ, can we find a continuous formula ψ satisfying $\mu p.\varphi \equiv \mu p.\psi$?

The answer is still unknown and this could be a first step for further study of the relation between continuity and constructivity.

Decidability of constructivity is also an interesting question. We would like to mention that in [8], Martin Otto proved that it is decidable in EXPTIME whether a basic modal formula $\varphi(p)$ is bounded. We recall that a basic modal formula $\varphi(p)$ is bounded if there is a natural number n such that for all models $\mathcal{M}$, we have $\varphi_p^n(\emptyset) = \varphi_p^{n+1}(\emptyset)$.

Continuity and PDL

We finish this section by few words about the connection between the continuous fragment and PDL. We start by defining the completely additive formulas.

Definition 10. *Let P be a subset of $Prop$ and let X be a subset of Var. The set of* completely additive formulas *with respect to $P \cup X$ is defined by induction in the following way:*

$$\varphi ::= \top \mid p \mid x \mid \psi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid \mu y.\chi,$$

where p is in P, x is in X, ψ is a formula of the μ-calculus such that the proposition letters of ψ and the variables of ψ do not belong to $P \cup X$ and χ is completely additive with respect to $P \cup X \cup \{y\}$.

We define the *completely additive fragment* as the fragment of the μ-calculus in which $\mu x.\varphi$ is allowed only if φ is completely additive with respect to x. As mentioned in the introduction, it was observed by Yde Venema that this fragment coincides with test-free PDL.

Similarly, we define the *continuous fragment* as the fragment of the μ-calculus in which $\mu x.\varphi$ is allowed only if φ is continuous in x. It is routine to check that any completely additive formula with respect to p is continuous in p (and the proof is similar to the proof of Lemma 1 below). In particular, the completely additive fragment is included in the continuous fragment. That is, PDL is a subset of the continuous fragment. We remark that this inclusion is strict. An example is the formula $\varphi = \mu x.\,(\Diamond(p \wedge x) \wedge \Diamond(q \wedge x))$. This formula belongs to the continuous fragment but is not equivalent to a formula in PDL. Roughly, the sentence φ is true at a point s if there is a finite binary tree-like submodel starting from s, such that each non-terminal node of the tree has a child at which p is true and a child at which q is true. This example was given by Johan van Benthem in [9].

4 Syntactic characterization of the continuous fragment

In this section, we give a characterization of the continuous fragment of the μ-calculus. The main result states that the sentences which are continuous in p are exactly the sentences such that p and the variables are only in the scope of the operators $\vee$, $\wedge$, $\Diamond$ and μ. These formulas are formally defined as the set $CF(p)$.

Definition 11. *Let P be a subset of $Prop$ and let X be a subset of Var. The set of formulas $CF(P \cup X)$ is defined by induction in the following way:*

$$\varphi ::= \top \mid p \mid x \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \mu y.\chi,$$

where p is in P, x is in X, ψ is a formula of the μ-calculus such that the proposition letters of ψ and the variables of ψ do not belong to $P \cup X$ and χ belongs to $CF(P \cup X \cup \{y\})$. We abbreviate $CF(\{p\})$ to $CF(p)$.

As a first property, we mention that the formulas in $CF(P \cup X)$ are closed under composition.

Proposition 4. *If φ is in $CF(P \cup X \cup \{p\})$ and ψ is in $CF(P \cup X)$, then $\varphi[\psi/p]$ belongs to $CF(P \cup X)$.*

Proof. By induction on φ.

Next we observe that the sentences in $CF(p)$ are continuous.

Lemma 1. *A sentence φ in $CF(p)$ is continuous in p.*

Proof. We prove by induction on φ that for all sets $P \subseteq Prop$ and $X \subseteq Var$, $\varphi \in CF(P \cup X)$ implies that φ is continuous in p and in x, for all $p \in P$ and all $x \in X$. We focus on the inductive step $\varphi = \mu y.\chi$, where χ is in $CF(P \cup X \cup \{y\})$. We also restrict ourselves to show that φ is continuous in p, for a proposition letter p in P.

Fix a proposition letter $p \in P$. First we introduce the following notation. For a model $\mathcal{M} = (M, R, V)$, an assignment τ and a subset U of W, we let $\chi_y^U : \mathcal{P}(M) \to \mathcal{P}(M)$ be the map defined by $\chi_y^U(W) = [\![\chi]\!]_{\mathcal{M}[p:=U],\tau[y:=W]}$, for all $W \subseteq M$. We also denote by $f(U)$ the least fixpoint of χ_y^U.

Now we show that φ is monotone in p. That is, for all models $\mathcal{M} = (M, R, V)$, all assignments τ and all subsets U, U' of M such that $U \subseteq U'$, we have $\mathcal{M}[p := U], s \Vdash_\tau \mu y.\chi$ implies $\mathcal{M}[p := U'], s \Vdash_\tau \mu y.\chi$. Fix a model $\mathcal{M} = (M, R, V)$, an assignment τ and sets $U, U' \subseteq M$ satisfying $U \subseteq U'$. Suppose $\mathcal{M}[p := U_0], s \Vdash_\tau \mu y.\chi$. That is, s belongs to the least fixpoint $f(U)$of the map χ_y^U. Since χ is monotone in p, we have that for all $W \subseteq M$, $\chi_y^U(W) \subseteq \chi_y^{U'}(W)$. It follows that the least fixpoint $f(U)$ of the map χ_y^U is a subset of the least fixpoint of the map $\chi_y^{U'}$. Putting this together with $s \in f(U)$, we get that s belongs to the least fixpoint of $\chi_y^{U'}$. That is, $\mathcal{M}[p := U'], s \Vdash_\tau \mu y.\chi$ and this finishes the proof that φ is monotone in p.

Next suppose that $\mathcal{M}, s \Vdash_\tau \mu y.\chi$, for a point s in a model $\mathcal{M} = (M, R, V)$. That is, s belongs to the least fixpoint L of the map χ_y. Now let $\mathcal{F}$ be the set of finite subsets of $V(p)$. We need to find a set $F \in \mathcal{F}$ satisfying $\mathcal{M}[p := F], s \Vdash_\tau \mu y.\chi$. Or equivalently, we have to show that there exists $F \in \mathcal{F}$ such that s belongs to the least fixpoint $f(F)$ of the map χ_y^F.

Let $\mathcal{G}$ be the set $\{f(F) : F \in \mathcal{F}\}$. It is routine to show that $\mathcal{G}$ is a directed family. Since L is the least fixpoint of χ_y, we have that for all $U \subseteq M$, $\chi_y(U) \subseteq U$ implies $L \subseteq U$. So if we can prove that $\chi_y(\bigcup \mathcal{G}) \subseteq \bigcup \mathcal{G}$, we will obtain $L \subseteq \bigcup \mathcal{G}$. Putting this together with $s \in L$, it will follow that $s \in \bigcup\{f(F) : F \in \mathcal{F}\}$. Therefore, in order to show that $s \in f(F)$ for some $F \in \mathcal{F}$, it is sufficient to prove that $\chi_y(\bigcup \mathcal{G}) \subseteq \bigcup \mathcal{G}$.

