On the proof theory of the modal mu-calculus

Thomas Studer

January 24, 2007

Abstract

We study the proof theoretic relationship between several deductive systems for the modal mu-calculus. This results in a completeness proof for a system that is suitable for deciding the validity problem of the mu-calculus. Moreover, this provides a new proof theoretic proof for the finite model property of the mu-calculus.

1 Introduction

The propositional modal μ-calculus has been introduced by Kozen [10]. It is the extension of (multi-)modal logic by least and greatest fixed point operators. The μ-calculus has a very expressive language which allows for arbitrary nestings of (possibly interleaved) fixed points. Still it has good computational properties: the model checking problem is in NP $\cap$ co-NP and the validity problem is in EXPTIME.

The μ-calculus is important in many logic approaches to computer science, mainly because its language is suitable for stating properties about the behavior of processes. For a first overview and as a guide to the literature see for instance Bradfield and Stirling [3]. There are also many connections to neighboring areas in mathematics and theoretical computer science such as automata theory, game theory, universal algebra, and lattice theory.

Let us mention some articles dealing with the proof theory of the μ-calculus. In his initial study, Kozen [10] proposes an axiomatization for the μ-calculus which he shows to be sound and complete for the so-called aconjunctive fragment. Furthermore, Kozen [11] establishes the finite model property of the μ-calculus by relating it to the theory of well-quasi-orders. This allows him to introduce an infinitary deduction rule and to claim soundness and completeness for for a system with this rule, however, making crucial use of a cut rule. Completeness of finitary axiomatizations with respect to the full language of the μ-calculus is addressed in Walukiewicz [17, 18]. However, automata- rather than proof-theoretic methods are at the core of his approach.

Jäger, Kretz and Studer [7] introduce a cut-free infinitary system $\mathsf{T}^{\omega}_{\mu+}$ for the μ-calculus. They also make use of an infinitary deduction rule which derives the validity of a greatest fixed point from the validity of all its (infinitely many) finite approximations. To show soundness of this rule, the finite model property is employed. Completeness of $\mathsf{T}^{\omega}_{\mu+}$ is established by a canonical counter-model construction.

The finite model property of the μ-calculus is readdressed by Streett and Emerson [16] who show that if a formula A is satisfiable, then it is already satisfiable in a finite Kripke structure whose number of worlds is exponentially bounded in the length of A. Making use of this result and of an idea from [8, 9, 13], it becomes possible to finitize $\mathsf{T}^{\omega}_{\mu+}$. This leads to a finitary, sound and complete, cut-free deductive system for the modal μ-calculus [7].

An interesting direction of proof theoretic research on the subject is also represented by Miclan [14] who studies natural deduction style translations of Kozen's system and their implementation in interactive theorem provers. More remotely, the approach taken by Andersen, Stirling and Winskel [1] as well as the follow-up work by Berezin and Gurov [2] also studies proof systems for the μ-calculus. However, the purpose of their systems is to derive local satisfaction statements of the form *A holds for process p* and not the global validity of a given formula.

Dax, Hofmann, and Lange [4] present an infinitary proof system for the linear time μ-calculus which is suitable for deciding the validity problem of the linear time μ-calculus. They also mention a related system for the full modal μ-calculus which we call $\mathsf{T}^{\mathsf{pre}}_{\mu}$. A proof in $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is a finitely branching tree which may have branches of infinite length. These infinite branches must satisfy a global criterion saying (roughly) that there must be an outermost greatest fixed point unfolded infinitely many often in this branch.

The main contribution of the present paper is to show that if a formula A of the μ-calculus is derivable in $\mathsf{T}^{\omega}_{\mu+}$, then it is also derivable in $\mathsf{T}^{\mathsf{pre}}_{\mu}$. This provides

1. completeness of $\mathsf{T}^{\mathsf{pre}}_{\mu}$ (since $\mathsf{T}^{\omega}_{\mu+}$ is complete),
2. a soundness proof for $\mathsf{T}^{\omega}_{\mu+}$ that does not refer to the finite model property (since $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is sound),
3. a proof-theoretic proof of the finite model property of the μ-calculus (since our observations make it possible to adapt the canonical counter-

model construction presented in [7] such that the constructed model is finite).

Sprenger and Dam [15] also compare two proof systems for the μ-calculus each using a different type of induction. However, their systems include a cut-rule which allows for a straightforward translation from local to global induction. We study cut-free systems which makes the construction more involved.

Let us illustrate our approach by an example. Consider the $\mathsf{T}^{\omega}_{\mu+}$ proof of the $\mathcal{L}_\mu$ formula $(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}$ shown below. Starting from this given proof, we can construct a $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof for as follows.

$$
\dfrac{
\dfrac{\dfrac{\dfrac{(\mu\mathsf{X})\Box\mathsf{X}, \top}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu^1\mathsf{Y})\Diamond\mathsf{Y}}
\qquad
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(\mu\mathsf{X})\Box\mathsf{X}, \top}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu^1\mathsf{Y})\Diamond\mathsf{Y}}}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond((\nu^1\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond((\nu^1\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu^2\mathsf{Y})\Diamond\mathsf{Y}}
\qquad \cdots
}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}
$$

We take the branch through the premise $(\nu^2\mathsf{Y})\Diamond\mathsf{Y}$ of the infinitary greatest fixed point rule. In this branch, we drop all the iteration numbers. That is we replace all subexpressions of the form $(\nu^k\mathsf{X})\mathcal{C}$ by $(\nu\mathsf{X})\mathcal{C}$. This gives us the following.

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(\mu\mathsf{X})\Box\mathsf{X}, \top}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu^1\mathsf{Y})\Diamond\mathsf{Y}}}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond((\nu^1\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond((\nu^1\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu^2\mathsf{Y})\Diamond\mathsf{Y}}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}
\qquad \Longrightarrow \qquad
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(\mu\mathsf{X})\Box\mathsf{X}, \top}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond\top}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond((\nu\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond((\nu\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}
$$

We note that droping the iteration numbers in the sequents

$$(\mu\mathsf{X})\Box\mathsf{X}, (\nu^2\mathsf{Y})\Diamond\mathsf{Y} \text{ and } (\mu\mathsf{X})\Box\mathsf{X}, (\nu^1\mathsf{Y})\Diamond\mathsf{Y}$$

makes them identical. Therefore we can loop between these two sequents and get the following infinitary $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof.

$$
\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}}{\Box((\mu\mathsf{X})\Box\mathsf{X}), \Diamond((\nu\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, \Diamond((\nu\mathsf{Y})\Diamond\mathsf{Y})}}{(\mu\mathsf{X})\Box\mathsf{X}, (\nu\mathsf{Y})\Diamond\mathsf{Y}}
$$

In this example, we could choose the branch through the second premise of the $(\nu.\omega)$ rule in order to find two identical sequents. To show that this approach works in general, we have to guarantee that if we derive $\Gamma, (\nu\mathsf{X})\mathcal{A}$ by a $(\nu.\omega)$ rule, then there is a branch providing two identical sequents to build a loop.

