

Demystifying μ

Bahareh Afshari*

Department of Philosophy, Linguistics and Theory of Science

University of Gothenburg

Graham E. Leigh*

Department of Philosophy, Linguistics and Theory of Science

University of Gothenburg

Guillermo Menéndez Turata

Institute for Logic, Language and Computation

University of Amsterdam

Abstract. We explore the theory of illfounded and cyclic proofs for the propositional modal μ -calculus. A fine analysis of provability for classical and intuitionistic modal logic provides a novel bridge between finitary, cyclic and illfounded conceptions of proof and re-enforces the importance of two normal form theorems for the logic: guardedness and disjunctiveness.

Keywords: Modal mu-Calculus, Illfounded Proofs, Cyclic Proofs, Cut Elimination, Constructive Modal Logic

Address for correspondence: B. Afshari, University of Gothenburg, Box 200, 405 30 Göteborg, Sweden.

ACM CCS • Theory of computation \Rightarrow Modal and temporal logics • Theory of computation \Rightarrow Proof theory

*This work was supported by the Knut and Alice Wallenberg Foundation [2020.0199] and Dutch Research Council [OCENW.M20.048]. The authors are deeply grateful for the substantial feedback and highly constructive comments received from the referees during the reviewing process of the article; special thanks also go to Sebastian Enqvist for feedback on an earlier version.

1. Introduction

Formal proofs are an important tool in the mathematical study of computational systems, such as certifying the correctness of an abstract model or verifying that an implementation adheres to a specification. The certification and verification role of proofs boils down to questions of *proof existence* and *proof synthesis*: Does a given formula (sequent, judgement, etc.) admit a proof? Can a proof be generated for each provable formula (sequent, etc.)? Both questions rest heavily on a third, often unspoken: What constitutes a proof?

A *proof* is normally understood as a finite tree with vertices labelled from a class of *deduction elements* (formulas, sequents, judgements, etc.) where each vertex together with its children match the conclusion and premises of one of a fixed collection of *inferences*, the decision of which should be complexity-theoretically simple, at worst polynomial-time decidable. This definition is computationally sufficient for many logics such as propositional logic, modal logics, even predicate logic. In each case, soundness and completeness theorems for the corresponding notion of proof confirm that proof existence agrees with semantic validity, and normal form theorems – notably admissibility of cut – give rise to proof synthesis.

Logics and theories incorporating inductive and co-inductive concepts strain the traditional notion of proof, however. Proof existence relies on more complex completeness theorems which coax the infinitary behaviour of the inductive and co-inductive constructions into finitary induction principles. Synthesis likewise suffers, becoming a task in discerning induction invariants from mere provability or elaborating the proof calculus to recover desirable normal forms. The latter is the typical approach of sequent calculi for modal logics where inductive properties – transitivity, factivity, well-foundedness, etc. – are internalised in the inferences rules.

Illfounded proofs offer to alleviate the proof-theorist’s burden not by complicating the notion of proofs but by *relaxing* it. Proofs are no longer constrained to finite trees; they can be infinite trees, or even graphs, at the cost of infinite branches fulfilling some pre-determined correctness condition. Illfounded proofs provide an alternative to the traditional notion of proof that treats logics and theories of inductive concepts as manifestations of a general infinitary framework rather than diverging extensions of finitary logic. With a change in perspective comes the need to revisit the encompassing theory. Indeed, the inductive construction of formal proof plays a non-trivial role in most – if not all – fundamental results of well-founded proof theory: proof transformations, cut elimination, computational interpretations, syntactic approaches to interpolation and so on.

This article seeks to further the development of the theory of illfounded proofs. We do so in the context of the *modal μ -calculus*, the extension of propositional modal logic by quantifiers for inductive and co-inductive properties. The logic encompasses the vast majority of modal and temporal logics used in specification and verification. As well as being highly expressive, it enjoys good meta-logical properties, namely, decidability, finite model property and uniform interpolation. It serves as a natural metalanguage for comprehending the dynamics and dependencies between induction and co-induction without obscuring its core fragments.

The first syntactic characterisation of the modal μ -calculus was, in retrospect, a system of illfounded proofs [19, 26]. Completeness of a *finitary* calculus materialised only as a consequence of an intricate analysis of the infinitary system [35, 36]. Kozen’s finitary axiomatisation [23] may provide

the ‘correct’ notion of proof according to the traditional view, yet it is the infinitary notion which has proved more relevant for, and amenable to, investigation. The median of these extremes are the *cyclic proofs*, illfounded proofs whose infinite expression is confined to periodic patterns, that is, unfoldings of finite graphs into trees. Such proofs display many of the advantages of both formalisms. Cyclic illfounded proofs are known to be complete for the modal μ -calculus and can be used to show decidability and the small model property [26]. With refinement, cyclic proof calculi have been developed for deterministic proof-search and representation [20, 32, 2, 14, 25] and interpolation [5, 8, 30]. The first purely proof-theoretic investigation of illfounded, cyclic and wellfounded notions of proof for the modal μ -calculus was undertaken by Studer [34].

Reflecting on the results and techniques mentioned above, we examine the connection between illfounded, cyclic and finitary proof systems for the modal μ -calculus in the context of completeness for each notion of proof. Two novel ingredients of the study are a focus on the *constructive* modal μ -calculus as a framework for proving completeness of the classical logic proof systems and the role of *admissibility*, notably cut admissibility, in tying the role of standard normal form theorems to completeness arguments.

By *constructive modal μ -calculus* we mean Wijesekera’s *constructive modal logic* [37], conceived by Fitch [17], extended by operators for least and greatest fixed point of all positive propositional functions. Constructive modal logic is an intuitionistic realisation of classical modal logic K with a natural proof-theoretic rendering via a sequent calculus for modal logic restricted to ‘intuitionistic’ sequents $\Gamma \Rightarrow \alpha$. Indeed, we take this ‘proof-theoretic’ presentation as the definition of the constructive counterpart of the illfounded and cyclic proof systems. The logic is strictly contained in what is commonly called *intuitionistic modal logic* [27, 31].¹ As our employment of constructive modal logic is as a tool for analysing provability in *classical* logic, we restrict attention to a fragment expressively complete over the classical modal μ -calculus, namely, (intuitionistic) sequents of implication-free formulas allowing negated atoms. We refer to this as the *constructive fragment* of the classical modal μ -calculus and leave open the extent to which the presented results, specifically completeness and normal form theorems, can be lifted to the ‘full’ constructive modal μ -calculus. Our main results concerning the constructive fragment specifically is the admissibility of cut over illfounded proofs and the extension of known normal form theorems from classical modal logic—guardedness and disjunctivity—to the constructive fragment. Admissibility of cut is proved via a co-inductive argument following established techniques [18, 12, 29, 11, 28, 1, 7]. The *disjunctive normal form theorem*, first established in [19], is a powerful tool in the study of classical modal fixed point logics. The result states that every formula is equivalent to a formula in which all uses of conjunction are constrained to specific combinations of the modal operators and atoms. Satisfiability and validity checking for disjunctive formulas (over classical logic) is linear-time, and completeness of both the illfounded and finitary proof systems is trivial for this fragment. The classical definition of disjunctive formulas is insufficient as a normal form over intuitionistic logic. A modest generalisation, however, yields a notion of disjunctive formula which is expressively adequate over the constructive fragment. As an immediate corollary of the expressive adequacy result is that the constructive modal μ -calculus satisfies the uniform (Lyndon) interpolation property.

The main advantage of constructive modal logic and, in particular, its robust sequent calculi lies

¹For a comparison of modal logics with an intuitionistic base see, e.g., [13].

in an unexpected place, namely the reduction of illfounded provability to cyclic provability, specifically in isolating well-structured finite representations of illfounded proofs. This way we establish a partial completeness theorem for the system of *constructive* cyclic proofs (sans implication) which, paired with the aforementioned admissibility and normal form theorems, bootstraps the completeness argument for the full, classical, logic. A relatively straightforward translation confirms that our classical system of cyclic proofs coincides with the finitary one, namely Kozen’s axiomatisation of the propositional modal μ -calculus.

The reader familiar with proofs of completeness for μ -calculi will notice a similarity in our account to other completeness theorems for μ -calculi [36, 16, 21, 15]. In broad strokes, our argument follows the Kozen–Walukiewicz strategy ([36]) of establishing completeness for an expressively adequate fragment of the logic (negated disjunctive formulas) and showing that each formula is provably equivalent (over Kozen’s axiomatisation) to a formula in the fragment. Of the two parts of the proof, the second is the more challenging and involves identifying a tight connection between formulas and carefully chosen disjunctive equivalents. The connection induces a consequence relation, referred to as *tableaux consequence* in [36], which refines provable implication in certain contexts via a reduction already isolated by Kozen [23]. Our completeness argument forgoes such a consequence relation by showing that the reduction of the full logic to the disjunctive fragment is *constructive* in the sense described above. Indeed, we suspect that Walukiewicz’ tableaux consequence, as well as the closely related *game consequence* relation of Enqvist et al [16], realise intensional forms of constructive implication. Likewise, it is reasonable to ask whether the completeness proofs for linear-time μ -calculus due to Kaivola [21] and Doumane [15], which follow strategies similar to Walukiewicz [36], can also be reconceptualised as completeness via constructive logics.

Outline of paper

Section 2 introduces the modal μ -calculus, its semantics and the three fragments of importance to this article: guarded formulas, disjunctive formulas and the Π -fragment which comprises all formulas lacking the least fixed point quantifier. The finitary and illfounded sequent calculus are defined in Section 3. The finitary system is a sequent calculus rendering of Kozen’s axiomatisation of the modal μ -calculus. The system of illfounded proofs is analogous to Niwinski and Walukiewicz’ refutation tableaux system from [26].

Section 4 introduces the constructive fragment of the illfounded calculus and establishes preliminary results, including the admissibility of cut and the expressive adequacy of the guarded fragment over intuitionistic logic. The cyclic calculus is introduced in Section 5, where it is related to the finitary and illfounded proof systems. It is shown that cyclic proofs for classical logic coincides with Kozen’s axiomatisation, and that the (cut-free) constructive fragment coincides with the illfounded proofs on implications $\alpha \rightarrow \beta$ where α is restricted to the Π -closure of the disjunctive formulas.

Section 6 treats the second normal form theorem, the expressive adequacy of the disjunctive formulas. We establish a strong form of the result in that every formula is provably equivalent to a disjunctive formula via a constructive illfounded proof. We remark on two corollaries of the disjunctive normal form theorem. The first is completeness of the constructive illfounded proofs for negative sequents (sequents $\alpha \Rightarrow \perp$), which will be utilised later. The second corollary is the observation that

the method of assigning disjunctive formulas can be relativised to a specific selection of atoms, giving a novel proof of uniform Lyndon interpolation for the constructive fragment.

The three calculi finally converge in Section 7 where the normal form theorems and partial completeness results for constructive logics combine to establish completeness of the classical illfounded, cyclic and finitary calculi.

2. Modal μ -Calculus: Syntax and Semantics

Fix a countable set Var of *variables*. Symbols x, y, \dots serve as meta-variables for variables. We assume a set $\text{Var}^\perp = \{\neg x \mid x \in \text{Var}\}$ of symbols, called *negated variables*, disjoint from Var . Elements of $\text{Lit} = \text{Var} \cup \text{Var}^\perp$ are *literals*. The *formulas* and *predicates* of the modal μ -calculus are constructed from literals via the grammars

Formulas: $\alpha ::= x \mid \neg x \mid \Box\alpha \mid \Diamond\alpha \mid \bigwedge\Gamma \mid \bigvee\Gamma \mid \mu\varphi \mid \nu\varphi$

Finite sets: $\Gamma ::= \emptyset \mid \Gamma \cup \{\alpha\}$

Predicates: $\varphi ::= \lambda x \alpha$ provided α is x -positive

We assume the usual notion of free and bound variable occurrences. A formula is V -positive (or x -positive if $V = \{x\}$) if $\neg x$ does not occur in the formula for every variable $x \in V$. Predicates and formulas are identified up to α -equivalence, understood in the usual sense. We refer to \bigwedge and \bigvee as *connectives*, \Box and \Diamond as *modal operators*, and μ and ν as *quantifiers*. Formulas $\bigvee\Gamma$ ($\bigwedge\Gamma$) are called *disjunctions* (*conjunctions*) and elements of Γ as *disjuncts* (*conjuncts*). A formula of the kind $\mu\varphi$ or $\nu\varphi$ is called *quantified*. The set of formulas is denoted L_μ . A set L of literals is *consistent* if $x \notin L$ or $\neg x \notin L$ for every $x \in \text{Var}$.

We adopt the following naming convention. Symbols α, β, \dots range over formulas, φ, ψ, χ over predicates and Γ, Δ, \dots over finite sets of formulas. The symbol \bigcirc is used exclusively as a meta-variable for the two connectives \bigvee and \bigwedge , the symbol \triangle as a meta-variable for the modal operators, and σ for the quantifiers. The following abbreviations for common formulas and constructions are utilised:

$$\begin{aligned} \alpha \wedge \beta &:= \bigwedge\{\alpha, \beta\} & \Box\Gamma &:= \{\Box\alpha \mid \alpha \in \Gamma\} & \mu x \alpha &:= \mu(\lambda x. \alpha) & \top &:= \bigwedge\emptyset \\ \alpha \vee \beta &:= \bigvee\{\alpha, \beta\} & \Diamond\Gamma &:= \{\Diamond\alpha \mid \alpha \in \Gamma\} & \nu x \alpha &:= \nu(\lambda x. \alpha) & \perp &:= \bigvee\emptyset \end{aligned}$$

For the *complexity* of a formula α , in symbols $|\alpha|$, we employ the following definition:

$$\begin{aligned} |x| &= |\neg x| = 0 & |\triangle\alpha| &= |\alpha| + 1 & |\bigcirc\Gamma| &= \max\{|\alpha| + 1 \mid \alpha \in \Gamma\} \\ |\lambda x. \alpha| &= |\alpha| & |\sigma\varphi| &= |\varphi| + 1 \end{aligned}$$

Negation is simulated via an involution on formulas and predicates:

$$\begin{aligned} x^\perp &= \neg x & (\bigcirc\Gamma)^\perp &= \bigcirc^\perp \Gamma^\perp & \Gamma^\perp &= \{\alpha^\perp \mid \alpha \in \Gamma\} \\ (\neg x)^\perp &= x & (\triangle\alpha)^\perp &= \triangle^\perp \alpha^\perp & (\sigma\varphi)^\perp &= \sigma^\perp \varphi^\perp & (\lambda x. \alpha)^\perp &= \lambda x \alpha[x \mapsto \neg x]^\perp \end{aligned}$$

where $\{\bigcirc, \bigcirc^\perp\} = \{\bigwedge, \bigvee\}$ and similarly for \triangle^\perp and σ^\perp , and $\alpha[x \mapsto \neg x]$ is the result of substituting $\neg x$ for all free occurrences of x in α . Observe that $\alpha[x \mapsto \neg x]^\perp$ is x -positive iff α is x -postive.

With the negation operation in place, the primitive operation of substitution given above can be lifted to a general form in a natural way. We call a function $\theta: V \rightarrow L_\mu$ where $V \subseteq \text{Var}$ a *substitution*. Given a formula or predicate η and a substitution θ , define $\eta[\theta]$ to be the result of simultaneously substituting θx for x in η modulo renaming of bound variables in η . The operation commutes with the connectives and quantifiers and is given by the following clauses for literals and predicates:

$$x[\theta] = \begin{cases} \theta x, & x \in \text{dom } \theta, \\ x, & \text{otherwise,} \end{cases} \quad \neg x[\theta] = (\theta x)^\perp,$$

$$(\lambda y. \alpha)[\theta] = \lambda y. \alpha[\theta], \text{ for } y \notin \text{dom } \theta \text{ and } y \text{ not free in } \text{Rng } \theta.$$

We employ notation $\varphi\beta$ as shorthand for $\alpha[x \mapsto \beta]$ where $\varphi = \lambda x \alpha$.

Lemma 2.1. For any predicate or formula η , $|\eta| \leq |\eta[\theta]| \leq |\eta| + \max\{|\theta(x)| \mid x \in \text{dom } \theta\}$.

2.1. Subformulas and threads

There are two natural choices for a notion of *subformula* in μ -calculi, depending on the treatment of quantified formulas. We say that β is a *literal subformula* of α if β is among the formulas constructed in the generation of α according to the grammar generating L_μ . In particular, β is a literal subformula of $\sigma x. \beta$. The second view takes its definition from the inference rules used to derive quantified (and other) formulas and specifies $\varphi(\sigma\varphi)$ as a subformula of $\sigma\varphi$. The complexity of a formula always bounds the complexity of its literal subformulas, but not of its derived subformulas.

In the context of proof systems, derived subformula plays a more important role than the literal one and, hence, is the primary notion in this paper and we drop the adjective ‘derived’ in most cases. The length of a sequence s is denoted $|s|$. For a sequence of formulas $\tau = (\alpha_i)_i$ define $\tau^\perp = (\alpha_i^\perp)_i$.

Definition 2.2. (Subformula)

The *formula graph* is the directed graph $\langle L_\mu, \triangleright \rangle$ where $\alpha \triangleright \beta$ iff (i) $\alpha = \bigcirc \Gamma$ and $\beta \in \Gamma$, (ii) $\alpha = \triangle \beta$, or (iii) $\alpha = \sigma\varphi$ and $\beta = \varphi\alpha$. If $\alpha \triangleright \beta$ then β is called an *immediate subformula* of α . The reflexive transitive closure of \triangleright is called the *(derived) subformula relation* and is denoted \triangleright^* . A *(derived) subformula* of α is any β such that $\alpha \triangleright^* \beta$. The set of subformulas of α is denoted $\text{Sub}(\alpha)$.

Witnesses of the subformula relation, namely sequences $(\alpha_i)_{i < k}$ such that $\alpha_i \triangleright \alpha_{i+1}$ for each i are called *threads*. A sequence following the reflexive closure of \triangleright is called a *weak thread*. Formulas may admit both infinite and infinitely many threads, such as in the formula $\alpha = \mu x (\Diamond x \vee \Box y)$ for which there is one infinite thread and infinitely many distinct threads witnessing $\alpha \triangleright^* y$. A thread that starts at α is called an α -thread. We write $\tau : \alpha \triangleright^* \beta$ to express that τ is a finite α -thread ending in β , that is, $|\tau| > 0$, $\tau(0) = \alpha$ and $\tau(|\tau| - 1) = \beta$. If $\beta = \alpha$ and $|\tau| > 1$ the thread is called *cyclic*. Writing $\tau : \alpha \triangleright_\eta^* \beta$ conveys, in addition, that η is a formula of lowest complexity on the thread τ , namely $\tau(i) = \eta$ for some $i < |\tau|$ and $|\eta| \leq |\tau(i)|$ for every $i \leq |\tau|$. Note that if τ a thread then so is $\tau^\perp = (\tau(i)^\perp)_{i < |\tau|}$.

The following proposition relates threads with substitution.

Proposition 2.3. Let θ be a substitution. If $\tau : \alpha[\theta] \triangleright^* \beta$ then either:

1. $\tau = (\tau'[\theta])\tau''$ for threads $\tau' : \alpha \triangleright^* x \in \text{dom } \theta$ and $\tau'' : \theta(x) \triangleright^* \beta$.
2. $\tau = \tau'[\theta]$ for some thread $\tau' : \alpha \triangleright^* \beta'$ where $\beta = \beta'[\theta]$.

The next proposition sums up the central property of threads.

Proposition 2.4. Let τ be an infinite thread. There is a unique formula η satisfying: (i) $\eta = \tau(i)$ for infinitely many i and (ii) $|\eta| \leq |\tau(j)|$ for all but finitely many j . Moreover, $\eta = \sigma\varphi$ for some predicate φ and quantifier σ .

Definition 2.5. An infinite thread is ν if the formula associated to the thread by Proposition 2.4 has the form $\nu\varphi$, and is said to be μ otherwise.

2.2. Three fragments of L_μ

We identify three important fragments of L_μ . The first fragment is often the focus of studies in modal μ -calculi due to its well-structured threads. This is the *guarded* fragment comprising the formulas for which every cyclic thread among subformulas contains a modal formula. The guarded fragment is especially relevant in the context of proof-search.

Definition 2.6. (Guarded fragment)

A formula α is *guarded* if every cyclic thread in the subformula graph of α passes through a modality: for every $\beta \in \text{Sub}(\alpha)$ and $\tau : \beta \triangleright^* \beta$, if $|\tau| > 1$ then $\tau(i) = \Delta\tau(i+1)$ for some $i < |\tau|$ and $\Delta \in \{\Box, \Diamond\}$.

The second fragment of interest consists of all formulas in which ν is the only quantifier. We call this the Π -fragment. In practise, however, greater stock will be put on the ‘closure’ of certain sets of formulas under the generating operations of the Π -fragment. Thus, the Π -fragment is defined as $\Pi(\emptyset)$ where, for any $F \subseteq L_\mu$, the set $\Pi(F)$ is given by

Definition 2.7. ($\Pi(F)$ -fragment)

Let $F \subseteq L_\mu$. The set $\Pi(F)$ is the smallest set of formulas satisfying:

1. $F \cup \text{Lit} \subseteq \Pi(F)$,
2. If $\Gamma \subseteq \Pi(F)$, then $\bigvee \Gamma, \bigwedge \Gamma \in \Pi(F)$,
3. If $\alpha \in \Pi(F)$, then $\Box\alpha, \Diamond\alpha \in \Pi(F)$,
4. If $\varphi x \in \Pi(F)$ then $\nu\varphi \in \Pi(F)$,
5. If $\alpha \in \Pi(F)$ is V -positive and $\theta : V \rightarrow \Pi(F)$, then $\alpha[\theta] \in \Pi(F)$.

A $\Pi(\emptyset)$ -formula is called a Π -formula. A Σ -formula is a formula α^\perp for $\alpha \in \Pi(\emptyset)$.

The closure operation $F \mapsto \Pi(F)$ induces a quantifier hierarchy for L_μ known as the *alternation depth* or *Niwiński hierarchy*, given by $\Pi_1 = \Pi(\emptyset)$ and $\Pi_{n+1} = \Pi(\Sigma_n)$ where $\Sigma_n = \Pi_n^\perp$. The hierarchy exhausts L_μ . We will mostly be concerned with $\Pi(F)$ -formulas for two specific choices of F however: the empty set and the set of formulas in the final fragment to be introduced.

Lemma 2.8. If $\mu\varphi \in \Pi(F)$, then $\varphi = \psi[\theta]$ for $\mu\psi \in F$ and a substitution $\theta: V \rightarrow \Pi(F)$.

