

This is a "preproof" accepted article for *The Journal of Symbolic Logic*.
This version may be subject to change during the production process.
DOI: 10.1017/jsl.2023.60

Proof Systems for Two-way Modal μ -Calculus*

Bahareh Afshari¹, Sebastian Enqvist², Graham E. Leigh¹,
Johannes Marti³, and Yde Venema⁴

¹Department of Philosophy, Linguistics and Theory of Science,
University of Gothenburg

²Department of Philosophy, Stockholm University

³Department of Informatics, University of Zurich

⁴Institute for Logic, Language and Computation, University of
Amsterdam

August 31, 2023

Abstract

We present sound and complete sequent calculi for the modal μ -calculus with converse modalities, aka two-way modal μ -calculus. Notably, we introduce a cyclic proof system wherein proofs can be represented as finite trees with back-edges, i.e., finite graphs. The sequent calculi incorporate ordinal annotations and structural rules for managing them. Soundness is proved with relative ease as is the case for the modal μ -calculus with explicit ordinals. The main ingredients in the proof of completeness are isolating a class of non-wellfounded proofs with sequents of bounded size, called slim proofs, and a counter-model construction that shows slimness suffices to capture all validities. Slim proofs are further transformed into cyclic proofs by means of re-assigning ordinal annotations.

1 Introduction

The modal μ -calculus is an extension of basic modal logic with least and greatest fixpoint operators. The additional operators are given an interpretation that breaks the locality of modal logic. Notably, the calculus can express all bisimulation-invariant monadic second-order properties [10]. As a consequence,

ACM CCS • Theory of computation - Logic - Proof theory • Theory of computation - Logic - Modal and temporal logics • Theory of computation - Logic - Logic and verification

*This work was supported by the Knut and Alice Wallenberg Foundation [2020.0199], Swedish Research Council [2020-01873, 2017-05111], and Dutch Research Council [OCENW.M20.048, 617.001.857].

well-studied modal logics such as the temporal logics LTL, CTL, CTL* and the program logic PDL can be translated into the μ -calculus. Many theoretical results on the modal μ -calculus have been established through its connection with automata theory and the theory of infinite games [7], the central observation being that every formula can be represented as an alternating tree automaton, and vice versa, such that the automaton accepts an infinite tree if and only if the tree is a model of the formula [9, 22].

The two-way modal μ -calculus, also known as the full μ -calculus, is an extension of the μ -calculus with modal operators for converses of accessibility relations. Thus, in addition to the standard modalities $[a]$ and $\langle a \rangle$ that quantify over a -successors (states reachable via a single a -labeled transition), the two-way μ -calculus includes modalities $[\check{a}]$ and $\langle \check{a} \rangle$ quantifying over a -predecessors. A central result, due to Vardi [21], is that the satisfiability problem for the two-way μ -calculus is decidable in exponential time. To prove this result, Vardi introduces a notion of (alternating) two-way automaton and shows that for every formula of the two-way μ -calculus there is a two-way automaton that accepts an infinite tree if and only if the tree *encodes* a model of the formula. The decidability result then follows with a construction that provides for every two-way automaton an equivalent nondeterministic parity tree automaton. Vardi's construction does not induce a translation of two-way μ -calculus formulas into equivalent μ -calculus formulas. The translation merely preserves satisfiability and validity. Indeed, two-way μ -calculus is strictly more expressive than its 'one-way' fragment, for instance it lacks the finite model property [19].

We present a sound and complete sequent calculus for the two-way μ -calculus. The proof theory of the logic has not been extensively explored. It is still an open question whether the calculus is complete with respect to a Hilbert-style axiomatisation that includes Kozen's induction rule for the fixpoint operators. A complete finitary Hilbert-style axiomatisation of "flat" fragments is given in [6] and a cyclic system for the alternation-free fragment in [14]. A sound and complete infinitary proof system for the full calculus is provided in [2]. The system we present here is finitary. More precisely, proofs are represented as finite (cyclic) graphs with a local correctness criterion on simple cycles.

The proof system we introduce is a variant of systems developed for the modal μ -calculus by Dam and Gurov [4] and for the first-order μ -calculus by Dam and Sprenger [17]. It further incorporates ideas developed by Jungteerapanich [11] and Stirling [18], using a derivation rule influenced by the Safra-construction for ω -automata to formulate the correctness criterion on cycles. This approach has been utilised in [1] to introduce a path-based cyclic proof system for first-order μ -calculus that is complete for the fragment corresponding to the one-way μ -calculus.

The distinguishing feature of the proof systems developed by Dam and Gurov is to work in an extended syntax with explicit variables referring to ordinal approximants of least fixed-points, permitting the expression of propositions like "the least fixed point of the map $x \mapsto \varphi(x)$ is reached at a smaller ordinal than the least fixed point of $x \mapsto \psi(x)$ ". This added expressive power plays a crucial role in our completeness proof, allowing us to build small saturated sets

of formulas, called *tiles*, from which a tree-like counter-model for an unprovable formula is constructed step by step. The added difficulty in the counter-model construction (compared to the modal μ -calculus) is that what is true at a vertex in the tree-model depends on both successors and predecessors of the vertex, a condition that needs to be taken into account in the saturation process.

The heart of Vardi's decision procedure for the two-way μ -calculus is the use of auxiliary second-order variables as part of an extended alphabet for simulating non-determinism, allowing the simulating automaton to guess partial information about "loops" that can occur as a result of an alternating two-way automaton traversing the tree in both directions, i.e., ancestor to descendant or descendant to ancestor. These variables can then be projected away. In our setting, the "guessing" happens in the form of cuts on formulas that encode (via ordinal variables) information about fixpoint unfoldings. We do not know whether our system is complete without the cut rule. A crucial part of our completeness argument, however, shows that the cut formulas can be chosen to belong to a relatively small finite set. Therefore, although the proof system is not cut-free, it does support automatic proof search.

Outline The structure of this paper is as follows: In Section 2 we discuss the necessary preliminaries related to the two-way μ -calculus and annotated formulas. Section 3 contains the definition of our cyclic proof system. In Section 4 we prove that this system is sound. The completeness proof consists of two parts: In Section 5 we first show completeness for a particular class of non-wellfounded proofs, which we call slim proofs. In Section 6 we then show that every slim proof can be transformed into a cyclic proof in our system.

2 Two-way μ -calculus

The syntax of the two-way μ -calculus makes use of the following non-logical symbols: a countably infinite set Prop of propositional constants or proposition letters (denoted p, q, p_0, \dots), with an involution $p \mapsto \bar{p}$; a countably infinite set Act of action symbols (denoted a, b, a_0, \dots), with an involution $a \mapsto \check{a}$; and a countably infinite set FV of fixed point variable symbols (denoted x, y, z, \dots).

It will be convenient for us to work with formulas that are in negation normal form. That is, the set of (plain) two-way μ -calculus formulas is given by the following grammar.

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

We will need the following basic syntactic definitions. The set $Sfor(\varphi)$ of subformulas of φ is defined as usual. The set of variables that occur in φ is denoted as $Var(\varphi)$. Since the fixpoint operators μ and ν bind the variables that they occur with, we can define in a standard way the notions of free and bound variables; a sentence is a formula without free variables. The unfolding of a fixpoint formula $\sigma x \varphi$ is the formula $\varphi[\sigma x \varphi/x]$ that we obtain by substituting $\sigma x \varphi$ for x in φ ; here we ensure that this substitution never causes variable capture (so that no renaming of variables is needed). The *closure* $Clos(\Gamma)$ of

a set Γ of formulas is defined as the smallest set of formulas that is closed under taking boolean subformulas, modal subformulas, unfoldings of fixpoint formulas, and single negations. It is well-known that for every formula φ the closure $\text{Clos}(\{\varphi\})$ is a finite set.

The semantics of formulas is given in terms of *models* $M = (W, R, V)$, where W is any set, R provides for every $a \in \text{Act}$ a relation $R_a \subseteq W \times W$ such that $R_{\bar{a}}$ is the converse $R_{\bar{a}} = \{(v, w) \mid (w, v) \in R_a\}$ of R_a for all $a \in \text{Act}$, and $V : \text{Prop} \rightarrow \mathcal{P}W$ is any function. The elements of W are called *worlds*, the relation R_a is the *accessibility relation* for a and the function V is the *valuation function*.

The semantic clauses for the two-way μ -calculus are completely standard. The boolean and modal connectives are interpreted as in modal logic, where the relation R_a is used for the modality $\langle a \rangle$. The semantic value of the fixpoint formulas $\mu x \varphi$ and $\nu x \varphi$ are the least and greatest fixpoints of the monotone map that describes the interpretation of φ as depending on an interpretation of the variable x . A precise formulation of the semantic clauses is given in the following subsection. The two-way μ -calculus introduced here is a fragment of the language of annotated formulas that is discussed in the following subsection.

2.1 Annotated formulas

The proof systems that we shall introduce here admit formulas with quantified versions of the fixpoint operators, involving a countable set OV of ordinal variables (denoted $\kappa, \lambda, \kappa_0, \dots$). Since we will also allow quantifiers over these ordinal variables, the formulas that we work with will be of the following form:

$$\begin{aligned} \varphi := & p \mid \bar{p} \mid x \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid [a]\varphi \mid \langle a \rangle \varphi \mid \mu x \varphi \mid \nu x \varphi \mid \\ & \mu x^\kappa \varphi \mid \nu x^\kappa \varphi \mid \forall \kappa < \lambda \varphi \mid \exists \kappa < \lambda \varphi \end{aligned}$$

A formula which does not contain any ordinal variables, i.e., a formula of the two-way modal μ -calculus, is called *plain*. The *underlying plain formula* of a formula φ , denoted $u(\varphi)$, is the plain formula obtained from erasing all ordinal annotations and quantifiers from φ :

$$\begin{aligned} u(x) &= x & u(\varphi \wedge \psi) &= u(\varphi) \wedge u(\psi) & u([a]\varphi) &= [a]u(\varphi) \\ u(p) &= p & u(\varphi \vee \psi) &= u(\varphi) \vee u(\psi) & u(\langle a \rangle \varphi) &= \langle a \rangle u(\varphi) \\ u(\eta x^\kappa \varphi) &= \eta x u(\varphi) & u(\forall \lambda < \kappa \varphi) &= u(\varphi) \\ u(\eta x \varphi) &= \eta x u(\varphi) & u(\exists \lambda < \kappa \varphi) &= u(\varphi) \end{aligned}$$

The semantics of this language can be defined as follows. If $f : \mathcal{P}W \rightarrow \mathcal{P}W$ is a monotone function on the powerset of W we identify two ways of iterating f along ordinals: $f_{\top}^\kappa \in \mathcal{P}W$ denotes the result of iterating f κ -many times on the starting from W , and $f_{\perp}^\kappa \in \mathcal{P}W$ the κ -th iterant of f starting from \emptyset :

$$f_{\top}^\kappa = \bigcap_{\xi < \kappa} f(f_{\top}^\xi) \qquad f_{\perp}^\kappa = \bigcup_{\xi < \kappa} f(f_{\perp}^\xi)$$

Note, $f_{\top}^0 = W$ and $f_{\perp}^0 = \emptyset$. Given a model $M = (W, R, V)$, an ordinal assignment is a map o assigning an ordinal $o(\kappa)$ to each ordinal variable κ . Then the meaning $\llbracket \varphi \rrbracket_M^o$ of a formula φ in this model and under this assignment is inductively defined as follows. We write $\llbracket \lambda x. \varphi \rrbracket$ to express the monotone map $Z \mapsto \llbracket \varphi \rrbracket_{M[x \mapsto Z]}^o$ on $\mathcal{P}W$.

- For a propositional variable p , $\llbracket p \rrbracket_M^o = V(p)$.
- Standard clauses for booleans and modalities.
- $\llbracket \mu x^\kappa. \varphi \rrbracket_M^o$ is the $o(\kappa)$ -th iterant of $\llbracket \lambda x. \varphi \rrbracket$ on \emptyset , namely $\llbracket \mu x^\kappa. \varphi \rrbracket_M^o = \llbracket \lambda x. \varphi \rrbracket_{\perp}^{o(\kappa)}$.
- $\llbracket \nu x^\kappa. \varphi \rrbracket_M^o$ is the $o(\kappa)$ -th iterant of $\llbracket \lambda x. \varphi \rrbracket$ on W , namely $\llbracket \nu x^\kappa. \varphi \rrbracket_M^o = \llbracket \lambda x. \varphi \rrbracket_{\top}^{o(\kappa)}$.
- $\llbracket \mu x. \varphi \rrbracket_M^o$ is the least fixpoint $\llbracket \lambda x. \varphi \rrbracket$, namely $\llbracket \mu x. \varphi \rrbracket_M^o = \bigcup_{\xi} \llbracket \lambda x. \varphi \rrbracket_{\perp}^{\xi}$.
- $\llbracket \nu x. \varphi \rrbracket_M^o$ is the greatest fixpoint of $\llbracket \lambda x. \varphi \rrbracket$, namely $\llbracket \nu x. \varphi \rrbracket_M^o = \bigcap_{\xi} \llbracket \lambda x. \varphi \rrbracket_{\top}^{\xi}$.
- $\llbracket \exists \lambda < \kappa. \varphi \rrbracket_M^o = \{u \in W \mid \exists \xi < o(\kappa) : u \in \llbracket \varphi \rrbracket_M^{o[\lambda \mapsto \xi]}\}$.
- $\llbracket \forall \lambda < \kappa. \varphi \rrbracket_M^o = \{u \in W \mid \forall \xi < o(\kappa) : u \in \llbracket \varphi \rrbracket_M^{o[\lambda \mapsto \xi]}\}$.

We write $M, u \models_o \varphi$ if $u \in \llbracket \varphi \rrbracket_M^o$. If φ is a plain formula we may write simply $M, u \models \varphi$.

We think of *negation* as an operation on sentences, extending the involution on proposition letters, and determined by connective duality. Inductively we define the operation $\varphi \mapsto \bar{\varphi}$ for all formulas:

$$\begin{aligned} \bar{x} &= x & \overline{\varphi \wedge \psi} &= \bar{\varphi} \vee \bar{\psi} & \overline{\langle a \rangle \varphi} &= \langle a \rangle \bar{\varphi} & \overline{\mu x^\kappa \varphi} &= \nu x^\kappa \bar{\varphi} & \overline{\forall \lambda < \kappa \varphi} &= \exists \lambda < \kappa \bar{\varphi} \\ \overline{\varphi \vee \psi} &= \bar{\varphi} \wedge \bar{\psi} & \overline{\langle a \rangle \varphi} &= \langle a \rangle \bar{\varphi} & \overline{\nu x^\kappa \varphi} &= \mu x^\kappa \bar{\varphi} & \overline{\exists \lambda < \kappa \varphi} &= \forall \lambda < \kappa \bar{\varphi} \end{aligned}$$

It is routine to show that on the set of sentences this operation indeed behaves as classical negation. Furthermore, observe that this operation is an involution on the set of formulas, and that the negation of a plain formula is plain.

2.2 Subsumption, well-annotated formulas and active variables

In this subsection we define some notions that are not needed in the definition of the proof systems but play a key role in our reasoning about the properties of the proof system.

The *subsumption order* $<_{\rho}$ associated with a formula ρ is defined as the smallest preorder on $\text{Var}(\rho)$ such that $x <_{\rho} y$ if ρ has a subformula $\sigma y \psi$ of which x is a free variable. Observe that the subsumption order of a fixpoint formula is identical to that of its unfolding. By taking, if needed, alphabetic

variants (i.e., renaming the bound variables in ρ) we may always assume that the subsumption order of a given formula is a strict partial order. It may occasionally be convenient to make the following assumption which is possible without loss of generality.

Convention 2.1. In some parts of this paper we will restrict attention to a fixed finite set Γ_0 of plain formulas in which distinct occurrences of fixed point quantifiers are associated with distinct variables, and its closure. We then may assume a strict partial order $<$ on the set of fixpoint variables occurring in Γ_0 which is such that $< \supseteq <_{\varphi} \cap (\text{Var}(\varphi) \times \text{Var}(\varphi))$, for all $\varphi \in \text{Clos}(\Gamma_0)$. The order $<$ will be referred to as the subsumption order. If $x < y$ we refer to x as being *higher ranked* than y .

Definition 2.2. An *annotation* is a partial function from fixed point variables to ordinal variables. Let x_0, x_1, \dots enumerate the fixed point variables in decreasing order with respect to subsumption. Given an annotation o and $n \in \omega$, $o \upharpoonright n$ is the restriction of o to the domain $\{x_i \mid i < n\}$. Given a plain formula φ and annotation $o: \text{FV} \rightarrow \text{OV}$ define a (nonplain) formula φ^o as follows:

$$\begin{aligned} x^o &= x & (\psi \wedge \theta)^o &= \psi^o \wedge \theta^o & ([a]\varphi)^o &= [a]\varphi^o \\ p^o &= p & (\psi \vee \theta)^o &= \psi^o \vee \theta^o & (\langle a \rangle \varphi)^o &= \langle a \rangle \varphi^o \\ (\eta x_i \psi)^o &= \begin{cases} \eta x_i^{o(x_i)} \psi^{o \upharpoonright i}, & \text{if } x_i \in \text{dom } o, \\ \eta x_i \psi^{o \upharpoonright i}, & \text{otherwise.} \end{cases} \end{aligned}$$

Definition 2.3. A formula φ is *well-annotated* if there exists an annotation o such that $\varphi = u(\varphi)^o$. The annotation o satisfying this equation with smallest domain is named o_{φ} . φ is *positively annotated* if it is well-annotated and $\text{dom}(o_{\varphi})$ consists only of ν -fixed point variables of φ . The negation of a positively annotated formula is said to be *negatively annotated*.

Note that plain formulas are positively annotated, and that well-annotated formulas do not contain quantifiers.

Definition 2.4. Given a set of formulas Γ , we say that an ordinal variable κ is *active* in Γ if κ occurs free in some positively annotated formula in Γ . The set of active variables in Γ is denoted as $\text{Act}(\Gamma)$.

2.3 Game semantics

In this section we briefly review the game-theoretic semantics for the two-way μ -calculus; this will prove to be a useful approach in the completeness argument further on.

The *evaluation game* $\mathcal{E}(M, \varphi)$ of a formula φ on a model $M = (W, R, V)$ is an infinite board game, the players of which we shall call Verifier and Falsifier. The positions of the game are all pairs of the form (w, ψ) , where $w \in W$ and $\psi \in \text{Clos}(\varphi)$. The player to move at a given position and the moves at his or her disposal are listed in Table 1.

Position	Player	Moves
(w, p) for $w \in V(p)$	Falsifier	\emptyset
(w, p) for $w \notin V(p)$	Verifier	\emptyset
(w, \bar{p}) for $w \in V(p)$	Verifier	\emptyset
(w, \bar{p}) for $w \notin V(p)$	Falsifier	\emptyset
$(w, \varphi \wedge \psi)$	Falsifier	$\{(w, \varphi), (w, \psi)\}$
$(w, \varphi \vee \psi)$	Verifier	$\{(w, \varphi), (w, \psi)\}$
$(w, \mu x \varphi)$	–	$\{(w, \varphi[\mu x \varphi/x])\}$
$(w, \nu x \varphi)$	–	$\{(w, \varphi[\nu x \varphi/x])\}$
$(w, \langle a \rangle \varphi)$	Verifier	$\{(v, \varphi) \mid wR_a v\}$
$(w, [a] \varphi)$	Falsifier	$\{(v, \varphi) \mid wR_a v\}$

Table 1: The evaluation game $\mathcal{E}(M, \varphi)$

Any match or play of this game consists of a (finite or infinite) sequence $(w_n, \varphi_n)_{n < \kappa}$ of positions (with $\kappa \leq \omega$). A finite match, i.e., with $\kappa < \omega$, is *won* by a player if it is their opponent who is supposed to move at the last position $(w_{\kappa-1}, \varphi_{\kappa-1})$, while there is no move available.

To determine the winner of an infinite match $(w_n, \varphi_n)_{n < \omega}$ we observe that the induced sequence $(\varphi_n)_{n < \omega}$ of formulas is an infinite *trace*, that is: for every $i < \omega$, either φ_i is a fixpoint formula and φ_{i+1} is its unfolding, or else φ_{i+1} is a direct (modal or boolean) subformula of φ_i . It is well known that for every infinite trace $\tau = (\varphi_n)_{n < \omega}$ there is a unique formula that occurs infinitely often on τ and is a subformula of φ_n for cofinitely many n . We will call this formula, which must be a fixpoint formula, the *most significant formula* of τ , and we declare Verifier (Falsifier) to be the winner of an infinite match $(w_n, \varphi_n)_{n < \omega}$ if the most significant formula of the induced trace $(\varphi_n)_{n < \omega}$ is a ν -formula (a μ -formula, respectively). It is well known that this winning condition can be formulated as a parity condition and that consequently the game $\mathcal{E}(M, \varphi)$ has positional determinacy.

Theorem 2.5 (Adequacy of evaluation games). For any plain formula φ , model M and world w , we have $M, w \models \varphi$ if and only if the position (w, φ) is winning for Verifier in $\mathcal{E}(M, \varphi)$.

For a proof of the theorem see, e.g., [5].

3 Proof systems

In this section we first define the finitary, cyclic proof system that is the subject of this paper and then discuss infinitary, non-wellfounded proofs that are needed for our completeness argument.

3.1 Sequents and constraints

The sequents in our proof system contain a constraint that describes the relative size of ordinal variables and keeps track of the order in which they are introduced.

Definition 3.1. A *constraint* is a tuple $\mathcal{O} = (O, <, \triangleleft)$ where

1. O is a finite set of ordinal variable symbols, called the *domain*,
2. $<$ is an irreflexive, transitive and upwards linear ordering $<$ (so $(O, >)$ is a finite forest), called the *descendant* relation,
3. \triangleleft is a total linear order on \mathcal{O} , called the *age* relation, consistent with the ancestor relation: $\kappa < \lambda$ implies $\lambda \triangleleft \kappa$ for all $\kappa, \lambda \in O$.

Given a constraint $\mathcal{O} = (O, <_{\mathcal{O}}, \triangleleft_{\mathcal{O}})$, we write $\text{OV}(\mathcal{O})$ for O , the set of ordinal variables appearing in \mathcal{O} . When there is no risk of confusion, we identify the constraint \mathcal{O} with the set $\text{OV}(\mathcal{O})$, writing $\kappa \in \mathcal{O}$ rather than the formally precise $\kappa \in \text{OV}(\mathcal{O})$. The reflexive closure of $<_{\mathcal{O}}$ is denoted $\leq_{\mathcal{O}}$.

Definition 3.2. A *sequent* is an expression $\mathcal{O} : \Gamma$ where \mathcal{O} is a constraint and Γ is a finite set of formulas whose free ordinal variables are elements of \mathcal{O} . We sometimes write a sequent $\mathcal{O} : \Gamma$ as just Γ , denoting \mathcal{O} by $\mathcal{O}(\Gamma)$.

Given a constraint \mathcal{O} and $\kappa \in \mathcal{O}$, a descendant of κ is any $\lambda \in \mathcal{O}$ such that $\lambda <_{\mathcal{O}} \kappa$, in which case κ is called an ancestor of λ (in \mathcal{O}). We say λ is a child of κ , or that κ is the parent of λ , if $\lambda <_{\mathcal{O}} \kappa$ and there is no $\rho \in \mathcal{O}$ such that $\lambda <_{\mathcal{O}} \rho$ and $\rho <_{\mathcal{O}} \kappa$. Every $\kappa \in \mathcal{O}$ has at most one parent, but may have multiple children. If $\kappa \triangleleft_{\mathcal{O}} \lambda$ we say that κ is *older than* λ (relative to \mathcal{O}).

A substitution on ordinal variables is simply a map $\sigma : \text{OV} \rightarrow \text{OV}$. With respect to a constraint \mathcal{O} we call a substitution σ increasing if $\lambda \leq_{\mathcal{O}} \sigma(\lambda)$ for all λ , and decreasing if $\sigma(\lambda) \leq_{\mathcal{O}} \lambda$ for all λ .

