

On the proof theory of the modal μ -calculus

Thomas Studer

April 10, 2008

Abstract

We study the proof-theoretic relationship between two deductive systems for the modal μ -calculus. First we recall an infinitary system which contains an omega rule allowing to derive the truth of a greatest fixed point from the truth of each of its (infinitely many) approximations. Then we recall a second infinitary calculus which is based on non-well-founded trees. In this system proofs are finitely branching but may contain infinite branches as long as some greatest fixed point is unfolded infinitely often along every branch. The main contribution of our paper is a translation from proofs in the first system to proofs in the second system. Completeness of the second system then follows from completeness of the first, and a new proof of the finite model property also follows as corollary.

1 Introduction

The propositional modal μ -calculus has been introduced by Kozen [13]. It is the extension of (multi-)modal logic by least and greatest fixed point operators. This provides a very expressive language which allows for arbitrary nestings of (possibly interleaved) fixed points. 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 [2].

There are two approaches to give infinitary axiomatizations for the modal μ -calculus. The first approach is to make use of a so-called ω -rule that has infinitely many premises to ensure that a fixed point is a least (respectively greatest) one. In the context of the μ -calculus, such a rule has first been introduced in [14]. There, Kozen 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 ω -rule which derives the validity of a greatest fixed point from the validity of all its (infinitely many) finite approximations. The resulting system is sound and complete. However, note that it makes crucial use of a cut rule. Jäger, Kretz and Studer [12] introduce the cut-free system $\mathsf{T}_{\mu+}^{\omega}$ which is also based on the ω -rule. Completeness of $\mathsf{T}_{\mu+}^{\omega}$ is established by a canonical counter-model construction.

The second approach is to define a deductive system $\mathsf{T}_{\mu}^{\text{pre}}$ such that in a proof search procedure fixed points are simply unfolded (which corresponds to closure of fixed points). This results in a so-called preproof which may have infinitely long branches. A global condition is then added which (roughly) says that in each infinite branch, there must be an outermost greatest fixed point unfolded infinitely many often. A tableau version of such a system has first been proposed by Niwinski and Walukiewicz [16]. They establish a completeness result for their system which is the starting point for the completeness proof of the finitary axiomatizations carried out by Walukiewicz [19, 20]. Dax, Hofmann, and Lange [7] present a proof system with infinitely long branches for the linear time μ -calculus. They also mention a related system for the full modal μ -calculus. We will employ their formulation of such an infinitary proof system.

The main contribution of the present paper is the embedding of $\mathsf{T}_{\mu+}^{\omega}$ in $\mathsf{T}_{\mu}^{\text{pre}}$. That means we provide a translation from proofs in the first system to proofs in the second. This provides completeness of $\mathsf{T}_{\mu}^{\text{pre}}$ since $\mathsf{T}_{\mu+}^{\omega}$ is complete. Moreover, we get a new proof of the finite model property of the μ -calculus. Note that these two results are not new. Already Niwinski and Walukiewicz [16] established a completeness result for a tableau version of $\mathsf{T}_{\mu}^{\text{pre}}$. Moreover, we do not get the exponential bound for the size of the model obtained by Emerson and Jutla [8]. However, our proof translation is a novel construction. We hope that it contributes to a better understanding of the proof theory of modal fixed point logics.

The field of proof theory for the modal μ -calculus and similar circular logics has been and still is very active. For instance, Sprenger and Dam [18] also compare two proof systems for the μ -calculus each using a different type of induction. Their starting system uses a local induction rule on ordinal variables (thus it is finitely branching) together with a cut-rule which allows for a straightforward translation from local to global induction. We study cut-free systems with an ω -rule (instead of an induction rule) which makes the construction more involved.

The proof theory of deductive systems with circular rules is studied by Aldwinckle and Cockett [1] who claim a cut-elimination result which implies a completeness theorem similar to ours. However, they only give an example of their cut-elimination procedure; but no precise description and also no

proofs are provided. Santocanale [17] also investigates a calculus of circular proofs and establishes a form of cut-elimination by exploring the categorical semantics. Therefore, his result applies to systems that are based on intuitionistic logic. Closely related to the modal μ -calculus are the systems for inductive definitions which Brotherston introduces in his recent PhD thesis [3]. He studies the proof theory of systems with induction rules, of infinitary systems, as well as of cyclic systems.

2 Language

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

Definition 2.1 (Language \mathcal{L}_μ). Let Φ be a countable set of atomic propositions and their negations $\mathbf{p}, \sim \mathbf{p}, \mathbf{q}, \sim \mathbf{q}, \mathbf{r}, \sim \mathbf{r}, \dots$, let \mathbf{V} be a set containing countably many variables and their negations $\mathbf{X}, \sim \mathbf{X}, \mathbf{Y}, \sim \mathbf{Y}, \mathbf{Z}, \sim \mathbf{Z}, \dots$, let $\mathbf{T} = \{\top, \perp\}$ be a set containing symbols for truth and falsehood and \mathbf{M} a set of indices. Define the formulae of the language \mathcal{L}_μ inductively as follows:

1. If P is an element of $\Phi \cup \mathbf{V} \cup \mathbf{T}$, then P is a formula of \mathcal{L}_μ .
2. If A and B are formulae of \mathcal{L}_μ , then so are $(A \wedge B)$ and $(A \vee B)$.
3. If A is a formula of \mathcal{L}_μ and $i \in \mathbf{M}$, then so are $\Box_i A$ and $\Diamond_i A$.
4. If A is a formula of \mathcal{L}_μ and the negated variable $\sim X$ does not occur in A , then $(\mu X)A$ and $(\nu X)A$ are also formulae of \mathcal{L}_μ .

In case there is no danger of confusion, we will omit parentheses in formulae. If the negated variable $\sim X$ does not occur in a formula A of \mathcal{L}_μ , 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 $\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$. Furthermore, we will call a formula A of \mathcal{L}_μ closed, if $fv(A) = \emptyset$ where $fv(A)$ are the free variables occurring in A . We write $\mathcal{A}[B/X]$ for the formula \mathcal{A} where the variable X has been substituted with B . If X is clear from the context, we simply write $\mathcal{A}[B]$.

Definition 2.2 (Language \mathcal{L}_μ^+). The formulae of the extended language \mathcal{L}_μ^+ are defined by adding the following clause to Definition 2.1:

5. If A is a formula of \mathcal{L}_μ^+ and the negated variable $\sim X$ does not occur in A , then for every natural number $k > 0$, $(\nu^k X)A$ is also a formula of \mathcal{L}_μ^+ .

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

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

We make use of the standard Kripke semantics for multi-modal fixed point logics to give meaning to \mathcal{L}_μ^+ formulae. That is we consider soundness and completeness with respect to a standard notion of validity, see for instance [2, 12, 13, 16].

