A circular proof system for the hybrid μ-calculus

Sebastian Enqvist [1]

Stockholm University, Department of Philosophy
Universitetsvägen 10 D
Frescati, Stockholm

Abstract

We present a cut-free circular proof system for the hybrid μ-calculus, and prove its soundness and completeness. The system is an adaptation of a circular proof system for the modal μ-calculus due to Stirling, and uses a system of annotations to keep track of fixpoint unfoldings. The language considered here extends the μ-calculus with nominals and satisfaction operators, but not the converse modality. This version of the hybrid μ-calculus is known to have the finite model property, unlike the version that includes converse. The presence of nominals and satisfaction operators causes some non-trivial difficulties to deal with in the completeness proof. In particular we need to be careful about what information attached to nominals to keep and what to discard, and furthermore the structure of traces in a proof-tree becomes more complicated. Still, it turns out that the proof system is complete with the same global condition for validity as Stirling's system. The key tool that we develop for the completeness proof is a proof-search game, in which one of the players attempts to construct a proof in a restricted normal form making use of certain derived rules. We conclude the paper with some tasks for future research, which include proving completeness of a cut-free non-circular sequent calculus, and extending the system developed here to incorporate converse modalities.

Keywords: Hybrid logic, μ-calculus, circular proofs, completeness, automata

1 Introduction

Circular and non-wellfounded proofs are a powerful method for reasoning with fixpoints, and have been considered in a number of contexts [19,6,22,3,4,21]. For the modal μ-calculus, a circular proof system with names for keeping track of fixpoint unfoldings was developed by Stirling [23], building on work by Jungteerapanich [12] and bearing similarities with earlier systems using variables for ordinal approximations [6]. Recently Stirling's system has been simplified and used by Afshari and Leigh to give a cut-free complete sequent system for the

[1] This research was supported by the Swedish Research Council grant 2015-01774.

modal μ-calculus [2]. This provides a novel completeness proof for Kozen's axiomatization [15] that avoids the intricate detour via disjunctive normal forms in Walukiewicz's proof [27]. Building on this work, circular proofs were used in [8] to settle the open problem of completeness of Parikh's logic of games [17].

The present work is intended as a step towards exploring the use of circular proofs to provide complete finitary proof systems for richer *extensions* of the modal μ-calculus. A number of such extensions have been presented in the literature, including the two-way or "full" μ-calculus [26], hybrid μ-calculus [20] and guarded fixpoint logic [11]. In many cases such extensions remain decidable. However, complete proof systems mostly appear to be lacking. Some work in this area does exist: a generic completeness result for coalgebraic versions of the μ-calculus (including extensions like the graded μ-calculus) was presented in [10]. This general result does not cover the hybrid μ-calculus however, since the sort of global conditions that are expressible in hybrid logics are out of scope for the framework in [10]. An infinitary proof system for the two-way μ-calculus was proved complete in [1].

As a proof of concept, we shall develop a cut-free Stirling-style circular proof system for the hybrid μ-calculus. Orignally introduced by Sattler and Vardi in [20], the hybrid μ-calculus features *nominals*, which are used to name points in a model, and *satisfaction operators* that describe what is true at a named point in a model. We shall follow Tamura [24] by not including converse modalities in our language, unlike Sattler and Vardi. Tamura shows that the hybrid μ-calculus without converse has the finite model property, unlike the more expressive version considered by Sattler and Vardi. We also mention a version of the hybrid μ-calculus involving a binder modality, which was investigated in [13] under the name "fully hybrid μ-calculus". This logic is undecidable, and therefore seems out of scope for the kind of methods that we consider here.

The presence of nominals and satisfaction operators already presents some non-trivial challenges for the completeness proof, and addressing these difficulties gives some guidelines on how to deal with proof theory for fixpoint logics that lack the tree model property. In a manner of speaking, we are continuing here along Sattler and Vardi's line of working with logics that lack the tree model property "as if they had the tree model property" [20], but taking the idea in a proof-theoretic direction.

Proofs have been removed or shortened due to page limitations. For a longer version of this paper including detailed proofs, see the preprint availabe online at `https://arxiv.org/abs/2001.04971`.

2 Preliminaries

2.1 The hybrid μ-calculus

The hybrid μ-calculus was initially introduced by Sattler and Vardi in [20]. Their version of the language included a global modality and converse modalities. Here, we shall be considering the weaker version of the hybrid μ-calculus that was studied by Tamura in [24]. For ease of notation we consider the language with only a single box and diamond, but all the results and proofs presented here easily extend to a multi-modal version of the language.

The language $\mathcal{L}$ of the hybrid μ-calculus is given by the following grammar:

$$\varphi := p \mid \neg p \mid \mathsf{i} \mid \neg\mathsf{i} \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \Box\varphi \mid \mathsf{i}\!:\!\varphi \mid \mu x.\varphi \mid \nu x.\varphi$$

Here, p and x are members of a fixed countably infinite supply Prop of propositional variables, and i comes from a fixed countably infinite supply Nom of nominals. For $\eta x.\varphi$ with $\eta \in \{\mu, \nu\}$, we impose the usual constraint that no occurrence of x in φ is in the scope of a negation, and we also require that each occurrence of x in φ is within the scope of some modality ($\Box$ or $\Diamond$). This latter extra constraint means that we restrict attention to *guarded* formulas. This is a fairly common assumption, and it is well known that removing the constraint of guardedness does not increase the expressive power of the language. It is not an entirely innocent assumption however, since putting a formula in its guarded normal form may cause an exponential blow-up in the size of a formula [5].

Note also that the language is presented in negation normal form. It is routine to verify, given the semantics presented below, that the language is semantically closed under negation, and furthermore there is a simple effective procedure for converting formulas in the extended language with explicit negation of all formulas into formulas in negation normal form.

Free and bound variables of a formula are defined in the usual manner. A *literal* is a formula of the form p or $\neg p$ where $p \in \mathsf{Prop}$, or of the form i or $\neg\mathsf{i}$ where $\mathsf{i} \in \mathsf{Nom}$. We introduce the following abbreviations:

$$\mathsf{i} \approx \mathsf{j} := \mathsf{i}\!:\!\mathsf{j} \qquad \mathsf{i} \not\approx \mathsf{j} := \mathsf{i}\!:\!\neg\mathsf{j}$$

These formulas express identity and non-identity, respectively, of the values assigned to the nominals i, j in a model.

Definition 1 *Let φ be any formula in $\mathcal{L}$ and let $x, y \in \mathsf{Prop}$ be bound variables in φ. We say that y is* dependent *on x, written $x <_\varphi y$, if there is a subformula of φ of the form $\eta y.\psi$ in which there is a free occurrence of x. We denote the reflexive closure of $<_\varphi$ by $\leq_\varphi$.*

Definition 2 *We say that a formula φ is* locally well-named *if $<_\varphi$ is irreflexive, no variable occurs both free and bound in φ, and no variable is bound by both μ and ν in φ.*

Note that every formula is equivalent to a locally well-named one up to renaming of bound variables (α-equivalence).

Proposition 2.1 (Afshari & Leigh -17) *If $\eta x.\varphi(x)$ is locally well-named then so is $\varphi(\eta x.\varphi)$.*

Convention 1 *We shall assume throughout the paper that all formulas are locally well-named. Given a locally well-named formula we refer to a bound variable x as a μ-variable if it is bound (only) by μ in φ, and a ν-variable if it is bound (only) by ν.*

Semantics of the hybrid μ-calculus is a simple extension of the usual Kripke semantics for the modal μ-calculus.