For later use we introduce two auxiliary relations on a constraint \mathcal{O} . First of all, we say that λ is *to the left of* ρ if λ and ρ are incomparable with respect to $\leq_{\mathcal{O}}$ and $\lambda' \triangleleft_{\mathcal{O}} \rho'$, where λ' is the $<$ -greatest ancestor of λ that is not an ancestor of ρ and ρ' is the $<$ -greatest ancestor of ρ that is not an ancestor of λ . We then define $\lambda \prec_{\mathcal{O}} \rho$ if $\lambda <_{\mathcal{O}} \rho$ or λ is to the left of ρ , and we sometimes denote $\lambda \prec_{\mathcal{O}} \rho$ as $\lambda \prec_{\mathcal{O}} \rho$. As we will see later, $\prec_{\mathcal{O}}$ is in fact a strict linear order.

Here is an example to illustrate the different orders on ordinal variables and how they are related. Consider a constraint \mathcal{O} containing seven ordinal variables $\kappa_0, \dots, \kappa_6$, where $\kappa_0 \triangleleft_{\mathcal{O}} \dots \triangleleft_{\mathcal{O}} \kappa_6$ and $<_{\mathcal{O}}$ is the transitive closure of:

$$\{(\kappa_3, \kappa_1), (\kappa_4, \kappa_1), (\kappa_5, \kappa_2), (\kappa_6, \kappa_2), (\kappa_1, \kappa_0), (\kappa_2, \kappa_0)\}.$$

Represented as a tree the $<_{\mathcal{O}}$ -relation is shown in Figure 1. The figure is drawn so that, with the age relation as specified, the “left-of” relation between $<_{\mathcal{O}}$ -incomparable variables can be read off directly from the diagram. So κ_3 is to




Figure 1: A sample constraint.

the left of $\kappa_2, \kappa_4, \kappa_5, \kappa_6$, while κ_1 is to the left of $\kappa_2, \kappa_5, \kappa_6$, etc. The $\prec_{\mathcal{O}}$ relation is thus the strict linear order given by:

$$\kappa_3 \prec_{\mathcal{O}} \kappa_4 \prec_{\mathcal{O}} \kappa_1 \prec_{\mathcal{O}} \kappa_5 \prec_{\mathcal{O}} \kappa_6 \prec_{\mathcal{O}} \kappa_2 \prec_{\mathcal{O}} \kappa_0$$

To simplify notation we introduce a special symbol \star and write $o(x) = \star$ for an annotation o if $x \notin \text{dom}(o)$. Given an ordinal constraint \mathcal{O} , we extend the order $\prec_{\mathcal{O}}$ to $\mathcal{O} \cup \{\star\}$ by setting $\kappa \prec_{\mathcal{O}} \star$ for every ordinal variable κ in \mathcal{O} . Note that $\prec_{\mathcal{O}}$, with its definition extended to incorporate \star , is still a linear order over $\mathcal{O} \cup \{\star\}$ and $\kappa \prec_{\mathcal{O}} \star$ for every ordinal variable κ in \mathcal{O} .

The semantics of sequents is given as follows.

Definition 3.3. Let $M = (W, R, V)$ be a model, A sequent $\mathcal{O} : \Gamma$ holds in M if for all ordinal assignments o such that $o(\kappa) < o(\lambda)$ whenever $\kappa \prec_{\mathcal{O}} \lambda$ we have that $\bigcup\{\llbracket \varphi \rrbracket_M^o \mid \varphi \in \Gamma\} = W$. We say that an ordinal assignment o refutes $\mathcal{O} : \Gamma$ in M if $o(\kappa) < o(\lambda)$ whenever $\kappa \prec_{\mathcal{O}} \lambda$, but $\bigcup\{\llbracket \varphi \rrbracket_M^o \mid \varphi \in \Gamma\} \neq W$.

3.2 Rules and derivations

The sequent calculus we introduce makes use of three operations on constraints: The operation denoted $\mathcal{O} + \lambda$ extends \mathcal{O} by a fresh variable λ as the youngest element and makes no change to the descendant relation. As a variation, in $\mathcal{O} +_{\kappa} \lambda$ the variable λ is also added as a child of κ . That is, for $\mathcal{O} = (O, <, \triangleleft)$, $\lambda \in O$ and $\kappa \in O$:

$$\begin{aligned} \mathcal{O} + \lambda &= (O \cup \{\lambda\}, <, \triangleleft \cup \{(\rho, \lambda) \mid \rho \in O\}) \\ \mathcal{O} +_{\kappa} \lambda &= (O \cup \{\lambda\}, < \cup \{(\lambda, \kappa') \mid \kappa \leq \kappa'\}, \triangleleft \cup \{(\rho, \lambda) \mid \rho \in O\}) \end{aligned}$$

In both the above constructions it is a requirement that λ does not occur already in \mathcal{O} .

The third construction is the restriction of a constraint to a set of ordinal variables. Given \mathcal{O} as above and $V \subseteq O$, we define $\mathcal{O} \setminus V$ to be the constraint

$$\mathcal{O} \setminus V = (O', < \cap (O' \times O'), \triangleleft \cap (O' \times O')) \quad \text{where } O' = O \setminus V$$

Using these operations we can define the inference rules of our sequent calculus. These are presented in Table 2, where we use the expression $\mathcal{O}(\lambda < \kappa)$ to

$\frac{}{\mathcal{O} : p, \bar{p}} \top$	$\frac{\mathcal{O} : \Gamma, \varphi \quad \mathcal{O} : \Gamma, \psi}{\mathcal{O} : \Gamma, \varphi \wedge \psi} \wedge$
$\frac{\mathcal{O} + \kappa : \Gamma, \nu x^\kappa \psi}{\mathcal{O} : \Gamma, \nu x \psi} \kappa : x$	$\frac{\mathcal{O} : \Gamma, \varphi, \psi}{\mathcal{O} : \Gamma, \varphi \vee \psi} \vee$
$\frac{\mathcal{O} : \Gamma, \varphi[\eta x \varphi/x]}{\mathcal{O} : \Gamma, \eta x \varphi} \eta \in \{\mu, \nu\}$	$\frac{\mathcal{O} : [\check{a}] \vee \Gamma, \varphi, \Psi}{\mathcal{O} : \Gamma, [a] \varphi, \langle a \rangle \Psi} \text{mod}$
$\frac{\mathcal{O} + \kappa \lambda : \Gamma, \varphi[\nu x^\lambda \varphi/x]}{\mathcal{O} : \Gamma, \nu x^\kappa \varphi} \nu(\kappa)$	$\frac{\mathcal{O}(\lambda < \kappa) : \Gamma, \varphi[\mu x^\lambda \varphi/x]}{\mathcal{O}(\lambda < \kappa) : \Gamma, \mu x^\kappa \varphi} \mu(\kappa)$
$\frac{\mathcal{O} + \kappa \lambda : \Gamma, \varphi[\lambda/\rho]}{\mathcal{O} : \Gamma, \forall \rho < \kappa : \varphi} \forall$	$\frac{\mathcal{O}(\lambda < \kappa) : \Gamma, \varphi[\lambda/\rho]}{\mathcal{O}(\lambda < \kappa) : \Gamma, \exists \rho < \kappa : \varphi} \exists$
$\frac{\mathcal{O} \setminus V : \Gamma}{\mathcal{O} : \Gamma} \text{lw}$	$\frac{\mathcal{O} : \Gamma}{\mathcal{O} : \Gamma, \varphi} \text{rw}$
$\frac{\mathcal{O} : \Gamma, \varphi \quad \mathcal{O} : \Gamma, \bar{\varphi}}{\mathcal{O} : \Gamma} \text{cut}$	

Table 2: Rules of sequent calculus

refer to a constraint \mathcal{O} such that $\lambda <_{\mathcal{O}} \kappa$. A tree constructed by applications of these rules, and labelled with sequents and names of the rules applied, will be called a *derivation*. We shall reserve the term *proof* for derivations satisfying one of several syntactic criteria guaranteeing validity, defined in the following sections. Given a derivation Π and a vertex m in Π , the sequent labelling m is denoted $\Pi(m)$.

Note that these rules feature no side conditions, besides the obvious constraint that all sequents involved in a rule instance must be bona fide *sequents*. This means that the \forall -rule satisfies the usual eigenvariable condition: the variable λ appearing in the premise cannot occur in the conclusion. If it had occurred in some formula, then by definition of a sequent it should also appear in the constraint, meaning that the constraint $\mathcal{O} + \kappa \lambda$ appearing in the premise would not be well-defined. A similar eigenvariable constraint holds for the rule $\nu(\kappa)$, and for the same reason. In the same vein, observe that in an application of lw , for the premise to be a well-formed sequent, no ordinal variable in the set V may occur free in Γ .

The reader may have noticed that the cut rule also has a hidden restriction due to the definition of a sequent: the cut formula will never contain an ordinal variable that does not occur in the conclusion. This is not as restrictive as it may at first seem. Cuts can still be used to introduce new ordinal variables in proof search, since the cut formulas may contain quantifiers which can then be

instantiated, such as in the following:

$$\frac{\frac{\vdots}{\mathcal{O} +_{\kappa} \lambda : \Gamma, \varphi[\lambda/\rho]} \quad \frac{\vdots}{\mathcal{O} : \Gamma, \exists \rho < \kappa : \bar{\varphi}}}{\mathcal{O} : \Gamma} \text{ cut}$$

The calculus does not feature a rule for introducing an unapproximated least fixed point from an approximant, for instance,

$$\frac{\mathcal{O} : \Gamma, \mu x^{\kappa} . \varphi}{\mathcal{O} : \Gamma, \mu x . \varphi} \mu$$

Although sound, the above rule is not necessary for a complete sequent calculus. Without such a rule, derivations have the property that explicitly approximated least fixed points only arise from premises of the cut rule.

Proposition 3.4. Let R be an instance of any of the derivation rules *except cut*, with conclusion $\mathcal{O} : \Gamma$. If every formula in Γ is well-annotated, then every formula appearing in a premise of R is well-annotated.

Proof. By inspection of the proof rules. □

Further on we shall frequently need to employ the following *minimisation* rule.

$$\frac{\mathcal{O} : \Gamma, \exists \lambda < \kappa : \bar{\varphi}[\mu x^{\lambda} \bar{\psi}/z] \quad \mathcal{O} +_{\kappa} \kappa_0 : \Gamma, \varphi[\nu x^{\kappa_0} \psi/z]}{\mathcal{O} : \Gamma} \text{ min}$$

This rule is easily derived via a cut:

$$\frac{\mathcal{O} : \Gamma, \exists \lambda < \kappa : \bar{\varphi}[\mu x^{\lambda} . \bar{\psi}/z] \quad \frac{\mathcal{O} +_{\kappa} \kappa_0 : \Gamma, \varphi[\nu x^{\kappa_0} \psi/z]}{\mathcal{O} : \Gamma, \forall \lambda < \kappa : \varphi[\nu x^{\lambda} \psi/z]} \vee}{\mathcal{O} : \Gamma} \text{ cut}$$

Intuitively, we can view the minimisation rule as expressing a case distinction for how we may construct a counter-model to Γ , in particular when Γ contains the formula $\varphi[\nu x^{\kappa} \psi/z]$: either we decide that κ is the smallest ordinal approximant of the fixpoint $\nu x \psi$ for which the formula fails to hold, or we introduce a name κ_0 for smaller ordinal approximant such that $\varphi[\nu x^{\kappa_0} \psi/z]$ fails to hold. In the sequel we shall assume that the minimisation rule is a rule of our proof systems.

3.3 Cyclic proofs

We now present our main notion of a valid proof, which is based on ‘cyclic’ derivations, which can be seen as finite representations of infinite but regular, non-wellfounded proof trees. It is straightforward to show that having no restrictions on allowed cycles trivialises the calculus, allowing any sequent to be

derived by a cyclic proof. The notion of cyclic proof therefore comes equipped with a condition on cycles, called the *correctness criterion*, ensuring only valid sequents are derivable. First, we define more precisely what is meant by a “cycle”.

Definition 3.5. A *derivation with back-edges* is a pair (Π, c) where Π is a derivation and c is a partial function from the leaves of Π to inner nodes of Π , called the *companion function*, such that that for every leaf $l \in \text{dom } c$,

- l and $c(l)$ are labelled by the same sequent, i.e., $\Pi(l) = \Pi(c(l))$, and
- l is an ancestor of $c(l)$.

We call $c(l)$ the *companion* of l .

Note that our requirement that a leaf l is labelled by the same sequent as its companion $c(l)$ entails more than just the same formulas occurring in the leaf and companion. It also implies that the *constraints* are identical, which means in particular that the relative *ages* of ordinal variables in leaf and companion are the same. This plays a dual role. On one hand, it facilitates proof search, and will be used extensively in our completeness argument. For this purpose however, it is not clear that the age ordering needs to be incorporated directly into the definition of a valid proof. We shall see that, in fact, the age ordering also plays a crucial role in ensuring that the proof system is *sound*. It is for this reason that it is explicitly part of the proof system, rather than just an auxiliary technical device used to prove completeness.

It remains to formulate a criterion for distinguishing valid cyclic proofs from invalid derivations. For our purposes, the correctness criterion can be formulated as a particular instance of the structural rule *lw* occurring on the path between companion and leaf. We thus begin with isolating these special instances of weakening, henceforth called *resets*.

Definition 3.6. The *reset rule* is any instance of constraint weakening of the form

$$\frac{\mathcal{O} \setminus K : \Gamma}{\mathcal{O} : \Gamma} \text{reset}(\kappa)$$

such that

1. $\kappa \in \mathcal{O}$ and κ does not occur in Γ ;
2. K is the set of children of κ in \mathcal{O} .

Note, since derivations are assumed to be well formed, the children of κ , namely the variables in K above, do not occur in Γ .

Definition 3.7. Let (Π, c) be a derivation with back-edges and $l \in \text{dom } c$ be labelled by a sequent $\mathcal{O} : \Gamma$. An ordinal variable $\kappa \in \mathcal{O}$ is a *reset variable for l* if

1. κ occurs in the constraint at every vertex on the path from $c(l)$ to l ;

$$\begin{array}{c}
\vdots \\
\frac{\emptyset : \bar{Y}, Y}{\emptyset : [a]\bar{Y}, \langle a \rangle Y} \text{ mod} \\
\frac{\emptyset : [a]\bar{Y}, Y}{\emptyset : [a]\bar{Y}, Y} \text{ rw, } \vee, \mu \\
\frac{\emptyset : [a]\bar{Y}, Y}{\kappa\kappa_1 : [a]\bar{Y}, Y} \text{ lw} \\
\hline
\frac{\kappa\kappa_1 : [a]\bar{Y}, X^{\kappa_1}}{\kappa\kappa_1 : [a]\bar{Y}, X^{\kappa_1}} \text{ cut} \\
\frac{\kappa\kappa_1 : [a]\bar{Y}, X^{\kappa_1}}{\kappa\kappa_1 : \bar{Y}, [\check{a}]X^{\kappa_1}} \text{ mod} \\
\frac{\kappa\kappa_1 : \bar{Y}, [\check{a}]X^{\kappa_1}}{\kappa\kappa_0\kappa_1 : \bar{Y}, [\check{a}]X^{\kappa_1}} \text{ reset}(\kappa) \\
\vdots \\
\frac{\kappa\kappa_0\kappa_1 : \bar{Y}, Y}{\kappa\kappa_0 : \bar{Y}, X^{\kappa_0}} \wedge, \nu(\kappa_1) \\
\hline
\frac{\kappa\kappa_0 : [a]\bar{Y}, Y}{\kappa\kappa_0 : \bar{Y}, X^{\kappa_0}} \text{ cut} \\
\frac{\kappa\kappa_0 : [a]\bar{Y}, X^{\kappa_0}}{\kappa\kappa_0 : \bar{Y}, [\check{a}]X^{\kappa_0}} \text{ mod} \\
\frac{\dagger \kappa\kappa_0 : \bar{Y}, [\check{a}]X^{\kappa_0}}{\kappa\kappa_0 : \bar{Y}, [\check{a}]X^{\kappa_0} \wedge Y} \nu(\kappa) \\
\hline
\frac{\emptyset : \bar{p}, p}{\emptyset : \bar{p}, Y} \text{ rw, } \mu, \vee \\
\frac{\emptyset : \bar{p}, Y}{\emptyset : \bar{p}, X} \text{ cut} \\
\hline
\frac{\emptyset : \bar{p}, X}{\emptyset : \bar{p} \vee X} \vee
\end{array}$$

Figure 3: Cyclic proof of the formula $\bar{p} \vee \nu x. [\check{a}]x \wedge \mu y. p \vee \langle a \rangle y$. Subproofs of the sequent $\emptyset : \bar{Y}, Y$ are as in Figure 2 and omitted. The relation of non-axiomatic leaves to companions is denoted by \dagger . The proof employs the same abbreviations as Figure 2 with, in addition, X and X^κ denoting formulas $\nu x. [\check{a}]x \wedge Y$ and $\nu x^\kappa. [\check{a}]x \wedge Y$ respectively.

a single non-axiomatic leaf whose companion is the inner sequent marked as \dagger , though each of the five omitted subproofs of \bar{Y}, Y also involves an internal leaf-companion pair as per Figure 2. As the variable κ occurs in the constraint of every sequent along the connecting path, this leaf is successful. Note that in both examples, the repeating cycles require two applications of the reset rule. The reason for this is that we require leaves and companions to be identical as sequents. This is just a technical convenience; one could instead formulate a weaker condition on leaf-companion pairs, requiring identity only up to a renaming of ordinal variables. Some care is needed to ensure soundness however; the precise formulation of such a condition was worked out in [1].

For certain proofs in the following, it will be convenient to consider a conditional notion of cyclic proof, i.e., cyclic proofs from assumptions. For this notion of proof it is important to restrict the application of structural rules on paths from the conclusion to assumptions.

Definition 3.9. Let $\mathcal{S} \cup \{\mathcal{O} : \Gamma\}$ be a set of sequents. A (cyclic) proof of $\mathcal{O} : \Gamma$ from assumptions \mathcal{S} is a finite derivation with back-edges (Π, c) such that every leaf $l \in \text{dom } c$ is successful and every leaf $l \in \text{dom } \Pi$ is either an axiom, or else (i)

it is labelled by a sequent in \mathcal{S} , and (ii) there is no application of lw (including reset rules) on the path from the root of Π to l .

We are now ready to state our main result:

Theorem 3.10 (Soundness and Completeness). Let ρ be any plain formula of the two-way μ -calculus. Then ρ is valid if, and only if, it has a cyclic proof.

The remainder of the paper is devoted to proving Theorem 3.10. Section 4 culminates in the proof of soundness. Completeness is the subject of Sections 5 and 6.

The proof of completeness of the cyclic proof system relies on completeness of non-wellfounded proofs, where the correctness condition on infinite branches is stipulated in terms of infinite descending chains of ordinal variables in the controls of sequents along infinite paths. In their most general formulation, completeness for such proofs is easy to prove, but not very helpful towards proving the main theorem since non-wellfounded proofs can be highly non-regular without any bound on the size of sequents in a proof tree.

Therefore, we isolate a sub-class of non-wellfounded proofs, called *slim proofs*, which, although not finitely presentable in general, are in an important sense closer to the finitary notion of provability: in slim proofs, the number of formulas that can appear in sequents is bounded. Constraints, however, can grow without bound. Finitising the constraints is the second step towards obtaining a cyclic proof for a valid sequent. Completeness with respect to slim proofs is established via a two-player game that we call the *mosaic game*, in which one player (Prover) tries to construct a proof of the root formula, and the opposing player (Refuter) attempts to build a counter-model by selecting certain ‘saturated’ sequents called *tiles*.

With completeness of slim proofs in place, the next step is to insert uses of the reset rule in order to bound the size of constraints. This transformation alters the correctness condition on infinite paths from an infinite descent condition to an infinitary version of the reset condition from cyclic proofs: every infinite branch features some ordinal variable that is reset infinitely often. The final step of the completeness argument is to show that any non-wellfounded slim derivation satisfying the infinite reset condition can be pruned to a finite, cyclic proof.

3.4 Non-wellfounded proofs

In the proof of completeness for cyclic derivations we will make extensive use of an intermediate notion of proof based on infinite, or non-wellfounded, derivations. In analogy with the case of cyclic proofs, an infinite derivation will be considered a proof provided every infinite path in the derivation fulfils a syntactic criterion. Later we will consider infinite proofs where correctness is based on an (infinitary) notion of reset variable. For now we introduce the concept of an *infinite descent proof* where the requirement on infinite paths is that the sequence of constraints in the path induces an infinite descending sequence of ordinal variables.

Definition 3.11. Let $P = (\mathcal{O}_i)_{i < \omega}$ be an infinite sequence of constraints.

We say that P has an *infinite $<$ -descending chain of ordinals* if there are an infinite sequence $(\kappa_i)_{i < \omega}$ of ordinal variables and an increasing function $\sigma : \omega \rightarrow \omega$ such that for every i it holds that $\kappa_{i+1} <_{\mathcal{O}_{\sigma(i)}} \kappa_i$ and $\kappa_{i+1} \in \text{OV}(\mathcal{O}_l)$ for all $l \in \{\sigma(i), \sigma(i)+1, \dots, \sigma(i+1)\}$. An *infinite \prec -descending chain of ordinals* is defined analogously.

An infinite derivation is said to be an *infinite descent proof* if every leaf is an axiom, and every infinite path is such that the sequence of constraints through this path has an infinite $<$ -descending chain of ordinal variables.

Definition 3.12. A finite proof tree is said to be a *wellfounded proof* if every leaf is labelled by a sequent of the form $\mathcal{O} : \varphi, \bar{\varphi}$.

4 Soundness

In this section we show how to prove soundness for the cyclic and non-wellfounded variants of the proof system that is defined in the previous section. We start with the soundness proof for the cyclic proof system.

Definition 4.1. A *strongly connected component*, abbreviated *SCC*, in a cyclic proof tree is a set X of nodes which is connected, seen as a subgraph of the proof tree, where the (directed) edge relation is given as the union of the parent-child and the back-edge relation.

It is clear that every strongly connected component X of a cyclic proof tree has a lowest element, i.e. a unique element with the shortest path to the root of the proof tree, which must be the companion node of some leaf. We call this the *root* of X . (Note that we imagine trees as growing upwards, with the root at the bottom, in line with the way that we depict proof trees.)

The following characterization of cyclic proofs will be useful:

Proposition 4.2. Let Π be a cyclic proof tree. Then Π is a valid cyclic proof if, and only if, for every SCC X of Π there is some ordinal variable κ_X that appears in the constraint of each sequent in X , and is reset in at least one vertex of X .

Proof. For right to left, it suffices to note that for every leaf l , the path from the companion of l to l is a strongly connected component of Π .

For left to right, let Π be a valid cyclic proof.

Claim 1. Let X be any SCC of Π . Then there exists an ordinal variable κ_X that belongs to the constraint of every vertex in X , and is equal to the variable κ_l associated with some leaf l .

Proof of claim. We prove this by induction on the number of non-axiom leaves in X . The base case consists of the case where X is a single cycle (comprising all nodes on the path to a leaf l from its companion), and in this case we just take $\kappa_X = \kappa_l$.

Now, suppose that X has more than one non-axiom leaf. Consider the set S of vertices in X that are ancestors of all non-axiom leaves in X . It is clear that any two vertices in this set are comparable with respect to the descendant relation, so we can pick the maximum element s of the set, i.e. $s \in S$ and s is a descendant of every member of S . We call s the *splitting point* of X . Because all proof rules have at most two premises it follows that s has exactly two children, a left child and a right child. Every non-axiom leaf is a descendant of either the left child of s or the right child of s . In the first case we speak of a *left leaf* and in the second case a *right leaf*. Note that some leaf of X must have the root of X as companion, and this leaf must be either a left or right leaf. We may assume without loss of generality that it is a left leaf, the other case is symmetrical. Note that some right leaf must have a companion that is an ancestor of s , since X is strongly connected. Let c denote the lowest ancestor of s (closest to the root of X) that is the companion of some right leaf. Let X_0 be the set of vertices that lie on a path from the root of X to some left leaf, and let X_1 be the set of vertices that lie on a path from c to some right leaf. Then $X_0 \cup X_1 = X$, and both X_0 and X_1 are SCC's with fewer non-axiom leaves than X (note that c is the root of X_1). So by the induction hypothesis, there are variables λ_0, λ_1 such that λ_0 appears in the constraint of every vertex in X_0 and λ_1 appears in the constraint of every vertex in X_1 . In particular, since $c \in X_0 \cap X_1$, both λ_0 and λ_1 appear in the constraint of c . We now consider two cases.