Assume $t \in \chi_y(\bigcup \mathcal{G})$. Since $\mathcal{G}$ is a directed family and χ is Scott continuous in y, we have $\chi_y(\bigcup \mathcal{G}) = \bigcup \chi_y(\mathcal{G})$. Thus, there exists $F_0 \in \mathcal{F}$ such that $t \in \chi_y(f(F_0))$. Now since χ is continuous in p, there exists a finite set $F_1 \subseteq V(p)$ such that $t \in \chi_y^{F_1}(f(F_0))$. Let F be the set $F_0 \cup F_1$. Since χ is monotone in p, $t \in \chi_y^{F_1}(f(F_0))$ implies $t \in \chi_y^F(f(F_0))$. It also follows from the monotonicity in p that for all $U \subseteq M$, $\chi_y^{F_0}(U) \subseteq \chi_y^F(U)$. Therefore, the least fixpoint $f(F_0)$ of $\chi_y^{F_0}$ is a subset of the least fixpoint $f(F)$ of χ_y^F. Using the fact that χ is monotone in y and the inclusion $f(F_0) \subseteq f(F)$, we obtain $\chi_y^F(f(F_0)) \subseteq \chi_y^F(f(F))$. Putting this together with $t \in \chi_y^F(f(F_0))$, we get $t \in \chi_y^F(f(F))$. Moreover, since $f(F)$

is a fixpoint of χ_y^F, we have $\chi_y^F(f(F)) = f(F)$. Hence, t belongs to $f(F)$. In particular, t belongs to $\bigcup \mathcal{G}$ and this finishes the proof.

We also prove the converse: the sentences in $CF(p)$ are enough to characterize the continuous fragment of the μ-calculus. The proof is inspired by the one given by Marco Hollenberg in [4], where he shows that a sentence is distributive in p over unions iff it is equivalent to $\langle \pi \rangle p$, for some p-free μ-program π.

Theorem 2. *A sentence φ is continuous in p iff it is equivalent to a sentence in $CF(p)$.*

Proof. By Lemma 1, we only need to prove the implication from left to right. Let φ be a sentence continuous in p. We need to find a formula χ in $CF(p)$ that is equivalent to φ.

The proof consists in constructing a finite set $\Pi \subseteq CF(p)$ such that

$$\varphi \equiv \bigvee \{\psi : \psi \in \Pi \text{ and } \psi \models \varphi\}. \tag{1}$$

Indeed, if there is such a set Π, we can define χ as the formula $\bigvee\{\psi : \psi \in \Pi \text{ and } \psi \models \varphi\}$. Clearly, χ belongs to $CF(p)$ and is equivalent to ψ.

We define Π as the set of sentences in $CF(p)$, which correspond to μ-automata with at most k states, where k is a natural number that we will define later and which depends on φ. In order to define k, we introduce the following notation. First, let $\mathbb{A} = (Q, q_0, \delta, \Omega)$ be a μ-automaton corresponding to φ. For $q \in Q$, let φ_q denote the sentence corresponding to the automaton we get from $\mathbb{A}$ by changing the initial state from q_0 to q.

Next we denote by Sort0 be the set of sentences of the form

$$\bigwedge\{p' : p' \in Prop\backslash\{p\}, p' \in \sigma\} \wedge \bigwedge\{\neg p' : p' \in Prop\backslash\{p\}, p' \notin \sigma\},$$

where σ is a subset of $Prop\backslash\{p\}$. For a point s in a model, there is a unique formula in Sort0 true at s. This formula gives us exactly the set of proposition letters in $Prop\backslash\{p\}$ which are true at s. Sort1 is the set of all sentences of the form

$$\bigwedge\{\varphi_q[\bot/p] : q \in S\} \wedge \bigwedge\{\neg\varphi_q[\bot/p] : q \notin S\},$$

where S is a subset of Q. Finally Sort2 contains all sentences of the form $\chi \wedge \nabla \Psi$, where $\chi \in$ Sort0 and Ψ is a subset of Sort1. As for the formulas in Sort0, it is easy to see that given a model $\mathcal{M}$ and a point s in $\mathcal{M}$, there is exactly one sentence in Sort1 and one sentence in Sort2 which are true at s. Remark finally that Sort0, Sort1 and Sort2 are sets of sentences which do not contain p.

Now we can define Π as the set of sentences in $CF(p)$, which correspond to μ-automata with at most $|\text{Sort2}| \cdot 2^{|Q|+1}$ states. Since there are only finitely many such automata modulo equivalence, Π is finite (up to equivalence). It is also immediate that Π is a subset of $CF(p)$. Thus it remains to show that equivalence (1) holds.

From right to left, equivalence (1) is obvious. For the direction from left to right, suppose that $\mathcal{M} = (M, R, V)$ is a model such that $\mathcal{M}, s \Vdash \varphi$, for some point s. We need to find a sentence $\psi \in \Pi$ satisfying $\psi \models \varphi$ and such that $\mathcal{M}, s \Vdash \psi$. Equivalently, we can construct an automaton $\mathbb{A}'$ corresponding to a formula $\psi \in \Pi$ such that $\psi \models \varphi$ and $\mathcal{M}, s \Vdash \psi$. That is, we can construct an automaton $\mathbb{A}'$ with at most $|\mathrm{Sort2}| \cdot 2^{|Q|+1}$ states, corresponding to a sentence in $CF(p)$, such that $\mathbb{A}'$ accepts $(\mathcal{M}, s)$ and satisfying $\mathbb{A}' \Vdash \mathbb{A}$ (that is, for all models $\mathcal{M}'$ and all $s' \in \mathcal{M}'$, if $\mathbb{A}'$ accepts $(\mathcal{M}', s')$, then $\mathbb{A}$ accepts $(\mathcal{M}', s')$).

By Proposition 1, we may assume that $\mathcal{M}$ is a tree with root s and that $\mathcal{M}$ is ω-expanded. Since φ is continuous, there is a finite subset F of $V(p)$ such that $\mathcal{M}[p := F], s \Vdash \varphi$. Let T be the minimal downward closed set that contains F. Using T, we define the automaton $\mathbb{A}'$. Roughly, the idea is to define the set of states of $\mathbb{A}'$ as the set T together with an extra point $a_\top$. However, we need to make sure that the set of states of $\mathbb{A}'$ contains at most $|\mathrm{Sort2}| \cdot 2^{|Q|+1}$ elements. There is of course no guarantee that $T \cup \{a_\top\}$ satisfies this condition. The solution is the following. We define for every point in T its representation, which encodes the information we might need about the point. Then we can identify the points having the same representation in order to “reduce” the cardinality of T.

