

Focus-style proofs for the two-way alternation-free μ -calculus^{*}

Jan Rooduijn^{**} and Yde Venema

ILLC, University of Amsterdam, The Netherlands

Abstract. We introduce a cyclic proof system for the two-way alternation-free modal μ -calculus. The system manipulates one-sided Gentzen sequents and locally deals with the backwards modalities by allowing analytic applications of the cut rule. The global effect of backwards modalities on traces is handled by making the semantics relative to a specific strategy of the opponent in the evaluation game. This allows us to augment sequents by so-called trace atoms, describing traces that the proponent can construct against the opponent's strategy. The idea for trace atoms comes from Vardi's reduction of alternating two-way automata to deterministic one-way automata. Using the multi-focus annotations introduced earlier by Marti and Venema, we turn this trace-based system into a path-based system. We prove that our system is sound for all sequents and complete for sequents not containing trace atoms.

Keywords: two-way modal μ -calculus · alternation-free · cyclic proof theory

1 Introduction

The modal μ -calculus, introduced in its present form by Kozen [10], is an extension of modal logic by least and greatest fixed point operators. It retains many of the desirable properties of modal logic, such as bisimulation invariance, and relatively low complexity of the model-checking and satisfiability problems. Nevertheless, the modal μ -calculus achieves a great gain in expressive power, as the fixed point operators can be used to capture a form of recursive reasoning. This is illustrated by the fact that the modal μ -calculus embeds many well-known extensions of modal logic, such as Common Knowledge Logic, Linear Temporal Logic and Propositional Dynamic Logic.

A natural further extension is to add a converse modality \check{a} for each modality a . The resulting logic, called *two-way modal μ -calculus*, can be viewed as being able to reason about the past. As such, it can interpret the past operator of Tense Logic, and moreover subsumes PDL with converse. In this paper we are concerned with the proof theory of the two-way modal μ -calculus.

^{*} A version of this paper including an appendix with full proofs can be found on arXiv.

^{**} The research of this author has been made possible by a grant from the Dutch Research Council NWO, project number 617.001.857.

Developing good proof systems for the modal μ -calculus is notoriously difficult. In [10], Kozen introduced a natural Hilbert-style axiomatisation, which was proven to be complete only more than a decade later by Walukiewicz [24]. Central to this proof is the use of tableau systems introduced by Niwiński and Walukiewicz in [17]. One perspective on these tableau systems is that they are cut-free Gentzen-style sequent systems allowing infinite branches. A proof in such a system, called a *non-well-founded proof*, is accepted whenever every infinite branch satisfies a certain progress condition. In case this progress condition is ω -regular (as it is in the case of the modal μ -calculus), automata-theoretic methods show that for every non-well-founded proof there is a *regular* proof, *i.e.* a proof tree containing only finitely many non-isomorphic subtrees. Since these kind of proofs can be naturally presented as finite trees with back edges, they are called *cyclic proofs*. As an alternative to non-well-founded proofs, one can use proof rules with infinitely many premisses. We will not take this route, but note that it has been applied to the two-way modal μ -calculus by Afshari, Jäger and Leigh in [2].

In [12] Lange and Stirling, for the logics LTL and CTL, annotate formulas in sequents with certain automata-theoretic information. This makes it possible to directly construct cyclic proof systems, without the detour through automata theory. This technique has been further developed by Jungteerapanich and Stirling [21, 7] for the modal μ -calculus. Moreover, certain fragments of the modal μ -calculus, such as the alternation-free fragment [14] and modal logic with the master modality [19] have received the same treatment. Encoding automata-theoretic information in cyclic proofs, through annotating formulas, makes them more amenable to proof-theoretic applications, such as the extraction of interpolants from proofs [13, 3].

The logic at hand, the two-way modal μ -calculus, poses additional difficulties. Already without fixed point operators, backwards modalities are known to require more expressivity than offered by a cut-free Gentzen system [18]. A common solution is to add more structure to sequents, as *e.g.* the nested sequents of Kashima [8]. This approach, however, does not combine well with cyclic proofs, as the number of possible sequents in a given proof becomes unbounded. We therefore opt for the alternative approach of still using ordinary sequents, but allowing analytic applications of the cut rule (see [6] for more on the history of this approach). The combination of analytic cuts and cyclic proofs has already been shown to work well in the case of Common Knowledge Logic [20]. Choosing analytic cuts over sequents with extended structure has recently also been gaining interest in the proof theory of logics without fixed point operators [4].

Although allowing analytic cuts handles the backwards modalities on a local level, further issues arise on a global level in the combination with non-well-founded branches. The main challenge is that the progress condition should not just hold on infinite branches, but also on paths that can be constructed by moving both up and down a proof tree. Our solution takes inspiration from Vardi's reduction of alternating two-way automata to deterministic one-way automata [22]. Roughly, the idea is to view these paths simply as upwards paths,

only interrupted by several detours, each returning to the same state as where it departed. One of the main insights of the present research is that such detours have a natural interpretation in terms of the game semantics of the modal μ -calculus. We exploit this by extending the syntax with so-called *trace atoms*, whose semantics corresponds with this interpretation. Our sequents will then be one-sided Gentzen sequents containing annotated formulas, trace atoms, and negations of trace atoms.

For the sake of simplicity we will restrict ourselves to the alternation-free fragment of the modal μ -calculus. This roughly means that we will allow no entanglement of least and greatest fixed point operators. In this setting it suffices to annotate formulas with just a single bit of information, distinguishing whether the formula is *in focus* [14]. This is a great simplification compared to the full language, where annotations need to be strings and a further global annotation, called the *control*, is often used [21, 7]. Despite admitting simple annotations, the trace structure of the alternation-free modal μ -calculus remains intricate. This is mainly caused by the fact that disjunctions may still appear in the scope of greatest fixed point operators, causing traces to split.

While this paper was under review, the preprint [1] by Enqvist et al. appeared, in which a proof system is presented for the two-way modal μ -calculus (with alternation). Like our system, their system is cyclic. Moreover, they also extend the syntax in order to apply the techniques from Vardi in a proof-theoretical setting. However, their extension, which uses so-called *ordinal variables*, is substantially different from ours, which uses trace atoms. It would be interesting to see whether the two approaches are intertranslatable.

In Section 2 we define the two-way alternation-free modal μ -calculus. Section 3 is devoted to introducing the proof system, after which in Section 4 we show that proofs correspond to winning strategies in a certain parity game. In Section 5 we prove soundness and completeness. The concluding Section 6 contains a short summary and some ideas for further research.

2 The (alternation-free) two-way modal μ -calculus

For the rest of this paper we fix the countably infinite sets P of *propositional variables* and D of *actions*. Since we want our modal logic to be *two-way*, we define an involution operation $\check{\cdot} : D \rightarrow D$ such that for every $a \in D$ it holds that $\check{\check{a}} \neq a$ and $\check{\check{\check{a}}} = a$. We work in negation normal form, where the language $\mathcal{L}_{2\mu}$ of the *two-way modal μ -calculus* is generated by the following grammar:

$$\varphi ::= p \mid \bar{p} \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \langle a \rangle \varphi \mid [a] \varphi \mid \mu x \varphi \mid \nu x \varphi$$

where $p, x \in P$, $a \in D$ and in the formation of $\eta x \varphi$ ($\eta \in \{\mu, \nu\}$) the formula \bar{x} does not occur in φ . The language $\mathcal{L}_{2\mu}$ expresses \top and \perp , e.g. as $\nu x.x$ and $\mu x.x$. For the reader familiar with the ordinary modal μ -calculus, note that the only distinctive feature of $\mathcal{L}_{2\mu}$ is the assumed involution operator on D .