3 The system $\mathbb{T}_{\mu+}^\omega$

The infinitary calculus $\mathbb{T}_{\mu+}^\omega$ is introduced in [12]. This deductive system provides a cut-free, sound and complete axiomatization for the modal μ -calculus. $\mathbb{T}_{\mu+}^\omega$ is formulated as a Tait-style system which derives finite sets $\Gamma, \Delta, \Sigma, \dots$ of \mathcal{L}_μ^+ 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}_μ^+ formulae, then $\diamond_i \Gamma := \{\diamond_i A_1, \dots, \diamond_i A_n\}$. We say a formula B is a *subformula of a sequent* Γ if B is a subformula of some formula $A \in \Gamma$.

Definition 3.1. The system $\mathbb{T}_{\mu+}^\omega$ is defined by the following inference rules:

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

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

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

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

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

$$\frac{\Gamma, A}{\diamond_i \Gamma, \square_i A, \Sigma} \text{ (}\square\text{)}$$

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

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

Fixed point rules: For all sequents Γ and X -positive formulae \mathcal{A} of \mathcal{L}_μ^+

$$\frac{\Gamma, \mathcal{A}[(\mu\mathsf{X})\mathcal{A}]}{\Gamma, (\mu\mathsf{X})\mathcal{A}} \quad (\mu) \qquad \frac{\Gamma, (\nu^k\mathsf{X})\mathcal{A} \quad \text{for all } k > 0}{\Gamma, (\nu\mathsf{X})\mathcal{A}} \quad (\nu.\omega)$$

Jäger, Kretz and Studer [12] present a canonical counter model construction which provides completeness of $\mathsf{T}_{\mu+}^\omega$.

Theorem 3.2. *The system $\mathsf{T}_{\mu+}^\omega$ is sound and complete for closed \mathcal{L}_μ formulae.*

Remark 3.3. The soundness proof given in [12] makes essential use of the finite model property of the modal μ -calculus. The crucial point is to show that if all premises of an instance of $(\nu.\omega)$ are valid, then also its conclusion is valid, see also [14]. Assume that the conclusion is not valid. By the finite model property, there exists a finite counter-model for it. Since in a finite model all closure ordinals of positive inductive definitions are finite, we obtain that there is a premise which is not satisfied in that counter-model. Thus not all premises are valid.

In the sequel we are going to introduce the notion of a thread in a branch of the proof tree.

Definition 3.4. 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 3.5. 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.

1. Assume r is not an instance of (\Box) . We have $(A, B) \in \mathsf{Con}(r)$ if $A = B$ 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 .
2. Assume r is an instance of (\Box) . We have $(\Box_i A, A) \in \mathsf{Con}(r)$ if $\Box_i A$ is the active formula in the conclusion of r and $(\Diamond_i B, B) \in \mathsf{Con}(r)$ if $\Diamond_i B \in \Diamond_i \Gamma$.

Definition 3.6. Assume we are given a branch $\Gamma_0, \Gamma_1, \dots$ 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, \dots such that $(A_i, A_{i+1}) \in \text{Con}(r_i)$ and $A_i \in \Gamma_i$ for every i .

Definition 3.7. An \mathcal{L}_μ formula A is called *well-named* if every variable is bound at most once. Note that for a bound variable X in a well-named formula A , there exists exactly one subformula of A that has the form $(\sigma X)\mathcal{B}$. We then call $(\sigma X)\mathcal{B}$ the *binding formula* of X . If the binding formula of a variable X is of the form $(\nu X)\mathcal{B}$, then X is called a ν *variable* in A . Let A be formula containing two bound variables X and Y . We say X is *higher* than Y if the binding formula of Y is a subformula of the binding formula of X .

In the sequel we consider only proofs for sequents of well-named formulae.

4 About threads

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

Definition 4.1 (Fischer–Ladner closure). Let D be a closed formula of \mathcal{L}_μ . 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 $(\mu X)\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A}[(\mu X)\mathcal{A}] \in \mathbb{FL}(D)$.
5. If $(\nu X)\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A}[(\nu X)\mathcal{A}] \in \mathbb{FL}(D)$ and $\mathcal{A}[\top/X] \in \mathbb{FL}(D)$.

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

Lemma 4.2 (see for instance [9]). *The Fischer-Ladner closure $\mathbb{FL}(D)$ of a closed formula D is finite.*

Definition 4.3 (Strong closure). Let D be a closed formula of \mathcal{L}_μ^+ . 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 $(\mu X)\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A}[(\mu X)\mathcal{A}] \in \mathbb{SC}(D)$.
5. If $(\nu X)\mathcal{A} \in \mathbb{SC}(D)$, then $(\nu^n X)\mathcal{A} \in \mathbb{SC}(D)$ for every natural number $n > 0$.
6. If $(\nu^1 X)\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A}[\top/X] \in \mathbb{SC}(D)$.
7. If n is a natural number greater than 0 and $(\nu^{n+1} X)\mathcal{A} \in \mathbb{SC}(D)$, then $\mathcal{A}[(\nu^n X)\mathcal{A}] \in \mathbb{SC}(D)$.

Lemma 4.4. *Let D be a closed formula of \mathcal{L}_μ . Then for all formulae A of \mathcal{L}_μ^+ we have*

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

In the sequel, we need to know whether two formulae of the form $(\nu^k X)\mathcal{A}$, that both occur in a $\mathbb{T}_{\mu+}^\omega$ proof, originate from the same greatest fixed point. Therefore, we introduce annotated threads that support the bookkeeping about applications of $(\nu.\omega)$.

Definition 4.5. Assume that we are given a $\mathbb{T}_{\mu+}^\omega$ proof S for an \mathcal{L}_μ formula D . Further we are given a thread B_1, B_2, \dots in that proof. We *annotate* all subformulae of the form $(\nu^k X)\mathcal{A}$ occurring in it as follows.

1. B_1 is an \mathcal{L}_μ formula. Thus it cannot contain a subformula of the form $(\nu^k X)\mathcal{A}$.
2. If $B_i = (\nu X)\mathcal{A}$ and $B_{i+1} = (\nu^k X)\mathcal{A}$, then we annotate $(\nu^k X)\mathcal{A}$ with the subtree T of S that is given by the node of S in which B_{i+1} occurs.
3. If $(\nu^k X)\mathcal{A}$ is a subformula of both B_i and B_{i+1} and $(\nu^k X)\mathcal{A}$ is annotated with a tree T in B_i , then $(\nu^k X)\mathcal{A}$ is also annotated with T in B_{i+1} .
4. If $B_i = (\nu^{k+1} X)\mathcal{A}$ and $B_{i+1} = \mathcal{A}[(\nu^k X)\mathcal{A}]$ and $(\nu^{k+1} X)\mathcal{A}$ is annotated with a tree T in B_i , then $(\nu^k X)\mathcal{A}$ is also annotated with T in B_{i+1} .

We write $(\nu_T^k X)\mathcal{A}$ if $(\nu^k X)\mathcal{A}$ is annotated with T .