Case 1: λ_0 is older than λ_1 in the constraint of c . By assumption λ_0 appears in the constraint of every vertex in X_0 . We define the *depth* of a right leaf l as the number of leaves passed on the shortest path from the leaf to c . More precisely, if the companion of l is c then l has depth 0. Otherwise the depth of l is the smallest number $k + 1$ such that the companion of l belongs to the path from the companion of l' to l' , where l' is some right leaf of depth k . We prove by induction on the depth of a leaf l that λ_0 belongs to the constraint of every vertex on the path from the companion of l to l , and furthermore is older than λ_1 in every such constraint.

For depth = 0, suppose l has c as companion. If there is some u on the path from c to l in which λ_0 does not appear in the constraint, then it has to be re-introduced later since the label of l is equal to that of c . But since λ_1 is on the constraint of every vertex on the path from c to l , λ_0 can only be re-introduced as a younger variable than λ_1 , and remain so. So the label of l cannot equal that of c after all, contradiction.

For depth = $k + 1$, suppose l has companion v that is on the path from the companion of l' to l' , where l' has depth k . Then again, λ_0 is introduced as a younger variable than λ_1 in the companion v , and we can repeat the same argument.

Case 2: λ_1 is older than λ_0 in c . Then in fact, λ_1 must appear in the constraint of the root of X , since λ_0 belongs to the constraint of every vertex on the path from the root of X to c , and so if λ_1 were introduced on this path it would have to be younger than λ_0 . With this observation in place we can repeat the whole argument from the previous case, but with the root of X in place of c . \square

To finish the proof, take any connected component X of Π , and let κ_X be the variable provided by the Claim. Then $\kappa_X = \kappa_l$ for some leaf l . By definition of a valid cyclic proof, κ_l is reset at least once on the path from the companion of l to l . But this path is contained in X , so $\kappa_X = \kappa_l$ is reset in some vertex in X . \square

Proposition 4.3. If ρ has a cyclic proof then ρ is valid.

Proof. Let Π be a cyclic proof of ρ .

Assume for a contradiction that there is some model M , a world w in M and an ordinal assignment o such that $M, w \not\models_o \rho$. Our strategy is to find an infinite walk through the cyclic proof Π , inducing a series of ordinal assignments, from which we can read off an infinitely descending sequence of ordinals. This contradiction then gives the proposition. We construct the walk by induction as follows. We define a vertex v_n of Π , an element w_n of W and an ordinal assignment o_n by induction on n . We shall maintain the invariant that for each n , o_n refutes the label of v_n (henceforth denoted Γ_n) at w_n . Note that this means that we never reach a leaf labelled by an axiom. For the base case, we define v_0 to be the root of the proof tree labelled $\vdash \rho$, set $w_0 = w$, and set $o_0 = o$. Note that o can be taken to be the empty assignment if ρ does not contain any ordinal variables. The assignment o_0 refutes the sequent $\vdash \rho$ at w_0 by the assumption that $M, w \not\models_o \rho$. The inductive step is by case distinction depending on the rule applied at the vertex v_n . We distinguish between applications of the left weakening rule that correspond to the reset of a variable from other applications of left weakening. These need to be treated separately as they allow us to build the decreasing chain of ordinals.

Case v_n is the conclusion of an application of $\kappa : x$. Here, κ is a fresh variable and the principal formula of the rule is of the form $\nu x \varphi$, which is replaced by $\nu x^\kappa \varphi$. We define o_{n+1} to be the extension of o_n obtained by mapping κ to the closure ordinal of the map $\varphi_M^{o_n}$, and set v_{n+1} to be the premise of v_n and $w_{n+1} = w_n$.

Case v_n is the conclusion of an application of the modal rule. Let the principal formula be $[a]\varphi$. Since o_n refutes the label of Γ_n at w_n , it is clear that there is some a -successor w' at which o_n refutes the premise. So we set $w_{n+1} = w'$, v_{n+1} to be the premise of the rule application and $o_{n+1} = o_n$.

Case v_n is the conclusion of an application of the cut rule. Since o_n refutes Γ_n at w_n , it has to refute one of the premises at w_n also. We pick v_{n+1} to be this premise, and set $w_{n+1} = w_n$, $o_{n+1} = o_n$.

Case v_n is the conclusion of an application of the \wedge -rule. Since o_n refutes Γ_n at w_n , it has to refute one of the premises at w_n also. We pick v_{n+1} to be this premise, and set $w_{n+1} = w_n$, $o_{n+1} = o_n$.

Case v_n is the conclusion of an application of the \vee -rule, η -rule or \exists -rule. We set v_{n+1} to be the premise, and set $w_{n+1} = w_n$ and $o_{n+1} = o_n$. The invariant is easily seen to be maintained.

Case v_n is the conclusion of an application of $\mu(\kappa)$. We define $w_{n+1} = w_n$ and $o_{n+1} = o_n$. By assumption we have that $w_n \notin \llbracket \mu x^\kappa \varphi \rrbracket^{o_n} = f^{o_n(\kappa)}(\emptyset)$,

where f is the monotone map with $f(Z) = \llbracket \varphi \rrbracket_{M[x \mapsto Z]}^{o_n}$. If we now consider any λ such that $o_n(\lambda) < o_n(\kappa)$ we have that $f^{o_n(\lambda)+1}(\emptyset) \subseteq f^{o_n(\kappa)}(\emptyset)$, because $o_n(\lambda) + 1 \leq o_n(\kappa)$. Thus it follows that $w_{n+1} = w_n \notin f^{o_n(\lambda)+1}(\emptyset) = f(f^{o_n(\lambda)}(\emptyset)) = \llbracket \varphi[\mu x^\lambda. \varphi/x] \rrbracket^{o_n}$. Thus the invariant is maintained.

Case v_n is the conclusion of an application of the \forall -rule. Let $\lambda < \kappa$ be the fresh variable introduced. Since o_n refutes Γ_n at w_n there must be some $\xi < o_n(\kappa)$ such that $o_n[\lambda \mapsto \xi]$ refutes the premise at w_n . So we set v_{n+1} to be the premise of v_n , $w_{n+1} = w_n$ and $o_{n+1} = o_n[\lambda \mapsto \xi]$.

Case v_n is the conclusion of an application of the $\nu(\kappa)$ -rule. Let $\lambda < \kappa$ be the fresh variable introduced. Since o_n refutes Γ_n at w_n we have that $w_n \notin \llbracket \nu x^\kappa. \varphi \rrbracket^{o_n}$. If we write f for the monotone map with $f(Z) = \llbracket \varphi \rrbracket_{M[x \mapsto Z]}^{o_n}$ this means that $w_n \notin f^{o_n(\kappa)}(W)$. By the definition of $f^{o_n(\kappa)}$ we have that

$$f^{o_n(\kappa)}(W) = \bigcap_{\zeta < \kappa} f(f^{o_n(\zeta)}(W)).$$

Thus there exists some ordinal $\zeta < o_n(\kappa)$ such that $w_n \notin f(f^\zeta(W))$. We set $w_{n+1} = w_n$ and $o_{n+1} = o_n[\lambda \mapsto \zeta]$. This refutes the premise because $w_{n+1} = w_n \notin f(f^\zeta(W)) = \llbracket \varphi[\nu x^\lambda. \varphi/x] \rrbracket^{o_{n+1}}$.

Case v_n is a leaf with companion v' . We set $o_{n+1} = o_n$, and set $w_{n+1} = w_n$, $v_{n+1} = v'$. The invariant is obviously maintained.

Case v_n is the conclusion of an application of left weakening in which the variable κ is reset. This means that neither κ nor any of its children appear on the right-hand side of Γ_n , and all children of κ are removed. List the children of κ as $\lambda_1, \dots, \lambda_m$. We define the new assignment o_{n+1} by setting:

$$o_{n+1}(\kappa) = \max(o_n(\lambda_1), \dots, o_n(\lambda_m)),$$

and $o_{n+1}(\lambda) = o_n(\lambda)$ for $\lambda \neq \kappa$. We set $w_{n+1} = w_n$ and we set v_{n+1} to be the premise of the rule application. We need to check that the new assignment o_{n+1} refutes Γ_{n+1} at w_{n+1} . Since none of the variables $\kappa, \lambda_1, \dots, \lambda_m$ appear on the right-hand side of Γ_n or Γ_{n+1} , it suffices to show that the new assignment is consistent with the constraint \mathcal{O}_{n+1} of Γ_{n+1} .

First, note that $o_{n+1}(\kappa) < o_n(\kappa)$, since we have $o_n(\lambda_i) < o_n(\kappa)$ for each $i \in \{1, \dots, m\}$ and since $o_{n+1}(\kappa)$ was defined as $\max(o_n(\lambda_1), \dots, o_n(\lambda_m))$ (this observation will be important later!). Hence if κ' is the parent of κ in \mathcal{O}_{n+1} then:

$$o_{n+1}(\kappa) < o_n(\kappa) < o_n(\kappa') = o_{n+1}(\kappa')$$

as required. Now, suppose ξ is a child of κ in \mathcal{O}_{n+1} . Then there must be some λ_i , $i \in \{1, \dots, m\}$, such that ξ was a descendant of λ_i in the constraint of Γ_n . Hence we get:

$$\begin{aligned} o_{n+1}(\xi) &= o_n(\xi) \\ &< o_n(\lambda_i) \\ &\leq \max(o_n(\lambda_1), \dots, o_n(\lambda_m)) \\ &= o_{n+1}(\kappa), \end{aligned}$$

so $o_{n+1}(\xi) < o_{n+1}(\kappa)$ as required.

Other cases. The other cases of left or right weakening are trivial.

With this construction in place, consider the infinite walk $v_0v_1v_2\dots$ through Π that we obtain in the limit. Let X be the set of all vertices that are visited infinitely many times on Π . This is obviously a strongly connected component, so since Π was a valid cyclic proof, by Proposition 4.2 there is some κ that appears in the constraint of every vertex in X , is reset on the path from the root of X to one of its leaves. Hence, it is clear that the value assigned to κ by the ordinal assignments o_i never increases, and decreases every time the walk passes through the vertex at which κ is reset. So the successive values that these ordinal assignments give to κ produce an infinite descending series of ordinals, which is impossible. This contradiction concludes the proof. \square

We now briefly sketch to adapt the soundness proof for the cyclic system to non-wellfounded proofs.

Proposition 4.4. If ρ has a non-wellfounded proof then ρ is valid.

Proof sketch. The argument is similar to the proof of Proposition 4.3: Let Π be a non-wellfounded proof of ρ . We assume for contradiction that there is some model M , world w in M and an ordinal assignment o such that $M, w \not\models_o \rho$. We find an infinite path v_0, v_1 through Π and ordinal assignments o_0, o_1, \dots such that o_i refutes $\mathcal{O}_i : \Gamma_i$, where $\mathcal{O}_i : \Gamma_i$ is the sequent at v_i . The construction of this infinite path is similar as in the proof of Proposition 4.3. We can omit the case for leafs with companions and treat all instances of left weakening as left weakenings that do not correspond to a reset. In the end a contradiction arises from an infinite descending chain of ordinals that can be read off from the sequence o_1, o_2, \dots using the condition from Definition 3.11. \square

5 Completeness for infinite proofs

In this section we will prove completeness for infinite non-wellfounded proofs, as the first step towards our main result, i.e. completeness for cyclic proofs. More precisely, we introduce a special class of non-wellfounded proofs that we call *slim proofs* and prove that any valid formula has a slim proof. Slim proofs are genuinely infinite (non-regular), and can contain infinitely many different sequents. However, they do have the property that there is a fixed bound on the number of *formulas* occurring in any sequent, although the constraints can grow without bound. This property will be important later when we transform slim proofs into finite, cyclic proofs.

5.1 Slim proofs

We start by defining the notion of a slim proof.

Definition 5.1. Let $\mathcal{O} : \Gamma$ be any sequent and let o_0 and o_1 be any annotations that have their range is in $OV(\mathcal{O})$. We write $o_0 \prec_{\mathcal{O}}^x o_1$ if $o_0(x) \prec_{\mathcal{O}} o_1(x)$ while, for all y higher ranking than x , we have $o_0(y) = o_1(y)$. We write $o_0 \prec_{\mathcal{O}} o_1$ if there is some ordinal variable x such that $o_0 \prec_{\mathcal{O}}^x o_1$.

Definition 5.2. Let $\mathcal{O} : \Gamma$ be any sequent, and let $\varphi, \psi \in \Gamma$ be positively annotated formulas. We write $\varphi \prec_{\mathcal{O}} \psi$, or just $\varphi \prec \psi$ when \mathcal{O} is clear from context, if $u(\varphi) = u(\psi)$ and, $o_{\varphi} \prec_{\mathcal{O}} o_{\psi}$.

We write $\preceq_{\mathcal{O}}$ as shorthand for “ $\prec_{\mathcal{O}}$ or $=$ ”. A positively annotated formula φ is *minimal* with respect to a constraint \mathcal{O} if there is no formula ψ such that $\psi \prec_{\mathcal{O}} \varphi$. A sequent $\mathcal{O} : \Gamma$ is said to be minimal if every positively annotated formula occurring in Γ is minimal with respect to \mathcal{O} .

Definition 5.3. Let Γ be a set of plain formulas. The set $OC(\Gamma)$ consists of all positively or negatively annotated formulas φ such that $u(\varphi) \in \text{Clos}(\Gamma)$. The set $QC(\Gamma)$ consists of $OC(\Gamma)$, together with all formulas of the form $Q\kappa < \lambda : \varphi$, where $Q \in \{\forall, \exists\}$ and $\varphi \in OC(\Gamma)$. The extended closure $EC(\Gamma)$ of Γ consists of $QC(\Gamma)$, together with all formulas of the form $[a]\bigvee \Delta$ or $\bigvee \Delta$, where Δ is a subset of $QC(\Gamma)$.

A proof tree Π is said to be slim if:

1. All formulas in Π belong to the extended closure of the root formula.
2. For any formula φ that occurs in some sequent Γ in Π , all free ordinal variables of φ are active in Γ . Equivalently, for any sequent Γ in Π , any ordinal variable that occurs freely in some formula in Γ also occurs freely in some positively annotated formula in Γ .
3. Any sequent in Π which is not minimal with respect to its constraint is the conclusion of an application of right weakening.
4. There is no application of left weakening in Π .
5. In any application of the cut rule, $\mu(\kappa)$ or \exists , all ordinal variables occurring free in the cut formula or minor formula of the rule application are active variables in the conclusion.

Observe that because infinite branches in slim proofs do not contain applications of *lw* the set of ordinal variables in the constraint only grows as we move along the branch. We can thus define the following notion of the limit constraint of a branch:

Definition 5.4. Fix an infinite branch $\beta = v_0, v_1, \dots$ of a slim proof and let $\mathcal{O}_i = (O_i, <_i, \triangleleft_i)$ be the constraint at v_i for all $i \in \omega$. Define the infinite set $O_{\beta} = \bigcup_{i \in \omega} O_i$, and the orders $<_{\beta}$, \triangleleft_{β} and \prec_{β} over O_i such that $\ll_{\beta} = \bigcup_i \ll_i$ for $\ll \in \{<, \leq, \triangleleft\}$.

It is clear that in slim proofs an infinite $<$ -descending chain of ordinals in a branch β , according to Definition 3.11 is the same as an infinite descending chain in the order $<_{\beta}$, according to the usual definition of an infinite descending chain in some order. The same holds for infinite \prec -descending chains.

Proposition 5.5. If $\beta = v_0, v_1, \dots$ is an infinite branch in a slim proof, then

1. $<_\beta$ is irreflexive, transitive and upwards linear.
2. \triangleleft_β is a linear order.
3. \prec_β is a linear order.
4. $<_\beta$ is conversely well-founded and \triangleleft_β is well-founded.
5. $\lambda \prec_\beta \kappa$ iff $\lambda <_\beta \kappa$ or λ is to the left of κ with respect to the orders $<_\beta$ and \triangleleft_β .

Proof. First observe that left weakening is never applied on β because by assumption it is a branch in a slim proof. It follows then from an inspection of the proof rules that for all $i \in \omega$ it holds that either $\mathcal{O}_{i+1} = \mathcal{O}_i$, $\mathcal{O}_{i+1} = \mathcal{O}_i + \lambda$, or $\mathcal{O}_{i+1} = \mathcal{O}_i +_\kappa \lambda$, for some ordinal variable λ . Thus, in every step there is at most one variable λ added; this λ is made maximal in \triangleleft_{i+1} , and it is either a new leaf in $<_{i+1}$ or $<_{i+1}$ -incomparable to all existing variables. Using this observation it is relatively straightforward to verify items 1–4.

For item 5 first recall that λ is to the left of κ with respect to the orders $<_\beta$ and \triangleleft_β if λ and κ are $<_\beta$ -incomparable and $\lambda' \triangleleft_\beta \kappa'$ holds, where λ' is the $<_\beta$ -greatest $<_\beta$ -ancestor of λ that is not an $<_\beta$ -ancestor of κ and κ' is the $<_\beta$ -greatest $<_\beta$ -ancestor of κ that is not an $<_\beta$ -ancestor of λ . Note that the $<_\beta$ -greatest node with some property is well-defined because $<_\beta$ is conversely well-founded.

The details of the proof of item 5 are left to the reader. The crucial step in the argument is to show that the following are equivalent:

1. λ' is the $<_\beta$ -greatest $<_\beta$ -ancestor of λ that is not a $<_\beta$ -ancestor of κ ;
2. for some i , $\lambda', \lambda, \kappa$ exist and λ' is the $<_i$ -greatest $<_i$ -ancestor of λ that is not a $<_i$ -ancestor of κ .

The proof of this equivalence relies on the observation that when moving from \mathcal{O}_i to \mathcal{O}_{i+1} only new leaves or new roots are added to the order $<_{i+1}$. Thus, variables that are comparable at stage i stay comparable at stage $i + 1$, and variables that exist in \mathcal{O}_i and are incomparable stay incomparable. \square

Proposition 5.6. Let β be an infinite branch of a slim proof tree Π . Then it contains an infinite \prec -descending chain of ordinal variables if, and only if, it contains an infinite $<$ -descending chain of ordinal variables.

Proof. It is clear that it suffices to prove this result for the orders \prec_β and $<_\beta$ from Definition 5.4. We write $\mathcal{O}, \prec, <$ and \triangleleft for $\mathcal{O}_\beta, \prec_\beta, <_\beta$, and \triangleleft_β , respectively. The direction from right to left follows directly from the definitions of $<_\beta$ and \prec_β and the relation between $<_i$ and \prec_i at any finite $i \in \omega$.

For the direction from left to right assume that we have an infinite descending chain $\kappa_0 \succ \kappa_1 \succ \dots$ of ordinal variables in \mathcal{O} . The aim is to use König's Lemma to prove that there is an infinite descending chain in $<_{\mathcal{O}_\beta}$.

Consider the following set of variables from \mathcal{O} :

$$F = \{\lambda \in \mathcal{O} \mid \lambda \geq \kappa_i \text{ for some } i \in \omega\}.$$

Because by Proposition 5.5 $<$ is upwards linear and conversely well-founded we can consider F to be a forest under the order $<$. The set F is infinite because it contains the infinitely many distinct κ_i for $i \in \omega$. We first show that there is an infinite subset $T \subseteq F$ that is a tree under $<$.

Claim 2. There are only finitely many $<$ -maximal elements in F .

Proof of Claim 2. First, note that because F is upwards closed in $<$ the concepts of being $<$ -maximal and of being $<$ -maximal in F coincide. Assume then for a contradiction that there are infinitely many $<$ -maximal elements in F . Because $<$ is a linear order and it is well-founded this means that we can build an infinite chain $\delta_0 \triangleleft \delta_1 \triangleleft \dots$ of $<$ -maximal elements in F . By definition of F we have for each $i \in \omega$ some $k_i \in \omega$ such that $\kappa_{k_i} \leq \delta_i$. It follows by item 5 of Proposition 5.5 that there is an infinite ascending $\kappa_{k_0} \prec \kappa_{k_1} \prec \dots$. But because of the infinite descending chain $\kappa_0 \succ \kappa_1 \succ \dots$, this contradicts the fact that \prec is a partial order. \square

It follows by the pigeonhole principle that there must be one $<$ -maximal element r below which there are infinitely many elements from F . We define the set $T \subseteq F$ to be all those elements from F that are $<$ -smaller than r . The set T is an infinite tree under the order $<$. The claim of the proposition then follows by König's Lemma if we can show that this tree T is finitely branching. This is shown in the following claim:

Claim 3. Every variable $\lambda \in T$ has only finitely many $<$ -children in T .

Proof of Claim 3. Assume for a contradiction that some $\lambda \in T$ has infinitely many $<$ -children in T . Because $<$ is a linear order and it is well-founded this means that we can build an infinite chain $\delta_0 \triangleleft \delta_1 \triangleleft \dots$ of $<$ -children of λ with $\delta_i \in T$ for all i . By the definition of T we can find, for every $i \in \omega$, some $k_i \in \omega$ such that $\kappa_{k_i} \leq \delta_i$. It follows by item 5 of Proposition 5.5 that $\kappa_{k_0} \prec \kappa_{k_1} \prec \dots$, because as the δ_i are $<$ -siblings they must be $<$ -greatest ancestors of the κ_{k_i} without being each others ancestors. But this contradicts the fact that \prec is a partial order and the assumption that $\kappa_0 \succ \kappa_1 \succ \dots$. \square

This finishes the proof of Proposition 5.6. \square

A basic property of slim proofs that will be useful is the provability of the following generalized version of excluded middle:

Proposition 5.7. Let φ be a positively annotated formula, and let σ be an increasing substitution. Then there is slim proof for any sequent $\mathcal{O} : \varphi, \overline{\varphi[\sigma]}$.

Proof. The proposition is proved by a formula induction, which we illustrate by considering the case where $\varphi = \nu x^{\kappa_0} \psi$.

By the inductive hypothesis there is a slim proof Π for any sequent $\mathcal{O}' : \psi', \overline{\psi'[\sigma]}$, where ψ' is the formula $\psi[p/x]$, for some fresh proposition letter p . Let Π' be the derivation we obtain from Π by replacing every occurrence of p with the formula $\nu x^{\kappa_1} \psi$, and every occurrence of \overline{p} with the formula $\overline{\mu x^{\kappa_0} \psi[\sigma]}$. Now consider the following derivation Π_1 (where we write λ for $\sigma(\kappa_0)$):

$$\frac{\frac{\frac{\Pi'}{\vdots}}{\mathcal{O} +_{\kappa_0} \kappa_1 : \psi(\nu x^{\kappa_1} \psi), \overline{\psi[\sigma]}(\overline{\mu x^{\kappa_0} \psi[\sigma]})} \nu(\kappa_0)}{\mathcal{O} : \nu x^{\kappa_0} \psi, \overline{\psi[\sigma]}(\overline{\mu x^{\kappa_0} \psi[\sigma]})} \mu}{\mathcal{O} : \nu x^{\kappa_0} \psi, \overline{\mu x^\lambda \psi[\sigma]}}$$

Note that Π_1 need not be a proof, since there may be leaves in Π_1 , that were labeled in Π' by an axiom $\mathcal{O}' : p, \overline{p}$, but in π_1 are labeled with the sequent $\mathcal{O}' : \nu x^{\kappa_1} \psi, \overline{\mu x^{\kappa_0} \psi[\sigma]}$. However, to any such leaf we may apply the same construction that we just described, with $\sigma[\kappa_0/\kappa_1]$ taking the role of σ . Iterating this procedure we arrive at a slim proof in which any newly created infinite branch carries an infinitely descending sequence of variables $\dots < \kappa_2 < \kappa_1 < \kappa_0$. \square

The following is an immediate corollary of Proposition 5.7

Proposition 5.8 (Excluded middle). Let φ be a positively annotated formula. Then there is slim proof for any sequent $\mathcal{O} : \varphi, \overline{\varphi}$.