We use standard terminology for the binding of variables by a fixpoint operator η . In particular, we write $FV(\varphi)$ for the set of variables $x \in P$ that occur

freely in φ and $BV(\varphi)$ for the set of those that are bound by some fixpoint operator. Note that for every \bar{x} occurring in φ , we have $x \in FV(\varphi)$. For technical convenience, we assume that each formula φ is *tidy*, i.e. that $FV(\varphi) \cap BV(\varphi) = \emptyset$. The *unfolding* of a formula $\psi = \eta x \varphi$ is the formula $\varphi[\psi/x]$, obtained by substituting every free occurrence of x in φ by ψ . No free variables of ψ are captured by this procedure, because $FV(\psi) \cap BV(\varphi) \subseteq FV(\varphi) \cap BV(\varphi) = \emptyset$. The *closure* of a formula $\xi \in \mathcal{L}_{2\mu}$ is the least set $\text{Clos}(\xi) \subseteq \mathcal{L}_{2\mu}$ such that $\xi \in \text{Clos}(\xi)$ and:

- (i) $\varphi \circ \psi \in \text{Clos}(\xi)$ implies $\varphi, \psi \in \text{Clos}(\xi)$ for each $\circ \in \{\vee, \wedge\}$;
- (ii) $\Delta \varphi \in \text{Clos}(\xi)$ implies $\varphi \in \text{Clos}(\xi)$ for every $\Delta \in \{\langle a \rangle, [a] \mid a \in \mathbf{D}\}$;
- (iii) $\eta x \varphi \in \text{Clos}(\xi)$ implies $\varphi[\eta x \varphi/x] \in \text{Clos}(\xi)$ for every $\eta \in \{\mu, \nu\}$.

It is well known that $\text{Clos}(\xi)$ is always finite and that all formulas in $\text{Clos}(\xi)$ are tidy if ξ is so (see e.g. [23]).

Formulas of $\mathcal{L}_{2\mu}$ are interpreted in *Kripke models* $\mathbb{S} = (S, (R_a)_{a \in \mathbf{D}}, V)$, where S is a set of *states*, for each $a \in \mathbf{D}$ we have an *accessibility relation* $R_a \subseteq S \times S$, and $V : \mathbf{P} \rightarrow \mathcal{P}(S)$ is a *valuation function*. We assume that each model is *regular*, i.e. that R_a is the converse relation of $R_{\bar{a}}$ for every $a \in \mathbf{D}$. Recall that the converse relation of a relation R consists of those (y, x) such that $(x, y) \in R$.

We set $R_a[s] := \{t \in S : s R_a t\}$ and let $\mathbb{S}[x \mapsto X]$ be the model obtained from \mathbb{S} by replacing the valuation function V by $V[x \mapsto X]$, defined by setting $V[x \mapsto X](x) = X$ and $V[x \mapsto X](p) = V(p)$ for every $p \neq x$. The *meaning* $\llbracket \varphi \rrbracket^{\mathbb{S}} \subseteq S$ of a formula $\xi \in \mathcal{L}_{2\mu}$ in \mathbb{S} is inductively on the complexity of ξ :

$$\begin{aligned} \llbracket p \rrbracket^{\mathbb{S}} &:= V(p) & \llbracket \bar{p} \rrbracket^{\mathbb{S}} &:= S \setminus V(p) \\ \llbracket \varphi \vee \psi \rrbracket^{\mathbb{S}} &:= \llbracket \varphi \rrbracket^{\mathbb{S}} \cup \llbracket \psi \rrbracket^{\mathbb{S}} & \llbracket \varphi \wedge \psi \rrbracket^{\mathbb{S}} &:= \llbracket \varphi \rrbracket^{\mathbb{S}} \cap \llbracket \psi \rrbracket^{\mathbb{S}} \\ \llbracket \langle a \rangle \varphi \rrbracket^{\mathbb{S}} &:= \{s \in S \mid R_a[s] \cap \llbracket \varphi \rrbracket^{\mathbb{S}} \neq \emptyset\} & \llbracket [a] \varphi \rrbracket^{\mathbb{S}} &:= \{s \in S \mid R_a[s] \subseteq \llbracket \varphi \rrbracket^{\mathbb{S}}\} \\ \llbracket \mu x \varphi \rrbracket^{\mathbb{S}} &:= \bigcap \{X \subseteq S \mid \llbracket \varphi \rrbracket^{\mathbb{S}[x \mapsto X]} \subseteq X\} & \llbracket \nu x \varphi \rrbracket^{\mathbb{S}} &:= \bigcup \{X \subseteq S \mid X \subseteq \llbracket \varphi \rrbracket^{\mathbb{S}[x \mapsto X]}\} \end{aligned}$$

We will use the definable (see [23]) negation operator $\bar{\cdot}$ on $\mathcal{L}_{2\mu}$, for which it holds that $\llbracket \bar{\xi} \rrbracket^{\mathbb{S}} = S \setminus \llbracket \xi \rrbracket^{\mathbb{S}}$.

In this paper we shall only work with an alternative, equivalent, definition of the semantics, given by the *evaluation game* $\mathcal{E}(\xi, \mathbb{S})$. We refer the reader to the appendix below for the basic notions of (parity) games. The game $\mathcal{E}(\xi, \mathbb{S})$ is played on the board $\text{Clos}(\xi) \times S$, and its ownership function and admissible moves are given in the following table.

Position	Owner	Admissible moves
$(p, s), s \in V(p)$	\forall	\emptyset
$(p, s), s \notin V(p)$	\exists	\emptyset
$(\varphi \vee \psi, s)$	\exists	$\{(\varphi, s), (\psi, s)\}$
$(\varphi \wedge \psi, s)$	\forall	$\{(\varphi, s), (\psi, s)\}$
$(\langle a \rangle \varphi, s)$	\exists	$\{\varphi\} \times R_a[s]$
$([a] \varphi, s)$	\forall	$\{\varphi\} \times R_a[s]$
$(\eta x \varphi, s)$	$-$	$\{(\varphi[\eta x \varphi/x], s)\}$

The following proposition is standard in the literature on the modal μ -calculus. See [11, Proposition 6.7] for a proof.

Proposition 1. *For every infinite $\mathcal{E}(\xi, \mathbb{S})$ -match $\mathcal{M} = (\varphi_n, s_n)_{n \in \omega}$, there is a unique fixpoint formula $\eta x \chi$ which occurs infinitely often in \mathcal{M} and is a subformula of φ_n for cofinitely many n .*

The winner of an infinite match $\mathcal{E}(\xi, \mathbb{S})$ -match is \exists if in the previous proposition $\eta = \nu$, and \forall if $\eta = \mu$. It is well known that $\mathcal{E}(\xi, \mathbb{S})$ can be realised as a parity game by defining a suitable priority function on $\text{Clos}(\xi) \times S$ (we again refer the reader to [11] for a detailed proof of this fact). Because of this we may, by Theorem 1 in Appendix A, assume that winning strategies are optimal and positional. Finally, we state the known fact that the two approaches provide the same meaning to formulas. For every $\varphi \in \text{Clos}(\xi)$: $(\varphi, s) \in \text{Win}_{\exists}(\mathcal{E}(\xi, \mathbb{S}))@(\varphi, s)$ if and only if $s \in \llbracket \varphi \rrbracket^{\mathbb{S}}$. If either side of the bi-implication holds, we say that φ is satisfied in \mathbb{S} at s and write $\mathbb{S}, s \Vdash \varphi$.

In this paper we are concerned with a fragment of $\mathcal{L}_{2\mu}$ containing only those formulas ξ which are *alternation free*, i.e. such that for every subformula $\eta x \varphi$ of ξ it holds that no free occurrence of x in φ is in the scope of an $\bar{\eta}$ -operator in φ (where $\bar{\eta}$ denotes the opposite fixed point operator of η). This fragment is called *the alternation-free two-way modal μ -calculus* and denoted by $\mathcal{L}_{2\mu}^{\text{af}}$. We close this section by stating some typical properties of the alternation-free fragment. For $\eta \in \{\mu, \nu\}$ we use the term η -formula for a formula of the form $\eta x \varphi$.

Proposition 2. *Let $\xi \in \mathcal{L}_{2\mu}^{\text{af}}$ be an alternation-free formula. Then:*

- *Every formula $\varphi \in \text{Clos}(\xi)$ is alternation free.*
- *The negation $\bar{\xi}$ is alternation free.*
- *An infinite $\mathcal{E}(\xi, \mathbb{S})$ -match is won by \exists precisely if it contains infinitely many ν -formulas, and by \forall precisely if it contains infinitely many μ -formulas.*

3 The proof system

We will call a set Σ of formulas *negation-closed* if for every $\xi \in \Sigma$ it holds that $\bar{\xi} \in \Sigma$ and $\text{Clos}(\xi) \subseteq \Sigma$. For the remainder of this paper we fix a finite and negation-closed set Σ of $\mathcal{L}_{2\mu}^{\text{af}}$ -formulas. For reasons of technical convenience, we will assume that every formula is drawn from Σ . This does not restrict the scope of our results, as any formula is contained in some finite negation-closed set.

