

An infinitary treatment of full μ -calculus

Bahareh Afshari^{1,2}, Gerhard Jäger³, and Graham E. Leigh²

¹ University of Amsterdam, The Netherlands
b.afshari@uva.nl

² University of Gothenburg, Sweden
graham.leigh@gu.se

³ University of Bern, Switzerland
gerhard.jaeger@inf.unibe.ch

Abstract We explore the proof theory of the modal μ -calculus with converse, aka the ‘full μ -calculus’. Building on nested sequent calculi for tense logics and infinitary proof theory of fixed point logics, a cut-free sound and complete proof system for full μ -calculus is proposed. As a corollary of our framework, we also obtain a direct proof of the regular model property for the logic: every satisfiable formula has a tree model with finitely many distinct subtrees. To obtain the results we appeal to the basic theory of well-quasi-orderings in the spirit of Kozen’s proof of the finite model property for μ -calculus without converse.

1 Introduction

Modal logic provides an effective language for expressing properties of state-based systems. When equipped with operators that can test for infinite behaviour like looping and reachability, the logic becomes a powerful tool for specifying correctness of nonterminating reactive processes such as communication protocols and control systems. An elegant example of such a logic is the *modal μ -calculus*, an extension of modal logic which captures the essence of inductive and co-inductive reasoning.

In modal μ -calculus two quantifiers, μ and ν , binding propositional variables, are added to the syntax of modal logic. The formulæ $\mu x\phi$ and $\nu x\phi$ are interpreted over directed graphs as, respectively, the least and greatest fixed points of the monotone function $x \mapsto \phi(x)$. The calculus can thus be thought of as a logic that allows for restricted second-order quantification while still maintaining decidability. Indeed all standard computational problems, such as model-checking and satisfiability, are decidable for this logic (see e.g. [4, 15]).

Despite its importance, many fundamental questions regarding μ -calculus, and in particular its intricate proof theory, remain open. There are two notable proof systems for modal μ -calculus. Kozen [19] proposed extending the axioms of basic modal logic K with the fixed point axioms

$$\phi(\mu x\phi) \rightarrow \mu x\phi(x) \qquad \phi(\psi) \rightarrow \psi \vdash \mu x\phi(x) \rightarrow \psi$$

asserting that $\mu x\phi$ is a pre-fixed point of $\phi(x)$ and that it is the least such.

$$\begin{array}{c}
\mu x\phi, \nu x\bar{\phi} \quad \frac{\Gamma, \phi, \psi}{\Gamma, \phi \vee \psi} \vee \quad \frac{\Gamma, \phi \quad \Gamma, \psi}{\Gamma, \phi \wedge \psi} \wedge \quad \frac{\Gamma, \phi}{\langle \mathbf{a} \rangle \Gamma, [\mathbf{a}] \phi} \text{mod} \\
\frac{\Gamma}{\Gamma, \phi} \text{weak} \quad \frac{\Gamma, \phi(x/\mu x\phi)}{\Gamma, \mu x\phi} \mu \quad \frac{\Gamma, \phi(x/\bar{\Gamma})}{\Gamma, \nu x\phi} \text{ind} \quad \frac{\Gamma, \bar{\phi} \quad \Gamma, \phi}{\Gamma} \text{cut} \\
\hline
\langle \mathbf{a} \rangle \Gamma := \{ \langle \mathbf{a} \rangle \phi \mid \phi \in \Gamma \} \quad \bar{\Gamma} := \bigwedge \{ \bar{\phi} \mid \phi \in \Gamma \} \\
\hline
\end{array}$$

Figure 1. Axioms and rules of Koz.

Completeness for the aconjunctive fragment of the language was established by Kozen [19], but full completeness of this axiomatisation was not proved until Walukiewicz’s seminal work [31]. Walukiewicz’s proof combines an analysis of tableaux, games and automata which, it is generally agreed, is highly complex [3, 9]. A natural sequent representation of Kozen’s axiomatisation, denoted **Koz**, is given in Figure 1. The fixed point rule, μ , and the induction rule, ind , capture the two fixed point axioms above.

The second important axiomatisation for μ -calculus is a cut-free infinitary system due to Jäger, Kretz and Studer [16]. The system, denoted $K_\omega^+(\mu)$ in [16], is **Koz** with the cut and ind rules replaced by the single inference

$$\frac{\Gamma, \nu^0 x\phi \quad \Gamma, \nu^1 x\phi \quad \cdots}{\Gamma, \nu x\phi} \nu_\omega$$

The formula $\nu^n x\phi$ denotes the finite approximation to the greatest fixed point: $\nu^0 x\phi = \top$ and $\nu^{n+1} x\phi = \phi(\nu^n x\phi)$ for each $n < \omega$. The proof of completeness for the system $K_\omega^+(\mu)$ is established by adapting the method of canonical model construction for modal logics to the fixed point extension. To demonstrate soundness of the system, more specifically that of ν_ω -rule, the finite model property of μ -calculus [18, 29] is invoked.

In this paper we are interested in the proof theory of μ -calculus extended by converse modalities. The extension, known as the *two-way μ -calculus* or *full μ -calculus*, assumes each action \mathbf{a} is associated a “converse” action $\bar{\mathbf{a}}$ and that a transition system has an $\bar{\mathbf{a}}$ -edge from vertices u to v iff it has an \mathbf{a} -edge from v to u . Axiomatically, one stipulates $\phi \rightarrow [\mathbf{a}]\langle \bar{\mathbf{a}} \rangle \phi$ for every formula and action.

Checking satisfiability for μ -calculus with converse was proven to be decidable by Vardi in [30] (see also [5]) where he introduces the *two-way automata* characterising this extension and shows the emptiness problem is decidable. In contrast to pure modal μ -calculus, the finite model property fails.

To the best of our knowledge, a sound and complete axiomatisation for full μ -calculus has not been given. This can seem somehow surprising if one were to speculate that the presence of converse can simplify completeness results – such as is the case for the computational tree logic CTL* with past [24, 25]. One can add to **Koz** suitable converse axioms and ask whether the resulting system is complete. Walukiewicz’s completeness proof for modal μ -calculus does not easily lend itself to this question because its machinery, particularly the parts

based on tableaux, fall short of converse. Similarly, for the alternative proof of completeness via cyclic proofs given in [1] it is unclear how modalities operating in both directions can be incorporated.

There is another possibility for obtaining a sound and complete axiomatisation for full μ -calculus, namely an adaptation of the infinitary system $K_\omega^+(\mu)$ of [16], which is undertaken in this paper. There are two obstacles to this approach: accommodating converse in the canonical model construction and recovering any structural properties that remain in the absence of the finite model property that are needed to show soundness of an infinitary ν -rule.

We overcome the first issue by stepping into the framework of nested sequents in the style of Kashima's work for tense logics [17]. The failure of the finite model property shows the ν_ω -rule is unsound in the presence of converse. We establish soundness for the infinitary ν -rule with a premise for each approximant below ω^ω and prove that the ensuring nested sequent calculus is complete for the full μ -calculus. Moreover, we observe that this bound is optimal over trees: the greatest fixed point cannot be identified with its transfinite approximant for any ordinal below ω^ω .

Related work The history of modal logic with converse goes back to Prior and his introduction of tense logics.⁴ Temporal logics with past have been widely studied. For example, completeness of converse PDL was first shown in [23], and a sound and complete axiomatisation of PCTL* (computation tree logic with past) is given in [24]. More recent work relevant to this paper include the treatment of tense logics in [14] and the completeness proof for the flat fragment of μ -calculus [8]. The literature on nested sequents is also rich: they have been used to establish algorithmic properties on a wide range of logics (e.g. [2, 10, 11, 13, 26]). The explicit use of ordinal approximations in the language of μ -calculus is a feature that has been used by other authors studying fixed-point logics. Of particular note is the work of [6] in which they are utilised for the correspondence between circular proofs and induction.

2 Full μ -calculus

Fix finite sets \mathbf{Act} and \mathbf{Var} of **actions** and **variables**, respectively. The μ -calculus **formulæ** are given by the following grammar, where \mathbf{a} ranges over actions and x over variables.

$$\phi := x \mid \phi \wedge \phi \mid \phi \vee \phi \mid [\mathbf{a}]\phi \mid \langle \mathbf{a} \rangle \phi \mid \mu x \phi \mid \nu x \phi$$

The two propositional quantifiers μ and ν are called the **least**, and **greatest, fixed point quantifier** respectively. The syntax above omits both negation and proposition constants. Constants for 'true' and 'false' can be defined via the quantifiers: $\top = \nu x. x$ and $\perp = \mu x. x$; other (unspecified) propositional constants can be represented via additional actions. Negation is representable

⁴ For a comprehensive account see, e.g., [12].

via De Morgan duality in the usual way; see e.g., [7, chap. 8]. We write $\phi(x/\psi)$ for the result of substituting the formula ψ for every free occurrence of x in ϕ subject to the proviso that no free variable of ψ becomes bound.