The remainder of this section is devoted to proving completeness for slim proofs (Theorem 5.28). The proof is carried out in several steps, so a high-level sketch may be helpful:

The first step is to define a particular kind of sequents, called *tiles*, that will be used as building blocks from which we construct a counter-model to a given formula that is not provable. Tiles are saturated sequents of a bounded size, where the saturation conditions are intended to mimic the conditions under which all formulas in the sequent are false; for example if a disjunction is present in the sequent, both disjuncts should be false, and if a formula $\forall \lambda < \kappa \varphi(\lambda)$ is present in the sequent then there should be some witness ξ for which $\varphi(\xi)$ is false. However, adding such saturation conditions naïvely - e.g. if the sequent contains $\forall \lambda < \kappa \varphi(\lambda)$ then it contains $\varphi(\xi)$ for some $\xi < \kappa$ - would quickly destroy any hopes of maintaining a bound on the number of formulas appearing in a sequent. The appropriate saturation conditions have to be formulated carefully; the precise definition is Definition 5.13.

Completeness is proved in Section 5.4. A determined two-player game called the *mosaic game* is defined, in which “Prover” tries to produce a slim proof of a given formula, and the opposing player “Refuter” tries to construct a counter-model from tiles. The most difficult part of the completeness proof is to show

that a winning strategy for Refuter actually provides a counter-model. To get some intuition about what the difficulty is, and how it is handled, consider a formula $[a]\nu x^\kappa \langle \check{a} \rangle \varphi(x)$, and say that Refuter is trying to construct a counter-model for some tile Γ containing this formula. In the mosaic game, Prover can “attack” the box $[a]$ and ask Refuter to present some tile Ψ which is to be a -accessible from Γ , and contains $\nu x^\kappa \langle \check{a} \rangle \varphi(x)$. In order to falsify $\nu x^\kappa \langle \check{a} \rangle \varphi(x)$ at Ψ , Refuter has to saturate by adding some child λ of the variable κ to the constraint of Ψ , along with the formula $\langle \check{a} \rangle \varphi(\nu x^\lambda \langle \check{a} \rangle \varphi(x))$. So here, the variable κ has progressed to the “smaller” λ . But if Ψ was to be a -accessible from Γ , this means that Refuter has to make sure that the formula $\nu x^\lambda \langle \check{a} \rangle \varphi(x)$ is false at Γ . In other words, already when the tile Γ was constructed, Refuter has to anticipate that the variable κ will progress due to later steps in the counter-model construction.

The solution to handling this sort of back-and-forth behaviour is to make sure that Γ already contains some formula which is “as good as” the formula $\nu x^\lambda \langle \check{a} \rangle \varphi(x)$, in some sense. It is quite tricky to formulate an appropriate saturation condition corresponding to this intuition. The notion of a *minimality assumption* introduced in Definition 5.9 will be the key concept used to handle this, with the saturation condition given in Definition 5.13. Definition 5.17 defines a notion of a -successor of a tile, which is intended to capture the coherence conditions we require of the a -accessibility relation in the counter-model construction used in the completeness proof. The key technical result that is used to handle the difficulty sketched above is Proposition 5.24, which captures the sense in which a tile anticipates progressions of ordinal variables that may occur in its a -successors.

5.2 Pre-tiles and tiles

Our completeness proof for slim proofs will build a counter-model for an unprovable formula out of building blocks that we call *tiles*, which are essentially “small” but saturated sequents. The precise notion of “saturated” that we will use is quite technical, since a naïve concept of saturation would be incompatible with maintaining a bound on the number of formulas appearing in a sequent. It will be convenient to first define the notion of a “pre-tile” to capture the general structure of the sequents that we will be working with, and then define a tile as a pre-tile that satisfies appropriate saturation conditions.

Definition 5.9. A ρ -pre-tile is a sequent $\mathcal{O} : \Gamma$ where Γ admits a partition into three disjoint subsets Γ^+ , Γ^- , Γ^μ such that:

- Every formula $\varphi \in \Gamma^+$ is positively annotated and $u(\varphi)$ belongs to the closure of ρ .
- Every formula $\varphi \in \Gamma^-$ is the negation of a positively annotated formula ψ such that $u(\psi)$ belongs to the closure of ρ , and every ordinal variable that occurs freely in ψ is active in Γ^+ .

- Every formula $\varphi \in \Gamma^\mu$ is of the form $\exists \lambda < \kappa : \bar{\psi}[\lambda/\kappa]$ where $\psi \in \Gamma^+$ and κ appears in ψ .

The sets Γ^+ , Γ^- , Γ^μ are called the *positive part*, the *negative part* and the *minimality assumptions* of Γ respectively. The ρ -pre-tile is said to be *small* if every formula in Γ^+ is minimal with respect to \mathcal{O} .

Observe that, for every plain formula φ , adding the empty constraint yields an example of a pre-tile, viz., $\emptyset : \varphi$. Note that this pre-tile is of a special shape: the negative part and the set of minimality assumptions are empty. In general we can think of the partition of pre-tiles into these three parts - positive, negative and minimality assumptions - as keeping track of the different roles that formulas play in the completeness proof. The positive part of a pre-tile is the “main” part in a sense; our goal from the start is to construct a counter-model to the positive part of a pre-tile of the form $\emptyset : \varphi$. Along the way, assumptions may be introduced via cuts. In a two-sided sequent notation we can think of pre-tiles as sequents of the shape:

$$\mathcal{O} : \Gamma, \Delta \Rightarrow \Theta$$

where Δ, Θ consist of positively annotated formulas, and Γ consists of formulas of the form $\forall \lambda < \kappa \psi$ where $\psi \in \Theta$. The function of the sets Γ, Δ is essentially to constrain what formulas can belong to the positive part (Θ) without making the sequent provable. A formula $\psi \in \Delta$ belonging to the negative part ensures that ψ cannot belong to the positive part, or “on the right-hand side” of the sequent, since that would result in a trivially provable sequent. Similarly, a minimality assumption $\forall \lambda < \kappa \psi(\lambda) \in \Gamma$ ensures that Θ cannot contain the formula $\psi(\xi)$ for any $\xi <_{\mathcal{O}} \kappa$, since that would also make the sequent provable.

Proposition 5.10. For every constraint $\mathcal{O} = (O, <_{\mathcal{O}}, \triangleleft_{\mathcal{O}})$ the order $\triangleleft_{\mathcal{O}}$ is a strict linear order over ordinal variables.

Proof. It is obvious from the definition that $\triangleleft_{\mathcal{O}}$ is irreflexive.

To show that $\triangleleft_{\mathcal{O}}$ is transitive assume that $\kappa \triangleleft_{\mathcal{O}} \lambda$ and $\lambda \triangleleft_{\mathcal{O}} \delta$. Because of the disjunctive definition of $\triangleleft_{\mathcal{O}}$ one has to distinguish four cases. We only consider the case where κ is to the left of λ and λ is to the left of δ . Let κ' be the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of κ that is not a $<_{\mathcal{O}}$ -ancestor of λ , α be the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of λ that is not a $<_{\mathcal{O}}$ -ancestor of κ , β the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of λ that is not a $<_{\mathcal{O}}$ -ancestor of δ , and δ' the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of δ that is not a $<_{\mathcal{O}}$ -ancestor of λ . From the assumption we have that $\kappa' \triangleleft_{\mathcal{O}} \alpha$ and $\beta \triangleleft_{\mathcal{O}} \delta'$. Moreover, we have that κ' and α either have the same parent in $<_{\mathcal{O}}$ or they are both $<$ -maximal and similarly for β and δ' . Because $<_{\mathcal{O}}$ is upwards-linear we have that either $\alpha = \beta$, $\alpha <_{\mathcal{O}} \beta$ or $\beta <_{\mathcal{O}} \alpha$. We consider all these cases.

If $\alpha = \beta$ then it follows that either κ' and δ' have the same parent in $<_{\mathcal{O}}$ or that they are both $<_{\mathcal{O}}$ -maximal. It follows that κ' is also the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of κ that is not a $<_{\mathcal{O}}$ -ancestor of δ and similarly δ' is also the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of δ that is not a $<_{\mathcal{O}}$ -ancestor of κ . Moreover, by

the transitivity of $\triangleleft_{\mathcal{O}}$ we also get that $\kappa' \triangleleft_{\mathcal{O}} \lambda'$. Hence, κ is to the left of λ and $\kappa \prec_{\mathcal{O}} \lambda$ follows by the definition of $\prec_{\mathcal{O}}$.

If $\alpha <_{\mathcal{O}} \beta$ then β is a $<_{\mathcal{O}}$ -ancestor of κ because α was the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of λ that is not an $<_{\mathcal{O}}$ -ancestor of κ . It follows that either the common parent of β and δ' is a $<_{\mathcal{O}}$ -ancestor of both κ and δ , or that β and δ' are both $<_{\mathcal{O}}$ -maximal. Because we already know that β and δ' have the same parent or are both $<_{\mathcal{O}}$ -maximal it follows that β is the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of κ that is not a $<_{\mathcal{O}}$ -ancestor of δ , and δ' is the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of δ that is not a $<_{\mathcal{O}}$ -ancestor of κ . Hence we get $\kappa \prec_{\mathcal{O}} \delta$ because $\beta \triangleleft_{\mathcal{O}} \delta'$.

If $\beta <_{\mathcal{O}} \alpha$ then α is a $<_{\mathcal{O}}$ -ancestor of δ because β was the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of λ that is not an $<_{\mathcal{O}}$ -ancestor of δ . It follows that either the common parent of κ' and α is a $<_{\mathcal{O}}$ -ancestor of both κ and δ , or that κ' and β are both $<_{\mathcal{O}}$ -maximal. Because we already know that κ' and α have the same parent or are both $<_{\mathcal{O}}$ -maximal it follows that κ' is the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of κ that is not a $<_{\mathcal{O}}$ -ancestor of δ , and α is the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor of δ that is not a $<_{\mathcal{O}}$ -ancestor of κ . Hence we get $\kappa \prec_{\mathcal{O}} \delta$ because $\kappa' \triangleleft_{\mathcal{O}} \beta$.

To show that $\prec_{\mathcal{O}}$ is linear consider any $\kappa \neq \lambda$. We need to show that either $\kappa \prec_{\mathcal{O}} \lambda$ or $\lambda \prec_{\mathcal{O}} \kappa$. First, distinguish cases depending on whether κ and λ are comparable in the strict partial order $<_{\mathcal{O}}$. If they are comparable then we have either $\kappa <_{\mathcal{O}} \lambda$ or $\lambda <_{\mathcal{O}} \kappa$. It thus follows immediately from the definition of $\prec_{\mathcal{O}}$ that then $\kappa \prec_{\mathcal{O}} \lambda$ or $\lambda \prec_{\mathcal{O}} \kappa$. In the other case κ and λ are $<_{\mathcal{O}}$ -incomparable it follows also that they are $\leq_{\mathcal{O}}$ -incomparable since we assume $\kappa \neq \lambda$. There then must exist the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor κ' of κ that is not an $<_{\mathcal{O}}$ -ancestor of λ and, similarly, there exists the $<_{\mathcal{O}}$ -maximal $<_{\mathcal{O}}$ -ancestor λ' of λ that is not a $<_{\mathcal{O}}$ -ancestor of κ . It is clear that λ' and κ' must be distinct. Because $\triangleleft_{\mathcal{O}}$ is a strict linear order it follows that either $\kappa \triangleleft_{\mathcal{O}} \lambda$ or $\lambda \triangleleft_{\mathcal{O}} \kappa$. In the former case we get that $\kappa \prec_{\mathcal{O}} \lambda$ and in the latter that $\lambda \prec_{\mathcal{O}} \kappa$. \square

Proposition 5.11. Let $\mathcal{O} : \Gamma$ be any pre-tile and φ any propositional formula. Then $\prec_{\mathcal{O}}$ is a strict linear order over the set $\{\psi \in \Gamma^+ \mid u(\psi) = \varphi\}$.

Proof. Fix φ and restrict $\prec_{\mathcal{O}}$ to the set $\{\psi \in \Gamma^+ \mid u(\psi) = \varphi\}$.

Observe that if $\psi, \chi \in \Gamma^+$ with $u(\psi) = u(\chi)$ then the domains of \mathfrak{o}_{ψ} and \mathfrak{o}_{χ} must be the same. This is the case because by Definition 5.9 $\psi, \chi \in \Gamma^+$ entails that both ψ and χ are positively annotated, which by Definition 2.3 means that all greatest fixpoint variables are annotated, and both formulas contain the same greatest fixpoint variables because $u(\psi) = u(\chi)$.

It follows from this observation that the restriction of the order $\prec_{\mathcal{O}}$ from Definition 5.2 to formulas from $\{\psi \in \Gamma^+ \mid u(\psi) = \varphi\}$ is just the lexicographic extension of the order $\prec_{\mathcal{O}}$ over ordinal variables to finite sequences, where each position in the sequence corresponds to one greatest fixpoint variable in φ . It is clear that this yields a strict linear order, because as shown in Proposition 5.10 $\prec_{\mathcal{O}}$ is a strict linear order of ordinal variables. \square

We can now prove that small ρ -pre-tiles are indeed small in the sense that there is a bound on the number of formulas that they contain.

Proposition 5.12. Let ρ be any formula. Then there is a bound n such that, for any small ρ -pre-tile $\mathcal{O} \vdash \Gamma$, there are at most n distinct formulas in Γ (up to renaming of bound ordinal variables).

Proof. Because the ρ -pre-tile Γ is small, Γ^+ contains only positively annotated formulas that are minimal with respect to the order $\prec_{\mathcal{O}}$. From Proposition 5.11 we know that $\prec_{\mathcal{O}}$ is a strict linear order over any set Ψ of annotated formulas with the same underlying plain formula, i.e. such that $u(\varphi) = u(\psi)$ whenever $\varphi, \psi \in \Psi$. It follows that Γ^+ contains at most one annotated formula φ for every underlying plain formula $u(\varphi)$, equal to $u(\varphi)^o$ for some annotation o . This means that Γ^+ is bounded by the size m of the closure of ρ . This also means that the number of active ordinal variables in Γ has a fixed upper bound, since every annotated formula θ^o in Γ^+ can at most have one ordinal variable assigned to each fixpoint variable by the annotation o . Since every formula in Γ^- is the negation of some annotated formula, whose underlying plain formula is in the closure of ρ , it follows that the number of formulas in Γ^- is bounded by $m \cdot k^{|\text{Var}(\rho)|}$, where k is the fixed upper bound on the number of active ordinal variables. (Note that $k^{|\text{Var}(\rho)|}$ bounds the number of annotations of fixpoint variables in ρ by active ordinal variables in Γ .) For Γ^μ , since each of the formulas in Γ^μ is of the form $\exists \lambda < \kappa : \varphi[\lambda/\kappa]$ where $\varphi \in \Gamma^+$, the number of distinct such formulas up to renaming of bound variables is bounded by $m \cdot k$. Since the formula part of Γ is $\Gamma^+ \cup \Gamma^- \cup \Gamma^\mu$, the result follows. \square

Note that essentially the same proof with some minor modifications can be used to show that, for any slim proof Π , there is a bound on the number of formulas that can occur in any sequent in Π up to renaming of bound ordinal variables. We sometimes abuse notation by writing just Γ instead of $\mathcal{O} : \Gamma$. In such cases we let $\mathcal{O}(\Gamma)$ denote \mathcal{O} . From now on, without further mention we will equate formulas that differ only by a renaming of bound variables, and similarly we count two pre-tiles as the same if they are equal up to renaming of bound ordinal variables.

Definition 5.13. Let $\mathcal{O} : \Gamma$ be a (small) ρ -pre-tile. We say that this pre-tile is *propositionally saturated* if it satisfies the following condition:

1. If φ is positively annotated, $u(\varphi)$ is in the closure of ρ and all ordinal variables appearing in φ are active in Γ^+ , then either $\bar{\varphi} \in \Gamma^-$ or $\varphi' \in \Gamma^+$ for some $\varphi' \preceq_{\mathcal{O}} \varphi$.

If, in addition, $\mathcal{O} : \Gamma$ satisfies the conditions below, we call it a (small) ρ -tile:

2. $\nu x \varphi \notin \Gamma^+$ for all formulas of the form $\nu x \varphi$ (outermost ν -fixpoints must be annotated).
3. If $\nu x^\kappa \varphi \in \Gamma^+$ then for some κ_0 with $\kappa_0 \prec_{\mathcal{O}} \kappa$ and for some $\varphi' \preceq_{\mathcal{O}} \varphi(\nu x^{\kappa_0} \varphi)$, we have $\varphi' \in \Gamma^+$.
4. If all variables in $\varphi(\nu x^\kappa \psi)$ are active in Γ and the underlying plain formula is in the closure of ρ , and $\theta \in \Gamma^+$ for some $\theta \preceq_{\mathcal{O}(\Gamma)} \varphi(\nu x^\kappa \psi)$, then either

$\exists \lambda < \kappa : \overline{\varphi}(\mu x^\lambda \overline{\psi}) \in \Gamma^\mu$, or for some κ_0 with $\kappa_0 <_{\mathcal{O}} \kappa$ and some $\theta' \preceq_{\mathcal{O}} \varphi(\nu x^{\kappa_0} \psi)$ we have $\theta' \in \Gamma^+$.

The collection of small ρ -tiles is denoted as T_ρ .

The last condition in Definition 5.13 expresses a saturation condition corresponding to the minimisation rule. We can understand it semantically as follows: suppose we wish to construct a counter-model for the unprovable sequent $\mathcal{O} : \Gamma$, where $\varphi(\nu x^\kappa \psi) \in \Gamma$ (or some \preceq -smaller formula belongs to Γ). By the minimisation rule, either $\mathcal{O} : \Gamma, \exists \lambda < \kappa : \overline{\varphi}(\mu x^\lambda \overline{\psi})$ is unprovable, or $\mathcal{O} +_\kappa \kappa_0 : \Gamma, \varphi(\nu x^{\kappa_0} \psi)$ is unprovable. In either of these two sequents, we have more information about how the ordinal variable κ must be evaluated in a refuting model. In any counter-model for $\mathcal{O} : \Gamma, \exists \lambda < \kappa : \overline{\varphi}(\mu x^\lambda \overline{\psi})$, we know that κ has to be evaluated as the *smallest* ordinal that refutes $\varphi(\nu x^\kappa \psi)$. Otherwise it would not be a counter-model, as the formula $\exists \lambda < \kappa : \overline{\varphi}(\mu x^\lambda \overline{\psi})$ would be true. In a counter-model for the sequent $\mathcal{O} +_\kappa \kappa_0 : \Gamma, \varphi(\nu x^{\kappa_0} \psi)$, the opposite holds: there is some ordinal smaller than the value of κ that refutes the formula $\varphi(\nu x^\kappa \psi)$, and we have introduced a fresh name κ_0 for this ordinal.

In the next subsection we will show how an unprovable pre-tile can be “saturated” to an unprovable tile. For technical reasons it will be necessary to be careful with how this saturation is carried out, so that a given pre-tile stands in a suitable relationship to the corresponding tiles. This relationship \preceq is spelled out in the following definition; intuitively, we can read $\Psi \preceq \Gamma$ as “ Ψ is at least as saturated as Γ ”.

Definition 5.14. Let Γ, Ψ be pre-tiles, and let L be a set of ordinal variables. We write $\Psi \preceq_L \Gamma$ if the following conditions hold:

1. $\text{OV}(\mathcal{O}(\Psi)) = \text{OV}(\mathcal{O}(\Gamma)) \uplus L$.
2. For all $\lambda, \lambda' \in \text{OV}(\mathcal{O}(\Gamma))$ it holds that $\lambda <_{\mathcal{O}(\Psi)} \lambda'$ iff $\lambda <_{\mathcal{O}(\Gamma)} \lambda'$, and that $\lambda <_{\mathcal{O}(\Psi)} \lambda'$ iff $\lambda <_{\mathcal{O}(\Gamma)} \lambda'$.
3. $\text{Act}(\Psi) \subseteq \text{Act}(\Gamma) \uplus L$.
4. $\kappa <_{\mathcal{O}(\Psi)} \lambda$ for all $\kappa \in \text{OV}(\mathcal{O}(\Gamma))$ and $\lambda \in L$.
5. No variable in L is an $<_{\mathcal{O}(\Gamma)}$ -ancestor of any variable in $\mathcal{O}(\Gamma)$.
6. If $\lambda \in L$, $\kappa \in \mathcal{O}(\Gamma)$ and $\lambda <_{\mathcal{O}(\Psi)} \kappa$, then there is some $\kappa' \in \text{Act}(\Gamma)$ such that $\lambda <_{\mathcal{O}(\Psi)} \kappa' \leq_{\mathcal{O}(\Psi)} \kappa$.
7. For every formula $\varphi \in \Gamma^+$ there is some formula $\varphi' \in \Psi^+$ such that $\varphi' \preceq_{\mathcal{O}(\Psi)} \varphi$.

We write $\Psi \preceq \Gamma$ if there is some L such that $\Psi \preceq_L \Gamma$.

Proposition 5.15. The relation \preceq over tiles is a pre-order.

Proof. Reflexivity is trivial ($\Gamma \preceq_{\emptyset} \Gamma$), so we check transitivity. Let $\Gamma_2 \preceq_{L_1} \Gamma_1 \preceq_{L_0} \Gamma_0$. We show that $\Gamma_2 \preceq_{L_0 \uplus L_1} \Gamma_0$, checking each condition individually.

1. We have:

$$\begin{aligned} \text{OV}(\mathcal{O}(\Gamma_2)) &= \text{OV}(\mathcal{O}(\Gamma_1)) \uplus L_1 \\ &= \text{OV}(\mathcal{O}(\Gamma_0)) \uplus L_0 \uplus L_1 \end{aligned}$$

2. Straightforward.

3. We have:

$$\begin{aligned} \text{Act}(\Gamma_2) &\subseteq \text{Act}(\Gamma_1) \uplus L_1 \\ &\subseteq \text{Act}(\Gamma_0) \uplus L_0 \uplus L_1 \end{aligned}$$

4. No variable in L_0 is older than any variable in $\mathcal{O}(\Gamma_0)$, and no variable in L_1 is older than any variable in $\mathcal{O}(\Gamma_1)$. But $\mathcal{O}(\Gamma_1) = \mathcal{O}(\Gamma_0) \uplus L_0$, so this means that no variable in L_1 is older than any variable in $\mathcal{O}(\Gamma_0)$. Hence no variable in $L_0 \uplus L_1$ is older than any variable in $\mathcal{O}(\Gamma_0)$.

5. No variable in L_0 is an ancestor of any variable in $\mathcal{O}(\Gamma_0)$, and no variable in L_1 is an ancestor of any variable in $\mathcal{O}(\Gamma_1)$. But $\mathcal{O}(\Gamma_1) = \mathcal{O}(\Gamma_0) \uplus L_0$, so this means that no variable in L_1 is ancestor of any variable in $\mathcal{O}(\Gamma_0)$. Hence no variable in $L_0 \uplus L_1$ is an ancestor of any variable in $\mathcal{O}(\Gamma_0)$.