Definition 3 *A* Kripke model *is a tuple* $\mathcal{M} = (W, R, V, A)$ *where* W *is a non-empty set members of which will be referred to as* points, $R \subseteq W \times W$ *is the* accessibility relation *over* W, $V : \mathsf{Prop} \to \mathcal{P}(W)$ *is a* valuation *of the propositional variables and* $A : \mathsf{Nom} \to W$ *is an* assignment *of a value in* W *to each nominal.*

Given a Kripke model $\mathcal{M} = (W, R, V, A)$, the interpretation $[\![\varphi]\!]_{\mathcal{M}}$ of a formula φ is defined by the usual recursive clauses for boolean connectives and modalities. Semantics of least fixpoint operators is given according to the Knaster-Tarski Theorem [14,25] as:

$$[\![\mu x.\varphi(x)]\!]_{\mathcal{M}} := \bigcap\{Z \subseteq W \mid [\![\varphi]\!]_{\mathcal{M}[Z/x]} \subseteq Z\},$$

where $\mathcal{M}[Z/x]$ is like $\mathcal{M}$ except that its valuation maps the variable x to Z. For greatest fixpoint operators we have the dual definition:

$$[\![\nu x.\varphi(x)]\!]_{\mathcal{M}} := \bigcup\{Z \subseteq W \mid Z \subseteq [\![\varphi]\!]_{\mathcal{M}[Z/x]}\}.$$

For nominals and satisfaction operators, we have the following clauses: $[\![\mathsf{i}]\!]_{\mathcal{M}} = \{A(\mathsf{i})\}$ and $[\![\mathsf{i} : \varphi]\!]_{\mathcal{M}} = \{w \in W \mid A(\mathsf{i}) \in [\![\varphi]\!]\}$. In other words, $[\![\mathsf{i} : \varphi]\!]_{\mathcal{M}} = W$ if $A(\mathsf{i}) \in [\![\varphi]\!]$, and $[\![\mathsf{i} : \varphi]\!]_{\mathcal{M}} = \emptyset$ otherwise. Given a formula φ and a pointed Kripke model $(\mathcal{M}, w)$ (a model with a distinguished point), we write $\mathcal{M}, w \Vdash \varphi$ to say that $w \in [\![\varphi]\!]_{\mathcal{M}}$.

This semantics may be referred to as the *denotational* semantics of the μ-calculus. The μ-calculus also has an *operational* semantics in the form of a game semantics, which is often easier to work with and neatly captures the intuitive meaning of least and greatest fixpoints (i.e. “finite looping” vs “infinite looping”). In this game semantics it is convenient to work with the (Fischer-Ladner) *closure* $c(\varphi)$ of a formula. The precise definition is a straightforward adaptation of that in [20], with the new clause that $\mathsf{i} : \theta \in c(\varphi)$ imples $\theta \in c(\varphi)$.

Throughout the paper we assume familiarity with basic notions concerning board games and parity games (see [9] for a very brief introduction). Given a Kripke model $\mathcal{M} = (W, R, V, A)$, the *evaluation game* for a formula $\rho \in \mathcal{L}$ in the model $\mathcal{M}$ is a two-player board game between players **Ver**, **Fal**, the set of positions of which is $W \times c(\rho)$, with player assignments and moves defined as follows:

- For a position of the form (w, l) where l is a literal, the set of available moves is $\emptyset$. The position is assigned to **Fal** if $\mathcal{M}, w \Vdash l$ and is assigned to **Ver** otherwise.
- For a position of the form $(w, \varphi O \psi)$ where $O \in \{\wedge, \vee\}$, the available moves are (w, φ) and (w, ψ). The position is assigned to **Ver** if $O = \vee$ and is assigned to **Fal** if $O = \wedge$.
- For a position of the form $(w, O\varphi)$ where $O \in \{\Diamond, \Box\}$, the set of available moves is $\{(v, \varphi) \in W \mid wRv\}$. The position is assigned to **Ver** if $O = \Diamond$ and is assigned to **Fal** if $O = \Box$.

- For a position of the form $(w, \mathsf{i} : \varphi)$, the unique avaliable move is $(A(\mathsf{i}), \varphi)$. The player assignment is arbitrary in this case since there is only one move, but as a convention we assign such positions to player **Ver**.
- For a position of the form $(w, \eta x.\varphi(x))$, the unique available move is $(w, \varphi(\eta x.\varphi))$. By convention we assign such positions to **Ver**.

Partial plays, full plays and strategies for players are defined as usual. Note that if a full play is finite, then the player to which the last position is assigned must be "stuck", i.e. the set of available moves is empty. So the winning condition of finite full plays is defined by declaring the player who got stuck to be the loser of the play. For infinite plays $(w_1, \varphi_1)(w_2, \varphi_2)(w_3, \varphi_3)...$, say that a fixpoint variable x is *unfolded* at the index i if φ_i is of the form $\eta x.\psi(x)$.

Proposition 2.2 *For any (locally well-named) formula ρ and any infinite play π in the evaluation game in $\mathcal{M}$, there is a unique $<_\rho$-minimal variable x that is unfolded infinitely many times on π.*

We shall often refer to the $<_\rho$-minimal variable unfolded infinitely often on π as the *highest ranking* variable that is unfolded infinitely often. We can now define the winning condition of infinite plays: the winner is **Ver** if the highest ranking variable that gets unfolded infinitely often is a ν-variable (relative to ρ), and the winner is **Fal** otherwise.

A strategy is called *positional* if it only depends on the last position of a play, i.e. it can be described as a choice function from positions to available moves. Since the evaluation game is a parity game, and parity games have positional determinacy [7,28], we have:

Proposition 2.3 *The evaluation game of any formula in a model is determinate, and the winning player at any given position has a positional winning strategy.*

As expected the operational semantics agrees with the denotational one:

Proposition 2.4 *Given a pointed Kripke model $(\mathcal{M}, w)$ and a formula ρ, we have $\mathcal{M}, w \Vdash \rho$ if and only if the position (w, ρ) is winning for **Ver** in the evaluation game.*

3 Infinite proofs

In this section we define an infinite sequent-style proof system **Inf** for the hybrid μ-calculus. This proof system will be used as a tool to prove completeness of the finite circular proof system that will be introduced in Section 5.1. The infinite system presented here is essentially dual to an infinite tableau system for the hybrid μ-calculus. An important difference from the tableaux developed by Sattler and Vardi in [20] is that the system is cut-free, which is required since the finitary circular system we shall present later will also be cut-free. Sattler and Vardi's automata-theoretic approach relies on "guessing" all the relevant information about some nominals at the start of the tableau construction. In the dual setting of sequent calculi this amounts to starting the proof construction with a series of cuts.

3.1 The system Inf

We will work with a sequent style proof system, where a sequent is a finite set of formulas interpreted as an implicit disjunction. It will be convenient to require that every formula in a sequent starts with some satisfaction operator, so each sequent has the form:

$$\mathsf{i}_1 : \varphi_1, ..., \mathsf{i}_n : \varphi_n$$

This means that our proof system will only prove formulas of this shape. However, this is not a serious restriction: given a formula φ that is not in the required format, we can always replace it with the formula $\mathsf{i} : \varphi$ where i is some arbitrarily chosen, fresh nominal not appearing in φ. Clearly $\mathsf{i} : \varphi$ is then semantically valid if and only if φ is, and we may regard any proof of $\mathsf{i} : \varphi$ as a proof of φ.

The system has two axioms, which are the law of exluded middle and an identity axiom:

$$\mathsf{i} : p, \mathsf{i} : \neg p \qquad \mathsf{i} \approx \mathsf{i}$$

Here, p is a nominal or a propositional variable. Rules of inference are given in Figure 1. We remark that, in the modal rule Mod, the nominal j must be fresh, i.e. it cannot appear in any formula in the conclusion of the rule.

$$\frac{\Gamma, \mathsf{i} : \varphi \wedge \psi, \mathsf{i} : \varphi \qquad \Gamma, \mathsf{i} : \varphi \wedge \psi, \mathsf{i} : \psi}{\Gamma, \mathsf{i} : \varphi \wedge \psi} \wedge \qquad \frac{\Gamma, \mathsf{i} : \varphi \vee \psi, \mathsf{i} : \varphi, \mathsf{i} : \psi}{\Gamma, \mathsf{i} : \varphi \vee \psi} \vee$$

$$\frac{\Gamma, \mathsf{i} : \eta x.\phi(x), \mathsf{i} : \phi(\eta x.\phi(x))}{\Gamma, \mathsf{i} : \eta x.\phi(x)} \eta x$$

$$\frac{\Gamma, \mathsf{i} : \phi, \mathsf{j} : \phi, \mathsf{i} \not\approx \mathsf{j}}{\Gamma, \mathsf{i} : \phi, \mathsf{i} \not\approx \mathsf{j}} \mathsf{Eq} \qquad \frac{\Gamma, \mathsf{i} : (\mathsf{j} : \varphi), \mathsf{j} : \varphi}{\Gamma, \mathsf{i} : (\mathsf{j} : \varphi)} \mathsf{Glob} \qquad \frac{\Gamma, \mathsf{i} \not\approx \mathsf{j}, \mathsf{j} \not\approx \mathsf{i}}{\Gamma, \mathsf{i} \not\approx \mathsf{j}} \mathsf{Com}$$

$$\frac{\Gamma, \mathsf{i} : \Box\varphi, \mathsf{i} : \Diamond\Psi, \mathsf{j} : \varphi, \mathsf{j} : \Psi}{\Gamma, \mathsf{i} : \Box\varphi, \mathsf{i} : \Diamond\Psi} \mathsf{Mod} \qquad \frac{\Gamma}{\Gamma \cup \Psi} \mathsf{Weak}$$

