Chapter 11

Interpolation via Cyclic Proofs

Bahareh Afshari
University of Gothenburg, Sweden

Graham E. Leigh
University of Gothenburg, Sweden

Abstract

This chapter provides an expository account of cyclic proofs for fragments of the modal μ-calculus and demonstrates how the Maehara–Pitts method of interpolant construction can be lifted to the formalism. Existing results on separation and decidability in the modal μ-calculus are recast as consequences of the form of obtained interpolants. The wider applicability of the methods to other modal fixed point logics and μ-calculi is discussed.

Contents

In Balder ten Cate, Jean Christoph Jung, Patrick Koopmann, Christoph Wernhard, and Frank Wolter, editors. *Theory and Applications of Craig Interpolation*. Ubiquity Press, 2026. To appear.
https://www.doi.org/10.5334/bdg.k

1 Introduction

Many modal and temporal logics can be realised as extensions of a basic modal logic by certain inductive and co-inductive definitions. Common examples are *linear-time temporal logic* which extends the basic logic of a linear modality by the inductively defined 'until' operator, *propositional dynamic logic*, extending a multi-agent version of modal logic K by inductive operators for program recursion, and the *logic of common knowledge* which extends multi-agent epistemic logic by the co-inductive 'common knowledge' operator, as well as many *description logics* which might extend various modal logics by transitive concepts or even fixed point definitions. At the extreme are the so-called μ-calculi which incorporate explicit quantifiers for defining inductive and co-inductive operators of any (definable) specification. Linear-time temporal logic is, for example, a definable fragment of the linear-time μ-calculus and propositional dynamic logic is definable in the multi-modal μ-calculus. Other notable logics which are definable in the (single-agent) μ-calculus are *Gödel–Löb provability logic* and *compositional tree logic*.

Axiomatic systems for fixed point logics and μ-calculi invariably include axioms or rules of induction (and co-induction) to enforce the correct interpretation of operators. These rules present a challenge for the proof-theoretic approach to interpolation. Whether it be the Maehara method for Craig–Lyndon interpolation [26] or Pitts' construction of uniform interpolants [32], a clear obstacle for *any* logic is the isolation of a finitary sound and complete analytic[1] proof system. For the construction of uniform interpolants one expects, in addition, the analytic calculus to admit a terminating proof-search algorithm. General properties of a proof calculus that are sufficient for interpolation were explored in [17].

If the addition of induction principles to a logic does not outright block completeness of analytic proofs (as it does over first-order logic), it certainly complicates their design and theory. For many modal logics, analyticity can be regained by making structural changes

[1] For the purpose of the present chapter, it suffices to take *analytic* to mean that the premises of all rules employ the same (or smaller) signature than the conclusion, though stronger conditions are needed to ensure the existence of interpolants (cf. [17]). Notably, this simple constraint already eliminates the most problematic inference rule from the perspective of interpolation, the cut rule, which permits the derivation of an implication $\alpha \to \beta$ from implications $\alpha \to \iota$ and $\iota \to \beta$ for arbitrary ι.

to the proof calculus. In some cases it suffices to switch from a pure sequent calculus to a calculus of *labelled*, *hyper-* or *nested* sequents (see [17] for detail). When the fixed point logic is given by adding only inductively and co-inductively defined operators to a robust logic such as K, it is possible to obtain an analytic calculus by embedding induction into the *structure* of proofs. Proofs typically become infinite objects, either as well-founded trees over rules that may have infinitely many premises or as finitely branching *illfounded* trees. In the latter class, soundness of proofs is ensured by a syntactic correctness condition on admissible infinite branches.

This chapter derives interpolation from illfounded proofs for the modal μ-calculus and the various fixed point logics it encompasses. The expressive 'simplicity' of modal logic ensures that the focus can remain on *finite* proof objects, specifically, the illfounded proofs that are represented as tree unravellings of finite graphs, known as the *cyclic proofs*. In this representation, the correctness condition requires that every infinite path through the cyclic proof fulfils a decidable property known as the *global trace condition*. The cyclic proofs are sound and complete for the modal μ-calculus [30] and linear-time μ-calculus [14]. Analyticity of the inference rules means that the cyclic proofs are also sound and complete for any fragment that is determined by inductive and co-inductive (modal) operators.[2]

The construction of Craig, Lyndon and uniform interpolants from cyclic proofs proceeds in the usual fashion for analytic sequent calculi: candidate interpolants are assigned to each leaf of the proof-tree (in the case of uniform interpolation a proof-search-tree) and, recursively, to inner vertices. Unlike the purely finitary proofs, a cyclic proof can contain leaves that are not initial sequents and merely cycle back to 'earlier' points in the proof. These leaves, called 'buds', cannot be interpolated directly. Instead, buds are assigned placeholder 'interpolants' which are explicitly defined once the recursive construction has traversed the entire associated cycle. Correctness of the construction (that the candidate interpolant fulfils the requirements of a Craig/Lyndon/uniform interpolant) follows from a normal form theorem for cyclic proofs which ensures that the relative ordering of cycles in a proof internalises salient properties of the global trace condition.

As the definition and required properties of cyclic proofs for the modal μ-calculus is quite involved, it is introduced in two stages. First, we treat the fragment involving *only* inductive (or co-inductive) operators where the notion of cyclic proof is close to the naive one and the aforementioned normal form theorem is vacuous (Section 3). Lyndon and uniform interpolation are established for this fragment in Sections 4 and 5 respectively. Cyclic proofs for the full calculus are defined in Section 6 and followed by proofs of Lyndon and uniform interpolation in the general case. The chapter closes with a discussion of the wider applicability of the cyclic Maehara method to modal fixed point logics and other μ-calculi.

---

[2] Presumably, the result also holds for robust fragments of guarded first order fixed point logic (cf [11]) though to the best of the authors' knowledge this has not been examined.

Beyond proofs of Lyndon and uniform interpolation, the chapter explores the complexity of constructed interpolants in terms of alternations between inductive and co-inductive operators. Several well-known theorems on the expressivity of the modal μ-calculus can be realised as corollaries of interpolation augmented by structural observations on the form of cyclic proofs for specific quantifier classes. The construction of Craig–Lyndon interpolants witnesses, for example, the collapse of the quantifier class Δ_1 to the modal fragment and the collapse of Δ_2 to the alternation-free fragment. Corollaries of the uniform interpolation theorem include Otto's theorem on the decidability of equivalence to the modal fragment [31], Küsters and Wilkie's corresponding theorem for the co-induction-free fragment [24] and the disjunctive normal form theorem of Janin and Walukiewicz [20].

This chapter does not contribute any new results. Uniform interpolation for the modal μ-calculus was established by D'Agostino and Hollenberg in 2000 via a semantic argument [13]. The argument of [13] is discussed in [18] and will not be rehearsed here. The approach of this chapter is a generalisation of the proof-theoretic verification of Lyndon interpolation and uniform Lyndon interpolation in, respectively, [3] and [4]. The first explicit use of cyclic proofs to establish interpolation is Shamkanov's proof of the Lyndon interpolation property for Gödel–Löb provability logic [35]. Predating most of the above results, however, are various separation theorems for the modal μ-calculus and fragments [31, 24, 34, 7] which employ tableaux or related syntactic methods.

2 The modal μ-calculus

The following basic notions are used throughout this work. The set of finite sequences over a set X is denoted $X^{<\omega}$ and assumed to be canonically ordered by the (strict) prefix relation $<$. A *tree* is any non-empty set $T \subset \omega^{<\omega}$ closed under the prefix relation. Elements of T are called *vertices* with the empty sequence referred to as the *root* of T. By $<_T$ and $\leq_T$ we denote the strict and weak prefix relations restricted to vertices of T. The set of *immediate successors (in T)* of a vertex u is denoted $\mathrm{Suc}_T(u)$. If $|\mathrm{Suc}_T(u)| = 1$, the single immediate successor of u will often be denoted u^+. A *branch* of T is a (finite or infinite) sequence $(u_i)_i$ of vertices such that u_0 is the root and $u_{i+1} \in \mathrm{Suc}_T(u_i)$ for every i. A tree *over (a set)* X is a tree T equipped with a *labelling function* $\lambda\colon T \to X$ over X. Every finite tree F is assigned a *height*, in symbols $h(F)$, defined in the expected way with the trivial tree having height 0.