Proof:

By induction on the generation of $\Pi(F)$. If $\mu\varphi \in F$ then we are done. Otherwise, $\mu\varphi = \alpha[\theta]$ for some substitution $\theta: V \rightarrow \Pi(F)$ and V -positive α . We may assume that $\alpha \notin \text{Var}$ as otherwise $\mu\varphi = \theta\alpha$ and the result follows from the induction hypothesis. Therefore, $\alpha = \mu\chi$ for some predicate χ , whence the induction hypothesis provides $\mu\psi \in F$ and $\theta': V' \rightarrow \Pi(F)$ such that $\mu\chi = \mu\psi[\theta']$. Without loss of generality we may assume that each variable in V' occurs free in ψ and, therefore, $\theta'x$ is V -positive (as α is). Define $\theta'' = \theta' \circ \theta: V \cup V' \rightarrow \Pi(F)$ by $\theta''x = (\theta'x)[\theta]$ for $x \in V'$ and $\theta''x = \theta x$ otherwise. Then $\mu\varphi = \mu\psi[\theta'']$. \square

The final fragment is the *disjunctive* fragment introduced by Janin and Walukiewicz [19] and employed heavily in Walukiewicz' completeness proof [36]. It is the class of formulas wherein use of conjunction is restricted to specific arrays of modal formulas and variables. One of the many results concerning disjunctive formulas established in [19] is that every formula is equivalent to a formula in the disjunctive fragment.

For our rendition of the disjunctive normal form theorem it is necessary to weaken the definition of disjunctive formula slightly so that the expressive adequacy of disjunctive formulas holds constructively rather than merely under classical logic. The modification also has the advantage of being more amenable to standard techniques from structural proof theory while maintaining the core properties that make the fragment indispensable. Our modification takes the form of a liberalisation of the modal operator at the heart of the disjunctive fragment.

Definition 2.9. (Expanded cover modality)

Given a set $\Gamma \cup \Delta$ of formulas, we introduce formulas $\nabla(\Gamma, \Delta)$ and, more generally, $\nabla_L(\Gamma, \Delta)$ for $L \subseteq \text{Lit}$ defined by

$$\nabla_L(\Gamma, \Delta) := \bigwedge (L \cup \Diamond\Gamma \cup \{\Box\bigvee\Delta\}) \quad \nabla(\Gamma, \Delta) := \nabla_\emptyset(\Gamma, \Delta) = \bigwedge (\Diamond\Gamma \cup \{\Box\bigvee\Delta\})$$

The *expanded cover modality* comprises the formulas $\nabla_L(\Gamma, \Delta)$ where $\Gamma \subseteq \Delta$ and L is consistent.

The ‘standard’ cover modality is the special case $\nabla\Gamma := \nabla(\Gamma, \Gamma)$. The two model operators \Box and \Diamond have simple representations via the expanded cover modality: $\Box\alpha$ is equivalent to $\nabla(\emptyset, \{\alpha\})$ and $\Diamond\alpha$ to $\nabla(\{\alpha\}, \{\alpha, \top\})$. Moreover, these equivalences hold over constructive modal logic (cf. Proposition 4.5 below). Indeed, the representation of $\Box\alpha$ already demonstrates a difference with the unary cover modality base on which $\Box\alpha$ is expressed as $\nabla\emptyset \vee \nabla\{\alpha\}$, involving an external disjunction. More generally, the expanded cover modality $\nabla(\Gamma, \Delta)$ is equivalent to the disjunction $\nabla\Gamma \vee \nabla\Gamma'$ where $\Gamma' = \Gamma \cup \{\bigvee(\Delta \setminus \Gamma)\}$, but this equivalence does not hold constructively in general.

The restriction of consistency on L does not reduce the expressivity of the cover modality as any formula $\nabla_L(\Gamma, \Delta)$ where L is *not* consistent is clearly unsatisfiable. Nor is it simply a matter

of convenience: Theorem 4.9 depends specifically on the restriction to show that for a certain class of formulas employing the dual of the cover modality, provability relative to a constructive calculus coincides with validity against the classical semantics. In particular, it is necessary that no instance of the excluded middle can be expressed as the dual of a cover modality.

Definition 2.10. (Disjunctive fragment)

A formula is *disjunctive* iff every subformula that is a conjunction is an instance of an expanded cover modality. The *disjunctive fragment* is the set \mathbb{D} of disjunctive formulas. A formula α *conjunctive* iff α^\perp is disjunctive.

Although \perp is disjunctive by definition, \top is not as every conjunction contains at least one formula. In the context of disjunctive formulas, therefore, we take in its place the disjunctive rendition of $\nu x \Box x$. It is with remarking that the disjunctive fragment is more constrained than merely utilising the operator ∇ in place of conjunction in the inductive definition of formulas. Specifically, variables occurring alongside the modal operators (the ‘ V ’ in $\nabla_V(\Gamma, \Delta)$) can not be bound in a disjunctive formula. The formula $\nabla_{\{x\}}(\emptyset, \emptyset)$, for instance, is disjunctive yet $\alpha = \nu x. \nabla_{\{x\}}(\emptyset, \emptyset)$ is not because $\bigwedge \{\alpha, \Box \top\} \in \text{Sub}(\alpha)$ can not be realised as an instance of the expanded cover modality. However, $\alpha \in \Pi(\mathbb{D})$, a fact that will be useful later.

Combining the restriction on conjunctions in disjunctive formulas and on quantifiers in the Π fragments is:

Proposition 2.11. Let $\tau : \bigwedge \Gamma \triangleright_{\mu\varphi}^* \bigwedge \Gamma$ be a cyclic thread. If $\bigwedge \Gamma \in \Pi(\mathbb{D})$ then $\bigwedge \Gamma$ is an instance of the expanded cover modality.

Proof:

Suppose $\Gamma \subseteq \Pi(\mathbb{D})$ and $\tau : \bigwedge \Gamma \triangleright_{\mu\varphi}^* \bigwedge \Gamma$. In particular, $\mu\varphi \in \Pi(\mathbb{D})$, whereby Proposition 2.8 implies that $\mu\varphi = \mu\psi[\theta]$ for $\mu\psi \in \mathbb{D}$ and a substitution $\theta : V \rightarrow \Pi(\mathbb{D})$. Without loss of generality we may assume that all variables in V occur free in ψ . In particular, $|\theta(x)| < |\mu\varphi|$ for each $x \in V$. As $|\tau(i)| \geq |\mu\varphi|$ for all $i < |\tau|$ it follows that $\Gamma = \Gamma'[\theta]$ for some $\bigwedge \Gamma' \in \text{Sub}(\mu\psi)$ (Proposition 2.3(2)), whence $\bigwedge \Gamma = \nabla_V(\Delta, \Delta')$ for some Δ, Δ' and V . \square

2.3. Fixed-point semantics

Formulas of the modal μ -calculus are evaluated against labelled trees, where an *L-labelled tree* (henceforth simply *tree*) is a pair $\langle T, \lambda \rangle$ where $T \subseteq \omega^{<\omega}$ is a non-empty and prefix-closed set of finite sequences of natural numbers, called *vertices*, and $\lambda : T \rightarrow L$ is function assigning each $v \in T$ a *label* $\lambda(v)$. The empty sequence is called the *root*. A proper prefix of a vertex $u \in T$ is called the *predecessor* of u ; conversely, a *successor* of u is a vertex that has u as a predecessor. Denoting the prefix relation on sequences by \leq and the proper prefix relation by $<$, we have that u is a successor of v iff $v < u$. Proper successors of minimal length are called *children*, namely the set of children of $u \in T$ is the set $\text{child}_T(u) := T \cap \{ui \mid i \in \omega\}$.

The label function associated to a tree T is written as λ_T . We abuse notation and identify labelled trees with the set of their vertices, writing $T = \langle T, \lambda_T \rangle$. Given $\langle T, \lambda \rangle$ and $s \in T$ define $\langle T \uparrow s, \lambda \uparrow s \rangle$ as

the subtree of T at s , given by $T \uparrow s = \{v \in T \mid sv \in T\}$ and $\lambda \uparrow s: v \mapsto \lambda(sv)$. The set of immediate subtrees of T is denoted $\langle T, \lambda \rangle_* = \{\langle T \uparrow s, \lambda \uparrow s \rangle \mid s \in T \text{ and } |s| = 1\}$.

A *path* in T is a finite or infinite sequence $P = v_0 v_1 \dots$ of vertices such that $v_{i+1} \in \text{child}_T(v_i)$ for all $i < |P|$. An infinite path is also called a *branch*. Given a branch B we often write $v \in B$ if $v \leq B(i)$ for some i .

Definition 2.12. Set $E = 2^{\text{Var}}$. A (tree) *model* is a E -labelled tree $\langle T, A \rangle$. The set of models is denoted Mod . Given a model $M = \langle T, A \rangle$ and a set of models $U \subseteq \text{Mod}$, we write $A[x \mapsto U]$ for the function $A': T \rightarrow E$ such that $A'(v) = A(v) \cup \{x\}$ if $\langle T, A \rangle \uparrow v \in U$ and $A'(v) = A(v) \setminus \{x\}$ otherwise.

2.3.1. Kripke semantics

In Kripke semantics for the modal μ -calculus each formula α is assigned a set $\llbracket \alpha \rrbracket \subseteq \text{Mod}$ of models, called the *denotation* of α , and each predicate φ associated a function $\llbracket \varphi \rrbracket: 2^{\text{Mod}} \rightarrow 2^{\text{Mod}}$, defined by mutual recursion:

$$\begin{aligned} \llbracket x \rrbracket &= \{\langle T, A \rangle \in \text{Mod} \mid x \in A(\varepsilon)\} \\ \llbracket \Box \alpha \rrbracket &= \{M \in \text{Mod} \mid M_* \subseteq \llbracket \alpha \rrbracket\} & \llbracket \Diamond a \rrbracket &= \{M \in \text{Mod} \mid M_* \cap \llbracket a \rrbracket \neq \emptyset\} \\ \llbracket \bigwedge \Gamma \rrbracket &= \bigcap \llbracket \Gamma \rrbracket & \llbracket \mu \varphi \rrbracket &= \bigcap \{U \subseteq \text{Mod} \mid \llbracket \varphi \rrbracket U \subseteq U\} \\ \llbracket \bigvee \Gamma \rrbracket &= \bigcup \llbracket \Gamma \rrbracket & \llbracket \nu \varphi \rrbracket &= \bigcup \{U \subseteq \text{Mod} \mid U \subseteq \llbracket \varphi \rrbracket U\} \end{aligned}$$

where

$$\llbracket \Gamma \rrbracket = \{\llbracket \gamma \rrbracket \mid \gamma \in \Gamma\} \quad \llbracket \varphi \rrbracket: U \mapsto \{\langle T, A \rangle \mid \langle T, A[x \mapsto U] \rangle \in \llbracket \varphi \rrbracket\}.$$

If $\llbracket \alpha \rrbracket = \text{Mod}$ we call α *valid*.

The proofs of the following are standard; the reader can consult, for instance, [10].

Proposition 2.13. Suppose $\langle T, A \rangle$ and $\langle T, A' \rangle$ are models and $V \subseteq \text{Var}$ is such that for all $v \in T$, $A(v) \cap V = A'(v) \cap V$. Then for every formula α whose free variables are among V , $\langle T, A \rangle \in \llbracket \alpha \rrbracket$ iff $\langle T, A' \rangle \in \llbracket \alpha \rrbracket$.

Proposition 2.14. The denotation of predicates are monotone functions: If φ is a predicate and $U \subseteq V \subseteq \text{Mod}$, then $\llbracket \varphi \rrbracket U \subseteq \llbracket \varphi \rrbracket V$.

Lemma 2.15. $\llbracket \alpha^\perp \rrbracket = \text{Mod} \setminus \llbracket \alpha \rrbracket$ for every formula α .

Proposition 2.16. $\llbracket \varphi \rrbracket \llbracket \sigma \varphi \rrbracket = \llbracket \sigma \varphi \rrbracket$.

2.3.2. Game semantics

Rather than inductively ascribing a set of models to each formula we can instead view the connectives, modalities and quantifiers as operations that ‘reduce’ the claim ‘ α is true in M ’ to analogous claims about immediate subformulas and submodels. In the case of the purely modal vocabulary, it is straightforward to read the Kripke semantics in this style. However, from this local perspective the statement ‘ $\mu\varphi$ is true at M ’ is re-expressed as the just as complex claim ‘ $\varphi(\mu\varphi)$ is true at M ’; the global constraint ensures that this reduction is permitted only finitely often.

This semantics has a natural formulation as a two-player game, one player’s aim being to ‘verify’ the claim ‘ α is true in M ’ against moves by their opponent. For our purposes, we need only the ‘certificate’ of this verification process, i.e., a winning strategy, which we call a *verification*:

Definition 2.17. (Verification)

A *pre-verification* is an $(L_\mu \times \text{Mod})$ -labelled tree $V = \langle T, \lambda \rangle$ satisfying, for every $v \in T$:

1. If $\lambda(v) = (x, M) \in \text{Var} \times \text{Mod}$ then $M \in \llbracket x \rrbracket$.
2. If $\lambda(v) = (\bigvee \Gamma, M)$ there exists $\gamma \in \Gamma$ and $v' \in \text{child}_V(v)$ with $\lambda(v') = (\gamma, M)$.
3. If $\lambda(v) = (\Diamond \alpha, M)$ there exists $M' \in M_*$ and $v' \in \text{child}_V(v)$ with $\lambda(v') = (\alpha, M')$.
4. If $\lambda(v) = (\bigwedge \Gamma, M)$ then for all $\gamma \in \Gamma$ there exists $v' \in \text{child}_V(v)$ with $\lambda(v') = (\gamma, M)$.
5. If $\lambda(v) = (\Box \alpha, M)$ then for all $M' \in M_*$ there exists $v' \in \text{child}_V(v)$ with $\lambda(v') = (\alpha, M')$.
6. If $\lambda(v) = (\sigma\varphi, M)$ there exists $v' \in \text{child}_V(v)$ with $\lambda(v') = (\varphi(\sigma\varphi), M)$.

A *verification* is a pre-verification V such that every branch $(\alpha_i, M_i)_i$ of V , the thread $(\alpha_i)_i$ is ν . By $\llbracket \alpha \rrbracket$ is denoted the set of models M for which there exists a verification V with root (α, M) .

Unlike Kripke semantics, each (pre-)verification only references subtrees of the tree modal labelling the root. But following that semantics, predicates can be interpreted as monotone functions over the set of models:

$$\llbracket \varphi \rrbracket : U \mapsto \{ \langle T, A \rangle \mid \langle T, A[x \mapsto U] \rangle \in \llbracket \varphi x \rrbracket \}.$$

It can be shown that the game semantics ascribes the appropriate fixed points to the two quantifiers:

Proposition 2.18. If $\llbracket \varphi \rrbracket U \subseteq U$ then $\llbracket \mu\varphi \rrbracket \subseteq U$ and if $U \subseteq \llbracket \nu\varphi \rrbracket$ then $U \subseteq \llbracket \nu\varphi \rrbracket$.

The equivalence of game and Kripke semantics, often referred to as the ‘fundamental semantic theorem’, was first established by Streett and Emerson [33]. Contrary to some sources, however, the theorem admits an easy proof on the basis of a few elementary lemmas relating the two semantics. Indeed, we have already all the necessary ingredients.

Theorem 2.19. (Fundamental Semantic Theorem)

For all α , $\llbracket \alpha \rrbracket = \llbracket \alpha \rrbracket$.

Proof:

By induction on $|\alpha|$ using Propositions 2.16 and 2.18. □

Corollary 2.20. $\langle \alpha^\perp \rangle = \text{Mod} \setminus \langle \alpha \rangle$.

3. Finitary and Illfounded Proofs

The basic object of deduction in the finitary and illfounded proof system is the *sequent*, expressions of the form $\Phi \Rightarrow \Psi$ where Φ and Ψ are finite sequences of formulas. In section 5 we introduce a third proof system which operates on a slightly broader notion of sequent in which ‘formula’ are constructed via a richer grammar and the sequent itself carries an additional, non-formula, component that interacts with some inferences of the calculus. A uniform definition of ‘proof system’ should, therefore, be broad enough to accommodate both choices of deduction elements. In the interests of space, however, we formulate the concept in terms of ‘ordinary’ sequents hoping that the general intension of the definition is clear.

The set of sequents is denoted Seq . We employ symbols Φ, Ψ , etc. as metavariables for finite sequences of formulas. In a sequent $\Phi \Rightarrow \Psi$, the terms *antecedent* and *consequent* refer to Φ and Ψ respectively, and the sequent is *intuitionistic* if $|\Psi| \leq 1$. In an intuitionistic sequent we refer to the formula in Ψ (if it is non-empty) as the *head* and to Φ as the *context*. In the context of illfounded proofs it is important to have a well-defined notion of formula *occurrence*. Once established, however, the ordering of formulas within sequents will be largely ignored and left implicit.

A *position* is a pair $p = (s, k) \in \{\text{L}, \text{R}\} \times \omega$. If $s = \text{L}$ we refer to p as a *left* position; otherwise, p is a *right* position. A *formula occurrence* in a sequent $\mathcal{S} = \Phi \Rightarrow \Psi$ is a position $p = (s, k)$ such that $k < |\Phi|$ if p is a left position, and $k < |\Psi|$ otherwise. The formula at position $p = (s, k)$ in \mathcal{S} is denoted $\mathcal{S}(p)$, where

$$(\Phi \Rightarrow \Psi)(p) = \begin{cases} \Phi(k), & \text{if } s = \text{L}, \\ \Psi(k), & \text{if } s = \text{R}. \end{cases}$$

An *inference* is a pair (\mathcal{S}, S) where \mathcal{S} is a sequent and S is a finite set of sequents, usually presented in the form:

$$\frac{S}{\mathcal{S}} \quad \text{or} \quad \frac{\mathcal{S}_0 \quad \cdots \quad \mathcal{S}_n}{\mathcal{S}} \quad \text{if } S = \{\mathcal{S}_0, \dots, \mathcal{S}_n\}.$$

The sequent \mathcal{S} is referred to as the *conclusion* of the inference and elements of S as *premises*. A set of inferences constitutes a *rule*. Typically, rules are specified by an inference scheme such as the rule (\vee/R) comprising the inference

$$\frac{\Phi \Rightarrow \Psi, \gamma}{\Phi \Rightarrow \Psi, \vee \Gamma}$$

for all Φ, Ψ, Γ and $\gamma \in \Gamma$. An inference is an *instance* of a rule iff it is an element of that rule. For our purposes each rule is further identified as being either a *logical* or *structural* rule, the only difference being how they may be utilised within derivations and proofs.

Definition 3.1. (Derivation)

A *derivation* over a set of rules \mathcal{R} is a tree d with labelings $\mathcal{S}_d: d \rightarrow \text{Seq}$ and $\mathcal{R}_d: d \rightarrow \mathcal{R}$ satisfying:

1. for each vertex $v \in d$, $(\mathcal{S}_d(v), \{\mathcal{S}_d(u) \mid u \in \text{child}(v)\}) \in \mathcal{R}_d(v)$.
2. for every branch $B \subseteq d$ there are infinitely many $i < \omega$ such that $\mathcal{R}_d(B(i))$ is a logical rule.

The *endsequent* of a derivation is the sequent labelling its root. A derivation is *constructive* if every sequent occurring in the derivation is intuitionistic.

The second condition in the definition is primarily of significance for the illfounded proof system introduced in section 3.3 where certain derivations with infinite paths qualify as proofs. As a consequence of the first condition, leaves of a derivation are labelled by rule instances with no premises.

Not every derivation constitutes a *proof*. In finitary proof systems, only finite derivations will be permitted as proofs, while in the system of illfounded proofs infinite derivations are recognised as proofs *if* their branches fulfil a syntactic correctness condition.

3.1. Basic fixed-point logic

The first system of derivations and proofs we introduce forms the core of the other calculi in this paper. This is the system $K\bar{\mu}$ of *basic fixed-point logic*. The structural and logical rules of $K\bar{\mu}$ are listed in fig. 1 and 2 respectively. Rules denoted $(\star L)$ are called *left* rules and $(\star R)$ are *right* rules. Other rules, namely (b) and (id), are both left and right rules. Leaves of a derivation labelled by the structural rule (b) are called *buds* and serve to allow conditional proofs, i.e., proofs from *assumptions*. Thus a finite derivation in the rules of $K\bar{\mu}$ will be called a $(K\bar{\mu})$ -*proof from assumptions* with a $(K\bar{\mu})$ -*proof* being a finite derivation *without* buds.

The structural rules (eL) and (eR) are jointly called the *exchange* rules; (wL) and (wR) are *weakening* rules, and (cL) and (cR) *contraction* rules. The rules (idL), (idR) and (id) are collectively the *identity* rules. By (e) (respectively, (w) and (c)) we denote the rule expressing closure under applications of the exchange (resp., weakening and exchange, contraction and exchange) rules, namely (e) denotes the collection of instances $(\mathcal{S}, \{\mathcal{S}'\})$ for which there is a (possibly trivial) proof from assumptions using (eL) and (eR) only whose conclusion is \mathcal{S} and (only) bud is \mathcal{S}' . When there is no cause for confusion we leave instances of the structural rules implicit.

In each of the rules of $K\bar{\mu}$ with the exception of (b) and (id), a single formula (occurrence) in the conclusion is designated the *principal* formula (occurrence). This is the formula occurrence $(L, 0)$ for a left rule and $(R, |\Psi|)$ for a right rule with Ψ as in the figures. In the case of (b) and (id), all formula occurrences in the conclusion are designated as principal. The ‘matching’ formula occurrence in the premise(s) of the rule are *minor* formula(occurrence)s. In the logical rules, the minor formula is the formula occurrence at the position of the principal formula; for the structural rules it is the positions inhabited by α in fig. 1. Other formula occurrences are *side* formulas. With the exception of the modal inferences $(\Box R)$ and $(\Diamond L)$, the sequent of side formulas occurs in both conclusion and each premise in the same order.