3.1 Sequents

Syntax Inspired by [14], we annotate formulas by a single bit of information.

Definition 1. *An annotated formula is a formula with an annotation in $\{\circ, \bullet\}$.*

The letters b, c, d, \dots are used as variables ranging over the annotations \circ and \bullet . An annotated formula φ^b is said to be *out of focus* if $b = \circ$, and *in focus* if $b = \bullet$. The focus annotations will keep track of so-called *traces* on paths through proofs. Roughly, a trace on a path is a sequence of formulas, such that the i -th formula occurs in the i -th sequent on the path, and the $i + 1$ -th formula ‘comes from’ the

i -th formula in a way which we will define later. In Section 4 we will construct a game in which the winning strategies of one player correspond precisely to the proofs in our proof system. The focus mechanism enables us to formulate this game as a parity game. This is essentially also the approach taken in [14].

Where traces usually only moves upwards in a proof, the backwards modalities of our language will be enable them to go downwards as well. We will handle this in our proof system by further enriching our sequents with the following additional information.

Definition 2. *For any two formulas φ, ψ , there is a trace atom $\varphi \rightsquigarrow \psi$ and a negated trace atom $\varphi \not\rightsquigarrow \psi$.*

The idea for trace atoms will become more clear later, but for now one can think of $\varphi \rightsquigarrow \psi$ as expressing that there is some kind of trace going from φ to ψ , and of $\varphi \not\rightsquigarrow \psi$ as its negation. Finally, our sequents are built from the above three entities.

Definition 3. *A sequent is a finite set consisting of annotated formulas, trace atoms, and negated trace atoms.*

Whenever we want to refer to general elements of a sequent Γ , without specifying whether we mean annotated formulas or (negated) trace atoms, we will use the capital letters A, B, C, \dots

Semantics We will now define the semantics of sequents. Unlike annotations, which do not affect the semantics but only serve as bookkeeping devices, the trace atoms have a well-defined interpretation. We will work with a refinement of the usual satisfaction relation that is defined with respect to a strategy for \forall in the evaluation game. Most of the time, this strategy will be both *optimal* and *positional* (see Appendix A for the precise definition of these terms). Because we will frequently need to mention such optimal positional strategies, we will refer to them by the abbreviation *ops*. We first define the interpretation of annotated formulas. Note that the focus annotations play no role in this definition.

Definition 4. *Let \mathbb{S} be a model, let f be an ops for \forall in $\mathcal{E}(\wedge \Sigma, \mathbb{S})$ and let φ^b be an annotated formula. We write $\mathbb{S}, s \Vdash_f \varphi^b$ if f is not winning for \forall at (φ, s) .*

The following proposition, which is an immediate consequence of Theorem 1 of the appendix, relates \Vdash_f to the usual satisfaction relation \Vdash .

Proposition 3. $\mathbb{S}, s \Vdash \varphi$ iff for every ops f for \forall in $\mathcal{E}(\wedge \Sigma, \mathbb{S})$: $\mathbb{S}, s \Vdash_f \varphi^b$.

The semantics of trace atoms is also given relative to an ops for \forall in the game $\mathcal{E}(\wedge \Sigma, \mathbb{S})$ (in the following often abbreviated to \mathcal{E}).

Definition 5. *Given an ops f for \forall in \mathcal{E} , we say that $\varphi \rightsquigarrow \psi$ is satisfied in \mathbb{S} at s with respect to f (and write $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \psi$) if there is an f -guided match*

$$(\varphi, s) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\psi, s) \quad (n \geq 0)$$

such that for no $i < n$ the formula φ_i is a μ -formula. We say that \mathbb{S} satisfies $\varphi \not\rightsquigarrow \psi$ at s with respect to f (and write $\mathbb{S}, s \Vdash_f \varphi \not\rightsquigarrow \psi$) iff $\mathbb{S}, s \not\Vdash_f \varphi \rightsquigarrow \psi$.

The idea behind the satisfaction of a trace atom $\varphi \rightsquigarrow \psi$ at a state s is that \exists can take the match from (φ, s) to (ψ, s) without passing through a μ -formula. This is good for the player \exists . For instance, if $\varphi \rightsquigarrow \psi$ and $\psi \rightsquigarrow \varphi$ are satisfied at s with respect to f for some $\varphi \neq \psi$, then f is necessarily losing for \forall at the position (φ, s) . We will later relate trace atoms to traces in infinitary proofs.

We interpret sequents disjunctively, that is: $\mathbb{S}, s \Vdash_f \Gamma$ whenever $\mathbb{S}, s \Vdash_f A$ for some $A \in \Gamma$. The sequent Γ is said to be *valid* whenever $\mathbb{S}, s \Vdash_f \Gamma$ for every model \mathbb{S} , state s of \mathbb{S} , and ops f for \forall in \mathcal{E} .

Remark 1. There is another way in which one could interpret sequents, which corresponds to what one might call *strong validity*, and which the reader should note is different from our notion of validity. Spelling it out, we say that Γ is *strongly valid* if for every model \mathbb{S} and state s there is an A in Γ such that for every ops f for \forall in \mathcal{E} it holds that $\mathbb{S}, s \Vdash_f A$. While these two notions coincide for sequents containing only annotated formulas, the sequent given by $\{\varphi \wedge \psi \rightsquigarrow \varphi, \varphi \wedge \psi \rightsquigarrow \psi\}$ shows that they do not in general.

We finish this subsection by defining three operations on sequents that, respectively, extract the formulas contained annotated in some sequent, take all annotated formulas out of focus, and put all formulas into focus.

$$\begin{aligned} \Gamma^- &:= \{\chi \mid \chi^b \in \Gamma \text{ for some } b \in \{\circ, \bullet\}\}, \\ \Gamma^\circ &:= \{\varphi \rightsquigarrow \psi \mid \varphi \rightsquigarrow \psi \in \Gamma\} \cup \{\varphi \not\rightsquigarrow \psi \mid \varphi \not\rightsquigarrow \psi \in \Gamma\} \cup \{\chi^\circ \mid \chi \in \Gamma^-\}, \\ \Gamma^\bullet &:= \{\varphi \rightsquigarrow \psi \mid \varphi \rightsquigarrow \psi \in \Gamma\} \cup \{\varphi \not\rightsquigarrow \psi \mid \varphi \not\rightsquigarrow \psi \in \Gamma\} \cup \{\chi^\bullet \mid \chi \in \Gamma^-\}. \end{aligned}$$

3.2 Proofs

In this subsection we give the rules of our proof system. Because the rule for modalities is quite involved, its details are given in a separate definition.

Definition 6. Let Γ be a sequent and let $[a]\varphi^b$ be an annotated formula. The jump $\Gamma^{[a]\varphi^b}$ of Γ with respect to $[a]\varphi^b$ consists of:

1. (a) $\varphi^{s([a]\varphi, \Gamma)}$;
 (b) $\psi^{s(\langle a \rangle \psi, \Gamma)}$ for every $\langle a \rangle \psi^c \in \Gamma$;
 (c) $[\check{a}]\chi^\circ$ for every $\chi^d \in \Gamma$ such that $[\check{a}]\chi \in \Sigma$;
2. (a) $\varphi \rightsquigarrow \langle \check{a} \rangle \chi$ for every $[a]\varphi \rightsquigarrow \chi \in \Gamma$ such that $\langle \check{a} \rangle \chi \in \Sigma$;
 (b) $\langle \check{a} \rangle \chi \not\rightsquigarrow \varphi$ for every $\chi \not\rightsquigarrow [a]\varphi \in \Gamma$ such that $\langle \check{a} \rangle \chi \in \Sigma$;
 (c) $\psi \rightsquigarrow \langle \check{a} \rangle \chi$ for every $\langle a \rangle \psi \rightsquigarrow \chi \in \Gamma$ such that $\langle \check{a} \rangle \chi \in \Sigma$;
 (d) $\langle \check{a} \rangle \chi \not\rightsquigarrow \psi$ for every $\chi \not\rightsquigarrow \langle a \rangle \psi \in \Gamma$ such that $\langle \check{a} \rangle \chi \in \Sigma$,

where $s(\xi, \Gamma)$ is defined by:

$$s(\xi, \Gamma) = \begin{cases} \bullet & \text{if } \xi^\bullet \in \Gamma, \\ \bullet & \text{if } \theta \not\rightsquigarrow \xi \in \Gamma \text{ for some } \theta^\bullet \in \Gamma, \\ \circ & \text{otherwise.} \end{cases}$$

Before we go on to provide the rest of the proof system, we will give some intuition for the modal rule, by proving the lemma below. This lemma essentially expresses that the modal rule is sound. Since the annotations play no role in the soundness of an individual rule, we suppress the annotations in the proof below for the sake of readability. Intuition for the annotations in the modal rule, and in particular for the function s , is given later.