Fig. 1. Rules of **Inf**

In an application of the modal rule as shown in Figure 1, we refer to $\mathsf{i} : \Box\varphi$ as the principal formula. The expression $\mathsf{i} : \Diamond\Psi$ is short-hand for $\{\mathsf{i} : \Diamond\psi \mid \psi \in \Psi\}$, and likewise $\mathsf{j} : \Psi$ abbreviates $\{\mathsf{j} : \psi \mid \psi \in \Psi\}$. The intuition behind the modal rule is that, if the formulas $\Box\varphi, \Diamond\psi_1, ..., \Diamond\psi_n$ are all false at a point named i, then this must be witnessed by some point that we can give an arbitrary name j, and at which all the formulas $\varphi, \psi_1, ..., \psi_n$ are false. In an application of the rule Eq as shown in the figure, the formula $\mathsf{i} : \phi$ is called the principal formula and $\mathsf{i} \not\approx \mathsf{j}$ the *side formula*. In all other cases where a notion of principal formula makes sense, it should be clear from the form of the rules what the principal formula is. Note that we do allow that the premises and conclusion of a rule application are all the same sequent.

Definition 4 *A rule application is said to be* repeating *if all premises are equal to the conclusion.*

Definition 5 *A* **Inf**-proof*, or proof-tree, is a ranked labelled tree where the label of a node specifies the sequent appearing at the node, the rule application of which the node is the conclusion (if any), and the principal formula (if any), and such that the labels of children of a node are the premises of the specified rule application.*

We shall often abuse terminology slightly by referring to the sequent appearing at a node in a proof as the label of the node. To distinguish *valid* proofs from invalid ones, we need a notion of trace.

Definition 6 *A* partial trace $\boldsymbol{t}$ *(of length $k \leq \omega$) on a branch β of an* ***Inf****-proof Π is a sequence $(u_j, \mathsf{i}_j : \psi_j)_{j<k}$ such that for each j, u_j is a node on β whose label contains $\mathsf{i}_j : \psi_j$, u_{j+1} is the unique successor of u_j in β whenever $j+1 < k$, and one of the following conditions holds if $j+1 < k$:*

(i) $\mathsf{i}_j : \psi_j = \mathsf{i}_{j+1} : \psi_{j+1}$. *We sometimes refer to such parts of traces as "silent steps".*

(ii) $\mathsf{i}_j : \psi_j = \mathsf{i}_j : (\theta_1 \vee \theta_2)$ *is the principal formula in an application of the $\vee$-rule, and $\mathsf{i}_{j+1} : \psi_{j+1} \in \{\mathsf{i}_j : \theta_1, \mathsf{i}_j : \theta_2)\}$.*

(iii) $\mathsf{i}_j : \psi_j = \mathsf{i}_j : (\theta_1 \wedge \theta_2)$ *is the principal formula in an application of the $\wedge$-rule, and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{i}_j : \theta_1$ or $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{i}_j : \theta_2$ depending on whether u_{j+1} is the left or right premise of the rule.*

(iv) $\mathsf{i}_j : \psi_j = \mathsf{i}_j : \mathsf{i}' : \theta$ *is the principal formula in an application of the* Glob*-rule, and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{i}' : \theta$.*

(v) $\mathsf{i}_j : \psi_j$ *is the principal formula in an application of the* Eq*-rule with side formula $\mathsf{i}_j \not\approx \mathsf{i}'$, and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{i}' : \psi_j$.*

(vi) $\mathsf{i}_j : \psi_j = \mathsf{i}_j : \eta x.\theta(x)$ *is the principal formula in an application of the η-rule, and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{i}_j : \theta(\eta x.\theta(x))$. In this case we say that an* unfolding *of variable x occurs on the trace $\boldsymbol{t}$ at the index j.*

(vii) *u_j is the conclusion of an application of the* Mod*-rule labelled $\Gamma, \mathsf{i} : \Box\theta, \mathsf{i} : \Diamond\Psi$, the premise is labelled $\Gamma, \mathsf{j} : \theta, \mathsf{j} : \Psi, \mathsf{i} : \Box\theta, \mathsf{i} : \Diamond\Psi$, $\mathsf{i}_j : \psi_j = \mathsf{i} : \Box\theta$, and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{j} : \theta$.*

(viii) *u_j is the conclusion of an application of the* Mod*-rule labelled $\Gamma, \mathsf{i} : \Box\theta, \mathsf{i} : \Diamond\Psi$, the premise is labelled $\Gamma, \mathsf{j} : \theta, \mathsf{j} : \Psi, \mathsf{i} : \Box\theta, \mathsf{i} : \Diamond\Psi$, and for some $\psi \in \Psi$, $\mathsf{i}_j : \psi_j = \mathsf{i} : \Diamond\psi$ and $\mathsf{i}_{j+1} : \psi_{j+1} = \mathsf{j} : \psi$.*

A trace is said to be infinite if it is of length ω. We say that the infinite trace $\boldsymbol{t}$ is trivial *if for some $j < \omega$, $\mathsf{i}_j : \psi_j = \mathsf{i}_m : \psi_m$ for all $j \leq m < \omega$. A non-trivial infinite trace is said to be* good *if the highest ranking fixpoint variable that is unfolded infinitely many times on $\boldsymbol{t}$ is a ν-variable.*

Note that traces move along branches in the direction from conclusions to premises, i.e. traces travel away from the root, and not the other way around. Note also that we do not require traces to start at the root, but adding this constraint would make no substantial difference since every formula appearing in a sequent somewhere in a proof can be connected to a trace starting at the root.

Definition 7 *An **Inf**-proof is said to be* valid *if every infinite branch contains a good trace, and every leaf is labelled by an axiom.*

In order to produce finite circular proofs later on it will be important to carefully apply the weakening rule to discard formulas that are no longer needed and so maintain an upper bound on the size of sequents. The following terminology will play an important role in this regard.

Definition 8 *Given an **Inf**-proof* Π *for some formula* $\mathsf{r} : \rho$*, a nominal* j *appearing in* Π *is said to be* original *if it appears in* $\mathsf{r} : \rho$*. A formula appearing in* Π *is said to be a* ground formula *if it is of the form* $\mathsf{j} : \psi$ *where* j *is an original nominal.*

Definition 9 *An **Inf**-proof is said to be* frugal *if at most finitely many sequents appear in the proof.*

3.2 Derived rules

We shall allow the use of derived rules in proof constructions, as abbreviations of their derivations. These derived rules will be used to formulate a proof search game, which is the main technical tool needed for our completeness proof for **Inf**, and are based on two ideas:

- For all rules except Weak we define what we will call its *narrow* counterpart, which is a derived rule of **Inf**. These rules will be used to automatically discard formulas that will no longer be needed (using Weak), but keep those formulas that might be needed later in the proof construction.
- Two additional derived rules that we call the *deterministic* rule and the *ground rule* will be used to isolate the "essential" choices for the player that tries to construct a proof. These choices will be restricted to two types:

(i) Applications of the Mod-rule to introduce new nominals.
(ii) *Repeating* applications of other rules, which only serve to introduce traces.