The following argument shows that this indeed is the case. Let n be the cardinality of the Fischer-Ladner closure of $\Gamma, (\nu\mathsf{X})\mathcal{A}$. Droping the iteration numbers in a sequent of the $\mathsf{T}^{\omega}_{\mu+}$ proof of $\Gamma, (\nu\mathsf{X})\mathcal{A}$, gives us a sequent which belongs to the Fischer-Lander closure of $\Gamma, (\nu\mathsf{X})\mathcal{A}$. Note that there are only $2^n - 1$ non-empty subset of this Fischer-Ladner closure. Therefore, we can chose the branch through the premise $\Gamma, (\nu^{2^n}\mathsf{X})\mathcal{A}$. This branch must contain two sequents that are identical if the iteration numbers are dropped. Thus we can build the corresponding $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof tree from that branch. An additional observation concerning threads in $\mathsf{T}^{\omega}_{\mu+}$ proofs shows that the global criterion on $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof branches is also satisfied.

The paper is organized as follows. Sections 2 and 3 introduce the language and semantics of the modal μ-calculus. We recall the definition of the system $\mathsf{T}^{\omega}_{\mu+}$ in Section 4. The deductive system $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is presented in Section 5. Section 6 proves some lemmata about threads which are needed in the completeness proof. They guarantee that the $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof tree we construct satisfies the global criterion on infinite branches. We establish completeness of $\mathsf{T}^{\mathsf{pre}}_{\mu}$ in Section 7 where we show how to construct the $\mathsf{T}^{\mathsf{pre}}_{\mu}$ proof from a given $\mathsf{T}^{\omega}_{\mu+}$ proof. Section 8 deals with soundness issues. It is shown that soundness of $\mathsf{T}^{\mathsf{pre}}_{\mu}$ makes it possible to finitize $\mathsf{T}^{\omega}_{\mu+}$ which results in a finitary cut-free system $\mathsf{T}_{\mu+}$. All three systems $\mathsf{T}^{\omega}_{\mu+}$, $\mathsf{T}_{\mu+}$, and $\mathsf{T}^{\mathsf{pre}}_{\mu}$ are sound and complete. In Section 9, we present some applications of our construction. The main result is a new proof theoretic proof for the finite model property of the μ-calculus. Section 10 concludes the paper.

2 Language

We will introduce the language $\mathcal{L}_\mu$ of the modal μ-calculus. In addition, we will need an extension $\mathcal{L}^+_\mu$ of $\mathcal{L}_\mu$ that contains formulae to explicitly represent the finite approximations $(\nu^k\mathsf{X})\mathcal{A}$ of a greatest fixed point νXA.

Definition 1 (Language $\mathcal{L}_\mu$, free variables). Let

$$\Phi = \{\mathsf{p}, \sim\!\mathsf{p}, \mathsf{q}, \sim\!\mathsf{q}, \mathsf{r}, \sim\!\mathsf{r}, \ldots\}$$

be a countable set of atomic propositions,

$$\mathsf{V} = \{\mathsf{X}, \sim\!\mathsf{X}, \mathsf{Y}, \sim\!\mathsf{Y}, \mathsf{Z}, \sim\!\mathsf{Z}, \ldots\}$$

a set containing countably many variables and their negations, $\mathsf{T} = \{\top, \bot\}$ a set containing symbols for truth and falsehood and M a set of indices. Define the formulae of the language $\mathcal{L}_\mu$ as well as the set *fv* of free variables of each formula inductively as follows:

1. If P is an element of $\Phi \cup \mathsf{V} \cup \mathsf{T}$, then P is a formula of $\mathcal{L}_\mu$. Furthermore, if $P = \mathsf{X}$ or $P = \sim\!\mathsf{X}$ for some X or $\sim\!\mathsf{X}$ from V, then $fv(P) := \{\mathsf{X}\}$, otherwise $fv(P) := \emptyset$.

2. If A and B are formulae of $\mathcal{L}_\mu$, then so are $(A \wedge B)$ and $(A \vee B)$. Furthermore, we define $fv((A \wedge B)) := fv((A \vee B)) := fv(A) \cup fv(B)$.

3. If A is a formula of $\mathcal{L}_\mu$ and $i \in \mathsf{M}$, then $\Box_i A$ and $\Diamond_i A$ are also formulae of $\mathcal{L}_\mu$. Furthermore, let $fv(\Box_i A) := fv(\Diamond_i A) := fv(A)$.

4. If A is a formula of $\mathcal{L}_\mu$ and the negated variable $\sim\!\mathsf{X}$ does not occur in A, then $(\mu\mathsf{X})A$ and $(\nu\mathsf{X})A$ are also formulae of $\mathcal{L}_\mu$. Furthermore, we define $fv((\mu\mathsf{X})A) := fv((\nu\mathsf{X})A) := fv(A) \setminus \{\mathsf{X}\}$.

In case there is no danger of confusion, we will omit parentheses in formulae. If the negated variable $\sim\!\mathsf{X}$ does not occur in a formula A of $\mathcal{L}_\mu$, we say that A is X–positive or alternatively positive in X. Formulae which are positive in a certain variable determined by the context will henceforth be denoted by letters $\mathcal{A}, \mathcal{B}, \mathcal{C}, \ldots$. Furthermore, we will call a formula A of $\mathcal{L}_\mu$ closed, if $fv(A) = \emptyset$.

Definition 2 (Language $\mathcal{L}_\mu^+$). The formulae of the extended language $\mathcal{L}_\mu^+$ (and their free variables) are defined by adding the following clause to the induction of Definition 1

5. If A is a formula of $\mathcal{L}_\mu^+$ and the negated variable $\sim\!\mathsf{X}$ does not occur in A, then for every natural number $k > 0$ $(\nu^k\mathsf{X})A$ is also a formula of $\mathcal{L}_\mu^+$. Furthermore, we define $fv((\nu^k\mathsf{X})A) := fv(A) \setminus \{\mathsf{X}\}$.

We define X–positive and closed formulae of $\mathcal{L}_\mu^+$ analogously to those of $\mathcal{L}_\mu$. Given a formula B of $\mathcal{L}_\mu^+$ we define B^- as the formula obtained from B by first replacing all subexpressions of the form $(\nu^k\mathsf{X})\mathcal{C}$ by $(\nu\mathsf{X})\mathcal{C}$ and then all free variables by $\top$. Clearly B^- is a formula of $\mathcal{L}_\mu$. For a set Γ of $\mathcal{L}_\mu^+$ formulae, we define Γ^- as $\bigcup_{B \in \Gamma} B^-$.

We use $(\sigma\mathsf{X})\mathcal{A}$ to denote formulae of the form $(\mu\mathsf{X})\mathcal{A}$, $(\nu\mathsf{X})\mathcal{A}$, and $(\nu^k\mathsf{X})\mathcal{A}$ for all k. Moreover, we write $B \in \mathsf{sub}(A)$ if B is a subformula of A.

3 Semantics of $\mathcal{L}_\mu^+$

We make use of the standard Kripke semantics for multi-modal fixed point logics to give meaning to $\mathcal{L}_\mu^+$ formulae.

Definition 3 (Kripke structure). A Kripke structure for $\mathcal{L}_\mu^+$ is a triple $\mathsf{K} = (S, R, \pi)$, where S is a non–empty set, $R : \mathsf{M} \to \mathcal{P}(S \times S)$ and $\pi : (\Phi \cup \mathsf{V}) \to \mathcal{P}(S)$ is a function such that $\pi({\sim}\mathsf{X}) = S \setminus \pi(\mathsf{X})$ for all ${\sim}\mathsf{X} \in \mathsf{V}$ and $\pi({\sim}\mathsf{p}) = S \setminus \pi(\mathsf{p})$ for all ${\sim}\mathsf{p} \in \Phi$. The function R assigns an accessibility relation to each $i \in \mathsf{M}$ where we write R_i for the relation $R(i)$. Furthermore, given a set $T \subset S$ and a variable $\mathsf{X} \in \mathsf{V}$ we define the Kripke structure $\mathsf{K}[\mathsf{X} := T]$ as the triple (S, R, π'), where $\pi'(\mathsf{X}) = T$, $\pi'({\sim}\mathsf{X}) = S \setminus T$ and $\pi'(P) = \pi(P)$ for all other $P \in \Phi \cup \mathsf{V}$.