In the full μ -calculus every action $\mathbf{a} \in \mathbf{Act}$ has an associated converse which we denote as $\bar{\mathbf{a}}$. Thus we assume the presence of an involution, $\bar{\cdot}: \mathbf{Act} \rightarrow \mathbf{Act}$, on the set of actions: for every $\mathbf{a} \in \mathbf{Act}$ we have $\mathbf{a} \neq \bar{\mathbf{a}}$, and $\bar{\bar{\mathbf{a}}} = \mathbf{a}$. If a formula ϕ contains at most one of \mathbf{a} or $\bar{\mathbf{a}}$ for each $\mathbf{a} \in \mathbf{Act}$ we may call ϕ **pure**.

We recall the standard Kripke semantics for the modal μ -calculus.

Definition 1. A **labelled frame** is a pair (S, E) where S is a non-empty set of **vertices** and $E \subseteq \mathbf{Act} \times S \times S$ is a set of (**labelled**) **edges**. The **symmetric closure** of a frame (S, E) is the frame (S, E') where $E' = E \cup \{(\bar{\mathbf{a}}, v, u) \mid (\mathbf{a}, u, v) \in E\}$. A frame is **symmetric** if it is identical to its symmetric closure.

If the labelled frame (S, E) is clear from the context, we write $u \xrightarrow{\mathbf{a}} v$ (or simply $u \longrightarrow v$) if $(\mathbf{a}, u, v) \in E$. Labelled frames provide a semantics for μ -calculus formulæ via the possible worlds interpretation of the modal connectives:

Definition 2. Let ϕ be a formula, possibly with free variables, $\mathcal{S} = (S, E)$ a labelled frame, and $V \subseteq \mathbf{Var} \times S$. The **denotation** of ϕ in \mathcal{S} relative to V , in symbols $\|\phi\|_V^{\mathcal{S}}$, is the subset of S defined with the standard semantics for boolean and modal operators (e.g. [7]) and the following equations for fixed point quantifiers, where

$$\begin{aligned} V[x \mapsto T] &= \{(y, u) \in V \mid y \neq x\} \cup (\{x\} \times T). \\ \|\mu x \phi\|_V^{\mathcal{S}} &= \bigcap \{T \subseteq S \mid \|\phi\|_{V[x \mapsto T]}^{\mathcal{S}} \subseteq T\} \\ \|\nu x \phi\|_V^{\mathcal{S}} &= \bigcup \{T \subseteq S \mid T \subseteq \|\phi\|_{V[x \mapsto T]}^{\mathcal{S}}\} \end{aligned}$$

We write $(\mathcal{S}, u) \models \phi$ to express $u \in \|\phi\|_0^{\mathcal{S}}$. A formula ϕ is **satisfied** by \mathcal{S} if $(\mathcal{S}, u) \models \phi$ for some vertex u and is **true** in \mathcal{S} if $(\mathcal{S}, u) \models \phi$ for every vertex u . Given a class of frames \mathcal{C} , a formula ϕ has a **model** in \mathcal{C} if ϕ is satisfied in some frame $\mathcal{S} \in \mathcal{C}$, and is **valid** over \mathcal{C} if ϕ is true in every frame in \mathcal{C} . If mention of \mathcal{C} is omitted we have in mind the class of all countable frames.

Definition 3. Let ϱ be a well-formed formula. The **Fischer–Ladner closure** of ϱ , denoted $\mathbb{FL}(\varrho)$, is the smallest set of formulæ containing ϱ satisfying the following conditions.

- If $\phi \circ \psi \in \mathbb{FL}(\varrho)$ for $\circ \in \{\wedge, \vee\}$ then $\{\phi, \psi\} \subseteq \mathbb{FL}(\varrho)$.
- If $\Delta \phi \in \mathbb{FL}(\varrho)$ for $\Delta \in \{[\mathbf{a}], \langle \mathbf{a} \rangle \mid \mathbf{a} \in \mathbf{Act}\}$ then $\phi \in \mathbb{FL}(\varrho)$.
- If $\sigma x \phi \in \mathbb{FL}(\varrho)$ for $\sigma \in \{\mu, \nu\}$ then $\phi(x/\sigma x \phi) \in \mathbb{FL}(\varrho)$.

In what follows we utilise an extended language where the greatest fixed point quantifier ν spawns an infinite hierarchy of ‘approximation’ quantifiers indexed by ordinals. Fix an ordinal κ . The κ -**formulæ** are generated as follows

$$\phi := x \mid \phi \wedge \phi \mid \phi \vee \phi \mid [\mathbf{a}]\phi \mid \langle \mathbf{a} \rangle \phi \mid \mu x \phi \mid \nu x \phi \mid \nu^\alpha x \phi \quad (\alpha < \kappa)$$

The intended reading of the quantifier ν^α is of α -times unfolding the matrix, taking conjunctions at limits with $\nu^0 x \phi$ equivalent to \top . When κ is fixed it is convenient to identify the unannotated quantifier ν with ν^κ . From a κ -formula ϕ we derive a μ -calculus formula ϕ^- , called the **template** of ϕ , by removing the approximation of every quantifier.

The class of ω -formulae corresponds to the language \mathcal{L}_μ^+ of [16]. The following definition expands the generalisation of the Fischer–Ladner closure in [16] to a form appropriate for κ -formulae.

Definition 4. *The **strong closure** of a κ -formula ϱ , denoted $\text{SC}_\kappa(\varrho)$, is the smallest set containing ϱ closed under the formation rules for the Fischer–Ladner closure and the following clause.*

- If $\nu^\alpha x \phi \in \text{SC}_\kappa(\varrho)$ ($\alpha \leq \kappa$) then $\{\phi(x/\nu^\beta x \phi) \mid \beta < \alpha\} \subseteq \text{SC}_\kappa(\varrho)$.

Semantics for κ -formulae is obtained by extending the definition for μ -calculus formulae to accommodate the approximating quantifiers subject to the equation $(\mathcal{S}, u) \models \nu^\alpha x \phi$ iff $(\mathcal{S}, u) \models \phi(x/\nu^\beta x \phi)$ for every $\beta < \alpha$. Standard arguments on the fixed point semantics show there exists κ s.t. $\nu x \phi \leftrightarrow \nu^\kappa x \phi$ is true in \mathcal{S} .

An important concept in μ -calculus is the relation of subsumption between variables occurring in a given formula, a syntactic constraint that mirrors the priority of quantifiers implicit in the semantics. In the present article, we take a pragmatic approach to subsumption, assuming a fixed strict partial order on Var , called the **subsumption order**, and constrain considerations to formulae whose variables respect this relation, in the sense that if y occurs free in a sub-formula $\sigma x \phi$ then y subsumes x . We call such formulae **well-formed**. The subsumption order must be irreflexive, asymmetric, transitive and for every $x \in \text{Var}$ the set of variables subsuming x should be linearly ordered by subsumption.

If ϕ is well-formed then for a substitution $\phi(x/\psi)$ to be ‘correct’ it suffices that x does not subsume any free variable of ψ . Thus, every formula occurring in the Fischer–Ladner closure of a well-formed formula is well-formed. It is common to assume that each quantified formula uniquely determines a variable symbol that is bound; we call such a formula **well-named**. Note, however, that, unlike the notion of well-formed, the Fischer–Ladner closure conditions do not preserve well-namedness.

Definition 5. *L_μ denotes the set of closed formulae that appear in the Fischer–Ladner closure of some well-named formula.*

We likewise need to isolate a class of κ -formulae to assist the presentation of our results. The class of κ -formulae with templates in L_μ is natural, but there is a strict sub-class of these formulae that we should restrict attention to. This turns out to be the collection of κ -formulae that arise when evaluating the denotation of an $\phi \in L_\mu$ subject to the identification of ν with ν^κ . These formulae, which we call **well-annotated**, satisfy the following three conditions:

1. Their template is well-formed.
2. If $\nu^\alpha x$ and $\nu^\beta x$ are two quantifiers binding the same variable then $\alpha = \beta$.