2.1 Syntax

Fix disjoint sets $\mathrm{Var} = \{x, y, z, \ldots\}$ and $\mathrm{Cons} = \{a, b, c, \ldots\}$ of, respectively, *(propositional) variables* and *(propositional) constants*. The *formulas* of the modal μ-calculus are the expressions α generated by the grammar:

$$\text{Formulas: } \alpha ::= x \mid a \mid \neg a \mid \bigwedge\Gamma \mid \bigvee\Gamma \mid \Box\alpha \mid \Diamond\alpha \mid \nu x\, \alpha \mid \mu x\, \alpha$$

$$\text{Cedents: } \Gamma ::= \emptyset \mid \Gamma \cup \{\alpha\}.$$

That is, formulas are constructed from variables, constants and negated constants by arbitrary finite conjunctions and disjunctions, the (unary) modal operators $\Box$ and $\Diamond$ and two propositional quantifiers,[3] μ and ν. The set of formulas is denoted L_μ. Symbols α, β, etc. range over formulas, and Γ, Δ, Θ, etc. over finite sets of formulas, called *cedents*. A constant or negated constant is called a *literal* and the set of literals is denoted $\mathrm{Lit} = \mathrm{Cons} \cup \{\neg a \mid a \in \mathrm{Cons}\}$. Expressions l, l', etc. range over literals, and L, L' over cedents of literals. Variables are not literals and are never negated. Formulas $\mu x\, \alpha$ and $\nu x\, \alpha$ are called *quantified* and express, respectively, the *least* and *greatest* fixed points of the propositional function $\varphi_\alpha \coloneqq \lambda x\, \alpha$. Symbols σ, σ', etc., act as metavariables for the two quantifiers. A formula without quantifiers is called *quantifier-free*.

▶ **Definition 2.1** (Signature). *The* signature *of a formula/cedent is the set of freely occurring literals and variables in the formula/cedent, given by* $\mathrm{sig}(\alpha) = \{\alpha\}$ *for* $\alpha \in \mathrm{Lit} \cup \mathrm{Var}$*,* $\mathrm{sig}(\sigma x\, \alpha) = \mathrm{sig}(\alpha) \setminus \{x\}$*,* $\mathrm{sig}(\Box\alpha) = \mathrm{sig}(\Diamond\alpha) = \mathrm{sig}(\alpha)$*, and* $\mathrm{sig}(\bigwedge\Gamma) = \mathrm{sig}(\bigvee\Gamma) = \mathrm{sig}(\Gamma) = \bigcup_{\gamma\in\Gamma} \mathrm{sig}(\gamma)$*. A formula* α *is* closed *if* $\mathrm{sig}(\alpha) \cap \mathrm{Var} = \emptyset$*.*

We introduce abbreviations for commonly used formulas and sets:

$$\begin{array}{lll} \top \coloneqq \bigwedge\emptyset & \alpha_1 \wedge \cdots \wedge \alpha_n \coloneqq \bigwedge\{\alpha_1, \dots, \alpha_n\} & \Box\Gamma \coloneqq \{\Box\gamma \mid \gamma \in \Gamma\} \\ \bot \coloneqq \bigvee\emptyset & \alpha_1 \vee \cdots \vee \alpha_n \coloneqq \bigvee\{\alpha_1, \dots, \alpha_n\} & \Diamond\Gamma \coloneqq \{\Diamond\gamma \mid \gamma \in \Gamma\} \end{array}$$

Substitution of formulas for variables is expressed by notation $\alpha(\beta/x)$, meaning the result of replacing each free occurrence of x in α by the formula β and renaming bound variables in α to avoid variable capture. As variables do not occur in literals, $\alpha(\beta/x)$ is always a formula. Negation is implemented through logical duality by an operation $\alpha \mapsto \alpha^\perp \in \mathsf{L}_\mu$ defined below. The exception is variables which, having no negated form, are considered self-dual. This approach is motivated by the action of negation on quantified formulas, which should identify the formula $\mu x(\Diamond x \vee a)$ with its 'complement' $\nu x(\Box x \wedge \neg a)$.

$$\begin{array}{lllll} a^\perp \coloneqq \neg a & \left(\bigwedge\Gamma\right)^\perp \coloneqq \bigvee \Gamma^\perp & (\Box\alpha)^\perp \coloneqq \Diamond\alpha^\perp & (\nu x\, \alpha)^\perp \coloneqq \mu x\, \alpha^\perp & \Gamma^\perp \coloneqq \{\gamma^\perp \mid \gamma \in \Gamma\} \\ \neg a^\perp \coloneqq a & \left(\bigvee\Gamma\right)^\perp \coloneqq \bigwedge \Gamma^\perp & (\Diamond\alpha)^\perp \coloneqq \Box\alpha^\perp & (\mu x\, \alpha)^\perp \coloneqq \nu x\, \alpha^\perp & x^\perp \coloneqq x \end{array}$$

There are two notions of subformula pertinent to the intension of quantified formulas. The *direct subformulas* of α are the formulas encountered during the construction of α in the generating grammar. So α is a direct subformula of $\Box\alpha$ and $\Diamond\alpha$, of $\bigwedge\Gamma$ and $\bigvee\Gamma$ if $\alpha \in \Gamma$, and of $\mu x\alpha$ and $\nu x\alpha$ for all $x \in \mathrm{Var}$. Literals and variables have no direct subformulas. The *complexity* of a formula α, in symbols $|\alpha|$, is the length of the longest chain of direct subformulas starting from α.

---

[3] The term 'quantifier' refers to the role of the operators in Tarskian semantics in which they are interpreted as quantifiers ranging over truth values; cf Proposition 2.11.

A second notion of subformula is more relevant to the proof systems and semantics utilised in this chapter, where the provability or validity of a quantified formula $\sigma x \alpha$ is made contingent on the provability/validity of the fixed point 'unfolding' $\alpha(\sigma x \alpha / x)$. This notion is called the Fischer–Ladner subformula relation, or *FL-subformula*, after its introduction in [16]. Unlike the relation of direct subformula, the FL-subformula relation is not well-founded: $\alpha = \mu x \Box x$ is a proper FL-subformula of $\Box\alpha$, which is a proper FL-subformula of α. A formula α is an *immediate FL-subformula* of β iff either

1. $\beta = \triangle\alpha$ for $\triangle \in \{\Box, \Diamond\}$,
2. $\beta = \bigcirc\Gamma$ for $\bigcirc \in \{\bigwedge, \bigvee\}$ and $\alpha \in \Gamma$, or
3. $\beta = \sigma x \gamma$ and $\alpha = \gamma(\beta/x)$ for some $\sigma \in \{\mu, \nu\}$, $x \in \mathrm{Var}$ and γ.

The relation of *FL-subformula* is the reflexive and transitive closure of the *immediate FL-subformula*. The set of FL-subformulas of α is denoted $\mathsf{FL}(\alpha)$. An induction on the complexity of α confirms that $\mathsf{FL}(\alpha)$ is finite.

▶ **Proposition 2.2.** $\mathsf{FL}(\alpha)$ *is bounded in cardinality by the number of direct subformulas of* α.

A *thread* is a finite or infinite sequence of formulas weakly following the FL-subformula relation, that is, a sequence $t = (\alpha_i)_{i<K}$ for $K \leq \omega$ such that for all $i + 1 < K$, either $\alpha_{i+1} = \alpha_i$ or α_{i+1} is an immediate FL-subformula of α_i. In the second case, the thread is said to *progress at* i. Of the infinite threads, we are primarily interested in those that progress infinitely often. These we call *progressing* threads.

The following present some important properties of threads.

▶ **Proposition 2.3.** $\mathrm{sig}(\alpha) = \mathsf{FL}(\alpha) \cap (\mathrm{Var} \cup \mathrm{Lit})$.

▶ **Proposition 2.4.** *If* β *is an FL-subformula of* α *then* $\mathsf{FL}(\beta) \subset \mathsf{FL}(\alpha)$ *and* $\mathrm{sig}(\beta) \subset \mathrm{sig}(\alpha)$.

▶ **Proposition 2.5.** *Let* $(\alpha_i)_{i\leq k}$ *be a finite thread and suppose* $|\alpha_k| < |\alpha_i|$ *for all* $i < k$. *Then there exists a thread* $(\beta_i)_{i\leq k}$ *such that* $\beta_k \in \mathrm{Var}$ *and* $\alpha_i = \beta_i(\alpha_k/\beta_k)$ *for all* $i \leq k$.

Proof. By induction on k. Let $(\alpha_i)_{i\leq k}$ be a finite thread and $|\alpha_k| < |\alpha_i|$ for all $i < k$. Let $k > 0$ and suppose $\alpha_0 = \sigma y\, \alpha'$. The induction hypothesis applied to the thread $(\alpha_1, \ldots, \alpha_k)$ yields a corresponding thread $(\beta_1, \ldots, \beta_k)$ such that $\alpha_i = \beta_i(\alpha_k/x)$ for all $1 \leq i \leq k$ where $x = \beta_k$. In particular, $\alpha'(\alpha_0/y) = \beta_1(\alpha_k/x)$. As $|\alpha_0| > |\alpha_k|$, there is a formula β' such that $\beta_1 = \beta'(\sigma y\, \beta'/y)$ and $\alpha' = \beta'(\alpha_k/x)$. Setting $\beta_0 = \sigma y\, \beta'$ completes the argument. The other cases of α_0 are straightforward. ◀

▶ **Proposition 2.6.** *Let* $(\alpha_i)_{i<\omega}$ *be a progressing thread. There is a unique formula* α *satisfying*

1. $\alpha = \alpha_i$ *for infinitely many* i,
2. *there exists* $k < \omega$ *such that* $|\alpha| \leq |\alpha_i|$ *for all* $i > k$.