Lemma 4.6. *Assume that we are given a $\mathbb{T}_{\mu+}^\omega$ proof for an \mathcal{L}_μ formula C_1 . Further assume that C_1, \dots, C_n, \dots is an annotated thread in this proof. Let $(\nu_T^k X)\mathcal{A}$ be a subformula of C_n . Then $(\nu_T^k X)\mathcal{A}$ does not have a strict subformula of the form $(\nu_T^h X)\mathcal{B}$.*

Proof. By induction on n . Since C_1 is an \mathcal{L}_μ formula, it cannot have a subformula of the form $(\nu_T^k \mathbf{X})\mathcal{A}$. To show the induction step, we assume that $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}(C_n)$ for $n > 1$. We distinguish the cases for the rule of which C_n belongs to a premise.

1. $(\nu.\omega)$. If the rule introduces the annotation T , then no strict subformula will be annotated by T and the claim holds. If another annotation is introduced, then $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}(C_{n-1})$ and the claim immediately follows by the induction hypothesis.
2. (μ) . If $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}(C_{n-1})$, then apply the induction hypothesis. If $(\nu_T^k \mathbf{X})\mathcal{A} \notin \text{sub}(C_{n-1})$, then C_{n-1} is of the form $(\mu \mathbf{Y})\mathcal{B}$ and we have $(\mu \mathbf{Y})\mathcal{B} \in \text{sub}((\nu_T^k \mathbf{X})\mathcal{A})$. Let $(\nu_T^k \mathbf{X})\hat{\mathcal{A}}$ be $(\nu_T^k \mathbf{X})\mathcal{A}$ where $(\mu \mathbf{Y})\mathcal{B}$ is replaced by \mathbf{Y} . Then $(\nu_T^k \mathbf{X})\hat{\mathcal{A}} \in \text{sub}(C_{n-1})$ and the claim follows by the induction hypothesis.
3. $(\nu.1)$ and $(\nu.k+1)$. Reasoning similar to the previous case provides the desired result.
4. For all other rules we have that $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}(C_{n-1})$ and the claim immediately follows by the induction hypothesis. \square

Assume that we are given a thread A_0, A_1, \dots . We write this thread as $D, \dots, (\nu_T^{k+1} \mathbf{X})\mathcal{A}, \mathcal{A}[(\nu_T^k \mathbf{X})\mathcal{A}] = B_1, \dots, B_n, \dots$ if $A_0 = D$ and if there is an i such that $A_i = (\nu_T^{k+1} \mathbf{X})\mathcal{A}$, $A_{i+1} = \mathcal{A}[(\nu_T^k \mathbf{X})\mathcal{A}]$, and $A_{i+j} = B_j$ for all natural numbers $j \geq 1$.

Lemma 4.7. *Assume that we are given a $\top_{\mu+}^\omega$ proof for an \mathcal{L}_μ formula D . Further assume that $D, \dots, (\nu_T^{k+1} \mathbf{X})\mathcal{A}, \mathcal{A}[(\nu_T^k \mathbf{X})\mathcal{A}] = B_1, \dots, B_n, \dots$ is an annotated thread in this proof. We find that $(\nu_T^h \mathbf{X})\mathcal{B} \notin \text{sub}(B_n)$ for any $h > k$ and any \mathcal{B} .*

Proof. Induction on n . Case $n = 1$. Assume $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(B_1)$ with $h > k$. Then either (i) $(\nu_T^h \mathbf{X})\mathcal{B}$ is a strict subformula of $(\nu_T^k \mathbf{X})\mathcal{A}$, (ii) $(\nu_T^k \mathbf{X})\mathcal{A}$ is a strict subformula of $(\nu_T^h \mathbf{X})\mathcal{B}$, or (iii) none of the two holds and $(\nu_T^h \mathbf{X})\mathcal{B}$ is a subformula of \mathcal{A} . (i) and (ii) are not possible by Lemma 4.6. (iii) implies that $(\nu_T^h \mathbf{X})\mathcal{B}$ is a strict subformula of $(\nu_T^{k+1} \mathbf{X})\mathcal{A}$ which again cannot be by Lemma 4.6.

Case $n > 1$. Assume $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(B_n)$ with $h > k$. Consider the rule ρ in which B_n belongs to a premise and B_{n-1} to the conclusion. If B_n is a side formula in ρ , then $B_n = B_{n-1}$ and thus $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(B_{n-1})$. This cannot be by the induction hypothesis. Therefore, we assume that B_n is active in ρ . We have the following cases:

1. ρ is an instance of (\vee) , (\wedge) , or (\square) . Again, we find $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(B_{n-1})$ which cannot be by the induction hypothesis.
2. ρ is an instance of $(\nu.k + 1)$ with the distinguished formula $(\nu_T^{h+1} \mathbf{X})\mathcal{B}$. This means $(\nu_T^{h+1} \mathbf{X})\mathcal{B} \in \text{sub}(B_{n-1})$ which cannot be by the induction hypothesis.
3. ρ is an instance of $(\nu.\omega)$ with the distinguished formula $(\nu \mathbf{X})\mathcal{B}$. This cannot be since the annotation T has been introduced before B_{n-1} .
4. ρ is an instance of $(\nu.1)$. That means $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(\mathcal{C}[\top/\mathbf{Y}])$. There are two cases: (i) $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(\mathcal{C})$. Then $(\nu_T^h \mathbf{X})\mathcal{B} \in \text{sub}(B_{n-1})$ which cannot be by the induction hypothesis. (ii) There is a $\hat{\mathcal{B}}$ such that $\hat{\mathcal{B}}[\top/\mathbf{Y}] = \mathcal{B}$ and $(\nu_T^h \mathbf{X})\hat{\mathcal{B}} \in \text{sub}(\mathcal{C})$. Then $(\nu_T^h \mathbf{X})\hat{\mathcal{B}} \in \text{sub}(B_{n-1})$ which cannot be by the induction hypothesis.
5. The remaining cases are similar to the previous case. \square

Lemma 4.8. *Assume that we are given a $\top_{\mu+}^\omega$ proof for an \mathcal{L}_μ formula D . Further assume that $D, \dots, (\nu_T^{k+1} \mathbf{X})\mathcal{A}, \mathcal{A}[(\nu_T^k \mathbf{X})\mathcal{A}] = B_1, \dots, B_n, \dots$ is an annotated thread in this proof such that*

$$(\mu \mathbf{Y})\mathcal{B} \in \text{sub}(B_n) \text{ as well as } (\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}((\mu \mathbf{Y})\mathcal{B}).$$

Then there is a formula $(\mu \mathbf{Y})\mathcal{B}' \in \text{sub}((\nu_T^{k+1} \mathbf{X})\mathcal{A})$ with $\mathbf{X} \in \text{fv}(\mathcal{B}')$.