Before defining the automaton $\mathbb{A}'$, we introduce some notation. Given a point t in $\mathcal{M}[p := F]$, there is a unique sentence in Sort2 that is true at t. We denote it by $s2(t)$. Next if t belongs to F, we define the color $col(t)$ of t as 1 and otherwise, the color of t is 0. We let $Q(t)$ be the set $\{q \in Q : \mathcal{M}[p := F], t \Vdash \varphi_q\}$. Finally, we define the representation map $r : M \to (\mathrm{Sort2} \times Q \times \{0,1\}) \cup \{a_\top\}$ by

$$r(t) = \begin{cases} (s2(t), Q(t), col(t)) & \text{if } t \in T, \\ a_\top & \text{otherwise.} \end{cases}$$

The automaton $\mathbb{A}' = (Q', q'_0, \delta', \Omega')$ is a μ-automaton over the alphabet $\mathrm{Sort2} \times \{0,1\}$. Its set of states Q' is given by

$$Q' = \{r(t) : t \in T\} \cup \{a_\top\},$$

and its initial state q'_0 is $r(s)$. Next for all $(\sigma, i) \in \mathrm{Sort2} \times \{0,1\}$, the set $\delta'(q', (\sigma, i))$ is defined by

$$\delta'(q', (\sigma, i)) = \begin{cases} \big\{r[R(u)] : u \in T \text{ and } r(u) = r(t)\big\} & \\ & \text{if } q' = r_t,\ \sigma = s2(t) \text{ and } i = col(t), \\ \{\{a_\top\}, \emptyset\} & \text{if } q' = a_\top, \\ \emptyset & \text{otherwise.} \end{cases}$$

Intuitively, when the automaton is in the state $q' = r(t)$ and it reads the label $(s2(t), col(t))$, the Duplicator has to pick a successor u of t that is in T and this induces a description in $\delta'(r(t), (s2(t), col(t)))$. As soon as the automaton reaches the state $a_\top$, either the match is finite or the automaton stays in the state $a_\top$. In all other cases, the Duplicator gets stuck.

Finally, the map Ω' is such that $\Omega'(a_\top) = 0$ and $\Omega'(q') = 1$, for all $q' \neq a_\top$. In other words, the only way the Duplicator can win an infinite match is to reach the state $a_\top$ and to stay there.

Remark that a model $\mathcal{M}' = (M', R', V')$ can be seen as a frame (M', R') with a labeling $L' : M' \to \mathrm{Sort2} \times \{0, 1\}$ defined by $L'(t') = (s2(t'), 1)$ if p is true at t' and $L'(t') = (s2(t'), 0)$ otherwise. Thus, the automaton $\mathbb{A}'$ can operate on models.

In order to extract the formula ψ from this automaton, it will be convenient to think of the alphabet of $\mathbb{A}'$ not being the set $\mathrm{Sort2} \times \{0, 1\}$ but the set $\mathcal{P}(\mathrm{Sort2} \cup \{p\})$. The idea is to see a pair $(\sigma, i) \in \mathrm{Sort2} \times \{0, 1\}$ as the set $\{\sigma\}$ if $i = 0$ or as the set $\{\sigma, p\}$ if $i = 1$. More precisely, if $\rho \subseteq \mathrm{Sort2} \cup \{p\}$, the transition map would associate to the pair (q', ρ) the set $\delta'(q', (\sigma, 0))$ if $\rho = \{\sigma\}$ for some $\sigma \in \mathrm{Sort2}$, the set $\delta'(q', (\sigma, 1))$ if $\rho = \{\sigma, p\}$ for some $\sigma \in \mathrm{Sort2}$ and the empty set otherwise.

Now if we think to the formulas of Sort2 as proposition letters, it follows from Theorem 1 that $\mathbb{A}'$ is equivalent to a sentence ψ whose proposition letters belong to $\mathrm{Sort2} \cup \{p\}$. Such a formula ψ is also a sentence with proposition letters in *Prop*, in an obvious way. To finish the proof, we need to show that ψ is equivalent to a sentence which is in Π, ψ is true at s and $\psi \models \varphi$.

Claim 1. ψ is equivalent to a sentence in Π.

Proof. The intuition is the following. In order to win an $\mathbb{A}'$-match, the Duplicator has to reach the state $a_\top$ and then, the match is basically over. It seems natural that such a property can be expressed using only least fixpoints (and no greatest fixpoint).

Next we also need to make sure that in a formula corresponding to $\mathbb{A}'$, neither p nor any variable is in the scope of the operator $\Box$. This is guaranteed by the presence of the state $a_\top$ in any non-empty description that the Duplicator might pick. Very informally, each description corresponds to a subformula (of the sentence corresponding to the automaton) which starts with the operator ∇. Using the fact that $a_\top$ belongs to any of these descriptions (except the empty one) and corresponds to the sentence $\top$, we can show that the ∇ operator can be replaced by the modal operator $\Diamond$.

Formally the proof is the following. First observe that $\mathbb{A}'$ has at most $|\mathrm{Sort2}| \cdot 2^{|Q|+1}$ states. Thus in order to show that ψ is equivalent to a formula in Π, it is sufficient to show that ψ is equivalent to a sentence in $CF(p)$.

For $q' \in Q'$ and $S' \subseteq Q'$, we define the translation $tr(S', q')$ of q' with respect to S'. The translation $tr(S', q')$ is a formula in the language whose set of proposition letters is *Prop* and whose set of variables is $Var \cup Q'$. For those q' that are equal to $r(t) = (s2(t), Q(t), col(t))$ and $S' \subseteq Q'$, we have

$$
\begin{aligned}
tr(S', q') := {} & s2(t) \wedge col(t).p \wedge \\
& \bigvee \Big\{ \bigwedge \Big\{ \Diamond \mu q''.tr\,(S' \backslash \{q''\}, s') : q'' \in r[R(u)] \text{ and } q'' \in S' \Big\} \\
& \quad \wedge \bigwedge \Big\{ \Diamond q'' : q'' \in r[R(u)] \text{ and } q'' \notin S' \Big\} : u \in T \text{ and } r(u) = q' \Big\}
\end{aligned}
$$

where $col(t).p$ is p if $col(t) = 1$ and $\top$ if $col(t) = 0$. By convention, $\bigwedge \emptyset = \top$. For all $S' \subseteq Q'$, we define $tr(S', a_\top)$ by $\top$.

It is routine to show that $tr(S', q')$ is a well-defined sentence with proposition letters in $Prop \cup (Q' \backslash S')$ and that it belongs to $CF(\{p\} \cup (Q' \backslash S'))$. The proofs are by induction on the cardinality of S'. In particular, $tr(Q', q'_0)$ belongs to $CF(p)$. Therefore, in order to prove the claim, it is enough to show that ψ is equivalent to $tr(Q', q'_0)$. The proof is in the appendix.

Claim 2. $\mathcal{M}, s \Vdash \psi$.

Proof. The proof is rather straightforward. Details are given in the appendix.

Claim 3. $\psi \models \varphi$.

Proof. Suppose $\mathcal{M}' = (M', R', V')$ is a model such that $\mathcal{M}', s' \Vdash \psi$, for some point s'. That is, the Duplicator has a winning strategy in the $\mathbb{A}'$-game in $\mathcal{M}'$ with starting position (s', q'_0). We have to show that $\mathcal{M}', s' \Vdash \varphi$. That is, we need to find a winning strategy for the Duplicator in the $\mathbb{A}$-game in $\mathcal{M}'$ with starting position (s', q'_0).