Lemma 1. *Given a model \mathbb{S} , a state s of \mathbb{S} , and an ops f for \forall in \mathcal{E} such that $\mathbb{S}, s \not\vdash_f [a]\varphi^b, \Gamma$, there is an a -successor t of s , such that $\mathbb{S}, t \not\vdash_f \Gamma^{[a]\varphi^b}$.*

Proof. Let $\mathbb{S}, s \not\vdash_f [a]\varphi$ be the state t chosen by $f([a]\varphi, s)$. We claim that $\mathbb{S}, t \not\vdash_f \Gamma^{[a]\varphi^b}$. To start with, since f is winning, we have $\mathbb{S}, t \not\vdash_f \varphi$. Moreover, if $\langle a \rangle \psi$ belongs to Γ , then $\mathbb{S}, s \not\vdash_f \langle a \rangle \psi$ and thus $\mathbb{S}, s \not\vdash_f \psi$. Thirdly, if χ belongs to Γ and $[\check{a}]\chi \in \Sigma$, then, by optimality, it holds that $\mathbb{S}, t \not\vdash_f [\check{a}]\chi$.

The above shows all conditions under item 1. For the conditions under item 2, suppose that $\langle \check{a} \rangle \chi \in \Sigma$. We only show 2(d), because the others are similar. Suppose that $\chi \not\rightsquigarrow \langle a \rangle \psi \in \Gamma$. Then $\mathbb{S}, s \not\vdash_f \chi \not\rightsquigarrow \langle a \rangle \psi$, whence $\mathbb{S}, s \Vdash_f \chi \rightsquigarrow \langle a \rangle \psi$. That means that there is an f -guided \mathcal{E} -match

$$(\chi, s) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\langle a \rangle \psi, s) \quad (n \geq 0)$$

such that none of the φ_i 's is a μ -formula. But then the f -guided \mathcal{E} -match

$$(\langle \check{a} \rangle \chi, t) \cdot (\varphi_0, s_0) \cdots (\varphi_n, s_n) \cdot (\psi, t)$$

witnesses that $\mathbb{S}, t \not\vdash_f \langle \check{a} \rangle \chi \not\rightsquigarrow \psi$, as required.

The rules of the system Focus^2 are given in Figure 1. In each rule, the annotated formulas occurring in the set Γ are called *side formulas*. Moreover, the rules in $\{\mathbf{R}_\vee, \mathbf{R}_\wedge, \mathbf{R}_\mu, \mathbf{R}_\nu, \mathbf{R}_{[a]}\}$ have precisely one *principal formula*, which by definition is the annotated formula appearing to the left of Γ in the conclusion. Note that, due to the fact that sequents are taken to be sets, an annotated formula may at the same time be both a principal formula and a side formula.

We will now define the relation of *immediate ancestry* between formulas in the conclusion and formulas in the premisses of some arbitrary rule application. For any side formula in the conclusion of some rule, we let its *immediate ancestors* be the corresponding side formulas in the premisses. For every rule except $\mathbf{R}_{[a]}$, if some formula in the conclusion is a principal formula, its *immediate ancestors* are the annotated formulas occurring to the left of Γ in the premisses. Finally, for the *modal rule* $\mathbf{R}_{[a]}$, we stipulate that $\varphi^{s([a]\varphi, \Gamma)}$ is an *immediate ancestor* of the principal formula $[a]\varphi^b$, and that each $\psi^{s(\langle a \rangle \psi, \Gamma)}$ contained in $\Gamma^{[a]\varphi^b}$ due to clause 1(b) of Definition 6 is an *immediate ancestor* of $\langle a \rangle \psi^b \in \Gamma$.

As mentioned before, the purpose of the focus annotations is to keep track of *traces* of formulas on branches. Usually, a trace is a sequence of formulas $(\varphi_n)_{n < \omega}$ such that each φ_k is an immediate ancestor of φ_{k+1} . The idea is then that whenever an infinite branch has cofinitely many sequents with a formula in focus, this branch contains a trace on which infinitely many formulas are ν -formulas.

$\frac{}{\varphi^b, \bar{\varphi}^c, \Gamma} \text{Ax1}$	$\frac{}{\varphi \rightsquigarrow \psi, \varphi \not\rightsquigarrow \psi, \Gamma} \text{Ax2}$	$\frac{}{\varphi \rightsquigarrow \varphi, \Gamma} \text{Ax3}$
$\frac{(\varphi \vee \psi) \not\rightsquigarrow \varphi, (\varphi \vee \psi) \not\rightsquigarrow \psi, \varphi^b, \psi^b, \Gamma}{\varphi \vee \psi^b, \Gamma} \text{R}_\vee$		$\frac{\varphi^\circ, \Gamma \quad \bar{\varphi}^\circ, \Gamma}{\Gamma} \text{cut}$
$\frac{(\varphi \wedge \psi) \not\rightsquigarrow \varphi, \varphi^b, \Gamma \quad (\varphi \wedge \psi) \not\rightsquigarrow \psi, \psi^b, \Gamma}{\varphi \wedge \psi^b, \Gamma} \text{R}_\wedge$		$\frac{\varphi[\mu x \varphi/x]^\circ, \Gamma}{\mu x \varphi^b, \Gamma} \text{R}_\mu$
$\frac{\nu x \varphi \not\rightsquigarrow \varphi[\nu x \varphi/x], \varphi[\nu x \varphi/x] \rightsquigarrow \nu x \varphi, \varphi[\nu x \varphi/x]^b, \Gamma}{\nu x \varphi^b, \Gamma} \text{R}_\nu$		$\frac{\Gamma^{[a] \varphi^b}}{[a] \varphi^b, \Gamma} \text{R}_{[a]}$
$\frac{\Gamma^\bullet}{\Gamma^\circ} \text{F}$	$\frac{\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi, \varphi \not\rightsquigarrow \chi, \Gamma}{\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi, \Gamma} \text{trans}$	$\frac{\varphi \rightsquigarrow \psi, \Gamma \quad \varphi \not\rightsquigarrow \psi, \Gamma}{\Gamma} \text{tc}$

Fig. 1. The proof rules of the system Focus^2 .

Disregarding the backwards modalities for now, this can be seen as follows. As long as the focus rule is not applied, any focussed formula is an immediate ancestor of some earlier focussed formula. Since the principal formula of R_μ loses focus, while the principal formula of R_ν preserves focus, a straightforward application of Kőnig's Lemma shows that every infinite branch contains a trace with infinitely many ν -formulas. We refer the reader to [14] for more details.

Our setting is slightly more complicated, because the function s in Definition 6 additionally allows the focus to transfer along negated trace atoms, rather than just from a formula to one of its immediate ancestors. This is inspired by [22], as are the conditions in the second part of Definition 6. The main idea is that, because of the backwards modalities, traces may move not only up, but also down a proof tree. To get a grip on these more complex traces, we cut them up in segments consisting of upward paths, which are the same as ordinary traces, and loops, which are captured by the negated trace atoms. This intuitive idea will become explicit in the proof of completeness in Section 5.

We are now ready to define a notion of infinitary proofs in Focus^2 .

Definition 7. A Focus_∞^2 -proof is a (possibly infinite) derivation in Focus^2 with:

1. All leaves are axioms.
2. On every infinite branch cofinitely many sequents have a formula in focus.
3. Every infinite branch has infinitely many applications of $\text{R}_{[a]}$.