The calculus $K\bar{\mu}$ is of little significance beyond it being the common core of the other proof systems we introduce. We collect some observations on the rules of $K\bar{\mu}$ for later use.

$$\begin{array}{cccc}
\frac{\Phi, \alpha, \Lambda \Rightarrow \Psi}{\alpha, \Phi, \Lambda \Rightarrow \Psi} \text{eL} & \frac{\alpha, \alpha, \Phi \Rightarrow \Psi}{\alpha, \Phi \Rightarrow \Psi} \text{cL} & \frac{\Phi \Rightarrow \Psi}{\alpha, \Phi \Rightarrow \Psi} \text{wL} & \frac{}{\Phi \Rightarrow \Psi} \text{b} \\
\frac{\Phi \Rightarrow \Psi, \alpha, \Lambda}{\Phi \Rightarrow \Psi, \Lambda, \alpha} \text{eR} & \frac{\Phi \Rightarrow \Psi, \alpha, \alpha}{\Phi \Rightarrow \Psi, \alpha} \text{cR} & \frac{\Phi \Rightarrow \Psi}{\Phi \Rightarrow \Psi, \alpha} \text{wR} &
\end{array}$$

Figure 1. Structural rules of $K\bar{\mu}$. Rules denoted $\star L$ ($\star R$) are *left* (resp. *right*) rules.

$$\begin{array}{cccc}
(\gamma \in \Gamma) \frac{\gamma, \Phi \Rightarrow \Psi}{\bigwedge \Gamma, \Phi \Rightarrow \Psi} \wedge L & \frac{\{\gamma, \Phi \Rightarrow \Psi \mid \gamma \in \Gamma\}}{\bigvee \Gamma, \Phi \Rightarrow \Psi} \vee L & \frac{\varphi(\sigma\varphi), \Phi \Rightarrow \Psi}{\sigma\varphi, \Phi \Rightarrow \Psi} \sigma L & \frac{\alpha, \Phi \Rightarrow \Psi}{\diamond\alpha, \Box\Phi \Rightarrow \diamond\Psi} \diamond L \\
\frac{\{\Phi \Rightarrow \Psi, \gamma \mid \gamma \in \Gamma\}}{\Phi \Rightarrow \Psi, \bigwedge \Gamma} \wedge R & (\gamma \in \Gamma) \frac{\Phi \Rightarrow \Psi, \gamma}{\Phi \Rightarrow \Psi, \bigvee \Gamma} \vee R & \frac{\Phi \Rightarrow \Psi, \varphi(\sigma\varphi)}{\Phi \Rightarrow \Psi, \sigma\varphi} \sigma R & \frac{\Phi \Rightarrow \Psi, \alpha}{\Box\Phi \Rightarrow \diamond\Psi, \Box\alpha} \Box R \\
\frac{}{\alpha, \alpha^\perp \Rightarrow} \text{idL} & \frac{}{\Rightarrow \alpha^\perp, \alpha} \text{idR} & \frac{}{\alpha \Rightarrow \alpha} \text{id} &
\end{array}$$

Figure 2. Logical rules of system $K\bar{\mu}$. Rules denoted $\star L$ ($\star R$) are *left* (resp. *right*) rules.

Proposition 3.2. Any $K\bar{\mu}$ -derivation whose endsequent has empty consequent is constructive.

Proposition 3.3. There is a $K\bar{\mu}$ -proof from $\Phi \Rightarrow \Psi$ assumptions $\{\Phi_i \Rightarrow \Psi_i \mid i \in A\}$ iff there is a $K\bar{\mu}$ -proof of $\Psi^\perp, \Phi \Rightarrow$ from assumptions $\{\Psi_i^\perp, \Phi_i \Rightarrow \mid i \in A\}$.

Proof:

Replace each sequent in the proof in the way described. In the ‘left to right’ direction, this causes each instance of a right rule ($\star R$) to become an instance of the corresponding left rule ($\star L$) up to exchange, and each left rule ($\star L$) to remain as a left rule (up to exchange). \square

An important derived rule in the various calculi to follow expresses that predicates induce monotone functions:

$$\frac{\alpha \Rightarrow \beta}{\varphi\alpha \Rightarrow \varphi\beta} \varphi$$

Proposition 3.4. The monotonicity rules are derivable in $K\bar{\mu}$: For each φ, α and β , there is a $K\bar{\mu}$ -proof of $\varphi\alpha \Rightarrow \varphi\beta$ from assumptions $\alpha \Rightarrow \beta$.

3.2. Finitary proofs

The first extension of $K\bar{\mu}$ is a rendition of Kozen’s axiomatisation of the modal μ -calculus in the sequent calculus. This calculus, called $K\mu$, extends basic fixed point logic by *initial* sequents, three forms of sequent expressing $\alpha \Rightarrow \alpha$, the structural rule of cut and rules of induction. The new rules are presented in fig. 3. Proofs (with assumptions) in the sense of $K\bar{\mu}$ are, again, the finite derivations (with buds). We refer to proofs in $K\mu$ as *finitary* proofs. The constructive fragment of $K\mu$, i.e., the restriction to rules employing only intuitionistic sequents, is denoted $\text{c}K\mu$.

The transformation of $K\bar{\mu}$ -derivations from Proposition 3.3 carries over to $K\mu$ though Proposition 3.2 does not due to the presence of cut.

$$\frac{\Phi \Rightarrow \Psi, \varphi(\bigwedge(\Phi \cup \Psi^\perp))}{\Phi \Rightarrow \Psi, \nu\varphi} \text{indR} \quad \frac{\varphi(\bigvee(\Phi^\perp \cup \Psi)), \Phi \Rightarrow \Psi}{\mu\varphi, \Phi \Rightarrow \Psi} \text{indL} \quad \frac{\Phi \Rightarrow \Psi, \alpha \quad \alpha, \Lambda \Rightarrow \Xi}{\Phi, \Lambda \Rightarrow \Psi, \Xi} \text{cut}$$

Figure 3. Induction and cut rules of $K\mu$. The induction rules (indL) and (indR) are denoted jointly by (ind).

Proposition 3.5. $K\mu \vdash \Phi \Rightarrow \Psi$ iff $K\mu \vdash \Psi^\perp, \Phi \Rightarrow$.

Proof:

By Proposition 3.3, it suffices to treat the rules in fig. 3. The induction rules commute with the translation by design. The transformed cut rule is replaced by (\vee L) with principal formula $\bigvee\{\alpha, \alpha^\perp\}$ and cut against a $K\mu$ -proof of the same formula. \square

Later we will have need for a more general formulation of the induction rules, introduced in [4]:

$$\frac{\Phi \Rightarrow \Psi, \nu x \varphi(\bigwedge(\Phi \cup \Psi^\perp) \vee x)}{\Phi \Rightarrow \Psi, \nu\varphi} \text{indR} \quad \frac{\mu x \varphi(\bigvee(\Phi^\perp \cup \Psi) \wedge x), \Phi \Rightarrow \Psi}{\mu\varphi, \Phi \Rightarrow \Psi} \text{indL}$$

Instances of the two rules immediately above are referred to as *strong induction*. By (sind) we denote the union of (sindL) and (sindR). The reader is referred to [4, Thm VI.3] for a proof of the next proposition.

Proposition 3.6. The strong induction rules are derivable in $K\mu$. Instances of strong induction with intuitionistic conclusion are provable in $cK\mu$.

Soundness of $K\mu$ with respect to Kripke semantics is straightforward:

Theorem 3.7. If $K\mu \vdash \Phi \Rightarrow \Psi$ then $\bigwedge\Phi \rightarrow \bigvee\Psi$ is valid.

3.3. Illfounded proofs

Rather of expanding basic fixed point logic by induction and cut rules, a sound and complete calculus can be obtained by taking inspiration from the game semantics and admitting infinite derivations as proofs that express winning strategies. These will be the derivations for which every branch carries an infinite thread of a suitable kind. The system $K\mu_\omega$ of *illfounded proofs* – *co-proofs* for short, standing for *co-inductive proofs* – is a refinement of Niwiński and Walukiewicz’s refutation calculus from [26].

Definition 3.8. Let B be a branch of a $K\mu$ -derivation d , and let \mathcal{S}_i be the sequent at $B(i)$. A finite trace through B is a sequence of positions $(p_i)_{i \leq K}$ such that p_0 is a position in \mathcal{S}_0 and for all $i < K$,

1. if $\mathcal{R}_d(B(i))$ is not structural, then $p_{i+1} = p_i$,
2. if $\mathcal{R}_d(B(i))$ is structural, then p_{i+1} is the formula occurrence corresponding to p_i .

We refer to p_i as a *descendent* of p_0 at $B(i)$. An infinite trace is a sequence all of whose finite prefixes are traces. A (finite or infinite) trace $t = (p_i)_{i < K}$ through B induces a weak thread $w = (\mathcal{S}_i(p_i))_{i < K}$ selecting the formula at position p_i in each \mathcal{S}_i . The thread *supported* by t (along B) is the unique thread $\tau = w \circ g$ where $g: L \rightarrow K$ is the strictly increasing function such that $j \in g[L]$ iff p_j is principal at $B(j)$. We say that B *carries* τ if τ is an infinite thread that is support by some trace in B .

For a set of positions P , set $\text{Thd}_P(d, B)$ to be the set of infinite threads supported by a trace t (along B) satisfying $t(0) \in P$. If $t(0)$ is a left (right) position we refer to the thread supported by t as a *left (right)* thread. Note, a thread can be both left and right, even relative to the same branch. We write $\text{Thd}_p(d, B)$ in place of $\text{Thd}_{\{p\}}(d, B)$, and $\text{Thd}(d, B)$ for $\text{Thd}_P(d, B)$ where P is the set of all positions in the root sequent. Given a subsequence \mathcal{S} of the endsequent of d , we may write $\text{Thd}_{\mathcal{S}}(d, B)$ in place of $\text{Thd}_P(d, B)$ where P is the set of positions corresponding to the formulas in \mathcal{S} .

Proposition 3.9. Let d be a derivation in $K\bar{\mu}$ with a branch B . If B carries an infinite thread starting in a guarded formula, then every infinite trace through B supports an infinite thread.

Proof:

Fix a branch B and let $\tau \in \text{Thd}(d, B)$. As $\tau(0)$ is guarded yet infinite, it follows that B passes through a modal rule infinitely often. Thus, for any trace t through B , the weak thread supported by t steps through the subformula relation \triangleright infinitely often, i.e., t supports an infinite thread. \square

Definition 3.10. The calculus $K\mu_\omega$ comprises the rules of $K\bar{\mu}$ (fig. 1 and 2) with the identity rules restricted to variables. A proof in $K\mu_\omega$ is a derivation such that every branch carries either a right ν -thread or a left μ -thread. Proofs in $K\mu_\omega$ are called *illfounded proofs* or *co-proofs*.

We commonly refer to a branch satisfying the condition above as *good*. A branch carrying neither a left μ -, nor right μ -thread may be called *bad*.

Unlike the finitary calculus, the rule of cut is not permitted in co-proofs. The rule can be added, though the condition on proofhood must be weakened to allow for the good threads to initiate from cut formulas: simply replace ‘every branch carries’ by ‘every branch has a suffix which carries’. Our choice of cut-free illfounded proofs is motivated by the reduction of co-proofs to cyclic proofs in Theorem 5.12 which requires a bound on the number of distinct formulas occurring in the initial co-proof that need not exist for illfounded proofs with cut. For the same reason, although we prove admissibility of the cut rule for $K\mu_\omega$ (in the next section), this fact does not provide a method of *eliminating* arbitrary cuts in the illfounded context.

Our first result concerning co-proofs is soundness.

Theorem 3.11. (Soundness of Co-inductive Proofs)

If $K\mu_\omega \vdash \Phi \Rightarrow \Psi$ then $\Phi \Rightarrow \Psi$ is valid.

Proof:

Without loss of generality we may assume that Ψ is empty and Φ comprises a single formula α . Let π be a co-proof of $\alpha \Rightarrow$ and let M be any model. We show that there is no verification of (α, M) .

Assume the contrary, namely that V is a verification with root (α, M) . The verification induces a branch B of π and associated sequence of vertices $(v_i)_i: \omega \rightarrow M$ with the property that for every weak thread τ induced by a trace along B , the sequence $(\tau(j), v_j)_{j < |\tau|}$ is a weak path in V . As B carries a left μ -thread, this contradicts that V is a verification. \square

The proof above easily generalises to co-proofs with cut.

4. The Constructive Fragment

By restricting the rules of $K\mu_\omega$ to intuitionistic sequents, a co-inductive sequent calculus can be given which respects intuitionistic logic. As with the previous calculi, we refer to the calculus as the *constructive fragment* of $K\mu_\omega$ and denoted it $\text{c}K\mu_\omega$. The robustness of the constructive fragment is central to our proof of completeness of (classical) $K\mu$, and is the topic of much of the remainder of this article.

As remarked in the introduction, the term ‘fragment’ is something of a misnomer as the classical calculus $K\mu_\omega$ can be identified with the negative fragment of the constructive μ -calculus under the transformation described in Proposition 3.3. That is, we may reinterpret classical co-proofs as the constructive co-proofs with endsequents $\Phi \Rightarrow \emptyset$.

The following observation sums up the syntactic properties of constructive co-proofs.

Proposition 4.1. Every branch of a constructive co-proof carries at most one right thread and every co-proof with empty consequent is constructive.

Definition 4.2. By $\Phi \Vdash \alpha$ we express that there exists a constructive co-proof of the sequent $\Phi \Rightarrow \alpha$ in which the identity rules have a variable as a principal formula. For formulas α, β , define $\alpha \dashv\vdash \beta$ as $\alpha \Vdash \beta$ and $\beta \Vdash \alpha$.

We begin with some simple closure properties on \Vdash .

Proposition 4.3. Let α, β be formulas, φ, ψ be predicates and x a fresh variable.

1. If $\alpha \Vdash \beta$ and $\varphi x \Vdash \psi x$ then $\varphi \alpha \Vdash \psi \beta$.
2. If $\varphi x \Vdash \psi x$ then $\sigma \varphi \Vdash \sigma \psi$.

Proof:

1. Throughout a constructive co-proof of $\varphi x \Rightarrow \psi x$, uniformly substitute α for x occurring in the context and β for x in the head. The positivity constraints on abstraction ensure that the result is a constructive derivation of $\varphi \alpha \Rightarrow \psi \beta$ in which buds are of the form $\alpha \Rightarrow \beta$. So $\varphi \alpha \Vdash \psi \beta$.

2. Using the same construction as in part 1, from a constructive co-proof of $\varphi x \Rightarrow \psi x$ a constructive derivation of $\sigma \varphi \Rightarrow \sigma \psi$ can be obtained, wherein all buds are identical to the root sequent $\sigma \varphi \Rightarrow \sigma \psi$. Identifying such buds with the root and unravelling to an illfounded derivation witnesses $\sigma \varphi \Vdash \sigma \psi$ for both $\sigma = \nu$ and $\sigma = \mu$. \square

The next result offers our first example of a constructive co-proof.

Proposition 4.4. $\alpha \Vdash \alpha$.

Proof:

By induction on the syntactic complexity of α . Suppose $\alpha = \sigma\varphi$. The induction hypothesis implies $\varphi x \Vdash \varphi x$, whence $\alpha \Vdash \alpha$ by Proposition 4.3. The other cases are straightforward. \square

It was remarked in section 2 that the expanded cover modality allows for an intuitionistic treatment of disjunctive formulas not offered by unary form of the modality. The natural disjunctive formulas arising from the normal-form construction (detailed in Section 6 and implicit in [36, 8]) are intuitionistically equivalent to the starting formula when the expanded cover modality is utilised but not for the non-expanded modality. The distinction occurs at certain stages in the construction of the disjunctive formula where in place of $\nabla(\Gamma, \Delta)$ a formula $\nabla\Gamma \vee \nabla\Gamma'$ is utilised (for Γ' dependent on Γ and Δ). The equivalence of the two forms holds classically but is not derivable in the constructive fragment.

Proposition 4.5. There exist formulas not equivalent over $\text{cK}\mu_\omega$ to any disjunctive formula employing only the non-expanded (unary) cover modality.

Proof:

Suppose β is a disjunctive formula satisfying $\Box x \Vdash \beta$. Inspecting the witnessing co-proof we can, without loss of generality, assume $\beta = \nabla_L(\Gamma, \Delta)$. Soundness implies that $L = \Gamma = \emptyset$ and $\llbracket x \rrbracket \subseteq \llbracket \nabla\Delta \rrbracket$, i.e., $\Gamma \subsetneq \Delta$. In particular, β cannot be formed by use of the unary cover modality only. \square

On the other hand, the expanded cover modality does provide equivalent representations of the usual modalities:

Proposition 4.6. $\Box x \dashv\vdash \nabla(\emptyset, \{x\})$ and $\Diamond x \dashv\vdash \nabla(\{x\}, \{x, \top\})$.

The next observation confirms the claimed interpretation of classical co-proofs in the constructive fragment.

Proposition 4.7. $\text{K}\mu_\omega \vdash \Phi \Rightarrow \Psi$ iff $\Phi, \Psi^\perp \Vdash \perp$.

Proof:

The proof transformation described in Proposition 3.3 preserves the co-proof acceptance condition on branches as it simply dualises traces and threads. So we can assume that Ψ is empty. Moreover, it is clear that $\Phi \Vdash \perp$ iff $\text{cK}\mu_\omega \vdash \Phi \Rightarrow$ by Proposition 4.4 and the definition of \perp . Proposition 3.2 connects the two equivalences. \square

The following result will be useful later.

Proposition 4.8. If $\Phi \Vdash \alpha$ then there is a constructive co-proof of $\Phi \Rightarrow \alpha$ in which principal formulas of contractions are not disjunctive.

Proof:

By permuting contractions away from the root, we can assume that the principal formula any contraction is a conjunction. In particular, we can assume we are working with a contraction-free constructive co-proof using an expanded $(\wedge L)$ rule:

$$(\Delta \subseteq \Gamma) \frac{\Delta, \Phi \Rightarrow \Psi}{\wedge \Gamma, \Phi \Rightarrow \Psi} \wedge^* L$$

Consider a vertex v labelled by the instance of $(\wedge^* L)$ above where the principal formula $\wedge \Gamma$ is disjunctive. In particular, $\Gamma = L \cup \diamond \Gamma' \cup \{\square \vee \Gamma''\}$ for $L \subseteq \text{Lit}$ and $\Gamma' \subseteq \Gamma'' \subseteq \mathbb{D}$. We can assume that $|\Delta| > 1$ as otherwise an instance of either $(\wedge L)$ or (wL) would suffice. By permuting the rule with that above if possible, we may also assume that the child of v is an instance of either an identity or modal rule. In the former case, as L is consistent at most one element of L occurs in the premise, and the instance of $(\wedge^* L)$ can be replaced by $(\wedge L)$. In the latter case, we can assume that $\Delta \subseteq \diamond \Gamma' \cup \{\square \vee \Gamma''\}$. As $\Delta, \Phi \Rightarrow \Psi$ is the conclusion of a modal rule, we have that $\Delta = \{\diamond \gamma, \square \vee \Gamma''\}$ for some $\gamma \in \Gamma'$ and that $\diamond \gamma$ is principal in the premise. As $\gamma \in \Gamma''$, it suffices to choose $\Delta = \{\diamond \gamma\}$. Although this change has the potential to alter left threads by replacing subsequences $\wedge \Gamma \triangleright \square \vee \Gamma'' \triangleright \vee \Gamma'' \triangleright \gamma$ by $\wedge \Gamma \triangleright \diamond \gamma \triangleright \gamma$, it does not change whether the property of carrying a good thread. \square

Theorem 4.9. Let α be a valid conjunctive formula. Then $\emptyset \Vdash \alpha$.

The theorem may appear at odds with our referring to cK_{μ_ω} as an *intuitionistic* sequent calculus given that it derives all classical validities of certain syntactic shape. The answer lies in the restrictive notion of conjunctive formulas. Recall that disjunctive formulas only permit occurrences of conjunction in the form the expanded cover modality. Whence, by definition, if $\beta = \nabla_L(\Gamma, \Delta)$ is an instance of the expanded cover modality, then β^\perp is valid iff $L = \emptyset$ (since L must be consistent) and Γ^\perp contains a valid formula. In particular, no instance of the classical law of excluded middle $\beta \vee \beta^\perp$ qualifies as conjunctive.

Proof:

We perform an exhaustive proof search starting with the sequent $\Rightarrow \alpha$ while staying within the constructive fragment. That is, we build a derivation tree in (an expansion of) cK_{μ_ω} such each vertex is annotated by a intuitionistic sequent. At the case that the head formula is a disjunction, it will be an instance of the binary co- ∇ modality, in which case the proof-search branches to explore the various premises that could be used in some derivation of the sequent. We will show that if α is valid then there is a pruning of the proof-search tree to a constructive co-proof. In the following, we denote by ∇^\perp the modality dual to ∇ :

$$\nabla_L^\perp(\Gamma, \Delta) = \vee(L \cup \square \Gamma \cup \{\diamond \wedge \Delta\})$$

subject to the same constraints as uses of ∇ , namely L is consistent and $\Gamma \subseteq \Delta$. Treatment of ∇^\perp formulas is given by the branching rule:

$$\frac{\Rightarrow \gamma_1 \quad \cdots \quad \Rightarrow \gamma_k}{\Rightarrow \nabla_L^\perp(\{\gamma_1, \dots, \gamma_k\}, \Delta)} \nabla^\perp$$

Note that given a co-proof of some premise, say $\Rightarrow \gamma_i$, a co-proof of the conclusion is obtained via the $(\Box R)$ and $(\bigvee R)$ rules. Also, if Γ is empty then the rule (∇^\perp) has no premises.

Let d be a maximal derivation with endsequent $\Rightarrow \alpha$ in the specified rules. The assumption that α is guarded entails that every branch of d passes through infinitely many instances of (∇^\perp) . Letting \mathcal{S}_u denote the sequent labelling the vertex $u \in d$, on every branch B of d there contains a pair of vertices $u < v \in B$ such that both are conclusions of (∇^\perp) and \mathcal{S}_u and \mathcal{S}_v are identical sequents. Such a pair of vertices is called a *repetition*. In fact, d will be a regular tree: for every repetition (u, v) , $d \uparrow v = d \uparrow u$, i.e., the subderivations of d with root u and v are identical.

We claim that if α is valid then a subtree of d is a $K\mu_\omega$ -proof, where the subtree is determined by collapsing instances of (∇^\perp) to a single premise and replacing the rule as described above. To that aim, suppose that d cannot be pruned in such a way to form a $K\mu_\omega$ -proof. From the regularity assumption we derive a (regular) subtree $d' \subseteq d$ such that for all $v \in d'$,

1. if $\mathcal{R}_d(v) = \nabla^\perp$ then all children of v are present in d' ,
2. if $\mathcal{R}_d(v) = \bigwedge R$ then d' contains exactly one child of v ,
3. The unique thread carried by each branch of d' is μ .

From d' a model $M = \langle T, A \rangle$ of α^\perp can be defined. Let \sim be the equivalence relation on d' given by $u \sim v$ if $u \leq v$ are not separated by an instance of (∇^\perp) . The domain of M is the set of \sim -equivalence classes with $[u] \leq [v]$ iff $u \leq v$. We say that a variable x occurs at $w \in T$ if there exists $v \in w$ such that $\mathcal{S}_d(v)$ either contains x or a formula $\nabla_L^\perp(\Gamma, \Delta)$ with $x \in L$. The label function $A: T \rightarrow E$ is given by $x \in A(w)$ iff x does not occur at w . As α is conjunctive, if $v \in d'$ is the sequent $\Rightarrow \nabla_L^\perp(\Gamma, \Delta)$ then L^\perp is, by assumption, consistent and thus $x \in A([v])$ iff $x \notin L$. If, furthermore, $\Gamma = \emptyset$ then $[v]$ is a bud of M and $M \uparrow [v] \notin \langle \nabla_L^\perp(\Gamma, \Delta) \rangle$ by design.