We say that a point t' is marked with a state q' if there is an $\mathbb{A}'$-match during which the Duplicator plays according to his winning strategy and the point t' is marked with q'. Let T' be the set of points marked with a state $q' \neq a_\top$. When we define the strategy for the Duplicator in the $\mathbb{A}$-game, the idea is roughly to make sure that if $t' \in T'$ and $q'(t') = r(t)$, then positions of the form (t', q) are played only if $q \in Q(t)$. Details are given in the appendix.

As a corollary of this last proof, we obtain that it is decidable whether a formula is continuous in p.

Theorem 3. *It is decidable whether a formula is continuous in p.*

Proof. Fix a proposition letter p. Let Π be the set of sentences in $CF(p)$ which correspond to μ-automata with at most $|\mathrm{Sort2}| \cdot 2^{|Q|+1}$ states. Now there are only finitely many such automata (modulo equivalence). There is also an effective translation from μ-automata to μ-sentences. Finally it is easy to verify whether a formula is in $CF(p)$. Therefore, we can compute Π.

It follows from the proof of Theorem 2 that a sentence φ is continuous in p iff $\varphi \equiv \{\psi : \psi \in \Pi \text{ and } \psi \models \varphi\}$. That is, φ is continuous in p iff there exists a subset Ψ of Π such that $\varphi \equiv \bigvee \Psi$. Therefore, in order to decide if φ is continuous in p, we can compute all the subsets Ψ of Π and check whether φ is equivalent to a disjunct $\bigvee \Psi$. Since the μ-calculus if finitely axiomatizable and has the finite model property, it is decidable whether φ is equivalent to a disjunct $\bigvee \Psi$ and this completes the proof.

Looking at the decision procedure presented in the proof of Theorem 3, we can see that the complexity is at most 4EXPTIME. That is, it involves four interlocked checking procedures, each of them being of complexity at most EXPTIME. This result is not very satisfying and we are looking for a better algorithm.

Finally, we mention that a similar syntactic characterization can be obtained in the case of basic modal logic. More precisely, a basic modal formula is continuous in p iff it belongs to the modal fragment $CF_m(p)$ of $CF(p)$. We give a formal definition of $CF_m(p)$ and a sketch of the proof in the appendix.

Definition 12. *Let P be a subset of Prop. $CF_m(P)$ is defined by induction in the following way:*

$$\varphi ::= \top \mid p \mid \psi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond \varphi,$$

where p belongs to P and no proposition letters of ψ is in P.

Corollary 1. *A basic modal formula is continuous in p iff it belongs to $CF_m(p)$.*

5 Conclusion and further work

We defined the continuous fragment of the μ-calculus and showed how it relates to Scott continuity. We also started to investigate the relation between continuity and constructivity. Finally, we gave a syntactic characterization of the continuous formulas and we proved that it decidable whether a formula is continuous.

This work can be continued in various directions. To start with, it would be interesting to clarify the link between continuity and constructivity. In particular, we could try to answer the following question: given a constructive formula φ, can we find a continuous formula ψ satisfying $\mu p.\varphi \equiv \mu p.\psi$?

Next we observe that in the proof of Theorem 2, the construction of the automaton $\mathbb{A}'$ depends on the model $\mathcal{M}$ and the point s at which φ is true. Is it possible to construct an automaton $\mathbb{A}'$ by directly transforming the automaton $\mathbb{A}$ that is equivalent to φ? Such a construction might help us to find a better lower complexity bound for the decision procedure (for the membership of a formula in the continuous fragment).

We believe that it might be interesting to generalize our approach. As mentioned earlier, similar results to our characterization have been obtained by Giovanna D'Agostino and Marco Hollenberg in [3]. Is there any general pattern that can be found in all these proofs?

We could also extend this syntactic characterization to other settings. For example, we can try to get a similar result if we restrict our attention to the class of finitely branching models.

Finally, we would like to mention that in [10], Daisuke Ikegami and Johan van Benthem proved that the μ-calculus is closed under taking product update. Using their method together with our syntactic characterization, it is possible to show that the set of continuous formulas is closed under taking product update.

Acknowledgments. I am grateful to Johan van Benthem for proposing this work, interesting comments and suggestions for the introduction. Many thanks also to my supervisor Yde Venema for his valuable help and support. Finally special thanks to Balder ten Cate for comments and inspiring emails.

References

1. Abramsky, S., Jung, A.: Domain Theory. In S., A., M., G.D., T.S.M., M., eds.: Handbook for Logic in Computer Science. Clarendon Press, Oxford (1994)

2. Chang, C., Keisler, H.: Model Theory. North Holland (1973)
3. D'Agostino, G., Hollenberg, M.: Logical Questions Concerning The mu-Calculus: Interpolation, Lyndon and Los-Tarski. Journal of Symbolic Logic (2000) 310–332
4. Hollenberg, M.: Logic and Bisimulation. PhD thesis, Utrecht University (1998) Zeno Institute of Philosophy.
5. van Benthem, J.: Exploring Logical Dynamics. Studies in Logic, Language and Information. CSLI Publications and Cambridge University Press (1996)
6. Janin, D., Walukiewicz, I.: Automata for the modal μ-calculus and related results. In: Proceedings MFCS'95, Springer (1995) 552–562
7. Gierz, G., Hofman, K., Keimel, K., Lawson, J., M., M., Scott, D.: A Compendium of Continuous Lattices. Springer-Verlag (1980)
8. Otto, M.: Eliminating Recursion in the mu-Calculus. In: Proceedings of STACS'99, Springer (1999)
9. van Benthem, J.: Modal Frame Correspondences and Fixed-Points. Studia Logica (2006) 133–155
10. Ikegami, D., van Benthem, J.: Modal Fixed-Point Logic and Changing Models (2008) ILLC Preprint PP-2008-19.

Appendix

Proof of Claim 1

Before giving details for the proof of Claim 1 in Theorem 2, we recall the rules of the evaluation game.

Definition 13. *Let φ be a sentence such that $\neg\psi$ is a subformula only if ψ is a proposition letter or the formula $\top$. We also assume that each variable is bound by at most one fixpoint. Thus, for every variable x occurring in φ, there is a unique formula ψ_x such that $\mu x.\psi_x$ or $\nu x.\psi_x$ is a subformula of φ.*

We fix a model $\mathcal{M} = (M, R, V)$ and a point $s \in M$. We define the evaluation game *for the sentence φ in the model $\mathcal{M}$ with starting position (s, φ) as the following game. The game is played between two players, the Duplicator and the Spoiler. The starting position is (s, φ). The rules, determining the admissible moves, together with the player who is supposed to make a move, are given in the table above.*

If the match is finite, the player who gets stuck looses. If the match is infinite, we let Inf be the set of variables x such that positions of the form (t, x) are reached infinitely often. Let x_0 be a variable in Inf such that for all variables y in Inf, ψ_y is a subformula of ψ_{x_0}. If x_0 is bound by a μ-operator, then the Spoiler wins the match. Otherwise the Duplicator wins.