3. The set of variable symbols bound by a quantifier ν^α with $\alpha < \kappa$ is linearly ordered by the subsumption relation.

The set of well-annotated κ -formulae is denoted L_μ^κ . It is a simple exercise to check that every formula in the strong closure of a well-annotated formula is well-annotated.⁵

2.1 Nested sequent calculi

Nested sequents were utilised by Kashima to establish canonical completeness for tense logics [17]. In the following we adapt Kashima's approach to $L_\mu^{\omega_1}$. For the present section κ is an arbitrary ordinal $\leq \omega_1$.

Definition 6. A *sequent* is a finite set of closed L_μ^κ formulae. The *nested sequents* (ns) are defined inductively:

1. every plain sequent is a nested sequent,
2. if Γ is a nested sequent and \mathbf{a} is an action then $\mathbf{a}\{\Gamma\}$ is a ns,
3. if Γ, Δ are ns then so is $\Gamma \cup \Delta$.

As is usual, we use comma to abbreviate the union of two (nested) sequents and identify singleton sequents with their unique element. Hence, every nested sequent can be presented in the form

$$\Gamma = \phi_1, \dots, \phi_m, \mathbf{a}_1\{\Delta_1\}, \dots, \mathbf{a}_n\{\Delta_n\} \quad (1)$$

where $\phi_1, \dots, \phi_m \in L_\mu^\kappa$, $\Delta_1, \dots, \Delta_n$ are nested sequents and $\mathbf{a}_1, \dots, \mathbf{a}_n \in \text{Act}$. The intended interpretation of the nested sequent Γ in (1) is the formula

$$\iota(\Gamma) = \bigvee_{i=1}^m \phi_i \vee \bigvee_{i=1}^n [\mathbf{a}_i] \iota(\Delta_i).$$

A **sequent with context** (simply **context**) is a nested sequent built from an additional unit \square , called the context, which must have exactly one occurrence within the nested sequent. If Γ is a sequent with context and Δ is a nested sequent $\Gamma[\Delta]$ is the nested sequent given by substituting Δ for \square in Γ .

Definition 7. Fix $\kappa \leq \omega_1$. $\mathsf{K}_{\mu+}^\kappa$ is the calculus deriving nested sequents given by the inferences in Figure 2. K_μ^κ denotes the subsystem without the inference $\text{con}_\mathbf{a}$.

A special case of the $\nu.\alpha$ inference is when $\alpha = 0$, whereby the sequent $\Gamma[\nu^0 x \phi]$ is derivable without premises. Hence, $\mathsf{K}_\mu^\kappa \vdash \Gamma[\top]$ for any sequent context $\Gamma[\square]$. Clearly, a smaller value of κ makes introducing greatest fixed points easier. The following properties can be established by induction on the length of derivations.

- Lemma 1.**
1. If $\mathsf{K}_{\mu+}^\alpha$ is complete so is $\mathsf{K}_{\mu+}^\beta$ for every $\beta \leq \alpha$; similarly for K_μ^κ .
 2. For all $\phi \in L_\mu$ and contexts $\Gamma[\square]$, $\mathsf{K}_\mu^\kappa \vdash \Gamma[\phi, \bar{\phi}]$ where $\bar{\phi}$ denotes the De Morgan dual of ϕ .
 3. If $\mathsf{K}_{\mu+}^\kappa \vdash \Gamma[\nu^\alpha x \phi]$ then $\mathsf{K}_{\mu+}^\kappa \vdash \Gamma[\nu^\beta x \phi]$ for every $\beta < \alpha$; similarly for K_μ^κ .

⁵ See Appendix A for precise definitions of the concepts of this section.

$$\begin{array}{c}
 \frac{\Gamma[\phi, \psi]}{\Gamma[\phi \vee \psi]} \vee \qquad \frac{\Gamma[\phi] \quad \Gamma[\psi]}{\Gamma[\phi \wedge \psi]} \wedge \qquad \frac{\Gamma[\phi(x/\mu x\phi)]}{\Gamma[\mu x\phi]} \mu \\
 \frac{\Gamma[\mathbf{a}\{\phi\}]}{\Gamma[[\mathbf{a}]\phi]} [\mathbf{a}] \qquad \frac{\Gamma[\mathbf{a}\{\Delta, \phi\}]}{\Gamma[\mathbf{a}\{\Delta\}, \langle \mathbf{a} \rangle \phi]} \langle \mathbf{a} \rangle \qquad \frac{\Gamma[\mathbf{a}\{\Delta\}, \phi]}{\Gamma[\mathbf{a}\{\Delta, \langle \bar{\mathbf{a}} \rangle \phi]} \text{con}_{\mathbf{a}} \\
 \frac{\Gamma[\phi(x/\nu^\beta x\phi)]}{\Gamma[\nu^\alpha x\phi]} \text{ for all } \beta < \alpha \leq \kappa \quad \nu.\alpha \qquad \frac{\Gamma[\phi(x/\nu^\alpha x\phi)]}{\Gamma[\nu x\phi]} \text{ for all } \alpha < \kappa \quad \nu_\kappa
 \end{array}$$

Figure 2. System $K_{\mu^+}^\kappa$; K_μ^κ is $K_{\mu^+}^\kappa$ without $\text{con}_{\mathbf{a}}$.

3 Completeness: building canonical models

Definition 8. A κ -*system* (in ϱ) is a tuple (S, E, λ) where (S, E) is a frame and $\lambda: S \rightarrow \text{Pow}(\text{SC}_\kappa(\varrho))$ assigns to each vertex of S a set of κ -formulæ from the strong closure of ϱ . A system (T, E, λ) **expands** a system (S, F, ρ) if $S \subseteq T$, $F \subseteq E$, and $\rho(u) \subseteq \lambda(u)$ for every $u \in S$.

Explicit mention of κ and ϱ will be dropped if they can be inferred from context and, when there is no cause for confusion, vertices of a system will be identified with their labels: $\phi \in u$ in place of $\phi \in \lambda(u)$. Recall that edges of a labelled frame (and so of a system) are labelled by actions and that symmetry is not assumed. A nested sequent $\Gamma = \Delta_0, \mathbf{a}_1\{\Delta_1\}, \dots, \mathbf{a}_l\{\Delta_l\}$ ($\Delta_0 \subseteq L_\mu^\kappa$) has a natural representation as finite κ -system, $\text{tree}(\Gamma)$, comprising a root with label Δ_0 and, for each $0 < i \leq l$, an \mathbf{a}_i -child with immediate subtree $\text{tree}(\Delta_i)$.

For the proof of completeness, starting from an assumption that a sequent Γ is undervivable we will construct a system expanding Γ by saturating the sequent through the $K_{\mu^+}^\kappa$ rules applied from conclusion to premise. Deconstructing a modality corresponds to creating, or saturating, other vertices in the system. This method combines saturation arguments for the (pure) modal μ -calculus [16] and the tableau-style constructions for tense logic [17]. If we obtain two different annotations of the same formula, say $\nu^\alpha x\phi$ and $\nu^\beta x\phi$, then clearly, from the perspective of non-derivability, the smaller approximation suffices. Thus, to maintain some control on the κ -formulæ enumerated via the process, we desire an ordering on L_μ^κ formulæ based on the ordinal approximations.

Recall a **quasi-order** is a reflexive, transitive relation. Let \sqsubseteq be the quasi-order on L_μ^κ determined by $\phi \sqsubseteq \psi$ iff $\phi^- = \psi^-$ and for every maximal chain x_1, \dots, x_n of the ν -quantified variables in ϕ such that x_i subsumes x_{i+1} , we have $(\alpha_1, \dots, \alpha_n) \leq (\beta_1, \dots, \beta_n)$, where α_i (β_i) is the ordinal assigned to x_i in ϕ (resp. ψ) and \leq is the lexicographic ordering on sequences of ordinals.⁶

Definition 9. A κ -system \mathcal{S} is **saturated** if the following hold for every $u \in \mathcal{S}$, $\phi, \psi \in L_\mu^\kappa$, $\mathbf{a} \in \text{Act}$ and $\alpha \leq \kappa$.

- a) $\phi \wedge \psi \in u$ implies $\phi \in u$ or $\psi \in u$,
- b) $\phi \vee \psi \in u$ implies $\phi \in u$ and $\psi \in u$,

⁶ Cf. Appendix A.