We are now ready to assign a meaning to the formulae of $\mathcal{L}_\mu^+$ in terms of Kripke structures. This is achieved in a straightforward way by induction on the structure of formulae, with a side induction on all natural numbers greater than 0 to treat finite greatest fixed point approximations.

Definition 4 (Denotation). Let $\mathsf{K} = (S, R, \pi)$ be a Kripke structure. For every $A \in \mathcal{L}_\mu^+$ we define the set $\|A\|_{\mathsf{K}} \subset S$ inductively as follows:

$$\|P\|_{\mathsf{K}} := \pi(P) \text{ for all } P \in \Phi \cup \mathsf{V}, \quad \|\top\|_{\mathsf{K}} := S, \quad \|\bot\|_{\mathsf{K}} := \emptyset,$$
$$\|B \wedge C\|_{\mathsf{K}} := \|B\|_{\mathsf{K}} \cap \|C\|_{\mathsf{K}}, \quad \|B \vee C\|_{\mathsf{K}} := \|B\|_{\mathsf{K}} \cup \|C\|_{\mathsf{K}},$$
$$\|\Box_i B\|_{\mathsf{K}} := \{w \in S : v \in \|B\|_{\mathsf{K}} \text{ for all } v \text{ such that } wR_iv\},$$
$$\|\Diamond_i B\|_{\mathsf{K}} := \{w \in S : v \in \|B\|_{\mathsf{K}} \text{ for some } v \text{ such that } wR_iv\}.$$

For every formula $(\mu\mathsf{X})\mathcal{A}$ and $(\nu\mathsf{X})\mathcal{A}$ we define

$$\|(\mu\mathsf{X})\mathcal{A}\|_{\mathsf{K}} := \bigcap\{T \subset S : T \supset F^{\mathsf{K}}_{\mathcal{A},\mathsf{X}}(T)\} \text{ and}$$
$$\|(\nu\mathsf{X})\mathcal{A}\|_{\mathsf{K}} := \bigcup\{T \subset S : T \subset F^{\mathsf{K}}_{\mathcal{A},\mathsf{X}}(T)\}$$

where $F^{\mathsf{K}}_{\mathcal{A},\mathsf{X}}$ is the operator on S given by $F^{\mathsf{K}}_{\mathcal{A},\mathsf{X}}(T) := \|\mathcal{A}\|_{\mathsf{K}[\mathsf{X}:=T]}$ for every subset T of S. Furthermore, if $\mathcal{A}$ is an X–positive formula, then we define $\|(\nu^k\mathsf{X})\mathcal{A}\|_{\mathsf{K}}$ for every $k > 0$ by induction on k as follows:

$$\|(\nu^1\mathsf{X})\mathcal{A}\|_{\mathsf{K}} := \|\mathcal{A}[\top/\mathsf{X}]\|_{\mathsf{K}}$$
$$\|(\nu^{n+1}\mathsf{X})\mathcal{A}\|_{\mathsf{K}} := \|\mathcal{A}[(\nu^n\mathsf{X})\mathcal{A}]\|_{\mathsf{K}}.$$

We write $\mathsf{K}, s \models A$ for $s \in \|A\|_{\mathsf{K}}$. We call a $\mathcal{L}_\mu^+$ formula A *satisfiable* if there is a Kripke structure K such that $\|A\|_{\mathsf{K}}$ is non-empty. We say A is satisfiable by a finite Kripke structure if there is a Kripke structure $\mathsf{K} = (S, R, \pi)$ with S finite such that $\|A\|_{\mathsf{K}}$ is non-empty. The formula A is called *valid* if for every Kripke structure $\mathsf{K} = (S, R, \pi)$ we have $\|A\|_{\mathsf{K}} = S$.

Let us illustrate the semantics of $\mathcal{L}_\mu^+$ formulae by an example. First, we need the following definition.

Definition 5. Let $\mathsf{K} = (S, R, \pi)$ be a Kripke structure and let $j \in \mathsf{M}$. A *j-path in* K is a (possibly infinite) sequence $(s_0, \dots, s_n)$ of elements of S such that for every $0 \leq i < n$ we have $(s_i, s_{i+1}) \in R_j$. Note that we may have $s_m = s_n$ for $m \neq n$. We say the path $(s_0, \dots, s_n)$ has *length* n.

Example 6. Assume we are given a Kripke structure $\mathsf{K} = (S, R, \pi)$. The following holds (see [12] for details):

1. If $\|\Diamond_j \top\|_{\mathsf{K}}$ is non-empty, then K contains a j-path of length 1.
2. If $\|(\nu^l \mathsf{X})\Diamond_j \mathsf{X}\|_{\mathsf{K}}$ is non-empty, then K contains a j-path of length l.
3. If $\|(\nu \mathsf{X})\Diamond_j \mathsf{X}\|_{\mathsf{K}}$ is non-empty, then K contains a j-path of infinite length.
4. If $\|(\mu \mathsf{X})\Box_j \mathsf{X}\|_{\mathsf{K}} = S$, then K does not contain a j-path of infinite length.

4 The system $\mathsf{T}^\omega_{\mu+}$

The infinitary calculus $\mathsf{T}^\omega_{\mu+}$ is introduced in [7]. This deductive system provides a cut-free, sound and complete axiomatization for the modal μ-calculus. $\mathsf{T}^\omega_{\mu+}$ is formulated as Tait-style system which derives finite sets $\Gamma, \Delta, \Sigma, \dots$ of $\mathcal{L}_\mu^+$ formulae which we call *sequents*. These sequents are interpreted disjunctively. In general, we write Γ, A for $\Gamma \cup \{A\}$. Moreover, if Γ is the set $\{A_1, \dots, A_n\}$ of $\mathcal{L}_\mu^+$ formulae, then $\Diamond_i \Gamma := \{\Diamond_i A_1, \dots, \Diamond_i A_n\}$.