Narrow rules We define the narrow rule versions as follows. For the $\wedge$- and $\vee$-rules, the η-rules, the Com-rule, the Glob-rule and the Eq-rule, if the principal formula is a ground formula, then the narrow version of the rule is the same as the standard one. Otherwise, it is defined as follows: we first apply the standard version of the rule, and immediately after we apply the weakening rule to all premises in order to remove the principal formula. For example, if i is a non-original nominal then an instance of the narrow $\wedge$-rule is:

$$\frac{\Gamma, \mathsf{i} : \varphi \qquad \Gamma, \mathsf{i} : \psi}{\Gamma, \mathsf{i} : (\varphi \wedge \psi)}$$

corresponding to the derivation:

$$\dfrac{\mathsf{Weak}\,\dfrac{\Gamma, \mathsf{i}\!:\!\varphi}{\Gamma, \mathsf{i}\!:\!(\varphi \wedge \psi), \mathsf{i}\!:\!\varphi} \qquad \dfrac{\Gamma, \mathsf{i}\!:\!\psi}{\Gamma, \mathsf{i}\!:\!(\varphi \wedge \psi), \mathsf{i}\!:\!\psi}\,\mathsf{Weak}}{\Gamma, \mathsf{i}\!:\!(\varphi \wedge \psi)}\,\wedge$$

The narrow version of the rule Mod is a bit different from the others: if the principal formula $\mathsf{i} : \Box\varphi$ is a ground formula then the rule is the same as Mod. Otherwise, an instance of the narrow rule consists of an application of the modal rule immediately followed by an application of the weakening rule in order to remove *all* formulas of the form $\mathsf{k}\!:\!\theta$ that appear in the premise, and for which k is not an original nominal. For example, if i is a non-original nominal and j is original, then the following is an instance of the narrow Mod-rule:

$$\dfrac{\mathsf{k}\!:\!\varphi, \mathsf{k}\!:\!\psi, \mathsf{j}\!:\!\theta}{\mathsf{i}\!:\!\Box\varphi, \mathsf{i}\!:\!\Diamond\psi, \mathsf{i}\!:\!p, \mathsf{j}\!:\!\theta}$$

If j is non-original then the corresponding instance would be:

$$\dfrac{\mathsf{k}\!:\!\varphi, \mathsf{k}\!:\!\psi}{\mathsf{i}\!:\!\Box\varphi, \mathsf{i}\!:\!\Diamond\psi, \mathsf{i}\!:\!p, \mathsf{j}\!:\!\theta}$$

Note that what counts as an instance of the narrow rules depends on what nominals are considered original, which in turn depends on the root formula of the proof-tree. We therefore emphasize that these rules are not explicitly part of the proof system **Inf**, but only serve as tools for the completeness proof.

Next, we define the deterministic rule and the ground rule. To make these rules precise we need the following:

Convention 2 *Throughout the rest of the paper we fix an arbitrary well-ordering $\prec$ over all formulas (which restricts to a well-ordering over the set of nominals since each nominal is a formula). Furthermore we fix an arbitrary well-ordering over the set of all instances of rules in **Inf**. We overload the notation and denote also this well-ordering by $\prec$.*

The deterministic rule The *deterministic rule* is defined as follows: given a sequent Γ, if there are no applicable instances of the narrow $\wedge$-rule, the narrow $\vee$-rule, the narrow Glob or narrow η-rules *except repeating ones*, then the deterministic rule does not apply. Otherwise, the deterministic rule applies uniquely as follows: we pick the $\prec$-smallest formula in Γ which is the principal formula in an applicable non-repeating instance of one of these rules, we pick the $\prec$-smallest such rule instance for which it is the principal formula, and we apply that rule.

Note that if we repeatedly apply the deterministic rule starting from some sequent Γ until it no longer applies, then this process must eventually terminate. The assumption that all formulas are guarded plays an important role here, without guardedness the process could go on indefinitely via fixpoint unfoldings.

The ground rule The *ground rule* is designed to deterministically apply the Eq-rule and the Com-rule in the same way as the deterministic rule, but also to ensure that original nominals are given special treatment. It is defined as follows: we consider the original nominals appearing in a sequent Γ. If possible, apply the $\prec$-smallest applicable rule instance for which one of the following conditions holds:

(i) it is a non-repeating instance of the narrow Com-rule with principal formula $\mathsf{i} \not\approx \mathsf{j}$, where both i and j are original nominals, or:

(ii) it is a non-repeating instance of the narrow Eq-rule with principal formula $\mathsf{j} : \varphi$ and side formula $\mathsf{j} \not\approx \mathsf{i}$, where i is a $\prec$-minimal original nominal for which such a rule instance applies.

If there are no such rule instances available then the ground rule does not apply. Like the deterministic rule, the process of repeatedly applying the ground rule must eventually terminate.

4 Completeness for Inf

4.1 A game for building Inf-proofs

To prove completeness we shall make use of a proof search game, played between two players **Ver** (the proponent) and **Fal** (the opponent). We fix a root formula $\mathsf{r} : \rho$, so that what counts as a narrow rule is defined relative to this root formula as before, and similarly with the deterministic rule and the ground rule.

Definition 10 *An instance of the weakening rule is called* terminal *if its premise is an axiom.*

Definition 11 *The* **Inf**-game *is a board game, defined as usual by specifying its positions, player assignments and admissible moves for positions and winning conditions on full (finite or infinite) plays.*

Positions: *Game positions are of two types: sequents, which belong to* ***Ver****, and pairs of sequents, which belong to* ***Fal****. We sometimes refer to positions belonging to* ***Ver*** *as "basic positions".*

Moves for Fal: *Given a position belonging to* ***Fal****, consisting of a pair of sequents, the player simply chooses one of the sequents from the pair.*

Moves for Ver: *Given a position belonging to* ***Ver****, consisting of a sequent Γ, if Γ is an axiom then the game ends and* ***Ver*** *is declared the winner. Otherwise available moves are defined as follows:*

- *If the deterministic rule is applicable to Γ then this is the only move allowed for* ***Ver****.*
- *If the deterministic rule is not applicable, but the ground rule is applicable, then this is the only move allowed for* ***Ver****.*
- *If neither the deterministic rule nor the ground rule are applicable, then the possible moves of* ***Ver*** *are the narrow modal rule, terminal applications of the weakening rule,* repeating *applications of narrow rules or repeating applications of Weakening.*

If π is a partial play ending with some sequent Γ, then we often refer to Γ as the label *of π. Note that since we allow repeating applications of Weakening,* ***Ver*** *never gets stuck. So the only full finite plays are those that end in an* *axiom, and are won by **Ver**. Thus to finish the construction of the **Inf**-game it remains only to decide the winner of an infinite play.* Traces *on a play of the **Inf**-game are defined similarly as traces in proof trees, the only difference being that a trace on a play π of length $k \leq \omega$ is now an object of the form $(\pi_n, \mathsf{i}_n : \varphi_n)_{n<k}$ where each π_n is an initial segment of the play π, and for each $n+1 < k$ the initial segment π_{n+1} extends π_n with a single move. Given an infinite play π, **Ver** is then declared the winner if the play contains a good trace, and otherwise the winner is **Fal**.*

We now draw some simple consequences of how the **Inf**-game has been designed.

Proposition 4.1 *In any sequent of the form $\Gamma, \mathsf{i} : \psi$ appearing in a play of the **Inf**-game, ψ contains no non-original nominals.*

Proof. Just observe that all the admissible moves preserve this condition. $\Box$

From this proposition a few useful facts follow:

Proposition 4.2 *If a play of the **Inf**-game contains any sequent of the form $\Gamma, \mathsf{i} \not\approx \mathsf{j}$, then j is an original nominal.*

Proof. Special case of Proposition 4.1. $\Box$

Proposition 4.3 *For each nominal i, and each partial play π in the **Inf**-game, the label of π contains at most k formulas of the form $\mathsf{i} : \psi$, where k is linear in the size of the root formula.*

Proof. Easy using Proposition 4.1. . $\Box$

Proposition 4.4 *For any sequent Γ appearing in a play of the **Inf**-game, at most one non-original nominal appears in Γ.*

Proof. The only moves that can introduce new non-original nominals are applications of the narrow modal rule, and by design each instance of this rule erases all occurrences of non-original nominals other than the new nominal that was introduced. $\Box$

Like the games for satisfiability checking for the modal μ-calculus introduced in [16], the proof search game is determinate:

Proposition 4.5 *The **Inf**-game is determinate, i.e. at every position there is a player who has a winning strategy.*

A crucial part of proving completeness of **Inf** is to show that the standard "good trace" condition on valid proofs, in terms of traces going from the root up along a single branch, is not too strong. At first sight it may seem that we need to consider a more general condition, allowing traces to jump between different occurrences of the same nominal. In this subsection we prove a useful result that deals with this issue.