Moreover, α *is quantified, namely* $\alpha = \sigma x \beta$ *for some* $\sigma \in \{\mu, \nu\}$, $x \in \mathrm{Var}$ *and* β.

▶ **Definition 2.7** (μ-/ν-thread)**.** *Let* t *be a progressing thread and* $\alpha = \sigma x \beta$ *the formula identified by Proposition 2.6. If* $\sigma = \mu$ *we refer to* t *as a* μ*-thread, otherwise as a* ν*-thread.*

2.2 Semantics

Formulas of the modal μ-calculus are evaluated against *pointed frames* (shortened to *frames*), tuples $F = \langle W, E, r, \lambda\rangle$ where $\langle W, E\rangle$ is a countable directed graph over a non-empty set W, $r \in W$ and $\lambda \colon W \to 2^{\mathrm{Cons}}$ is a function mapping a set of constants to each element of W. We call W the *carrier*, elements of W *worlds* and λ the *assignment*. The distinguished world r is called the *root* of F. Frames will be identified with their carriers, writing $u \in F$ to express that u is a world of F, and the assignment of F denoted λ_F. For a world $u \in F$, let $\langle F, u\rangle$ be the pointed frame $\langle W, E, u, \lambda\rangle$ with root u. For $u \in F$, define $F(u) = \{v \in W \mid (u, v) \in E\}$.

There are two (equivalent) semantics for the modal μ-calculus. One is a lifting of Tarski–Kripke semantics for modal logic to the quantifiers of L_μ. The other is to view satisfiability of L_μ formulas as an infinite game between two players, called the *model-checking game*. We will employ the latter as it aligns most closely to the realm of cyclic proofs; for more on either choice of semantics we refer the reader to [10, 8].

In game-semantics, formulas are conceived as specifying the rules of a game between two players, *Verifier* and *Refuter*, whose aim is to verify/refute claims of the kind 'F validates α'. In the case that $\alpha = \bigvee\Gamma$, for example, the rules stipulate that it is Verifier's turn and that her move is to select as the next position in the game a statement 'F validates γ' for some $\gamma \in \Gamma$. If Verifier is able to win the game from one such position, she will have confirmed that F is a model of γ and, therefore, that F is a model of α. Refuter's aim is to win the game regardless of the move Verifier makes.

Formally, the model checking game is represented as a directed graph whose vertices express the positions in the game. For a fixed closed formula α and pointed frame $\langle F, r\rangle$, a position is a pair (γ, v) for $\gamma \in \mathsf{FL}(\alpha)$ and $v \in F$. Each position is assigned an owner – Verifier or Refuter – and a set of available next moves, called *continuations*. From the initial position (α, r), a play of the game unfolds by the player owning the current position selecting the next from the set of continuations. A player who cannot make a move when required loses. For example, Verifier immediately wins in the position $(\bigwedge\emptyset, v)$ since there are no continuations and Refuter has ownership of the position.

Ownership of positions and admissible continuations are detailed in Table 1. The continuations from $(\Diamond\gamma, v)$ and $(\Box\gamma, v)$ are the positions (γ, w) for $w \in F(v)$. Verifier owns $(\Diamond\gamma, v)$ and Refuter owns $(\Box\gamma, v)$. For a formula $\alpha = \sigma x \beta$ there is just one continuation which is to move to its unfolding $\beta(\alpha/x)$. As the formula may increase in complexity, plays in the model-checking game can be infinite. The winning player of an infinite play is determined by the quantifiers which are unfolded infinitely often during the course of the play.

▶ **Definition 2.8** (Model checking game). *Let α be a closed formula and F a frame. The* model checking game *for (α, F) is the directed graph $G(\alpha, F)$ over vertices $\mathsf{FL}(\alpha) \times F$ with an edge from p to p' iff p' is a continuation of p per Table 1. Each position is assigned an* owner*, either* Verifier *or* Refuter*, detailed in the table.*

A play *of the model checking game $G(\alpha, F)$ is a maximal path through the graph starting from the position (α, r) where r is the root of F.*

Table 1 The model-checking game $G(\alpha, F)$.

Positions owned by Verifier		Positions owned by Refuter	
Position	Continuations	Position	Continuations
(a, v) for $a \notin \lambda_F(v)$	none	(a, v) for $a \in \lambda_F(v)$	none
$(\neg a, v)$ for $a \in \lambda_F(v)$	none	$(\neg a, v)$ for $a \notin \lambda_F(v)$	none
$(\bigvee\Gamma, v)$	$\{(\gamma, v) \mid \gamma \in \Gamma\}$	$(\bigwedge\Gamma, v)$	$\{(\gamma, v) \mid \gamma \in \Gamma\}$
$(\Diamond\alpha, v)$	$\{(\alpha, w) \mid w \in F(v)\}$	$(\Box\alpha, v)$	$\{(\alpha, w) \mid w \in F(v)\}$
$(\mu x\alpha, v)$	$\{(\alpha(\mu x\alpha/x), v)\}$	$(\nu x\alpha, v)$	$\{(\alpha(\nu x\alpha/x), v)\}$

A finite play must end in a position that has no continuations, namely positions (l, v) for $l \in \text{Lit}$, $(\bigwedge\emptyset, v)$ and $(\bigvee\emptyset, v)$, as well as $(\Diamond\gamma, v)$ and $(\Box\gamma, v)$ if $F(v) = \emptyset$. In such cases we say that the play is *losing* for the player owning the final position and *winning* for the other player. An infinite play is an infinite sequence $p = (\alpha_i, v_i)_i$ of positions such that $(\alpha_i)_i$ is an infinite progressing thread. Definition 2.7 designates the thread as either μ or ν. Verifier is the winner of p iff the thread is ν.

The winning condition on plays is cofinal: for finite plays it depends only on the final position and for infinite plays is invariant under removing finite prefixes.

▶ **Definition 2.9** (Strategy). *A* strategy for Verifier (Refuter) *in* $G(\alpha, \langle F, r\rangle)$ *is a tree* $S = \langle T, \rho\rangle$ *over positions in the game with root labelled by* (α, r) *and such that for every* $u \in T$*,*

1. *if* $\rho(u)$ *is a position owned by Refuter (resp. Verifier) then every continuation from* $\rho(u)$ *is the label of a successor of* u *in* T*.*
2. *if* $\rho(u)$ *is a position owned by Verifier (resp. Refuter) and the set of continuations from* $\rho(u)$ *is non-empty then* u *has a successor in* T *labelled by a continuation of* $\rho(u)$*.*

A play is consistent *with* S *if the play enumerates a path through* S*. A strategy for Verifier (Refuter) is* winning *if every consistent play is winning for Verifier (Refuter).*

For a closed formula α, we set $\langle\!|\alpha|\!\rangle$ to be the set of pointed frames F for which there exists a winning strategy for Verifier in $G(\alpha, F)$. This definition gives rise to the following properties showing that the game semantics agrees with the Tarski–Kripke semantics.

▶ **Proposition 2.10.** *The following hold for every* $a \in \text{Cons}$*,* $\alpha \in \mathsf{L}_\mu$*, cedent* Γ *and quantifier* σ*.*

1. $\langle\!|a|\!\rangle = \{\langle F, r\rangle \mid a \in \lambda_F(r)\}$ *and* $\langle\!|\neg a|\!\rangle = \{\langle F, r\rangle \mid a \notin \lambda_F(r)\}$.
2. $\langle\!|\bigwedge\Gamma|\!\rangle = \bigcap_{\gamma\in\Gamma}\langle\!|\gamma|\!\rangle$ *and* $\langle\!|\bigvee\Gamma|\!\rangle = \bigcup_{\gamma\in\Gamma}\langle\!|\gamma|\!\rangle$.
3. $\langle\!|\Diamond\alpha|\!\rangle = \{\langle F, r\rangle \mid \exists u \in F(r)\ \langle F, u\rangle \in \langle\!|\alpha|\!\rangle\}$ *and* $\langle\!|\Box\alpha|\!\rangle = \{\langle F, r\rangle \mid \forall u \in F(r)\ \langle F, u\rangle \in \langle\!|\alpha|\!\rangle\}$.