6. Suppose $\kappa \in \mathcal{O}(\Gamma_0)$ and $\lambda \in L_0 \uplus L_1$, where $\lambda <_{\mathcal{O}(\Gamma_2)} \kappa$. If $\lambda \in L_0$ then since $\Gamma_1 \preceq_{L_0} \Gamma_0$ it follows that there is some $\lambda' \in \text{Act}(\Gamma_0)$ such that $\lambda <_{\mathcal{O}(\Gamma_1)} \lambda' \leq_{\mathcal{O}(\Gamma_1)} \kappa$. But then $\lambda <_{\mathcal{O}(\Gamma_2)} \lambda' \leq_{\mathcal{O}(\Gamma_2)} \kappa$ as well. If $\lambda \in L_1$, then since $\kappa \in \mathcal{O}(\Gamma_1)$ and $\Gamma_2 \preceq_{L_1} \Gamma_1$, there is some $\lambda' \in \text{Act}(\Gamma_1)$ such that $\lambda <_{\mathcal{O}(\Gamma_2)} \lambda' \leq \kappa$. But $\text{Act}(\Gamma_1) \subseteq \text{Act}(\Gamma_0) \uplus L_0$. If $\lambda' \in \text{Act}(\Gamma_0)$ then we are done. If $\lambda' \in L_0$, then $\lambda' <_{\mathcal{O}(\Gamma_1)} \kappa$ and so since $\Gamma_1 \preceq_{L_0} \Gamma_0$ there is some $\lambda'' \in \text{Act}(\Gamma_0)$ such that $\lambda' <_{\mathcal{O}(\Gamma_1)} \lambda'' \leq_{\mathcal{O}(\Gamma_1)} \kappa$. But then we have:

$$\lambda <_{\mathcal{O}(\Gamma_1)} \lambda' <_{\mathcal{O}(\Gamma_1)} \lambda'' \leq_{\mathcal{O}(\Gamma_1)} \kappa$$

and so we are done.

7. Let $\varphi \in \Gamma_0^+$. Since $\Gamma_1 \preceq \Gamma_0$ there is some formula $\varphi' \in \Gamma_1^+$ such that $\varphi' \preceq_{\mathcal{O}(\Gamma_1)} \varphi$. Since $\Gamma_2 \preceq \Gamma_1$ there is some formula $\varphi'' \in \Gamma_2^+$ such that $\varphi'' \preceq_{\mathcal{O}(\Gamma_2)} \varphi'$. But clearly $\varphi' \preceq_{\mathcal{O}(\Gamma_1)} \varphi$ entails $\varphi' \preceq_{\mathcal{O}(\Gamma_2)} \varphi$, so transitivity of $\preceq_{\mathcal{O}(\Gamma_2)}$ we get $\varphi'' \preceq_{\mathcal{O}(\Gamma_2)} \varphi$. \square

Proposition 5.16. Suppose $\Psi \preceq_L \Gamma$. If $\lambda \in L$ and $\kappa \in \text{Act}(\Gamma)$ are such that λ is to the left of κ in $\mathcal{O}(\Psi)$, then there is an ordinal variable $\lambda' \in \text{Act}(\Gamma)$ such that $\lambda <_{\mathcal{O}(\Psi)} \lambda'$ and λ' is also to the left of κ in $\mathcal{O}(\Psi)$.

Proof. Suppose $\lambda \in L$ and $\kappa \in \text{Act}(\Gamma)$, and λ is to the left of κ in $\mathcal{O}(\Psi)$. Then there are ancestors κ_0, λ_0 of κ, λ respectively such that $\lambda_0 \triangleleft_{\mathcal{O}(\Psi)} \kappa_0$. Since $\kappa \in \mathcal{O}(\Gamma)$, we have $\kappa_0 \in \mathcal{O}(\Gamma)$ since $\mathcal{O}(\Psi) = \mathcal{O}(\Gamma) \uplus L$ and no variable in L is an ancestor of any variable in $\mathcal{O}(\Gamma)$. Since $\lambda_0 \triangleleft_{\mathcal{O}(\Psi)} \kappa_0$ we have $\lambda_0 \in \mathcal{O}(\Gamma)$ since $\mathcal{O}(\Psi) = \mathcal{O}(\Gamma) \uplus L$ and no variable in L is older than any variable in $\mathcal{O}(\Gamma)$. It follows that $\lambda <_{\mathcal{O}(\Psi)} \lambda_0$ since $\lambda \in L$, $\lambda_0 \in \mathcal{O}(\Gamma)$ and λ_0 was an ancestor of λ . Since $\Psi \preceq_L \Gamma$, there is some $\lambda' \in \text{Act}(\Gamma)$ with $\lambda <_{\mathcal{O}(\Psi)} \lambda' \leq_{\mathcal{O}(\Psi)} \lambda_0$. It follows that λ' is to the left of κ , and so we are done. \square

Definition 5.17. Let Γ be a ρ -tile. A ρ -tile Ψ is said to be an (a, φ, L) -successor of Γ if $[a]\varphi \in \Gamma^+$ and

$$\Psi \preceq_L (\mathcal{O}(\Gamma) : \{\varphi\} \cup \{\psi \mid \langle a \rangle \psi \in \Gamma^+\}).$$

We say that Ψ is an a -successor of Γ if it is an (a, φ, L) -successor of Γ for some φ and L . An a -successor Ψ of Γ is said to be a *non-trivial* a -successor if the sequent $\mathcal{O}(\Psi) : [\check{a}] \vee \Gamma, \Psi$ is not provable by a slim proof.

Similarly, we say that Ψ is an a -successor of Γ , relative to some formula φ , and we write $\Gamma \rightarrow_{a,\varphi} \Psi$, if Ψ is an (a, φ, L) -successor of Γ for some L .

Proposition 5.18. Suppose Ψ is a (a, φ, L) -successor of Γ . If $\lambda \in L$ and $\kappa \in \text{Act}(\Gamma)$ are such that λ is to the left of κ in $\mathcal{O}(\Psi)$, then there is an ordinal variable $\lambda' \in \text{Act}(\Gamma)$ such that $\lambda < \lambda'$ and λ' is also to the left of κ in $\mathcal{O}(\Psi)$.

Proof. Immediate from Proposition 5.16. □

5.3 Technical results on tiles and pre-tiles

In this subsection we will prove several observations on the relations between tiles and pre-tiles that will be relevant to the mosaic game that we introduce later as the main tool for our completeness proof.

The following proposition shows that provability of a given pre-tile can be reduced to provability of a (finite) set of tiles. From the dual perspective, any non-provable pre-tile can be saturated to some non-provable tile. Intuitively, the idea behind the proof is to step by step construct a tree representing a non-deterministic procedure to saturate a given pre-tile, in such a way that in each step of the procedure some of the ordinal variables progress. For example, we may add a child ξ of κ together with the formula $\varphi(\xi)$ to a pre-tile that contains the formula $\forall \lambda < \kappa \varphi(\lambda)$ in order to meet the saturation condition associated with such formulas, which results in a progression of the variable κ . The result of this is that if a particular run of the procedure does not terminate but goes on forever, then it produces an infinitely descending chain of ordinal variables. The tree representing all possible runs of this non-deterministic procedure can then be viewed as a slim proof tree that derives the given pre-tile from a finite set of tiles as assumptions.

Proposition 5.19. Let $\mathcal{O} : \Gamma$ be any small ρ -pre-tile, and let Δ be any finite set of formulas. Then there is a (possibly infinite) slim proof tree Π with root labelled $\mathcal{O} : \Gamma, \Delta$, in which every non-axiom leaf is of the form $\mathcal{O}' : \Psi, \Delta$ for some small ρ -tile $\mathcal{O}' : \Psi$ with $\Psi \preceq \Gamma$, and every infinite branch in Π has an infinite descending $<$ -chain of ordinal variables.

Proof. Throughout the proof we assume $\Delta = \emptyset$, for notational convenience. The proof for the general case is not different in any significant respect.

We construct the proof tree Π in a step-by-step manner, as the limit of a series of approximations $(\Pi_i)_{i < \omega}$, where each leaf in each approximant Π_i is propositionally saturated. Let Π_0 be the finite proof tree having Γ as its root

sequent, followed by enough applications of the cut rule to ensure that all the leaves of Π_0 are propositionally saturated.

Given that the proof tree Π_i has been constructed, consider a leaf of Π_i labelled by the sequent Ψ . We say that χ in the closure of ρ is a *defect* of Ψ if one of the following conditions hold:

- χ is of the form $\nu x\varphi$.
- χ belongs to Ψ^+ , is of the form $\nu x^{\kappa}\varphi$, and there are no κ_0 and $\varphi \in \Psi^+$ such that $\kappa_0 <_{\mathcal{O}(\Psi)} \kappa$ and $\varphi' \preceq_{\mathcal{O}(\Gamma)} \varphi(\nu x^{\kappa_0}\varphi)$.
- χ is of the form $\varphi(\nu x^{\kappa}\psi)$, all free ordinal variables of χ are active in Γ , the formula $\exists\lambda < \kappa : \overline{\varphi(\mu x^{\lambda}\psi)}$ does not belong to Ψ^{μ} , and there are no κ_0 and $\varphi' \in \Psi^+$ with $\kappa_0 <_{\mathcal{O}(\Psi)} \kappa$ and $\varphi' \preceq_{\mathcal{O}(\Gamma)} \varphi(\nu x^{\kappa_0}\psi)$.

Clearly, a leaf of Π_i is a tile if it is propositionally saturated and has no defects. If every leaf in Π_i is a tile then we stop the process and set $\Pi = \Pi_i$. Otherwise, we pick a leaf l in Π_i and a defect of its label Ψ , and extend Π_i to Π_{i+1} according to the following steps:

1. Apply the appropriate rule (the minimisation rule, the rule $x : \kappa$ adding the formula $\nu x^{\kappa}\varphi \prec_{\mathcal{O}(\Pi_i)+\kappa} \varphi$, or the $\nu(\kappa)$ -rule) depending on the type of defect, adding one or two children for l in which the chosen formula is no longer a defect.
2. Use right weakening to remove all non-minimal formulas in the positive part of each new leaf. Note that this step removes any formula of the form $\nu x.\varphi$ that was the principal formula of the rule $x : \kappa$ applied in the first step.
3. Finally, apply cuts until all leaves above l are propositionally saturated.

To be precise, the choice of which defect to deal with in each step of the construction has to follow a fair scheduling procedure to ensure that each defect is eventually taken care of, but this is rather trivial since there are only finitely many defects to choose from at each stage. Given that we implement fair scheduling, we ensure that each leaf in the proof tree Π that we produce as the limit of the sequence $(\Pi_i)_{i < \omega}$ is a tile, so it remains only to show that every infinite branch in Π has an infinite descending chain of ordinal variables.

So let β be an infinite branch of Π . For each $i < \omega$, let $\beta|\Pi_i$ denote the initial segment of β that is a branch from the root to a leaf of Π_i , and let Γ_i denote the last sequent on $\beta|\Pi_i$. For each $i < \omega$ we write \mathcal{O}_i for $\mathcal{O}(\Gamma_i)$. Since β is infinite there are infinitely many i such that $\beta|\Pi_i$ is a proper initial segment of $\beta|\Pi_{i+1}$; we say that such an index i is a *growth point* for β .

The following claim shows that the procedure maintains the invariant that, if Ψ' is a descendant of Ψ in Π_i , then $\Psi' \preceq \Psi$.

Claim 4. Let $L = \mathcal{O}(\Gamma_{i+1}) \setminus \mathcal{O}(\Gamma_i)$. Then,

1. $\mathcal{O}(\Gamma_{i+1}) = \mathcal{O}(\Gamma_i) \uplus L$.

2. $\text{Act}(\Gamma_{i+1}) \subseteq \text{Act}(\Gamma_i) \uplus L$.
3. No variable in L is older than any variable in $\text{OV}(\mathcal{O}(\Gamma_i))$, and every variable in L is a root (has no ancestors) or a *child* of some variable in $\text{Act}(\Gamma_i)$.
4. If $\varphi \in \Gamma_i^+$, then there is some formula $\varphi' \in \Gamma_{i+1}^+$ with $\varphi' \preceq_{\mathcal{O}_{i+1}} \varphi$.

Proof of Claim 4. The proof of items (1) and (2) are trivial since the construction never removes variables from the constraint and never makes pre-existing non-active variables active. For the proof of item (3), we just need to inspect the construction of Γ_{i+1} : it is obvious that no variable in $\mathcal{O}(\Gamma_{i+1}) \setminus \mathcal{O}(\Gamma_i)$ is older than any variable in $\mathcal{O}(\Gamma_i)$, by inspection of the rules applied. Furthermore, we note that any variable in L must have been added either as a new root or as a child of some active variable of Γ_i ; the rule $x : \kappa$ adds κ as a new root, the minimisation rule is only applied to formulas all of whose variables are active, and the $\nu(\kappa)$ -rule is only applied to formulas that actually appear in Γ_i^+ , which certainly means all variables of that formula are active. Applications of right weakening or cuts do not introduce any new ordinal variables at all.

The proof of item (4) is easy, since the only steps in the construction of Γ_{i+1} that removes formulas in Γ_i are applications of weakening where we pick the minimal formulas. \square

Given an index i let $\mathbf{u}(\Gamma)_i$ denote the set $\{\mathbf{u}(\varphi) \mid \varphi \in \Gamma_i^+\}$. It is clear from Claim 4 that $\mathbf{u}(\Gamma)_i \subseteq \mathbf{u}(\Gamma)_j$ whenever $i \leq j$. So since the set of underlying plain formulas in Γ_i^+ is bounded by the size of the closure of the root formula ρ , there is some index i with $\mathbf{u}(\Gamma)_i = \mathbf{u}(\Gamma)_j$ for all $j \geq i$. We may as well assume $i = 0$ since otherwise we can just consider a final segment of β suitably re-indexed. With this assumption in place we can now prove:

Claim 5. For every growth point i of β , either there is some formula $\varphi \in \Gamma_i^+$ and some formula φ' in $\Gamma_{i+1}^+ \setminus \Gamma_i^+$ with $\varphi' \prec_{\mathcal{O}_{i+1}} \varphi$, or Γ_{i+1} has strictly fewer defects than Γ_i .

Proof of Claim 5. Let i be a growth point. We make a case distinction according to how Γ_{i+1} follows a fix of some defect. There are the following possible cases to consider:

Case Γ_i is the conclusion of an application of the $x : \kappa$ -rule in which the principal formula $\nu x \varphi$ is a defect. Then Γ_{i+1}^+ contains the formula $\nu x^\kappa \varphi \prec_{\mathcal{O}(\Gamma_{i+1})} \nu x \varphi$. Furthermore, $\nu x^\kappa \varphi$ cannot be in Γ_i^+ because then $\nu x \varphi$ would not be minimal.

Case Γ_i is the conclusion of an application of the $\nu(\kappa)$ -rule in which the principal formula $\nu x^\kappa \varphi$ is a defect. The premise of this rule application then contains $\varphi[\nu x^{\kappa'} \varphi/x]$ where $\kappa' <_{\mathcal{O}_{i+1}} \kappa$. Let ψ be the unique formula in Γ_i such that $\mathbf{u}(\psi) = \mathbf{u}(\varphi[\nu x^{\kappa'} \varphi/x])$. We cannot have $\psi \preceq_{\mathcal{O}_{i+1}} \varphi[\nu x^{\kappa'} \varphi/x]$, since then $\nu x^\kappa \varphi$ was not a defect of Γ_i . So $\varphi[\nu x^{\kappa'} \varphi/x] \prec_{\mathcal{O}_{i+1}} \psi$, and since $\psi \in \Gamma_i^+$ we are done.

Case Γ_i is the conclusion of an application of the minimisation rule, and Γ_{i+1} is a descendant of the left premise. Since no formulas in Γ_i^+ or Γ_i^- are removed in the left premise, and no new active variables are introduced, we see that the left premise is already a propositionally saturated small pre-tile and therefore equal to Γ_{i+1} . So in this case the defect is fixed and no new defects are introduced, meaning that Γ_{i+1} has strictly fewer defects than Γ_i .

Case Γ_i is the conclusion of an application of the minimisation rule, and Γ_{i+1} is a descendant of the right premise. The defect that is to be fixed is of the form $\varphi(\nu x^\kappa \psi)$, where all variables of this formula are active in Γ_i , $\theta \in \Gamma_i^+$ for some $\theta \preceq_{\mathcal{O}_i} \varphi(\nu x^\kappa \psi)$, and the right premise contains $\varphi(\nu x^{\kappa'} \psi)$ for some κ' with $\kappa' <_{\mathcal{O}_{i+1}} \kappa$. Either $\varphi(\nu x^{\kappa'} \psi) \in \Gamma_{i+1}^+$, or it is replaced by some $\prec_{\mathcal{O}_{i+1}}$ -smaller formula $\theta' \in \Gamma_{i+1}^+$. Note that in the latter case we cannot have $\theta \prec_{\mathcal{O}_{i+1}} \theta'$, since then θ' would not be minimal. By linearity of $\preceq_{\mathcal{O}_{i+1}}$ we get $\theta' \preceq_{\mathcal{O}_{i+1}} \theta$. If $\theta' = \theta$ then clearly Γ_{i+1} contains strictly fewer defects than Γ_i since no new formulas were introduced, and in the other case we have $\theta' \prec_{\mathcal{O}_{i+1}} \theta$ as required.

This concludes the proof of the Claim. \square

Let φ be the underlying plain formula of some formula in Γ_0^+ . Say that φ is *improved* at index i if, for φ' denoting the unique formula in Γ_i^+ with $u(\varphi') = \varphi$ and φ'' denoting the unique formula in Γ_{i+1}^+ with $u(\varphi'') = \varphi$, we have $\varphi'' \prec_{\mathcal{O}_{i+1}} \varphi'$.

From Claim 5 and the pigeonhole principle (there is only a fixed finite number of underlying plain formulas), it follows that some formula is improved at infinitely many indices. This gives us an infinite \prec -descending chain of formulas on β and therefore, by considering the highest ranking variable in φ whose annotation changes infinitely many times, an infinite \prec -descending chain of ordinal variables. By Proposition 5.6, this ensures that we find an infinite \prec -descending chain of ordinal variables on β .

Finally, from Claim 4 it clearly follows that $\Gamma_{i+1} \preceq \Gamma_i$ for all i . Together with Proposition 5.15, it follows that $\Psi \preceq \Gamma$ for any leaf Ψ in the proof tree Π that we have constructed. \square

As a corollary we obtain

Proposition 5.20. Let $\mathcal{O} : \Gamma$ be any small ρ -tile and suppose $[a]\varphi \in \Gamma$. Then there is a (possibly infinite) slim proof tree Π , in which every leaf Ψ is a non-trivial small a -successor of Γ containing φ' for some $\varphi' \prec_{\mathcal{O}(\Psi)} \varphi$, and every infinite branch has an infinite \prec -descending chain of ordinal variables.

Proof. We first note that, where $\{[a]\varphi\} \cup \{\langle a \rangle \psi \mid \psi \in \Psi\} \subseteq \Gamma$, the following is an instance of the modal rule:

$$\frac{\mathcal{O} : [a] \bigvee \Gamma, \varphi, \Psi}{\mathcal{O} : \Gamma}$$

Now, apply Proposition 5.19 to the sequent $\mathcal{O} : [a] \bigvee \Gamma, \varphi, \Psi$, producing a slim proof tree in which every leaf is labelled by some sequent of the form $\mathcal{O}' :$

$[\check{a}] \vee \Gamma, \Theta$ where $\mathcal{O}' : \Theta$ is an a -successor of Γ . Finally, for each such leaf labelled $\mathcal{O}' : [\check{a}] \vee \Gamma, \Theta$, if the a -successor $\mathcal{O}' : \Theta$ is non-trivial then extend it by an application of weakening, leaving $\mathcal{O}' : \Theta$ as the label of the new leaf, and otherwise, plug in a wellfounded proof of $\mathcal{O}' : [\check{a}] \vee \Gamma, \Theta$. The resulting proof tree now has the properties required by the proposition. \square

Next we draw some consequences of the saturation conditions of a tile. The following proposition shows that an unprovable tile behaves in accordance with the local conditions for a counter-model.

Proposition 5.21. Let Γ be a small tile which does not have a slim proof. The following properties hold.

1. If $\varphi \wedge \psi \in \Gamma^+$ then $\varphi' \in \Gamma^+$ for some $\varphi' \preceq_{\mathcal{O}(\Gamma)} \varphi$ or $\psi' \in \Gamma^+$ for some $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi$.
2. If $\varphi \vee \psi \in \Gamma$ then $\varphi' \in \Gamma^+$ for some $\varphi' \preceq_{\mathcal{O}(\Gamma)} \varphi$ and $\psi' \in \Gamma^+$ for some $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi$.
3. If $\mu x \varphi \in \Gamma$ then $\theta \in \Gamma^+$ for some $\theta \preceq_{\mathcal{O}(\Gamma)} \varphi[\mu x \varphi/x]$.

Proof. For item (1), suppose that $\varphi \wedge \psi \in \Gamma^+$ but there is no $\varphi' \preceq_{\mathcal{O}(\Gamma)} \varphi$ with $\varphi' \in \Gamma^+$ and no $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi$ with $\psi' \in \Gamma$. Then since all ordinal variables in both φ and ψ are active in Γ , we get $\bar{\varphi} \in \Gamma$ and $\bar{\psi} \in \Gamma$. Thus we can provide a slim proof of Γ as follows:

$$\frac{\frac{\text{Proposition 5.8}}{\mathcal{O}(\Gamma) : \varphi, \bar{\varphi}} \quad \frac{\text{Proposition 5.8}}{\mathcal{O}(\Gamma) : \psi, \bar{\psi}}}{\mathcal{O}(\Gamma) : \varphi, \bar{\varphi}, \bar{\psi}} \text{rw} \quad \frac{\mathcal{O}(\Gamma) : \psi, \bar{\psi}}{\mathcal{O}(\Gamma) : \psi, \bar{\varphi}, \bar{\psi}} \text{rw} \quad \wedge}{\frac{\mathcal{O}(\Gamma) : \varphi \wedge \psi, \bar{\varphi}, \bar{\psi}}{\mathcal{O}(\Gamma) : \Gamma} \text{rw}}$$

The proofs of items (2) and (3) are similar. \square

Next, we prove two crucial propositions about a -successors of a tile, both involving the “backwards modalities” for the converse action \check{a} .

Proposition 5.22. Let Γ be a small ρ -tile and let Ψ be a non-trivial small a -successor of Γ . Suppose $\langle \check{a} \rangle \psi \in \Psi^+$, and suppose every ordinal variable in ψ is active in Γ . Then $\psi' \in \Gamma^+$ for some $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi$.

Proof. Since $\langle \check{a} \rangle \psi \in \Psi^+$ the underlying plain formula of ψ belongs to the closure of ρ , and by assumption all its ordinal variables are active in Γ . Hence, since Γ is a tile, either $\psi' \in \Gamma^+$ for some $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi$ or $\bar{\psi} \in \Gamma^-$. But in the latter case,

we can easily construct a wellfounded proof of the sequent $\mathcal{O}(\Psi) : [\check{a}] \vee \Gamma, \Psi$:

$$\frac{\text{Proposition 5.8}}{\frac{\mathcal{O}(\Psi) : \psi, \bar{\psi}}{\mathcal{O}(\Psi) : \psi, \vee \Gamma} \text{ mod}} \frac{\mathcal{O}(\Psi) : \langle \check{a} \rangle \psi, [\check{a}] \vee \Gamma}{\mathcal{O}(\Psi) : \Psi, [\check{a}] \vee \Gamma}$$

This contradicts the assumption that Ψ was a non-trivial a -successor. \square

The next proposition is more subtle, and will play a key role in our completeness proof for slim proofs. First a definition:

Definition 5.23. Let Γ be a small ρ -tile and let Ψ be a non-trivial small a -successor of Γ . Suppose $\langle \check{a} \rangle \psi \in \Psi^+$ and suppose x is the highest ranking ν -variable of ψ such that $\lambda = o_\psi(x)$ is not active in Γ , assuming such a variable exists. Let κ be an active variable of Γ such that $\lambda \prec_{\mathcal{O}(\Psi)} \kappa$. Let $\sigma : \text{Act}(\Psi) \rightarrow \text{Act}(\Gamma)$ be some substitution on ordinal variables such that:

1. $\sigma(\xi) = \xi$ for each ξ that is active in Γ ,
2. $\sigma(\lambda) \prec_{\mathcal{O}(\Gamma)} \kappa$,
3. $\psi' \in \Gamma^+$ for some $\psi' \preceq_{\mathcal{O}(\Gamma)} \psi[\sigma]$.

Then we say that the substitution σ *preserves progression from κ to $\langle \check{a} \rangle \psi$* (relative to Γ and Ψ). If φ is a formula in Γ^+ and σ preserves progression from $o_\varphi(x)$ to $\langle \check{a} \rangle \psi$, then we say the substitution σ *preserves progression from φ to $\langle \check{a} \rangle \psi$* .