Definition 12 *Let S be a set of plays in the **Inf**-game. A* good trace loop *on S in the **Inf**-game is a sequence of partial traces $\langle \boldsymbol{t}_1 .... \boldsymbol{t}_n \rangle$ for which there exist $\pi_1, ..., \pi_n \in S$ such that:*

- *Each $\boldsymbol{t}_i$ is a partial trace on π_i,*
- *Each $\boldsymbol{t}_i$ starts and ends with ground formulas,*
- *Each $\boldsymbol{t}_{i+1}$ starts with the last formula of $\boldsymbol{t}_i$,*
- *The trace $\boldsymbol{t}_1$ starts with the last formula of $\boldsymbol{t}_n$, and*
- *At least one variable is unfolded on some trace $\boldsymbol{t}_i$ and the highest ranking such variable is a ν-variable.*

Note that in the following lemma, our focus is on analyzing winning strategies for **Fal** in the proof search game, rather than strategies for **Ver**. The explanation for this is as follows. As winning strategies for **Ver** correspond to proofs, we may think of strategies of **Fal** as providing *refutations*. The completeness proof for **Inf** will build counter-models from such refutations. Rather than building the counter-model from an arbitrary refutation of the root formula, we will start by showing that if such a refutation exists, then there is a refutation of some sequent containing the root formula, with certain properties that make it ideally suited for constructing a counter-model.

Lemma 4.6 *Suppose that **Fal** has a winning strategy in the **Inf**-game for $\mathsf{i}:\rho$. Then there exists a sequent Φ containing $\mathsf{i}:\rho$ and a winning strategy σ for **Fal** in the **Inf**-game with starting position Φ, such that the following conditions hold:*

(i) *For every sequent appearing in a σ-guided play, its ground formulas are the same as the ground formulas in Φ.*

(ii) *The set of σ-guided plays does not contain any good trace loops.*

Proof. We first prove the following claim:

Claim 4.7 *There exists some sequent Φ such that:*

- $\mathsf{r}:\rho \in \Phi$,
- ***Fal** has a winning strategy σ in the **Inf**-game at the starting position Φ,*
- *For every sequent Γ that appears in some σ-guided partial play, the ground formulas appearing in Γ are the same as the ground formulas in Φ.*

Proof of Claim Let τ be the winning strategy assumed to exist for **Fal**. First note that the ground formulas appearing in τ-guided partial plays in the **Inf**-game are increasing in the sense that, whenever Γ' appears later than Γ in a partial play, all ground formulas in Γ are also in Γ'. This is because the only admissible rule application that can remove a ground formula is a terminal application of the weakening rule, the premise of which is an axiom. Such applications of weakening never happen in any τ-guided partial play, since such a play would be a loss for **Fal**.

We construct a series of τ-guided partial plays $\pi_0, \pi_1, \pi_2 ...$, where each π_i is an initial segment of π_{i+1}. For each i we let G_i be the set of ground formulas appearing on the last position of π_i. We shall maintain the invariant that, for all proper initial segments π' of π_{i+1}, the ground formulas appearing in the last sequent of π' are contained in G_i. Let π_0 be the start position of the **Inf**-game. Suppose that π_i has been constructed. If there is no τ-guided partial play π' extending π_i in which the last sequent contains ground formulas not in G_i, then we are done: for all τ-guided partial plays extending this play, the ground formulas appearing in all sequents must be equal to G_i, and τ provides a winning strategy for **Fal** in the **Inf**-game for the label of π_i. If there is some τ-guided partial play π' extending π_i in which the last sequent contains ground formulas not in G_i, then just pick π_{i+1} to be its smallest initial segment for which this holds. This procedure must eventually terminate, since otherwise we get an infinite and strictly increasing sequence of sets of ground formulas $G_0 \subset G_1 \subset G_2 ...$, which is impossible since there are only finitely many possible ground formulas that can appear in any play. ◀

Now let Φ and σ be as in the previous claim. Given a σ-guided play π, let $\uparrow \pi$ be the set of partial plays π' such that $\pi \cdot \pi'$ is a σ-guided partial play. Our aim is to find a σ-guided play π such that $\uparrow \pi$ does not contain any good trace loops; we can then simply take the label of π to the sequent claimed to exist in the statement of the Proposition, and we obtain the required winning strategy for **Fal** by assigning the move $\sigma(\pi \cdot \pi')$ to a partial play π'.

Given a good trace loop $\langle \boldsymbol{t}_1, ..., \boldsymbol{t}_n \rangle$, let its *kind* be the set of triples:

$$\{(\mathsf{i}_1 : \varphi_1, X_1, \mathsf{j}_1 : \psi_1), ..., (\mathsf{i}_n : \varphi_n, X_n, \mathsf{j}_n : \psi_n)\}$$

such that for each $m \in \{1, ..., n\}$, the trace $\boldsymbol{t}_m$ begins with $\mathsf{i}_m : \varphi_m$, ends with $\mathsf{j}_m : \psi_m$, and the variables unfolded on $\boldsymbol{t}_m$ are precisely the members of the set X_m. Since each trace in a good trace loop begins and ends with a ground formula, and since there are only finitely many ground formulas, there are finitely many kinds of good trace loops. We shall show how to find a σ-guided play π such that $\uparrow \pi$ does not contain any good trace loops of a given kind. By simply repeating the argument, we can then kill off all the kinds of good trace loop one by one.

So let the kind K be $\{(\mathsf{i}_0 : \varphi_0, X_0, \mathsf{j}_0 : \psi_0), ..., (\mathsf{i}_{n-1} : \varphi_{n-1}, X_{n-1}, \mathsf{j}_{n-1} : \psi_{n-1})\}$. We construct a sequence of partial plays $\pi_0, \pi_1, \pi_2 ...$, where each π_i is an initial segment of π_{i+1}, as follows. If the set of all σ-guided plays does not contain any good trace loops of kind K, we are done. Otherwise, let π_0 be some play on which the part $(\mathsf{i}_0 : \varphi_0, X_0, \mathsf{j}_0 : \psi_0)$ appears, which must exist. Note that we have a partial trace $\boldsymbol{t}_0$ on π_0 leading from $\mathsf{i}_0 : \varphi_0$ to $\mathsf{j}_0 : \psi_0$ on which exactly the variables X_0 were unfolded; since the first formula is a ground formula, and these are the same in all positions in all σ-guided plays, we can simply "pad" the partial trace with silent steps repeating the same formula to extend it to a trace on the whole play π_0. Now we repeat the procedure: if $\uparrow \pi_0$ does not contain any good trace loop, then we are done. Otherwise, we can extend π_0 in the same way to a partial play π_1 containing a trace $\boldsymbol{t}_1$ appearing *after* π_0, such that $\boldsymbol{t}_1$ starts with $\mathsf{i}_1 : \varphi_1$, ends with $\mathsf{j}_1 : \psi_1$ and the variables unfolded are precisely X_1. Then since $\boldsymbol{t}_0$ and $\boldsymbol{t}_1$ end and start respectively with the same ground formula, and ground formulas stay the same, they can be linked together by "padding with silent steps" repeating this formula to form a trace on π_1. It is not hard to see that, if this procedure never terminates, then we end up building an infinite σ-guided play containing a good trace, which is a contradiction since σ was a winning strategy. So the procedure terminates with some π_m, and the proof is finished. $\square$

Note that, since the set of *finite* partial plays in the **Inf**-game is a countable set (being a set of finite sequences over a countable set), we can define a surjective mapping F from the set of nominals to the set of finite σ-guided partial plays, such that $F^{-1}[\pi]$ is infinite for each finite partial play π. We leave the little set theoretic exercise of proving this to the reader. Throughout the rest of this section we fix such a map F. Informally, we think of $F(\mathsf{i})$ as a *tag* attached to the nominal i to remember where it was introduced.

Definition 13 *We say that a full or partial play π of the **Inf**-game is* clean *if, for every initial segment π' of the play ending with the conclusion of an application of the (narrow) modal rule introducing a new nominal j, we have $F(\mathsf{j}) = \pi'$.*

When proving completeness of **Inf** we shall construct a counter-model to the root formula from a winning strategy for **Fal**, and it will then be convenient to restrict attention to clean plays. We are now ready for the main result about the system **Inf**.