All that is missing from the above is a characterisation of quantified formulas $\sigma x\,\alpha$. Let $\mathcal{F}$ be the set of frames. For a variable x and $U \subset \mathcal{F}$ of frames, let $\langle\!|\alpha|\!\rangle_{x\mapsto U}$ be the set of frames F for which Verifier has a winning strategy in the modified model-checking game $G(\alpha, F)$ where positions (x, u) are winning for Verifier iff $\langle F, u\rangle \in U$. The next result establishes the equivalence between game and Tarski–Kripke semantics.

▶ **Proposition 2.11.** *For every closed formula $\mu x\,\alpha$, $\langle\!|\mu x\,\alpha|\!\rangle = \bigcap\{U \subset \mathcal{F} \mid \langle\!|\alpha|\!\rangle_{x\mapsto U} \subset U\}$ and $\langle\!|\nu x\,\alpha|\!\rangle = \bigcup\{U \subset \mathcal{F} \mid U \subset \langle\!|\alpha|\!\rangle_{x\mapsto U}\}$.*

Proof sketch. Let S be a winning strategy for Verifier in $G(\mu x\,\alpha, F)$. Let $M \subseteq S$ be the set of vertices at which $\mu x\,\alpha$ occurs which, by assumption, is a well-founded subtree of S. Via transfinite induction on this tree it can be shown that if $\langle\!|\alpha|\!\rangle_{x\mapsto U} \subset U$ then $F \in U$. So $\langle\!|\mu x\,\alpha|\!\rangle \subset U$. As $\langle\!|\alpha|\!\rangle_{x\mapsto\langle\!|\mu x\,\alpha|\!\rangle} \subset \langle\!|\mu x\,\alpha|\!\rangle$, the first claim is complete. The second claim follows a dual argument: If $U \subset \langle\!|\alpha|\!\rangle_{x\mapsto U}$ then, by repeated application of the inclusion, a winning strategy in $G(\nu x\,\alpha, F)$ can be defined for every $F \in U$. ◀

In the model-checking game, negation acts as duality: exchanging ownership of all positions in $G(\alpha, F)$ is simply the game $G(\alpha^\perp, F)$ up to the naming of positions. Thus, winning strategies for Verifier in $G(\alpha, F)$ transfer to winning strategies for Refuter in $G(\alpha^\perp, F)$ and vice versa. The semantic properties above yield determinacy of the model-checking game: in every instance $G(\alpha, F)$ (exactly) one of the two players has a winning strategy:

▶ **Corollary 2.12** (Determinacy)**.** *For every closed α, $\langle\!|\alpha^\perp|\!\rangle = \mathcal{F} \setminus \langle\!|\alpha|\!\rangle$.*

▶ **Definition 2.13** (Validity, equivalence)**.** *A closed formula α is* satisfiable *if $\langle\!|\alpha|\!\rangle \neq \emptyset$ and* unsatisfiable *otherwise. We call α* valid *if $\langle\!|\alpha|\!\rangle$ comprises all frames. If $\langle\!|\alpha|\!\rangle = \langle\!|\beta|\!\rangle$ we say that α and β are* equivalent *and write $\alpha \equiv \beta$.*

2.3 Important Fragments

A variable x is *guarded* in α if every occurrence of x in α is under the scope of a modal operator. A formula is *guarded* if for every direct subformula $\sigma x\,\alpha$ the variable x is guarded in α. Replacing 'direct' in the above by 'FL-' yields an equivalent definition. Another equivalent formulation can be given in terms of cyclic threads:

▶ **Proposition 2.14** (Guarded fragment)**.** *A formula α_0 is guarded iff for every finite thread $(\alpha_i)_{i\leq k}$, if $\alpha_l = \alpha_k$ for some $l < k$ and the thread progresses between l and k, then for some $j \in [l, k)$ and $\triangle \in \{\Box, \Diamond\}$ it is the case that $\alpha_j = \triangle\alpha_{j+1}$.*

We define two closure operations on sets of formulas.

▶ **Definition 2.15** (Closure operation)**.** *For a set of formulas F, define Clos(F) to be the least set $C \subset \mathsf{L}_\mu$ satisfying*

1. *$F \subset C$,*
2. *If $\Gamma \subset C$ then $\{\bigwedge\Gamma, \bigvee\Gamma\} \subset C$,*
3. *If $\alpha \in C$ then $\{\Box\alpha, \Diamond\alpha\} \subset C$,*
4. *If $\alpha, \beta \in C$ and $x \in \mathrm{Var}$ then $\alpha(\beta/x) \in C$.*

For $\sigma \in \{\mu, \nu\}$ define Clos$^\sigma(F)$ to be the least set $C \subset \mathsf{L}_\mu$ satisfying above plus:

5. *If $\alpha \in C$ then $\sigma x\alpha \in C$.*

Note, in particular, that $\mathit{Clos}(F)$ (as well as $\mathit{Clos}^{\sigma}(F)$) is closed under substitution. So, for instance, $\mu x(x \vee \nu y(a \wedge y)) \in \mathit{Clos}(\{\mu x(x \vee z), \nu y(a \wedge y)\}) \subset \mathit{Clos}^{\mu}(\{x \vee z, \nu y(a \wedge y)\})$, whereas $\mu x(x \vee \nu y(x \wedge y)) \in \mathit{Clos}^{\mu}(\{x \vee z, \nu y(x \wedge y)\}) \setminus \mathit{Clos}(\{\mu x(x \vee z), \nu y(x \wedge y)\})$.

The closure operation determines a natural quantifier hierarchy of formulas.

▶ **Definition 2.16** (Niwiński quantifier hierarchy). *The sets $\Pi_n, \Sigma_n \subset \mathsf{L}_\mu$ are defined by mutual recursion:*

- $\Sigma_0 = \Pi_0$ *is the set of quantifier-free formulas,*
- $\Sigma_{n+1} = \mathit{Clos}^{\mu}(\Pi_n)$,
- $\Pi_{n+1} = \mathit{Clos}^{\nu}(\Sigma_n)$.

Formulas in $\mathit{Clos}(\Sigma_1 \cup \Pi_1)$ are called alternation-free.

▶ **Proposition 2.17** ([22]). *Every formula is equivalent to a guarded formula in the same quantifier classes.*

The winning condition of the model-checking game is trivial for Σ_1 and Π_1 formulas:

▶ **Lemma 2.18.** *Verifier wins no infinite plays in $G(\alpha, F)$ if $\alpha \in \Sigma_1$ and wins all infinite plays if $\alpha \in \Pi_1$.*

The following characterisation is straightforward.

▶ **Lemma 2.19.** $\Sigma_{n+1} \cap \Pi_{n+1} = \mathit{Clos}(\Sigma_n \cup \Pi_n)$.

A formula in $\Sigma_{n+1} \setminus \mathit{Clos}(\Sigma_n \cup \Pi_n)$ is called *genuine* Σ_{n+1}. Similarly, a formula in $\Pi_{n+1} \setminus \mathit{Clos}(\Sigma_n \cup \Pi_n)$ is *genuine* Π_{n+1}.

▶ **Definition 2.20** (Ambiguous classes). *For a set of formulas F, define $F^{\equiv}$ to be the set of formulas equivalent to some formula in F. The* ambiguous class at level n *is the set* $\Delta_n = \Sigma_n^{\equiv} \cap \Pi_n^{\equiv}$.

The counting of the ambiguous classes conventionally begins at the first level, so Δ_1 is called the *first* ambiguous class, Δ_2 the *second* and so on. From Lemma 2.19 we have that $\mathit{Clos}(\Sigma_n \cap \Pi_n)^{\equiv} \subset \Delta_{n+1}$ for all n. The two sets turn out to be identical for $n \in \{0, 1\}$. If α is equivalent to both a Σ_1 and Π_1 formula then α is equivalent to a quantifier-free formula, and if α is equivalent to formulas in Σ_2 and Π_2, then α is equivalent to a formula in the intersection $\Sigma_2 \cap \Pi_2$ [7]. Both results are corollaries of Craig interpolation, proven in Sections 4 (for $n = 0$) and 6 (for $n = 1$). A seminal result in the modal μ-calculus, established by Santocanale and Arnold [34], is that the comparison is strict at all other levels: for every $n > 2$ there exists a formula in Δ_n which is not equivalent to any formula in $(\Sigma_n \cap \Pi_n)^{\equiv}$.