Definition 7 (The system $\mathsf{T}^\omega_{\mu+}$). The system $\mathsf{T}^\omega_{\mu+}$ is defined by the following inference rules:

Axioms: For all sequents Γ of $\mathcal{L}_\mu^+$, all p in Φ, and all X in V

$$\frac{}{\Gamma, \mathsf{p}, {\sim}\mathsf{p}}\ (\text{ID1}), \qquad \frac{}{\Gamma, \mathsf{X}, {\sim}\mathsf{X}}\ (\text{ID2}), \qquad \frac{}{\Gamma, \top}\ (\text{ID3}).$$

Propositional rules: For all sequents Γ and formulae A and B of $\mathcal{L}_\mu^+$

$$\frac{\Gamma, A, B}{\Gamma, A \vee B}\ (\vee) \qquad \frac{\Gamma, A \qquad \Gamma, B}{\Gamma, A \wedge B}\ (\wedge)$$

Modal rules: For all sequents Γ and Σ and formulae A of $\mathcal{L}_\mu^+$ and all indices i from M

$$\frac{\Gamma, A}{\Diamond_i \Gamma, \Box_i A, \Sigma} \quad (\Box)$$

Approximation rules: For all sequents Γ and X–positive formulae $\mathcal{A}$ of $\mathcal{L}_\mu^+$ and all natural numbers $k > 0$

$$\frac{\Gamma, \mathcal{A}[\top/\mathsf{X}]}{\Gamma, (\nu^1\mathsf{X})\mathcal{A}} \quad (\nu.1) \qquad\qquad \frac{\Gamma, \mathcal{A}[(\nu^k\mathsf{X})\mathcal{A}]}{\Gamma, (\nu^{k+1}\mathsf{X})\mathcal{A}} \quad (\nu.k+1)$$

Fixed point rules: For all sequents Γ and X–positive formulae $\mathcal{A}$ of $\mathcal{L}_\mu^+$

$$\frac{\Gamma, \mathcal{A}[(\mu\mathsf{X})\mathcal{A}]}{\Gamma, (\mu\mathsf{X})\mathcal{A}} \quad (\mu) \qquad\qquad \frac{\Gamma, (\nu^k\mathsf{X})\mathcal{A} \quad \text{for all } k \in \omega}{\Gamma, (\nu\mathsf{X})\mathcal{A}} \quad (\nu.\omega)$$

The *distinguished formula* of a rule is the formula that is explicitly displayed in the conclusion of the rule. The *active formulae* of a rule are those formulae that are explicitly displayed in the rule. The formulae in Γ and Σ are called *side formulae* of a rule.

Definition 8 (Provability). Assume Γ is a sequent of $\mathcal{L}_\mu^+$ and α an ordinal. We define the provability of Γ in $\mathsf{T}_{\mu+}^\omega$ in α many steps, denoted by $\mathsf{T}_{\mu+}^\omega \stackrel{\alpha}{\vdash} \Gamma$, by induction as follows:

1. If Γ is obtained by one of the axioms of $\mathsf{T}_{\mu+}^\omega$, then $\mathsf{T}_{\mu+}^\omega \stackrel{\beta}{\vdash} \Gamma$ holds for all ordinals β.

2. If Γ is obtained by one of the propositional, modal, approximation or fixed point rules where Γ_i are the premises of the respective rule, $\mathsf{T}_{\mu+}^\omega \stackrel{\beta_i}{\vdash} \Gamma_i$ holds for all of these premises, and β is an ordinal such that $\beta_i < \beta$ for all β_i, then $\mathsf{T}_{\mu+}^\omega \stackrel{\beta}{\vdash} \Gamma$.

We say a sequent Γ is provable and write $\mathsf{T}_{\mu+}^\omega \vdash \Gamma$ if there exists an ordinal β such that Γ is provable in β many steps.

Jäger, Kretz and Studer [7] present a canonical counter model construction for $\mathsf{T}_{\mu+}^\omega$ which shows its completeness. Soundness of $\mathsf{T}_{\mu+}^\omega$ is shown in [7] by making use of the finite model property of the μ-calculus.

Theorem 9. *The system $\mathsf{T}_{\mu+}^\omega$ is sound and complete for $\mathcal{L}_\mu$ formulae.*

5 The system $\mathsf{T}_{\mu}^{\mathsf{pre}}$

Dax, Hofmann, and Lange [4] present an infinitary proof system for the linear time μ-calculus. A proof in their system is an infinite tree in which each branch satisfies an additional global condition concerning the existence of so-called threads. This condition can be checked by finite automata. Hence, they obtain a decision procedure for the validity problem. Note that this characterization of valid proof branches is closely related to Streett and Emerson's [16] notions of premodels and models.

In the section 'Further Work' of [4], it is mentioned how a corresponding infinitary system for the modal μ-calculus can be formulated. In this section, we present such a deductive system which we call $\mathsf{T}_{\mu}^{\mathsf{pre}}$.

Definition 10. A *preproof* for a sequent Γ of $\mathcal{L}_\mu$ formulae is a possibly infinite tree whose root is labeled with Γ and which is built according to the following rules.

Axioms: For all sequents Γ of $\mathcal{L}_\mu$, all p in Φ, and all X in V

$$\frac{}{\Gamma, \mathsf{p}, {\sim}\mathsf{p}}\ (\text{ID1}), \qquad \frac{}{\Gamma, \mathsf{X}, {\sim}\mathsf{X}}\ (\text{ID2}), \qquad \frac{}{\Gamma, \top}\ (\text{ID3}).$$

Propositional rules: For all sequents Γ and formulae A and B of $\mathcal{L}_\mu$

$$\frac{\Gamma, A, B}{\Gamma, A \lor B}\ (\lor) \qquad \frac{\Gamma, A \qquad \Gamma, B}{\Gamma, A \land B}\ (\land)$$

Modal rules: For all sequents Γ and Σ and formulae A of $\mathcal{L}_\mu$ and all indices i from M

$$\frac{\Gamma, A}{\Diamond_i \Gamma, \Box_i A, \Sigma}\ (\Box)$$

Fixed point rules: For all sequents Γ and X–positive formulae $\mathcal{A}$ of $\mathcal{L}_\mu$

$$\frac{\Gamma, \mathcal{A}[(\mu\mathsf{X})\mathcal{A}]}{\Gamma, (\mu\mathsf{X})\mathcal{A}}\ (\mu) \qquad \frac{\Gamma, \mathcal{A}[(\nu\mathsf{X})\mathcal{A}]}{\Gamma, (\nu\mathsf{X})\mathcal{A}}\ (\nu)$$

In the sequel, we use the term *proof tree* not only for $\mathsf{T}_{\mu}^{\mathsf{pre}}$ preproofs, but also for proofs in the system $\mathsf{T}_{\mu+}^{\omega}$.

Definition 11. Assume we are given a proof tree for some sequent. For all rule applications r occurring in this proof tree, we define a *connection relation* $\mathsf{Con}(r)$ on formulae as follows. $(A, B) \in \mathsf{Con}(r)$ iff $A = B$ is a side formula of r or A is an active formula in the conclusion and B is an active formula in a premise of r.

Note that if r is an instance of ($\Box$) and $A \in \Gamma$ in this rule, then we have $(\Diamond A, A) \in \mathsf{Con}(r)$.

Definition 12. Assume we are given a (possibly infinite) branch $\Gamma_0, \Gamma_1, \ldots$ in a proof tree and let r_i be the rule application that derived Γ_i from Γ_{i+1}. A *thread* in this branch is a sequence of formulae $A_0, A_1, \ldots$ such that $(A_i, A_{i+1}) \in \mathsf{Con}(r_i)$.

Definition 13. An $\mathcal{L}_\mu$ sequent Γ is called *well-named* if every variable is bounded at most once. Note that for a bound variable X in a well-named sequent $A_1, \ldots, A_n$, there exists exactly one formula B of the form $(\sigma \mathsf{X})\mathcal{A}$ which is a subformula of an A_i. We call B the *bounding formula* of X. If the bounding formula of a variable X is of the form $(\nu \mathsf{X})\mathcal{A}$, then X is called a *ν variable* in Γ. Let Γ be sequent containing two bound variables X and Y. We say X is *higher* that Y if the bounding formula of Y is a subformula of the bounding formula of X.

In the sequel we consider only proofs and preproofs for well-named sequents. We immediately get the following fact about threads.