As mentioned above, conditions 2 and 3 are meant to ensure that every infinite trace contains infinitely many ν -formulas. We will use this in Section 5 to show that infinitary proofs are sound. The key idea is to relate the traces in a proof to matches in a purported countermodel of its conclusion.

We leave it to the reader to verify that each rule, apart from the modal rule, is truth-preserving with respect to a given model \mathbb{S} , state s of \mathbb{S} , and ops f for Refuter in $\mathcal{E}(\wedge \Sigma, \mathbb{S})$. Since Lemma 1 already showed the soundness of the modal rule, we obtain:

Proposition 4. *Well-founded Focus $^2_\infty$ -proofs are sound.*

We close this section with two examples of Focus $^2_\infty$ -proofs. The first example demonstrates cut and item 1(c) of Definition 6. The second example demonstrates trace atoms.

Example 1. Define the following two formulas:

$$\varphi := \mu x(\langle \check{a} \rangle x \vee p), \quad \psi := \nu y([a]x \wedge \varphi).$$

The formula φ expresses ‘there is a backwards a -path to some state where p holds’. The formula ψ expresses ‘ φ holds at every state reachable by a forwards a -path’. As our context Σ we take least negation-closed set containing φ and ψ :

$$\{\varphi, \langle \check{a} \rangle \varphi \vee p, \langle \check{a} \rangle \varphi, p, \psi, [a]\psi \wedge \varphi, [a]\psi, \bar{\varphi}, [\check{a}]\bar{\varphi} \wedge \bar{p}, p, [\check{a}]\bar{\varphi}, \bar{\psi}, \langle a \rangle \bar{\psi} \vee \bar{\varphi}, \langle a \rangle \bar{\psi}\}.$$

The implication $p \rightarrow \psi$ is valid, and below we give a Focus $^2_\infty$ -proof. As this particular proof does not rely on trace atoms, we omit them for readability.

$$\frac{\frac{\frac{\frac{\bar{p}^\bullet, \psi^\bullet, \langle \check{a} \rangle \varphi^\circ, p^\circ}{\bar{p}^\bullet, \psi^\bullet, \langle \check{a} \rangle \varphi \vee p^\circ} \text{Ax1}}{\bar{p}^\bullet, \psi^\bullet, \varphi^\circ} \text{R}_\vee}{\bar{p}^\bullet, \psi^\bullet, \varphi^\circ} \text{R}_\mu}{\bar{p}^\bullet, \psi^\bullet} \text{cut} \quad \frac{\frac{\frac{\frac{\psi^\bullet, [\check{a}]\bar{\varphi}^\circ}{\bar{p}^\bullet, [a]\psi^\bullet, \bar{\varphi}^\circ} \text{R}_{[a]}}{\bar{p}^\bullet, \psi^\bullet, \bar{\varphi}^\circ} \text{R}_\wedge}{\bar{p}^\bullet, [a]\psi \wedge \varphi^\bullet, \bar{\varphi}^\circ} \text{R}_\nu}{\bar{p}^\bullet, \psi^\bullet, \bar{\varphi}^\circ} \text{cut}}{\bar{p}^\bullet, \psi^\bullet} \text{cut}}{\bar{p}^\bullet, \psi^\bullet} \text{cut}$$

In the above proof, the proof π is given by

$$\frac{\frac{\frac{\frac{\frac{\bar{\varphi}^\circ, \varphi^\circ}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ, \langle \check{a} \rangle \varphi^\circ, p^\circ} \text{Ax1}}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ, \langle \check{a} \rangle \varphi \vee p^\circ} \text{R}_{[\check{a}]}}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ, \varphi^\circ} \text{R}_\vee}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{R}_\mu}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{cut} \quad \frac{\frac{\frac{\psi^\bullet, [\check{a}]\bar{\varphi}^\circ}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ, \bar{\varphi}^\circ} \text{Ax1}}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{R}_{[\check{a}]}}{[a]\psi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{cut}}{\psi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{R}_\wedge}{\frac{\frac{\frac{\frac{\varphi^\circ, \bar{\varphi}^\circ}{\langle \check{a} \rangle \varphi^\circ, p^\circ, [\check{a}]\bar{\varphi}^\circ} \text{Ax1}}{\langle \check{a} \rangle \varphi \vee p^\circ, [\check{a}]\bar{\varphi}^\circ} \text{R}_\vee}{\varphi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{R}_\mu}{[a]\psi \wedge \varphi^\bullet, [\check{a}]\bar{\varphi}^\circ} \text{R}_\nu} \text{cut}$$

where the vertical dots indicate that the proof continues by repeating what happens at the root of π . The resulting proof of $\bar{p}^\bullet, \psi^\bullet$ has a single infinite branch, which can easily be seen to satisfy Condition 2 of Definition 7.

Example 2. Define $\varphi := \nu x \langle a \rangle \langle \check{a} \rangle x$, i.e. φ expresses that there is an infinite path of alternating a and \check{a} transitions. Clearly this holds at every state with an a -successor. Hence the implication $\langle a \rangle p \rightarrow \varphi$ is valid. As context Σ we consider the least negation-closed set containing both $\langle a \rangle p$ and φ , i.e.,

$$\{\langle a \rangle p, p, \varphi, \langle a \rangle \langle \check{a} \rangle \varphi, \langle \check{a} \rangle \varphi, [a]\bar{p}, \bar{p}, \bar{\varphi}, [a][\check{a}]\bar{\varphi}, [\check{a}]\bar{\varphi}\}.$$

The following is a Focus_∞^2 -proof of $\langle a \rangle p \rightarrow \varphi$.

$$\frac{\frac{\overline{\overline{p^\bullet}, \langle \check{a} \rangle \varphi^\bullet, \langle \check{a} \rangle \varphi \not\rightsquigarrow \langle \check{a} \rangle \varphi, \langle \check{a} \rangle \varphi \rightsquigarrow \langle \check{a} \rangle \varphi}}{[a]\overline{p^\bullet}, \langle a \rangle \langle \check{a} \rangle \varphi^\bullet, \varphi \not\rightsquigarrow \langle a \rangle \langle \check{a} \rangle \varphi, \langle a \rangle \langle \check{a} \rangle \varphi \rightsquigarrow \varphi}}{[a]\overline{p^\bullet}, \varphi^\bullet} \begin{array}{l} \text{Ax2} \\ \text{R}_{[a]} \\ \text{R}_\nu \end{array}$$

Note that it is also possible to use Ax3 instead of Ax2 in the above proof.

4 The proof search game

We will define a proof search game $\mathcal{G}(\Sigma)$ for the proof system Focus_∞^2 in the standard way. First, we require a slightly more formal definition of the notion of a rule instance.

Definition 8. A rule instance is a triple $(\Gamma, r, \langle \Delta_1, \dots, \Delta_n \rangle)$ such that

$$\frac{\Delta_1 \cdots \Delta_n}{\Gamma} r$$

is a valid rule application in Focus^2 .

The set of positions of $\mathcal{G}(\Sigma)$ is $\text{Seq}_\Sigma \cup \text{Inst}_\Sigma$, where Seq_Σ is the set of sequents and Inst_Σ is the set of valid rule instances (containing only formulas in Σ). Since Σ is finite, the game $\mathcal{G}(\Sigma)$ has only finitely many positions. The ownership function and admissible moves of $\mathcal{G}(\Sigma)$ are as in the following table:

Position	Owner	Admissible moves
$\Gamma \in \text{Seq}_\Sigma$	Prover	$\{i \in \text{Inst}_\Sigma \mid \text{conc}(i) = \Gamma\}$
$(\Gamma, r, \langle \Delta_1, \dots, \Delta_n \rangle) \in \text{Inst}_\Sigma$	Refuter	$\{\Delta_i \mid 1 \leq i \leq n\}$

In the above table, the expression $\text{conc}(i)$ stands for the conclusion (*i.e.* the first element of the triple) of the rule instance i . As usual, a finite match is lost by the player who got stuck. An infinite $\mathcal{G}(\Sigma)$ -match is won by Prover if and only if it has a final segment

$$\Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdots$$

on which each Γ_k has at least one formula in focus and the instance i_k is an application of $\text{R}_{[a]}$ for infinitely many k . The two main observations about $\mathcal{G}(\Sigma)$ that we will use are the following:

1. A Focus_∞^2 -proof of Γ is the same as a winning strategy for Prover in $\mathcal{G}(\Sigma)@ \Gamma$.
2. $\mathcal{G}(\Sigma)$ is a parity game, whence positionally determined.