$$\frac{\Theta, \Gamma \Rightarrow \Delta}{\bigwedge\Theta, \Gamma \Rightarrow \Delta}\bigwedge\mathsf{L} \qquad \frac{\alpha, \Gamma \Rightarrow \Delta \quad \text{for all } \alpha \in \Theta}{\bigvee\Theta, \Gamma \Rightarrow \Delta}\bigvee\mathsf{L} \qquad \frac{\alpha(\sigma x\, \alpha/x), \Gamma \Rightarrow \Delta}{\sigma x\, \alpha, \Gamma \Rightarrow \Delta}\sigma\mathsf{L}$$

$$\frac{\Gamma \Rightarrow \Delta, \alpha \quad \text{for all } \alpha \in \Theta}{\Gamma \Rightarrow \Delta, \bigwedge\Theta}\bigwedge\mathsf{R} \qquad \frac{\Gamma \Rightarrow \Delta, \Theta}{\Gamma \Rightarrow \Delta, \bigvee\Theta}\bigvee\mathsf{R} \qquad \frac{\Gamma \Rightarrow \Delta, \alpha(\sigma x\, \alpha/x)}{\Gamma \Rightarrow \Delta, \sigma x\, \alpha}\sigma\mathsf{R}$$

$$\frac{\alpha, \Gamma \Rightarrow \Delta}{\Gamma', \Diamond\alpha, \Box\Gamma \Rightarrow \Diamond\Delta, \Delta'}\Diamond\mathsf{L} \qquad \frac{\Gamma \Rightarrow \Delta, \alpha}{\Gamma', \Box\Gamma \Rightarrow \Diamond\Delta, \Box\alpha, \Delta'}\Box\mathsf{R}$$

$$\frac{}{\Gamma, l \Rightarrow l, \Delta}\mathsf{id} \qquad \frac{}{\Gamma, a, \neg a \Rightarrow \Delta}\mathsf{idL} \qquad \frac{}{\Gamma \Rightarrow \Delta, a, \neg a}\mathsf{idR}$$

$$\frac{}{\Gamma \Rightarrow \Delta}\mathsf{b} \qquad \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta}\mathsf{c}$$

Figure 1 Basic rules of inference. Rules denoted $\star\mathsf{L}/\star\mathsf{R}$ are called *left*/*right* rules respectively. Rules b and c are *structural* rules; others are *logical* rules. The rules $\Box\mathsf{R}$ and $\Diamond\mathsf{L}$ are, collectively, the *modal* rules. Instances of b and c are called *buds* and *companions* respectively.

3 Cyclic Derivations

We introduce a sequent calculus for the modal μ-calculus. The calculus derives expressions $\Gamma \Rightarrow \Delta$, called *sequents*, where Γ and Δ are finite sets of closed and guarded formulas, henceforth referred to as *cedents*. We refer to Γ as the *left* cedent of $\Gamma \Rightarrow \Delta$ and to Δ as the *right* cedent. In the context of sequents, formulas are identified with singleton sets and union is expressed by comma, so $\Gamma, \Delta = \Gamma \cup \Delta$ and $\Gamma, \alpha = \Gamma, \{\alpha\}$.

The sequent $\Gamma \Rightarrow \Delta$ is *valid* if $\langle\!|\bigwedge\Gamma|\!\rangle \subset \langle\!|\bigvee\Delta|\!\rangle$. That is, for every frame $\langle F, r\rangle$, if Verifier has a winning strategy from (γ, r) for each $\gamma \in \Gamma$, then she has a winning strategy from (δ, r) for some $\delta \in \Delta$. Corollary 2.12 implies that the sequents $\Gamma \Rightarrow \Delta$ and $\Gamma, \Delta^\perp \Rightarrow$ are equi-valid: one is valid iff both are.

Figure 1 presents the rules of our sequent calculus, named $\mathsf{K}\mu$. Dropping the structural and quantifier rules yields a natural sound and complete sequent calculus for basic modal logic K. A sequent which has been derived via the rule b will be called a *bud*. In the presence of b the calculus is clearly unsound, but incomplete without the rule. A sound and complete calculus for sequents in $\mathsf{K}\mu$ will be obtained by carefully constraining, in a purely syntactic manner, how buds may be utilised within a proof. In other words, the calculus of rules from Figure 1 should not be conceived of as a *proof* system per se, but as a system of *derivations* of which some will constitute *proofs*.

Before presenting the calculus in detail, we fix some basic concepts and notation.

▶ **Definition 3.1** (The sequent calculus $\mathsf{K}\mu$). *The sequent calculus $\mathsf{K}\mu$ is specified by the rules in Figure 1. For each rule, the sequent below the line is the* conclusion *and the sequents above are the* premises. *An instance of the structural rule b or c is called, respectively, a* bud *or* companion.

$$
\begin{prooftree}
\AxiomC{}
\RightLabel{b}
\UnaryInfC{$\Rightarrow \alpha, \beta, \alpha \vee \beta$}
\RightLabel{$\bigvee\mathsf{R}$}
\UnaryInfC{$\Rightarrow \alpha, \beta, \alpha \vee \beta$}
\RightLabel{c}
\UnaryInfC{$\Rightarrow \alpha, \beta, \alpha \vee \beta$}
\end{prooftree}
\qquad
\begin{prooftree}
\AxiomC{}
\RightLabel{id}
\UnaryInfC{$a, \alpha \Rightarrow a$}
\RightLabel{$\bigwedge\mathsf{L}$}
\UnaryInfC{$a \wedge \alpha \Rightarrow a$}
\RightLabel{$\Box\mathsf{R}$}
\UnaryInfC{$b, \Box(a \wedge \alpha) \Rightarrow \Box a, \Diamond\beta$}
\AxiomC{}
\RightLabel{b}
\UnaryInfC{$b, \alpha \Rightarrow \beta$}
\RightLabel{$\bigwedge\mathsf{L}$}
\UnaryInfC{$b \wedge \alpha \Rightarrow \beta$}
\RightLabel{$\Diamond\mathsf{L}$}
\UnaryInfC{$b, \Diamond(b \wedge \alpha) \Rightarrow \Box a, \Diamond\beta$}
\RightLabel{$\bigvee\mathsf{L}$}
\BinaryInfC{$b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \Box a, \Diamond\beta$}
\RightLabel{$\bigvee\mathsf{R}$}
\UnaryInfC{$b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \Box a \vee \Diamond\beta$}
\RightLabel{$\mu\mathsf{R}$}
\UnaryInfC{$b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \beta$}
\RightLabel{$\mu\mathsf{L}$}
\UnaryInfC{$b, \alpha \Rightarrow \beta$}
\RightLabel{c}
\UnaryInfC{$b, \alpha \Rightarrow \beta$}
\end{prooftree}
$$

Figure 2 Two examples of derivations. In the right-hand derivation, $\alpha = \mu x. \Box(a \wedge x) \vee \Diamond(b \wedge x)$ and $\beta = \mu x(\Box a \vee \Diamond x)$. The left-hand derivation is not a cyclic derivation because of the lack of a modal rule. The right-hand derivation qualifies as a cyclic derivation and, moreover, as a cyclic proof.

▶ **Definition 3.2** (Derivation). *A* derivation *in a sequent calculus is a tree over sequents and rules of the calculus such that every vertex and its immediate successors are the conclusion and premises respectively of the associated rule. The sequent labelling the root vertex is the* end-sequent *of the derivation.*

Relevant to subsequent definitions is the implicit constraint that every vertex of a derivation is a correct application of a rule of inference. In particular, leaves must be assigned rules with zero premises, such as id, idL, idR, b, or the specific instances of $\bigvee$L and $\bigwedge$R introducing the formulas $\bot = \bigvee\emptyset$ and $\top = \bigwedge\emptyset$. Such rule instances are called *initial*. The rules b and c are collectively called *structural* rules; all others rules are *logical*. The rules $\Box$R and $\Diamond$L are *modal* rules, and the rules id, idL and idR are *identity* rules.

In each instance of a logical rule one or more formulas in the conclusion sequent are distinguished and called the *principal* formula(s). This is the distinguished formula $\bigwedge\Theta$ in the conclusion of $\bigwedge$L and $\bigwedge$R, the formula $\sigma x\, \alpha$ in σL and σR, the formula $\Diamond\alpha$ in the left cedent of $\Diamond$L and $\Box\alpha$ in the right cedent of $\Box$R and, finally, the marked literals in the three identity rules. The *active* formulas are the formulas in the premise(s) of each rule that correspond to the principal formula: all formulas in Θ for $\bigwedge$L, the formula $\alpha(\sigma x\, \alpha/x)$ in the premise of σL, the formula α in the left cedent of the premise of $\Diamond$L, etc. All other formulas in a rule are called *side formulas*. Given a vertex v of a derivation, a formula is *principal* at v if it is the principal formula of the rule at v, and is *active* if it is an active formula in the sense of the rule at the parent vertex (to which v is a premise).