Lemma 14. *Assume we are given an infinite branch of a preproof for an $\mathcal{L}_\mu$ sequent Γ. For every thread in this branch there is a unique bound variable X such that*

1. *the bounding formula of X occurs infinitely often in the thread and*

2. *for every other formula of the form $(\sigma \mathsf{Y})\mathcal{A}$ which occurs infinitely often, we have that X is higher than Y.*

Definition 15. Assume we are given a preproof for an $\mathcal{L}_\mu$ sequent Γ. A thread in this proof is called *ν-thread* if the unique variable given by the previous lemma is a ν variable in Γ. A $\mathsf{T}_\mu^{\mathsf{pre}}$ *proof* for a sequent Γ of $\mathcal{L}_\mu$ formulae is a preproof of Γ such that every finite branch ends in an axiom and every infinite branch contains a ν-thread. We write $\mathsf{T}_\mu^{\mathsf{pre}} \vdash \Gamma$ if there exists a $\mathsf{T}_\mu^{\mathsf{pre}}$ proof for Γ.

6 About threads

Let us study some properties of threads in $\mathsf{T}_{\mu+}^{\omega}$ proofs. These properties will be needed later to show completeness of $\mathsf{T}_\mu^{\mathsf{pre}}$. We start with defining auxiliary sets of formulae satisfying certain closure conditions.

Definition 16 (Fischer–Ladner closure). Let D be a closed formula of $\mathcal{L}_\mu$. The Fischer–Ladner closure $\mathbb{FL}(D)$ of D is defined inductively as follows:

1. $D \in \mathbb{FL}(D)$
2. If $A \wedge B \in \mathbb{FL}(D)$ or $A \vee B \in \mathbb{FL}(D)$, then $A \in \mathbb{FL}(D)$ and $B \in \mathbb{FL}(D)$.
3. If $\Box_i A \in \mathbb{FL}(D)$ or $\Diamond_i A \in \mathbb{FL}(D)$, then $A \in \mathbb{FL}(D)$.
4. If $(\mu\mathsf{X})\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A} \in \mathbb{FL}(D)$ and $\mathcal{A}[(\mu\mathsf{X})\mathcal{A}] \in \mathbb{FL}(D)$.
5. If $(\nu\mathsf{X})\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A} \in \mathbb{FL}(D)$ and $\mathcal{A}[(\nu\mathsf{X})\mathcal{A}] \in \mathbb{FL}(D)$.

Let Γ be a sequent of closed formulae of $\mathcal{L}_\mu$. We define $\mathbb{FL}(\Gamma)$ as $\bigcup_{D\in\Gamma} \mathbb{FL}(D)$.

It is standard to show that the Fischer-Ladner closure of a formula is a finite set [6].

Lemma 17. *The cardinality of $\mathbb{FL}(D)$ of a formula D of $\mathcal{L}_\mu$ is linear in the length of D, thus in particular $\mathbb{FL}(D)$ is a finite set.*

Definition 18 (Strong closure). Let D be a closed formula of $\mathcal{L}_\mu^+$. The strong closure $\mathbb{SC}(D)$ of D is defined inductively as follows:

1. $D \in \mathbb{SC}(D)$
2. If $A \wedge B \in \mathbb{SC}(D)$ or $A \vee B \in \mathbb{SC}(D)$, then $A \in \mathbb{SC}(D)$ and $B \in \mathbb{SC}(D)$.
3. If $\Box_i A \in \mathbb{SC}(D)$ or $\Diamond_i A \in \mathbb{SC}(D)$, then $A \in \mathbb{SC}(D)$.
4. If $(\mu\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A} \in \mathbb{SC}(D)$ and $\mathcal{A}[(\mu\mathsf{X})\mathcal{A}] \in \mathbb{SC}(D)$.
5. If $(\nu\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A} \in \mathbb{SC}(D)$ and for every natural number $n > 0$ also $(\nu^n\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$.
6. If $(\nu^1\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A}[\top/\mathsf{X}] \in \mathbb{SC}(D)$.
7. If n is a natural number greater than 0 and $(\nu^{n+1}\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A}[(\nu^n\mathsf{X})\mathcal{A}] \in \mathbb{SC}(D)$.
8. If $\mathcal{A}$ is X–positive and $\mathcal{A} \in \mathbb{SC}(D)$, then for every variable Y also $\mathcal{A}[\mathsf{Y}/\mathsf{X}] \in \mathbb{SC}(D)$.

In the sequel, we need the following lemma from [7].

Lemma 19. *Let D be a closed formula of $\mathcal{L}_\mu$. Then for all formulae A of $\mathcal{L}_\mu^+$ we have*

$$A \in \mathbb{SC}(D) \implies A^- \in \mathbb{FL}(D).$$

The following two lemmata are corollaries from the definitions.

Lemma 20. *Let D be a well-named closed formula of $\mathcal{L}_\mu$ and let X be a variable occurring in D. Then for all $\mathcal{L}_\mu^+$ formulae $A \in \mathbb{SC}(D)$ we have: if $(\sigma \mathsf{Y})\mathcal{B} \in \mathsf{sub}(A)$ and $X \in fv((\sigma \mathsf{Y})\mathcal{B})$, then X is higher than Y in D.*

Lemma 21. *Assume we are given a thread $A_1, A_2, \ldots$. Let A_i, A_j be formulae of this thread with $i \leq j$. Then we have $A_j \in \mathbb{SC}(A_i)$.*

Lemma 22. *Assume we are given closed $\mathcal{L}_\mu^+$ formula $(\nu^{k+1}\mathsf{X})\mathcal{A}$ and a formula $B \in \mathbb{SC}((\nu^{k+1}\mathsf{X})\mathcal{A})$ such that $(\nu^k\mathsf{X})\mathcal{A} \in \mathbb{SC}(B)$, then B has the form $\mathcal{B}'[(\nu^k\mathsf{X})\mathcal{A}/\mathsf{X}]$ for some $\mathcal{B}'$ not containing $(\nu^k\mathsf{X})\mathcal{A}$.*

Proof. This lemma is shown by induction on the build up of $\mathbb{SC}((\nu^{k+1}\mathsf{X})\mathcal{A})$. □

Lemma 23. *Assume we are given a $\mathsf{T}_{\mu+}^{\omega}$ proof for an $\mathcal{L}_\mu$ formula D. Further assume there is a thread in this proof of the form*

$$D, \ldots, (\nu^{k+1}\mathsf{X})\mathcal{A}, \mathcal{A}[(\nu^k\mathsf{X})\mathcal{A}], \ldots, (\mu\mathsf{Y})\mathcal{B}, \mathcal{B}[(\mu\mathsf{Y})\mathcal{B}], \ldots, (\nu^k\mathsf{X})\mathcal{A}, \ldots.$$

Then we have that X is higher than Y in D.

Proof. Lemmata 21 and 22 imply that $(\mu\mathsf{Y})\mathcal{B}$ is of the form

$$(\mu\mathsf{Y})\mathcal{B}'[(\nu^k\mathsf{X})\mathcal{A}/\mathsf{X}].$$

This is only possible if $(\mu\mathsf{Y})\mathcal{B}'(X) \in \mathsf{sub}((\nu^{k+1}\mathsf{X})\mathcal{A})$. Now we find by Lemma 20 that X is higher than Y in D. □

7 Completeness of $\mathsf{T}_\mu^{\mathsf{pre}}$

We will prove completeness of $\mathsf{T}_\mu^{\mathsf{pre}}$ by showing how to construct a $\mathsf{T}_\mu^{\mathsf{pre}}$ proof for an $\mathcal{L}_\mu$ formula A given a $\mathsf{T}_{\mu+}^{\omega}$ proof of A.