We have that $M \notin \langle \alpha \rangle$, for if $M \in \langle \alpha \rangle$ with V the witnessing verification then there will be branches of V and d' carrying identical threads, which contradicts condition 3 above. Thus, if α is a valid conjunctive formula, then $\emptyset \Vdash \alpha$. \square

Reformulating Theorem 4.9 we have

Corollary 4.10. If α is disjunctive and α^\perp is valid, then $\alpha \dashv\vdash \perp$.

Proof:

The hypothesis and Theorem 4.9 yields $cK\mu_\omega \vdash \Rightarrow \alpha^\perp$, whence $\alpha \Vdash \perp$ by Proposition 4.7. That also $\perp \Vdash \alpha$ is immediate. \square

Our main result concerning the constructive fragment is:

Theorem 4.11. (Admissibility of Cut in the Constructive Fragment)

If $\Phi \Vdash \alpha$ and $\alpha, \Psi \Vdash \beta$, then $\Phi, \Psi \Vdash \beta$.

Theorem 4.11 is equivalent to statement that the intuitionistic rule of cut:

$$\frac{\begin{array}{c} \vdots \pi \\ \Phi \Rightarrow \alpha \end{array} \quad \begin{array}{c} \vdots \rho \\ \alpha, \Psi \Rightarrow \beta \end{array}}{\Phi, \Psi \Rightarrow \beta} \text{ cut}$$

is admissible within the constructive fragment. The proof of the theorem proceeds by using the derivation with cut as a first approximation to the constructive co-proof of $\Phi, \Psi \Rightarrow \beta$ and performs a co-inductive process of cut elimination to ultimately eliminate this instance of cut entirely.

The sequence of proof transformation that leads to elimination of the cut rule is, in essence, nothing more than Gentzen's method of 'reductive' cut elimination iterated ad infinitum, producing from co-proofs π and ρ above a constructive derivation d with endsequent $\Phi, \Psi \Rightarrow \beta$. Following Mints [24], reductive cut elimination is conceived as a continuous transformation on the space of maximal derivations (relative to a suitable topology) and the challenge is not in describing the derivation d but in confirming that it is an illfounded *proof*.

Preservation of proofhood under cut elimination means identifying in each branch of d either a left μ -thread or right ν -thread. We will demonstrate that the threads carried by any branch of d are all carried by particular branches of π or ρ and that the manner in which cut elimination transforms these two proofs into d necessarily preserves sufficiently many threads to confirm that d is a co-proof.

Continuous and co-inductive cut elimination has been established for numerous systems of illfounded proof including fixed point logics and μ -calculi, for example [18, 12, 29, 11, 28, 1, 7]. The approach presented here builds on Fortier and Santocanale's cut elimination argument [18], adding modalities and lifting the restriction on the form of sequents. Viewed as a cut-elimination method for an intuitionistic logic, our argument is closely related to Arai's well-founded cut-elimination for fixed point predicates over Heyting arithmetic [9].

Rather than focusing on a binary cut rule it is convenient to utilise a multi-premise version encoding a sequence of cuts called the *multicut* rule:

$$\frac{\Phi_1 \Rightarrow \alpha_1 \quad \cdots \quad \Phi_n \Rightarrow \alpha_n \quad \Psi_0, \alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{ mcut}$$

The formulas $\alpha_1, \dots, \alpha_n$, namely the heads of all but the final premise, are referred to as *cut formulas*. The rightmost premise of a multicut is called the *major* premise; other premises are *minor*. A *multicut derivation* is a derivation in the rules of $\text{cK}\bar{\mu} + \text{mcut}$ such that every branch passes through exactly one instance of multicut. If the multicut is at the root we refer to the derivation as *multicut rooted*. It is important to note that the multicut rule includes the *single* premise instance:

$$\frac{\Psi \Rightarrow \beta}{\Psi \Rightarrow \beta} \text{ mcut}$$

which provides a technical advantage in tracking the multicut reductions.

Although the above form of multicut is adequate for the present setting, to account for the reduction rules necessary for either classical or full intuitionistic logic (with implication) a more liberal conception of multi-cut would be required. A version of multicut for intuitionistic μ -calculus with implication is presented in [7], with cut elimination (not merely admissibility) established for a class of

illfounded proofs. For continuous cut-elimination for classical logic, we refer the reader to [28]. The argument we present below can be seen as a special (and simpler) case of [7] for the logic expanded by modal operators.

4.1. Multicut reductions

The local transformations on derivations consist of a collection of ‘reductions’ $d \mapsto d'$ where d and d' are multicut derivations and d is multicut rooted. Thus, we assume that d combines via a multicut the constructive derivations:

- d_i with endsequent $\Phi_i \Rightarrow \alpha_i$ for $1 \leq i \leq n$, constituting the minor premises, and
- e with endsequent $\Psi_0, \alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta$, constituting the major premise.

In addition, we impose the restriction that the minor premise derivations do not contain any right structural rules. This constraint is technically motivated to simplify some cases of the multicut reductions and proof of cut admissibility. In the proof of Theorem 4.11 below we show that the restriction does not reduce the generality of the theorem.

The reduction $d \mapsto d'$ described above is visualised as:

$$\frac{\frac{d_1}{\Phi_1 \Rightarrow \alpha_1} \quad \dots \quad \frac{d_n}{\Phi_n \Rightarrow \alpha_n} \quad \frac{e}{\Psi_0, \alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{mcut} \mapsto d' \quad (1)$$

The multicut reductions fall into three categories:

External reductions These apply if the root inference of any premise is a non-modal rule and the cut formula is not principal.

Internal reductions These require that the final rule of e is a left rule and that α_1 is principal. Some internal reductions additionally require that the final rule of d_1 is a right rule.

Modal reductions These apply if the final rule of every premise is a modal inference.

Internal reductions ‘consume’ the final inference in the appropriate derivations and, possibly, restructure the multicut, whereas external reductions ‘permute’ the final inference from premise(s) with the multicut. The modal reductions exhibit features of both reductions.

Modal reductions We first treat the modal reductions, which apply if all premises end in an instance of one of the two modal rules. There are two reductions, depending on whether the final inference of e is $\Diamond L$ or $\Box R$. The latter case induces the following reduction

$$\begin{array}{c} \frac{\frac{d_1^*}{\Xi_1 \Rightarrow \gamma_1} \Box R \quad \dots \quad \frac{d_n^*}{\Xi_n \Rightarrow \gamma_n} \Box R \quad \frac{e^*}{\gamma_1, \Lambda_1, \dots, \gamma_n, \Lambda_n \Rightarrow \delta} \Box R}{\Xi_1, \Box \Lambda_1, \dots, \Box \Lambda_n \Rightarrow \Box \delta} \text{mcut} \\ \mapsto \frac{\frac{d_1^*}{\Xi_1 \Rightarrow \gamma_1} \quad \dots \quad \frac{d_n^*}{\Xi_n \Rightarrow \gamma_n} \quad \frac{e^*}{\gamma_1, \Lambda_1, \dots, \gamma_n, \Lambda_n \Rightarrow \delta}}{\Xi_1, \Lambda_1, \dots, \Lambda_n \Rightarrow \delta} \Box R \end{array}$$

Following the schema of (1) we have that $\Psi_0 = \emptyset$, and $\Phi_i = \Box \Xi_i$ and $\Psi_i = \Box \Lambda_i$ for each $1 \leq i \leq n$. The case of e ending in $\Diamond L$ is similar but structurally slightly more complex.

External reductions The external reductions all permute the multicut with a rule at one of the premises provided that the principal formula is not a cut formula. One such example is that the major premise is a (non-modal) right rule:

$$\frac{\dots \frac{d_n^*}{\Phi_n \Rightarrow \alpha_n} \quad \frac{\dots \frac{e^*}{\Psi, \alpha_1, \dots \Rightarrow \gamma} \dots}{\Psi_0, \alpha_1, \dots \Rightarrow \beta} \star R}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{mcut} \mapsto \dots \frac{\frac{d_n^*}{\Phi_n \Rightarrow \alpha_n} \quad \frac{e^*}{\Psi_0, \alpha_1, \dots \Rightarrow \gamma}}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \gamma} \text{mcut} \dots \star R$$

The other two scenarios are that some d_i comprises a left rule or e a right rule. The reductions in these cases is clear.

Internal reductions The internal reductions cover the case that the cut formula α_1 is principal in the major premise. We begin with reductions related to structural rules. Suppose, therefore, that e ends with in a left structural rule and $\Psi_0 = \varepsilon$. This may be (eL), (wL) or (cL). We present the latter. The associated internal reduction duplicates the first minor premise of the rule and inserts a (left) contraction below the multicut unless $\Phi_1 = \varepsilon$:

$$\frac{\frac{d_1}{\Phi_1 \Rightarrow \alpha_1} \quad \dots \quad \frac{\frac{e^*}{\alpha_1, \alpha_1, \Psi_1, \dots \Rightarrow \beta} \text{cL}}{\alpha_1, \Psi_1, \dots \Rightarrow \beta} \text{mcut}}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \mapsto \frac{\frac{d_1}{\Phi_1 \Rightarrow \alpha_1} \quad \frac{d_1}{\Phi_1 \Rightarrow \alpha_1} \quad \dots \quad \frac{e^*}{\alpha_1, \alpha_1, \Psi_1, \dots \Rightarrow \beta}}{\frac{\Phi_1, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{(c)}} \text{mcut}$$

If the final inference in e is a logical rule, then a reduction is applicable provided that d_1 also concludes with a logical rule. The case of that e ends in (id) induces the reduction:

$$\frac{\frac{x \Rightarrow x \text{ id}}{x \Rightarrow x} \quad \frac{x \Rightarrow x \text{ id}}{x \Rightarrow x} \text{mcut}}{x \Rightarrow x} \mapsto \frac{x \Rightarrow x \text{ id}}{x \Rightarrow x}$$

The reduction is similar if e ends instead with (idL), though in that case there could be two cut formulas, in which case we require that d_1 ends in the identity rule also.

The remaining logical multicut reductions follow a common form where the first minor premise d_1 consists of a right rule and the major premise the corresponding left rule. For \star a logical operator other than a modality, the following reduction is available:

$$\frac{\dots \frac{d_1^*}{\Phi_1 \Rightarrow \alpha} \dots \star R \quad \dots \frac{e^*}{\alpha, \Psi_1, \dots \Rightarrow \beta} \dots \star L}{\frac{\Phi_1 \Rightarrow \alpha_1 \quad \dots \quad \alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{mcut}} \mapsto \frac{\frac{d_1^*}{\Phi_1 \Rightarrow \alpha} \quad \dots \quad \frac{e^*}{\alpha, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{mcut}$$

The number of premises to $\star R/\star L$ depends on the operator \star . For $\star = \bigvee$ the right rule $\star R$ is unary and $\star L$ is multi-premise whereas for $\star \in \{\mu, \nu\}$ both rules are unary. In any case, there will always be one premise of each rule with identical minor formula.

By design, every multicut with only non-trivial immediate subderivations is reducible:

Proposition 4.12. Let d be a multicut rooted derivation. If no immediate subderivation of d is a bud and no minor premise ends in (wR), then there exists a multicut derivation d' such that $d \mapsto d'$.

Proof:

If no external or modal reduction is applicable then every minor premise ends in a right rule and at least one cut formula is principal in the major premise. If the final rule of the major subderivation is structural, then a reduction is applicable. Otherwise, an internal logical reduction is applicable. \square

4.2. Admissibility of cut

The multicut reductions can be used to establish the admissibility of the cut rule in the constructive fragment of $K\mu_\omega$. As remarked, the main difficulty is showing that the maximal derivation obtained by co-recursively applying is indeed a co-proof. The argument proceeds by relating the branches and threads of the resultant derivation to the initial co-proofs. The desired connection is expressed by the next lemma.

Let $\sim \subseteq T_0 \times T_1$ be a relation between trees T_0 and T_1 . For branches B_0 and B_1 through T_0 and T_1 respectively, we write $B_0 \sim B_1$ iff for every $u \in B_0$ there exists $v \in B_1$ such that $u \sim v$ and for every $v \in B_1$ there exists $u \in B_0$ such that $u \sim v$. For T a set of threads, let $T^\perp := \{\tau^\perp \mid \tau \in T\}$. A function $h: T \rightarrow T'$ between trees is *order-preserving* if $h(u) \leq h(v)$ whenever $u \leq v \in T$. For an order-preserving function $h: T \rightarrow T'$ and a branch B of T , we write $h(B)$ to denote the branch of T' traversing all $h(B(i))$ if this is unique, or else $h(B) := h(B(k))$ if k is such that $h(B(k)) = h(B(k+i))$ for all i .

Lemma 4.13. Let $\pi \vdash \Phi \Rightarrow \alpha$ and $\rho \vdash \alpha, \Psi \Rightarrow \beta$ be constructive co-proofs and π not containing (wR). There exists a maximal constructive derivation d in $K\mu_\omega$ with endsequent $\Phi, \Psi \Rightarrow \beta$, an order-preserving function $h: d \rightarrow \rho$ and a relation $\sim \subseteq \pi \times d$ satisfying, for every branch $B \subseteq d$,

1. $\text{Thd}_{\Psi \Rightarrow \beta}(\rho, h(B)) = \text{Thd}_{\Psi \Rightarrow \beta}(d, B)$.
2. $\bigcup \{ \text{Thd}_{\Phi \Rightarrow}(\pi, E) \mid B \sim E \} = \text{Thd}_{\Phi \Rightarrow}(d, B)$.
3. $\text{Thd}_{\alpha \Rightarrow}(\rho, h(B)) = \bigcup \{ \text{Thd}_{\alpha \Rightarrow}^\perp(\pi, E) \mid E \sim B \}$.

Proof:

The derivation d is defined as the limit of a increasing sequence of finite $K\mu_\omega$ -derivations $(d_j)_{j < \omega}$. It is technically convenient to assume that the combining rules (w) and (c) are available as structural rules in the formation the d_j . Each derivation d_j will be associated a function m_j which assigns to the buds of d_j instances of the multicut rule whose conclusion matches the bud sequent. This data will be termed a *multicut-tipped* derivation, defined to be a pair (d, m) of a finite derivation d and a function m mapping each bud b of d to a sequence $m(b) = (u_1, \dots, u_n, u) \in \pi^n \times \rho$ such that the sequents $\mathcal{S}_\pi(u_1), \dots, \mathcal{S}_\pi(u_n)$ enumerate the minor premises of a multicut, $\mathcal{S}_\rho(u)$ is the major premise and the sequent at b is the conclusion of this multicut. Formally, m should also specify which formula

occurrences in $\mathcal{S}_\rho(u)$ are cut formulas but we leave this information implicit. Thus, the function m in a multicut-tipped derivation (d, m) associates to each bud b of the multicut-rooted derivation

$$\frac{\pi \uparrow u_1 \quad \cdots \quad \pi \uparrow u_n \quad \Psi_0, \alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{ mcut} \quad \text{where } m(b) = (u_1, \dots, u_n, u) \quad (2)$$

The construction begins with the trivial multicut-tipped derivation (d_0, m_0) where d_0 is a single bud assigned the initial two-premise multicut:

$$d_0 := \left\{ \frac{}{\Phi, \Psi \Rightarrow \beta}^b \text{ and } m_0: \varepsilon \rightarrow (\varepsilon, \varepsilon) \text{ representing } \frac{\pi \quad \alpha, \Psi \Rightarrow \beta}{\Phi, \Psi \Rightarrow \beta} \text{ mcut} \right.$$

Let b be a bud of d_j and suppose $m_j(b)$ describes the multicut in (2) above. Proposition 4.12 implies that some multicut reduction is applicable. Applying that reduction replaces the multicut at b by a finite derivation d in which each bud is assigned a multicut built from subderivations of $\pi \uparrow u_1, \dots, \pi \uparrow u_n, \rho \uparrow u$. Some reductions replace the multicut in (2) directly with another multicut while other reductions, such as the external reductions, insert a non-trivial derivation whose buds are multicuts. We describe the process for the three forms of multicut reduction.

Internal reduction We present the case of the non-modal logical rule. Assume that the reduction applied to the multicut in (2) for the connective \star is:

$$\begin{array}{c} \frac{\cdots \quad \frac{\pi \uparrow u_1 i}{\Phi_1 \Rightarrow \gamma} \quad \cdots}{\Phi_1 \Rightarrow \alpha_1} \star R \quad \cdots \quad \frac{\cdots \quad \frac{\rho \uparrow u j}{\gamma, \Psi_1, \dots \Rightarrow \beta} \quad \cdots}{\alpha_1, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta} \star L}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{ mcut} \\ \mapsto \frac{\pi \uparrow u_1 i \quad \cdots \quad \gamma, \Psi_1, \dots, \alpha_n, \Psi_n \Rightarrow \beta}{\Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{ mcut} \end{array}$$

In this case, b remains a bud in d_{j+1} but is assigned the multicut $m_{j+1}(b) = (u_1 i, u_2, \dots, u_n, u j)$. The internal reduction for structural rules is analogous.

External reduction Suppose the rule at u (in ρ) is $(\star R)$ and the following reduction is applied.

$$\begin{array}{c} \frac{\pi \uparrow u_1 \quad \cdots \quad \pi \uparrow u_n \quad \cdots \quad \frac{\Psi_0, \alpha_1, \dots \Rightarrow \beta_i \quad \cdots}{\Psi_0, \alpha_1, \dots \Rightarrow \beta} \star R}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta} \text{ mcut} \\ \mapsto \frac{\cdots \quad \frac{\pi \uparrow u_1 \quad \cdots \quad \pi \uparrow u_n \quad \delta_i, \Psi, \alpha_1, \dots \Rightarrow \beta_i}{\Psi_0, \Phi_1, \Psi_1, \dots, \Phi_n, \Psi_n \Rightarrow \beta_i} \text{ mcut} \quad \cdots}{\Psi_0, \alpha_1, \dots \Rightarrow \beta} \star R \end{array}$$

Let (d, m) be the finite multicut-tipped derivation comprising only the rule $(\star L)$ with the i -th premise assigned the multicut (u_1, \dots, u_n, ui) . Then d_{j+1} extends d_j by inserting d in place of the bud b , and m_{j+1} maps the new buds to the corresponding multicut.

Modal reduction The case of the modal reduction combines the two reduction forms above.

With $(d_j)_{j < \omega}$ defined, where d_{j+1} extends d_j , set d to be the limit of the sequence, defined in the obvious way. For $(u, v) \in \pi \times d$, set $u \sim v$ iff v is a bud of some d_j and u occurs as a minor premise of $m_j(v)$. The function $h: d \rightarrow \rho$ is defined as tracking the major premise of the multicut through the approximations of d . Every vertex in d is a bud of at least one of the derivations approximating d . For $v \in d$ set $h(v)$ to be the major premise of the multicut $m_j(v)$ where j is minimal such that $v \in d_j$. This function is order-preserving.

Let B be any branch of d , and let b_j be the bud of d_j along B . Tracing the sequence of multicut reductions that generate B , it is clear that the second condition of the lemma holds. To establish the first and third condition we must examine $h(B)$. We first consider the case that $h(B)$ is not a branch, i.e., $h(b_j) = h(b_{j+1})$ for all $j > k$. Recall that if $b_j < b_{j+1}$ then $h(b_{j+1})$ is the major premise of $m_{j+1}(b_{j+1})$ if this exists. Hence, for $h(B)$ not to be a branch, no subsequent reduction alters the major premise. That is, no branch E of π for which $E \sim B$ carries an α -trace due to only finitely many right rules along E , meaning that the three conditions of the lemma are met in this case.

Thus, we may assume that B is a branch of d and $h(B)$ defines a branch of ρ . Tracing the multicut reductions B , every left α -thread carried by $h(B)$ induces a branch $E \sim B$ of π carrying the dual thread, and vice versa. Hence, the third condition is met. Furthermore, each non-cut thread $\tau \in \text{Thd}_{\Psi \Rightarrow \beta}(\rho, h(B))$ is carried by B by virtue of the external reductions, whence the first condition.

All that remains is to show that d is maximal. Suppose the contrary, namely that d has a bud b . Let k be such that $b \in d_k$. In particular, b is a bud in d_{k+j} for all j meaning that the multicut reduction which transforms $m_{k+j}(b)$ to $m_{k+j+1}(b)$ must be internal for all j . But this scenario leads to a sequence of branches E_1, \dots, E_n of π and a branch F of ρ such that the E_i carry only right threads, F carries left α -threads only, and $\bigcup_{i \leq n} \text{Thd}(\pi, E_i) = \text{Thd}^\perp(\rho, F)$, contradicting that π and ρ are co-proofs. \square

Theorem 4.11 follows directly from Lemma 4.13:

Proof of Theorem 4.11: Suppose $\Phi \Vdash \alpha$ and $\alpha, \Psi \Vdash \beta$. Let $\hat{\alpha}$ be the result of augmenting each (derived) subformula of α by disjunction with \perp . Then $\Phi \Vdash \hat{\alpha}$ by replacing in the co-proof of $\Phi \Vdash \alpha$ every head formula γ by $\hat{\gamma}$ and inserting an instance of $(\vee R)$ after every right logical rule. We can, furthermore, assume that the co-proof of $\Phi \Rightarrow \hat{\alpha}$ does contain instance of (wR) since the principal in this case will be $\hat{\gamma} \vee \perp$ for some $\gamma \in \text{Sub}(\alpha)$, and a minor formula of \perp can be chosen instead. Also, $\hat{\alpha}, \Psi \Vdash \beta$ is easily established. Thus, we may assume constructive co-proofs $\pi \vdash \Phi \Rightarrow \hat{\alpha}$ and $\rho \vdash \hat{\alpha}, \Psi \Rightarrow \beta$ with π not containing any right structural rules.

Applying Lemma 4.13, we obtain a maximal derivation d satisfying properties 1–3. Let B be a branch of d and consider the branch $h(B)$ of ρ (if $h(B)$ is not a branch a similar argument applies). Towards a contradiction, suppose B is bad. As $h(B)$ is good, there must be a left μ -thread $\tau \in \text{Thd}_{\alpha \Rightarrow}(\rho, h(B))$ initiating from the cut formula α . By condition 3, $\tau \in \text{Thd}_{\Rightarrow \alpha}^\perp(\pi, E)$ for some $E \sim B$. As every path in a constructive co-proof carries at most one right trace (Proposition 4.1) and τ is μ , there is a branch $E \sim B$ carrying a left μ -thread. Condition 2 implies that $\text{Thd}(d, B)$ contains the same thread, whence B is good. \square

Theorem 4.11 can be generalised in two ways. The method of multicut admissibility can be lifted to a (partial) cut elimination result to the effect that any co-proof with finitely many distinct cut formulas (but, potentially, infinitely many cut *occurrences*) can be transformed into a cut-free co-proof. The proof, described in [7], requires a more complex conception of multi-cut than presented above and a more careful analysis of the limit co-derivation. The second generalisation is to a cut-admissibility theorem for the constructive modal μ -calculus including implication. The two generalisations are essentially equivalent: by projecting cuts into the root sequent a co-proof of $\Phi \Rightarrow \alpha$ with finitely many distinct cut-formulas β_0, \dots, β_n can be directly transformed into a cut-free co-proof of $\hat{\beta}, \Phi \Rightarrow \alpha$ where $\hat{\beta} = \nu x (\Box x \wedge \bigwedge \{\beta_i \rightarrow \beta_i \mid i \leq n\})$.