Finite derivations without buds are sound, in the sense that their end-sequent must be valid. The set of finite bud-free derivations is, however, incomplete as the rules of the calculus do not distinguish between the quantifiers μ and ν. If a formula α admits a bud-free derivation then so does the result of exchanging any ν for μ, and vice versa, in α. In particular, the valid sequent $\Rightarrow \nu x \Box x$ has no bud-free derivation.

A sound and complete calculus is obtained by permitting some (but not all) *infinite* derivations as proofs. The symmetry of μ and ν is broken by imposing a restriction on their occurrences along any infinite path. The valid formula $\nu x \Box x$ admits an infinite derivation by continually alternating applications of $\nu\mathsf{R}$ and $\Box\mathsf{R}$, and the invalid formula $\mu x \Box x$ has a derivation given by continually alternating $\mu\mathsf{R}$ and $\Box\mathsf{R}$. The infinitely many applications of the rule $\nu\mathsf{R}$ will qualify the former derivation as a 'proof' but not the latter.

In fact, for the modal μ-calculus (and many other modal/temporal logics), it is sufficient to treat only derivations which ultimately repeat themselves. These derivations will be called 'cyclic'. The special rules b and c provide a syntax for describing cyclic derivations. Given a derivation D, let $\mathrm{bud}(D)$ and $\mathrm{comp}(D)$ be the set of buds and, respectively, companions.

▶ **Definition 3.3** (Cyclic derivation)**.** *A* cyclic derivation *is a finite derivation D equipped with a function $\mathsf{c}\colon \mathrm{bud}(D) \to \mathrm{comp}(D)$ mapping each bud b to a companion $b^{\mathsf{c}} <_D b$ satisfying:*

1. *the sequents at b and b^{c} are identical,*
2. *there is an application of a modal rule on the path from b^{c} to b.*

We refer to b^{c} as the companion to b.

▶ **Example 3.4.** Figure 2 presents two derivations. The left-hand tree is a derivation because of the set-based perspective of sequents that allows the principal formula of $\bigvee\mathsf{R}$, namely $\alpha \vee \beta$, to occur also as a side formula. Although the single bud is repeating the root sequent and the latter is a companion, the tree is not a *cyclic* derivation because there is no modal rule between the (only) companion and bud. In contrast, the right-hand derivation in Figure 2 is a cyclic derivation: The companion to the single bud is also the root and there is a modal rule ($\Diamond\mathsf{L}$) between the two vertices. ⌟

Let D be a cyclic derivation. As the sequent at every bud in D is identical to that of its companion, any bud b of D could be extended by inserting a copy of the subderivation rooted at b^{c}. The result would be another cyclic derivation of the same sequent, which we call an *unfolding* of D. Repeatedly applying this operation so that every bud is eventually unfolded and removing superfluous companion rules, produces a bud-free derivation which we call the *unravelling* of D and denoted D_ω. If D has at least one bud, the unravelling of D will be an infinite derivation.

The branches of a cyclic derivation are all finite by assumption. We also want to consider sequences of vertices that describe branches of the *unravelling* of the cyclic derivation. We call these *paths* where a path through a cyclic derivation D is a finite or infinite sequence of vertices $(u_i)_{i<K}$ such that for every i, if $i+1 < K$ then either u_{i+1} is an immediate successor of u_i or u_i is a bud and $u_{i+1} = (u_i)^{\mathsf{c}}$ is the associated companion.

An important relation that we will make extensive use of later is a connectedness relation between companions taking into account the process of unravelling. We introduce a relation $\lhd_D \subset \mathrm{comp}(D) \times D$ defined by $u \lhd_D v$ iff $u \in \mathrm{comp}(D)$ is a companion and $u <_D v <_D w$ for some $w \in \mathrm{bud}(D)$ such that $u = w^{\mathsf{c}}$.

3.1 Cyclic Proofs at the First Level

Attending to sequents at the first level of the quantifier hierarchy, that is, $\Gamma \Rightarrow \Delta$ where $\Gamma \cup \Delta \subset \Sigma_1 \cup \Pi_1$, a simple syntactic constraint isolates a class of derivations that are sound and complete for the fragment. These are the derivations for which every bud is a repetition of a sequent on the same branch as the bud and, moreover, the bud contains either a Σ_1 formula in the left cedent or a Π_1 formula in the right cedent. The fact that a bud must be a repeated sequent means that a proof for this fragment can be conceived of as an infinite derivation where in place of the bud the derivation loops back to its repeat. Unsurprisingly, such derivations will be called *cyclic proofs*.

We treat the notion of cyclic proof for sequents at the first level in detail in this section. This will be sufficient to obtain various Lyndon and Uniform Interpolation theorems for sequents and formulas at the first quantifier level, covered in Sections 4 and 5. Sequents beyond the first level necessitate a more careful definition of proof to ensure both soundness and completeness. The same concept of a cyclic derivation will be used but the condition on when a cyclic derivation is a *proof* is no longer localised to a simple property on each bud. Rather, we must consider the various infinite paths through the cyclic derivation and the formula threads that each path carries. A definition of proofhood appropriate for the second level and higher is the topic of Section 6.

For sequents at the first level of the quantifier hierarchy, the sound derivations can be readily carved out based on the formulas featured in buds and companions.

▶ **Definition 3.5** (Cyclic proofs at the first level). *Let $\Gamma \Rightarrow \Delta$ be a sequent at the first level and D a cyclic derivation of $\Gamma \Rightarrow \Delta$. Let $\Gamma_u \Rightarrow \Delta_u$ denote the sequent at vertex $u \in D$. A bud u is:*

Left successful *if Γ_u contains a genuine Σ_1 formula.*

Right successful *if Δ_u contains a genuine Π_1 formula.*

Successful *if u is left or right successful.*

A cyclic proof *is a finite cyclic derivation D for which every bud is successful. If there exists a cyclic proof of $\Gamma \Rightarrow \Delta$, we write $\vdash \Gamma \Rightarrow \Delta$.*

▶ **Example 3.6.** The right-hand derivation in Figure 2 is a cyclic proof. The end-sequent is in $\Pi_1 \cup \Sigma_1$. As the single bud contains a genuine Σ_1 formula in the left cedent (α), it is successful. As the left-hand derivation in that figure is not a cyclic derivation, neither is it a proof. Figure 3 presents a cyclic derivation demonstrating the necessity of requiring that the Π_1/Σ_1 formula in the bud is *genuine*. ⌟

Duality of the sequent calculus yields

▶ **Proposition 3.7.** $\vdash \Gamma \Rightarrow \Delta$ *iff* $\vdash\ \Rightarrow \Gamma^\perp, \Delta$, *iff* $\vdash \Gamma, \Delta^\perp \Rightarrow$.

Weakening is the rule that allows deriving sequents $\Gamma', \Gamma \Rightarrow \Delta, \Delta'$ from $\Gamma \Rightarrow \Delta$. We denote the rule by w. The sequent calculus $\mathsf{K}\mu$ implicitly incorporates weakening at initial sequents and in applications of the modal rules but it does not allow uses of the rule more generally.

$$
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\Rightarrow \Box\bot, \Box\mu x(\Box\bot \lor \Box x)}\,\mathsf{b}}{\Rightarrow \Box\bot \lor \Box\mu x(\Box\bot \lor \Box x)}\,\textstyle\bigvee\mathsf{R}}{\Rightarrow \mu x(\Box\bot \lor \Box x)}\,\mu\mathsf{R}}{\Rightarrow \Box\bot, \Box\mu x(\Box\bot \lor \Box x)}\,\Box\mathsf{R}}{\Rightarrow \Box\bot, \Box\mu x(\Box\bot \lor \Box x)}\,\mathsf{c}
$$

Figure 3 A cyclic derivation of an invalid sequent. The single bud of the derivation contains a Π_1 formula ($\Box\bot$) but no *genuine* Π_1 formula.

On occasion it is convenient to assume that the rule is explicitly included in the calculus. Indeed, the addition of the rule to $\mathsf{K}\mu$, while keeping the definition of cyclic derivation/proof otherwise unchanged, proves the same sequents as $\mathsf{K}\mu$ itself. The proof of this fact proceeds by 'pushing' each application of weakening upwards through the cyclic derivation (permuting with the logical rules) until it is absorbed by a modal or identity rule. A subtlety arises in uses of weakening below buds or companions. These are treated by unfolding the derivation once and continuing the process. We leave the full proof as an exercise.

▶ **Proposition 3.8.** *A sequent admits a cyclic proof in* $\mathsf{K}\mu + \mathsf{w}$ *iff it does so in* $\mathsf{K}\mu$.

We prove soundness and completeness for this notion of proof with respect to the game semantics of Section 2. The proof of soundness, that conclusions of cyclic proofs are valid, proceeds by relating strategies in the model-checking game to paths through a cyclic derivation and the success condition on buds to winning conditions.