The following lemma follows immediately from the fact that in a $\mathsf{T}_{\mu+}^{\omega}$ derivation a formula of the form $(\nu^i\mathsf{X})\mathcal{A}$ can only be eliminated if $(\nu^{i+1}\mathsf{X})\mathcal{A}$ or $(\nu\mathsf{X})\mathcal{A}$ is introduced.

Lemma 24. *Let Γ be an $\mathcal{L}_\mu^+$ sequent. Assume we are given a $\mathsf{T}_{\mu+}^{\omega}$ proof for $\Gamma, (\nu\mathsf{X})\mathcal{A}$ where $(\nu\mathsf{X})\mathcal{A}$. Further, assume we are given a branch in the proof tree with a thread*

$$\ldots, (\nu\mathsf{X})\mathcal{A}, (\nu^k\mathsf{X})\mathcal{A}, \ldots, (\nu^1\mathsf{X})\mathcal{A}, \mathcal{A}[\top/\mathsf{X}], \ldots$$

Then this thread contains $(\nu^i\mathsf{X})\mathcal{A}$ for every $1 \leq i \leq k$.

For our construction, we need the function f computing the number of subsets of the Fischer-Ladner closure of Γ^- for a $\mathcal{L}_\mu^+$ sequent Γ.

Definition 25. Let f the function assigning to each $\mathcal{L}_\mu^+$ sequent Γ a natural number as follows:

$$\mathsf{f}(\Gamma) := 2^{|\mathbb{FL}(\Gamma^-)|}$$

where $|\mathbb{FL}(\Gamma^-)|$ is the cardinality of the Fischer-Ladner closure of Γ^-.

Definition 26. Assume we are given a $\mathsf{T}_{\mu+}^\omega$ proof for an $\mathcal{L}_\mu^+$ sequent Γ. The *pruned proof tree* $\mathcal{PPT}$ of this given proof is a tree labeled by $\mathcal{L}_\mu^+$ sequents. We define $\mathcal{PPT}$ by induction on the length of the given proof as follows where we distinguish the different cases for the last rule applied in the proof.

1. If the given proof consists only of an axiom, the $\mathcal{PPT}$ consists as well only of this axiom.

2. If the last rule was an instance of $(\vee)$, $(\wedge)$, $(\Box)$, $(\nu.1)$, $(\nu.k+1)$, or (μ), then we construct the pruned proof trees of the proofs for the premises. $\mathcal{PPT}$ is now given as the disjoint union of these pruned proof trees with the addition new root node labeled by Γ.

3. If the last rule was an instance of $(\nu.\omega)$, then $\mathcal{PPT}$ is given as the pruned proof tree for the premise $\Sigma, (\nu^k \mathsf{X})\mathcal{A}$ where $k = \mathsf{f}(\Gamma)$.

Note that $\mathcal{PPT}$ is a finite tree.

Lemma 27. *Assume we are given a pruned proof tree $\mathcal{PPT}$ of a $\mathsf{T}_{\mu+}^\omega$ proof of an $\mathcal{L}_\mu$ sequent Γ. Let $\Gamma_1, \dots \Gamma_n$ be a branch in $\mathcal{PPT}$ such that Γ_h has been derived from Γ_{h+1} by an application of a $(\nu.1)$ rule for some $1 \le h < n$. Then there are $1 \le i, j \le n$ with*

$$\Gamma_i = \Delta_i, (\nu^l \mathsf{X})\mathcal{A} \text{ and } \Gamma_j = \Delta_j, (\nu^k \mathsf{X})\mathcal{A}$$

for some natural number $k \neq l$ such that

(1) $(\nu^l \mathsf{X})\mathcal{A}$ is the distinguished formula of Γ_i, and

(2) $(\nu^k \mathsf{X})\mathcal{A}$ is the distinguished formula of Γ_j, and

(3) $\Gamma_i^- = \Gamma_j^-$, and

(4) there is a thread containing both $(\nu^l \mathsf{X})\mathcal{A}$ and $(\nu^k \mathsf{X})\mathcal{A}$.

Proof. Assume we are given a branch of $\mathcal{PPT}$ in which $(\nu^1\mathsf{X})\mathcal{A}$ occurs as distinguished formula in the label of a node. Since Γ is an $\mathcal{L}_\mu$ sequent, the corresponding branch in the $\mathsf{T}^\omega_{\mu+}$ proof must contain a node labeled by $\Delta, (\nu\mathsf{X})\mathcal{A}$ where $(\nu\mathsf{X})\mathcal{A}$ is the distinguished formula of the node. Thus Lemma 24 and the definition of pruned proof tree imply that for each $h \leq \mathsf{f}(\Delta, (\nu\mathsf{X})\mathcal{A})$ there exists a node in the branch $\Gamma_1, \ldots \Gamma_n$ with the distinguished formula $(\nu^h\mathsf{X})\mathcal{A}$. Therefore $n \geq \mathsf{f}(\Delta, (\nu\mathsf{X})\mathcal{A})$. Moreover, Lemma 19 implies $\emptyset \neq \Gamma_i^- \subseteq \mathbb{FL}(\Delta, (\nu\mathsf{X})\mathcal{A})$ for $1 \leq i \leq n$. Since there are only $\mathsf{f}(\Delta, (\nu\mathsf{X})\mathcal{A}) - 1$ many non-empty subsets of $\mathbb{FL}(\Delta, (\nu\mathsf{X})\mathcal{A})$, there must exist Γ_i and Γ_j such that (1), (2), (3), and (4) hold. □

Let d be a node in a pruned proof tree $\mathcal{PPT}$. We denote the label of d in $\mathcal{PPT}$ by $\mathsf{label}(d)$.

Definition 28. Assume we are given a pruned proof tree $\mathcal{PPT}$ of a $\mathsf{T}^\omega_{\mu+}$ proof for an $\mathcal{L}_\mu$ sequent Γ. We simultaneously construct a $\mathsf{T}^{\mathsf{pre}}_\mu$ preproof $\mathcal{PRE}$ for Γ and a function origin which relates nodes of $\mathcal{PRE}$ to nodes of $\mathcal{PPT}$.

1. Let a be the root of $\mathcal{PRE}$. We define $\mathsf{origin}(a) := b$ where b is the root of $\mathcal{PPT}$.

2. A node $a \in \mathcal{PRE}$ is labeled by the $\mathcal{L}_\mu$ sequent Δ^- where Δ is the label of $\mathsf{origin}(a)$ in $\mathcal{PPT}$.

3. A node $a \in \mathcal{PRE}$ has child nodes $c_1, \ldots, c_n$ if $\mathsf{origin}(a)$ has n child nodes $b_1, \ldots, b_n$ in $\mathcal{PPT}$. For $1 \leq i \leq n$, we define

   (a) $\mathsf{origin}(c_i) := d$ if b_i has an ancestor node $d \in \mathcal{PPT}$ such that

      i. there is a $\mathcal{L}_\mu$ sequent $\Delta, (\nu\mathsf{X})\mathcal{A}$ with

      $$(\mathsf{label}(b_i))^- = (\mathsf{label}(d))^- = \Delta, (\nu\mathsf{X})\mathcal{A},$$

      and

      ii. $(\nu^l\mathsf{X})\mathcal{A}$ is the distinguished formula of d, $(\nu^k\mathsf{X})\mathcal{A}$ is the distinguished formula of b_i, and there is a thread containing both of these formulae.