We recall that the Duplicator has a winning strategy in the evaluation game in $\mathcal{M}$ with starting position (s, φ) iff $\mathcal{M}, s \Vdash \varphi$.

Claim. ψ is equivalent to a sentence in Π.

Position	Player	Admissible moves
$(t, \varphi_1 \vee \varphi_2)$	Duplicator	$\{(t, \varphi_1), (t, \varphi_2)\}$
$(t, \varphi_1 \wedge \varphi_2)$	Spoiler	$\{(t, \varphi_1), (t, \varphi_2)\}$
$(t, \Diamond\psi)$	Duplicator	$\{(u, \psi) : tRu\}$
$(t, \Box\psi)$	Spoiler	$\{(u, \psi) : tRu\}$
(t, x)	-	$\{(t, \psi_x)\}$
$(t, \mu x.\psi)$	-	$\{(t, \psi)\}$
$(t, \nu x.\psi)$	-	$\{(t, \psi)\}$
$(t, \neg\top)$ or $[(t, p)$ with $t \notin V(p)]$ or $[(t, \neg p)$ with $t \in V(p)]$	Duplicator	$\emptyset$
$(t, \neg p)$ with $t \in V(p)$	Duplicator	$\emptyset$
$(t, \top)$ or $[(t, p)$ with $t \in V(p)]$ or $[(t, \neg p)$ with $t \notin V(p)]$	Spoiler	$\emptyset$
$(t, \neg p)$ with $t \notin V(p)$	Duplicator	$\emptyset$

Proof. Since $\mathbb{A}'$ has at most $|\mathrm{Sort2}| \cdot 2^{|Q|+1}$ states, it is enough to show that ψ is equivalent to a sentence in $CF(p)$.

Recall that for all $q' = r(t)$ in Q' and all $S' \subseteq Q'$, we define the translation $tr(S', q')$ by

$$tr(S', q') := s2(t) \wedge col(t).p \wedge \bigvee \Big\{ \bigwedge \Big\{ \Diamond \mu q''.tr\,(S' \backslash \{q''\}, s') : q'' \in r[R(u)] \text{ and } q'' \in S' \Big\} \wedge \bigwedge \Big\{ \Diamond q'' : q'' \in r[R(u)] \text{ and } q'' \notin S' \Big\} : u \in T \text{ and } r(u) = q' \Big\}$$

where $col(t).p$ is p if $col(t) = 1$ and $\top$ if $col(t) = 0$. By convention, $\bigwedge \emptyset = \top$. For all $S' \subseteq Q'$, we define $tr(S', a_\top)$ by $\top$.

Finally, we recall that for all $q' \in Q'$ and all $S' \subseteq Q'$, $tr(S', q')$ is a well-defined sentence with proposition letters in $Prop \cup (Q' \backslash S')$ and that it belongs to $CF(\{p\} \cup (Q' \backslash S'))$. In particular, $tr(Q', q'_0)$ belongs to $CF(p)$. Therefore, in order to prove the claim, it is enough to show that ψ is equivalent to $tr(Q', q'_0)$.

Before proving that ψ is equivalent to formula in $CF(p)$, we show that the translations satisfy the following property. For all $q' \in Q'$,

$$tr(Q', q') \text{ is equivalent to } \mu q'.tr(Q' \backslash \{q'\}, q'). \tag{2}$$

We give a sketch of the proof of (2). First, we observe that for all $S' \subseteq Q'$ and all $q' \notin S'$, we have

$$tr(S', q') = tr(S' \backslash \{q'\}, q')[\mu q'.tr(S' \backslash \{q'\}, q')/q']. \tag{3}$$

We skip the proof which is a standard proof by induction on the cardinality of S'. In order to prove (2), fix a state $q' \in Q'$. Now let χ be the formula $tr(Q' \backslash \{q'\}, q')$. It follows from the equality (3) that $tr(Q', q')$ is equivalent to $tr(Q' \backslash \{q'\}, q')[\mu q'.tr(Q' \backslash \{q'\}, q')/q']$. That is, $tr(Q', q')$ is equivalent to $\chi[\mu q'.\chi/q']$. Next, by definition of the fixpoint operator, we know that $\chi[\mu q'.\chi/q']$ is equivalent to $\mu q'.\chi$. Putting everything together, we obtain that $tr(Q', q')$ is equivalent to $\mu q'.\chi$. That is, $tr(Q', q')$ is equivalent to $\mu q'.tr(Q' \backslash \{q'\}, q')$ and this finishes the proof of (2).

Now we prove that ψ is equivalent to $tr(Q', q'_0)$. For let $\mathcal{M} = (M', R', V')$ be a model and let s' be a point in $\mathcal{M}'$. We need to show that

$$\mathcal{M}', s' \Vdash \psi \leftrightarrow tr(Q', q'_0). \tag{4}$$

By Proposition 1, we may assume that $\mathcal{M}'$ is a tree with root s', that is ω-expanded.

For the direction from left to right of (4), suppose that $\mathcal{M}', s' \Vdash \psi$. We know that the Duplicator has a winning strategy g in the $\mathbb{A}'$-game in $\mathcal{M}'$ with starting position (s', q'_0). Given a winning strategy f, we say that a point t' is marked with a state $q' \in Q'$ if there is a f-conform $\mathbb{A}'$-match (in $\mathcal{M}'$ with starting position (s', q'_0)) during which the point t' is marked with q'. Now we define a winning strategy g' (for the Duplicator in the $\mathbb{A}'$-game in $\mathcal{M}'$ with starting position $(s'q'_0)$) such that for all $t' \in \mathcal{M}'$, t' is marked with exactly one state $q'(t')$ and the set of successors marked with a state $q' \neq a_\top$ is finite.

We construct g' by induction on the "distance" to the root s'. It is immediate that s' is only marked with q'_0. Now assume that t' is only marked with q'. Thus there is an $\mathbb{A}'$-match where the Duplicator plays according to his winning strategy and during which the position (t', q') is reached. Then the Duplicator chooses a description D and a marking m. First we can modify the marking such that for each state $q'' \neq a_\top$, exactly one point is marked with q''. Next suppose that a successor u' of t' is marked with q'_1 and q'_2. Since $\mathcal{M}'$ is ω-expanded, we can pick a successor v' of t' that is bisimilar to u' and only marked with $a_\top$. We can then modify the marking such that u' is marked with q'_1 and v' is only marked with q'_2. It is not hard to see that this is still a winning strategy for the Duplicator and that, according to this strategy, every successor of t' is marked with a unique state and only finitely many of these successors are marked with a state $q'' \neq a_\top$. This completes the definition of g'.

Now in order to show the left to right implication of (4), we have to prove that $\mathcal{M}', s' \Vdash tr(Q', q'_0)$. The idea is to show that if t' is marked with q', then $\mathcal{M}', t' \Vdash tr(Q', q')$. In particular, this would imply that $\mathcal{M}', s' \Vdash tr(Q', q'_0)$, since s' is marked with q'_0. Thus, it is sufficient to prove that if t' is marked with q', then $\mathcal{M}', t' \Vdash tr(Q', q')$.