- c) $\mu x\phi \in u$ implies $\phi(x/\mu x\phi) \in u$,
- d) $\nu^\alpha x\phi \in u$ implies $\phi(x/\nu^\beta x\phi) \in u$ for some $\beta < \alpha$,
- e) $[\mathbf{a}]\phi \in u$ implies for some $u \xrightarrow{\mathbf{a}} v$ and $\psi \in v$ we have $\psi \sqsubseteq \phi$,
- f) $\langle \mathbf{a} \rangle \phi \in u$ and either $u \xrightarrow{\mathbf{a}} v$ or $v \xrightarrow{\bar{\mathbf{a}}} u$ implies $\psi \in v$ for some $\psi \sqsubseteq \phi$.

Our notion of saturation combines a number of features from other work. It is closely related to Kozen's well-annotations in [18] expanded to cover converse modalities in the style of Kashima [17]. Note, however, that our quasi-order differs from Kozen's. Dropping the two modal clauses e and f yields the definition of saturation in [16], for $\kappa = \omega$.

Lemma 2. *Let \mathcal{S} be a saturated κ -system.*

1. *The symmetric closure of \mathcal{S} is saturated.*
2. *For every $u \in \mathcal{S}$ and $\phi \in u$, $(\mathcal{S}, u) \not\models \phi$ (Truth Lemma).*

Proof. The first claim is immediate given the formulation of condition f. For 2, we refer the reader to [18, Lemma 4.2], noting that, like the quasi-order utilised in [18], denotation is monotone in \sqsubseteq : if $\phi \sqsubseteq \psi$ then $(\mathcal{S}, u) \models \psi$ implies $(\mathcal{S}, u) \models \phi$ for any u . A more detailed proof of the result, based on the assignment of a rank to each formula of L_μ^κ , is given in [16, Lemma 33].⁷

We establish weak completeness of the calculi $\mathsf{K}_{\mu^+}^\kappa$ and K_μ^κ , namely that every underivable sequent has a counter-model. In view of Lemma 2, it suffices to show that every underivable sequent expands to a saturated system. In contrast to the constructions in [16, 17] (for pure μ -calculus and tense logic respectively), we cannot expect the result to be a finite system (i.e. a nested sequent); in general, an infinite tree will result.

Lemma 3 (Saturation Lemma). *Suppose $\mathsf{K}_{\mu^+}^\kappa \not\vdash \Gamma$. There exists a saturated κ -system \mathcal{S} expanding Γ such that every formula occurring in the label of a vertex of \mathcal{S} is an element of $\mathsf{SC}_\kappa(\varrho)$ for some formula ϱ .*

Proof. We require an auxiliary notion of saturation. Let us call a κ -system 0-saturated if the saturation conditions hold with the possible exception of the clauses for modalities, e and f. Every underivable nested sequent can be expanded to a 0-saturated nested sequent that remains underivable. The proof of this fact follows the argument of Lemma 24 in [16].

Suppose $\mathsf{K}_{\mu^+}^\kappa \not\vdash \Gamma$. We define a sequence of nested sequents $\Gamma = \Gamma_0, \Gamma_1, \Gamma_2, \dots$ such that Γ_{i+1} expands Γ_i and Γ_i is underivable. Given Γ_i , obtain Γ_{i+1} by

1. expanding Γ_i to a 0-saturated nested sequent Γ'_i ;
2. expanding Γ'_i to a ns Γ_{i+1} by correcting any failure of conditions e or f:
 - a. For any $v \in \Gamma'_i$ and formula $[\mathbf{a}]\phi \in u$ for which there is no \mathbf{a} -child of u in Γ'_i containing ϕ , create a \mathbf{a} -child with label $\{\phi\}$;

⁷ As already remarked, [16] deals only with the case $\kappa = \omega$. However, their notion of rank and the proof of the Truth Lemma readily generalises to arbitrary κ .

- b. If v is a \mathbf{a} -child of u in Γ'_i , expand the label of u to include $\{\phi \mid \langle \bar{\mathbf{a}} \rangle \phi \in v\}$ and the label of v to include $\{\phi \mid \langle \mathbf{a} \rangle \phi \in u\}$.

The process of 0-saturation preserves underivability. Moreover, Γ'_i can be derived from Γ_{i+1} by a sequence of $\text{con}_{\mathbf{a}}$, $[\mathbf{a}]$ and $\langle \mathbf{a} \rangle$ inferences, hence Γ_{i+1} is underivable. Let \mathcal{T} be the limit of trees $\text{tree}(\Gamma_i)$ for $i < \omega$. By construction, \mathcal{T} is a saturated κ -system fulfilling the requirements of the lemma. \square

As a consequence of the Saturation and Truth lemmas we deduce completeness for full μ -calculus. An analogous argument establishes completeness for the pure fragment.

Theorem 1. $K_{\mu^+}^{\kappa}$ is complete over symmetric frames. K_{μ}^{κ} is complete for arbitrary frames.

Proof. Suppose Γ is underivable in $K_{\mu^+}^{\kappa}$ and let \mathcal{S} be the symmetric closure of the κ -system expanding Γ provided by Lemma 3. As a consequence of Lemma 2, $(\mathcal{S}, r) \not\models \iota(\Gamma)$ where r is the root of \mathcal{S} . Hence, Γ is not valid. An analogous argument establishes completeness for K_{μ}^{κ} .

4 Soundness: refining canonical models

We now turn to soundness theorems for the systems $K_{\mu^+}^{\kappa}$ and K_{μ}^{κ} for certain κ . It can be easily confirmed that for either system the only inference we need be concerned with is the introduction rule for the greatest fixed point, ν_{κ} .

Some cases of soundness can be inferred from known properties of the μ -calculus. For instance, the pure μ -calculus (without converse modalities) has the finite model property: every satisfiable formula has a finite model [18, 29]. On the class of finite models the greatest fixed point coincides with the ω -th approximation, ν^{ω} . Thus soundness of K_{μ}^{ω} obtains.

Theorem 2. K_{μ}^{ω} is sound and complete for arbitrary frames.

The above theorem can also be deduced without directly appealing to the finite model property, by manipulating saturated systems. This argument was already made by Kozen [18] and will be extended below.

The full μ -calculus lacks the finite model property (there are satisfiable formulae with no finite models) but every satisfiable formula has a model which is (the symmetric closure of) a finitely branching tree [30]. As a consequence we deduce $K_{\mu^+}^{\kappa}$ is unsound for $\kappa \leq \omega$ but sound for $\kappa \geq \omega_1$.

Theorem 3. $K_{\mu^+}^{\omega_1}$ is sound and complete for arbitrary (symmetric) frames.

In the sequel we prove a strengthening of Theorem 3: the calculus $K_{\mu^+}^{\omega}$ is sound and complete for symmetric frames; and observe that, over trees, $K_{\mu^+}^{\kappa}$ is unsound for every $\kappa < \omega^{\omega}$. Our argument relies on a particular property of the quasi-order \sqsubseteq we introduced earlier, which we now state. Given a set $X \subseteq \text{SC}_{\kappa}(\varrho)$ let $\text{Ker } X = \{\phi \in X \mid \forall \psi \in X \psi \not\sqsubseteq \phi\}$ be the set of \sqsubseteq -minimal elements of X . Recall, a quasi-order \leq on a set Q is a **well-quasi-order** (wqo for short) if for every function $f: \omega \rightarrow Q$ there exists $i < j$ such that $f(i) \leq f(j)$.

Lemma 4. $(\mathbb{S}\mathbb{C}_\kappa(\varrho), \sqsubseteq)$ is a wqo. Moreover, there exists k such that for every $X \subseteq \mathbb{S}\mathbb{C}_\kappa(\varrho)$, $|\text{Ker } X| < k$.

That \sqsubseteq is a well-quasi-order follows from the observation that the ordering can be expressed as a sum of products of well-orders. Being a wqo we immediately deduce that $\text{Ker } X$ is finite for every $X \subseteq \mathbb{S}\mathbb{C}_\kappa(\varrho)$. The stronger result stated follows from the constraints we imposed in the definition of L_μ^κ , namely condition 3 (on page 6).⁸ Specifically, it is this property that marks the essential difference between \sqsubseteq and the wqo \preceq in [18].