Proof. Induction on n . Case $n = 1$. We have $B_1 = \mathcal{A}[(\nu_T^k \mathbf{X})\mathcal{A}]$. Note that by Lemma 4.6 $(\nu_T^k \mathbf{X})\mathcal{A} \notin \text{sub}((\nu_T^{k+1} \mathbf{X})\mathcal{A})$. Therefore $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}((\mu \mathbf{Y})\mathcal{B})$ is only possible if there is a \mathcal{B}' as required.

Case $n > 1$. $(\mu \mathbf{Y})\mathcal{B} \in \text{sub}(B_n)$ originates from a $(\mu \mathbf{Y})\hat{\mathcal{B}} \in \text{sub}(B_{n-1})$. Moreover, we have that $(\nu_T^k \mathbf{X})\mathcal{A} \in \text{sub}((\mu \mathbf{Y})\mathcal{B})$ cannot have been generated by an application of $(\nu.k + 1)$. If this were the case, then $B_{n-1} = (\nu_T^{k+1} \mathbf{X})\mathcal{A}$ which contradicts Lemma 4.7. Thus there exists $(\nu_T^k \mathbf{X})\hat{\mathcal{A}} \in \text{sub}((\mu \mathbf{Y})\hat{\mathcal{B}})$ and the claim follows by the induction hypothesis. \square

Lemma 4.9. *Assume that we are given a $\top_{\mu+}^\omega$ proof for a well-named closed \mathcal{L}_μ formula B_1 . Let B_1, \dots, B_n, \dots be an annotated thread in this proof. Let \mathbf{X} be a variable occurring in B_1 . If $(\sigma \mathbf{Y})\mathcal{B} \in \text{sub}(B_n)$ and $\mathbf{X} \in \text{fv}((\sigma \mathbf{Y})\mathcal{B})$, then \mathbf{X} is higher than \mathbf{Y} in B_1 .*

Proof. Induction on n . If $n = 1$, then the claim follows from the fact that B_1 is closed and the definition of free variable. For $n > 1$, we distinguish the cases for the rule of which B_n belongs to a premise.

1. (μ) . In this case we have $B_n = \mathcal{C}[(\mu Z)\mathcal{C}]$ as well as $B_{n-1} = (\mu Z)\mathcal{C}$. If $(\sigma Y)\mathcal{B} \in \mathbf{sub}((\mu Z)\mathcal{C})$, then we can apply the induction hypothesis. If $(\sigma Y)\mathcal{B} \notin \mathbf{sub}((\mu Z)\mathcal{C})$, then $(\mu Z)\mathcal{C}$ is a strict subformula of $(\sigma Y)\mathcal{B}$. Let $\hat{\mathcal{B}}$ be such that $\hat{\mathcal{B}}[(\mu Z)\mathcal{C}/Z] = \mathcal{B}$. Then $(\sigma Y)\hat{\mathcal{B}} \in \mathbf{sub}(B_{n-1})$. Note that B_{n-1} is closed since B_1 is closed. Thus $X \in \mathit{fv}((\sigma Y)\mathcal{B})$ implies $X \in \mathit{fv}((\sigma Y)\hat{\mathcal{B}})$. Now the claim follows by the induction hypothesis.
2. $(\nu.k + 1)$. Similar to the case for (μ) .
3. For all other rules the claim immediately follows by the induction hypothesis. \square

Lemma 4.10. *Assume we are given a $\mathbb{T}_{\mu+}^\omega$ proof for an \mathcal{L}_μ formula D . Further assume that $D, \dots, (\nu_T^{k+1}X)\mathcal{A}, \mathcal{A}[(\nu_T^kX)\mathcal{A}] = B_1, \dots, B_i, \dots, B_n, \dots$ is an annotated thread in this proof. If $(\nu_T^kX)\mathcal{C} \in \mathbf{sub}(B_n)$ for some \mathcal{C} , then for every B_i with $1 \leq i \leq n$ there is a \mathcal{B} such that $(\nu_T^kX)\mathcal{B} \in \mathbf{sub}(B_i)$.*

Proof. Induction on n : The case $n = 1$ is trivial. To show the induction step, let $n > 1$ and $(\nu_T^kX)\mathcal{C} \in \mathbf{sub}(B_n)$. Again, distinguish the cases for the rule of which B_n belongs to a premise.

1. $(\nu.k + 1)$ with distinguished formula $(\nu_T^{k+1}X)\mathcal{C}$. This cannot be by Lemma 4.7.
2. $(\nu.\omega)$ that has $(\nu_T^kX)\mathcal{C}$ as a premise. Then, according to Definition 4.5, T is the subtree given by that premise. This cannot be since T already occurs earlier in the thread.
3. For all other rules the claim immediately follows by the induction hypothesis. \square

Lemma 4.11. *Assume that we are given a $\mathbb{T}_{\mu+}^\omega$ proof for an \mathcal{L}_μ formula E . Further assume there is an annotated thread in this proof of the form*

$$E, \dots, (\nu_T^iX)\mathcal{A}, \dots, (\mu Y)\mathcal{B}, \dots, (\nu_T^jX)\mathcal{C}, \dots$$

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

Proof. First, observe that in a thread of the form

$$E, \dots, (\nu_T^iX)\mathcal{A}, B_1, \dots, B_n, (\nu_T^jX)\mathcal{C}, \dots$$

for each $j < k < i$ there is a $1 \leq l \leq n$ such that $B_l = (\nu_T^kX)\mathcal{C}'$ for some \mathcal{C}' . Thus the given thread is of the form

$$E, \dots, (\nu_T^{k+1}X)\mathcal{A}', \mathcal{A}'[(\nu_T^kX)\mathcal{A}'], \dots, (\mu Y)\mathcal{B}, \mathcal{B}[(\mu Y)\mathcal{B}], \dots, (\nu_T^kX)\mathcal{C}', \dots$$

for some natural number k and formulae \mathcal{A}' and \mathcal{C}' . By Lemma 4.10 we know $(\nu_T^k \mathbf{X})\mathcal{D} \in \text{sub}((\mu \mathbf{Y})\mathcal{B})$ for some \mathcal{D} . Thus, by Lemma 4.8 there is a formula $(\mu \mathbf{Y})\mathcal{B}' \in \text{sub}((\nu_T^{k+1} \mathbf{X})\mathcal{A}')$ with $\mathbf{X} \in \text{fv}(\mathcal{B}')$. By Lemma 4.9, we conclude that \mathbf{X} is higher than \mathbf{Y} in E . \square

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

Dax, Hofmann, and Lange [7] present an infinitary proof system for the linear time μ -calculus. In the section ‘Further Work’ of their paper, they mention how a corresponding infinitary system for the modal μ -calculus can be formulated. Let us now present such a deductive system which we call $\mathsf{T}_\mu^{\text{pre}}$. This section is basically taken from [7].