   (b) $\mathsf{origin}(c_i) := b_i$ if no such node d exists.

Definition 28 indeed constructs a $\mathsf{T}^{\mathsf{pre}}_\mu$ preproof. The only critical point is if $\mathcal{PPT}$ contains a branch with an instance of $(\nu.1)$. However, Lemma 27 guarantees that such a branch is always transformed into an infinite branch in the $\mathsf{T}^{\mathsf{pre}}_\mu$ preproof.

Theorem 29. *For all closed $\mathcal{L}_\mu$ formulae A we have*

$$\mathsf{T}^{\omega}_{\mu+} \vdash A \quad \Longrightarrow \quad \mathsf{T}^{\mathsf{pre}}_{\mu} \vdash A.$$

Proof. Given the $\mathsf{T}^{\omega}_{\mu+}$ proof of A, we can construct the corresponding pruned proof tree and from that a preproof of A according to the Definitions 26 and 28. It remains to show that every infinite path of the preproof contains a ν-thread. Assume we are given such an infinite branch. Looking at the construction of the preproof, we notice that such a branch can only occur because of Condition 3a in Definition 28. Thus, there is a thread in this branch that contains $(\nu \mathsf{X})\mathcal{A}$ infinitely often. Moreover, Lemma 23 guarantees that this thread is a ν-thread. □

Corollary 30. *The system $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is complete for $\mathcal{L}_\mu$ formulae.*

8 Soundness

Dax et al. [4] provide a simple soundness proof of their system for the linear time μ-calculus. A straightforward adaption of this proof shows the soundness of $\mathsf{T}^{\mathsf{pre}}_{\mu}$. Simply replace the case for the 'next'-rule by an appropriate treatment of $(\Box)$.

Theorem 31. *The system $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is sound.*

The only infinitary rule in $\mathsf{T}^{\omega}_{\mu+}$ is $(\nu.\omega)$ which introduces greatest fixed points. This rule can be collapsed to a finitary rule for greatest fixed points by making use the the finite model property of the μ-calculus [7].

Our completeness proof for $\mathsf{T}^{\mathsf{pre}}_{\mu}$ provides a different method to obtain a finitary version of $(\nu.\omega)$. Namely, the proof of Lemma 27 shows that only $\mathsf{f}(\Gamma, (\nu \mathsf{X})\mathcal{A})$-many premises of $(\nu.\omega)$ are needed to correctly infer $(\nu \mathsf{X})\mathcal{A}$. This implies that the following rule for deriving greatest fixed points is sound.

Definition 32. The definition of the system $\mathsf{T}_{\mu+}$ is analogous to that of $\mathsf{T}^{\omega}_{\mu+}$ except that the rule $(\nu.\omega)$ is replaced by the following finitary rule for greatest fixed points:

For all sequents Γ and X–positive formulae $\mathcal{A}$ of $\mathcal{L}^{+}_{\mu}$

$$\frac{\Gamma, (\nu^k \mathsf{X})\mathcal{A} \quad \text{for all } 0 < k \leq \mathsf{f}(\Gamma, (\nu \mathsf{X})\mathcal{A})}{\Gamma, (\nu \mathsf{X})\mathcal{A}} \quad (\nu.\mathrm{FL}).$$

Note that every derivation in $\mathsf{T}^{\omega}_{\mu+}$ collapses to a derivation in $\mathsf{T}_{\mu+}$. This can easily be shown by induction on the derivations in $\mathsf{T}^{\omega}_{\mu+}$ since every application of $(\nu.\omega)$ in $\mathsf{T}^{\omega}_{\mu+}$ can be replaced by an instance of $(\nu.\mathrm{FL})$ in $\mathsf{T}_{\mu+}$.

Corollary 33. *Let A be an $\mathcal{L}_\mu$ formula. We have*

$$A \text{ is valid} \Longrightarrow \mathsf{T}^\omega_{\mu+} \vdash A \Longrightarrow \mathsf{T}_{\mu+} \vdash A \Longrightarrow \mathsf{T}^{\mathsf{pre}}_\mu \vdash A \Longrightarrow A \text{ is valid.}$$

That is the systems $\mathsf{T}^{\mathsf{pre}}_\mu$, $\mathsf{T}_{\mu+}$, and $\mathsf{T}^\omega_{\mu+}$ are sound and complete for $\mathcal{L}_\mu$ formulae.

9 Applications

As we have seen in the previous section, our construction allows us to collapse the infinitary rule (ν) of $\mathsf{T}^\omega_{\mu+}$ to a finitary rule for greatest fixed points. In this section, we make use of this observation to establish the finite model property by purely proof theoretic means. Before doing so, let us mention another theorem which shows how the finitary rule for greatest fixed point can directly be used to get a theorem about the existence of certain models.

Theorem 34. *Let A be a satisfiable $\mathcal{L}_\mu$ formula. For every $j \in \mathsf{M}$ there exists a Kripke structure K such that $\|A\|_\mathsf{K}$ is non-empty and K either contains an infinite j-path or the length of every j-path is less than $\mathsf{f}(A) + n$ for a fixed natural number n.*

Proof. The formula $(\mu\mathsf{X})\Box_j\mathsf{X} \vee (\nu\mathsf{X})\Diamond_j\mathsf{X}$ is valid. Thus

$$A \wedge ((\mu\mathsf{X})\Box_j\mathsf{X} \vee (\nu\mathsf{X})\Diamond_j\mathsf{X})$$

is satisfiable and its negation

$$\neg A, (\nu\mathsf{X})\Diamond_j\mathsf{X} \wedge (\mu\mathsf{X})\Box_j\mathsf{X}$$

is not valid. Therefore (1) $\neg A, (\nu\mathsf{X})\Diamond_j\mathsf{X}$ is not valid or (2) $\neg A, (\mu\mathsf{X})\Box_j\mathsf{X}$ is not valid. We distinguish these two cases.

(1) The Soundness of $(\nu.\mathrm{FL})$ implies that $\neg A, (\nu^h\mathsf{X})\Diamond_j\mathsf{X}$ is not valid where $h := \mathsf{f}(\neg A, (\nu\mathsf{X})\Diamond_j\mathsf{X})$ (note that $\mathsf{f}(A) = \mathsf{f}(\neg A)$). Therefore, there exists a Kripke structure K' with a world s such that both $\mathsf{K}', s \models A$ and $\mathsf{K}', s \not\models (\nu^h\mathsf{X})\Diamond_j\mathsf{X}$. Let K be the part of K' that is reachable from s. This Kripke structure satisfies our claim.

(2) There exists a Kripke structure K with a world s such that $\mathsf{K}, s \models A$ and $\mathsf{K}, s \not\models (\mu\mathsf{X})\Box_j\mathsf{X}$. Hence the claim is shown.