We require a lifting of \sqsubseteq to sets of κ -formulae. A natural candidate is the Smyth powerdomain introduced in [27] and given by $X \sqsubseteq Y$ iff for every $\psi \in Y$ there exists $\phi \in X$ such that $\phi \sqsubseteq \psi$. In general, this lifting does not preserve well-quasi-orders [20] but, rather, the stronger notion of better-quasi-order due to Nash-Williams [21, 22]; $(\mathbb{S}\mathbb{C}_\kappa(\varrho), \sqsubseteq)$ is readily seen to be a better-quasi-order.

For our strengthening of Theorem 3, however, we depend on a refinement of the Smyth powerdomain whereby X is bounded by Y if Y can be realised as the image of X under an endomorphism on $(\kappa, <)$. This choice is motivated by the observation that saturation is preserved under any change of annotating ordinals by a strictly monotone function on κ . The main technical result is to establish that this notion of boundedness is a well-quasi-order on $\text{Pow}(\mathbb{S}\mathbb{C}_\kappa(\varrho))$ for every $\varrho \in L_\mu$. We begin making the above definitions precise.

Let $\mathbb{I}(\kappa)$ be the set of strictly monotone functions on ordinals $\leq \kappa$. Note that such functions are increasing, so $\alpha \leq f(\alpha) \leq \kappa$ for every $\alpha \leq \kappa$. Each $f \in \mathbb{I}(\kappa)$ induces an operation on L_μ^κ mapping ϕ to the result of replacing each annotated quantifier ν^α by $\nu^{f(\alpha)}$, which we denote as ϕ^f . Similarly, for $X \subseteq L_\mu^\kappa$, define $X^f = \{\phi^f \mid \phi \in X\}$ and for a system $\mathcal{T} = (S, E, \lambda)$ we let \mathcal{T}^f be the system (S, E, λ^f) where $\lambda^f: w \mapsto \lambda(w)^f$. The following is straightforward to verify.

Lemma 5. Let $f \in \mathbb{I}(\kappa)$. If $X \sqsubseteq Y$ then $X^f \sqsubseteq Y^f$. Hence, if \mathcal{T} is a saturated κ -system, so is \mathcal{T}^f .

We are now in a position to define the quasi-order on $\text{Pow}(\mathbb{S}\mathbb{C}_\kappa(\varrho))$:

$$X \sqsubseteq^* Y := \exists f \in \mathbb{I}(\kappa) \text{ s.t. } \text{Ker } Y = \text{Ker}(X^f)$$

Since $(\text{Ker } X)^f = \text{Ker } X^f$ for every $X \subseteq \mathbb{S}\mathbb{C}_\kappa(\varrho)$ and $f \in \mathbb{I}(\kappa)$, like the Smyth powerdomain, \sqsubseteq^* is determined by its restriction to kernels: $X \sqsubseteq^* \text{Ker } X \sqsubseteq^* X$. The fact that kernels are bounded (Lemma 4) is crucial for the following result.

We call κ **principal** if $\kappa = \omega^\alpha$ for some α .

Theorem 4. If κ is principal then $(\text{Pow}(\mathbb{S}\mathbb{C}_\kappa(\varrho)), \sqsubseteq^*)$ is a wqo.

Proof. Let $\text{Var}_\nu(\phi)$ be the set of ν -quantified variables in ϕ . To each $\phi \in L_\mu^\kappa$ we may associate a function $o_\phi: \text{Var}_\nu(\phi) \rightarrow \kappa + 1$ such that ϕ can be obtained from its template by replacing each quantifier νx in ϕ^- by $\nu^{o_\phi(x)}x$. We consider finite sequences in $\mathbb{F}\mathbb{L}(\varrho) \times \text{Var} \times (\kappa + 1)$, ordered pointwise by $(\phi_i, x_i, \alpha_i)_{i < m} \leq_{\text{pw}}$

⁸ See Appendix B for a proof of this fact.

$(\psi_i, y_i, \beta_i)_{i < n}$ iff $m = n$ and for all $i < m$, $\phi_i = \psi_i$, $x_i = y_i$ and $\alpha_i \leq \beta_i$. When restricted to a set of sequences of bounded length, \leq_{pw} is a wqo. For $X \subseteq \text{SC}_\kappa(\varrho)$, let $X^* = (\phi_i, x_i, \delta_i)_{i < k}$ be a sequence in $\mathbb{FL}(\varrho) \times \text{Var} \times (\kappa + 1)$ such that

$$\text{Ker } X = \{\phi \in \text{SC}_\kappa(\varrho) \mid \forall x \in \text{Var}_\nu(\phi) \exists i < k (\phi_i = \phi^- \wedge x_i = x \wedge o_\phi(x) = \sum_{j \leq i} \delta_j)\}.$$

Without loss of generality, we assume a total ordering $<$ of $\mathbb{FL}(\varrho) \times \text{Var}$ and that $\delta_{i+1} = 0$ implies $(\phi_i, x_i) < (\phi_{i+1}, x_{i+1})$. By Lemma 4, k can be chosen independent of X . Hence it remains only to observe that for principal κ , $X \sqsubseteq^* Y$ iff $X^* \leq_{\text{pw}} Y^*$. \square

Given systems \mathcal{T} and \mathcal{T}' , write $\mathcal{T} \sqsubseteq^* \mathcal{T}'$ if \mathcal{T}' is isomorphic to \mathcal{T}^f for some $f \in \mathbb{I}(\kappa)$. If \mathcal{T} is a tree, \mathcal{T}_u denotes the sub-tree rooted at $u \in \mathcal{T}$. Suppose \mathcal{T} is a system over a finite tree. We say \mathcal{T} is **quasi-saturated** if:

1. \mathcal{T} validates the saturation conditions for all vertices with the exception of a finite set L of leaves;
2. every $l \in L$ may fail the saturation requirements only in condition f;
3. for every $l \in L$ there exists a non-leaf vertex u in \mathcal{T} such that $u \sqsubseteq^* l$.

Theorem 5. *Let Γ be a nested sequent. TFAE*

1. *There exists a saturated expansion of Γ .*
2. *There exists a finite quasi-saturated expansion of Γ .*
3. *There exists a saturated expansion of Γ , \mathcal{T} , which is a tree, and a finite set $U \subseteq \mathcal{T}$ such that for every $v \in \mathcal{T}$ there exists $u \in U$ satisfying $\mathcal{T}_u \sqsubseteq^* \mathcal{T}_v$.*
4. *There exists a saturated expansion of Γ with a regular underlying frame.*

Proof. The implications $3 \Rightarrow 4$ and $4 \Rightarrow 1$ follow from the definitions. Moreover, Theorem 4 yields $1 \Rightarrow 2$. We show $2 \Rightarrow 3$. Suppose $\mathcal{S} = (S, E, \lambda)$ is quasi-saturated and let $U = \mathcal{S} \setminus L$ be the vertices of \mathcal{S} that fulfil all the saturation conditions. Fix a vertex $l \in L$. By assumption there exists $u \in U$ with $u \sqsubseteq^* l$. Let $\mathcal{S}_u = (S_u, E_u, \lambda|_{S_u})$ be the sub-tree of \mathcal{S} rooted at u , and $f \in \mathbb{I}(\kappa)$ be such that $u^f = l$. Consider the system $\mathcal{S}' = (\mathcal{S}_u)^f = (S_u, E_u, \lambda')$. In particular, $\lambda'(u) = \lambda(l)$. Define \mathcal{T} to be the system comprising the disjoint union of \mathcal{S} and \mathcal{S}' where the leaf l in \mathcal{S} is identified with the root u of \mathcal{S}' . We claim \mathcal{T} is quasi-saturated. Let l' be a leaf in $\mathcal{T} \setminus \mathcal{S}$ which fails the saturation conditions and let $u' \in U$ be such that $\lambda_{\mathcal{S}}(u') \sqsubseteq^* \lambda_{\mathcal{S}}(l')$. By construction $\lambda_{\mathcal{T}}(u') = \lambda_{\mathcal{S}}(u')$ and $\lambda_{\mathcal{T}}(l') = \lambda_{\mathcal{S}}(l')^f$, so $\lambda_{\mathcal{T}}(u') \sqsubseteq^* \lambda_{\mathcal{T}}(l')$ by transitivity. Repeating the method of unravelling the unsaturated leaves and considering the limit system yields a saturated system with the desired properties. \square

The following, due to Vardi [30], is an immediate consequence of Theorem 5.

Corollary 1. *The full μ -calculus has the regular model property.*

We claim the above result enables us to lower the bound on $K_{\mu^+}^{\omega_1}$. The idea is to find a refinement of Theorem 5 that controls the approximations appearing in a saturated system. This is the role of the next proposition.