Theorem 1 *Let ρ be any formula. The following are equivalent: (a) ρ is valid, (b) **Ver** has a winning strategy in the **Inf**-game for $\mathsf{r}:\rho$, where r is some fresh nominal, (c) ρ has a valid and frugal **Inf**-proof, (d) ρ has a valid **Inf**-proof.*

Proof. We sketch the proof of the implication (a) $\Rightarrow$ (b), which is the most difficult part of the proof. We prove this by contraposition: suppose there is a winning strategy for **Fal** in the **Inf**-game for $\mathsf{r}:\rho$. Let Φ be a set of ground formulas containing $\mathsf{r}:\rho$ and let σ be a winning strategy for **Fal** in the **Inf**-game for premise Φ such that the ground formulas stay the same in every σ-guided play, and the set of σ-guided plays contains no good trace loops. Such Φ and σ exist by Lemma 4.6. We shall construct a countermodel to (the disjunction of) Φ, which gives a countermodel to ρ since $\mathsf{r}:\rho \in \Phi$.

We construct the model $M = (W, R, A, V)$ using the strategy σ as follows. Let N be the set of nominals i such that i appears in some position in some clean σ-guided play π, and let $\equiv$ be the smallest equivalence relation over N containing all pairs (i, j) for which $\mathsf{i} \not\approx \mathsf{j}$ appears in some position in some clean σ-guided play π.

Claim 4.8 *For each i, the equivalence class $[\mathsf{i}]$ modulo $\equiv$ is either a singleton or contains at least one of the nominals in φ.*

Motivated by this claim, we call a nominal *representative* if its equivalence class is a singleton, or it is the $\prec$-smallest original nominal belonging to its equivalence class. We let W be the set of representative members of N. We set $\mathsf{i}R\mathsf{j}$ iff there is some $\mathsf{j}' \equiv \mathsf{j}$ and a clean σ-guided play in which j' is introduced by an application of the modal rule to the nominal i. Set $A(\mathsf{i})$ to be the representative of $[\mathsf{i}]$. Finally, for a representative i set $\mathsf{i} \in V(p)$ iff $\mathsf{i}:\neg p$ appears on some clean σ-guided play. We shall show that M is a counter-model to the sequent Φ. The key claims used to prove this are the following:

Claim 4.9 *Suppose that* $\mathsf{i} \equiv \mathsf{j}$ *and that* j *is an original nominal. Then for any basic position* u *appearing in a clean* σ*-guided play, and any* θ*, if* $\mathsf{i}:\theta$ *belongs to* u *then so does* $\mathsf{j}:\theta$*.*

Claim 4.10 *Suppose that* $\boldsymbol{t}$ *is some partial trace on a clean partial* σ*-guided play* π*, such that the last element of the trace* $\boldsymbol{t}$ *is of the form* $(\pi, \mathsf{k}':\psi)$ *where* $A(\mathsf{k}') = \mathsf{k}$*. If* ψ *is of the form* $\Box\theta$ *or* $\Diamond\theta$*, then there is a clean* σ*-guided play* υ *extending* π *and a partial trace on* υ *of the form* $(\pi, \mathsf{k}':\psi) \cdot \boldsymbol{u} \cdot (\upsilon, \mathsf{k}:\psi)$*, which contains no fixpoint unfoldings.*

We now proceed to show that the sequent Φ is not valid in M. Pick any formula $\mathsf{i} : \varphi \in \Phi$. We shall construct a winning strategy σ' for **Fal** in the evaluation game in M at the starting position $(A(\mathsf{i}), \varphi)$. Inductively, as an invariant we associate with each partial σ'-guided partial play of π' of the form:

$$(\mathsf{j}_1, \psi_1) \cdot \boldsymbol{p} \cdot (\mathsf{j}_n, \psi_n)$$

a sequence of non-empty partial traces $\langle \boldsymbol{t}_1, ..., \boldsymbol{t}_n \rangle$ such that each of these traces $\boldsymbol{t}_k$ belongs to some clean σ-guided partial play π_k, and such that the following conditions hold:

I1: The last element of each trace $\boldsymbol{t}_k$ is of the form $(\pi_k, \mathsf{j}'_\mathsf{k} : \psi_k)$ where $A(\mathsf{j}'_k) = \mathsf{j}_k$. Furthermore, if ψ_k is of the form $\Box\theta$ or $\Diamond\theta$, then $\mathsf{j}'_k = \mathsf{j}_k$.

I2: For each $k < n$, if the last element of $\boldsymbol{t}_k$ is $(\pi_k, \mathsf{j}'_\mathsf{k} : \psi_k)$ then the first element of $\boldsymbol{t}_{k+1}$ is of the form $(\pi_{k+1}, \mathsf{j}'_\mathsf{k} : \psi_k)$. Furthermore, if j_k is not an original nominal then $\pi_k = \pi_{k+1}$.

I3: For $k < n$, a fixpoint unfolding occurs on the trace $\boldsymbol{t}_{k+1}$ iff the same fixpoint is unfolded on $(\mathsf{j}_k, \psi_k) \cdot (\mathsf{j}_{k+1}, \psi_{k+1})$.

Suppose we are given a clean σ'-guided partial play π' of the form $(\mathsf{j}_1, \psi_1) \cdot \boldsymbol{p} \cdot (\mathsf{j}_n, \psi_n)$, and that the associated sequence of partial traces $\langle \boldsymbol{t}_1, ..., \boldsymbol{t}_n \rangle$ has been constructed. Then we can show that, if the last position on π' belongs to **Fal**, then we can define a move for which the invariant (I1) – (I3) can be maintained, and if the last position belongs to **Ver** then the invariant can be maintained for any possible move. This is proved by a case distinction as to the shape of the last position, and uses the claims 4.9, and 4.10. The details of the argument are omitted here.

To finish the proof of (a) $\Rightarrow$ (b), we have given a strategy σ' to **Fal** in the evaluation game such that the invariant (I1) – (I3) is maintained. The strategy σ' ensures that **Fal** never gets stuck, and any lost infinite σ'-guided play produces either an infinite clean σ-guided shadow-play in the **Inf**-game containing a good infinite trace, or a good trace loop on the set of all clean σ-guided plays. In either case we get a contradiction, so σ' is winning for **Fal** and therefore we have found a falsifying model for ρ. $\Box$

5 Finite proofs with names

5.1 The system Saf

In this section we introduce the finitary proof system **Saf**, which is an annotated circular proof system in Stirling's style [23]. We will borrow a more streamlined version of the rules for manipulating annotations from Afshari and Leigh [2]. For each fixpoint variable x we assume that we have a countably infinite supply $\mathsf{x}_0, \mathsf{x}_1, \mathsf{x}_2...$ of *names* for that variable. We assume that we have a fixed enumeration of the set of variable names for each variable x, so that we can speak of the n-th variable name for x. The system will be defined taking a strict linear order $<$ over fixpoint variables as a parameter, and in the presentation we assume such an order as given. Given $<$ we write $\mathsf{x} < \mathsf{y}$ for names x, y of variables x, y respectively if $x < y$. Given a word a over the set of variable names and a variable x, we write $\mathsf{a} \leq x$ if there is no variable $y > x$ for which a contains a name of y. Given two words a, b over the set of variable names we write $\mathsf{a} \sqsubseteq \mathsf{b}$ to say that b contains a as a subsequence. For example $\mathsf{xy} \sqsubseteq \mathsf{xzy}$. We write $\mathsf{a} \sqcap \mathsf{b}$ for the longest word c such that $\mathsf{c} \sqsubseteq \mathsf{a}$ and $\mathsf{c} \sqsubseteq \mathsf{b}$ (provided that a longest word with this property exists, otherwise $\mathsf{a} \sqcap \mathsf{b}$ is undefined).