The first observation is immediate when viewing a winning strategy as a subtree of the full game tree. To make the second observation more explicit, we give the parity function Ω for $\mathcal{G}(\Sigma)$. On Seq_Σ , we simply set $\Omega(\Gamma) := 0$ for every $\Gamma \in \text{Seq}_\Sigma$. On Inst_Σ , we define:

$$\Omega(\Gamma, r, \langle \Delta_1, \dots, \Delta_n \rangle) := \begin{cases} 3 & \text{if } \Gamma \text{ has no formula in focus,} \\ 2 & \text{if } \Gamma \text{ has a formula in focus and } r = \text{R}_{[a]}, \\ 1 & \text{if } \Gamma \text{ has a formula in focus and } r \neq \text{R}_{[a]}. \end{cases}$$

As a result we immediately obtain a method to reduce general non-well-founded proofs to cyclic proofs. Indeed, if Prover has a winning strategy, she also has *positional* winning strategy, which clearly corresponds to a *regular* Focus_∞^2 -proof (that is, a proof containing only finitely many non-isomorphic subtrees.)

5 Soundness and completeness

In this section we will prove the soundness and completeness of the system Focus_∞^2 . More specifically, for soundness we will show that if Γ is invalid, then Refuter has a winning strategy in $\mathcal{G}(\Sigma)@ \Gamma$. Our completeness result is slightly less wide in scope, showing only that if Refuter has a winning strategy in $\mathcal{G}(\Sigma)@ \Gamma$, then Γ^- is invalid.

5.1 Soundness

For soundness, we assume an ops f for \forall in $\mathcal{E} := \mathcal{E}(\wedge \Sigma, \mathbb{S})$ for some \mathbb{S} and s such that $\mathbb{S}, s \not\Vdash_f \Gamma$. The goal is to construct from f a strategy T_f for Refuter in $\mathcal{G} := \mathcal{G}(\Sigma)$. The key idea is to assign to each position p reached in \mathcal{G} a state s such that whenever $p = \Delta \in \text{Seq}_\Sigma$ it holds that $\mathbb{S}, s \not\Vdash_f \Delta$. For $p \in \text{Inst}_\Sigma$, the choice of T_f is then based on $f(\varphi, s)$ where φ is a formula determined by the rule instance p . The existence of such an s implies that p cannot be an axiom and thus that Refuter never gets stuck. For infinite matches, the proof works by showing that a T_f -guided $\mathcal{G}@ \Gamma$ -match lost by Refuter induces an f -guided $\mathcal{E}@ \varphi$ -match lost by \forall . As mentioned above, the key idea here is to relate an f -guided $\mathcal{E}@ \varphi$ -match to a trace through the T_f -guided $\mathcal{G}@ \Gamma$ -match. If the $\mathcal{G}@ \Gamma$ -match is losing for Refuter, it must contain a trace with infinitely many ν -formulas, which gives us an $\mathcal{E}@ \varphi$ -match lost by \forall . A novel challenge here is that not all steps in a trace necessarily go from a formula to one of its immediate ancestors, but may instead transfer along a negated trace atom. When this happens, say from φ_n to φ_{n+1} , it holds for Δ as above that both φ_n^\bullet and $\varphi_n \not\rightsquigarrow \varphi_{n+1}$ belong to Δ . Since, by the above, it holds that $\mathbb{S}, s \not\Vdash_f \Delta$, we use the fact that $\mathbb{S}, s \Vdash_f \varphi_n \rightsquigarrow \varphi_{n+1}$ to take the $\mathcal{E}@ \varphi$ -match from (φ_n, s) to (φ_{n+1}, s) . In the end, we obtain:

Proposition 5. *If Γ is the conclusion of a Focus_∞^2 -proof, then Γ is valid.*

5.2 Completeness

For completeness we conversely show that from a winning strategy T for Refuter in $\mathcal{G}@ \Gamma$, we can construct a model \mathbb{S}^T and a positional strategy f_T for \forall in $\mathcal{E}(\wedge \Sigma, \mathbb{S}^T)$ such that \mathbb{S}^T falsifies Γ^- with respect to f_T . The strategy f_T we construct will not necessarily be optimal, but by Theorem 1 of Appendix A it follows that there must also be an ops $\overline{f_T}$ such that $\mathbb{S}^T \not\Vdash_{\overline{f_T}} \Gamma^-$. We will view T as a tree, and restrict attention a certain subtree. We first need to define two relevant properties of rule applications.

Definition 9. *A rule application is cumulative if all of the premisses are supersets of the conclusion. A rule application is productive if all of the premisses are distinct from the conclusion.*

Without renaming T , we restrict T to its subtree where Prover adheres to the following (non-deterministic) strategy:

1. Exhaustively apply productive instances of cut and tc.
2. If applicable, apply the focus rule.
3. Exhaustively take applications of R_\vee , R_\wedge , R_μ , R_ν , trans that are both cumulative and productive.
4. If applicable, apply an axiom.
5. If applicable, apply a modal rule and loop back to stage (1).

It is not hard to see that each of the above phases terminates. More precisely, phases (2), (4) and (5) either terminate immediately or after applying a single rule. By the productivity requirement and the finiteness of Σ , phases (1) and (3) must terminate after a finite number of rule applications as well. Note also that non-cumulative rule applications can only happen in phases (2) or (5).

We will now define the model \mathbb{S}^T . The set S^T of states consists of maximal paths in T not containing a modal rule. We write $\Gamma(\rho)$ for $\bigcup\{\Gamma : \Gamma \text{ occurs in } \rho\}$. Note that, since the only possibly non-cumulative rule application in ρ is the focus rule, $\Gamma(\rho)^\bullet = \text{last}(\rho)^\bullet$ for every state ρ of \mathbb{S}^T . Moreover, we write $\rho_1 \xrightarrow{a} \rho_2$ if ρ_2 is directly above ρ_1 in T , separated only by an application of $R_{[a]}$ (we assume that trees grow upwards). We write \rightarrow for the union $\bigcup\{\xrightarrow{a} : a \in D\}$. Clearly, under the relation \rightarrow the states of \mathbb{S}^T form a forest (not necessarily a tree!). We write $\rho \leq \tau$ if τ is a descendant of ρ in this forest, *i.e.* \leq is the reflexive-transitive closure of \rightarrow . The relations R_a^T of \mathbb{S}^T are defined as follows:

$$\rho_1 R_a^T \rho_2 \text{ if and only if } \rho_1 \xrightarrow{a} \rho_2 \text{ or } \rho_2 \xrightarrow{\bar{a}} \rho_1.$$

Note that \mathbb{S}^T is clearly regular. We define the valuation $V^T : S^T \rightarrow \mathcal{P}(P)$ by

$$V^T(\rho) := \{p : \bar{p} \in \Gamma(\rho)^-\}.$$

The restriction on T , together with the fact that it is winning for Refuter, guarantees that each $\Gamma(\rho)$ satisfies certain saturation properties, which are spelled out in the following lemma. We will later use these saturation conditions to construct our positional strategy f_T for \forall in $\mathcal{E}(\wedge \Sigma, \mathbb{S}^T)$ and to show that \mathbb{S}^T falsifies Γ with respect to f_T .

Lemma 2. *For every state ρ of \mathbb{S}^T , the set $\Gamma(\rho)$ is saturated. That is, it satisfies all of the following conditions:*

- For no φ it holds that $\varphi, \bar{\varphi} \in \Gamma(\rho)^-$.
- For all φ it holds that $\varphi^\circ \in \Gamma(\rho)$ if and only if $\bar{\varphi}^\circ \notin \Gamma(\rho)$.
- For all φ it holds that $\varphi \rightsquigarrow \psi \in \Gamma(\rho)$ if and only if $\varphi \not\rightsquigarrow \psi \notin \Gamma(\rho)$.
- For no φ it holds that $\varphi \rightsquigarrow \varphi \in \Gamma(\rho)$.