Proposition 1. *Let \mathcal{T} be a κ -system satisfying condition 3 in Theorem 5. Suppose for every $u \in U$ and $\phi \sqsubseteq \psi \in u$, if $(\mathcal{T}, u) \not\models \phi$ then $\phi \in u$. Then $\text{Ker } u \subseteq L_\mu^{\omega^\omega}$ for every $u \in U$.*

Proof. Suppose $\mathcal{T} = (S, E, \lambda)$ is as stated. We may assume U is closed downwards in the accessibility relation on \mathcal{T} . Let \mathcal{T}_0 be the finite sub-system restricted to vertices in U and their immediate successors. By assumption, \mathcal{T}_0 is quasi-saturated. Let $L = \{l_0, \dots, l_{n-1}\}$ be the vertices of \mathcal{T}_0 not in U . These are leaves and for each $i < n$, let $u_i \in U$ be such that $\lambda(u_i) \sqsubseteq^* \lambda(l_i)$. Consider the κ -system $\mathcal{S}_0 = (S_0, E_0, \lambda_0)$ where $S_0 = U \cup L$, $E_0 = E|_{S_0}$ and $\lambda_0 = (\text{Ker} \circ \lambda)|_{S_0}$ with $|$ denoting restricting the domain of the function/relation.

We have that $\lambda(u) \sqsubseteq^* \lambda(v)$ implies $\lambda_0(u) \sqsubseteq^* \lambda_0(v)$, so \mathcal{S}_0 is quasi-saturated. Let O be the set of ordinals occurring in the sets $\lambda_0(u)$ for $u \in S_0$, which is finite, and $(\alpha_i)_{i < |O|}$ enumerate the elements in O in increasing order. Define $f: O \rightarrow \omega^\omega$ by $f(\alpha_i) = \min\{\alpha_i, \omega^i\}$. We claim \mathcal{S}_0^f is quasi-saturated. Since f is strictly monotone it suffices to check, for each $i < n$, that

$$\lambda_0(u_i)^f \sqsubseteq^* \lambda_0(l_i)^f. \quad (2)$$

Fixing $i < n$, let $a_0 < \dots < a_k$ be such that $\alpha_{a_0}, \dots, \alpha_{a_k}$ enumerates the ordinals in $\lambda_0(u_i)$ and let $b_0 < \dots < b_l < |O|$ be the analogous sequence for $\lambda_0(l_i)$. Given $\lambda_0(u_i) \sqsubseteq^* \lambda_0(l_i)$ we must have $k = l$ and $a_j \leq b_j$ for each $j \leq k$. By induction on $j \leq k$ we may define $h \in \mathbb{I}(\omega^\omega)$ such that $h(f(\alpha_{a_j})) = f(\alpha_{b_j})$ for every $j \leq k$. In other words, h witnesses (2). So \mathcal{S}_0^f is quasi-saturated. Moreover, for every vertex u of \mathcal{S}_0 , $\lambda_0(u)^f \sqsubseteq \lambda_0(u) \sqsubseteq \lambda(u)$ by the choice of f , hence $\lambda_0(u)^f \subseteq \lambda(u)$ by the Truth Lemma and the additional assumption on \mathcal{T} . But then $\text{Ker } u \subseteq \lambda_0(u)^f \subseteq L_\mu^{\omega^\omega}$. \square

Thus we obtain the following theorem.

Theorem 6. $\mathcal{K}_{\mu^+}^{\omega^\omega}$ is sound and complete system over symmetric frames.

Proof. Suppose $\Gamma = \Delta[\nu x \phi]$ is not valid. Applying Theorem 5 we obtain a κ -system expanding Γ , which can be further expanded to a system \mathcal{S} satisfying the assumptions of Proposition 1 with the additional property that the vertex u which contains the formula $\nu x \phi$ specified by the context is an element of the designated finite set U . As a consequence of the proposition, $\text{Ker } u \subseteq L_\mu^{\omega^\omega}$. By saturation, $\nu^\alpha x \phi' \in u$ for some $\alpha < \omega^\omega$ and $\phi' \sqsubseteq \phi$, whence the Truth Lemma implies $\Gamma[\nu^\alpha x \phi]$ is not valid. Thus the rule ν_{ω^ω} is sound. Completeness is given by Theorem 1. \square

It is not difficult (though rather technical) to show that the ordinal ω^ω is optimal for obtaining soundness over trees by leveraging the failure of the finite model property. For instance, to observe that the inference ν_{ω^2} is unsound (over trees), consider the sequent $\Gamma = \varrho, \psi, \phi$ where ψ expresses the existence of a finite $\{\mathbf{a}, \mathbf{b}\}$ -path, $\varrho = \langle \bar{\mathbf{a}} \rangle \top \vee \langle \bar{\mathbf{b}} \rangle \top$ and

$$\phi = \nu x(\langle \bar{\mathbf{b}} \rangle(x \wedge \varrho) \wedge \mu y \langle \mathbf{b} \rangle(y \vee x) \wedge \mu y \langle \mathbf{a} \rangle(y \vee x)).$$

This observation can be readily generalised to show ν_{ω^n} is unsound for each n . Combining with the previous theorem we conclude

Theorem 7. $K_{\mu^+}^\kappa$ is unsound over trees for every $\kappa < \omega^\omega$.

5 Discussion

There is an interesting tradeoff between the difficulty in establishing soundness and completeness for different axiomatisations of μ -calculus. With Kozen's axiomatisation the difficulty lies in showing completeness (soundness being reasonably straightforward) whereas in the goal-oriented proof system of [28] or the circular axiomatisations proposed in [1] the proof of soundness is more involved. The infinitary proof system $K_{\mu^+}^{\omega^\omega}$ belongs to this second category.

Finally, we wish to remark on one further result contained in Vardi's seminal article: the tree languages definable by μ -calculus formulæ with converse modalities are precisely those definable by formulæ *without* converse. Suppose $\phi \mapsto \phi^*$ is an effective translation of formulæ into pure formulæ such that $\phi \leftrightarrow \phi^*$ is true in the symmetric closure of every tree. To re-phrase Vardi's result, an arbitrary tree can be endowed with a saturated ω^ω -system containing ϕ in the root iff it can be given a saturated ω -system with root containing ϕ^* . Since we know that the ordinals ω^ω and ω are optimal for the respective languages (over trees), this leads us to wonder what features of the interpretation give rise to this necessary collapse (the 'only if' direction) and expansion ('if' direction) of ordinals. We cannot say at this stage, but believe questions in this vein demonstrate a clear gap in our understanding of the proof theory of fixed point logic.

Acknowledgements The work was initiated during the authors' research visit to the Hausdorff Research Institute for Mathematics (HIM), University of Bonn, as part of the trimester program *Types, Sets and Constructions, May–Aug 2018*. Both the financial support and the hospitality of HIM, are gratefully acknowledged.

This research was supported by the the Swedish Research Council (grants 2016-03502 and 2017-05111) and the Knut and Alice Wallenberg Foundation.

The authors also wish to thank Rajeev Goré for his unsparing advice to look at nested sequents, Steve Simpson for his interest and references to the theory of better-quasi-orders, and Valentin Goranko for his suggestions which have improved the final presentation of the results.

A Well-annotated formulæ

We begin by making more precise the definition of well-annotated κ -formulæ, and the properties that this class satisfy.

Fix $\kappa \leq \omega_1$ and let \triangleleft denote the subsumption ordering on \mathbf{Var} , where $x \triangleleft y$ reads as x **subsumes** y . We assume \triangleleft is a strict partial order on \mathbf{Var} which is downwards linear. Recall that we consider \triangleleft fixed and that all formulæ respect \triangleleft . Hence, if $\mu y \phi$ is a formula with x free, then $x \triangleleft y$.