The proof is by induction on the distance $d_\top(t')$ that we define in the following way. For all $t' \in \mathcal{M}'$,

$$d_\top(t') := \begin{cases} 0 & \text{if } q'(t') = a_\top, \\ max\{d_\top(v') : t'R'v'\} + 1 & \text{otherwise.} \end{cases}$$

Remark that since the set of points marked with a state $q' \neq a_\top$ is finite, we have that $d_\top(t')$ is a natural number for all $t' \in \mathcal{M}'$.

For the basic case, we check that if $d_\top(t') = 0$ and t' is marked with q', then $\mathcal{M}', t' \Vdash tr(Q', q')$. This is immediate since $q' = a_\top$ and $tr(Q', a_\top) = \top$.

For the induction step, we fix a point t' marked with a state $r(t)$ $(t \in T)$ and we assume that for all v' such that $d_\top(v') < d_\top(t')$, we have $\mathcal{M}', v' \Vdash tr(Q', q')$

if v' is marked with q'. In particular, for all successors v' of t', we have $\mathcal{M}', v' \Vdash tr(Q', q')$ if v' is marked with q'. We need to show that $\mathcal{M}', t' \Vdash tr(Q', r(t))$. Since t' is marked with $r(t)$ (and hence the Duplicator cannot get stuck in position $(t', r(t)))$, we have $\delta'(r(t), L'(t')) \neq \emptyset$. Thus, it follows from the definition of δ' that $\mathcal{M}', t' \Vdash s2(t) \wedge col(t).p$. It remains to find $u \in T$ such that $r(u) = r(t)$ and such that $\bigwedge\{\Diamond \mu s'.tr\,(Q' \backslash \{s'\}, s') : s' \in r[R(u)]\}$ is true at t'.

Since t' is marked with $r(t)$, we know that if the Duplicator plays according to his winning strategy in an $\mathbb{A}'$-match, then the position $(t', r(t))$ can be reached. In the next move of the match, the Duplicator chooses a description in $\delta'(r(t), L'(t'))$. Let $u \in T$ be such that the Duplicator chooses the description $r[R(u)]$. It follows from the definition of δ' that $r(u) = r(t)$.

Next we prove that $\bigwedge\{\Diamond \mu q''.tr\,(Q' \backslash \{q''\}, q'') : q'' \in r[R(u)]\}$ is true at t'. For let q'' be a state in $r[R(u)]$. We have to show that $\mathcal{M}', t' \Vdash \Diamond \mu q''.tr\,(Q' \backslash \{q''\}, q'')$. Recall that by (2), we have that $\mu q''.tr\,(Q' \backslash \{q''\}, q'')$ is equivalent to $tr(Q', q'')$. Thus it is enough to check that $\mathcal{M}', t' \Vdash \Diamond tr(Q', q'')$. Since $r[R(u)]$ is the description chosen by the Duplicator, there exists a successor v' of t' that is marked with q''. By induction hypothesis, we know that $\mathcal{M}', v' \Vdash tr(Q', q'')$. It immediately follows that $\mathcal{M}', t' \Vdash \Diamond tr(Q', q'')$ and this finishes the proof from the left to right implication of (4).

For the converse direction of (4), assume that $\mathcal{M}', s' \Vdash tr(Q', q'_0)$. Thus, the Duplicator has a winning strategy h in the evaluation game with starting position $(s', tr(Q', q'_0))$. We say that a point t' is h-marked with a state q' if there is an h-conform evaluation game during which the Duplicator plays the position $(t', tr(S', q'))$ for some $S' \subseteq Q'$. Since $\mathcal{M}'$ is ω-expanded, we may assume that the strategy h is such that each point is h-marked with at most one state. We do not give details, as this is similar to the transformation from the strategy g to the strategy g'.

We define by induction a winning strategy for the Duplicator in the $\mathbb{A}'$-game in $\mathcal{M}'$ with starting position (s', q'_0). The idea is to ensure that if a point t' is marked with $q' \neq a_\top$ (in an $\mathbb{A}$'-match conform to our strategy), then t' is h-marked with q' in the evaluation game. The starting position of the $\mathbb{A}'$-game is (s', q'_0). It is immediate that s' is h-marked with q'_0, since any evaluation match starts with the position $(s', tr(Q', q'_0))$.

Suppose that we have defined the strategy until the $\mathbb{A}'$-match reaches the position (t', q'), where $q' \neq a_\top$ and t' is h-marked with q'. Thus there exists $t \in T$ such that $q' = r(t)$. Since t' is h-marked with q', there is $S' \subseteq Q'$ and an h-conform evaluation match during which we reach the position $(t', tr(S', q'))$. In particular, this position is winning for the Duplicator. By definition of the translation tr, we have that $\mathcal{M}', t' \Vdash s2(t) \wedge col(t)$. Thus, $\delta'(q', L'(t'))$ is $\{r[R(u)] : u \in T \text{ and } r(u) = q'\}$ and the Duplicator has to choose a description in this set. Since the position $(t', tr(S', q'))$ is played by the Duplicator in an h-conform evaluation match, we know that there exists $u \in T$ such that $r(u) = q'$ and the next position of this h-conform evaluation match is

$$(t', \bigwedge \Big\{ \Diamond \mu q''.tr\,(S' \backslash \{q''\}, s') : q'' \in r[R(u)] \text{ and } q'' \in S' \Big\} \wedge \\ \bigwedge \Big\{ \Diamond q'' : q'' \in r[R(u)] \text{ and } q'' \notin S' \Big\}). \quad (5)$$

Now we define the strategy such that the Duplicator chooses the description $r[R(u)]$. Next we need to define a marking that is legal with respect to this description. Fix a state q'' in the chosen description. That is, q'' belongs to $r[R(u)]$. It follows from (5) that there is an h-conform evaluation match during which either the position $(t', \Diamond \mu q''.tr(S' \backslash \{q''\}, q''))$ or the position $(t', \Diamond q'')$ is reached. Thus, there exists a successor v' of t' such that the next position of this h-conform match is $(v', \mu q''.tr(S' \backslash \{q''\}, q''))$ or (v', q''). This implies that one of the next positions of this h-conform match is $(v', tr(S'', q''))$ for some $S'' \subseteq Q'$. Therefore, we can define our strategy in this $\mathbb{A}'$-match such that the point v' is marked with the state q''. Next observe that since T is finite and $\mathcal{M}$ is ω-expanded, there are successors of u which do not belong to T. In particular, $a_\top$ belongs to $r[R(u)]$. Therefore we can also mark with $a_\top$ all the successors of t'. It is easy to check that our marking is a legal marking.

It remains to show that this strategy is indeed winning. Since the formula $tr(Q', q'_0)$ contains only least fixpoints, the Duplicator can only wins finite matches in the evaluation game (with starting position $(s', tr(Q', q'_0)))$. Thus, on every path starting from s', there are only finitely many points h-marked with a state in Q'. This implies that in any $\mathbb{A}'$-match conform to our strategy, we eventually reached the position $(t', a_\top)$ and the Duplicator wins the game.