Proposition 5.24. Let Γ be a small ρ -tile and let Ψ be a non-trivial small a -successor of Γ . Suppose $\langle \check{a} \rangle \psi \in \Psi^+$. Let o be any annotation whose range is contained in $\text{Act}(\Gamma)$ and suppose x is a fixpoint variable such that $o_\psi \prec_{\mathcal{O}(\Psi)}^x o$ and $o_\psi(y) \in \text{Act}(\Gamma)$ for all y higher ranking or equal to x . Then there is a formula θ such that:

- $\theta \in \Gamma^+$,
- $u(\theta) = u(\psi)$, and
- $o_\theta \prec_{\mathcal{O}(\Gamma)}^y o$ for some y higher ranking or equal to x .

Proof. We first define a substitution $\tau : \text{Act}(\Psi) \rightarrow \text{Act}(\Gamma)$ as follows: for each variable $\lambda \in \text{Act}(\Psi)$, τ maps λ to its $\prec_{\mathcal{O}(\Psi)}$ -smallest ancestor in $\text{Act}(\Gamma)$. Note that this is well-defined by definition of an a -successor, and note that $\tau(\kappa) = \kappa$ for $\kappa \in \text{Act}(\Gamma) \cap \text{Act}(\Psi)$, and therefore $\tau(o_\psi(y)) = o_\psi(y) = o(y)$ for all y higher ranking or equal to x .

Claim 6. There is some formula $\theta \in \Gamma^+$ such that $\theta \preceq_{\mathcal{O}(\Gamma)} \psi[\tau]$.

Proof of Claim 6. Otherwise we have $\overline{\psi[\tau]} \in \Gamma^-$, and we can derive $\mathcal{O}(\Psi) : \Psi, [\check{a}] \vee \Gamma$ as follows.

$$\frac{\frac{\text{Proposition 5.7}}{\mathcal{O}(\Psi) : \psi, \overline{\psi[\tau]}} \ (\tau \text{ increasing})}{\mathcal{O}(\Psi) : \psi, \vee \Gamma} \text{ (basic modal logic)}$$

$$\frac{\mathcal{O}(\Psi) : \langle \check{a} \rangle \psi, [\check{a}] \vee \Gamma}{\mathcal{O}(\Psi) : \Psi, [\check{a}] \vee \Gamma} \quad \square$$

Since $o_\psi(x) \prec_{\mathcal{O}(\Psi)} o(x)$ we have two different cases to consider.

Case 1: $o_\psi(x)$ is to the left of $o(x)$. Since Ψ was an a -successor, we have $o_\psi(x) \in L$ for some L such that $\mathcal{O}(\Psi) = \mathcal{O}(\Gamma) \uplus L$. By Proposition 5.18 there is some variable $\kappa \in \text{Act}(\Gamma)$ such that $o_\psi(x) \prec_{\mathcal{O}(\Psi)} \kappa$ and κ is to the left of $o(x)$ in $\mathcal{O}(\Psi)$ (equivalently in $\mathcal{O}(\Gamma)$). By definition of τ it follows that $\tau(o_\psi(x))$ is also to the left of $o(x)$ in $\mathcal{O}(\Gamma)$. Hence $\tau(o_\psi(x)) \prec_{\mathcal{O}(\Gamma)} o(x)$. By Claim 6 there is some formula $\theta \in \Gamma^+$ such that $\theta \preceq_{\mathcal{O}(\Gamma)} \psi[\tau]$. It easily follows that $o_\theta \prec_{\mathcal{O}(\Gamma)}^y o$ for some y higher ranking or equal to x .

Case 2: $o_\psi(x) \prec_{\mathcal{O}(\Psi)} o(x)$. Let us write ψ as $\alpha(\nu x^{o_\psi(x)} \beta)$ so that $\psi[\tau]$ is $\alpha[\tau](\nu x^{\tau(o_\psi(x))} \beta[\tau])$.

Claim 7. $\exists \xi \prec o(x) : \overline{\alpha[\tau]}(\mu x^\xi \beta[\tau]) \notin \Gamma^\mu$.

Proof of Claim 7. By assumption we have $o_\psi(x) \prec_{\mathcal{O}(\Psi)} o(x)$, so if the displayed formula were in Γ^μ then we could prove $\mathcal{O}(\Psi) : \Psi, [\check{a}] \vee \Gamma$ as follows:

$$\frac{\frac{\text{Proposition 5.7}}{\mathcal{O}(\Psi) : \alpha(\nu x^{o_\psi(x)} \beta), \overline{\alpha[\tau]}(\mu x^{o_\psi(x)} \beta[\tau])} \ \tau \text{ increasing}}{\mathcal{O}(\Psi) : \alpha(\nu x^{o_\psi(x)} \beta), \exists \xi \prec o(x) : \overline{\alpha[\tau]}(\mu x^\xi \beta[\tau])} \ o_\psi(x) \prec_{\mathcal{O}(\Psi)} o(x)}{\mathcal{O}(\Psi) : \alpha(\nu x^{o_\psi(x)} \beta), \vee \Gamma} \text{ (basic modal logic)}$$

$$\frac{\mathcal{O}(\Psi) : \langle \check{a} \rangle \alpha(\nu x^{o_\psi(x)} \beta), [\check{a}] \vee \Gamma}{\mathcal{O}(\Psi) : \Psi, [\check{a}] \vee \Gamma} \quad \square$$

Since $\theta \in \Gamma^+$ and $\theta \preceq_{\mathcal{O}(\Psi)} \psi[\tau]$, the saturation condition of a tile with respect to the minimisation rule ensures that there is some $\lambda \prec_{\mathcal{O}(\Gamma)} o(x)$ and some formula $\theta' \in \Gamma^+$ such that $\theta' \preceq_{\mathcal{O}(\Gamma)} \alpha[\tau](\nu x^\lambda \beta[\tau])$. Let o' be the annotation of the formula $\alpha[\tau](\nu x^\lambda \beta[\tau])$, so that $o'(x) = \lambda \prec_{\mathcal{O}(\Gamma)} o(x)$, hence $o'(x) \prec_{\mathcal{O}(\Gamma)} o(x)$. Furthermore, for all y ranking higher than or equal to x we have:

$$\begin{aligned} o'(y) &= o_{\psi[\tau]}(y) \\ &= \tau(o_\psi(y)) \\ &= o(y) \end{aligned}$$

where the last inequality is due to $o_\psi(y)$ being active in Γ (since x was the highest ranking variable with $o_\psi(x)$ non-active in Γ). Now let o'' be the annotation of

Position	Player	Moves
ρ	Refuter	$\{\Gamma \in T_\rho \mid \rho \in \Gamma^+\}$
$\Gamma \in T_\rho$	Prover	$\{(\Gamma, [a]\varphi) \mid [a]\varphi \in \Gamma^+\}$
$(\Gamma, [a]\varphi) \in T_\rho \times \text{Clos}(\rho)$	Refuter	$\{\Psi \in T_\rho \mid \Gamma \rightarrow_{a,\varphi} \Psi\}$

Table 3: Moves in the mosaic game

θ' . Since $\theta' \preceq_{\mathcal{O}(\Gamma)} \alpha[\tau](\nu x^\lambda \beta[\tau])$ it follows by definition that $o'' \prec_{\mathcal{O}(\Gamma)}^y o'$ for some y . Either y is higher ranking than or equal to x in which case $o''(y) \prec_{\mathcal{O}(\Gamma)} o'(y) \preceq_{\mathcal{O}(\Gamma)} o(x)$, or y is lower ranking than x in which case $o''(x) = o'(x) \prec_{\mathcal{O}(\Gamma)} o(x)$. In either case the desired conclusion follows. \square

5.4 Completeness of slim proofs

The goal of this section is to show that every valid formula ρ has a slim proof. We will do so using a two-player game that we call the *mosaic game*, played between two players called “Prover” and “Refuter”. Our strategy is as follows: We first prove the *determinacy* of the mosaic game, that is, the fact that either Prover or Refuter has a winning strategy. It then suffices to show that Refuter cannot have a winning strategy if ρ is valid, and ρ has a slim proof if the mosaic game is a win for Prover.

We now define the mosaic game for the given formula ρ . The two players in the mosaic game are called Prover and Refuter. There are three kinds of positions in this game: the fixed formula ρ , used as the starting position, ρ -tiles and pairs $(\Gamma, [a]\varphi)$ where Γ is a ρ -tile and $[a]\varphi \in \Gamma^+$. Ownership of positions and legal moves of the game are specified in Table 3, where we recall that T_ρ denotes the set of ρ -tiles, and we write $\Gamma \rightarrow_{a,\varphi} \Psi$ if Ψ is an a -successor of Γ , relative to φ (cf. Definition 5.17). In words, the mosaic game for ρ starts at the formula ρ itself, which is a position for Refuter. After Refuter’s initial move, which consists of a ρ -tile containing the formula ρ , the two players take turns. At a ρ -tile Γ , Prover has to select a formula $[a]\varphi \in \Gamma^+$, and at a position of the form $(\Gamma, [a]\varphi)$, Refuter has to pick a witness in the form of an a -successor of Γ , relative to φ .

Prover wins an infinite play if the play contains an infinite descending \prec -chain of ordinal variables, otherwise Refuter wins the play. More formally, if $(\mathcal{O}_i : \Gamma_i)_{i < \omega}$ enumerates the ρ -tiles occurring in an infinite play, then this play is won by Prover iff there exists $k < \omega$ and an infinite sequence of ordinal variables $(\kappa_i)_{i < \omega}$ such that $\kappa_{i+1} \preceq_{\mathcal{O}_{k+i}} \kappa_i$ for all i , and $\kappa_{i+1} \prec_{\mathcal{O}_{k+i}} \kappa_i$ for infinitely many i . (Note that if Γ_0, Γ_1 are tiles such that Γ_1 is reachable from Γ_0 via some partial play in the mosaic game, then the constraint of Γ_1 contains that of Γ_0 . This is a direct consequence of the definition of an a -successor and the relation \preceq_L over pre-tiles.) Finite plays are lost by the player who got stuck.

Proposition 5.25. The mosaic game is determined.

Proof. Fix a formula ρ and an injection of the set of positions into the natural numbers, so that every play in the mosaic game for ρ corresponds to a unique function from \mathbb{N} to \mathbb{N} , i.e., an element of the Baire space $\mathbb{N}^{\mathbb{N}}$. Let $P \subseteq \mathbb{N}^{\mathbb{N}}$ be the set of sequences in which Refuter makes an illegal move and (moreover) the first illegal move. Furthermore, let $W \subseteq \mathbb{N}^{\mathbb{N}}$ describe the set of infinite plays in the mosaic game for ρ that are winning for Prover.

Recall that Prover and Refuter take alternating turns in the mosaic game; it is then not hard to see that Prover has a winning strategy in the mosaic game for ρ iff Player I (i.e., the player that plays first) has a winning strategy in the Gale–Stewart game on the Baire space with winning set $P \cup W$. We claim that this winning set is Borel, whence determinacy of the mosaic game is a consequence of Martin’s Borel Determinacy Theorem [13]. Being a union of open sets, P is clearly Borel. We show that W is also Borel.

Given a play π , say that an ordinal variable κ is *covered* in π if it has some active descendant in every constraint in π in which κ occurs. Note that for each ordinal variable κ , the set of plays in which κ is covered is Borel. Let $\Gamma_0, \dots, \Gamma_{n-1}$ be a sequence of n tiles from a given play π in the mosaic game, with constraints $\mathcal{O}_0, \dots, \mathcal{O}_{n-1}$, such that $\Gamma_{n-1} \preceq \dots \preceq \Gamma_0$. We call this sequence a *stable progression of length n* if there are variables $\kappa_0, \dots, \kappa_{n-1}$ each of which is covered in π , such that:

- $\kappa_i \in \text{Act}(\mathcal{O}_i)$ for each $i < n$,
- $\kappa_j <_{\mathcal{O}_j} \kappa_i$ for $0 \leq i < j < n$.

Let W_n be the set of infinite plays that contain a stable progression of length n . Clearly each set W_n is Borel, so the set $\bigcap_{n \in \omega} W_n$ is Borel. We claim that $W = \bigcap_{n \in \omega} W_n$.

For the left-to-right inclusion: if $\pi \in W$ then it has an infinite \prec -descending chain. Repeating the proof of Proposition 5.6, we can show that π in fact has an infinite \prec -descending chain. Since as a sequence of tiles the play π is decreasing with respect to the order \preceq , by definition of the relation \preceq on tiles (Definition 5.14), it is not hard to see that every ordinal variable in this infinite \prec -descending chain has to be covered. Hence π contains stable progressions of arbitrary length.

For right-to-left, suppose π has a progression of length n , for all $n \in \omega$. We define a tree consisting of all the ordinal variables that are covered in π , where λ is a child of κ if this holds with respect to some constraint appearing in π . Note that this is a finite set of finitely branching rooted trees: since we never add \prec -ancestors to pre-existing variables (by Definition 5.14), the set of children of a variable is increasing along π , and since the number of active variables in a tile is bounded this means there is also a bound on the size of \prec -anti-chains of covered variables in any constraint in π . We can view the structure as a single finitely branching tree by appending a new root. The resulting tree is infinite since π contains stable progressions of arbitrary length, and so by König’s lemma it contains an infinite branch. This infinite branch is

an infinite $<$ -descending chain of ordinal variables in π , and therefore also an infinite \prec -descending chain. \square

With Proposition 5.25 in place, we want to prove that a winning strategy for Prover yields a slim proof, and on the other hand, a winning strategy for Refuter yields a counter-model. We begin with the first statement.

Proposition 5.26. If Prover has a winning strategy in the mosaic game for ρ , then the formula ρ has a slim proof.

Proof. Let σ be a winning strategy for Prover in the mosaic game for ρ . We show how to construct a slim proof for an arbitrary small ρ -tile $\Gamma \preceq \{\rho\}$ as Proposition 5.19 ensures this is sufficient to establish that ρ admits a slim proof. Consider the tree of maximal plays consistent with σ starting from Γ . By assumption, all such plays are winning for Prover. We now let Π be the tree of positions owned by Prover in these plays. These positions are precisely the small ρ -tiles in the σ -plays starting from Γ . We can thus view Π as a derivation with root Γ in the calculus consisting only of rules

$$\frac{\{\Psi \in T_\rho \mid \Phi \rightarrow_{a,\varphi} \Psi\}}{\Phi} [a]\varphi \in \Phi$$

where $\rightarrow_{a,\varphi}$ is the relation of a -successor relative to φ from Definition 5.17. Moreover, as σ is winning, every infinite path through the derivation contains an infinite \prec -decreasing chain of ordinal variables through the sequence of constraints on the path. All that remains is to replace each of above rules by the corresponding slim derivation guaranteed by Proposition 5.20 and Proposition 5.6: The restriction on applications of lw in slim derivations guarantees that the induced expansion of Π is a slim proof. \square

We now come to the key technical result of the paper, showing that winning strategies for Refuter indeed correspond to counter-models.

Proposition 5.27. Suppose Refuter has a winning strategy in the mosaic game for ρ . Then ρ is not valid.

Proof. Fix a winning strategy σ for Refuter, and let Γ be the tile containing ρ chosen by the first move of Refuter according to σ . We construct a counter-model $M^\sigma = (W^\sigma, R^\sigma, V^\sigma)$ as follows: W^σ is the set of σ -guided partial plays whose last position belong to Prover, i.e. the last position is a small tile. In the sequel we may write $\pi \sqsubseteq \pi'$ ($\pi \sqsubset \pi'$) if π is an initial segment of π' (a proper initial segment of π' , respectively). Given an action label a , and plays $\pi_0, \pi_1 \in W^\sigma$, we set $(\pi_0, \pi_1) \in R_a^\sigma$ if and only if either

1. π_1 extends π_0 with a choice of a formula $[a]\varphi$ by Prover, followed by a choice of a non-trivial a -successor Ψ of the last position on π_0 containing some $\varphi' \prec_{\mathcal{O}(\Psi)} \varphi$, according to the strategy σ of Refuter, or

2. π_0 extends π_1 with a choice of a formula $[\check{a}]\varphi$ by Prover, followed by a choice of a non-trivial \check{a} -successor Ψ of the last position on π_1 containing some $\varphi' \prec_{\mathcal{O}(\Psi)} \varphi$, according to the strategy σ of Refuter.

Note that $\pi_0 R_a^\sigma \pi_1$ iff $\pi_1 R_a^\sigma \pi_0$. For the valuation V^σ , we set $\pi \in V^\sigma(p)$ if and only if $\bar{p} \in \Theta$, where Θ is the last position of π .

Note that (Γ, ρ) is a position in the evaluation game for ρ in M^σ , since $\Gamma \in W^\sigma$ (more precisely, the partial play consisting only of the position Γ is in W^σ). We shall provide a winning strategy σ' for Falsifier in the evaluation game at this position. While constructing the strategy σ' we shall inductively maintain the condition that for any σ' -guided partial play $(\pi_0, \varphi_0) \dots (\pi_n, \varphi_n)$, the partial plays π_0, \dots, π_n are σ -guided, and furthermore there is a sequence of positively annotated formulas $(\varphi'_0, \dots, \varphi'_n)$ such that, if we denote the last sequent on each play π_i as Θ_i and its constraint set as \mathcal{O}_i , then:

1. $\varphi'_i \in \Theta_i^+$.
2. $u(\varphi'_i) = \varphi_i$.
3. If $i < n$, $\varphi_i = \psi_0 \vee \psi_1$ and $\varphi_{i+1} = \psi_j$ where $j \in \{0, 1\}$, then $\pi_i = \pi_{i+1}$, φ'_i is of the form $\psi'_0 \vee \psi'_1$ and $\varphi'_{i+1} \preceq_{\mathcal{O}_i} \psi'_j$.
4. If $i < n$, $\varphi_i = \psi_0 \wedge \psi_1$ and $\varphi_{i+1} = \psi_j$ where $j \in \{0, 1\}$, then $\pi_i = \pi_{i+1}$, φ'_i is of the form $\psi'_0 \wedge \psi'_1$ and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'_j$.
5. If $i < n$, $\varphi_i = \mu x. \psi$ and $\varphi_{i+1} = \psi[\mu x. \psi/x]$, then $\pi_i = \pi_{i+1}$, φ'_i is of the form $\mu x. \psi'$ and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'[\mu x. \psi'/x]$.
6. If $i < n$, $\varphi_i = \nu x. \psi$ and $\varphi_{i+1} = \psi[\nu x. \psi/x]$, then $\pi_i = \pi_{i+1}$, φ'_i is of the form $\nu x^{\kappa} \theta$ and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \theta[\nu x^{\kappa_0} \theta/x]$ for some κ_0 such that $\kappa_0 <_{\mathcal{O}_{i+1}} \kappa$.
7. If $i < n$, $\varphi_i = [a]\psi$ and $\varphi_{i+1} = \psi$, then $\pi_i \sqsubset \pi_{i+1}$, $(\pi_i, \pi_{i+1}) \in R_a^\sigma$, and φ'_i is of the form $[a]\psi'$ where $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'$.
8. If $i < n$, $\varphi_i = \langle a \rangle \psi$, $\varphi_{i+1} = \psi$, then $(\pi_i, \pi_{i+1}) \in R_a^\sigma$, φ'_i is of the form $\langle a \rangle \psi'$ and either
 - (a) $\pi_i \sqsubset \pi_{i+1}$ and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'$, or
 - (b) $\pi_{i+1} \sqsubset \pi_i$, and letting k be the greatest index less than $i + 1$ with $\pi_k = \pi_{i+1}$, we have, either:
 - i. all ordinal variables in ψ' are active in Θ_{i+1} , and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'$, or
 - ii. some ordinal variable in ψ' is not active in Θ_{i+1} , for the highest ranking fixpoint variable x with $o_{\psi'}(x) \notin \text{Act}(\Theta_{i+1})$ it does not hold that $o_{\psi'} \prec_{\mathcal{O}_i}^x o_{\varphi'_k}$, and $\varphi'_{i+1} \preceq_{\mathcal{O}_{i+1}} \psi'[\tau]$ for some substitution τ with $\tau(\kappa) = \kappa$ for all $\kappa \in \text{Act}(\Theta_{i+1})$, or
 - iii. some ordinal variable in ψ' is not active in Θ_{i+1} , for the highest ranking fixpoint variable x with $o_{\psi'}(x) \notin \text{Act}(\Theta_{i+1})$ it holds that $o_{\psi'} \prec_{\mathcal{O}_i}^x o_{\varphi'_k}$ where x is the highest ranking, and for some y higher ranking or equal to x we have $o_{\varphi_{i+1}} \prec_{\mathcal{O}_{i+1}}^y o_{\psi'}$.

Suppose the strategy σ' has been defined on all partial plays of length $< n$, and let $(\pi_0, \varphi_0) \dots (\pi_n, \varphi_n)$ be a play of length n . We show how to define the strategy σ' on this play, if the last position belongs to Falsifier, and we show how to maintain the inductive invariant for each σ' -guided play of length $n + 1$ extending $(\pi_0, \varphi_0) \dots (\pi_n, \varphi_n)$. The construction is carried out on a case-by-case basis depending on the shape of the last position.

Case φ_n is of the form p or \bar{p} for some propositional variable p . Then there are no available moves, so we need not define the strategy σ' here and the inductive invariant is trivially maintained for all extensions of the play.

Case φ_n is of the form $\alpha \wedge \beta$. This position belongs to Falsifier. By the induction hypothesis $\varphi'_n \in \Theta_n^+$, and $u(\varphi'_n) = \alpha \wedge \beta$. Hence φ'_n is of the form $\alpha' \wedge \beta'$ where $u(\alpha') = \alpha$ and $u(\beta') = \beta$. By Proposition 5.21, $\alpha'' \in \Theta_n^+$ for some $\alpha'' \preceq_{\mathcal{O}_n} \alpha'$ or $\beta'' \in \Theta_n^+$ for some $\beta'' \preceq_{\mathcal{O}_n} \beta'$ (or both). In the first case we define σ' so that Falsifier chooses (π_n, α) and set $\varphi'_{n+1} = \alpha''$, and in the second case we define σ' so that Falsifier chooses (π_n, β) and set $\varphi'_{n+1} = \beta''$ (and if both cases hold the choice can be defined arbitrarily). The inductive invariant is clearly maintained.

Case φ_n is of the form $\alpha \vee \beta$. This position belongs to Verifier. By the induction hypothesis $\varphi'_n \in \Theta_n^+$, and $u(\varphi'_n) = \alpha \vee \beta$. Hence φ'_n is of the form $\alpha' \vee \beta'$ where $u(\alpha') = \alpha$ and $u(\beta') = \beta$. By Proposition 5.21, $\alpha'' \in \Theta_n^+$ for some $\alpha'' \preceq_{\mathcal{O}_n} \alpha'$ and $\beta'' \in \Theta_n^+$ for some $\beta'' \preceq_{\mathcal{O}_n} \beta'$. So clearly, given any extension of the play $(\pi_0, \varphi_0) \dots (\pi_n, \varphi_n)$ in which Verifier chooses (π_n, γ) with $\gamma \in \{\alpha, \beta\}$, we can choose γ' so that the inductive invariant is maintained.

Case φ_n is of the form $[a]\alpha$. This position belongs to Falsifier. The shadow formula φ'_n is of the form $[a]\alpha'$ where $u(\alpha') = \alpha$, and we have $[a]\alpha' \in \Theta_n^+$ by the induction hypothesis. Let Prover play the formula $[a]\alpha'$, and let Ψ be the non-trivial small a -successor of Θ_n chosen by Refuter in response to this move according to the strategy σ . By definition of an admissible move for Refuter there is some $\alpha'' \in \Psi^+$ such that $\alpha'' \preceq_{\mathcal{O}(\Psi)} \alpha'$. We set σ' so that Falsifier chooses $(\pi_n \cdot \Psi, \alpha)$. Since $\alpha'' \in \Psi$ and $u(\alpha'') = u(\alpha') = \alpha$, the inductive invariant is maintained.