▶ **Proposition 3.9** (Soundness). *For* $\Gamma \cup \Delta \subset \Sigma_1 \cup \Pi_1$, *if* $\vdash \Gamma \Rightarrow \Delta$ *then* $\Gamma \Rightarrow \Delta$ *is valid.*

Proof. Determinacy (Corollary 2.12) and duality (Proposition 3.7) imply that it suffices to treat the case that $\Delta = \emptyset$. Let $\Gamma \subset \Sigma_1 \cup \Pi_1$ and let D be a cyclic derivation of $\Gamma \Rightarrow$. Let F be any frame. We show that if there exists a winning Verifier strategy in the model checking game $G(\bigwedge\Gamma, F)$, then D is not a proof in the sense of Definition 3.5. Let S be any Verifier strategy in this game. We call a set $X \subset S$ a *bar* if no two elements of X lie on the same path through S.

Under the assumption that S is a winning strategy, we show it is possible to define a path $(u_i)_i$ through D, a sequence of finite bars $(X_i)_i$ and a sequence of worlds $(w_i)_i$ in F satisfying

1. The set of positions in X_i is $\{(\gamma, w_i) \mid \gamma \in \Gamma_i\}$ where $\Gamma_i \Rightarrow$ is the sequent at u_i.
2. Every position in X_{i+1} is either in X_i or is an available move from a position in X_i.
3. For every $i > 0$ there exists $j > i$ such that $X_i \cap X_j = \emptyset$.

The root of D is associated the bar X_0 of vertices of S immediately above the root and w_0 is the single world assigned to these positions. As the root of S is the position $(\bigwedge\Gamma, w_0)$, the first two conditions are met. Note that property 3 is a consequence of the first two since every path through D contains infinitely many vertices labelled by modal rules.

$$L \subset \mathrm{Lit}\,\frac{\gamma_1, \Delta \Rightarrow \qquad \cdots \qquad \gamma_n, \Delta \Rightarrow \qquad \Delta \Rightarrow}{\Diamond\gamma_1, \ldots, \Diamond\gamma_n, \Box\Delta, L \Rightarrow}\,\Diamond^*\mathsf{L}$$

Figure 4 The 'branching' modality rule ($\Diamond^*\mathsf{L}$) used to obtain templates. Applications of this rule in the calculus are subject to the constraint that no other rule of the calculus has the same conclusion. In particular, L is a consistent set of literals. The rightmost premise of $\Diamond^*\mathsf{L}$ is the *minor* premise; other premises are *major*.

Suppose the finite sequences $(u_i)_{i\leq k}$, $(X_i)_{i\leq k}$ and $(w_i)_{i\leq k}$ have been defined. If u_k is a companion, set u_{k+1} to be the premise; if u_k is a bud, u_{k+1} is the associated companion. In either case, $X_{k+1} = X_k$ and $w_{k+1} = w_k$.

If the rule at u_k is $\bigvee\mathsf{L}$ with principal formula $\bigvee\Delta$ then there is a position $(\bigvee\Delta, w_k)$ in the bar X_k. Since S is a winning, $\Delta \neq \emptyset$ and S has a position (δ, w_k) for some $\delta \in \Delta$ labelling the immediate successor of a vertex in X_k. Select u_{k+1} to be the successor of u_k whose active formula is δ, let $w_{k+1} = w_k$ and X_{k+1} be the bar exchanging the position $(\bigvee\Delta, w_k)$ for (δ, w_k) in X_k. The other non-modal rules are similar. Note, u_k cannot be an identity rule because in that case there will be positions (a, w_k) and $(\neg a, w_k)$ in S in contradiction with S winning.

Suppose the rule at u_k is a modal rule, namely $\Diamond\mathsf{L}$, with principal formula $\Diamond\gamma$. Let $\gamma, \Delta \Rightarrow$ be the sequent of the premise, so $\Box\Delta$ is a subset of the left cedent at u_k. By assumption, the positions $(\Diamond\gamma, w_k)$ and $(\Box\delta, w_k)$ for each $\delta \in \Delta$ label vertices in X_k. From the first, and the assumption that S is winning, we find a world w_{k+1} such that (γ, w_{k+1}) is the next position in S after $(\Diamond\gamma, w_k)$. We also find the positions (δ, w_{k+1}) among the immediate successors of the vertices in X_k. Setting X_{k+1} to be the corresponding vertices in S, satisfies the two conditions.

Let $X = \bigcup_i X_i$. This is a finitely branching subtree of S. Condition 3 provides that X is infinite and so X has an infinite branch by König's Lemma. Now consider a bud b of D which occurs infinitely often in the path $(u_i)_i$. Each formula at b necessarily appears infinitely often along some infinite path of X. As S is winning for Verifier, these formulas must all be Π_1, meaning that b is not successful. ◀

3.2 Proof Search

Completeness of cyclic proofs is shown via a proof-search argument: Given a sequent we define a cyclic derivation within a variant of $\mathsf{K}\mu$ and show that the derivation either contains a cyclic proof of the sequent or encodes a witness of its invalidity.

The proof-search calculus comprises the structural and left rules of $\mathsf{K}\mu$ with a branching modal rule in place of $\Diamond\mathsf{L}$. The special modal rule, $\Diamond^*\mathsf{L}$, is presented in Figure 4. A set of L of literals is *inconsistent* if $\{a, \neg a\} \subset L$ for some constant a, and is *consistent* otherwise. The conclusion of the branching modal rule must be a sequent of the form $\Diamond\Gamma, \Box\Delta, L \Rightarrow$ where $L \subset \mathrm{Lit}$ is a consistent set of literals. It has as a premise the sequent $\gamma, \Delta \Rightarrow$ for every $\gamma \in \Gamma$, called the *major* premises, and a final premise with the sequent $\Delta \Rightarrow$ alone, called the *minor*

$$
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{}{a, \Box\nu x\alpha, \neg a \Rightarrow}\,\mathsf{idL}
\qquad
\dfrac{
\dfrac{}{\nu x\alpha, \mu x\beta \Rightarrow}\,\mathsf{b}
\quad
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{}{\nu x\alpha \Rightarrow}\,\mathsf{b}
}{a, \Box\nu x\alpha \Rightarrow}\,\Diamond^*\mathsf{L}
}{a \wedge \Box\nu x\alpha \Rightarrow}\,\bigwedge\mathsf{L}
}{\nu x\alpha \Rightarrow}\,\nu\mathsf{L}
}{\nu x\alpha \Rightarrow}\,\mathsf{c}
}{}
}{a, \Box\nu x\alpha, \Diamond\mu x\beta \Rightarrow}\,\Diamond^*\mathsf{L}
}{a, \Box\nu x\alpha, \neg a \vee \Diamond\mu x\beta \Rightarrow}\,\bigvee\mathsf{L}
}{a, \Box\nu x\alpha, \mu x\beta \Rightarrow}\,\mu\mathsf{L}
}{a \wedge \Box\nu x\alpha, \mu x\beta \Rightarrow}\,\bigwedge\mathsf{L}
}{\nu x\alpha, \mu x\beta \Rightarrow}\,\nu\mathsf{L}
}{\nu x\alpha, \mu x\beta \Rightarrow}\,\mathsf{c}
}{}
}{}
}{}
$$

Figure 5 A template for $\{\nu x\,\alpha, \mu x\,\beta\}$ where $\alpha = a \wedge \Box x$ and $\beta = \alpha^\perp$. The central bud is successful as it contains a genuine Σ_1 formula, whereas the rightmost bud is unsuccessful. Pruning the lower instance of $\Diamond^*\mathsf{L}$ to its (only) major premise witnesses provability of the endsequent. Any other ordering of the logical rules results in the same sequents at modal rules.

premise. The minor premise plays no role in the proof-search beyond ensuring that instances of modal rules always contain at least one premise, relevant if $\Gamma = \emptyset$. Note, however, that if $\Delta = \emptyset$ then the minor premise will be the empty sequent $\Rightarrow$, which can be 'derived' via a degenerate application of $\Diamond^*\mathsf{L}$ with both conclusion and (single) premise empty.

We call the sequent calculus described above the *template calculus* and a (finite) cyclic derivation in the calculus a *template*. Figure 5 presents an example of a template.