4.3. Guardedness

Building on the closure properties at the start of this section, we can show that every formula is constructively equivalent to a guarded formula. The argument follows the usual transformation (see, e.g., [26]) of partially unfolding quantifiers, permuting propositional connectives and trivialising the remaining unguarded variable occurrences, though with a reliance on constructive co-proofs. The next proposition, proved by syntactic induction on formulas, presents the main reductions.

Proposition 4.14. For all predicates φ, ψ , formulas α , sets $\Delta_0, \dots, \Delta_n$, and operators $\sigma \in \{\mu, \nu\}$ and $\bigcirc \in \{\wedge, \vee\}$:

1. $\sigma x. \varphi x \dashv\vdash \sigma x. \sigma \varphi$ where x is permitted to occur freely in φ .
2. $\sigma(\varphi \cdot \psi) \dashv\vdash \varphi(\sigma(\psi \cdot \varphi))$ where $\varphi \cdot \psi = \lambda x. \varphi(\psi x)$ for x not occurring in φ, ψ .
3. $\bigcirc\{\alpha\} \dashv\vdash \alpha$.
4. $\bigcirc(\Delta_0 \cup \{\bigcirc\Delta_1\}) \dashv\vdash \bigcirc(\Delta_0 \cup \Delta_1)$.
5. $\bigcirc\{\bigcirc^\perp \Delta_i \mid i \leq n\} \dashv\vdash \bigcirc^\perp \{\bigcirc\Delta \mid |\Delta \cap \Delta_i| = 1 \text{ for all } i \leq n\}$.

Proposition 4.15. Suppose $\alpha \dashv\vdash \hat{\alpha}$. If β contains α as a literal subformula and $\hat{\beta}$ is the result of replacing one of these occurrence of α by $\hat{\alpha}$, then $\beta \dashv\vdash \hat{\beta}$.

Proof:

Induction on the syntactic complexity of β using Proposition 4.3. □

Theorem 4.16. (Expressive Adequacy of the Guarded Fragment)

For every α there exists a guarded formula $\gamma \dashv\vdash \alpha$.

The proof of the theorem makes use of transitivity of $\dashv\vdash$ to iterate the ‘single-step’ reductions above.

Proof:

Call a predicate $\varphi = \lambda x. \alpha$ *guarded* if every occurrence of x in α is under the scope of a modal operator, and *strongly unguarded* if *no* occurrence of x in α is under the scope of a modal operator. Via the first equivalence of Proposition 4.14 and Proposition 4.15, every formula α can be converted

to a formula $\hat{\alpha} \dashv\vdash \alpha$ in which every subpredicate φ in $\hat{\alpha}$ is either guarded or strongly unguarded. To complete the reduction to guarded form it suffices to observe the following equivalences and apply Proposition 4.15 iteratively.

- If φ is strongly unguarded, then $\mu\varphi \dashv\vdash \varphi\perp$ and $\nu\varphi \dashv\vdash \varphi\top$.

To establish the claim, first apply proposition 4.14 to reduce φ to the forms $\lambda x. \bigvee\{\bigwedge\Gamma_i \mid i \leq m\}$ and $\lambda x. \bigwedge\{\bigvee\Delta_i \mid i \leq n\}$ in which all occurrences of x are as elements of the Γ_i s and Δ_j s. Propositional logic confirms that $\varphi(\varphi x) \dashv\vdash \varphi x$, whence the claim obtains. \square

5. Cyclic Proofs

We now introduce a calculus in which proofs are finite representations of co-inductive proofs. That is, proofs in the calculus are finite derivations for which each bud can be connected to an earlier occurring vertex, called the companion, labelled by, essentially, the same sequent whose unravelling to a derivation induces a co-proof. Our notion of *cyclic* proof is more constrained than merely asserting that the tree unravelling of the cyclic derivation is a co-proof. Assessing whether a cyclic derivation is a *cyclic proof* in our sense is a purely syntactic condition on the local cycle between companion and bud. This is in contrast to requiring that the tree unravelling is its self a cyclic proof, for which one would have to confirm the thread condition on every induced cycle, not merely the simple ones.

In the context of the proof of completeness of $K\mu$, the cyclic calculus $C\mu$ introduced below plays two roles. One is as an intermediate system between the illfounded and finitary calculi. The second is as a framework for the proof-search method on the path to the disjunctive normal theorem (cf. Section 6). In the latter role, we will only utilise the concept of a derivation in $C\mu$ and not cyclic *proofs* per se. The system $C\mu$ is essentially the calculus Clo of cyclic proofs of [4] extended with the rule of cut, though the presentation below simplifies some details compared to earlier works. It is, however, straightforward to show that the cut-free fragment of $C\mu$ derives the same sequents as Clo .

It was claimed in [4] that the cut-free fragment of $C\mu$ is complete for the modal μ -calculus. A counter-example was recently discovered by Johanness Kloibhofer [22] that confirms the contrary. As a consequence of the present results, the cut-free fragment of $C\mu$ is complete for sequents $\Phi \Rightarrow \alpha$ where Φ consists entirely of $\Pi(\mathbb{D})$ formulas. In Section 7 we prove that in the presence of cut, partial completeness can be lifted to full completeness.

As mentioned, the objects of deduction in $C\mu$ are *annotated* sequents. These are ordinary sequents in which quantified subformulas are annotated by symbols from a designated set and the sequent as a whole carries a global ‘control’ on the annotations contained within. Fix an infinite set N ; elements of N are called *names*. An *annotation* is a finite non-repeating sequence of names. The set of annotations is denoted A . Each annotation a is associated variants of the quantifiers, marked as ν^a and μ^a . The plain ‘unannotated’ quantifiers ν and μ are identified with the quantifiers associated to the empty annotation ε . Annotated formulas are generated by the grammar that admits these new quantifiers. To reduce confusion, in the following we often refer to formulas without annotations as *plain* formulas.

Definition 5.1. The annotated formulas are formed by expanding the syntax of plain formulas with

quantifiers ν^a and μ^a for each annotation a :

$$\begin{aligned}\alpha &::= x \mid \Box\alpha \mid \Diamond\alpha \mid \bigwedge\Gamma \mid \bigvee\Gamma \mid \mu^a\varphi \mid \nu^a\varphi \quad (a \in \mathbf{A}) \\ \Gamma &::= \emptyset \mid \Gamma \cup \{\alpha\} \\ \varphi &::= \lambda x. \alpha, \text{ for } x\text{-positive } \alpha.\end{aligned}$$

If the quantifier σ^a occurs in α we say that a , and $a(i)$ for each $i < |a|$, *occurs* in α . For $\sigma \in \{\mu, \nu\}$, the σ -annotated formulas are the annotated formulas in which only the quantifier σ has non-trivial annotations.

Negation is lifted to annotated formulas by setting $(\sigma^a\varphi)^\perp := \bar{\sigma}^a\varphi^\perp$ where $\bar{\sigma} = \sigma^\perp$. The (plain) formula obtained by eliminating all annotations from α is denoted α^- . We write that α^- *underlies* α , and α *annotates* α^- .

Definition 5.2. An *annotated sequent* is an expression $c : \Phi \Rightarrow \Psi$ where c is an annotation, called the *control*, and Φ and Ψ are finite sequences of μ -annotated and, respectively, ν -annotated formulas such that every annotation occurring in $\Phi \Rightarrow \Psi$ is a subsequence of c . Plain sequents are identified with annotated sequents with empty control.

Given an annotated sequent \mathcal{S} and a sequence or set of names c , let $\mathcal{S}^{\upharpoonright c}$ be the sequent resulting from removing all names in \mathcal{S} not present in c . Note, if \mathcal{S} is an annotated sequent, then so is $\mathcal{S}^{\upharpoonright c}$.

We now proceed with the definition of the cyclic sequent calculus $C\mu$. The logical rules of $C\mu$ are presented in fig. 4 and 5 and comprise the natural liftings of $K\mu_\omega + \text{cut}$ to annotated sequents with the addition of structural rules for manipulating annotations. Implicit in all rules is that the premises and conclusions are annotated sequents. In particular, in instances of $(\nu^n R)$ the name n acts as an eigenvariable as it cannot occur in the endsequent (since cn is an annotation), and cut formulas are always plain (as they are both μ - and ν -annotated). The rules (aR) and (aL) allows annotations within subformulas to be appended by names present in the control. We denote by (aw) the rule incorporated finite iterations of the annotation weakening and exchange rules. The rule (cw) expands the control by a fresh name, though in the following we assume this rule also incorporates finite iterations of the operation. The definition of traces lifts to derivations in $C\mu$ in the obvious way.

Definition 5.3. (Cyclic proofs)

The system $C\mu$ of *cyclic proofs* has the rules listed in fig. 4 and 5, with the former comprising the *structural* rules. A proof in $C\mu$, called a *cyclic proof*, is a finite derivation d whose root is an unannotated sequent and is such that for every path $(v_i)_{i \leq k}$ from the root of d to a bud, there exists $l < k$ and $n \in \mathbf{N}$ satisfying three conditions:

1. The rule at v_l is either $(\nu^n R)$ or $(\mu^n L)$,
2. the name n occurs in the control of v_i for all $l < i \leq k$, and
3. $\mathcal{S}_d(v_k)^{\upharpoonright c_l} = \mathcal{S}_d(v_l)$ where c_l is the control of v_l .

The vertex v_l above is referred to as a *companion* of the bud v_k .

The *constructive fragment* of $C\mu$ is the system $cC\mu$ restricting the rules of $C\mu$ to intuitionistic sequents, namely sequents with at most one head formula.

$$\begin{array}{cccc}
\frac{c : \Phi, \alpha, \Lambda \Rightarrow \Psi}{c : \alpha, \Phi, \Lambda \Rightarrow \Psi} \text{eL} & \frac{c : \alpha, \alpha, \Phi \Rightarrow \Psi}{c : \alpha, \Phi \Rightarrow \Psi} \text{cL} & \frac{c : \Phi \Rightarrow \Psi}{c : \Lambda, \Phi \Rightarrow \Psi} \text{wL} & \frac{c : \varphi(\mu^a \psi), \Phi \Rightarrow \Psi}{c : \varphi(\mu^{an} \psi), \Phi \Rightarrow \Psi} \text{aL} \\
\frac{c : \Phi \Rightarrow \Psi, \alpha, \Lambda}{c : \Phi \Rightarrow \Psi, \Lambda, \alpha} \text{eR} & \frac{c : \Phi \Rightarrow \Psi, \alpha, \alpha}{c : \Phi \Rightarrow \Psi, \alpha} \text{cR} & \frac{c : \Phi \Rightarrow \Psi}{c : \Phi \Rightarrow \Psi, \Pi} \text{wR} & \frac{c : \Phi \Rightarrow \Psi, \varphi(\nu^a \psi)}{c : \Phi \Rightarrow \Psi, \varphi(\nu^{an} \psi)} \text{aR} \\
\frac{}{c : \Phi \Rightarrow \Psi} \text{b} & \frac{c : \Phi \Rightarrow \Psi, \alpha \quad c : \alpha, \Phi \Rightarrow \Psi}{c : \Phi \Rightarrow \Psi} \text{cut} & \frac{cd : \Phi \Rightarrow \Psi}{cad : \Phi \Rightarrow \Psi} \text{cw}
\end{array}$$

Figure 4. Structural rules of $C\mu$.

$$\begin{array}{cccc}
(\gamma \in \Gamma) \frac{c : \gamma, \Phi \Rightarrow \Psi}{c : \bigwedge \Gamma, \Phi \Rightarrow \Psi} \bigwedge \text{L} & \frac{\{c : \gamma, \Phi \Rightarrow \Psi \mid \gamma \in \Gamma\}}{c : \bigvee \Gamma, \Phi \Rightarrow \Psi} \bigvee \text{L} & \frac{c : \alpha, \Phi \Rightarrow \Psi}{c : \diamond \alpha, \square \Phi \Rightarrow \diamond \Psi} \diamond \text{L} & \frac{c : \varphi(\sigma \varphi), \Phi \Rightarrow \Psi}{c : \sigma \varphi, \Phi \Rightarrow \Psi} \nu \text{L} \\
\frac{\{c : \Phi \Rightarrow \Psi, \gamma \mid \gamma \in \Gamma\}}{c : \Phi \Rightarrow \Psi, \bigwedge \Gamma} \bigwedge \text{R} & (\gamma \in \Gamma) \frac{c : \Phi \Rightarrow \Psi, \gamma}{c : \Phi \Rightarrow \Psi, \bigvee \Gamma} \bigvee \text{R} & \frac{c : \Phi \Rightarrow \Psi, \alpha}{c : \square \Phi \Rightarrow \diamond \Psi, \square \alpha} \square \text{R} & \frac{c : \Phi \Rightarrow \Psi, \varphi(\sigma \varphi)}{c : \Phi \Rightarrow \Psi, \sigma \varphi} \mu \text{R} \\
\frac{cn : \varphi(\mu^{an} \varphi), \Phi \Rightarrow \Psi}{c : \mu^a \varphi, \Phi \Rightarrow \Psi} \mu^n \text{L} & \frac{cn : \Phi \Rightarrow \Psi, \varphi(\nu^{an} \varphi)}{c : \Phi \Rightarrow \Psi, \nu^a \varphi} \nu^n \text{R} & (\Phi^\perp \cup \Psi = \{x, x^\perp\}) \frac{}{\varepsilon : \Phi \Rightarrow \Psi} \text{id}
\end{array}$$

Figure 5. Logical rules of $C\mu$.

Inserting instances of structural rules as necessary (namely (aw), (cw) and (e)), it can be assumed that bud–companion pairs are unique and have the form represented in fig. 6.

The motivation behind the term ‘cyclic proof’ is that identifying buds with their companions gives rise to a finite presentation of a certain illfounded proof. Since co-proofs do not utilise the cut rule, this reduction of $C\mu$ to $K\mu_\omega$ applies to cut-free proofs only. The next theorem illustrates the construction of a co-proof from a cyclic proof.

Theorem 5.4. If $C\mu \vdash \Phi \Rightarrow \Psi$ via a cut-free proof, then $K\mu_\omega \vdash \Phi \Rightarrow \Psi$. Likewise, for systems $cC\mu$ and $cK\mu_\omega$ respectively.

Proof:

Consider a cut-free cyclic proof d which, without loss of generality, we assume has (unique) bud–companion pairs matching the scheme of fig. 6. Inserting a further instance of either (νR) or (μL) , we can view the bud–companion pair in as a literal repetition of the sequent labelling the child of the companion. Identifying the new bud with its repeat induces a maximal derivation in $C\mu$ which, if all

$$\begin{array}{cc}
\frac{}{cn : \Phi \Rightarrow \Psi, \nu^{an} \varphi} \text{b} & \frac{}{cn : \mu^{an} \varphi, \Phi \Rightarrow \Psi} \text{b} \\
\vdots & \vdots \\
\frac{cn : \Phi \Rightarrow \Psi, \varphi(\nu^{an} \varphi)}{c : \Phi \Rightarrow \Psi, \nu^a \varphi} \nu^n \text{R} & \frac{cn : \varphi(\mu^{an} \varphi), \Phi \Rightarrow \Psi}{c : \mu^a \varphi, \Phi \Rightarrow \Psi} \mu^n \text{L}
\end{array}$$

Figure 6. Schematic form of bud–companion pairs (leaf and conclusion vertex respectively).

annotations are removed, describes a maximal derivation in $K\mu_\omega$. Let B be a branch of this derivation and let c be the longest annotation that is a prefix of infinitely many controls in B . Necessarily, c has the form dn where n is the name associated to the bud–companion pair whose companion is closest to the root among those that appear infinitely often along B . As every sequent in the suffix of B contains a formula with n in its annotation, this branch must be good. The same interpretation embeds cut-free $cC\mu$ in $cK\mu_\omega$. \square

The admissibility of cut in $cK\mu_\omega$ is not sufficient to describe an interpretation of arbitrary $cC\mu$ -proofs in $cK\mu_\omega$ because cuts in a cyclic proof which straddle the bud–companion pairing induce infinitely many cuts in tree unravelling. A more general argument establishing cut *elimination*, based on for instance [7], would be necessary for this result.

The converse to theorem 5.4, that every $K\mu_\omega$ -derivable sequent admits a cut-free proof in $C\mu$, was claimed in [4] but refuted in [14, 22]. In the presence of cut however, $C\mu$ is a complete calculus (cf. Theorems 5.10 and 7.7 below). A direct reduction of co-proofs to cyclic proofs is blocked by the ‘eigenvariable’ condition implicit in admissible bud–companion pairs: the constraint on the ‘ n ’ in fig. 6 requires that exactly one formula in the sequent is labelled by the name. Simply annotating the traces in a co-proof with the goal that every branch has a finite prefix satisfying the cyclic path condition, cannot work unless augmented with a method to separate conflicting annotations into distinct branches [22]. The cut rule offers this possibility but at an obvious cost.

More liberal notions of the bud–companion pair can be given for which the corresponding definition of cyclic proof is cut-free complete. The first complete system of cyclic proofs for the modal μ -calculus in the style of $C\mu$ is the Jungteerapanich–Stirling ‘tablueax’ system of [20, 32], realised as a sequent calculus in [4, 2]. That calculus operates on a slightly different notion of annotated sequent but can be adapted to the present form without difficulty. A version of the Jungteerapanich–Stirling calculus is utilised in the next section to establish the disjunctive normal form theorem.

An alternative cyclic proof system, closer in conception to $C\mu$, was introduced by the authors in [3]. That calculus is also cut-free complete [2, 6] and has been used to establish Lyndon interpolation for the modal μ -calculus [5]. Generic methods to derive cyclic proof systems from illfounded calculi are explored in [14, 25].

5.1. From cyclic to finitary proofs

Whereas the connection between cyclic and illfounded proofs is complex, the relation to finitary proofs is (relatively) simple: the systems $C\mu$ and $K\mu$ prove the same (plain) sequents, and every $cC\mu$ -provable sequent is provable in $cK\mu$. Furthermore, these result are witnessed by direct translations between cyclic and finitary proofs. Proving these reductions will be our first task. Some observations will be useful in this regard. The first concerns the derivability via cyclic proofs of arbitrary identity sequents, closure under negation and the monotonicity rule. The argument is identical to that of Proposition 4.4, but employs annotations to discharge the buds.

Proposition 5.5.

1. $cC\mu \vdash \alpha \Rightarrow \alpha$ for every plain formula α .

$$\begin{array}{c}
\frac{\frac{\varepsilon : \alpha \Rightarrow \varphi\alpha}{n : \alpha \Rightarrow \varphi\alpha} \text{cw} \quad \frac{\frac{\frac{*}{n : \alpha \Rightarrow \nu^n \varphi} \text{b} \quad \vdots \text{Prop. 5.5}}{n : \varphi\alpha \Rightarrow \varphi(\nu^n \varphi)} \text{cut}}{n : \alpha \Rightarrow \varphi(\nu^n \varphi)} \nu^n \text{R} \\
\varepsilon : \alpha \Rightarrow \nu\varphi (*)
\end{array}
\quad
\begin{array}{c}
\frac{\frac{\frac{*}{n : \mu^n \varphi \Rightarrow \alpha} \text{b} \quad \vdots \text{Prop. 5.5}}{n : \varphi(\mu^n \varphi) \Rightarrow \varphi\alpha} \text{cut} \quad \frac{\varepsilon : \varphi\alpha \Rightarrow \alpha}{n : \varphi\alpha \Rightarrow \alpha} \text{cw}}{n : \varphi(\mu^n \varphi) \Rightarrow \alpha} \nu^n \text{R} \\
\varepsilon : \mu\varphi \Rightarrow \alpha (*)
\end{array}$$

Figure 7. The induction rules (indR) and (indR) with a single side formula expressed as constructive cyclic proofs. The companion of the bud marked $*$ is the root.

2. $C\mu \vdash \Phi \Rightarrow \Psi$ iff $C\mu \vdash \Phi, \Psi^\perp \Rightarrow$.
3. If $cC\mu \vdash \alpha \Rightarrow \beta$ and $cC\mu \vdash \varphi x \Rightarrow \psi x$ then $cC\mu \vdash \varphi\alpha \Rightarrow \psi\beta$.

As in the illfounded system $K\mu_\omega$, there is no explicit induction rule in $C\mu$. The rule is not derivable in $K\mu_\omega$ but certain forms admit straightforward co-proofs in the presence of cut which can be annotated to obtain cyclic proofs:

Proposition 5.6.

1. If $cC\mu \vdash \alpha \Rightarrow \varphi\alpha$ then $cC\mu \vdash \alpha \Rightarrow \nu\varphi$.
2. If $cC\mu \vdash \varphi\alpha \Rightarrow \alpha$ then $cC\mu \vdash \mu\varphi \Rightarrow \alpha$.
3. The induction rules (indL) and (indR) are derivable in $C\mu$.

Proof:

Figure 7 presents constructive derivations of the first two claims. The final claim is a consequence of the others. For example, for $\delta = \bigwedge(\Phi \cup \Psi^\perp)$, if $C\mu \vdash \Phi \Rightarrow \Psi$, $\varphi\delta$, then $C\mu \vdash \delta \Rightarrow \varphi\delta$ by the previous proposition, whereby $C\mu \vdash \delta \Rightarrow \nu\varphi$ by 1. Since $C\mu \vdash \Phi \Rightarrow \Psi, \delta$ we are done. \square

With the admissibility of the induction rule, it is straightforward to show that every finitary proof induces a cyclic proof of the same sequent. The constraints imposed on the bud–companion pairs opens the converse by allowing a syntactic translation of annotated sequents into plain sequents provable in $K\mu$. An obvious hurdle in the reduction is the transformation of a co-inductive conception of proof into a purely inductive form. The cyclic structure itself encodes a ‘global’ (co-)inductive argument only loosely connected to the predicates occurring within the proof. It is precisely the annotations of $C\mu$ which make the *implicit* induction of cycles *explicit* through the identification buds to companions and and relative order between nested bud–companion pairs.

The embedding of cyclic into finitary proofs is detailed in [4] for cut-free cyclic proofs in classical logic employing a slightly different notion of annotations and sequents. Following [4], the reduction proceeds by associating a syntactic translation to each vertex of a given cyclic proof that translates the annotated sequents into provable plain sequents. In most cases, the translation associated to a vertex is identical to the translation of its immediate successors. The rules which manipulate annotations and,

notably, the annotated (μ L) and (ν R) rules provide the cases in which the translation associated to the premise and conclusion differ.

We present the embedding in three stages. First, we isolate the general structure of the syntactic embeddings and identify conditions under which the rules of $cK\mu$ are admissible in $cK\mu$. These observations bootstrap the second step of a global embedding of constructive cyclic proofs in $cK\mu$. Finally, we lift the embedding to the full, classical, logics.