Case φ_n is of the form $\langle a \rangle \alpha$. This position belongs to Verifier, and we have to show that the inductive invariant is maintained for every possible move of Verifier. Every choice of Verifier at this position is of the form (π', α) where one of the following two cases holds: (i) π' is of the form $(\pi_n \cdot \Psi)$ where Ψ is a non-trivial a -successor of Θ_n , or (ii) π_n is of the form $\pi' \cdot \Theta_n$ where Θ_n is a non-trivial \bar{a} -successor of the last tile on π' . In either case, φ'_n is of the form $\langle a \rangle \beta$ where $u(\beta) = \alpha$. In case (i), we must have some $\alpha' \in \Psi^+$ with $\alpha' \preceq_{\mathcal{O}(\Psi)} \beta$ by definition of an a -successor, and the inductive invariant is maintained. In case (ii), denote the last tile on π' by Θ' . Clearly, this means that there must be some index $k < n$ for which $\pi_k = \pi'$, so let k be the largest such index. If all ordinal variables of β are active in Θ' then pick α' to be some formula in $(\Theta')^+$ with $\alpha' \preceq_{\mathcal{O}(\Theta')} \beta$, which exists by Proposition 5.22. If some ordinal variable of β is non-active in Θ' then let x be the highest ranking variable with $o_\beta(x) \notin \text{Act}(\Theta')$. If it does not hold that $o_\beta \prec_{\mathcal{O}(\Theta_n)}^x o_{\varphi'_k}$ then let the substitution τ be as in the

proof of Proposition 5.24. Then we pick φ'_{n+1} so that $\varphi'_{n+1} \preceq_{\mathcal{O}(\Theta_n)} \beta[\tau]$ as provided in the same proof. Finally, if it holds that $o_\beta \prec_{\mathcal{O}(\Theta_n)}^x o_{\varphi'_k}$ then pick φ'_{n+1} so that $o_{\varphi'_{n+1}} \prec_{\mathcal{O}(\Theta')}^y o_{\varphi'_k}$ for some y higher ranking or equal to x , as provided by Proposition 5.24.

Case φ_n is of the form $\nu x \alpha$. Then the unique next position in any extension of this play is $(\pi_n, \alpha[\nu x \alpha/x])$. The formula $\varphi'_n \in \Theta_n^+$ is $\nu x^\kappa \beta$ for some κ and some β with $u(\beta) = \alpha$. By the definition of a tile there is some κ_0 such that $\kappa_0 <_{\mathcal{O}(\Theta_n)} \kappa$ and there is some formula $\gamma \in \Theta_n^+$ such that $\gamma \preceq_{\mathcal{O}(\Theta_n)} \beta[\nu x^{\kappa_0} \beta/x]$. Since clearly $u(\gamma) = u(\beta[\nu x^{\kappa_0} \beta/x]) = \alpha[\nu x \alpha/x]$, the invariant is maintained.

Case φ_n is of the form $\mu x \alpha$. Then the unique next position in any extension of this play is $(\pi_n, \alpha[\mu x \alpha/x])$. The formula $\varphi'_n \in \Theta_n$ is $\mu x \beta$ for some β with $u(\beta) = \alpha$. By the definition of a tile we have $\gamma \preceq_{\mathcal{O}_n} \beta[\mu x \beta/x]$ for some $\gamma \in \Theta_n$. Since clearly $u(\gamma) = u(\beta[\mu x \beta/x]) = \alpha[\mu x \alpha/x]$, the invariant is maintained.

We now show that σ' is indeed a winning strategy for Falsifier. First we show that every full *finite* σ' -guided play is won by Falsifier. The last position of such a play must take one of the following forms: (i) (π, p) for some propositional variable p with $\pi \in V^\sigma(p)$, (ii) (π, \bar{p}) for some propositional variable p with $\pi \notin V^\sigma(p)$, or (iii) $(\pi, [a]\varphi)$ with π not having an a -successor.

Now the latter case cannot occur — this is a more or less direct consequence of our construction of the relation R_a^σ , based on Refuter's winning strategy σ in the mosaic game. In the other two cases we reason as follows, denoting the last tile on π by Θ . If the last position of the play is (π, p) and $\pi \in V^\sigma(p)$ then $\bar{p} \in \Theta$ by definition of V^σ . But $p \in \Theta$ by the invariant of σ' -guided plays, so we get an instance of excluded middle which contradicts the definition of a non-trivial tile. On the other hand, if the last position is (π, \bar{p}) and $\pi \notin V^\sigma(p)$ then $\bar{p} \notin \Theta$ by definition. This contradicts the invariant of σ' -guided plays, which says that $\bar{p} \in \Theta$. So in each case the assumption that the finite full play is lost by Falsifier leads to a contradiction, and we have established that Falsifier wins every full finite σ -guided play.

In the remainder of the proof we will argue that every *infinite* σ' -guided play is won by Falsifier as well. Let $(\pi_i, \varphi_i)_{i < \omega}$ be an infinite σ' -guided play, and assume for a contradiction that this play is lost by Falsifier. As an immediate consequence of the following claim (cf. our discussion of the evaluation game in section 2.3), this means that the highest ranking variable that is unfolded infinitely many times in the play is a ν -variable.

Claim 8. Let $(\varphi_i)_{i < \omega}$ be an infinite trace starting from $\varphi_0 = \rho$, and let $\eta x \psi$ be its most significant formula. Then for any other formula $\lambda z \chi$ that occurs infinitely often on this trace, we have $x <_\rho z$, i.e., x ranks higher than z in the subsumption order of ρ .

Proof of Claim (sketch). While the statement of the claim is rather intuitive, the details of its proof are quite tedious. Define the relation \sqsubseteq_C on the fixpoint formulas in the set $\text{Clos}(\rho)$ by putting $\alpha \sqsubseteq_C \beta$ if there is a finite trace from α to β such that α is a subformula of every formula on the trace. It then follows

from the assumptions of the claim that the formula $\eta x \psi$ is \sqsubseteq_C -minimal among all fixpoint formulas appearing on the trace $(\varphi_i)_{i < \omega}$. Now consider Kozen's expansion map \mathbf{exp} from $Sfor(\rho)$ to $Clos(\rho)$ [12]; the key observation is that for any two bound variables $x, y \in \text{Var}(\rho)$ we have that $x <_\rho y$ iff $\mathbf{exp}(x) \sqsubseteq_C \mathbf{exp}(y)$. \square

To contradict the assumption that σ was a winning strategy for Refuter, we shall now construct an infinite σ -guided play containing an infinite progressing chain of ordinal variables, which is therefore won by Prover.

We denote the last sequent on each partial play π_i by Θ_i , and we denote $\mathcal{O}(\Theta_i)$ by \mathcal{O}_i . Consider the infinite sequence of "shadow formulas" $(\varphi'_i)_{i < \omega}$. Given a fixpoint variable x , say that the index i is a *progress point* for x if π_i is an initial segment of π_{i+1} and $\circ_{\varphi'_{i+1}}(x) \neq \circ_{\varphi'_i}(x)$. We refer to the highest ranking ν -variable that has infinitely many progress points on $(\varphi'_i)_{i < \omega}$ as the *dominant variable* of the sequence and denote it by Z . Note that at least one ν -variable has infinitely many progress points, since by clause (6) in the construction of the shadow formulas φ'_i , every index at which a ν -variable is unfolded on the play $(\pi_i, \varphi_i)_{i < \omega}$ must be a progress point either for that variable or some higher ranking variable, and since there is a ν -variable unfolded infinitely many times in the play $(\pi_i, \varphi_i)_{i < \omega}$.

Claim 9. There is an index $k < \omega$ such that:

- π_k is an initial segment of π_m for all $m > k$,
- no higher ranking ν -variable than Z has any progress point after k ,
- no higher ranking variable than Z is unfolded at any point after k .

Proof of Claim. Let k_0 be some index for which no higher ranking ν -variable than Z has any progress points after k_0 , and no higher ranking fixpoint variable than Z is unfolded after k_0 . Such an index exists by definition of Z (and since Z must be higher ranking or equal to the highest ranking variable that is unfolded infinitely many times on $(\pi_i, \varphi_i)_{i < \omega}$, since that variable is a ν -variable by assumption and has infinitely many progress points). Now consider the set of σ -guided partial plays of the form π_m with $m \geq k_0$. It is clear that this set is downwards directed with respect to the initial segment order, so we find a shortest σ -guided play π' belonging to this set. We set k to be the smallest index $\geq k_0$ for which $\pi_k = \pi'$. \square

We assume that $k = 0$ in the previous claim, since otherwise we can just re-index the play $(\pi_i, \varphi_i)_{i < \omega}$. For each index i we let ξ_i denote the ordinal variable $\circ_{\varphi'_i}(Z)$.

Claim 10. Suppose that $i \leq j < \omega$ are indices such that $\pi_i = \pi_k = \pi_j$ for all $i \leq k \leq j$. If Z has a progress point k with $i \leq k < j$ then $\xi_j \prec_{\mathcal{O}_j} \xi_i$, and otherwise $\xi_j = \xi_i$.

Proof of Claim. This is proved by induction on the difference $j - i$, the case for $j - i = 0$ being trivial. For the induction step, supposing that the induction hypothesis holds for $j - i$, consider the index $j + 1$. It suffices to prove that $\xi_{j+1} \preceq_{\mathcal{O}_j} \xi_j$; it then follows that $\xi_{j+1} \prec_{\mathcal{O}_j} \xi_j$ if j is a progress point. We have to consider several cases for the shape of φ_j and φ_{j+1} . Note that by our assumption that $\pi_i = \pi_j = \pi_{j+1}$, φ_j cannot be of the form $[a]\psi$ or $\langle a \rangle \psi$.

Case φ_j is of the form $\alpha \wedge \beta$ and φ_{j+1} is either α or β . Say $\varphi_{j+1} = \alpha$. Then the shadow formula φ'_j is of the form $\alpha' \wedge \beta'$ where $u(\alpha') = \alpha$ and $u(\beta') = \beta$, and $\varphi'_{j+1} \preceq_{\mathcal{O}_j} \alpha'$. By assumption no variable ranking higher than Z has a progress point at j , so we get:

$$\begin{aligned} \circ_{\varphi'_{j+1}}(Z) &\preceq_{\mathcal{O}_{j+1}} \circ_{\alpha'}(Z) \\ &= \circ_{\alpha' \wedge \beta'}(Z) \\ &= \circ_{\varphi'_j}(Z) \end{aligned}$$

Hence $\xi_{j+1} \preceq_{\mathcal{O}_{j+1}} \xi_j$ as required.

Case φ_j is of the form $\alpha \vee \beta$. This is similar to the previous case.

Case φ_n is of the form $\nu x \alpha$. Then the formula φ'_j is $\nu x^\kappa \beta$ for some κ and some β with $u(\beta) = \alpha$. Then there is some κ_0 such that $\varphi'_{j+1} \preceq_{\mathcal{O}_j} \beta[\nu x^{\kappa_0} \beta/x]$. Since no variable ranking higher than Z has a progress point at j , we get:

$$\begin{aligned} \circ_{\varphi'_{j+1}}(Z) &\preceq_{\mathcal{O}_{j+1}} \circ_{\beta[\nu x^{\kappa_0} \beta/x]}(Z) \\ &\preceq_{\mathcal{O}_{j+1}} \circ_{\nu x^\kappa \beta}(Z) \\ &= \circ_{\varphi'_j}(Z) \end{aligned}$$

Hence $\xi_{j+1} \preceq_{\mathcal{O}_{j+1}} \xi_j$ as required.

Case φ_n is of the form $\mu x \alpha$. This is similar to the previous case. □

Claim 11. Suppose that $i \leq j < \omega$ are indices such that for all $i \leq k < j$, π_k is an initial segment of π_{k+1} . If Z has a progress point k with $i \leq k < j$ then $\xi_j \prec_{\mathcal{O}_j} \xi_i$, and otherwise $\xi_j = \xi_i$.

Proof of Claim. This is proved by induction on the difference $j - i$, in the same manner as Claim 10. The case for $j - i = 0$ is trivial, and for the induction step we show that $\xi_{j+1} \preceq_{\mathcal{O}_j} \xi_j$. The only difference is that we now have to consider the cases where φ_j is of the form $[a]\psi$ or $\langle a \rangle \psi$. We focus on the latter case since they are treated in the same way: if $\varphi_j = \langle a \rangle \psi$ then φ'_j is of the form $\langle a \rangle \psi'$ where $u(\psi') = \psi$. Furthermore, since we assumed that π_j is an initial segment of π_{j+1} , the last tile on π_{j+1} must be an a -successor of the last tile on π_j , and $\varphi'_{j+1} \preceq_{\mathcal{O}_{j+1}} \psi'$. Hence we get:

$$\begin{aligned} \circ_{\varphi'_{j+1}}(Z) &\preceq_{\mathcal{O}_{j+1}} \circ_{\psi'}(Z) \\ &= \circ_{\langle a \rangle \psi'}(Z) \\ &= \circ_{\varphi'_j}(Z) \end{aligned}$$

So $\xi_{j+1} \preceq_{\mathcal{O}_{j+1}} \xi_j$ as required. □

Claim 12. Suppose that $i \leq j < \omega$ are indices such that $\pi_i = \pi_j$ and π_i is an initial segment of π_k for all $i < k < j$. If Z has a progress point k with $i \leq k < j$ then $\xi_j \prec_{\mathcal{O}_j} \xi_i$, and otherwise $\xi_j \preceq_{\mathcal{O}_j} \xi_i$.

Proof of Claim. We prove this by induction on the difference $j - i$. If $j - i = 0$ then $i = j$ and so $\xi_i = \xi_j$. Now suppose $j - i > 0$, and suppose the statement holds for all $i' \leq j'$ with $j' - i' < j - i$. If $\pi_i = \pi_k$ for all $i \leq k \leq j$ then the statement follows directly from Claim 10. Otherwise, it is clear that we can find indices i', j' such that $i < i' \leq j' < j$, $\pi_{i'} = \pi_{j'}$, $\pi_i = \pi_{i'-1}$, $\pi_j = \pi_{j'+1}$, $\Theta_{i'} = \Theta_{j'}$ is a non-trivial a -successor of Θ_i for some a , and the partial play $(\pi_i, \varphi_i) \dots (\pi_j, \varphi_j)$ has the shape:

$$(\pi, \varphi_i) \dots (\pi, \varphi_{i'-1}) (\pi \cdot \Theta, \varphi_{i'}) \dots (\pi \cdot \Theta, \varphi_{j'}) (\pi, \varphi_{j'+1}) \dots (\pi, \varphi_j)$$

where $\pi = \pi_i = \pi_{i'-1} = \pi_{j'+1} = \pi_j$, and $\Theta = \Theta_{i'} = \Theta_{j'}$.

We shall prove the following statements:

1. If Z has a progress point m with $i \leq m < i' - 1$ then $\xi_{i'-1} \prec_{\mathcal{O}_{i'-1}} \xi_i$ and otherwise $\xi_{i'-1} \preceq_{\mathcal{O}_{i'-1}} \xi_i$.
2. If Z has a progress point m with $j' + 1 \leq m < j$ then $\xi_j \prec_{\mathcal{O}_j} \xi_{j'+1}$ and otherwise $\xi_j \preceq_{\mathcal{O}_j} \xi_{j'+1}$.
3. If Z has a progress point m with $i' \leq m < j'$ then $\xi_{j'} \prec_{\mathcal{O}_{j'}} \xi_{i'}$ and otherwise $\xi_{j'} \preceq_{\mathcal{O}_{j'}} \xi_{i'}$.
4. If $i' - 1$ is a progress point for Z then $\xi_{i'} \prec_{\mathcal{O}_{i'}} \xi_{i'-1}$, and otherwise $\xi_{i'} \preceq_{\mathcal{O}_{i'}} \xi_{i'-1}$.
5. If $\xi_{j'} \prec_{\mathcal{O}_{j'}} \xi_{i'-1}$ then $\xi_{j'+1} \prec_{\mathcal{O}_{j'+1}} \xi_{i'-1}$.
6. If $\xi_{j'} \preceq_{\mathcal{O}_{j'}} \xi_{i'-1}$ then $\xi_{j'+1} \preceq_{\mathcal{O}_{j'+1}} \xi_{i'-1}$.

Items (1–3) are immediate from the induction hypothesis.

For the proof of item (4), first suppose that $i' - 1$ is a progress point for Z ; then $\xi_{i'} \prec_{\mathcal{O}_{i'}} \xi_{i'-1}$ by Claim 11 applied to the indices $i' - 1, i'$, so that we are done. On the other hand, if $i' - 1$ is not a progress point for Z , we reason as follows. First note that $\varphi'_{i'-1}$ must be of the form $[a]\psi$ or $\langle a \rangle \psi$ for some formula ψ such that either $\varphi'_{i'} = \psi$ or $\varphi'_{i'} \prec_{\mathcal{O}_{i'}} \psi$. If $\psi = \varphi'_{i'}$ then (with $\bigcirc \in \{[a], \langle a \rangle\}$):

$$\xi_{i'} = \mathfrak{o}_{\varphi'_{i'}}(Z) = \mathfrak{o}_{\psi}(Z) = \mathfrak{o}_{\bigcirc\psi}(Z) = \mathfrak{o}_{\varphi'_{i'-1}}(Z) = \xi_{i'-1}.$$

Finally, if $\varphi'_{i'} \prec_{\mathcal{O}_{i'}} \psi$, then there is some fixpoint variable y such that $\mathfrak{o}_{\varphi'_{i'}}(y) \prec_{\mathcal{O}_{i'}} \mathfrak{o}_{\psi}(y)$ and $\mathfrak{o}_{\varphi'_{i'}}(z) = \mathfrak{o}_{\psi}(z)$ for all z ranking higher than y . But then $i' - 1$ is a progress point for y , hence by our choice of Z , the variable Z is higher ranking than y . Hence $\mathfrak{o}_{\varphi'_{i'}}(Z) = \mathfrak{o}_{\psi}(Z)$, and by the same reasoning as above we can show that $\mathfrak{o}_{\psi}(Z) = \mathfrak{o}_{\varphi'_{i'-1}}(Z)$. Hence $\mathfrak{o}_{\varphi'_{i'}}(Z) = \mathfrak{o}_{\varphi'_{i'-1}}(Z)$ as required.

For the proof of item (5), suppose $\xi_{j'} \prec_{\mathcal{O}_{j'}} \xi_{i'-1}$. We recall that $\Theta_{j'+1} = \Theta_{i'-1}$, so $\xi_{i'-1}$ is an active variable of $\Theta_{j'+1}$, and $\Theta_{j'}$ is a non-trivial a -successor

of $\Theta_{j'+1}$. By our assumption on Z being the highest ranking fixpoint variable with any progress points, this means that $o_{\varphi'_{j'}} \prec_{\mathcal{O}_{j'}}^Z o_{\varphi'_{i'-1}}$, hence $o_{\varphi'_{j'+1}} \prec_{\mathcal{O}_{j'}}^y o_{\varphi'_{i'-1}}$ for some y higher ranking or equal to Z . But then $y = Z$, again by our assumption that Z is the highest ranking variable with any progress points. In other words we have $o_{\varphi'_{j'+1}}(Z) \prec_{\mathcal{O}_{j'+1}} o_{\varphi'_{i'-1}}(Z)$, i.e. $\xi_{j'+1} \prec_{\mathcal{O}_{j'+1}} \xi_{i'-1}$ as required.

For the proof of item (6) suppose $\xi_{j'} = \xi_{i'-1}$. Then $\xi_{j'}$ is an active variable of $\Theta_{i'-1} = \Theta_{j'+1}$. Write $\varphi'_{j'} = \langle a \rangle \psi'$. If all variables in ψ' are active in $\Theta_{j'+1}$ then $\varphi'_{j'+1} \preceq_{\mathcal{O}_{j'+1}} \psi'$, and since Z was the highest ranking variable with any progress point it follows that

$$\begin{aligned} \xi_{j'+1} &= o_{\varphi'_{j'+1}}(Z) \\ &\preceq_{\mathcal{O}_{j'+1}} o_{\psi'}(Z) \\ &= \xi_{j'} \\ &= \xi_{i'-1} \end{aligned}$$

as required. On the other hand if some variable $o_{\psi'}(y)$ of ψ' is not active in $\Theta_{j'+1} = \Theta_{i'-1}$ then either $\varphi'_{j'+1} \preceq_{\mathcal{O}_{j'+1}} \psi'[\tau]$ for some substitution τ that is the identity on all active variables of $\Theta_{j'+1}$, including $\xi_{i'-1}$, or $o_{\varphi'_{j'+1}} \prec_{\mathcal{O}_{j'+1}}^y o_{\varphi'_{i'-1}}$ for some fixpoint variable y . Since Z was the highest ranking variable with any progress point, in either case we get $o_{\varphi'_{j'+1}}(Z) \preceq_{\mathcal{O}_{j'+1}} o_{\psi'}(Z)$ and therefore $\xi_{j'+1} \preceq_{\mathcal{O}_{j'+1}} \xi_{i'-1}$ as before.

We now proceed to finish the proof of the claim using items (1–6). If Z has a progress point m with $i' - 1 \leq m < j'$ then either $m = i' - 1$ or $i' \leq m < j'$. So combining items (3) and (4):

7. If Z has a progress point m with $i' - 1 \leq m < j'$, then $\xi_{j'} \prec_{\mathcal{O}_{j'}} \xi_{i'-1}$, and otherwise $\xi_{j'} \preceq_{\mathcal{O}_{j'}} \xi_{i'-1}$.

Combining items (5–7), we immediately obtain:

8. If Z has a progress point m with $i' - 1 \leq m < j'$, then $\xi_{j'+1} \prec_{\mathcal{O}_{j'+1}} \xi_{i'-1}$, and otherwise $\xi_{j'+1} \preceq_{\mathcal{O}_{j'+1}} \xi_{i'-1}$.

The index j' cannot be a progress point for Z since π'_j is not an initial segment of $\pi_{j'+1}$. So if Z has a progress point m with $i' - 1 \leq m < j$, then either $i' - 1 \leq m < j'$ or $j' + 1 \leq m < j$. So combining items (8) and (2):

9. If Z has a progress point m with $i' - 1 \leq m < j$, then $\xi_j \prec_{\mathcal{O}_j} \xi_{i'-1}$, and otherwise $\xi_j \preceq_{\mathcal{O}_j} \xi_{i'-1}$.

Finally, if Z has a progress point m with $i \leq m < j$, then either $i \leq m < i' - 1$ or $i' - 1 \leq m < j$. So combining items (9) and (1), we have:

10. If Z has a progress point m with $i \leq m < j$, then $\xi_j \prec_{\mathcal{O}_j} \xi_i$, and otherwise $\xi_j \preceq_{\mathcal{O}_j} \xi_i$.

But item (10) is the conclusion we aimed to prove, so the proof is finished. \square

Claim 13. Let $N \subseteq \omega$ be a set of indices such that $\pi_i = \pi_j$ for all $i \in N$. Then N is finite.

Proof of Claim. Say that a play π' of the mosaic game is visited infinitely many times if the set $N = \{i < \omega \mid \pi_i = \pi'\}$ is infinite. We want to show that no π' is visited infinitely many times.

It is easy to see that the set of plays π' of the mosaic game that are visited infinitely many times is downwards directed with respect to the initial segment ordering. Assume for contradiction that this set is non-empty then it contains some π' that is visited infinitely many times, and for any π'' that is visited infinitely many times, π' is an initial segment of π'' . Hence for some $k < \omega$, π' is an initial segment of π_m for all $m \geq k$. Furthermore, by the pigeonhole principle there is a formula ψ such that $(\pi_m, \varphi_m) = (\pi', \psi)$ for infinitely many $m < \omega$. Consequently there must be infinitely many $k_0, k_1, k_2, \dots < \omega$ for which $(\pi_{k_i}, \varphi_{k_i}) = (\pi', \psi)$, π' is an initial segment of π_m for $k_i \leq m \leq k_{i+1}$, and Z has a progress point m with $k_i \leq m < k_{i+1}$. But then, it follows immediately from Claim 12 that for each k_i , we have $\xi_{k_{i+1}} \prec_{\mathcal{O}} \xi_{k_i}$, where \mathcal{O} is the constraint of the last sequent on π' . Hence we find an infinite set of distinct ordinal variables in the same tile, which is impossible since the constraint of a sequent is always finite. \square

Claim 14. There is a (unique) infinite σ -guided play π_∞ and a sequence $S = (l_i, r_i)_{i < K}$ of pairs of indices (where $K \leq \omega$), such that the following conditions hold for each $i < K$:

- $l_i < r_i$, and $r_i < l_{i+1}$ (if $i + 1 < K$);
- $\pi_{l_i} = \pi_{r_i} \sqsubseteq \pi_\infty$;
- $\pi_{l_i} \sqsubseteq \pi_m$ for each m with $l_i \leq m \leq r_i$;
- $\pi_m \sqsubseteq \pi_{m+1}$, for all m with $r_i \leq m < l_{i+1}$ (if $i + 1 < K$).