$\Box$

Definition 35. For every natural number n, we define a deductive system $\mathsf{T}^n_{\mu+}$ as follows. The definition of $\mathsf{T}^n_{\mu+}$ is analogous to that of $\mathsf{T}^\omega_{\mu+}$ except that the rule $(\nu.\omega)$ is replaced by the following finitary rule for greatest fixed points:

For all sequents Γ and X–positive formulae $\mathcal{A}$ of $\mathcal{L}^+_\mu$

$$\frac{\Gamma, (\nu^k\mathsf{X})\mathcal{A} \quad \text{for all } 0 < k \leq n}{\Gamma, (\nu\mathsf{X})\mathcal{A}} \quad (\nu.n).$$

Lemma 36. *An $\mathcal{L}_\mu$ formula B is valid if and only if it is derivable in $\mathsf{T}^n_{\mu+}$ where $n = \mathsf{f}(B)$.*

Proof. Let $\Delta, (\nu\mathsf{X})\mathcal{A}$ be as in the proof of Lemma 27. We only have to show that $\mathsf{f}(\Delta, (\nu\mathsf{X})\mathcal{A}) \leq \mathsf{f}(B)$. Assume Γ is a sequent occurring in a $\mathsf{T}^\omega_{\mu+}$ proof of B. We have $\Gamma^- \subset \mathbb{FL}(B)$. Therefore also $\mathbb{FL}(\Gamma^-) \subset \mathbb{FL}(B)$. Hence $\mathsf{f}(\Gamma) \leq \mathsf{f}(B)$. In particular, we conclude $\mathsf{f}(\Delta, (\nu\mathsf{X})\mathcal{A}) \leq \mathsf{f}(B)$. □

The completeness proof presented in [7] constructs a counter-model to any given non-provable $\mathcal{L}_\mu$ formula A. The universe of this counter-model consists of so-called A-saturated sets. An A-saturated set is a subset of $\mathbb{SC}(A)$ which satisfies certain closure conditions.

In view of Lemma 36 we can replace Clause 5 in the definition of the strong closure of D by

5'. If $(\nu\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A} \in \mathbb{SC}(D)$ and for every natural number $0 < n \leq \mathsf{f}(D)$ also $(\nu^n\mathsf{X})\mathcal{A} \in \mathbb{SC}(D)$.

With this new definition, the strong closure of a formula A is a finite set. Thus there can be only finitely many A-saturated sets. Hence, the construction in [7] gives us a finite counter-model. This results in a new proof-theoretic proof of the finite model property of the modal μ-calculus.

Theorem 37. *Every satisfiable $\mathcal{L}_\mu$ formula is satisfiable by a finite Kripke structure.*

The deductive systems $\mathsf{T}_{\mu+}$ and $\mathsf{T}^n_{\mu+}$ can be employed to perform a systematic proof search (with a loop check for the (μ)-rule). Thus we also have a proof-theoretic proof of the decidability of the validity problem of the modal μ-calculus. Note, however, that this procedure is by far not optimal since the validity problem of the μ-calculus is in EXPTIME [5].

Theorem 38. *It is decidable whether an $\mathcal{L}_\mu$ formula is valid.*

10 Conclusion

We study the proof theoretic relationship between two deductive systems for the modal μ-calculus. The infinitary system $\mathsf{T}^{\omega}_{\mu+}$ includes an ω-rule to derive the validity of greatest fixed points. In [7], $\mathsf{T}^{\omega}_{\mu+}$ is shown to be complete by a canonical counter-model construction. The infinitary system $\mathsf{T}^{\mathsf{pre}}_{\mu}$ is designed to decide the validity problem of the μ-calculus [4]. This system has a simple soundness proof.

The main technical result of the present paper is the following: if an $\mathcal{L}_\mu$ formula A is derivable in $\mathsf{T}^{\omega}_{\mu+}$, then A is also derivable in $\mathsf{T}^{\mathsf{pre}}_{\mu}$. This yields completeness of $\mathsf{T}^{\mathsf{pre}}_{\mu}$. Moreover, it provides a new soundness proof for $\mathsf{T}^{\omega}_{\mu+}$ which does not employ the finite model property.

Moreover, we introduce finitary cut-free systems $\mathsf{T}^{n}_{\mu+}$ such that an $\mathcal{L}_\mu$ formula A is valid if and only if it is derivable in $\mathsf{T}^{n}_{\mu+}$ where $n = \mathsf{f}(A)$. This result makes it possible to adapt the completeness proof in [7] such that finite counter-models are constructed. Hence, we obtain the finite model property of the μ-calculus by purely proof-theoretic means.

The size of the finite models we obtain is, however, not optimal. The crucial point is in the proof of Lemma 27 where we use a simple cardinality argument to prune the proof tree. Maybe that argument can be replaced by a more sophisticated one which considers the structure of $\mathcal{L}_\mu$ formulae. Then it might be possible to replace f by a function which is not exponential. That would finally provide better bounds on the size of the constructed models.

References

[1] Henrik R. Andersen, Colin Stirling, and Glynn Winskel. A compositional proof system for the modal μ–calculus. In *Proceedings of the IEEE Symposium on Logic in Computer Science LICS'94*, pages 144–153, 1994.

[2] Sergey Berezin and Dilian Gurov. A compositional proof system for the modal μ–calculus and CCS. Technical Report CMU–CS–97–105, Carnegie Mellon University, 1997.

[3] Julian Bradfield and Colin Stirling. Modal logics and mu-calculi: an introduction. In Jan Bergstra, Alban Ponse, and Scott Smolka, editors, *Handbook of Process Algebra*, pages 293–330. Elsevier, 2001.

[4] Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ-calculus. In *Proc. 26th Conf. on Foundations of Software Technology and Theoretical Computer Science, FSTTCS'06*, volume 4337 of *LNCS*, pages 274–285. Springer, 2006.

[5] E. Allen Emerson and Charanjit S. Jutla. The complexity of tree automata and logics of programs (extended abstract). In *29th Annual Symposium on Foundations of Computer Science FOCS*, pages 328–337. IEEE, 1988.

[6] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computing and System Science*, 18(2):194 – 211, 1979.

[7] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut-free systems for the propositional modal mu-calculus. submitted.

[8] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut-free axiomatizations for stratified modal fixed point logic. In Holger Schlingloff, editor, *Proceedings of the 4th Workshop Methods for Modalities*, pages 125–143. 2005.

[9] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut-free common knowledge. *Journal of Applied Logic*, to appear.

[10] Dexter Kozen. Results on the propositional modal μ–calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[11] Dexter Kozen. A finite model theorem for the propositional μ–calculus. *Studia Logica*, 47(3):233–241, 1988.

[12] Mathis Kretz. *Proof-theoretic aspects of modal logic with fixed points.* PhD thesis, Universisty of Bern, 2006.

[13] Daniel Leivant. A proof theoretic methodology for propositional dynamic logic. In *Proceedings of the International Colloquium on Formalization of Programming Concepts*, Springer LNCS, pages 356–373, 1981.

[14] Marino Miculan. On the formalization of the modal μ–calculus in the calculus of inductive constructions. *Information and Computation*, 164:199–231, 2001.

[15] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: Circular and tree-shaped proofs in the mu-calculus. In *Proc. FOSSACS'03*, Springer LNCS, pages 425–440, 2003.

[16] Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional modal mu–calculus. *Information and Computation*, 81:249–264, 1989.

[17] Igor Walukiewicz. A complete deductive system for the μ–calculus. In *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science*, pages 136–147. IEEE Computer Science Press, 1993.

[18] Igor Walukiewicz. Completeness of Kozen's axiomatization of the propositional μ–calculus. *Information and Computation*, 157:142–182, 2000.

Address
Thomas Studer
Institut für Informatik und angewandte Mathematik, Universität Bern
Neubrückstrasse 10, CH-3012 Bern, Switzerland
`tstuder@iam.unibe.ch`