Let $*$: $N \rightarrow L_\mu$ be a partial function assigning names to (plain) formulas. We call such functions *assignments*. An assignment induces an interpretation of annotated formulas as plain formulas in the following way. Given a plain predicate $\sigma\varphi$ and annotation a we introduce a plain formula $\sigma^{a*}\varphi$ defined by $\sigma^{e*}\varphi = \sigma\varphi$ and

$$\sigma^{an*}\varphi := \begin{cases} \sigma^{a*}\varphi, & \text{if } n \notin \text{dom } *, \\ *n \circ \sigma^{a*}x. \varphi(*n \circ x), & \text{if } n \in \text{dom } *, \end{cases} \quad \text{where } \circ = \begin{cases} \wedge, & \text{if } \sigma = \nu, \\ \vee, & \text{if } \sigma = \mu. \end{cases}$$

Given an annotated formula α , the plain formula α^* is defined by structural recursion compositionally with $(\sigma^a\varphi)^* = \sigma^{a*}\varphi^*$ and $(\lambda x. \alpha)^* = \lambda x. \alpha^*$. In particular, for $a = n_1 \cdots n_k$,

$$\begin{aligned} \varphi(\nu^a\psi)^* &= \psi(\nu\psi) \text{ where } \psi = \lambda x. \varphi^*(*n_k \vee \cdots \vee *n_1 \vee x) \\ \varphi(\mu^a\psi)^* &= \psi(\mu\psi) \text{ where } \psi = \lambda x. \varphi^*(*n_k \wedge \cdots \wedge *n_1 \wedge x) \end{aligned}$$

Compositionality of the $*$ -translation implies

Lemma 5.7. For any assignment $*$, the $*$ -translation of each $cK\mu$ -rule except the annotated ν - and μ -rules is admissible in $cK\mu$.

The $*$ -translation of the annotated quantifier rules (ν^n R) and (μ^n L) is derivable in $cK\mu$ if $n \in \text{dom } *$ and $*n$ carries the side formulas:

Lemma 5.8.

1. If $cK\mu \vdash \Phi \Rightarrow \varphi(\nu^{an}\varphi)^*$ and $*n = \bigwedge \Phi$, then $cK\mu \vdash \Phi \Rightarrow (\nu^a\varphi)^*$.
2. If $cK\mu \vdash \varphi(\mu^{an}\varphi)^*, \Phi \Rightarrow \alpha$ and $*n = \bigvee(\Phi \cup \{\alpha\})$, then $cK\mu \vdash (\mu^a\varphi)^*, \Phi \Rightarrow \alpha$.

Proof:

We treat the second claim. From $cK\mu \vdash \varphi(\mu^{an}\varphi)^*, \Phi \Rightarrow \alpha$ we obtain $cK\mu \vdash \mu^{a*}x. \varphi^*(*n \wedge x), \Phi \Rightarrow \alpha$ by an application of (μ L) and weakening on the left. An application of strong induction (Proposition 3.6) yields $K\mu \vdash \mu^{a*}\varphi^*, \Phi \Rightarrow \alpha$. \square

We require a final observation before presenting the embedding of $cK\mu$ in $cK\mu$. Let $\Phi \leq \Psi$ express that $\Phi = \Psi^{\uparrow c}$ for some $c \in N^{<\omega}$.

Lemma 5.9. If $\Phi \leq \Psi$ are μ -annotated and $\alpha \leq \beta$ are ν -annotated, then for every assignment $*$, (plain) predicate φ and annotation a ,

1. if $cK\mu \vdash \Phi^* \Rightarrow \alpha^*$, then $cK\mu \vdash \Psi^* \Rightarrow \beta^*$.
2. if $*n = \bigwedge \Phi^*$ for some n in a , then $cK\mu \vdash \Psi^* \Rightarrow \nu^{a*}\varphi$.
3. if $*n = \bigvee(\Phi^{*\perp} \cup \{\alpha^*\})$ for some n in a , then $cK\mu \vdash \mu^{a*}\varphi, \Psi^* \Rightarrow \beta^*$.

Proof:

For the first claim, it suffices to show derivability of $\nu\varphi \Rightarrow \nu x \varphi(\gamma \vee x)$ and $\mu x \varphi(\gamma \wedge x) \Rightarrow \mu\varphi$ for arbitrary φ and γ , which is straightforward. For the second claim we show $cK\mu \vdash \Phi^* \Rightarrow \nu^{a*}\varphi$ given the assumption on $*$ and invoke part 1. That derivation is trivial as $\nu^{a*}\varphi = \gamma_1 \vee \dots \vee \gamma_k \vee \bigwedge \Phi^* \vee \gamma$ for some $\gamma_1, \dots, \gamma_k, \gamma$. The final claim is analogous. \square

From the above, we can complete the reduction of cyclic to finitary proofs.

Theorem 5.10. $cC\mu \vdash \Phi \Rightarrow \alpha$ implies $cK\mu \vdash \Phi \Rightarrow \alpha$.

Proof:

Let π be a constructive $C\mu$ -proof of $\Phi \Rightarrow \alpha$. Let $c_v : \Phi_v \Rightarrow \alpha_v$ be the annotated sequent labelling $v \in \pi$. To each $u \in \pi$ we associate an assignment $[u] : N \rightarrow L_\mu$. The domain of $[u]$ is the set of names in $c_u = n_0 \dots n_{k-1}$. The formula $[u]n_i$ is defined by recursion on $|u|$ and $i < k$ in the following way.

Let $cn \leq c_u$ and suppose $[u]n_i$ is defined for all $i < |c|$. There is a unique application of the rule $(\nu^n R)$ or $(\mu^n L)$ on the path from u to the root such that cn is a prefix of the control of all sequents on the path from u to the premise of the rule. Let v name this vertex. Suppose the rule in question is $(\nu^n R)$, meaning that the sequent at v has the form $c_v : \Phi_v \Rightarrow \nu^{an}\varphi$ and its premise is labelled by $c_v n : \Phi_v \Rightarrow \varphi(\nu^{an}\varphi)$. As $v < u$ we can assume that $[v]$ is defined on the names in c_v . In particular, $[v]$ is defined on the names in c because c is a subword of c_v . We then define $[u]n = (\bigwedge \Phi_v)^{[v]}$.

It remains to show that $K\mu \vdash (\Phi_u)^{[u]} \Rightarrow (\alpha_u)^{[u]}$ for every $u \in \pi$. As the control of a sequent is only modified via the quantifier rules (cw), if no applications of either rules occurs on a path from $u < v$ to v , then the assignments $[u]$ and $[v]$ are identical. Thus, by the observations above, it suffices to consider the case that u is the conclusion of those rules or is a bud. Suppose that u is the conclusion of an instance of $(\nu^n R)$ with $\beta = \nu^a\varphi$ the principal formula. Let v be the premise and $\Psi = \Phi_u = \Phi_v$. As Ψ does not contain n and $[v]m = [u]m$ for all $m \neq n$, the induction hypothesis yields

$$cK\mu \vdash \Psi^{[v]} \Rightarrow \varphi(\nu^{an}\varphi)^{[v]}.$$

Furthermore, $[v]n = \bigwedge \Psi^{[v]}$, whereby Lemma 5.8 implies

$$cK\mu \vdash \Psi^{[v]} \Rightarrow \beta^{[v]}.$$

However, this sequent is identical to the desired sequent $\Psi^{[u]} \Rightarrow \beta^{[u]}$. The case of (cw) is immediate, and that of buds is covered by Lemma 5.9. \square

Starting from a cyclic proof in $C\mu$, the same translation returns a finitary proof in $K\mu$. Combining the argument with Proposition 5.6 yields

Corollary 5.11. $C\mu \vdash \Phi \Rightarrow \Psi$ iff $K\mu \vdash \Phi \Rightarrow \Psi$.

5.2. From co-proofs to cyclic proofs

The remainder of this section addresses a partial interpretation of co-proofs as cyclic proofs which, in turn, establishes a partial completeness theorem for $\text{C}\mu$:

Theorem 5.12. Suppose $\Phi \Vdash \alpha$ and $\Phi \subseteq \Pi(\mathbb{D})$. Then $\text{cC}\mu \vdash \Phi \Rightarrow \alpha$. Moreover, the cyclic proof is cut-free.

The proof of the theorem appeals to the special properties of constructive co-proofs and $\Pi(\mathbb{D})$ -formulas. By definition, every branch of a co-proof witnessing $\Phi \Vdash \alpha$ carries at most one right-thread. If $\Phi \subseteq \mathbb{D}$ comprised only disjunctive formulas, then we can assume that every branch carries at most one left-thread for each formula in the context (Proposition 4.8). This constraint on threads clearly fails for a context with arbitrary $\Pi(\mathbb{D})$ -formulas. For the reduction to cyclic proofs, however, it suffices to observe that for every branch of the co-proof there is at most one *good* thread per formula in the conclusion. Proposition 2.11 guarantees this property for a context comprising $\Pi(\mathbb{D})$ -formulas.

Proof:

Let Φ be a finite sequence of $\Pi(\mathbb{D})$ formulas and $\pi \vdash \Phi \Rightarrow \alpha$ be a constructive co-proof. By Proposition 4.8, we may assume that no contraction in π has a principal disjunctive formula. We work with an extension of $\text{cC}\mu$ by two new quantifier rules which preserve the principal annotation without extending it:

$$\frac{c : \varphi(\mu^a \varphi), \Phi \Rightarrow \alpha}{c : \mu^a \varphi, \Phi \Rightarrow \alpha} \hat{\mu}^a \text{L} \qquad \frac{c : \Phi \Rightarrow \varphi(\nu^a \varphi)}{c : \Phi \Rightarrow \nu^a \varphi} \hat{\nu}^a \text{R}$$

The above inferences are derivable in $\text{cC}\mu$ via a combination of the annotated quantifier rules and the structural rules (aw) and (cw). For the argument below it is, however, convenient to assume the above rules are part of the calculus. Via these rules we can consider π as a $\text{C}\mu$ -derivation in which all controls and annotations are empty. We can, therefore, assume that π contains at least one branch.

The first step towards a $\text{C}\mu$ -proof is to annotate π by names in a way that annotations have length at most 1 and the rules $(\nu^n \text{R})$, $(\mu^n \text{L})$ and (cw) are applied whenever possible. For this purpose, we fix two disjoint infinite sets of names, referred to as the *primary* and *secondary* names respectively, and initially annotate π by primary names only, yielding a derivation π_* satisfying the following conditions:

1. for every instance of the rule $(\nu^n \text{R})$ or $(\mu^n \text{L})$ the name n is primary, the principal formula is annotated as $\nu^\varepsilon \varphi$ or $\mu^\varepsilon \varphi$ and the minor formula as $\varphi(\nu^n \varphi)$ or $\varphi(\mu^n \varphi)$,
2. all other instances of the νR - and μL -rule are *preserving*, namely instances of $(\hat{\nu}^n \text{R})$ or $(\hat{\mu}^n \text{L})$ for a primary name n already present in the control,
3. every vertex labelled by a sequent $c : \Phi \Rightarrow \alpha$ for which a name occurs in c but not in $\Phi \Rightarrow \alpha$ is the conclusion of an instance of (cw),
4. the result of removing all annotations from π_* and merging the premise and conclusion of each instance of (cw) is the derivation π .

Since comprises only $\Pi(\mathbb{D})$ formulas and has no contractions on disjunctive formulas, Proposition 2.11 implies that every primary name occurs in at most one formula of each sequent. That assumption alone does not preclude that sequents may grow unboundedly in π_* however. Let B be a branch of π_* . Since B is good, there is a primary name n which occurs in all but finitely many formulas on this thread and for which the rule $(\hat{\nu}^n R)$ or $(\hat{\mu}^n L)$ is applied infinitely often along B . As n can occur in an annotation of at most one (sub)formula in each sequent, the thread witnessing that B is good passes through all the principal formulas of instances of the inferences $(\hat{\nu}^n R)$ or $(\hat{\mu}^n L)$ along some suffix of B . Let $c_B = cn$ be the longest annotation which is a prefix of all controls on some suffix of B . This is well-defined as new names are only appended to the control. Given an annotated sequent \mathcal{S} along B and annotation c , set $\mathcal{S}^{\upharpoonright c}$ to be the annotated sequent that results after removing all names in \mathcal{S} not present in c .

Let $\mathcal{S}_u = (c_u : \Phi_u \Rightarrow \alpha_u)$ be the annotated sequent labelling $u \in \pi$. We call a pair $(u, v) \in B^2$ a *repetition* if

- $u < v$ and there is a modal premise on the path from u to v ,
- c_B is a prefix of c_w for every w such that $u \leq w \in B$,
- $\alpha_u^{\upharpoonright c_B} = \alpha_v^{\upharpoonright c_B}$ and $\Phi_u^{\upharpoonright c_B}$ is a subsequence of $\Phi_v^{\upharpoonright c_B}$ up to permutation,
- u and v are both conclusion of an instance of $(\hat{\nu}^n R)$ or $(\hat{\mu}^n L)$.

We refer to c_B as the *invariant* of the repetition (u, v) , and write it as c_v . Observe that if (u, v) and (u', v') are repetitions associated to branches B and B' , then $v = v'$ implies $c_B = c_{B'}$.

As π is a co-proof, there is an infinite subsequence $P: \omega \rightarrow B$ such that for every $i < j$, the pair $(P(i), P(j))$ is a repetition with invariant c_B . We thus obtain a finite $C\mu$ -derivation $\rho \subseteq \pi_*$ satisfying

1. ρ is an initial segment of π_* ,
2. every bud $l \in \rho$ is associated a vertex $l^c < l$, called the *companion* to l , such that (l^c, l) is a repetition in π_* ,
3. if $l^c \leq o^c < l$ for buds $l, o \in \rho$, then $c_l \leq c_o$.

Inserting instances of weakening and exchange at buds we may assume, additionally, that $\Phi_u^{\upharpoonright c_u} = \Phi_{u^c}^{\upharpoonright c_u}$ for each bud u . The derivation ρ can be obtained by recursion on length of invariants for paths in π_* . König's lemma ensures that the process of pruning π_* terminates in a finite derivation. Since c_l is a prefix of every control on the path from l^c to l , if $l^c \leq o^c < l$ then c_l and c_o are comparable: indeed, $c_l \leq c_o$ iff $|c_l| \leq |c_o|$.

What remains is to expand certain annotations in ρ by secondary names and replace applications of the $\hat{\nu}R$ - and $\hat{\mu}L$ -rules by appropriately chosen annotated versions to transform ρ into a $C\mu$ -proof with no buds. This operation is carried out by recursion through ρ . First, extend the definition of the invariant to non-buds: For $u \in \rho$ let c_u be the longest invariant c_v for a bud v such that $v^c < u \leq v$. In particular, c_u is the empty annotation if all leaves above u are axiomatic, and $c_{v^c} = c_v$ iff there is a companion $u^c < v^c < u$ with $c_u = c_v$.

We define a $C\mu$ -derivation ρ_u for each $u \in \rho$. A bud $v > u$ is said to be *open* if $v^c < u$. The buds of ρ_u are the sequents labelling the buds of ρ that are open above u , with each restricted to the invariant; the endsequent of ρ_u is the sequent at u similarly restricted:

$$\begin{array}{c} \{c_v : \Phi_v^{\vdash c_v} \Rightarrow \alpha_v^{\vdash c_v} \mid v^c < u \leq v\} \\ \vdots \\ \rho_u \\ \vdots \\ c_u : \Phi_u^{\vdash c_u} \Rightarrow \alpha_u^{\vdash c_u} \end{array}$$

For $u \in \rho$ a bud ρ_u is simply the trivial derivation $c_u : \Phi_u^{\vdash c_u} \Rightarrow \alpha_u^{\vdash c_u}$. Suppose ρ_v is defined for each $v \in \text{child}_\rho(u)$. We define ρ_u by case distinction on the inference with conclusion u . In the case of unary rules, u^+ names the unique child of u (in ρ).

Case (id). In this case $\Phi_u \cup \{\alpha_u^\perp\}$ contains an inconsistent pair of variables. As $c_u = \varepsilon$ is the empty annotation, ρ_u is the $C\mu$ -axiom

$$\frac{}{\varepsilon : \Phi_u^\varepsilon \Rightarrow \alpha_u^\varepsilon} \text{id}$$

Case ($\wedge L$). We have $c_{u^+} = c_u$, $\alpha_{u^+} = \alpha_u$, $\Phi_{u^+} = \Gamma, \Psi$ for some Ψ and $\Phi_u = \wedge \Gamma, \Psi$. In particular, $c_u = c_{u^+}$. Define ρ_u as

$$\frac{\rho_{u^+} \quad c_u : \Gamma^{\vdash c_u}, \Psi^{\vdash c_u} \Rightarrow \alpha_u^{\vdash c_u}}{c_u : \Phi_u^{\vdash c_u} \Rightarrow \alpha_u^{\vdash c_u}} \wedge L$$

Case ($\wedge R$) or ($\vee L$). For each $v \in \text{child}_\rho(u)$ we have a derivation ρ_v of the sequent $\lambda_\rho(v)$ restricted to c_v . Form ρ_u by an instance of ($\wedge R$) or ($\vee L$) with same principal formula and applications of (aw) and (cw) above the premises, as needed.

Case ($\nu^n R$) or ($\mu^n L$). As n is a primary name, u is not a companion and $c_u = c_{u^+}$. ρ_u is formed by attaching an instance of ($\nu^n R$) or ($\mu^n L$) at the root of ρ_{u^+} with no associated bud.

Case ($\hat{\nu}^a R$) or ($\hat{\mu}^a L$). Without loss of generality, we treat the rule ($\hat{\mu}^a L$). There are three subcases to contend:

1. u is not the companion of a bud. In this case, $c_u = c_{u^+}$. ρ_u concludes ρ_{u^+} by an application of (μL) as appropriate.
2. $u = v^c$ is the companion of a bud v in ρ and $c_u < c_{u^+}$. In particular, $c_{u^+} = c_u d n$ for some d and $\Phi_u = \nu^n \varphi, \Psi$. Let \mathcal{O} be buds of ρ which are open above u . By construction, ρ_{u^+} is a proof in $cC\mu$ of the sequent

$$c_u d n : \varphi(\mu^n \varphi)^{\vdash c_u d n}, \Psi^{\vdash c_u d} \Rightarrow \alpha_u^{\vdash c_u d}$$

from the assumption $c_u d n : \nu \varphi^{\vdash c_u d n}, \Psi^{\vdash c_u d n} \Rightarrow \alpha_u^{\vdash c_u d n}$ and the sequents labelling the buds in \mathcal{O} :

$$\rho_{u^+} = \left\{ \begin{array}{c} c_u d n : (\mu^n \varphi)^{\vdash c_u d n}, \Phi^{\vdash c_u d} \Rightarrow \alpha_u^{\vdash c_u d} \quad \{(\mathcal{S}_o)^{\vdash c_o} \mid o \in \mathcal{O}\} \\ \vdots \\ c_u d n : \varphi(\mu^n \varphi)^{\vdash c_u d n}, \Psi^{\vdash c_u d} \Rightarrow \alpha_u^{\vdash c_u d} \end{array} \right.$$

The derivation ρ_u is formed by removing the names in d from ρ_{u+} before inserting an instance of $(\nu^n R)$ at the root:

$$\rho_u = \left\{ \begin{array}{c} \frac{c_u n : (\mu^n \varphi)^{\upharpoonright c_u n}, \Phi^{\upharpoonright c_u} \Rightarrow \alpha_u^{\upharpoonright c_u} \quad \{(\mathcal{S}_o)^{\upharpoonright c_o} \mid o \in \mathcal{O}\}}{\vdots (\rho_{u+})^{\upharpoonright c_u n}} \\ \frac{c_u n : \varphi(\mu^n \varphi)^{\upharpoonright c_u n}, \Psi^{\upharpoonright c_u} \Rightarrow \alpha_u^{\upharpoonright c_u}}{c_u : \mu \varphi^{\upharpoonright c_u}, \Psi^{\upharpoonright c_u} \Rightarrow \alpha_u^{\upharpoonright c_u}} \mu^n L \end{array} \right.$$

3. $u = v^c$ is the companion of a bud in ρ and $c_u = c_{u+}$. In particular, $c_u = c_v$. Let $c_u = cn$ and $\Phi_u = \nu^n \varphi, \Psi$. In this case there are no names in the control that need removal but it is necessary to maintain the name n in the conclusion of ρ_u . This is achieved by introducing a secondary name m that does not occur in ρ_u and first forming the derivation in which all occurrences of the name n are replaced by nm (working from the root and stopping once a vertex is reached which does not contain n). Let $c^{(m)}$ express the result of effecting this change to c ; likewise $\alpha^{(m)}$ and $\Phi^{(m)}$. Let \mathcal{O} be the open buds of ρ above u . The derivation ρ_u is formed by using m as the principal name for discharging the associated bud:

$$\rho_u = \left\{ \begin{array}{c} \frac{c_u m : (\mu^{nm} \varphi)^{\upharpoonright c_u m}, \Phi^{\upharpoonright c_u} \Rightarrow \alpha^{\upharpoonright c_u} \quad \left\{ \frac{(\mathcal{S}_o)^{\upharpoonright c_o}}{(\mathcal{S}_o)^{(m)} aw + cw} \mid o \in \mathcal{O} \right\}}{\vdots} \\ \frac{c_u m : \varphi(\mu^{nm} \varphi)^{\upharpoonright c_u m}, \Phi^{\upharpoonright c_u} \Rightarrow \alpha^{\upharpoonright c_u}}{c_u : (\mu^n \varphi)^{\upharpoonright c_u}, \Phi^{\upharpoonright c_u} \Rightarrow \alpha^{\upharpoonright c_u}} \mu^m L \end{array} \right.$$

Other cases. The other rules are straightforward.

By construction, ρ_ε is a $C\mu$ -derivation with endsequent $\varepsilon : \Phi \Rightarrow \alpha$ and no buds, that is, ρ_ε is a $cC\mu$ -proof. \square

Combining theorem 5.12 with our partial completeness result for $cK\mu_\omega$, yields:

Corollary 5.13. If α is valid and conjunctive, then $cC\mu \vdash \Rightarrow \alpha$.

6. Disjunctive Normal Form

We now tend to the task of showing that the disjunctive fragment of L_μ is expressively adequate for constructive modal μ -calculus, namely that every formula can be associated an equivalent disjunctive formula and the equivalence holds in $cK\mu_\omega$. As already mentioned, the proof of this result uses a proof-search procedure for annotated sequents. Due to the restrictive form of proofs in $C\mu$, rarely can we expect the proof-search to actually isolate a cyclic proof. Rather, annotations assist in bounding the search space relative to potential *illfounded* proofs.

Theorem 6.1. (Disjunctive normal form theorem)

For every formula α there exists a guarded disjunctive formula $\gamma \dashv\vdash \alpha$. In particular, $\llbracket \gamma \rrbracket = \llbracket \alpha \rrbracket$.