- If $\psi_1 \vee \psi_2 \in \Gamma(\rho)^-$, then for both i : $\psi_1 \vee \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$ and $\psi_i \in \Gamma(\rho)^-$.
- If $\psi_1 \wedge \psi_2 \in \Gamma(\rho)^-$, then for some i : $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$ and $\psi_i \in \Gamma(\rho)^-$.
- If $\mu x \varphi \in \Gamma(\rho)^-$, then $\varphi[\mu x \varphi/x] \in \Gamma(\rho)^-$.
- If $\nu x \varphi \in \Gamma(\rho)^-$, then $\nu x \varphi \not\rightsquigarrow \varphi[\nu x \varphi/x] \in \Gamma(\rho)$ and $\varphi[\nu x \varphi/x] \in \Gamma(\rho)^-$.
- If $\nu x \varphi \in \Gamma(\rho)^-$, then $\varphi[\nu x \varphi/x] \rightsquigarrow \nu x \varphi \in \Gamma(\rho)$.
- If $\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi \in \Gamma(\rho)$, then $\varphi \not\rightsquigarrow \chi \in \Gamma(\rho)$.

Now let ρ_0 be a state of \mathbb{S}^T containing the root Γ and let φ_0 be some formula such that $\varphi_0 \in \Gamma^-$. We wish to show that φ_0 is *not* satisfied at ρ_0 in \mathbb{S}^T . To this end, we will construct a winning strategy f_T for \forall in the game $\mathcal{E} := \mathcal{E}(\wedge \Sigma, \mathbb{S}^T)$ initialised at (φ_0, ρ_0) . The strategy f_T is defined as follows:

- At $(\psi_1 \wedge \psi_2, \rho)$, pick a conjunct $\psi_i \in \Gamma(\rho)^-$ such that $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$.
- At $([a]\varphi, \rho)$, choose (φ, τ) for some τ such that $\rho \xrightarrow{a} \tau$ by virtue of some application of $R_{[a]}$ with $[a]\varphi^b$ principal for some $b \in \{\circ, \bullet\}$.

Before we show that f_T is winning for \forall , we must first argue that it is well defined. By saturation, for every formula $\psi_1 \wedge \psi_2$ contained in $\Gamma(\rho)^-$, there is a $\psi_i \in \Gamma(\rho)^-$ with $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$. Likewise, for every formula $[a]\varphi^b \in \Gamma(\rho)$, there is a τ directly above ρ in T , separated only by an application of $R_{[a]}$ with $[a]\varphi^b$ principal. The following lemma therefore suffices. Its proof is by induction on the length of \mathcal{M} and heavily relies on the saturation properties of Lemma 2.

Lemma 3. *Let \mathcal{M} be an f_T -guided \mathcal{E} -match initialised at (φ_0, ρ_0) . Then for any position (φ, ρ) occurring in \mathcal{M} it holds that $\varphi \in \Gamma(\rho)^-$. Moreover, if (φ, ρ) comes directly after a modal step and the focus rule is applied in ρ , then $\varphi^\bullet \in \Gamma(\rho)$.*

The following lemma is key to the completeness proof. It shows that if an f_T -guided $\mathcal{E}@\!(\varphi_0, \rho_0)$ -match loops from some state ρ to itself, without passing through a μ -formula, then this information is already contained in ρ in the form of a negated trace atom. The proof goes by induction on the number of distinct states of \mathbb{S}^T occurring in \mathcal{N} . The base case, where only ρ is visited, can be shown by applying several instances of Lemma 2. For the inductive step, we crucially rely on the conditions 2(a) – 2(d) of Definition 6 to relate the trace atoms in two states τ and τ' such that $\tau R_a^T \tau'$.

Lemma 4. *Let $\rho \in \mathbb{S}^T$. Suppose that an f_T -guided $\mathcal{E}@\!(\varphi_0, \rho_0)$ -match \mathcal{M} has a segment \mathcal{N} of the form:*

$$(\varphi, \rho) = (\psi_0, s_0) \cdot (\psi_1, s_1) \cdots (\psi_n, s_n) = (\psi, \rho) \quad (n \geq 0)$$

such that for no $i < n$ the formula φ_i is a μ -formula. Then $\varphi \not\rightsquigarrow \psi \in \Gamma(\rho)$.

With the above lemmata in place, we are ready to prove that \forall wins every full f_T -guided $\mathcal{E}@\!(\varphi_0, \rho_0)$ -match \mathcal{M} . If \mathcal{M} is finite, it is not hard to show that it must be \exists who got stuck. If \mathcal{M} is infinite, the proof depends on whether \mathcal{M} visits some single state infinitely often. If it does, one can show that if \exists would win the match \mathcal{M} , then \mathcal{M} would visit some state ρ with $\nu x \varphi, \varphi[\nu x \varphi/x] \not\rightsquigarrow \varphi \in \Gamma(\rho)^-$, contradicting saturation. If, on the other hand, \mathcal{M} visits each state at most

finitely often, the proof works by showing that a win for \exists in \mathcal{M} would imply that T contains an infinite branch won by Prover, which is also a contradiction. In the end, we obtain the following proposition.

Proposition 6. *The strategy f_T is winning for \forall in $\mathcal{E}@\langle\varphi_0, \rho_0\rangle$.*

Since φ_0 was chosen arbitrarily from Γ^- , we find that $\mathbb{S}^T \not\Vdash_{f_T} \Gamma^-$. Hence, by Theorem 1 of Appendix A, we obtain completeness for the formulas in a sequent.

Proposition 7. *If Γ^- is valid, then Γ has a Focus_∞^2 -proof.*

6 Conclusion

We have constructed a non-well-founded proof system Focus_∞^2 for the two-way alternation-free modal μ -calculus $\mathcal{L}_{2\mu}^{af}$. This system naturally reduces to a cyclic system when restricting to positional strategies in the proof search game.

Using the proof search game and the game semantics for the modal μ -calculus, we have shown that the system is sound for all sequents, and complete for those sequents not containing trace atoms. A natural first question for future research is to see if a full completeness result can be obtained. For this, a logic of trace atoms would have to be developed. One could for instance think of a rule like

$$\frac{\varphi \rightsquigarrow \chi, \Gamma \quad \psi \rightsquigarrow \chi, \Gamma}{\varphi \wedge \psi \rightsquigarrow \chi, \Gamma} \text{R}_\wedge$$

Following on this, we think it would be interesting to properly include trace atoms in the syntax by allowing the Boolean, modal and perhaps even the fixed point operators to apply to trace atoms. An example of a valid formula in this syntax is given by $((\varphi \rightsquigarrow \langle a \rangle \psi) \wedge [a](\psi \rightsquigarrow \langle \check{a} \rangle \varphi)) \rightarrow \varphi$.

Another pressing question is whether our system could be used to prove interpolation, as has been done for language without backwards modalities in [14]. To the best of our knowledge it is currently an open question whether $\mathcal{L}_{2\mu}^{af}$ has interpolation. At the same time, it is known that analytic applications of the cut rule do not necessarily interfere with the process of extracting interpolants from proofs [9, 16].

Finally, it would be interesting to see if our system can be extended to the full language $\mathcal{L}_{2\mu}$. The main challenge would be to keep track of the most important fixed point variable being unfolded on a trace. Perhaps this could be done by employing an annotation system such as the one by Jungteerapanich and Stirling [21, 7], together with trace atoms that record the most important fixed point variable unfolded on a loop.

Acknowledgements We thank Johannes Marti for insightful conversations at the outset of the present research. We also thank the anonymous reviewers for their helpful comments.

A Parity games

Definition 10. A (two-player) game is a structure $\mathcal{G} = (B_0, B_1, E, W)$ where E is a binary relation on $B := B_0 + B_1$, and W is a map $B^\omega \rightarrow \{0, 1\}$.

The set B is called the *board* of \mathcal{G} , and its elements are called *positions*. Whether a position belongs to B_0 or B_1 determines which player *owns* that position. If a player $\Pi \in \{0, 1\}$ owns a position q , it is their turn to play and the set of their *admissible moves* is given by the image $E[q]$.

Definition 11. A match in $\mathcal{G} = (B_0, B_1, E, W)$ (or simply a \mathcal{G} -match) is a path \mathcal{M} through the graph (B, E) . A match is said to be *full* if it is a maximal path.