An κ -**assignment** is a partial function from \mathbf{Var} into ordinals $< \kappa$ whose domain is linearly ordered by \triangleleft . $\mathbb{A}(\kappa)$ is the set of κ -assignments and we let $\text{dom } o$ denote the domain of $o \in \mathbb{A}(\kappa)$. It proves convenient to occasionally treat assignments as total functions $o: \mathbf{Var} \rightarrow \kappa + 1$, and set $\text{dom } o = \{x \in \mathbf{Var} \mid o(x) < \kappa\}$. Given $o \in \mathbb{A}(\kappa)$ and $x \in \mathbf{Var}$, $o_{\triangleleft x}$ denotes the restriction of o to the variables subsuming x :

$$o_{\triangleleft x}(y) = \begin{cases} o(y), & \text{if } y \triangleleft x, \\ \kappa, & \text{otherwise.} \end{cases}$$

For $\phi \in L_\mu$ and $o \in \mathbb{A}(\kappa)$, ϕ^o is the κ -formula generated as follows.

$$\begin{aligned} x^o &= x & (\phi \wedge \psi)^o &= \phi^o \wedge \psi^o & ([\mathbf{a}]\phi)^o &= [\mathbf{a}]\phi^o & (\mu x \phi)^o &= \mu x \phi^{o_{\triangleleft x}} \\ (\phi \vee \psi)^o &= \phi^o \vee \psi^o & (\langle \mathbf{a} \rangle \phi)^o &= \langle \mathbf{a} \rangle \phi^o & (\nu x \phi)^o &= \nu^{o(x)} x \phi^{o_{\triangleleft x}} \end{aligned}$$

That is, ν -quantifiers in ϕ^o are approximated by their value under o (which is no approximation if the variable is outside the domain) except for variables occurring within the scope of a variable lower in the subsumption ordering. The significance of constraining $\text{dom } o$ to be linearly ordered will become apparent shortly when we consider a quasi-ordering of $\mathbb{A}(\kappa)$.

Example 1. Suppose $x \triangleleft y$ and $o(x) = \alpha$ and $o(y) = \beta$, with $\alpha, \beta < \kappa$. Let ϕ be a formula without quantifiers containing both x and y free. Then $(\nu y \nu x \phi)^o = \nu^\beta y \nu^\alpha \phi$, whereas $(\nu x \nu y \phi)^o = \nu^\alpha x \nu y \phi$. The requirement that $\text{dom } o$ is linear means that if $((\nu x \phi) \vee (\nu z \psi))^o = (\nu^\alpha x \phi') \vee (\nu^\gamma z \psi')$ then either one of α and γ is κ , or x and z are comparable in \triangleleft .

Definition 10. *The image of a well-formed formula under a κ -assignment is **well-annotated**. We let L_μ^κ be the set of well-annotated κ -formulæ.*

Recall that substitution is well-defined for well-formed formulæ.

Lemma 6. *If ϱ is well-named then every formula in $\mathbb{S}\mathbb{C}_\kappa(\varrho)$ is well-annotated.*

Proof. Suppose $\phi = (\nu x \psi)^o = \nu^\alpha x \psi^{o_{\triangleleft x}} \in \mathbb{S}\mathbb{C}_\kappa(\varrho)$. Then for each $\beta < \alpha$, we have $\phi' = \psi^{o_{\triangleleft x}}(x/\nu^\beta x \psi^{o_{\triangleleft x}}) \in \mathbb{S}\mathbb{C}_\kappa(\varrho)$ by the closure condition and we require to show that ϕ' is well-annotated. Assume $o(x) < \kappa$ (otherwise the result is immediate) and let o' be the assignment with domain $\{y \in \text{dom } o \mid y \triangleleft x \vee y = x\}$ determined by $o'(y) = o(y)$ for $y \triangleleft x$ and $o'(x) = \beta$. Given the fact that ϱ is well-named, x does not appear bound in ψ , whence it is easy to check that $\phi' = \psi(x/\nu x \psi)^{o'}$. The other closure conditions are straightforward.

As defined, κ -assignments do not uniquely determine the formulæ in L_μ^κ . Each $\psi \in L_\mu$ determines an obvious equivalence relation on $\mathbb{A}(\kappa)$, given by $o \sim_\psi o'$ iff $\psi^o = \psi^{o'}$. However, for each $\phi \in L_\mu^\kappa$ there exists a unique κ -assignment o with smallest domain such that $\phi = \psi^o$, where $\psi = \phi^-$ is the template of ϕ . We call this assignment the **ordinal assignment** of ϕ and denote it o_ϕ .

We can thus give the formal definition of the quasi-order \sqsubseteq introduced immediately prior to Definition 9. This starts with a quasi-order \leq on κ -assignments, defined by $o \leq \hat{o}$ iff $\text{dom } o \subseteq \text{dom } \hat{o}$ and for every maximal chain $x_0 \triangleleft x_1 \triangleleft \dots \triangleleft x_n \in \text{dom } o$ the sequence $(o(x_0), \dots, o(x_n))$ is lexicographically prior to $(\hat{o}(x_0), \dots, \hat{o}(x_n))$.

Lemma 7. $(\mathbb{A}(\kappa), \leq)$ is a well-quasi-order. Moreover, there exists k such that for every set $X \subseteq \mathbb{A}(\kappa)$ with $|X| \geq k$ there exists $o, \hat{o} \in X$ s.t. $o < \hat{o}$.

Proof. Transitivity of \leq is established by induction along \triangleleft in Var . So, \leq is a quasi-order. Moreover, this quasi-order is a well-order on sets of κ -assignments with the same domain since it reduces to the lexicographic ordering on κ^k for some k (as domains are linearly ordered by \triangleleft). Since Var is a finite set, both claims follow. \square

Definition 11. Fix $\varrho \in L_\mu$ and for $\phi, \psi \in \mathbb{S}\mathbb{C}_\kappa(\varrho)$ define $\phi \sqsubseteq \psi$ iff $\phi^- = \psi^-$ and $o_\phi \leq o_\psi$.

This relation is well-defined because of Lemma 6, which implies that every formula in the strong closure of an L_μ formula is well-annotated and, hence, has a defined ordinal assignment.

We consider it instructive to note that there is another natural quasi-order sitting strictly between Kozen's \preceq and our \sqsubseteq , obtained by dropping the restriction of linearity of annotated quantifiers but otherwise applying the lexicographic ordering in \sqsubseteq . This too is a wqo, but does not satisfy the second part of Lemma 4.

B Omitted proofs

We now present some missing arguments from the main text. We begin with Lemma 4 as this result follows directly from our work on ordinal assignments:

Lemma 4. $(\mathbb{S}\mathbb{C}_\kappa(\varrho), \sqsubseteq)$ is a wqo. Moreover, there exists k such that for every $X \subseteq \mathbb{S}\mathbb{C}_\kappa(\varrho)$, $|\text{Ker } X| < k$.

Proof. We need only remark that the quasi-order $(\mathbb{S}\mathbb{C}_\kappa(\varrho), \sqsubseteq)$ can be expressed as the disjoint union of finitely many copies of $(\mathbb{A}(\kappa), \leq)$, one for each formula in $\mathbb{F}\mathbb{L}(\varrho)$, an operation that preserves wqo-ness. The second claim follows from this fact and Lemma 7.

Lemma 1(2): For all $\phi \in L_\mu$ and contexts $\Gamma[\]$, $\mathbb{K}_\mu^\kappa \vdash \Gamma[\phi, \bar{\phi}]$.

Proof. Induction on $\phi = \phi(x_1, \dots, x_k)$ shows the inference

$$\frac{\Gamma[\psi_1, \chi_1] \quad \dots \quad \Gamma[\psi_k, \chi_k]}{\Gamma[\bar{\phi}(\psi_1, \dots, \psi_k), \phi(\chi_1, \dots, \chi_k)]}$$

is admissible in K_μ^κ and $\mathsf{K}_{\mu^+}^\kappa$. For the case $\phi = \nu y \phi_0$, we have a derivation of $\Gamma[\bar{\phi}, \nu^0 y \phi_0]$ by $\nu.0$, and from $\Gamma[\bar{\phi}, \nu^\alpha y \phi_0]$ we derive $\Gamma[\bar{\phi}, \nu^{\alpha+1} y \phi_0]$ via the induction hypothesis and inferences μ and $\nu.(\alpha + 1)$. Thus transfinite induction shows that $\Gamma[\bar{\phi}, \nu^\alpha y \phi_0]$ is derivable for every $\alpha < \kappa$, whence $\Gamma[\bar{\phi}, \phi]$ results.

Theorem 4. The proof of this theorem ends with a statement of the following equivalence:

$$\forall X, Y \subseteq \mathsf{SC}_\kappa(\varrho) : X \sqsubseteq^* Y \quad \text{iff} \quad X^* \leq_{\text{pw}} Y^*$$

On first appearance this result appears non-trivial. However, it is an easy consequence of the following result relating finite sets of ordinals, the verification of which is straightforward.