The associated disjunctive formula γ is designed to encode all information about how the formula α can be utilised in a co-proof of a sequent $\alpha \Rightarrow \beta$ independent of the choice of β . The construction of γ proceeds in two steps. First, the role of α in unspecified co-proof is expressed as a canonically chosen derivation in a modification of the cyclic calculus $C\mu$ that merges the proof-search calculus of Theorem 4.9 and the annotation method of Theorem 5.12. This derivation will be called the *template* for α . Second, we show how the template can be encoded as a (disjunctive) formula of μ -calculus which is provably equivalent to the initial formula α . Essentially the same two-step construction is employed in [8] for establishing uniform Lyndon interpolants for the modal μ -calculus. Indeed, we are able to recapture that result as a corollary of the construction.

The present section is dedicated to the proof of theorem 6.1. With the theorem in place we obtain as a corollary completeness of the illfounded proof system:

Corollary 6.2. If α is valid then $\alpha^\perp \Vdash \perp$ and $K\mu_\omega \vdash \Rightarrow \alpha$.

Proof:

Let α be valid and let γ be a disjunctive formula $\gamma \dashv\vdash \alpha^\perp$, entailed by theorem 6.1. Since $\llbracket \alpha \rrbracket \subseteq \llbracket \gamma^\perp \rrbracket$ (Theorem 3.11), γ^\perp is a valid conjunctive formula, whence $\gamma \Vdash \perp$ by Theorem 4.9. Transitivity (Theorem 4.11) implies $\alpha^\perp \Vdash \perp$ and, therefore, $K\mu_\omega \vdash \Rightarrow \alpha$. \square

The proof of theorem 6.1 is split into three parts. First, we present the construction of the disjunctive formula. Following that, we establish each of the two provability claims.

6.1. Construction

Without loss of generality (Theorem 4.16) we may assume that α is guarded. The construction utilises the notion of annotated formulas and sequents from the previous section. To that base, we add a few more concepts.

Definition 6.3. A name n is *covered* in an annotated sequent $c : \Phi \Rightarrow \Psi$ if it occurs in Φ or Ψ and for every $\sigma^a \varphi \in \text{Sub}(\Phi, \Psi)$, if $a_i = n$ for some i then $|a| > i + 1$.

Fix $c \in N^{<\omega}$. The subsequences of c form a finite tree $S(c)$ relative to the prefix relation. This tree is equipped with an ordering on non-comparable vertices, defined as a is *left of* b if there exists $k < |a|$ such that

1. $a(i) = b(i)$ for all $i < k$, and
2. $a(k) = c(k_0)$ and $b(k) = c(k_1)$ for some $k_0 < k_1 < |c|$.

The *Kleene–Brouwer ordering* on $S(c)$ is the relation \leq_c defined as $a \leq_c b$ iff b is a prefix of a , or a is to the left of b . Observe that ε is the maximal element of \leq_c for every $c \in N^{<\omega}$.

Proposition 6.4. If c is an annotation then \leq_c well-orders the subsequences of c .

Proposition 6.5. If $c \leq d$ then $\leq_c \subseteq \leq_d$.

The orderings \leq_c will be shortly lifted to an ordering on annotated formulas that will be used to constrain the size of sequents in the template. For this purpose it is convenient to focus attention on the particular annotated formulas that can occur in derivations. These are formulas for which the annotated subformulas are linearly ordered by substitution. We call such formulas *linearly annotated*:

Definition 6.6. An annotated formula α is linearly σ -annotated if there exist a plain formula β and predicates $\varphi_1, \dots, \varphi_n$ and non-trivial annotations a_1, \dots, a_n such that

1. $\alpha = (\lambda x_n \dots (\lambda x_1 \beta)(\sigma^{a_1} \varphi_1) \dots)(\sigma^{a_n} \varphi_n) = \beta[x_1 \mapsto \sigma^{a_1} \varphi_1] \dots [x_n \mapsto \sigma^{a_n} \varphi_n]$,
2. x_1 is free in β and x_{i+1} is free in $\sigma \varphi_i$ for $i \in [1, n]$,
3. a name occurs in both a_i and a_j iff $i = j$.

Given a plain formula α and annotation c , let $L_\sigma(\alpha, c)$ be the set of linear σ -annotations of α in which all annotations are subsequences of c .

The second condition of the definition implies that a decomposition of α in the form of 1 is unique, if it exists. The third condition states that each name occurs in a unique subformula. This assumption assists with keeping track of names during the construction of the template. By the first two conditions, it follows that x_i^\perp does not occur in $\beta, \varphi_1, \dots, \varphi_{i-1}$.

Given linearly σ -annotated formulas α and β with respective decompositions

$$\begin{aligned}\alpha &= \alpha_0[x_1 \mapsto \sigma^{a_1} \varphi_1] \dots [x_n \mapsto \sigma^{a_n} \varphi_n] \\ \beta &= \beta_0[x_1 \mapsto \sigma^{b_1} \psi_1] \dots [x_m \mapsto \sigma^{b_m} \psi_m]\end{aligned}$$

we set $\alpha <_c \beta$ if

1. $\alpha_0 = \beta_0$, $m = n$ and $\varphi_i = \psi_i$ for each $i \in [1, m]$, and
2. there exists $k \in [1, m]$ such that $a_k <_c b_k$ and $a_i = b_i$ for all $i \in [k+1, m]$.

The weak order is defined as $\alpha \leq_c \beta$ iff $\alpha <_c \beta$ or $\alpha = \beta$. In particular, if $\alpha \leq_c \beta$ then $\alpha^- = \beta^-$.

Example 6.7. $\nu^a x(\Diamond x \vee \nu^b y. \Box y)$ is linearly ν -annotated for all $a, b \in A$. The formula $\mu^a x(\Diamond x \vee \mu^b y. \Box(x \wedge y))$ is linearly μ -annotated iff $b = \varepsilon$ because this is the only case in which the formula can be expressed as a linear decomposition following Definition 6.6. Similarly, the formula $\bigwedge \{\nu^a x. \Box x, \nu^b y \Diamond y\}$ is linearly annotated iff at least one of a and b is trivial.

The ordering $<_c$ compares linearly ordered formulas with the comparable decompositions by a lexicographic ordering on the annotations. Supposing that φ, ψ are plain predicates and φ contains x free, $(\nu^a \varphi)[\nu^b \psi] <_c (\nu^{a'} \varphi)[\nu^{b'} \psi]$ iff $b <_c b'$ or $b = b'$ and $a <_c a'$.

As \leq_c well-orders the subsequences of c its lifting to $L_\sigma(\alpha, c)$ is both a partial order and a well-quasi order. Indeed, a bound on the number of pairwise \leq_c -incomparable elements set in $L_\sigma(\alpha, c)$ depends only on α (and not c).

$$\begin{array}{c}
\frac{c : \Phi, \alpha \Rightarrow \quad \alpha \leq_c \alpha'}{c : \Phi, \alpha, \alpha' \Rightarrow} \text{dup} \qquad \frac{c : \Phi^{\downarrow n} \Rightarrow \quad n \text{ covered in } \Phi}{c : \Phi \Rightarrow} \text{cvr}^n \\
\\
\frac{c : \Gamma, \Phi \Rightarrow \Psi}{c : \bigwedge \Gamma, \Phi \Rightarrow \Psi} \wedge^* \text{L} \qquad (L \subseteq \text{Lit}) \frac{\{c : \alpha, \Phi \Rightarrow \mid \alpha \in \Psi\} \quad c : \Phi \Rightarrow}{c : \Diamond \Psi, \Box \Phi, L \Rightarrow} \Diamond^* \text{L}
\end{array}$$

Figure 8. Structural and logical rules specific to the template. $\Phi^{\downarrow n}$ is the restriction of each annotation in Φ which contains n to the shortest prefix containing n .

Proposition 6.8. Let α be a plain formula. There exists $N < \omega$ such that for all $c \in A$ every set of \leq_c -incomparable elements of $L(\alpha, c)$ has cardinality bounded by N .

We now proceed with the definition of the template. The template of α is a certain finite derivation in a variant of $C\mu$ adjusted for proof-search of the sequent $\alpha \Rightarrow$. As the calculus employed is cut-free, all formulas will be linear μ -annotations of subformulas of α . The additional rules of the calculus are presented in fig. 8. We use (cvr) to denote the union of (cvr^n) for all $n \in \mathbb{N}$. The structural rules (dup) and (cvr) provide special instances of formula weakening and annotation weakening respectively that can be permitted in the context of proof-search.

Definition 6.9. Fix an enumeration of $\text{Sub}(\alpha)$ and an enumeration n_0, n_1, \dots of \mathbb{N} . The *pre-template* of α is the maximal cut-free derivation with root $\varepsilon : \alpha$ in the rules of $C\mu$ where the rules (aw) , (cut) , $(\Diamond \text{L})$ and $(\wedge \text{L})$ are replaced by the four rules in fig. 8 which satisfies the following for every vertex u :

1. If the sequent at u is the conclusion of an instance of (cw) , (dup) or (cvr) then one of these rules labels u with that priority (so the rule is (cw) if possible, otherwise (dup) , or (cvr) if neither other rule is applicable).
2. If the rule at u is a non-modal logical rule, then the principal formula is the first formula in the sequent according to the enumeration of $\text{Sub}(\alpha)$.
3. If the rule at u is $(\mu^n \text{L})$ then all names of lower index occur in the sequent at u , i.e., $n = n_i$ and n_j occurs in u for all $j < i$.

Existence of a pre-template is immediate, and the three constraints of Definition 6.9 ensure its uniqueness. As the pre-template is required to be maximal, leaves are instances of either (idL) or $(\vee \text{L})$ with principal formula $\bigvee \emptyset$. All other sequents are conclusions of a rule in the calculus. Also of note is that a sequent of the form $\varepsilon : \emptyset \Rightarrow$ may occur in the pre-template (possibly obtained as a premise of $(\Diamond^* \text{L})$ where $\Phi = \emptyset$) in which case this vertex lies on a single branch whose suffix is the sequence $(\varepsilon : \emptyset \Rightarrow, (\Diamond^* \text{L}))^\omega$. The next observation makes use of the fact that only linearly annotated formulas occur in a pre-template.

Proposition 6.10. The pre-template is a regular tree. That is, there are only finitely many maximal subtrees of the pre-template of a given formula α .

Proof:

Let P be the pre-template of α . It suffices to show that every branch of P contains a repeated annotated sequent as regularity then follows from the uniqueness of the pre-template. By induction through P we see that every formula is linearly μ -annotated (note that $\varphi(\mu^a \varphi)$ is linearly annotated if $\mu^a \varphi$ is). Proposition 6.8 and the priority in applications of (dup) in P ensure that each annotated sequent in the branch has bounded length. Coupled with the priority on applications of (cvr) there is, therefore, a bound on the length of all controls, whence a bound on the number of distinct annotated sequents in the branch. Hence, every branch contains a repeated annotated sequent. \square

A consequence of regularity is that the pre-template for α can be identified with a finite initial segment of itself. For our purposes it is convenient to isolate a particular initial segment of P , henceforth called the *template*, in which information about the eliminated parts of the tree is directly expressed in the buds. Identifying the template is similar to isolating the basis of the cyclic proof in Theorem 5.12.

A *repetition* is a pair $(u, v) \in P \times P$ of vertices labelled by the same annotated sequent such that $u < v$. As α is guarded, if (u, v) is a repetition then there is a vertex $w \in [u, v)$ such that $\mathcal{R}_P(w) = (\diamond^* L)$. As observed, every branch of the pre-template contains infinitely many repetitions. For the present construction, the *invariant* of a repetition (u, v) will be a pair $(c, k) \in A \times \{0, 1\}$ where c is a non-trivial prefix of all controls on the path interval (u, v) and k records whether the final name in c is covered in one of these sequents (in which case $k = 0$). If there is no such control, then c is the longest prefix of all controls and $k = 1$. Thus, the invariant of a repetition (u, v) is the pair $\text{inv}(u, v) = (c, k) \in A \times \{0, 1\}$ where:

- c is the longest annotation that is a prefix of every vertex in the path segment $[u, v]$ and is such that if $\mathcal{R}_P(w) = (\text{cvr}^{c(i)})$ for some $w \in [u, v)$ then $i = |c| - 1$,
- $k = 0$ iff $\mathcal{R}_P(w) = (\text{cvr}^{c(i)})$ for some $w \in [u, v)$ and $i < |c|$.

Invariants are compared lexicographically: $(c, k) \sqsubseteq (d, l)$ iff $c < d$ or $c = d$ and $k \leq l$. In particular, $(c, 1) \sqsubseteq (d, 0)$ iff c is a proper prefix of d . The relation \sqsubseteq forms a partial order on invariants.

Assume a companion function on P is available, i.e., a partial function $u \mapsto u^c$ such that if u^c is defined then (u^c, u) is a repetition. If u^c is defined we write $\text{inv}(u)$ for $\text{inv}(u^c, u)$. The companion function induces a more liberal ‘reachability’ relation \rightsquigarrow on P by $u \rightsquigarrow v$ iff u^c and v^c are both defined and $u^c \leq v^c < u$. Note, \rightsquigarrow is not transitive in general. However, if $u \rightsquigarrow v$ then $\text{inv}(u)$ and $\text{inv}(v)$ are comparable: $\text{inv}(u) \sqsubseteq \text{inv}(v)$ or $\text{inv}(v) \sqsubseteq \text{inv}(u)$. The template is given by a choice of companion function such that the two orderings are in agreement:

Definition 6.11. A *template* of α is a finite initial subderivation $T \subseteq P$ and a companion function $u \mapsto u^c$ defined on the buds of T such that if $u \rightsquigarrow v$ for buds u, v then $\text{inv}(u) \sqsubseteq \text{inv}(v)$.

Existence of a template T of α is straightforward and can be obtained by isolating the buds of T by recursion through the finitely many invariants in P . Any template implicitly witnesses regularity of the pre-template in the sense that for every bud $u \in T$ the maximal sub-derivations of P rooted at u and u^c are identical.

In the following we assume a fixed template of α , denoted T . A bud u is *successful* if its invariant $\text{inv}(u) = (c, 0)$ for some c ; otherwise u is *unsuccessful*. If u and v are buds with the same companion,

$$\begin{array}{c}
\frac{n : \Box\beta, \Diamond\mu^n\varphi \Rightarrow}{n : \Box\beta, \Diamond\beta, \Diamond\mu^n\varphi \Rightarrow} \text{dup} \\
\frac{n : \Box\beta, \varphi(\mu^n\varphi) \Rightarrow}{nm : \Box\beta, \varphi(\mu^{nm}\varphi) \Rightarrow} \text{cvr} \\
\frac{n : \Box\beta, \mu^n\varphi \Rightarrow}{n : \beta, \mu^n\varphi \Rightarrow} \nu \\
\frac{n : \Box\beta, \Diamond\mu^n\varphi \Rightarrow}{n : \Box\beta \wedge \Diamond\mu^n\varphi \Rightarrow} \wedge \\
\frac{n : \Box\beta \wedge \Diamond\mu^n\varphi \Rightarrow}{\varepsilon : \mu\varphi \Rightarrow} \nu^n
\end{array}
\quad
\begin{array}{c}
\varepsilon : \emptyset \Rightarrow \Box^* \\
\varepsilon : \beta \Rightarrow \Box^* \\
\varepsilon : \Box\beta \Rightarrow \nu \\
\varepsilon : \beta \Rightarrow \text{cw} \\
n : \beta \Rightarrow \Box^*
\end{array}
\quad
\begin{array}{c}
z \\
z \\
z \\
z \\
z \\
z \\
z \\
z \\
z \\
\mu z \nabla(\{z\}, \{z, \nu\psi\}) \\
\mu z \nabla(\{z\}, \{z, \nu\psi\}) \\
\mu z \nabla(\{z\}, \{z, \nu\psi\})
\end{array}
\quad
\begin{array}{c}
x \\
y \nabla(\emptyset, \{x\}) \\
y \nabla(\{y\}, \{y, \gamma\}) \\
\nu\psi \\
\nu\psi \\
\mu z \nabla(\{z\}, \{z, \nu\psi\}) \\
\mu z \nabla(\{z\}, \{z, \nu\psi\}) \\
\mu z \nabla(\{z\}, \{z, \nu\psi\})
\end{array}$$

Figure 9. An example of a template (left) for the formula $\mu\varphi$ where $\varphi = \lambda x. \Box(\nu x \Diamond x) \wedge \Diamond x$ and $\beta = \nu x \Diamond x$. On the right is the corresponding disjunctive formula associated to each vertex of the template, with γ and ψ abbreviating $\mu x \nabla(\emptyset, \{x\})$ and $\lambda y. \nabla(\{y\}, \{y, \gamma\})$ respectively.

then $\text{inv}(u) = \text{inv}(v)$, so u is successful iff v is successful. Thus we can unambiguously refer to a companion u^c as $(un)successful$ if u is $(un)successful$.

We now proceed with the definition of the disjunctive formula. As discussed, this formula simply encodes the pre-template of α via the chosen template. We assume a distinct variable x_u is available for each vertex $u \in T$, let $\mathcal{R}_P(u)$ be the rule labelling u and let $c_u : \Phi_u \Rightarrow$ be the sequent labelling $u \in T$. The construction of $\hat{\alpha}$ proceeds by associating to each $u \in T$ disjunctive formulas α_u^+ and α_u^- in the following way. An example of the association is given in fig. 9.

Definition 6.12. The formulas α_u^+ and α_u^- are defined by recursion through $u \in T$. Assuming that α_u^- is defined, set

$$\alpha_u^+ := \begin{cases} \mu x_u \alpha_u^-, & \text{if } u \text{ is successful,} \\ \nu x_u \alpha_u^-, & \text{otherwise.} \end{cases}$$

The formula α_u^- is built from the formulas $\{\alpha_v^+ \mid v \in \text{child}_T(u)\}$ according to the rule at $u \in T$:

- If u is a bud, $\alpha_u^- = x_{u^c}$.
- If $\mathcal{R}_P(u) = (\bigvee L)$, then $\alpha_u^- = \bigvee \{\alpha_v^+ \mid v \in \text{child}_t(u)\}$.
- If $\mathcal{R}_P(u) = (\Diamond^* L)$ and Φ_u contains both x, x^\perp for some variable x then $\alpha_u^- = \perp$.
- If $\mathcal{R}_P(u) = (\Diamond^* L)$ and Φ_u is not of the form above. Let $\Phi_u = \Diamond\delta_0, \dots, \Diamond\delta_{n-1}, \Box\Psi, L$ where $L \subseteq \text{Lit}$ and let $\text{child}_T(u) = \{u_1, \dots, u_n\}$ be such that $\Phi_{u_i} = \delta_i, \Psi$ for each $i < n$ and $\Phi_{u_n} = \Psi$. Set

$$\alpha_u^- = \nabla_L(\{\alpha_{u_0}^+, \dots, \alpha_{u_{n-1}}^+\}, \{\alpha_{u_0}^+, \dots, \alpha_{u_n}^+\}).$$

- In all other cases u has exactly one child u^+ , and $\alpha_u^- = \alpha_{u^+}^+$.

Definition 6.13. The *disjunctive companion* to a guarded formula α is the formula $\hat{\alpha} := \alpha_\varepsilon^+$ associated to the root of the template as per the previous definition. More generally, the disjunctive companion to a vertex $u \in T$ is the formula

$$\hat{\alpha}_u := \alpha_u^+[x_{u_n} \mapsto \alpha_{u_n}^+] \cdots [x_{u_1} \mapsto \alpha_{u_1}^+]$$

where $\varepsilon = u_1 < \cdots < u_n < u$ enumerate the ancestors of u . The disjunctive companion to a non-guarded formula is the disjunctive companion to the canonical guarded equivalent formula determined by proposition 4.16.

We begin with a few useful observations.

Proposition 6.14. For $\hat{\alpha}_u$ defined as above,

1. $\hat{\alpha}_u$ is a guarded disjunctive formula.
2. If $u \in T$ is axiomatic then $K\bar{\mu} \vdash \hat{\alpha}_u$.
3. If $u \rightsquigarrow v$ then $\hat{\alpha}_u \triangleright^* \hat{\alpha}_v$ for a thread in which $\hat{\alpha}_v$ is the shortest occurring formula.

6.2. Validation

With $\hat{\alpha}$ as the disjunctive companion to α , we now establish the two properties of Theorem 6.1: $\alpha \Vdash \hat{\alpha}$ and $\hat{\alpha} \Vdash \alpha$. The two witnessing co-proofs are similar in construction. As the former is a more involved argument, we present it in detail and sketch the verification of the other claim. Both claims are implicit in the main result of [8].

Both constructive co-proofs are obtained by modifying the pre-template for α and inserting the corresponding disjunctive formula $\hat{\alpha}_u$ at each vertex. To confirm that each branch is good it is necessary to project branches of the pre-template P into the template T . Regularity of the pre-template means that every $u \in P \setminus T$ has the form $u = vw$ where v is a bud of T and $v^c w \in P$. Thus, the companion function of T induces a canonical projection $p: P \rightarrow T: u \mapsto u^p$ of P into T defined as: if $u = vw$ for a bud $v \in T$ then $u^p = (v^c w)^p$; otherwise $u^p = u$. Note that u^p is never a bud of T .

We appeal to the following observation relating invariants to the existence of threads in the pre-template. Recall that a non-bud vertex is successful iff it is the companion to a successful bud.

Proposition 6.15. Let B be a branch of the pre-template of α and let $u_B \in T$ be the \leq -minimal vertex such that $u_B = B(i)^p$ for infinitely many i . Then B is good (in the sense of the unannotated calculus where (dup) is interpreted as weakening) iff u_B is a successful companion.

For the proof of the proposition it is helpful to fix invariants to companions. Since buds with the same companion are associated identical invariants we can, without confusion, refer to the invariant of a companion.

Proof:

Let B^p be the vertices in T that are projection of infinitely many vertices in B . Let C_B be the set of companions in B^p . There is a unique \leq -minimal element of B^p and this is also an element of C_B . Let

u_B be this vertex and (c, k) the associated invariant. By the choice of T , for every $v \in C_B$ we have $\text{inv}(u_B) \sqsubseteq \text{inv}(v)$. Let c_B be the longest prefix of all controls for vertices in B^p . In particular, $c \leq c_B$. Henceforth, we treat only the template and pre-template rooted at u_B and assume $B(0) = u_B$.