Furthermore, if $K < \omega$ then

- $\pi_m \sqsubseteq \pi_{m+1}$ and $\pi_m \sqsubset \pi_\infty$, for all $m \geq r_{K-1}$.

We can view the play π_∞ as a limit of the partial plays $(\pi_i)_{i < \omega}$ in the mosaic game appearing in the infinite play $(\pi_i, \varphi_i)_{i < \omega}$ in the evaluation game, not in the sense that $\pi_i \sqsubseteq \pi_\infty$ for all $i < \omega$, but for *infinitely many* $i < \omega$. Intuitively, Claim 14 says that the infinite play $(\pi_i, \varphi_i)_{i < \omega}$ looks as displayed in Figure 4 below. The dotted arrow shows how the play $(\pi_i, \varphi_i)_{i < \omega}$ traverses the tree of σ -guided plays, and the shaded branch to the left shows the infinite σ -guided play π_∞ .

Proof of Claim. We begin by showing that there is an infinite σ -guided play π_∞ such that every finite initial segment of π_∞ is equal to π_i for some $i < \omega$. Consider the set Π of partial σ -guided plays π' such that every finite initial




Figure 4: The play π_∞ of Claim 14.

segment of π' is equal to π_i for some $i < \omega$. This set is clearly downwards closed under the initial segment order, so it forms a finitely branching tree in which π'' is considered as a child of π' iff it is a minimal partial play in Π of which π' is a proper initial segment. Thus, by König's lemma it suffices to prove that this tree is infinite. But clearly, π_i belongs to Π for every $i < \omega$. So if the set Π were finite, by the pigeonhole principle it would have to contain some member π' which is equal to π_i for infinitely many $i < \omega$, and this contradicts claim 13.

We now show that the play π_∞ must satisfy the constraints listed in the Claim. We define, by induction on the length of an initial segment π' of π_∞ , a set $S(\pi')$ of pairs of indices as follows. Suppose that π' is an initial segment of π_∞ and that $S(\pi'')$ has been defined for all proper initial segments π'' of π' . Let S' be the union of all $S(\pi'')$ for π'' a proper initial segment of π' . We define $S(\pi')$ by a case distinction. If there are no two distinct indices $i < j < \omega$ for which $\pi' = \pi_i = \pi_j$, then we set $S(\pi') = S'$. Otherwise, let l be the *smallest* index such that $\pi' = \pi_l$, and let r be the *greatest* index for which $\pi' = \pi_r$. The latter must exist by Claim 13. Then $l < r$. We set $S(\pi) = S' \cup \{(l, r)\}$.

With this construction in place, we set S to be the union of all $S(\pi')$ where

π' is an initial segment of π_∞ . It is not hard to see that the constraints of the Claim are met with this definition of S . \square

We now show that the σ -guided play π_∞ is lost by Refuter, which gives a contradiction. As a consequence of Claim 12, for every $i < K$ we have $\xi_{r_i} \prec_{\mathcal{O}_{r_i}} \xi_{l_i}$ if Z has a progress point m with $l_i \leq m < r_i$, and $\xi_{r_i} = \xi_{l_i}$ otherwise. Furthermore, by Claim 11, it follows that for each i with $i + 1 < K$ we have $\xi_{l_{i+1}} \prec_{\mathcal{O}_{l_{i+1}}} \xi_{r_i}$ if Z has a progress point m with $r_i \leq m < l_{i+1}$, and $\xi_{l_{i+1}} = \xi_{r_i}$ otherwise. Also by Claim 11, if K is finite then for each m with $m \geq r_{K-1}$, if m is a progress point for Z then $\xi_{m+1} \prec_{\mathcal{O}_{m+1}} \xi_m$, and $\xi_{m+1} = \xi_m$ otherwise.

Putting these observations together, since Z has infinitely many progress points by assumption, there must be a strictly increasing sequence of indices i_0, i_1, i_2, \dots such that $\pi_{i_j} \sqsubseteq \pi_\infty$ and $\xi_{i_{j+1}} \prec_{\mathcal{O}_{i_{j+1}}} \xi_{i_j}$ for each $j < \omega$. Hence π_∞ is lost by Refuter, contradicting the assumption that σ was a winning strategy for Refuter. We can now conclude that Falsifier wins every σ' -guided infinite play. \square

Theorem 5.28. If ρ is valid then it has a slim proof.

Proof. If ρ is valid, then by Proposition 5.27, Refuter does not have a winning strategy in the mosaic game for ρ , so that by Proposition 5.25, it is Prover who has a winning strategy. It then follows by Proposition 5.26 that ρ is provable. \square

6 From slim proofs to cyclic proofs

With completeness for slim proofs in place, we proceed to show that the existence of a slim proof for root formula ρ implies that there also exists a (finite) cyclic proof of that formula. In other words, we want a procedure to transform a given slim proof into a cyclic proof. It will be convenient to split this transformation into two steps, first introducing an infinitary analogue of cyclic proofs that will be called *infinitely reset proofs*, and then showing how such an infinite proof can be pruned into a finite proof tree which, with suitable back edges added, forms a cyclic proof.

6.1 Infinitely reset proofs

We first define the notion of an infinitely reset proof.

Definition 6.1. An infinite proof tree Π for root formula ρ is said to be *infinitely reset* if, for every infinite branch β on Π , there is some ordinal variable κ that appears non-active in every sequent in some final segment of β , and is reset infinitely many times on β .

We want to show that a slim proof can be transformed into an infinitely reset proof containing only finitely many sequents. In order to do that, we have to tame the growth of constraints in the proof tree using the constraint weakening rule, and inserting instances of the reset rule when possible. In order for this to

work we have to keep track of which ordinal variables we want to keep in order to possibly be reset later, and which ones we want to simply throw away. The following definition captures this distinction:

Definition 6.2. Let $\mathcal{O} : \Gamma$ be a sequent and κ an ordinal variable in \mathcal{O} . Then κ is said to be *redundant* if it is non-active and all of its descendants are non-active.

We are now ready to state and prove the result.

Proposition 6.3. If ρ has a slim proof then it has an infinitely reset proof in which only finitely many distinct sequents appear.

Proof. Let Π be a slim proof of ρ . We construct an infinite proof tree Π' together with a map $f : \Pi' \rightarrow \Pi$ as the union of finite approximants Π'_i and f_i , for $i < \omega$, defined as follows. First, we set Π'_0 to consist of a single vertex, labelled with the end sequent of Π , and we let f_0 map this single vertex to the root of Π . Now suppose the approximants Π'_i and f_i have been defined, so that for every vertex v of Π'_i labelled $\mathcal{O}_v : \Gamma_v$, the label of $f_i(v)$ is of the form $\mathcal{O}_{f_i(v)} : \Gamma_{f_i(v)}$ where $\Gamma_v = \Gamma_{f_i(v)}$ and $\mathcal{O}_v = \mathcal{O}_{f_i(v)} \setminus V$ for some set V of non-active variables of $\mathcal{O}_{f_i(v)}$. We construct Π'_{i+1} and f_{i+1} by doing the following for each non-axiom leaf v of Π'_i :

- If the constraint of v contains redundant ordinal variables, make v the conclusion to an application of left weakening that removes all these redundant variables and let f_{i+1} map the new premiss to $f_i(v)$.
- If the constraint of v contains no redundant ordinal variables but there is an instance of the reset rule with conclusion matching v , make v the conclusion to such an instance of reset and let f_{i+1} map the new premiss to $f_i(v)$.
- If neither of the two previous cases apply, then make v the conclusion of the same rule instance as $f_i(v)$ and let f_{i+1} map each new premiss w to the corresponding premiss of $f_i(v)$.

Note that the last case of this construction can be carried out thanks to the assumption that Π was a slim proof. In particular, we are relying here on the restriction on the cut rule, $\mu(\kappa)$ or \exists , that all ordinal variables occurring free in the cut formula or minor formula of the rule application are active variables in the conclusion. For example, suppose $f_i(v)$ is the conclusion to an instance of the \exists -rule of the form:

$$\frac{\mathcal{O}(\lambda < \kappa) : \Gamma, \varphi[\lambda/\xi]}{\mathcal{O}(\lambda < \kappa) : \Gamma, \exists \xi < \kappa. \varphi}$$

By assumption the constraint of v is of the form $\mathcal{O} \setminus V$ for some set of non-active variables V of \mathcal{O} . The constraint on applications of the \exists -rule in slim proofs ensures that λ is active in Γ_v and thus cannot belong to V . Hence λ belongs to $\mathcal{O} \setminus V$ and we can safely make v the conclusion to the corresponding rule instance:

$$\frac{(\mathcal{O} \setminus V)(\lambda < \kappa) : \Gamma, \varphi[\lambda/\xi]}{(\mathcal{O} \setminus V)(\lambda < \kappa) : \Gamma, \exists \xi < \kappa. \varphi}$$

We claim that the limit Π' obtained from this construction is an infinitely reset proof.

Claim 15. Let $\mathcal{O} : \Gamma$ be a sequent such that \mathcal{O} contains at most n active variables, and suppose \mathcal{O} contains a chain $\kappa_0, \dots, \kappa_{2n+1}$ of $2n+2$ different ordinal variables such that $\kappa_{2n+1} <_{\mathcal{O}} \dots <_{\mathcal{O}} \kappa_0$. Then there is an instance of the reset rule with conclusion $\mathcal{O} : \Gamma$ in which one of the variables $\kappa_0, \dots, \kappa_{2n+1}$ is reset.

Proof of claim 15. Since the length of the chain is $2n + 2$ it contains a chain of $n + 2$ different non-active variables. Clearly, as these variables are all part of a chain, at least $n + 1$ of these non-active variables have a non-empty set of children. Finally, as the sets of children of two different variables are disjoint, at most n of these variables can have an active child. So the chain must contain at least one non-active variable κ with a non-empty set of children, all of which are non-active. This variable can be reset. \square

Claim 16. Let $\beta = \Gamma_0\Gamma_1\Gamma_2\dots$ be an infinite branch of Π' with corresponding sequence of constraints $\mathcal{O}_0\mathcal{O}_1\mathcal{O}_2\dots$ and let $f[\beta]$ be the corresponding branch of Π . If $f[\beta]$ has an infinite descending chain $\kappa_0, \kappa_1, \kappa_2\dots$ of ordinal variables, then there is an index $k < \omega$ and variable λ in \mathcal{O}_k such that:

1. λ belongs to \mathcal{O}_i for all $i \geq k$, and
2. λ is reset infinitely many times on β .

Proof of claim 16. We write $f[\beta] = \Gamma_0^f\Gamma_1^f\Gamma_2^f\dots$ and let $\mathcal{O}_0^f\mathcal{O}_1^f\mathcal{O}_2^f\dots$ denote the corresponding sequence of constraints. We assume that κ_0 occurs in \mathcal{O}_0^f since otherwise we can just drop the prefix of the sequence before the first occurrence of κ_0 and re-index. For each $i < \omega$ we denote by $g(i)$ the highest index such that $\kappa_{g(i)}$ appears in \mathcal{O}_i^f . Note that $\kappa_{g(i)}$ must be active in Γ_i^f , since there must be some descendant Γ_j^f with $j > i$ that is the conclusion to either a ν -rule or \forall -rule in which $\kappa_{g(i)}$ occurs in the principal formula. So $\kappa_{g(i)}$ is active in Γ_j^f and hence, because Π is a slim proof, in Γ_i^f . Note that g is monotone in the sense that $i \leq j$ implies $g(i) \leq g(j)$, since there are no applications of left weakening in Π .

For each $i < \omega$ we denote by S_i the set of variables λ in \mathcal{O}_i that fulfil the following two properties.

Stability: No proper ancestor of λ is reset in any constraint \mathcal{O}_j with $j \geq i$.

Ancestry: $\kappa_{g(i)} <_{\mathcal{O}_i} \lambda$.

We note that S_i is non-empty for all sufficiently large i , since eventually $\kappa_{g(i)}$ must have at least one proper ancestor, and the $<_{\mathcal{O}_i}$ -largest ancestor of $\kappa_{g(i)}$ must satisfy Stability since it has no proper ancestors. Note that if $\lambda \in S_i$ then the Ancestry condition entails that $\kappa_{g(j)} <_{\mathcal{O}_j} \lambda$ for all $j \geq i$ such that λ belongs to \mathcal{O}_j . Together with the Stability condition this entails that $\lambda \in S_i$ belongs

to \mathcal{O}_j for all $j \geq i$, since an ancestor of $\kappa_{g(j)}$ is never redundant and hence can only be removed by resetting its parent.

Now pick an index i_0 for which S_{i_0} is non-empty, and pick $\lambda_0 \in S_{i_0}$. Since, by Stability, λ_0 belongs to \mathcal{O}_j for all $j \geq i_0$, we are done if λ_0 is reset infinitely many times. Otherwise, there is some $k > i_0$ such that λ_0 is never reset in any constraint \mathcal{O}_j for $j \geq k$. Let i_1 be the first index after k such that λ_0 has a child in \mathcal{O}_{i_1} which is an ancestor of $\kappa_{g(i_1)}$; such an index obviously exists since $\kappa_0 \kappa_1 \kappa_2 \dots$ is an infinite descending chain. We let λ_1 denote this child of λ_0 in \mathcal{O}_{i_1} . Since λ_0 is never reset after i_1 , λ_1 satisfies the Stability and Ancestor conditions with respect to the index i_1 . Now repeat the same argument; eventually, we must find a variable that belongs to \mathcal{O}_j for all sufficiently large j and is reset infinitely many times. For, suppose not; then, let m be the maximum number of active variables that can appear in any constraint in Π' . By repeating the previous argument sufficiently many times we eventually find a chain of $2m + 2$ variables $\lambda_{2m+1}, \dots, \lambda_0$ such that $\lambda_{2m+1} <_{\mathcal{O}_{i_{2m+1}}} \dots <_{\mathcal{O}_{i_{2m+1}}} \lambda_0$, none of which can be reset. This contradicts Claim 15, so the proof is done. \square

It follows immediately from Claim 16 that the proof tree Π' has the desired property, that on every infinite branch there is some variable that eventually appears in every constraint and is reset infinitely many times. Furthermore, whenever the number of non-active variables in the constraint exceeds a certain bound, then since the number of active variables is bounded either there are redundant variables (that are immediately removed by left weakening) or there is a sufficiently long chain of variables so that the reset rule can be applied (by Claim 15). It follows that the proof tree Π' has only finitely many sequents up to renaming of ordinal variables. So by suitably renaming ordinal variables, we can produce an infinitely reset proof Π'' in which only finitely many sequents appear. \square

6.2 Completeness of cyclic proofs

We are now ready for the final step of the completeness proof, showing that any infinitely reset proof containing finitely many sequents can be folded to a cyclic proof. We recall that a cyclic proof tree is a finite proof tree together with a back edge for each non-axiom leaf l , the target of which is called the companion of the leaf, such that the label of each non-axiom leaf l is identical with the label of its companion. Furthermore, such a cyclic proof tree is a valid cyclic proof if for every non-axiom leaf l there is some ordinal variable κ_l such that κ_l belongs to the constraint of every vertex on the path from the companion of l to l , and is reset at least once on this path.

Proposition 6.4. Let Π be an infinite reset proof for the formula ρ , and assume that only finitely many distinct sequents appear in Π . Then ρ has a cyclic proof.

Proof. Let Π and ρ be as in the formulation of the Proposition.

Consider an arbitrary infinite branch β of Π . Since Π is an infinite reset proof, we may associate an ordinal variable κ_β with β which appears non-active

in every sequent in some final segment of β and is reset infinitely often in β . But then by the pigeonhole principle it follows from the additional assumption on Π that β contains two nodes, c_β and l_β , labelled with the same sequent, and such that l_β is a (proper) descendant of c_β , the variable κ_β is a non-active element of the constraint of every sequent on the path from c_β to l_β , and κ_β is reset at least once on this path.

Now we prune the underlying tree of Π as follows. Let L be the set of *first repeats* of Π , that is, those nodes of the form l_β for some infinite branch β that do not have a proper ancestor of the form l_α for some other infinite branch α . Define Π_L as the proof tree we obtain by pruning all descendants of nodes in L ; in other words, the leaves of Π_L are either axiomatic or elements of L .

The key observation is that Π_L is finite. To see this, suppose otherwise, then by König's Lemma Π_L contains an infinite branch β . By construction β must be an infinite branch of Π as well, so that we may consider the node l_β . Since l_β belongs to Π_L , it cannot have a proper ancestor in L . It follows that l_β itself must be a first repeat and thus belongs to L . But then Π_L cannot contain any descendant of l_β , which contradicts the assumption that β is an infinite branch in Π_L .

We claim that in fact Π_L is a cyclic proof for ρ . For this purpose we need to define a companion map on Π_L , so consider an arbitrary non-axiomatic leaf l of Π_L . By construction we may fix some infinite branch β of Π such that $l = l_\beta$. Now define the companion of l to be the node c_β . It is then straightforward to verify that with this companion map, Π_L is indeed a cyclic proof for ρ . \square

Putting our results together, we now obtain

Proposition 6.5. If a plain formula ρ is valid, then it has a cyclic proof.

Proof. Let ρ be valid, then it has a slim proof by Theorem 5.28. As a direct consequence of Proposition 6.3 and Proposition 6.4 we may conclude that ρ has a cyclic proof as well. \square

Finally, Theorem 3.10 now follows immediately from Proposition 4.3 and Proposition 6.5.

7 Conclusion

We have presented a sound and complete cyclic proof system for the two-way μ -calculus. The defining feature of the system is the use of ordinal variables for detecting successful branches and ordinal quantifiers for handling the effect of two-way traces. The sequent calculus builds on two strong and successful formalisms for modal μ -calculus: explicit ordinal approximants [17, 4, 16] and name signatures [11, 18].

The proof system we introduced is suitable for effective proof search, as our completeness proof shows that cuts can be restricted to a small set of formulas. However, our completeness proof relies on saturation techniques that make it difficult to calculate a bound on the size of proofs for valid formulas. We leave

open whether the proof system can be used to obtain an optimal decision procedure for the two-way μ -calculus, matching the exponential time bound obtained by Vardi [21].

One can view the move from modal μ -calculus to two-way μ -calculus as a first step towards finitary proof systems for guarded fragments of first-order fixed point logic [8, 20, 3]. The cyclic calculus of the present paper is subsumed in the cyclic proof system for the first-order μ -calculus of [1]. Looking forward, we thus hope that insights from the present work can yield complete systems for larger (guarded) fragments of the first-order μ -calculus.

There is an apparent connection between two-way μ -calculus and model-checking infinite-state processes. It is well known that the two-way μ -calculus lacks the finite model property, although any satisfiable formula has a model that can be represented as a regular tree [21]. But even a regular tree model is in general infinite-state from the perspective of the two-way μ -calculus, as the appropriate notion of bisimulation will distinguish infinitely many states, all satisfying different formulas in the language. A calculus for model-checking context-free infinite-state processes was proposed by Schöpp and Simpson [15] which is similar to our own in the way it builds on [4]. As an example of an infinite-state process, Schöpp and Simpson consider a context-free system that represents a transition system with two actions converse to each other.

To date, there are no known complete Hilbert-style axiomatisation of the two-way μ -calculus. Indeed, the only sound and complete proof systems are the ω -branching well-founded proofs of [2] and the cyclic proofs of this article. An obvious question is whether either system sheds light on a Hilbert axiomatisation. One can also ask how the two proof systems are related. Aside from the difference in language (ordinals are external to the logical language in [2]), the crucial use of cut in the present work is a non-trivial hurdle in any comparison. As noted, we leave open whether our calculi are complete in the absence of cut.

References

- [1] Bahareh Afshari, Sebastian Enqvist, and Graham E. Leigh. Cyclic proofs for the first-order μ -calculus. *Logic Journal of the IGPL*, 2022.
- [2] Bahareh Afshari, Gerhard Jäger, and Graham E. Leigh. An infinitary treatment of full μ -calculus. In *Logic, Language, Information, and Computation: 26th International Workshop, WoLLIC 2019, Utrecht, The Netherlands, July 2-5, 2019, Proceedings*, pages 17–34. Springer, 2019.
- [3] Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded negation. *Journal of the ACM*, 62(3):1–26, 2015.
- [4] Mads Dam and Dilian Gurov. μ -calculus with explicit points and approximations. *Journal of Logic and Computation*, 12(2):255–269, 2002.

- [5] Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.
- [6] Sebastian Enqvist. Flat modal fixpoint logics with the converse modality. *Journal of Logic and Computation*, 28(6):1065–1097, 2018.
- [7] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata Logics, and Infinite Games: A Guide to Current Research*. Springer-Verlag New York, Inc., New York, NY, USA, 2002.
- [8] Erich Grädel and Igor Walukiewicz. Guarded fixed point logic. In *Proceedings. 14th Symposium on Logic in Computer Science (LICS)*, pages 45–54, 1999.
- [9] David Janin and Igor Walukiewicz. Automata for the modal μ -calculus and related results. In *Proceedings of the Twentieth International Symposium on Mathematical Foundations of Computer Science, MFCS'95*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.
- [10] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic. In *Proceedings of the Seventh International Conference on Concurrency Theory, CONCUR '96*, volume 1119 of *LNCS*, pages 263–277, 1996.
- [11] Natthapong Jungteerapanich. A tableau system for the modal μ -calculus. In *Automated Reasoning with Analytic Tableaux and Related Methods: 18th International Conference, TABLEAUX 2009, Oslo, Norway, July 6-10, 2009. Proceedings 18*, pages 220–234. Springer, 2009.
- [12] Dexter Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [13] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- [14] Jan Rooduijn and Yde Venema. Focus-style proofs for the two-way alternation-free μ -calculus. In Helle Hvid Hansen, Andre Scedrov, and Ruy J.G.B. de Queiroz, editors, *Logic, Language, Information, and Computation*, pages 318–335, Cham, 2023. Springer Nature Switzerland.
- [15] Ulrich Schöpp and Alex Simpson. Verifying temporal properties using explicit approximants: Completeness for context-free processes. In *Foundations of Software Science and Computation Structures: 5th International Conference, FOSSACS 2002 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8–12, 2002 Proceedings 5*, pages 372–386. Springer, 2002.

- [16] Christoph Sprenger and Mads Dam. On global induction mechanisms in a μ -calculus with explicit approximations. *RAIRO Theor. Informatics Appl.*, 37(4):365–391, 2003.
- [17] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: Circular and tree-shaped proofs in the μ calculus. In *Foundations of Software Science and Computation Structures: 6th International Conference, FOSSACS 2003 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2003 Warsaw, Poland, April 7–11, 2003 Proceedings 6*, pages 425–440. Springer, 2003.
- [18] Colin Stirling. A tableau proof system with names for modal mu-calculus. In *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, pages 306–318. 2014.
- [19] Robert S. Streett. Propositional dynamic logic of looping and converse is elementarily decidable. *Information and Control*, 54(1):121–141, 1982.
- [20] Balder ten Cate and Luc Segoufin. Unary negation. In *28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)*, 2011.
- [21] Moshe Y. Vardi. Reasoning about the past with two-way automata. In *Automata, Languages and Programming: 25th International Colloquium, ICALP'98 Aalborg, Denmark, July 13–17, 1998 Proceedings 25*, pages 628–641. Springer, 1998.
- [22] Thomas Wilke. Alternating tree automata, parity games, and modal μ -calculus. *Bulletin of the Belgian Mathematical Society*, 8:359–391, 2001.