Proof of Claim 2

Claim. $\mathcal{M}, s \Vdash \psi$.

Proof. We need to provide a winning strategy for the Duplicator in the $\mathbb{A}'$-game in the model $\mathcal{M}$ with starting position (s, q_0). The strategy is defined by induction and we ensure that whenever a position (t, q') is played, then $q' = a_\top$ or $q' = r(t)$.

This certainly holds for the initial position (s, q_0). Now assume that the Duplicator has to respond to a position (t, q'). Assume first that $q' = a_\top$. If t has at least one successor, the Duplicator chooses the description $\{a_\top\}$ and he marks all the successors of t with $a_\top$. If t has no successor, the Duplicator picks the description $\emptyset$ and the match is over.

Now if $t \in T$ and $q' = r(t)$, the Duplicator chooses the description $r[R(t)]$. A successor v of t that belongs to T is marked with $r(v)$. All the successors of t are marked with $a_\top$. It is routine to show that such a strategy is well-defined and winning.

Proof of Claim 3

Claim. $\psi \models \varphi$.

Proof. Suppose $\mathcal{M}' = (M', R', V')$ is a model such that $\mathcal{M}', s' \Vdash \psi$, for some point s'. That is, the Duplicator has a winning strategy in the $\mathbb{A}'$-game in $\mathcal{M}'$ with starting position (s', q'_0). We have to show that $\mathcal{M}', s' \Vdash \varphi$.

As before, we may assume that $\mathcal{M}'$ is a tree with root s' and that $\mathcal{M}'$ is ω-expanded. Recall from Claim 1 that we say that a point t' is marked with a state q' if there is an $\mathbb{A}'$-match during which the Duplicator plays according to his winning strategy and the point t' is marked with q'. As in Claim 1, each point t' of $\mathcal{M}'$ may assumed to be marked with a unique state $q'(t')$ of $\mathbb{A}'$ and given a point t', we may suppose that the set of successors marked with a state $q' \neq a_\top$ is finite.

Let T' be the set of points marked with a state $q' \neq a_\top$. Let F' be the set of $t' \in T'$ such that $q'(t')$ is of the form $(s2(t), Q(t), 1)$ for some $t \in T$. For v' in $\mathcal{M}'$, we let $col(v')$ be 1 if $v' \in F'$ and $col(v') = 0$ otherwise. We also define $L'(V')$ as the pair $(s2(v'), col(v'))$ (where $s2(v')$ is the only sentence in Sort2 which is true at v' in $\mathcal{M}'[p := F']$).

It is routine to check that the strategy for the Duplicator in the model $\mathcal{M}'$ (with starting position (s', q'_0)) remains a winning strategy in $\mathcal{M}'[p := F']$ (with starting position (s', q'_0)). Moreover since the Duplicator cannot get stuck in position $(t', q'(t'))$, we have $\delta'(q'(t'), L'(t')) \neq \emptyset$. Hence, if $q'(t') = r(t)$ for some $t \in T$, it follows from the definition of δ' that $L'(t') = (s2(t), col(t))$. In particular, if $q'(t') = (s2(t), Q(t), 1)$ for some $t \in T$, then $col(t') = 1$. In other words, F' is a subset of $V'(p)$. Since φ is monotone, it is then enough to prove $\mathcal{M}'[p := F'], s' \Vdash \varphi$ in order to show that $\mathcal{M}', s' \Vdash \varphi$.

To prove this, we need to find a winning strategy for the Duplicator in the $\mathbb{A}$-game in $\mathcal{M}'[p := F']$ with starting position (s', q_0). The idea is to make sure that if $t' \in T'$ and $q'(t') = r(t)$, then positions of the form (t', q) are played only if $q \in Q(t)$. This holds for the initial position (s', q_0), as $q'(s') = r(s)$ and $q_0 \in Q(s)$ (since $\varphi_{q_0} = \varphi$ and φ is true at s in $\mathcal{M}[p := F]$). We will see that following our strategy, as soon as a position (t', q) is reached with $t' \notin T'$, then the Duplicator can win.

Suppose that the Duplicator has to respond to a position (t', q) with $t' \in T'$, $q'(t') = r(t)$ and $q \in Q(t)$. By definition of the map q', there is some match of the $\mathbb{A}'$-game (in the model $\mathcal{M}'[p := F']$ with starting position (s', q'_0)) which is conform to the Duplicator's strategy and during which the position $(t', r(t))$ is reached. Following his winning strategy, the Duplicator has then to choose a description in $\delta(r(t), L'(t'))$. Let $u \in T$ be a point such that the Duplicator chooses the description $r[R(u)]$. Remark that by definition of δ', we have $r(t) = r(u)$. Putting this together with $q \in Q(t)$, we obtain that q belongs to $Q(u)$. Since $q \in Q(u)$, we have $\mathcal{M}[p := F], u \Vdash \varphi_q$. Therefore, the Duplicator has a winning strategy for the position (u, q) (in the $\mathbb{A}$-game in the model $\mathcal{M}[p := F]$).

Let m be a marking chosen according to this strategy. Suppose m is legal with respect to $D \in \delta(q, L(u))$, where $L(u) = \{p' \in Prop : \mathcal{M}[p := F], u \Vdash p'\}$. Since the position $(t', r(t))$ is winning for the Duplicator, he cannot get stuck. Hence, $\delta'(r(u), L'(t')) \neq \emptyset$. It follows from the definition of δ' that t' and u satisfy the same sentence in Sort2. In particular, they satisfy the same propositions letters of $Prop \backslash \{p\}$. Moreover, by definition of F', we also have $t' \in F'$ iff $u \in F$. Putting everything together, we obtain that the sets $L(t') = \{p' \in Prop : \mathcal{M}'[p := F'], t' \Vdash p'\}$ and $L(u)$ are the same. Thus, $\delta(q, L(t')) = \delta(q, L(u))$ and the description D is also available in $\delta(q, L(t'))$.

We define the marking $\bar{m}$ in response to (t', q) in the following way:

$$\begin{aligned}\bar{m}(\bar{q}) = \{&v' \notin T' : t'R'v' \text{ and } \mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}\} \cup \\ &\{v' \in T' : t'R'v', q'(v') = r(v) \text{ and } \bar{q} \in Q(v)\},\end{aligned}$$

for all $\bar{q} \in Q$. We show that $\bar{m}$ is a legal marking with respect to D.

For suppose that $\bar{q} \in D$. We need to find a successor v'_0 of t' such that $v'_0 \in \bar{m}(\bar{q})$. Since m is a legal marking with respect to D, there exists a successor v_0 of u such that $v_0 \in m(\bar{q})$. There are two cases.

First, suppose that v_0 belongs to T. Recall that following his winning strategy, the Duplicator chooses the description $r[R(u)]$ at position $(t', r(u))$. So there is a successor v'_0 of t' that is marked with $r(v_0)$. Since v_0 belongs to $m(\bar{q})$, we have $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}$. That is, $\bar{q} \in Q(v_0)$. Gathering everything together, we have that $v'_0 \in T'$ is a successor of t' such that $q'(v'_0) = r(v_0)$ and $\bar{q} \in Q(v_0)$. Thus, $v'_0 \in \bar{m}(\bar{q})$ and we are done.