Note that a full match \mathcal{M} is either finite, in which case $E[\text{last}(\mathcal{M})] = \emptyset$, or infinite. For a $\Pi \in \{0, 1\}$, we write $\bar{\Pi}$ for the other player $\Pi + 1 \pmod 2$.

Definition 12. A full match \mathcal{M} in $\mathcal{G} = (B_0, B_1, E, W)$ is won by player Π if either \mathcal{M} is finite and $\text{last}(\mathcal{M}) \in B_{\bar{\Pi}}$, or \mathcal{M} is infinite and $W(\mathcal{M}) = \Pi$.

If a full match \mathcal{M} is finite, and $\text{last}(\mathcal{M})$ belongs to B_Π for $\Pi \in \{0, 1\}$, we say that the player Π got *stuck*. A *partial match* is a match which is not full.

Definition 13. In the context of a game \mathcal{G} , we denote by PM_Π the set of partial \mathcal{G} -matches \mathcal{M} such that $\text{last}(\mathcal{M})$ belongs to the player Π .

Definition 14. A strategy for Π in a game \mathcal{G} is a map $f : \text{PM}_\Pi \rightarrow B$. Moreover, a \mathcal{G} -match \mathcal{M} is said to be *f-guided* if for any $\mathcal{M}_0 \sqsubset \mathcal{M}$ with $\mathcal{M}_0 \in \text{PM}_\Pi$ it holds that $\mathcal{M}_0 \cdot f(\mathcal{M}_0) \sqsubseteq \mathcal{M}$.

For a position q , the set $\text{PM}_\Pi(q)$ contains all $\mathcal{M} \in \text{PM}_\Pi$ such that $\text{first}(\mathcal{M}) = q$.

Definition 15. A strategy f for Π in \mathcal{G} is *surviving at a position q* if $f(\mathcal{M})$ is admissible for every $\mathcal{M} \in \text{PM}_\Pi(q)$, and *winning at q* if in addition all full f -guided matches starting at q are won by Π . A position q is said to be *winning for Π* if Π has a strategy winning at q . We denote the set of all positions in \mathcal{G} that are winning for Π by $\text{Win}_\Pi(\mathcal{G})$.

We write $\mathcal{G}@q$ for the game \mathcal{G} initialised at the position q of \mathcal{G} . A strategy f for Π is *surviving (winning)* in $\mathcal{G}@q$ if it is surviving (winning) in \mathcal{G} at q .

Definition 16. A strategy f is *positional* if it only depends on the last move, i.e. if $f(\mathcal{M}) = f(\mathcal{M}')$ for all $\mathcal{M}, \mathcal{M}' \in \text{PM}_\Pi$ with $\text{last}(\mathcal{M}) = \text{last}(\mathcal{M}')$.

We will often present a positional strategy for Π as a map $f : B_\Pi \rightarrow B$.

Definition 17. A *priority map on some board B* is a map $\Omega : B \rightarrow \omega$ of finite range. A *parity game* is a game of which the winning condition is given by $W_\Omega(\mathcal{M}) = \max(\text{Inf}_\Omega(\mathcal{M})) \pmod 2$, where $\text{Inf}_\Omega(\mathcal{M})$ is the set of positions occurring infinitely often in \mathcal{M} .

The following theorem captures the key property of parity games: they are *positionally determined*. In fact, each player Π has a positional strategy f_Π that is *optimal*, in the sense that f_Π is winning for Π in $\mathcal{G}@q$ for every $q \in \text{Win}_\Pi(\mathcal{G})$.

Theorem 1 ([15, 5]). For any parity game \mathcal{G} , there are positional strategies f_Π for each player $\Pi \in \{0, 1\}$, such that for every position q one of the f_Π is a winning strategy for Π in $\mathcal{G}@q$.

References

1. Afshari, B., Enqvist, S., Leigh, G.E., Marti, J., Venema, Y.: Proof Systems for Two-way Modal μ -Calculus (2023), ILLC Prepublication Series. PP-2023-03
2. Afshari, B., Jäger, G., Leigh, G.E.: An infinitary treatment of full μ -calculus. In: Iemhoff, R., Moortgat, M., de Queiroz, R.J.G.B. (eds.) WoLLIC 2019. Lecture Notes in Computer Science, vol. 11541, pp. 17–34. Springer (2019)
3. Afshari, B., Leigh, G.E., Turata, G.M.: Uniform interpolation from cyclic proofs: The case of modal μ -calculus. In: Das, A., Negri, S. (eds.) TABLEAUX 2021. LNCS, vol. 12842, pp. 335–353. Springer (2021)
4. Ciabattoni, A., Lang, T., Ramanayake, R.: A theory of cut-restriction: first steps (2022), arXiv: 2203.01600
5. Emerson, E.A., Jutla, C.S.: Tree automata, μ -calculus and determinacy (extended abstract). In: 32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991. pp. 368–377. IEEE Computer Society (1991)
6. Goré, R.: Tableau methods for modal and temporal logics. In: D’Agostino, M., Gabbay, D.M., Hähnle, R., Posegga, J. (eds.) Handbook of Tableau Methods, pp. 297–396. Springer Netherlands, Dordrecht (1999)
7. Jungteerapanich, N.: A tableau system for the modal μ -calculus. In: Giese, M., Waaler, A. (eds.) TABLEAUX 2009. LNCS, vol. 5607, pp. 220–234. Springer (2009)
8. Kashima, R.: Cut-free sequent calculi for some tense logics. *Studia Logica* **53**(1), 119–136 (1994)
9. Kowalski, T., Ono, H.: Analytic cut and interpolation for bi-intuitionistic logic. *Review of Symbolic Logic* **10**(2), 259–283 (2017)
10. Kozen, D.: Results on the propositional μ -calculus. *Theoretical Computer Science* **27**, 333–354 (1983)
11. Kupke, C., Marti, J., Venema, Y.: Size matters in the modal μ -calculus (2020), arXiv: 2010.14430
12. Lange, M., Stirling, C.: Focus games for satisfiability and completeness of temporal logic. In: LICS 2001. pp. 357–365 (2001)
13. Marti, J., Venema, Y.: Focus-style proof systems and interpolation for the alternation-free μ -calculus (2021), arXiv: 2103.01671
14. Marti, J., Venema, Y.: A focus system for the alternation-free μ -calculus. In: Das, A., Negri, S. (eds.) TABLEAUX 2021. LNCS, vol. 12842, pp. 371–388. Springer (2021)
15. Mostowski, A.: Games with forbidden positions. Tech. rep., Instytut Matematyki, Uniwersytet Gdański, Poland (1991)
16. Nguyen, L.A.: Analytic tableau systems and interpolation for the modal logics KB, KDB, K5, KD5. *Studia Logica* **69**(1), 41–57 (2001)
17. Niwiński, D., Walukiewicz, I.: Games for the μ -calculus. *Theoretical Computer Science* **163**(1&2), 99–116 (1996)
18. Ohnishi, M., Matsumoto, K.: Gentzen method in modal calculi. *Osaka Mathematical Journal* **9**(2), 113 – 130 (1957)
19. Rooduijn, J.M.W.: Cyclic hypersequent calculi for some modal logics with the master modality. In: Das, A., Negri, S. (eds.) TABLEAUX 2021. LNCS, vol. 12842, pp. 354–370. Springer (2021)
20. Rooduijn, J.M.W., Zenger, L.: An analytic proof system for common knowledge logic over S5. In: David Fernández-Duque, A.P., Pinchinat, S. (eds.) AiML 2022. pp. 659–680. College Publications (2022)

21. Stirling, C.: A tableau proof system with names for modal μ -calculus. In: Voronkov, A., Korovina, M.V. (eds.) HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday, EPiC Series in Computing, vol. 42, pp. 306–318 (2014)
22. Vardi, M.Y.: Reasoning about the past with two-way automata. In: Larsen, K.G., Skyum, S., Winskel, G. (eds.) ICALP 1998. LNCS, vol. 1443, pp. 628–641. Springer (1998)
23. Venema, Y.: Lectures on the modal μ -calculus. Lecture notes. Institute for Logic, Language and Computation, University of Amsterdam (2020)
24. Walukiewicz, I.: Completeness of Kozen's axiomatisation of the propositional μ -calculus. *Information and Computation* **157**(1-2), 142–182 (2000)