Definition 14 Annotated sequents *will be structures of the form:*

$$\mathsf{a} \vdash \mathsf{i}_1 : \varphi_1^{\mathsf{b}_1}, ..., \mathsf{i}_n : \varphi_n^{\mathsf{b}_n}$$

where $\mathsf{a}, \mathsf{b}_1, ..., \mathsf{b}_n$ are non-repeating words over the set of variable names (i.e. no variable name appears twice in any of these words), each b_i is non-decreasing with respect to the order $<$, and $\mathsf{b}_i \sqsubseteq \mathsf{a}$ for each $i \in \{1, ..., n\}$. The tuple a is called the control *of the sequent.*

A formula ρ will be said to be provable in the system if the sequent $\varepsilon \vdash \mathsf{r} : \rho^{\varepsilon}$ is provable, where the order $<$ on variable names is some arbitrary linearization of $<_\rho$, ε is the empty word, and r is a fresh nominal. We will allow suppressing occurrences of the empty word in our notation, including the control, so that for example the sequent $\varepsilon \vdash \mathsf{r} : \rho^{\varepsilon}$ can be written simply as $\mathsf{r} : \rho$.

Definition 15 *A sequent in the sense of the system **Inf** will be called a* plain sequent. *Given an annotated sequent $\Gamma = \mathsf{a} \vdash \mathsf{i}_1 : \varphi_1^{\mathsf{b}_1}, ..., \mathsf{i}_n : \varphi_n^{\mathsf{b}_n}$, the* underlying plain sequent $\underline{\Gamma}$ *is the plain sequent $\mathsf{i}_1 : \varphi_1, ..., \mathsf{i}_n : \varphi_n$.*

The system has two axioms, which are the law of exluded middle and an identity axiom, which now have the form:

$$\varepsilon \vdash \mathsf{i} : p^{\varepsilon}, \mathsf{i} : \neg p^{\varepsilon} \qquad \varepsilon \vdash \mathsf{i} \approx \mathsf{i}^{\varepsilon}$$

Here, p is a nominal or a propositional variable. Rules of inference are given in Figure 2. Note that Γ, Ψ here denote sets of annotated formulas rather than plain formulas. The rules are subject to the following constraints:

Mod: The nominal j must be fresh.

ηx**:** $\mathsf{b} \leq x$.

Rec(x): $\mathsf{b} \leq x$, and x is a fresh variable name for x.

Exp**:** $\mathsf{a} \sqsubseteq \mathsf{a}'$, $\mathsf{b}_i \sqsubseteq \mathsf{b}'_i$ and $\mathsf{b}'_i \sqcap \mathsf{a} \sqsubseteq \mathsf{b}_i$ for each i [2].

$\mathsf{Reset(x)}$**:** The variable x does not appear in any formula in Γ.

$$\frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : \varphi \wedge \psi^{\mathsf{b}}, \mathsf{i} : \varphi^{\mathsf{b}} \qquad \mathsf{a} \vdash \Gamma, \mathsf{i} : \varphi \wedge \psi^{\mathsf{b}}, \mathsf{i} : \psi^{\mathsf{b}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} : \varphi \wedge \psi^{\mathsf{b}}} \wedge$$

$$\frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : \varphi \vee \psi^{\mathsf{b}}, \mathsf{i} : \varphi^{\mathsf{b}}, \mathsf{i} : \psi^{\mathsf{b}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} : \varphi \vee \psi^{\mathsf{b}}} \vee \qquad \frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : \phi^{\mathsf{b}}, \mathsf{j} : \phi^{\mathsf{b}}, \mathsf{i} \not\approx \mathsf{j}^{\mathsf{c}}}{\mathsf{a} \vdash \Gamma, \mathsf{j} : \phi^{\mathsf{b}}, \mathsf{i} \not\approx \mathsf{j}^{\mathsf{c}}} \,\mathsf{Eq}$$

$$\frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : \Box\varphi^{\mathsf{b}}, \mathsf{i} : \Diamond\Psi, \mathsf{j} : \varphi^{\mathsf{b}}, \mathsf{j} : \Psi}{\mathsf{a} \vdash \Gamma, \mathsf{i} : \Box\varphi^{\mathsf{b}}, \mathsf{i} : \Diamond\Psi} \,\mathsf{Mod} \qquad \frac{\mathsf{a} \vdash \Gamma, \mathsf{i} \not\approx \mathsf{j}^{\mathsf{b}}, \mathsf{j} \not\approx \mathsf{i}^{\mathsf{b}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} \not\approx \mathsf{j}^{\mathsf{b}}} \,\mathsf{Com}$$

$$\frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : \eta x.\phi(x)^{\mathsf{b}}, \mathsf{i} : \phi(\eta x.\phi(x))^{\mathsf{b}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} : \eta x.\phi(x)^{\mathsf{b}}} \,\eta x \qquad \frac{\mathsf{a} \vdash \Gamma, \mathsf{i} : (\mathsf{j} : \varphi)^{\mathsf{b}}, \mathsf{j} : \varphi^{\mathsf{b}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} : (\mathsf{j} : \varphi)^{\mathsf{b}}} \,\mathsf{Glob}$$

$$\frac{\mathsf{ax} \vdash \Gamma, \mathsf{i} : \nu x.\varphi(x)^{\mathsf{b}}, \mathsf{i} : \varphi(\nu x.\varphi(x))^{\mathsf{bx}}}{\mathsf{a} \vdash \Gamma, \mathsf{i} : \nu x.\varphi(x)^{\mathsf{b}}} \,\mathsf{Rec(x)} \qquad \frac{\mathsf{a} \vdash \Gamma}{\mathsf{a} \vdash \Gamma \cup \Psi} \,\mathsf{Weak}$$

$$\frac{\mathsf{a} \vdash \Gamma, \mathsf{i}_1 : \varphi_1^{\mathsf{bx}}, ..., \mathsf{i}_\mathsf{n} : \varphi_n^{\mathsf{bx}}}{\mathsf{a} \vdash \Gamma, \mathsf{i}_1 : \varphi_1^{\mathsf{bxx_1c_1}}, ..., \mathsf{i}_\mathsf{n} : \varphi_n^{\mathsf{bxx_nc_n}}} \,\mathsf{Reset(x)} \qquad \frac{\mathsf{a} \vdash \mathsf{i}_1 : \varphi_1^{\mathsf{b_1}}, ..., \mathsf{i}_n : \varphi_n^{\mathsf{b_n}}}{\mathsf{a}' \vdash \mathsf{i}_1 : \varphi_1^{\mathsf{b}'_1}, ..., \mathsf{i}_n : \varphi_n^{\mathsf{b}'_n}} \,\mathsf{Exp}$$

Fig. 2. Rules of **Saf**

A **Saf**-proof is a labelled tree where the labels specify a sequent assigned to a node and the last rule application (for non-leaf nodes), and such that the children of a node are labelled with the premises of the specified rule application. Although valid proofs will always be finite it will be useful to consider infinite **Saf**-proofs as well. We say that the variable name x is *reset* in an instance of the rule $\mathsf{Reset(x)}$.

Definition 16 *A **Saf**-proof will be considered* valid *if it is a finite proof-tree, and there is a map f from non-axiom leaves to non-leaves, such that:*

[2] Note that $\mathsf{b}'_i \sqcap \mathsf{a}$ is well-defined here: since $\mathsf{b}'_i \sqsubseteq \mathsf{a}'$ and $\mathsf{a} \sqsubseteq \mathsf{a}'$, and since a' is non-repeating, any two variable names occurring in both b'_i and a must appear only once and in the same order in both words. From this follows that the set of words c such that $\mathsf{c} \sqsubseteq \mathsf{b}'_i$ and $\mathsf{c} \sqsubseteq \mathsf{a}$ is a $\sqsubseteq$-directed finite set, so it contains a $\sqsubseteq$-maximal word.

- *$f(l)$ is an ancestor of l, and has the same label.*
- *There is a variable name x that is contained in the control of every node in the path from $f(l)$ to l, and is reset at least once on this path.*

A map f from non-axiom leaves to non-leaves satisfying the first of these conditions is called a back-edge map*, and is* good *if it satisfies the second condition as well. So a finite proof-tree is a valid proof iff it has a good back-edge map.*

We can now state the main result of the paper.

Theorem 2 (Completeness of Saf) *A formula $\mathsf{i}:\varphi$ has a valid **Saf**-proof if and only if it is semantically valid.*

Proof. We only sketch the proof here. For the soundness part, it is a fairly simple exercize to "unfold" a valid **Saf**-proof to an infinite proof-tree in which every infinite branch has a good trace. By forgetting the annotations we can view this as a valid **Inf**-proof, and soundness thus follows from Theorem 1.