Next suppose that v_0 does not belong to T. Thus in $\mathcal{M}[p := F]$, there is no point from v_0 on where p holds. In particular, $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}$ implies $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}[\bot/p]$. Since $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}[\bot/p]$, we have $\mathcal{M}[p := F], u \Vdash \Diamond\varphi_{\bar{q}}[\bot/p]$. As u and t' satisfy the same sentences in Sort2, it follows that $\mathcal{M}'[p := F'], t' \Vdash \Diamond\varphi_{\bar{q}}[\bot/p]$. Thus there is a successor v' of t' such that $\mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}[\bot/p]$. Since $\varphi_{\bar{q}}[\bot/p]$ does not contain any p, this also means that $\mathcal{M}', v' \Vdash \varphi_{\bar{q}}[\bot/p]$.

Next observe that by construction of T' and by definition of the parity condition Ω', T' is finite. As $\mathcal{M}'$ is ω-expanded, we can choose a successor v'_0 of t' that is bisimilar in $\mathcal{M}'$ to v' and such that $v'_0 \notin T'$. Putting this together with $\mathcal{M}', v' \Vdash \varphi_{\bar{q}}[\bot/p]$, we obtain $\mathcal{M}', v'_0 \Vdash \varphi_{\bar{q}}[\bot/p]$. Using again the fact that $\varphi_{\bar{q}}[\bot/p]$ does not contain any p, this implies $\mathcal{M}'[p := F'], v'_0 \Vdash \varphi_{\bar{q}}[\bot/p]$. Now remark that by definition of δ', T' is downward closed. In particular, since $v'_0 \notin T'$, no point in the model generated by v'_0 belongs to T'. It follows that in the submodel of $\mathcal{M}'[p := F']$ generated by v'_0, p holds nowhere. Therefore, $\mathcal{M}'[p := F'], v'_0 \Vdash \varphi_{\bar{q}}[\bot/p]$ implies $\mathcal{M}'[p := F'], v'_0 \Vdash \varphi_{\bar{q}}$. Hence, v'_0 belongs to $\bar{m}(\bar{q})$.

To prove that $\bar{m}$ is a legal marking with respect to D, it remains to show that for all successors v' of t', there is a state $\bar{q} \in D$ such that $v' \in \bar{m}(\bar{q})$. Let v' be a successor of t'. There are again two cases.

First, suppose that v' belongs to T'. Thus, v' is marked with a state $r(v)$, for some successor v of u. Since m is a legal marking with respect to D, there is a state $\bar{q} \in D$ such that $v \in m(\bar{q})$. That is, $\bar{q} \in Q(v)$. By definition of $\bar{m}$, this means that $v' \in \bar{m}(\bar{q})$.

Finally assume that v' does not belong to T'. Let $s1(v')$ be the unique sentence in Sort1 that is true at v' in $\mathcal{M}'[p := F']$. Hence, $\Diamond s1(v')$ is true at t' in $\mathcal{M}'[p := F']$. Since u and t' satisfy the same sentence in Sort2, we get that $\mathcal{M}[p := F], u \Vdash \Diamond s1(v')$. Thus there is a successor v of u such that $\mathcal{M}[p := F], v \Vdash s1(v')$. As $s1(v')$ does not contain any p, we also have that $\mathcal{M}, v \Vdash s1(v')$. Since T is finite and $\mathcal{M}$ is ω-expanded, we can choose a successor v_0 of u that does not belong to T and that is bisimilar in $\mathcal{M}$ to v. In particular, $\mathcal{M}, v \Vdash s1(v')$ implies $\mathcal{M}, v_0 \Vdash s1(v')$. Using again the fact that $s1(v')$ does not contain any p, we get $\mathcal{M}[p := F], v_0 \Vdash s1(v')$.

Now since m is a legal marking with respect to D, there exists $\bar{q} \in D$ such that $v_0 \in m(\bar{q})$. That is, $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}$. As before, we can show that in the submodel of $\mathcal{M}[p := F]$ generated by v_0, p holds nowhere. Therefore, $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}$ implies $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}[\bot/p]$. As $\mathcal{M}[p := F], v_0 \Vdash s1(v')$, v' and v_0 satisfy the same sentence in Sort1. Thus, from $\mathcal{M}[p := F], v_0 \Vdash \varphi_{\bar{q}}[\bot/p]$, we get $\mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}[\bot/p]$. Moreover, since $v' \notin T'$, we know that in the submodel of $\mathcal{M}'[p := F']$ generated by v', p does not hold. Thus $\mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}$ and $v' \in \bar{m}(\bar{q})$.

The Spoiler may respond to $\bar{m}$ in two ways. First, he may pick a position $(v', \bar{q})$ with $v' \notin T'$ and $\mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}$. Then the Duplicator has a winning strategy from this point on. We continue with this strategy.

Next, the Spoiler may pick a position $(v', \bar{q})$ with $v' \in T'$, $q'(v') = r(v)$ and $\bar{q} \in Q(v)$. Then we continue with the strategy we have described here. Recall now that by definition of the parity condition Ω', T' is finite. Therefore, in any match played according to our strategy, the Spoiler will end up picking a position $(v', \bar{q})$ with $v' \notin T'$ and $\mathcal{M}'[p := F'], v' \Vdash \varphi_{\bar{q}}$. The Duplicator can then win and this completes the proof.

Proof of Corollary 1

Corollary. *A basic modal formula is continuous in p iff it belongs to $CF_m(p)$.*

Proof. We focus on the direction from left to right. Let φ be a basic modal formula continuous in p. It follows from Theorem 2 that φ is equivalent to a μ-sentence ψ in $CF(p)$. Remark that if we look carefully at the proof of Theorem 2, we can see that ψ is guarded (that is, each variable x in ψ is in the scope of a modal operator). We may also assume that each variable x in ψ is bound at most once in the formula. Thus for a variable x in ψ, there exists a unique subformula of the form $\mu x.\alpha_x$.

Recall that the dependency order $\preceq$ on the bound variables of ψ is the least partial order such that if x occurs free in $\mu y.\alpha_y$, then $x \preceq y$. Now we define the formula ψ^i ($i \in \mathbb{N}$) by induction on i. We let ψ^0 be the formula obtained by deleting all the μ-operators in ψ and we let ψ^{i+1} be the formula $\psi^i[\alpha_{x_n}/x_n] \dots [\alpha_{x_1}/x_1]$, where the sequence $x_1, \dots, x_n$ is a linear ordering of all bound variables of ψ such that if $x_i \preceq x_j$, then $i \leq j$.

Let n be the modal depth of φ. Consider the formula ψ^n. Now let χ be the basic modal formula obtained by replacing all the variables by $\top$. It is clear that χ is a basic modal formula in $CF_m(p)$. We skip the proof but using the fact that ψ is guarded and equivalent to a basic modal formula of depth n, we can show that χ is equivalent to ψ. Thus we found a basic modal formula χ in $CF_m(p)$ that is equivalent to φ.