Definition 5.1. A *preproof* for a sequent Γ of \mathcal{L}_μ 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}_μ , all \mathbf{p} in Φ , and all \mathbf{X} in \mathbf{V}

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

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

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

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

$$\frac{\Gamma, A}{\diamond_i \Gamma, \square_i A, \Sigma} \quad (\square)$$

Fixed point rules: For all sequents Γ and \mathbf{X} -positive formulae \mathcal{A} of \mathcal{L}_μ

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

We make use of the notions of *distinguished*, *active*, and *side* formulae for $\mathsf{T}_\mu^{\text{pre}}$ preproofs, too. We will also consider threads in the context of $\mathsf{T}_\mu^{\text{pre}}$ preproofs. Note that $\mathsf{T}_\mu^{\text{pre}}$ preproofs may have infinitely long branches and thus also threads may be infinite sequences. We have the following fact about threads.

Lemma 5.2. *Assume we are given an infinite branch of a preproof for an \mathcal{L}_μ sequent Γ . Assume we are given a thread in this branch such that infinitely many of its formulae are distinguished formulae of applications of (μ) and (ν) . Then there is a unique bound variable X such that*

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

Definition 5.3. *Assume we are given an infinite branch of a preproof for an \mathcal{L}_μ sequent Γ . Assume we are given a thread in this branch such that infinitely many of its formulae are distinguished formulae of applications of (μ) and (ν) . Such a thread is called a ν -thread if the unique variable given by the previous lemma is a ν variable in Γ .*

Definition 5.4. *A $\mathbb{T}_\mu^{\text{pre}}$ proof for a sequent Γ of \mathcal{L}_μ 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 $\mathbb{T}_\mu^{\text{pre}} \vdash \Gamma$ if there exists a $\mathbb{T}_\mu^{\text{pre}}$ proof for Γ .*

6 Embedding $\mathbb{T}_{\mu+}^\omega$ in $\mathbb{T}_\mu^{\text{pre}}$

We show how to obtain a $\mathbb{T}_\mu^{\text{pre}}$ proof for an \mathcal{L}_μ sequent Γ from given a $\mathbb{T}_{\mu+}^\omega$ proof of Γ . Let us first illustrate our approach by the following simple example. Assume we are given the following $\mathbb{T}_{\mu+}^\omega$ proof of $(\mu X)\Box X, (\nu Y)\Diamond Y$:

$$\begin{array}{c}
 \frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top} \\
 \frac{\frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top}}{(\mu X)\Box X, \Diamond \top} \\
 \frac{\frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top}}{(\mu X)\Box X, (\nu^1 Y)\Diamond Y} \\
 \frac{\frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top}}{(\mu X)\Box X, \Diamond((\nu^1 Y)\Diamond Y)} \\
 \frac{\frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top}}{(\mu X)\Box X, \Diamond((\nu^1 Y)\Diamond Y)} \\
 \frac{\frac{(\mu X)\Box X, \top}{\Box((\mu X)\Box X), \Diamond \top}}{(\mu X)\Box X, (\nu^2 Y)\Diamond Y} \quad \dots \\
 \hline
 (\mu X)\Box X, (\nu Y)\Diamond Y
 \end{array}$$

Starting from this proof we can construct a $\mathbb{T}_\mu^{\text{pre}}$ proof as follows. We take the branch through the premise $(\nu^2 Y)\Diamond Y$ of the infinitary greatest fixed point rule. In that branch we drop all the iteration numbers. That is we replace

all subexpressions of the form $(\nu^k \mathbf{X})\mathcal{C}$ by $(\nu \mathbf{X})\mathcal{C}$. This gives us the following:

$$\frac{\frac{\frac{\frac{(\mu \mathbf{X})\Box \mathbf{X}, \top}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond \top}}{(\mu \mathbf{X})\Box \mathbf{X}, \Diamond \top}}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu^1 \mathbf{Y})\Diamond \mathbf{Y}}}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond((\nu^1 \mathbf{Y})\Diamond \mathbf{Y})}}{\frac{(\mu \mathbf{X})\Box \mathbf{X}, \Diamond((\nu^1 \mathbf{Y})\Diamond \mathbf{Y})}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu^2 \mathbf{Y})\Diamond \mathbf{Y}}}}{\frac{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}}}} \quad \Rightarrow \quad \frac{\frac{\frac{\frac{(\mu \mathbf{X})\Box \mathbf{X}, \top}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond \top}}{(\mu \mathbf{X})\Box \mathbf{X}, \Diamond \top}}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}}{\frac{(\mu \mathbf{X})\Box \mathbf{X}, \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}}}{\frac{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}}}}$$

Note that dropping the iteration numbers in the sequents $(\mu \mathbf{X})\Box \mathbf{X}, (\nu^2 \mathbf{Y})\Diamond \mathbf{Y}$ and $(\mu \mathbf{X})\Box \mathbf{X}, (\nu^1 \mathbf{Y})\Diamond \mathbf{Y}$ makes them identical. Therefore we can loop between these two sequents which results in the following infinite $\mathbb{T}_\mu^{\text{pre}}$ proof:

$$\frac{\frac{\frac{\vdots}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}{\Box((\mu \mathbf{X})\Box \mathbf{X}), \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}}{(\mu \mathbf{X})\Box \mathbf{X}, \Diamond((\nu \mathbf{Y})\Diamond \mathbf{Y})}}{(\mu \mathbf{X})\Box \mathbf{X}, (\nu \mathbf{Y})\Diamond \mathbf{Y}}}}$$

A crucial ingredient to this construction is a cardinality argument which shows that after dropping the iteration numbers, there will be two identical sequents with the same distinguished formula. The following function provides an upper bound on the number of different sequents (taking also into account the different possibilities for the distinguished formula) that may occur in a proof of Γ after dropping the iteration numbers.

Definition 6.1. Let f the function assigning to each \mathcal{L}_μ^+ sequent Γ a natural number as follows:

$$f(\Gamma) := |\mathbb{FL}(\Gamma^-)| \cdot 2^{|\mathbb{FL}(\Gamma^-)|} + 1$$

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

Definition 6.2. Assume that we are given an $\mathbb{T}_{\mu+}^\omega$ proof for an \mathcal{L}_μ^+ sequent Γ . The *pruned proof tree* \mathcal{PPT} of this given proof is a tree labeled by \mathcal{L}_μ^+ 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) , (\square) , $(\nu.1)$, $(\nu.k+1)$, or (μ) , then we construct the pruned proof trees of the proofs for the premises of this last rule. \mathcal{PPT} is now given as the disjoint union of these pruned proof trees with the addition of a new root node labeled by Γ .
3. If the last rule was an instance of $(\nu.\omega)$ with the conclusion $\Sigma, (\nu\mathbf{X})\mathcal{A}$, then \mathcal{PPT} is given as the pruned proof tree for the premise $\Sigma, (\nu^k\mathbf{X})\mathcal{A}$ where $k = f(\Gamma)$.