For the completeness proof, we reason as follows: first, any valid formula $\mathsf{i}:\varphi$ has a valid and frugal **Inf**-proof Π by Theorem 1. We need to add annotations to the sequents in this proof, possibly inserting some rules of **Saf** for updating annotations, in such a way that we produce an infinite **Saf**-proof for the same conclusion which is still frugal, i.e. contains only finitely many annotated sequents, and satisfies the constraint that on every infinite branch there is some variable name that is reset infinitely many times. The construction of this infinite **Saf**-proof essentially uses annotations to mimick the Safra construction for automata on infinite words, and follows the same reasoning as in [12]. Since the class of proof trees satisfying these criteria is definable in monadic second-order logic, we may apply Rabin's Basis Theorem [18] to find a *regular* infinite **Saf**-proof for the same conclusion, in which every infinite branch has a variable name that is reset infinitely often. Finally, we note that any such regular proof can be "folded back" into a finite proof-tree with a back-edge map that yields a valid finite **Saf**-proof. $\Box$

Example 1 *We show a valid **Saf**-proof of the formula $\mathsf{i}:(\Box\neg\mathsf{i}\vee\nu x.\Diamond x)$, which is equivalent to:*

$$\mathsf{i}:(\Diamond\mathsf{i}\to\nu x.\Diamond x)$$

The "†" labels show how the back-edge map connects the leaf to an ancestor.

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\mathsf{x}_0\vdash\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{i}:\Diamond\nu x.\Diamond x^{\mathsf{x}_0}\quad\dagger}{\mathsf{x}_0\mathsf{x}_1\vdash\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{i}:\Diamond\nu x.\Diamond x^{\mathsf{x}_0\mathsf{x}_1}}\ \mathsf{Reset}(\mathsf{x}_0)}{\mathsf{x}_0\vdash\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{i}:\nu x.\Diamond x^{\mathsf{x}_0}}\ \mathsf{Rec}(\mathsf{x}_1)+\mathsf{Weak}}{\mathsf{x}_0\vdash\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{j}\not\approx\mathsf{i},\ \mathsf{j}:\nu x.\Diamond x^{\mathsf{x}_0}}\ \mathsf{Eq}+\mathsf{Weak}}{\mathsf{x}_0\vdash\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{i}:\Diamond\nu x.\Diamond x^{\mathsf{x}_0}\quad\dagger}\ \mathsf{Mod}+\mathsf{Weak}}{\mathsf{i}:\Box\neg\mathsf{i},\ \mathsf{i}:\nu x.\Diamond x}\ \mathsf{Rec}(\mathsf{x}_0)+\mathsf{Weak}}{\mathsf{i}:(\Box\neg\mathsf{i}\vee\nu x.\Diamond x)}\ \vee+\mathsf{Weak}$$

6 Concluding remarks

We conclude with some directions for future work. First of all, with the Stirling-style proof system in place for the hybrid μ-calculus, we should be able to prove cut-free completeness of a sequent system for the hybrid μ-calculus by following the same method of translation between proof systems as in [2]. The proof should not involve any substantial novelties, although the details remain to be checked.

We hope that the methods developed here can be extended to other extended μ-calculi, like guarded fixpoint logic. A first task in this direction is to consider converse modalities, and obtain a Stirling-style circular system for the hybrid μ-calculus including converse modalities. In Vardi's automata-theoretic decision procedure for the two-way μ-calculus, the key component is a finite data structure for encoding generalized traces that can go upwards or downwards along branches in a tableau. This extra component is then removed through a projection operation on automata recognizing valid tableaux. It would be interesting to investigate this construction from a proof-theoretic perspective.

References

[1] Afshari, B., G. Jäger and G. E. Leigh, *An infinitary treatment of full mu-calculus*, in: *International Workshop on Logic, Language, Information, and Computation*, Springer, 2019, pp. 17–34.

[2] Afshari, B. and G. E. Leigh, *Cut-free completeness for modal mu-calculus*, in: *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, IEEE, 2017, pp. 1–12.

[3] Brotherston, J. and A. Simpson, *Complete sequent calculi for induction and infinite descent*, in: *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*, IEEE, 2007, pp. 51–62.

[4] Brotherston, J. and A. Simpson, *Sequent calculi for induction and infinite descent*, Journal of Logic and Computation **21** (2010), pp. 1177–1216.

[5] Bruse, F., O. Friedmann and M. Lange, *On guarded transformation in the modal μ-calculus*, Logic Journal of the IGPL **23** (2015), pp. 194–216.

[6] Dam, M. and D. Gurov, *μ-calculus with explicit points and approximations*, Journal of Logic and Computation **12** (2002), pp. 255–269.

[7] Emerson, E. and C. Jutla, *Tree automata, mu-calculus and determinacy (extended abstract)*, in: *Proceedings of the 32nd Symposium on the Foundations of Computer Science* (1991), pp. 368–377.

[8] Enqvist, S., H. H. Hansen, C. Kupke, Y. Venema and J. Marti, *Completeness for game logic*, in: *2019 34th Annual ACM/IEEE Symposium on Logic in computer Science (LICS)*, 2019.

[9] Enqvist, S., F. Seifan and Y. Venema, *Completeness for the modal μ-calculus: separating the combinatorics from the dynamics*, Theoretical Computer Science **727** (2018), pp. 37–100.

[10] Enqvist, S., F. Seifan and Y. Venema, *Completeness for μ-calculi: a coalgebraic approach*, Annals of Pure and Applied Logic **170** (2019), pp. 578–641.

[11] Grädel, E. and I. Walukiewicz, *Guarded fixed point logic*, in: *Proceedings 14th IEEE Symposium on Logic in Computer Science LICS'99*, 1999.

[12] Jungteerapanich, N., *Tableau systems for the modal μ-calculus*, The University of Edinburgh (2010).

[13] Kernberger, D. and M. Lange, *The fully hybrid mu-calculus*, in: *24th International Symposium on Temporal Representation and Reasoning (TIME 2017)*, Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[14] Knaster, B., *Un théorème sur les fonctions des ensembles*, Annales de la Societé Polonaise de Mathematique **6** (1928), pp. 133–134.

[15] Kozen, D., *Results on the propositional μ-calculus*, Theoretical Computer Science **27** (1983), pp. 333–354.

[16] Niwiński, D. and I. Walukiewicz, *Games for the μ-calculus*, Theoretical Computer Science **163** (1996), pp. 99–116.

[17] Parikh, R., *The logic of games and its applications*, Annals of Discrete Mathematics **24** (1985), pp. 111–139.

[18] Rabin, M., *Automata on infinite objects and church's problem*, American Mathematical Soc. **13** (1972).

[19] Santocanale, L., *A calculus of circular proofs and its categorical semantics*, in: *International Conference on Foundations of Software Science and Computation Structures*, Springer, 2002, pp. 357–371.

[20] Sattler, U. and M. Y. Vardi, *The hybrid μ-calculus*, in: *International Joint Conference on Automated Reasoning*, Springer, 2001, pp. 76–91.

[21] Shamkanov, D. S., *Circular proofs for the Gödel-Löb provability logic*, Mathematical Notes **96** (2014), pp. 575–585.

[22] Sprenger, C. and M. Dam, *On the structure of inductive reasoning: Circular and tree-shaped proofs in the μcalculus*, in: *International Conference on Foundations of Software Science and Computation Structures*, Springer, 2003, pp. 425–440.

[23] Stirling, C., *A tableau proof system with names for modal mu-calculus*, EPiC Series in Computing **42** (2014), pp. 306–318.

[24] Tamura, K., *A small model theorem for the hybrid μ-calculus*, Journal of Logic and Computation **25** (2013), pp. 405–441.

[25] Tarski, A., *A lattice-theoretical fixpoint theorem and its applications*, Pacific Journal of Mathematics **5** (1955), pp. 285–309.

[26] Vardi, M. Y., *Reasoning about the past with two-way automata*, in: *International Colloquium on Automata, Languages, and Programming*, Springer, 1998, pp. 628–641.

[27] Walukiewicz, I., *Completeness of Kozen's axiomatisation of the propositional mu-calculus*, Information and Computation **157** (2000), pp. 142 – 182.
URL `http://www.sciencedirect.com/science/article/pii/S0890540199928365`

[28] Zielonka, W., *Infinite games on finitely coloured graphs with applications to automata on infinite trees*, Theoretical Computer Science **200** (1998), pp. 135–183.