Suppose B is good and, seeking a contradiction, that $k = 1$. So $c_B = c$. Consider the plain formula $\mu\varphi$ of lowest complexity for which annotations occur infinitely often in the thread and let τ be the suffix of the thread in which all formulas have complexity at least $|\mu\varphi|$. Let $(\mu^{a_i}\varphi_i)_i$ enumerate the annotations of $\mu\varphi$ through which τ passes. Linearity of the annotations means that $\varphi_i = \varphi_j$ for all i, j . Let $\psi = \varphi_0$. Let a be the longest common prefix among the a_i . Consider the vertex at which $\mu^{a_i}\psi$ is principal. The minor formula of this rule is the annotated formula $\psi(\mu^{a_i n}\psi)$ for some name n that is not in the control of the conclusion. As $\mu^{a_{i+1}}\psi$ is a descendent of $\psi(\mu^{a_i n}\psi)$ in the pre-template P , we have that either $a_i < a_{i+1}$ or there is a sequent on the path between the two formulas at which a name in a_i is covered. Given that there are finitely many annotations in the template, there is a shortest annotation a such that $a = a_i$ for infinitely many i and the last name in a , say m , must be covered in some sequent in B^p . Since a is the prefix of an annotation that occurs in every sequent in B^p it follows that a is a subsequence of c_B . But then the invariant of the repetition (v^c, v) that surrounds the instance of (cvr^m) must have invariant $(c, 0)$, contradicting that $\text{inv}(u_B) \sqsubseteq \text{inv}(v)$.

For the converse direction, suppose $\text{inv}(u_B) = (c, 0)$. Let n be the final name in c and consider the tree of annotated formula traces through B that follow only formulas in which n occurs. By the choice of u_B , n occurs in all controls of B^p , whence this tree of traces is infinite. By weak König's lemma there is therefore an infinite such annotated formula trace. Call this t . From the assumption that u_B is successful, some instance of (cvr^n) is applied infinitely often in B . For this to be the case, t must infinitely often pass through a formula in which n is covered, whence t must pass through a formula $\mu^{an}\varphi$ followed by $\varphi(\mu^{anm}\varphi)$ infinitely often for some a and m . Linearity of the annotations enforces that it is specifically the formula $\mu^{an}\varphi$ which witnesses that t supports a left μ -thread. \square

Proposition 6.16. $\alpha \Vdash \hat{\alpha}$.

Proof:

The desired co-proof is obtained from the template of α by inserting $\hat{\alpha}_u$ as the consequent of the sequent at $u \in T$ and inserting new inferences where appropriate. The definition of $\hat{\alpha}_u$ entails that buds continue to have the same sequent as their companion, and identifying buds with companions yields a maximal $K\mu_\omega$ -derivation π . That every branch of π is good is guaranteed by the choice of $\hat{\alpha}$: If a branch of π does not carry a left μ -thread then, by virtue of proposition 6.15, the unique right thread carried by the branch will be ν . \square

The converse provability claim, presented next, employs a similar argument but is complicated by the disjunctive interpretation of modalities. For $u \in T$, let α_u^+ , α_u^- be as in definition 6.12. For each $u \in T$, let $\hat{\alpha}_u^-$ be the subformula of $\hat{\alpha}_u$ corresponding to the choice of α_u^- , namely the following formula where $\varepsilon = u_1 < \dots < u_n < u$ is an enumeration of the ancestors of u :

$$\begin{aligned} \hat{\alpha}_u^- &:= \alpha_u^- [x_u \mapsto \alpha_u^+] [x_{u_n} \mapsto \alpha_{u_n}^+] \cdots [x_{u_1} \mapsto \alpha_{u_1}^+] \\ &= \alpha_u^- [x_{u_1} \mapsto \hat{\alpha}_{u_1}] \cdots [x_{u_n} \mapsto \hat{\alpha}_{u_n}] [x_u \mapsto \hat{\alpha}_u] \quad \text{by definition of the } \hat{\alpha}_v. \\ \hat{\alpha}_u &= \sigma_u x_u \hat{\alpha}_u^- \end{aligned}$$

Proposition 6.17. $\hat{\alpha} \Vdash \alpha$.

Proof:

Let P be the pre-template of α with $c_u : \Phi_u \Rightarrow$ being the sequent at $u \in P$. The witnessing co-proof is constructed in two steps. First, we co-recursively construct a co-proof of $\alpha \Rightarrow \alpha$ maintaining the invariant that each bud of the partially determined derivation is a sequent $(\Phi_u)^- \Rightarrow \gamma$ for an associated $u \in P$ satisfying $\gamma \in (\Phi_u)^-$. In place of the usual right modal rule $(\Box R)$, however, is an amalgamation of $(\Box R)$ and the branching rule $(\Diamond^* L)$:

$$\frac{\beta_1, \Phi \Rightarrow \delta \quad \dots \quad \beta_n, \Phi \Rightarrow \delta \quad \Phi \Rightarrow \delta}{\Diamond \beta_1, \dots, \Diamond \beta_n, \Box \Phi, L \Rightarrow \Box \delta} \Box^* R$$

The rule is sound because each premise is a weakening of the final one. The motivation for the rule comes from the second step in the process, wherein each antecedent $(\Phi_u)^-$ is replaced by the corresponding subformula of $\hat{\alpha}$, with the modal case necessitating the new rule above.

The co-proof in the first part is easily to construct. At any bud whose sequent satisfies the desired invariant, the derivation is continued by copying the left rules from the pre-template P until either the rule $(\Diamond^* L)$ is reached or the next rule in P will destroy the invariant. We first treat the case of a non-modal rule in P which destroys the invariant. This means that the construction of the desired co-proof has reached a sequent $(\Phi_u)^- \Rightarrow \gamma$ where γ is the principle formula at u . In all scenarios we can continue by applying the left rule from P followed by the corresponding right rule. For example, if $\gamma = \bigvee \{\gamma_1, \dots, \gamma_n\}$ and $(\Phi_u)^- = \gamma, \Psi$, the following derivation is inserted.

$$\frac{\frac{\gamma_1, \Psi \Rightarrow \gamma_1}{\gamma_1, \Psi \Rightarrow \gamma} \bigvee R \quad \dots \quad \frac{\gamma_n, \Psi \Rightarrow \gamma_n}{\gamma_n, \Psi \Rightarrow \gamma} \bigvee R}{\gamma, \Psi \Rightarrow \gamma} \bigvee L$$

Now suppose $\mathcal{R}_P(u) = \Diamond^* L$ and that $(\Phi_u)^- = \Diamond \Psi, \Box \Phi, L$ where L is a set of literals. If L is inconsistent or $\gamma \in L$, the sequent is initial, and if $\gamma \in \Diamond \Psi$ the rule $(\Diamond L)$ is inserted with the corresponding successor of u . For $\gamma = \Box \delta \in \Box \Phi$ we insert the special modal rule $(\Box^* R)$ presented above with $\Psi = \beta_1, \dots, \beta_n$. The buds of the extended derivation each satisfy the invariant and correspond to the immediate successors of u in P .

The second step of the argument is now simply the matter of replacing each antecedent $(\Phi_u)^-$ by the formula $\hat{\alpha}_{p(u)}$ where $p : P \rightarrow T$ is the projection of the pre-template P into the template from which $\hat{\alpha}$ is constructed. This step requires inserting instances of the left rules (μL) and (νL) to accommodate the quantifiers in $\hat{\alpha}$. But the logical steps transfer directly, with instances of the new rule $(\Box^* R)$ corresponding to

$$\frac{\begin{array}{c} \vdots \\ \hat{\alpha}_{p(u_0)} \Rightarrow \delta \end{array} \quad \dots \quad \begin{array}{c} \vdots \\ \hat{\alpha}_{p(u_{n-1})} \Rightarrow \delta \end{array} \quad \begin{array}{c} \vdots \\ \hat{\alpha}_{p(u_n)} \Rightarrow \delta \end{array}}{\bigvee \{\hat{\alpha}_{p(u_0)}, \dots, \hat{\alpha}_{p(u_n)}\} \Rightarrow \delta} \bigvee L$$

$$\frac{\hat{\alpha}_{p(u)}^- \Rightarrow \delta}{\hat{\alpha}_{p(u)} \Rightarrow \delta} \Box R, \bigwedge L$$

$$\frac{\hat{\alpha}_{p(u)}^- \Rightarrow \delta}{\hat{\alpha}_{p(u)} \Rightarrow \delta} \sigma L$$

That the constructed derivation is a proof is a consequence of Proposition 6.15. \square

Let $\text{Lit}(\alpha)$ denote the set of literals occurring free in α , i.e., $\text{Lit}(\alpha) = \{x \in \text{Lit} \mid \alpha \triangleright^* x\}$. As a corollary of the proof of Theorem 6.1, we obtain

Corollary 6.18. (Uniform Lyndon interpolation for constructive co-proofs)

For every formula α and set of literals L , there exists a formula ι satisfying

1. $\text{Lit}(\iota) \subseteq \text{Lit}(\alpha) \cap L$.
2. For every β with $\text{Lit}(\beta) \subseteq L$, $\alpha \Vdash \beta$ iff $\iota \Vdash \beta$.

Proof:

Let γ be the disjunctive formula constructed by the proof of Theorem 6.1 and define ι by replacing each subformula $\nabla_K(\Gamma, \Delta)$ of γ by $\nabla_{K \cap L}(\Gamma, \Delta)$. Then $\text{Lit}(\iota) \subseteq \text{Lit}(\gamma) \cap L \subseteq \text{Lit}(\alpha) \cap L$. Also, $\iota \Vdash \beta$ iff $\gamma \Vdash \beta$ provided $\text{Lit}(\beta) \subseteq L$. Since $\alpha \dashv\vdash \gamma$, we are done. \square

7. Completeness

We can now establish completeness of the finitary proof system $K\mu$. The result arises as a corollary of the stronger result that the constructive calculus $K\mu$ is complete for ‘negative’ sequents, i.e., $cK\mu \vdash \alpha^\perp \Rightarrow \perp$ if α is valid. This partial completeness result for the constructive μ -calculus arises from the partial interpretation of constructive co-proofs as cyclic proofs (Theorem 5.12) by applying the result to a valid sequent $\delta \Rightarrow \perp$ where $\delta \in \Pi(\mathbb{D})$ is such that $cC\mu \vdash \alpha^\perp \Rightarrow \delta$. As the cyclic and finitary calculi coincide, (full) completeness for $C\mu$ and $K\mu$ obtains.

A natural candidate for the $\Pi(\mathbb{D})$ -formula δ is the disjunctive normal form given by Theorem 6.1. While this choice is sufficient (and is the choice in [36]), the argument behind $C\mu \vdash \alpha^\perp \Rightarrow \delta$ is simplified by a more careful definition.

Definition 7.1. Each formula α is associated a *companion* formula $\alpha_d \in \Pi(\mathbb{D})$ as follows. Let $\hat{\alpha}$ be the disjunctive companion of α .

- If $\alpha \in \text{Lit}$ atomic, $\alpha_d := \alpha$.
- If $\alpha = \bigcirc \Gamma$ for $\bigcirc \in \{\bigvee, \bigwedge\}$, then $\alpha_d := \bigcirc \{\gamma_d \mid \gamma \in \Gamma\}$.
- If $\alpha = \triangle \beta$ for $\triangle \in \{\Box, \Diamond\}$, then $\alpha_d := \triangle \beta_d$.
- If $\alpha = \nu \varphi$, then $\alpha_d := \nu \varphi_d$ where $\varphi_d = \lambda x. (\varphi x)_d$.
- If $\alpha = \mu \varphi$, then $\alpha_d := \hat{\alpha}$.

The definition of α_d applies the disjunctive normal form theorem to at most on formula on each α -thread. An induction on the complexity of α , applying Proposition 4.15, transitivity of \Vdash and the disjunctive normal form theorem, yields

Proposition 7.2. $\alpha_d \dashv\vdash \alpha$.

Inserting Proposition 7.2 into Theorem 5.12 gives rise to

Proposition 7.3. $cC\mu \vdash \alpha_d \Rightarrow \alpha$.

Provability of the other implication is by induction on formulas:

Proposition 7.4. $cC\mu \vdash \alpha \Rightarrow \alpha_d$.

Proof:

By induction on α . We treat the quantifier cases as the others are straightforward. Let \vdash denote provability in $cC\mu$. Suppose $\alpha = \nu\varphi$, so $\alpha_d = \nu\varphi_d$. The induction hypothesis yields $\vdash \varphi x \Rightarrow \varphi_d x$. So $\vdash \alpha \Rightarrow \varphi_d \alpha$ by substitution and $\vdash \alpha \Rightarrow \alpha_d$ by derived ν -induction (Proposition 5.6).

Now suppose $\alpha = \mu\varphi$. Proposition 7.2 implies $\varphi_d x \Vdash \varphi x$, Proposition 4.3 gives rise to

$$\varphi_d \alpha_d \Vdash \varphi \alpha \Vdash \mu\varphi \Vdash \alpha_d.$$

Therefore, $\vdash \varphi_d \alpha_d \Rightarrow \alpha_d$ by Theorem 5.12 as $\varphi_d \alpha_d \in \Pi(\mathbb{D})$. Also, $\vdash \varphi \alpha_d \Rightarrow \varphi_d \alpha_d$ via the induction hypothesis, so $\vdash \alpha \Rightarrow \varphi_d \alpha_d$ by derived μ -induction. An application of cut yields $\vdash \alpha \Rightarrow \alpha_d$. \square

The preceding propositions are the final ingredient to the reduction of co-proofs to cyclic proofs.

Theorem 7.5. If $\Phi \Vdash \alpha$ then $cC\mu \vdash \Phi \Rightarrow \alpha$.

Proof:

Let $\Phi_d = \{\beta_d \mid \beta \in \Phi\}$. Transitivity of \Vdash implies $\Phi_d \Vdash \alpha$, whence Theorem 5.12 gives rise to $cC\mu \vdash \Phi_d \Rightarrow \alpha$. Proposition 7.4 yields $cC\mu \vdash \Phi \Rightarrow \alpha$. \square

Combining Corollary 6.2 and Theorems 7.5 and 5.10, we deduce

Corollary 7.6. If α is valid, then $\alpha^\perp \Rightarrow \perp$ is provable in $cC\mu$ and $cK\mu$.

The closure of the classical calculi under negation yields completeness these logics:

Corollary 7.7. $K\mu$ and $C\mu$ are complete: If α is valid then $K\mu \vdash \Rightarrow \alpha$ and $C\mu \vdash \Rightarrow \alpha$.

References

- [1] Acclavio M, Curzi G, Guerrieri G. Infinitary cut-elimination via finite approximations (extended version), 2024. doi:10.48550/arXiv.2308.07789. ArXiv:2308.07789.
- [2] Afshari B, Leigh GE. Finitary proof systems for Kozen's μ . Oberwolfach preprints (OWP) ISSN 1864-7596, Mathematisches Forschungsinstitut Oberwolfach, 2016.
- [3] Afshari B, Leigh GE. Circular proofs for the modal mu-calculus. *Proceedings in Applied Mathematics and Mechanics*, 2016. **16**(1):893–894. doi:10.1002/pamm.201610435.

- [4] Afshari B, Leigh GE. Cut-free completeness for modal μ -calculus. In: 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017. IEEE Computer Society, 2017 pp. 1–12. doi:10.1109/LICS.2017.8005088.
- [5] Afshari B, Leigh GE. Lyndon interpolation for modal μ -calculus. In: Özgün A, Zinova Y (eds.), Language, Logic, and Computation. Springer International Publishing, Cham, 2022 pp. 197–213. doi:10.1007/978-3-030-98479-3_10.
- [6] Afshari B, Leigh GE. Normal forms for cyclic proofs, 2023. Unpublished.
- [7] Afshari B, Leigh GE. A metapredicative study of μ -arithmetics, 2025. To appear.
- [8] Afshari B, Leigh GE, Menéndez Turata G. Uniform interpolation from cyclic proofs: The case of modal μ -calculus. In: Das A, Negri S (eds.), Automated Reasoning with Analytic Tableaux and Related Methods. TABLEAUX 2021, volume 12842 of *Lecture Notes in Computer Science*. Springer, Cham, 2021 doi:10.1007/978-3-030-86059-2_20.
- [9] Arai T. Quick cut-elimination for strictly positive cuts. *Annals of Pure and Applied Logic*, 2011. **162**(10):807–815. doi:10.1016/j.apal.2011.03.002.
- [10] Arnold A, Niwiński D. Rudiments of μ -calculus, volume 146 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2001. ISBN 978-0-444-50620-7.
- [11] Baelde D, Doumane A, Kuperberg D, Saurin A. Bouncing Threads for Circular and Non-Wellfounded Proofs: Towards Compositionality with Circular Proofs. In: Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science. ACM. ISBN 978-1-4503-9351-5, 2022 pp. 1–13. doi:10.1145/3531130.3533375.
- [12] Baelde D, Doumane A, Saurin A. Infinitary proof theory: the multiplicative additive case. In: Talbot J, Regnier L (eds.), 25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 - September 1, 2016, Marseille, France, volume 62 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016 pp. 42:1–42:17.
- [13] Das A, Marin S. On Intuitionistic Diamonds (and Lack Thereof). In: Ramanayake R, Urban J (eds.), Automated Reasoning with Analytic Tableaux and Related Methods. Springer Nature Switzerland, Cham. ISBN 978-3-031-43513-3, 2023 pp. 283–301. doi:10.1007/978-3-031-43513-3_16.
- [14] Dekker M, Klobhofer J, Marti J, Venema Y. Proof Systems for the Modal μ -Calculus Obtained by Determinizing Automata. In: Ramanayake R, Urban J (eds.), Automated Reasoning with Analytic Tableaux and Related Methods - 32nd International Conference, TABLEAUX 2023, Prague, Czech Republic, September 18-21, 2023, Proceedings, volume 14278 of *Lecture Notes in Computer Science*. Springer, 2023 pp. 242–259. doi:10.1007/978-3-031-43513-3_14.
- [15] Doumane A. Constructive completeness for the linear-time μ -calculus. In: 2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). IEEE, Reykjavik, Iceland. ISBN 978-1-5090-3018-7, 2017 pp. 1–12. doi:10.1109/LICS.2017.8005075.
- [16] Enqvist S, Seifan F, Venema Y. Completeness for the modal μ -calculus: Separating the combinatorics from the dynamics. *Theoretical Computer Science*, 2018. **727**:37–100. doi:10.1016/j.tcs.2018.03.001.
- [17] Fitch F. Intuitionistic modal logic with quantifiers. *Port. Math.*, 1948. **7**(2):113–118.
- [18] Fortier J, Santocanale L. Cuts for circular proofs: semantics and cut-elimination. In: Rocca SRD (ed.), Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy, volume 23 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013 pp. 248–262. doi:10.4230/LIPICs.CSL.2013.248.

- [19] Janin D, Walukiewicz I. Automata for the modal μ -calculus and related results. In: Wiedermann J, Hájek P (eds.), *Mathematical Foundations of Computer Science 1995: 20th International Symposium, MFCS '95* Prague, Czech Republic, August 28–September 1, 1995 Proceedings. Springer Berlin Heidelberg, Berlin, Heidelberg, 1995 pp. 552–562.
- [20] Jungteerapanich N. A tableau system for the modal μ -calculus. In: Giese M, Waaler A (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, 18th International Conference, TABLEUX 2009, Oslo, Norway, July 6–10, 2009. Proceedings, volume 5607 of *Lecture Notes in Computer Science*. Springer, 2009 pp. 220–234.
- [21] Kaivola R. Axiomatising linear time μ -calculus. In: Goos G, Hartmanis J, Leeuwen J, Lee I, Smolka SA (eds.), *CONCUR '95: Concurrency Theory*, volume 962 of *Lecture Notes in Computer Science*, pp. 423–437. Springer Berlin Heidelberg, Berlin, Heidelberg. ISBN 978-3-540-60218-7 978-3-540-44738-2, 1995. doi:10.1007/3-540-60218-6_32.
- [22] Kloibhofer J. A note on the incompleteness of Afshari & Leigh's system Clo, 2023. doi:10.48550/arXiv.2307.06846. ArXiv:2307.06846 [math].
- [23] Kozen D. Results on the propositional μ -calculus. *Theoretical Computer Science*, 1983. **27**:333–354.
- [24] Mints GE. Finite investigations of transfinite derivations. *Journal of Soviet Mathematics*, 1978. **10**(4):548–596. doi:10.1007/BF01091743.
- [25] Leigh GE, Wehr D. From GTC to Reset: Generating reset proof systems from cyclic proof systems. *Annals of Pure and Applied Logic*, 2024. **175**(10):103485. doi:10.1016/j.apal.2024.103485.
- [26] Niwinski D, Walukiewicz I. Games for the μ -calculus. *Theoretical Computer Science*, 1996. **163**(1&2):99–116.
- [27] Plotkin G, Stirling C. A framework for intuitionistic modal logics. In: *Proceedings of the 1986 Conference on Theoretical aspects of reasoning about knowledge*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. ISBN 978-0-934613-04-0, 1986 pp. 399–406.
- [28] Saurin A. A linear perspective on cut-elimination for non-wellfounded sequent calculi with least and greatest fixed-points. In: Ramanayake R, Urban J (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods - 32nd International Conference, TABLEUX 2023*, Prague, Czech Republic, September 18–21, 2023, Proceedings, volume 14278 of *Lecture Notes in Computer Science*. Springer, 2023 pp. 203–222. doi:10.1007/978-3-031-43513-3_12.
- [29] Savateev Y, Shamkanov D. Cut-Elimination for the Modal Grzegorczyk Logic via Non-well-founded Proofs. In: Kennedy J, De Queiroz RJ (eds.), *Logic, Language, Information, and Computation*, volume 10388 of *Lecture Notes in Computer Science*, pp. 321–335. Springer Berlin Heidelberg, Berlin, Heidelberg. ISBN 978-3-662-55385-5 978-3-662-55386-2, 2017. doi:10.1007/978-3-662-55386-2_23.
- [30] Shamkanov D. Circular Proofs for the Gödel-Löb Provability Logic. *Mathematical Notes*, 2014. **96**(3-4):575–585. doi:10.1134/S0001434614090326.
- [31] Simpson AK. *The Proof Theory and Semantics of Intuitionistic Modal Logic*. PhD, University of Edinburgh, School of Informatics, 1994. URL <http://hdl.handle.net/1842/407>.
- [32] Stirling C. A tableau proof system with names for modal μ -calculus. In: Voronkov A, Korovina MV (eds.), *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, volume 42 of *EPiC Series in Computing*, pp. 306–318. EasyChair, 2014.
- [33] Streett RS, Emerson EA. An automata theoretic decision procedure for the propositional μ -calculus. *Information and Computation*, 1989. **81**:249–264.

- [34] Studer T. Proof-Theoretic Contributions to Modal Fixed Point Logics. Habilitationsschrift, Universität Bern, 2009.
- [35] Walukiewicz I. Completeness of Kozen's axiomatisation of the propositional μ -calculus. In: Proceedings of Tenth Annual IEEE Symposium on Logic in Computer Science. 1995 pp. 14–24. doi:10.1109/LICS.1995.523240.
- [36] Walukiewicz I. Completeness of Kozen's Axiomatisation of the Propositional μ -Calculus. *Information and Computation*, 2000. **157**(1-2):142–182. doi:10.1006/inco.1999.2836.
- [37] Wijesekera D. Constructive modal logics I. *Annals of Pure and Applied Logic*, 1990. **50**(3):271–301. doi:10.1016/0168-0072(90)90059-B.