We will annotate subformulae of the form $(\nu^k\mathbf{X})\mathcal{A}$ occurring in a thread in \mathcal{PPT} as in the corresponding thread of the $\mathbb{T}_{\mu+}^\omega$ proof.

Example 6.3. If π is a $\mathbb{T}_{\mu+}^\omega$ proof, then we denote the pruned proof tree of π by $\mathcal{PPT}(\pi)$. Assume we are given the following $\mathbb{T}_{\mu+}^\omega$ proof:

$$\begin{array}{c}
 \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_1 \\ \text{---} \\ \Gamma, A \end{array} \qquad \qquad \qquad \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_{2,i} \\ \text{---} \\ \Gamma, (\nu^i\mathbf{X})\mathcal{B} \end{array} \\
 \vdots \qquad \qquad \qquad \vdots \\
 \hline
 \Gamma, A \wedge (\nu\mathbf{X})\mathcal{B}
 \end{array}$$

Let us now construct the corresponding pruned proof tree. In a first step we obtain:

$$\begin{array}{c}
 \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \mathcal{PPT}(\pi_1) \\ \text{---} \\ \Gamma, A \end{array} \qquad \qquad \qquad \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \mathcal{PPT}(\pi_2) \\ \text{---} \\ \Gamma, (\nu\mathbf{X})\mathcal{B} \end{array} \\
 \hline
 \Gamma, A \wedge (\nu\mathbf{X})\mathcal{B}
 \end{array}$$

where π_2 is the subproof deriving $\Gamma, (\nu\mathbf{X})\mathcal{B}$. When we construct $\mathcal{PPT}(\pi_2)$, then we get:

$$\begin{array}{c}
 \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \mathcal{PPT}(\pi_1) \\ \text{---} \\ \Gamma, A \end{array} \qquad \qquad \qquad \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \mathcal{PPT}(\pi_{2,k}) \\ \text{---} \\ \Gamma, (\nu^k\mathbf{X})\mathcal{B} \end{array} \\
 (*) \quad \hline
 \Gamma, A \wedge (\nu\mathbf{X})\mathcal{B}
 \end{array}$$

where $k = f(\Gamma, (\nu\mathbf{X})\mathcal{B})$. We make the following observations.

1. The pruned proof tree is a finite tree. When an instance of $(\nu.\omega)$ is treated, then a branch is selected and only that branch contributes to the construction of the pruned proof tree. Therefore there is no infinite branching in the pruned proof tree.
2. In the construction of $\mathcal{PPT}(\pi_2)$, the end-sequent $\Gamma, (\nu\mathbf{X})\mathcal{B}$ has been dropped. The pruned proof tree $\mathcal{PPT}(\pi_2)$ ends with $\Gamma, (\nu^k\mathbf{X})\mathcal{B}$. Therefore, at this point, $(*)$ is not an instance of (\wedge) .
3. In the sequel we will construct a $\mathbb{T}_\mu^{\text{pre}}$ preproof from a given pruned proof tree \mathcal{PPT} . In the course of this construction we will drop all the iteration numbers occurring in the sequents of \mathcal{PPT} . Note that dropping the iterations number in the above example makes $(*)$ an instance of (\wedge) .
4. If we had kept both the end-sequent $\Gamma, (\nu\mathbf{X})\mathcal{B}$ and its premise $\Gamma, (\nu^k\mathbf{X})\mathcal{B}$ in $\mathcal{PPT}(\pi_2)$, then dropping the iteration numbers would leave us with an inference where the premise and the conclusion are equal. Thus we can drop the end-sequent.

Lemma 6.4. *Assume that we are given a pruned proof tree \mathcal{PPT} of a $\mathbb{T}_{\mu+}^\omega$ proof of an \mathcal{L}_μ 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 $(\nu.1)$ for some $1 \leq h < n$. Then there are $1 \leq i, j \leq n$ and natural numbers $k \neq l$ with*

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

such that

- (1) there is a thread containing both $(\nu_T^l \mathbf{X})\mathcal{A}$ and $(\nu_T^k \mathbf{X})\mathcal{A}$, and
- (2) $(\nu_T^l \mathbf{X})\mathcal{A}$ is the distinguished formula of Γ_i , and
- (3) $(\nu_T^k \mathbf{X})\mathcal{A}$ is the distinguished formula of Γ_j , and
- (4) $\Gamma_i^- = \Gamma_j^-$.

Proof. Assume that we are given a branch of \mathcal{PPT} in which $(\nu_T^l \mathbf{X})\mathcal{B}$, for some \mathcal{B} , occurs as distinguished formula in the label of a node. Since Γ is an \mathcal{L}_μ sequent, the corresponding branch in the $\mathbb{T}_{\mu+}^\omega$ proof must contain a node labeled by $\Delta, (\nu\mathbf{X})\mathcal{C}$ such that $(\nu\mathbf{X})\mathcal{C}$ is the distinguished formula of the node and its premises are of the form $\Delta, (\nu_T^g \mathbf{X})\mathcal{C}$, that is $\Delta, (\nu\mathbf{X})\mathcal{C}$ is the conclusion of an instance of $(\nu.\omega)$. By the definition of pruned proof tree, there is a Γ_m in the given branch with $\Gamma_m = \Delta, (\nu_T^{f(\Delta, (\nu\mathbf{X})\mathcal{C})} \mathbf{X})\mathcal{C}$. Because of the annotation

we know that there is a thread in the given branch containing $(\nu_T^{\mathbf{f}(\Delta, (\nu\mathbf{X})\mathcal{C})}\mathbf{X})\mathcal{C}$ in Γ_m and $(\nu_T^1\mathbf{X})\mathcal{B}$. Thus for each $o \leq \mathbf{f}(\Delta, (\nu\mathbf{X})\mathcal{C})$ there exists $m \leq q_o \leq n$ such that

$$\Gamma_{q_o} = (\nu_T^o\mathbf{X})\mathcal{A}_o, \Delta_o \text{ where } (\nu_T^o\mathbf{X})\mathcal{A}_o \text{ is the distinguished formula of } \Gamma_{q_o} \quad (1)$$

and

$$\text{there is a thread containing all these } (\nu_T^o\mathbf{X})\mathcal{A}_o. \quad (2)$$