Lemma 8. *Given a non-empty finite set of ordinals A , let $A^* = (\delta_i^A)_{i < |A|}$ denote the unique sequence such that $A = \{\sum_{j \leq i} \delta_j^A \mid i < |A|\}$. Fix a principal ordinal κ and let $A, B \subset \kappa$ be non-empty finite sets of the same cardinality. There exists $f \in \mathbb{I}(\kappa)$ such that $B = \{f(\alpha) \mid \alpha \in A\}$ iff $A^* \leq_{\text{pw}} B^*$.*

References

- [1] B. Afshari and G. E. Leigh. ‘Cut-free completeness for modal mu-calculus’. In: **32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS’17)**. IEEE Computer Society, 2017, pp. 1–12. DOI: 10.1109/LICS.2017.8005088.
- [2] R. Arisaka, A. Das and L. Straßburger. ‘On nested sequents for constructive modal logics’. In: **Logical Methods in Computer Science** 11.3 (2015). DOI: 10.2168/LMCS-11(3:7)2015.
- [3] J. Bradfield and C. Stirling. ‘12 Modal mu-calculi’. In: **Handbook of Modal Logic**. Ed. by P. Blackburn, J. V. Benthem and F. Wolter. Vol. 3. Studies in Logic and Practical Reasoning. Elsevier, 2007, pp. 721–756. DOI: 10.1016/S1570-2464(07)80015-2.
- [4] J. Bradfield and I. Walukiewicz. ‘The mu-calculus and model checking’. In: **Handbook of Model Checking**. Ed. by E. M. Clarke, T. A. Henzinger, H. Veith and R. Bloem. Cham: Springer International Publishing, 2018, pp. 871–919. DOI: 10.1007/978-3-319-10575-8_26.
- [5] T. Cachat. ‘Two-way tree automata solving pushdown games’. In: **Automata Logics, and Infinite Games: A Guide to Current Research**. Ed. by E. Grädel, W. Thomas and T. Wilke. New York, NY, USA: Springer-Verlag New York, Inc., 2002, pp. 303–317. DOI: 10.1007/3-540-36387-4_17.

- [6] M. Dam and D. Gurov. ‘ μ -calculus with explicit points and approximations’. In: **Journal of Logic and Computation** 12.2 (2002), pp. 255–269. DOI: 10.1093/logcom/12.2.255.
- [7] S. Demri, V. Goranko and M. Lange. **Temporal Logics in Computer Science: Finite-State Systems**. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016. DOI: 10.1017/CB09781139236119.
- [8] S. Enqvist. ‘Flat modal fixpoint logics with the converse modality’. In: **Journal of Logic and Computation** 28.6 (June 2018), pp. 1065–1097. DOI: 10.1093/logcom/exy016.
- [9] S. Enqvist, F. Seifan and Y. Venema. ‘Completeness for the modal μ -calculus: Separating the combinatorics from the dynamics’. In: **Theoretical Computer Science** 727 (2018), pp. 37–100. DOI: 10.1016/j.tcs.2018.03.001.
- [10] M. Fitting. ‘Nested sequents for intuitionistic logics’. In: **Notre Dame Journal of Formal Logic** 55.1 (2014), pp. 41–61. DOI: 10.1215/00294527-2377869.
- [11] M. Fitting and R. Kuznets. ‘Modal interpolation via nested sequents’. In: **Annals of Pure and Applied Logic** 166.3 (2015), pp. 274–305. DOI: 10.1016/j.apal.2014.11.002.
- [12] V. Goranko and A. Galton. ‘Temporal Logic’. In: **The Stanford Encyclopedia of Philosophy**. Ed. by E. N. Zalta. Winter 2015. Metaphysics Research Lab, Stanford University, 2015.
- [13] R. Goré. ‘And-or tableaux for fixpoint logics with converse: LTL, CTL, PDL and CPDL’. In: **Automated Reasoning - 7th International Joint Conference, IJCAR 2014**. Ed. by S. Demri, D. Kapur and C. Weidenbach. Vol. 8562. Lecture Notes in Computer Science. Springer, 2014, pp. 26–45. DOI: 10.1007/978-3-319-08587-6_3.
- [14] R. Goré, L. Postniece and A. Tiu. ‘On the correspondence between display postulates and deep inference in nested sequent calculi for tense logics’. In: **Logical Methods in Computer Science** 7.2 (2011). DOI: 10.2168/LMCS-7(2:8)2011.
- [15] E. Grädel, W. Thomas and T. Wilke, eds. **Automata Logics, and Infinite Games: A Guide to Current Research**. New York, NY, USA: Springer-Verlag New York, Inc., 2002. DOI: 0.1007/3-540-36387-4.
- [16] G. Jäger, M. Kretz and T. Studer. ‘Canonical completeness of infinitary μ ’. In: **Journal of Logic and Algebraic Programming** 76.2 (2008), pp. 270–292. DOI: 10.1016/j.jlap.2008.02.005.
- [17] R. Kashima. ‘Cut-free sequent calculi for some tense logics’. In: **Studia Logica** 53.1 (1994), pp. 119–135. DOI: 10.1007/BF01053026.
- [18] D. Kozen. ‘A finite model theorem for the propositional μ -calculus’. In: **Studia Logica** 47.3 (1988), pp. 233–241. DOI: 10.1007/BF00370554.
- [19] D. Kozen. ‘Results on the propositional μ -calculus’. In: **Theoretical Computer Science** 27 (1983), pp. 333–354. DOI: 10.1016/0304-3975(82)90125-6.

- [20] A. Marcone. ‘Fine analysis of the quasi-orderings on the power set’. In: **Order** 18.4 (Dec. 2001), pp. 339–347. DOI: 10.1023/A:1013952225669.
- [21] C. S. J. A. Nash-Williams. ‘On better-quasi-ordering transfinite sequences’. In: **Mathematical Proceedings of the Cambridge Philosophical Society** 64.2 (1968), pp. 273–290. DOI: 10.1017/S030500410004281X.
- [22] C. S. J. A. Nash-Williams. ‘On well-quasi-ordering transfinite sequences’. In: **Mathematical Proceedings of the Cambridge Philosophical Society** 61.1 (1965), pp. 33–39. DOI: 10.1017/S0305004100038603.
- [23] R. Parikh. ‘The completeness of propositional dynamic logic’. In: **Mathematical Foundations of Computer Science 1978**. Ed. by J. Winkowski. Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 403–415. DOI: 10.1007/3-540-08921-7_88.
- [24] M. Reynolds. ‘An axiomatization of PCTL*’. In: **Information and Computation** 201.1 (2005), pp. 72–119. DOI: 10.1016/j.ic.2005.03.005.
- [25] M. Reynolds. ‘More past glories’. In: **15th Annual IEEE Symposium on Logic in Computer Science (LICS’00)**. IEEE Computer Society, 2000, pp. 229–240. DOI: 10.1109/LICS.2000.855772.
- [26] D. S. Shamkanov. ‘Nested sequents for provability logic GLP’. In: **Logic Journal of the IGPL** 23.5 (2015), pp. 789–815. DOI: 10.1093/jigpal/jzv029.
- [27] M. B. Smyth. ‘Power domains’. In: **Journal of Computer System Sciences** 16 (1978), pp. 23–36. DOI: 10.1016/0022-0000(78)90048-X.
- [28] C. Stirling. ‘A tableau proof system with names for modal mu-calculus’. In: **HOWARD-60: A Festschrift on the Occasion of Howard Barringer’s 60th Birthday**. Ed. by A. Voronkov and M. V. Korovina. Vol. 42. EPiC Series in Computing. EasyChair, 2014, pp. 306–318.
- [29] R. S. Streett and E. A. Emerson. ‘An automata theoretic decision procedure for the propositional mu-calculus’. In: **Information and Computation** 81 (1989), pp. 249–264. DOI: 10.1016/0890-5401(89)90031-X.
- [30] M. Y. Vardi. ‘Reasoning about the past with two-way automata’. In: **Automata, Languages and Programming**. Ed. by K. G. Larsen, S. Skyum and G. Winskel. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 628–641. DOI: 10.1007/BFb0055090.
- [31] I. Walukiewicz. ‘Completeness of Kozen’s axiomatisation of the propositional mu-calculus’. In: **Proceedings, 10th Annual IEEE Symposium on Logic in Computer Science (LICS’95)**. IEEE Computer Society, 1995, pp. 14–24. DOI: 10.1109/LICS.1995.523240.