Lemma 4.4 implies $\Delta_h^- \subseteq \mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})$ for $m \leq h \leq n$. Thus there are at most $2^{|\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})|}$ different possibilities what the sets Δ_o^- can be. Moreover, there are at most $|\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})|$ different possibilities what the formula $((\nu_T^o\mathbf{X})\mathcal{A}_o)^-$ can be. Thus there are at most $|\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})| \cdot 2^{|\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})|}$ different possibilities what $(\nu_T^o\mathbf{X})\mathcal{A}_o^-, \Delta_o^-$ can be. Since

$$\mathbf{f}(\Delta, (\nu\mathbf{X})\mathcal{C}) > |\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})| \cdot 2^{|\mathbb{FL}(\Delta, (\nu\mathbf{X})\mathcal{C})|},$$

and for each $o \leq \mathbf{f}(\Delta, (\nu\mathbf{X})\mathcal{C})$ there exists q_o satisfying (1), there must exist $k, l < \mathbf{f}(\Delta, (\nu\mathbf{X})\mathcal{C})$ with $k \neq l$ such that $\Gamma_{q_l}^- = \Gamma_{q_k}^-$ holds. Thus we let $i = q_l$ and $j = q_k$. By (1) we obtain that $(\nu_T^l\mathbf{X})\mathcal{A}$ is the distinguished formula of Γ_i and $(\nu_T^k\mathbf{X})\mathcal{A}$ is the distinguished formula of Γ_j . Moreover (2) guarantees the existence of the required thread. \square

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

Definition 6.5. Assume that we are given a pruned proof tree \mathcal{PPT} of a $\Gamma_{\mu+}^\omega$ proof for an \mathcal{L}_μ sequent Γ . We simultaneously construct a Γ_μ^{pre} 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 $\text{origin}(a) := b$ where b is the root of \mathcal{PPT} .
2. A node $a \in \mathcal{PRE}$ is labeled by the \mathcal{L}_μ sequent Δ^- where Δ is the label of $\text{origin}(a)$ in \mathcal{PPT} .
3. A node $a \in \mathcal{PRE}$ has child nodes c_1, \dots, c_n if $\text{origin}(a)$ has n child nodes b_1, \dots, b_n in \mathcal{PPT} . For $1 \leq i \leq n$, we define

- (a) $\text{origin}(c_i) := d$ if b_i has an ancestor node $d \in \mathcal{PPT}$ and there are \mathcal{L}_μ^+ sequents Δ_{b_i} and Δ_d such that $(\nu_T^l\mathbf{X})\mathcal{A}$ is the distinguished formula of d , $(\nu_T^k\mathbf{X})\mathcal{A}$ is the distinguished formula of b_i , there is a thread containing both of these formulae, and $\Delta_{b_i}^- = \Delta_d^-$ as well as $\text{label}(b_i) = \Delta_{b_i}, (\nu_T^k\mathbf{X})\mathcal{A}$ and $\text{label}(d) = \Delta_d, (\nu_T^l\mathbf{X})\mathcal{A}$.

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

Definition 6.5 indeed constructs a $\mathbb{T}_\mu^{\text{pre}}$ preproof. The only critical point is if \mathcal{PPT} contains a branch with an instance of $(\nu.1)$. However, Lemma 6.4 guarantees that such a branch is always transformed into an infinite branch in the $\mathbb{T}_\mu^{\text{pre}}$ preproof.

Theorem 6.6. *For all closed \mathcal{L}_μ formulae D we have*

$$\mathbb{T}_{\mu+}^\omega \vdash D \implies \mathbb{T}_\mu^{\text{pre}} \vdash D.$$

Proof. Given the $\mathbb{T}_{\mu+}^\omega$ proof of D , we can construct the corresponding pruned proof tree and from that a preproof of D according to the Definitions 6.2 and 6.5. It remains to show that every infinite path of the preproof contains a ν -thread. First, we notice that an infinite branch can only occur because of Condition 3a in Definition 6.5. Assume that we are given an infinite branch. Let B_1, B_2, \dots be a thread of this branch that contains a formula of the form $(\nu X)\mathcal{A}$ for which Condition 3a has been applied. Suppose that this thread contains the formula $(\mu Y)\mathcal{B}$ infinitely often. Then this thread must be of the form

$$\dots, (\nu Z)\mathcal{C}, \dots, (\mu Y)\mathcal{B}, \dots, (\nu Z)\mathcal{C}, \dots \quad (3)$$

such that there is a loop because of Condition 3a for $(\nu Z)\mathcal{C}$. Thus there must be a thread of the form

$$\dots, (\nu_T^i Z)\mathcal{C}, \dots, (\mu Y)\mathcal{B}, \dots, (\nu_T^j Z)\mathcal{C}, \dots$$

in the original $\mathbb{T}_{\mu+}^\omega$ proof of D (note that this thread need not be the same as (3), there may be different formulae at the \dots positions). Applying Lemma 4.11 to this thread yields that Z is higher than Y in D . Thus the infinite branch contains a ν -thread. \square

7 Applications

Dax et al. [7] provide a simple soundness proof of their system for the linear time μ -calculus. A straightforward adaptation of this proof shows the soundness of $\mathbb{T}_\mu^{\text{pre}}$. Simply replace the case for the 'next'-rule by an appropriate treatment of (\square) .

Theorem 7.1. *The system $\mathbb{T}_\mu^{\text{pre}}$ is sound.*

Completeness of $\mathbb{T}_{\mu+}^\omega$ is established in [12] by a canonical counter-model construction. We immediately obtain the following corollary about soundness and completeness of $\mathbb{T}_{\mu+}^\omega$ and $\mathbb{T}_\mu^{\text{pre}}$ with respect to \mathcal{L}_μ formulae.

Corollary 7.2. *Let A be an \mathcal{L}_μ formula. We have*

$$A \text{ is valid} \implies \mathbb{T}_{\mu+}^\omega \vdash A \implies \mathbb{T}_\mu^{\text{pre}} \vdash A \implies A \text{ is valid.}$$

Note that Corollary 7.2 provides soundness of $\mathbb{T}_{\mu+}^\omega$ without referring to the finite model property of the modal μ -calculus. This is interesting insofar as the soundness proof for $\mathbb{T}_{\mu+}^\omega$ in [12] essentially uses the finite model property to show that ω many premises are enough in the $(\nu.\omega)$ -rule, see Remark 3.3.

We can employ Corollary 7.2 to obtain the finite model property of the modal μ calculus.

Definition 7.3. For every natural number n , we define a deductive system $\mathbb{T}_{\mu+}^n$ as follows. The definition of $\mathbb{T}_{\mu+}^n$ is analogous to that of $\mathbb{T}_{\mu+}^\omega$ 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}_μ^+

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

Lemma 7.4. *An \mathcal{L}_μ formula B is valid if and only if it is derivable in $\mathbb{T}_{\mu+}^n$ where $n = \mathfrak{f}(B)$.*

Proof. Assume B is valid. Then it is provable in $\mathbb{T}_{\mu+}^\omega$. Then there is also a proof in $\mathbb{T}_{\mu+}^n$ since the two systems are the same except that $(\nu.n)$ has fewer premises than $(\nu.\omega)$.

For the other direction we show that from a given $\mathbb{T}_{\mu+}^n$ proof of a formula B , we can construct a $\mathbb{T}_\mu^{\text{pre}}$ proof. We only have to observe that the construction of the pruned proof tree can be carried out even if one starts from a $\mathbb{T}_{\mu+}^n$ proof (instead of a $\mathbb{T}_{\mu+}^\omega$ proof). The only critical case is clause 3 in Definition 6.2. We have to make sure that in an instance of $(\nu.n)$ with conclusion $\Delta, (\nu \mathsf{X})\mathcal{A}$ there is a premise $\Delta, (\nu^{\mathfrak{f}(\Delta, (\nu \mathsf{X})\mathcal{A})} \mathsf{X})\mathcal{A}$. That is we have to show

$$\mathfrak{f}(\Delta, (\nu \mathsf{X})\mathcal{A}) \leq n = \mathfrak{f}(B) \tag{4}$$

Let Γ be a sequent occurring in a $\mathbb{T}_{\mu+}^n$ proof of B . We have $\Gamma^- \subseteq \mathbb{FL}(B)$. Therefore also $\mathbb{FL}(\Gamma^-) \subseteq \mathbb{FL}(B)$ and hence $\mathfrak{f}(\Gamma) \leq \mathfrak{f}(B)$. In particular, this implies that (4) holds. Soundness finally yields then that B is valid. \square

The completeness proof presented in [12] constructs a counter-model to any given non-provable \mathcal{L}_μ 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 7.4 we can replace Clause 5 in the definition of the strong closure of D by

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

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 [12] gives us a finite counter-model. This results in a proof-theoretic proof of the finite model property of the modal μ -calculus. However, since f is exponential and the canonical counter-model construction takes all subsets of $\mathbb{SC}(D)$, the best we get is a double exponential bound for the size of the model (compare with the exponential bound provided by [8]).

8 Conclusion

The starting point of this paper are the two infinitary systems $\mathbb{T}_{\mu+}^\omega$ and $\mathbb{T}_\mu^{\text{pre}}$ for the modal μ -calculus. The first system includes an ω -rule to derive the truth of a greatest fixed point whereas the second system features infinitely long branches that have to satisfy a certain global criterion about the unfolding of greatest fixed points.

Our main technical contribution is a novel translation from $\mathbb{T}_{\mu+}^\omega$ proofs to $\mathbb{T}_\mu^{\text{pre}}$ proofs: given a $\mathbb{T}_{\mu+}^\omega$ proof of a sequent Γ we can construct a $\mathbb{T}_\mu^{\text{pre}}$ proof of Γ . The basic idea is that if an instance of $(\nu.\omega)$ has been applied in the $\mathbb{T}_{\mu+}^\omega$ proof, then we can choose a branch through a ‘big enough’ premise which guarantees that a certain repetition will occur on that branch. This then allows one to construct the corresponding $\mathbb{T}_\mu^{\text{pre}}$ proof. Two applications follow as corollaries. First, the completeness of $\mathbb{T}_{\mu+}^\omega$ implies the completeness of $\mathbb{T}_\mu^{\text{pre}}$. Second, we obtain a novel proof of the finite model property. However, note that the bound on the size of the model is not optimal and that Niwinski and Walukiewicz already showed completeness for a tableau version of $\mathbb{T}_\mu^{\text{pre}}$.

We have examined the proof-theoretic relationship between two infinitary sequent systems for the μ -calculus. Still, there are many important open questions concerning cut-free sequent systems for modal fixed point logics. First of all the question whether there are ‘nice’ such systems. Recently, an infinitary deep sequent system for logic of common knowledge has been presented [6]. That system is ‘nice’ in the sense that weakening and contraction are admissible, all its rules are invertible and there is a syntactic cut-elimination procedure for it. A similar system for the μ -calculus can also be defined. However, that system contains an ω -rule and makes essential use of so-called deep sequents. It is not known whether there is a ‘nice’ finitary sequent system.

It is possible to ‘finitize’ the system $\mathbb{T}_{\mu+}^\omega$ (either by making use of the finite model property, see [12, 15], or in the way shown in Lemma 7.4). The

resulting systems include rules for greatest fixed points that have only finitely many premises. However, such systems are not very natural and the usual proof theoretic techniques cannot be applied to them, see for instance [11].

Another approach to obtain finitary cut-free systems for modal fixed point logics is to reformulate focus games as sequent calculi [5]. This works for CTL and LTL but it is not clear yet whether it will extend to the full μ -calculus. Moreover, again it seems that the usual proof theoretic methods do not work for such calculi. The same also holds for the deductive system for LTL presented in [10]. Also very interesting are the systems for inductive definitions presented in [4] which have some nice properties. However, cut-elimination is only established semantically and not with a syntactic cut-elimination procedure.

Acknowledgements

We would like to thank the anonymous referees for the detailed reviews which were very helpful to improve the quality of the present paper.

References

- [1] John Aldwinckle and Robin Cockett. The proof theory of modal μ logics. In *Proc. Fixed Points in Computer Science*, 2001.
- [2] Julian Bradfield and Colin Stirling. Modal mu-calculi. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, pages 721–756. Elsevier, 2007.
- [3] James Brotherston. *Sequent Calculus Proof Systems for Inductive Definitions*. PhD thesis, University of Edinburgh, 2006.
- [4] James Brotherston and Alex Simpson. Complete sequent calculi for induction and infinite descent. In *Proceedings of LICS-22*, pages 51–60, 2007.
- [5] Kai Brännler and Martin Lange. Cut-free sequent systems for temporal logics. *Journal of Logic and Algebraic Programming*, to appear.
- [6] Kai Brännler and Thomas Studer. Syntactic cut-elimination for common knowledge. In *Methods for Modalities 5*, 2007.
- [7] 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.
- [8] E. Allen Emerson and Charanjit S. Jutla. The complexity of tree automata and logics of programs. In *29th Annual Symposium on Foundations of Computer Science FOCS*, pages 328–337. IEEE, 1988.
 - [9] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computing and System Science*, 18(2):194 – 211, 1979.
 - [10] Joxe Gaintzarain, Montserrat Hermo, Paqui Lucio, Marisa Navarro, and Fernando Orejas. A cut-free and invariant-free sequent calculus for PLTL. In Jacques Duparc and Thomas A. Henzinger, editors, *Computer Science Logic CSL 2007*, pages 481–495. Springer, 2007.
 - [11] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut-free common knowledge. *Journal of Applied Logic*, 5:681–689, 2007.
 - [12] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Canonical completeness of infinitary μ . *Journal of Logic and Algebraic Programming*, to appear.
 - [13] Dexter Kozen. Results on the propositional modal μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
 - [14] Dexter Kozen. A finite model theorem for the propositional μ -calculus. *Studia Logica*, 47(3):233–241, 1988.
 - [15] 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.
 - [16] Damian Niwinski and Igor Walukiewicz. Games for the mu-calculus. *Theoretical Computer Science*, 163(1&2):99–116, 1996.
 - [17] Luigi Santocanale. A calculus of circular proofs and its categorical semantics. In *FoSSaCS '02: Proceedings of the 5th International Conference on Foundations of Software Science and Computation Structures*, pages 357–371. Springer, 2002.
 - [18] 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.

- [19] 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.
- [20] 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