▶ **Definition 3.10** (Template calculus). *The* template calculus *is the sequent calculus comprising the structural rules and logical left rules of* $\mathsf{K}\mu$*, with the modal rule* $\Diamond\mathsf{L}$ *replaced by the rule* $\Diamond^*\mathsf{L}$ *in Figure 4.*

A template *is a cyclic derivation T in this calculus such that between each bud and its associated companion there is an instance of the modal rule $\Diamond^*\mathsf{L}$. If the end sequent of T is $\Gamma \Rightarrow$, we refer to T as a* template for Γ *or an* α-template *if $\Gamma = \{\alpha\}$.*

Let T be a derivation in the template calculus. A *pruning* of T is a subtree P obtained by discarding all but one premise of each modal rule (and the entire subtree above a removed vertex). If the pruning always selects a major premise of an undeleted modal rule $\Diamond^*\mathsf{L}$ then by labelling these rules as $\Diamond\mathsf{L}$ the pruning can be conceived as a derivation in $\mathsf{K}\mu$. If T is a template then any pruning to a $\mathsf{K}\mu$ derivation will be a cyclic derivation.

▶ **Proposition 3.11.** *Let T be a template for $\Gamma \subset \Pi_1 \cup \Sigma_1$. If $\Gamma \Rightarrow$ is valid then T admits a pruning to a cyclic proof.*

Proof. Let T be a template for $\Gamma \subset \Sigma_1 \cup \Pi_1$. We recursively label each vertex of T as either *successful* or *unsuccessful*. A bud is successful iff it fulfils the criterion of Definition 3.5, namely the left cedent contains a genuine Σ_1 formula. A vertex labelled by $\bigvee\mathsf{L}$ is successful iff all of its premises are successful and a vertex labelled by $\Diamond^*\mathsf{L}$ is successful iff at least one major premise is successful. Initial sequents are always successful. All other rules are unary and are successful iff the single premise is successful.

If the root of T is successful then there is pruning of T to a cyclic proof. For the purpose of the claim we may assume that the root of T is unsuccessful. It suffices to show that $\langle\!|\bigwedge\Gamma|\!\rangle \neq \emptyset$, from which it follows that $\Gamma \Rightarrow$ is not valid. Let U be a subtree of T comprising the unsuccessful vertices but with only one successor chosen for any instance of $\bigvee\mathsf{L}$. By recursion through U we construct a frame $F = \langle W, E, r, \lambda\rangle$ and a winning strategy for Verifier in the game $G(\bigwedge\Gamma, F)$. The root of U is associated the frame comprising a single root world r with empty assignment and the Verifier strategy consists of just the starting position $(\bigwedge\Gamma, w_r)$ and the next available moves $\{(\gamma, w_r) \mid \gamma \in \Gamma\}$. For non-modal rules, the current world is unchanged from conclusion to premise and the leaf in the strategy corresponding to the principal formulas is extended by one step. In all cases but disjunction, this is straightforward. If the principal formula is a disjunction $\bigvee\Delta$ then, as the present vertex of U is unsuccessful, a premise of the rule is in U and the strategy is extended by selecting the active formula for this premise. Upon reaching a modal rule, we know that all major premises are present in U. Let $\Diamond\Gamma, \Box\Delta, L$ be the left cedent at this vertex of T and let w be the world which has been associated to this vertex. The first task is to specify the assignment at w which we take to be the set of constants present in the sequent, i.e., $\lambda(w) = L \cap \mathrm{Cons}$. This definition is well-formed as every world created in F gets associated to exactly one conclusion of a modal rule. Next, the frame is extended by adding a fresh world, w_γ, and edge (w, w_γ) to F for each $\gamma \in \Gamma$ and associating the world w_γ to the major premise whose active formula is γ. The extension of the Verifier strategy is the obvious one: positions $(\Diamond\gamma, w)$ in S are continued to (γ, w_γ), a position $(\Box\delta, w)$ is extended by all available moves $\{(\delta, w_\gamma) \mid \gamma \in \Gamma\}$ and positions (l, w) (for $l \in L$) are terminating in the strategy.

Once the recursive construction of F and S has reached all buds of U, the intended frame and strategy is completed by identifying leaves of F with the world associated to the companion in T and likewise for S. Consider any play consistent with S. If the play is finite, the final position was created from a modal rule in U and, by design, the play is winning for Verifier. If the play is infinite then it infinitely often contains a formula in an unsuccessful bud of T. As such formulas are all Π_1, the plays are winning for Verifier. ◀

▶ **Proposition 3.12** (Completeness). *For $\Gamma \cup \Delta \subset \Sigma_1 \cup \Pi_1$, if $\Gamma \Rightarrow \Delta$ is valid then $\vdash \Gamma \Rightarrow \Delta$ with a cyclic proof of height $\leq 2^{|\mathsf{FL}(\Gamma)|+|\mathsf{FL}(\Delta)|+1}$.*

Proof. By Proposition 3.11 it suffices to confirm that a template for Γ exists with the required height. Let us call a logical rule of $\mathsf{K}\mu$ *conservative* if its principal formula is not also a side formula. For the logical rules in Figure 1, the conservative instances are those where the principal formula of a left/right rule is not also a side in the same cedent. Thus, of the two following instances of $\mu\mathsf{L}$, the first is conservative if $\mu x\,\alpha \notin \Gamma$ and the second is an example of a non-conservative instance of the same rule regardless of Γ.

$$\frac{\Gamma, \alpha(\mu x\,\alpha/x) \Rightarrow \Delta}{\Gamma, \mu x\,\alpha \Rightarrow \Delta}\,\mu\mathsf{L} \qquad \frac{\Gamma, \mu x\,\alpha, \alpha(\mu x\,\alpha/x) \Rightarrow \Delta}{\Gamma, \mu x\,\alpha \Rightarrow \Delta}\,\mu\mathsf{L}$$

Starting with root $\Gamma \Rightarrow$, systematically apply conservative instances of non-modal (left) rules in any order. Each application of a conservative non-modal rule reduces a particular complexity measure on cedents given by $|\Gamma|_{\Diamond^*\mathsf{L}} = \sum_{\gamma\in\Gamma} |\gamma|_{\Diamond^*\mathsf{L}}$ where $|\gamma|_{\Diamond^*\mathsf{L}}$ is defined as $|\gamma|$ except that $|\Box\delta|_{\Diamond^*\mathsf{L}} = |\Diamond\delta|_{\Diamond^*\mathsf{L}} = 0$ for all formulas δ. The only non-obvious case is that of the quantifiers, where we recall that sequents comprise only guarded formulas and that if $\mu x\,\alpha$ is locally guarded then $|\alpha(\mu x\,\alpha/x)|_{\Diamond^*\mathsf{L}} < |\mu x\,\alpha|_{\Diamond^*\mathsf{L}}$.

Thus, the above process concludes in a finite derivation in non-modal left rules wherein every leaf is a sequent of the form $\Diamond\Theta, \Box\Delta, L \Rightarrow$ for $L \subset \mathrm{Lit}$. If L is inconsistent, this sequent is an instance of idL; otherwise, apply the branching modal rule $\Diamond^*\mathsf{L}$ and repeat the process from each premise. As every sequent produced will be a subset of the Fischer–Ladner closure of Γ, which is finite, every branch will eventually repeat. Closing the derivation at the first repeated sequent as a bud and inserting an instance of the companion rule yields a template. The bound on height arises from the observation that after the insertion of companions, each branch of the template duplicates a sequent along it at most once, except for the bud which is duplicated twice. ◀

Combining the above we obtain:

▶ **Theorem 3.13.** *For $\Gamma, \Delta \subset \Pi_1 \cup \Sigma_1$, the following are equivalent*

1. $\Gamma \Rightarrow \Delta$ *is valid.*
2. $\vdash \Gamma \Rightarrow \Delta$ *is witnessed by a cyclic proof with height* $\leq 2^{|\mathsf{FL}(\Gamma)|+|\mathsf{FL}(\Delta)|+1}$.

The arguments underlying the proof of soundness and completeness are not constrained to the first quantifier level. Given a cyclic derivation D of a sequent $\Gamma \Rightarrow$, the soundness argument translates Verifier strategies in a game $G(\bigwedge\Gamma, F)$ for some F into paths through D that, in a precise sense, carries plays consistent with the strategy. If the strategy is winning for Verifier then the path will be infinite. Dually, the proof of Proposition 3.11 describes a translation of certain subtrees of templates into Verifier strategies over certain frames, with a similar correspondence between paths through the derivation and plays consistent with the strategy. The role of the syntactic restriction on sequents is to link the *winning* strategies to cyclic *proofs*. In both proofs, the success condition on the buds of a cyclic proof ensures, via Lemma 2.18, that the described Verifier strategies are winning.

Two notable corollaries to Theorem 3.13 arise from specialising the form of the sequent.

▶ **Corollary 3.14.** *For $\alpha \in \Pi_1$ and $\beta \in \Sigma_1$, $\alpha \Rightarrow \beta$ is valid iff there exists a finite bud-free derivation of $\alpha \Rightarrow \beta$ of height at most* $2^{|\mathsf{FL}(\alpha)|+|\mathsf{FL}(\beta)|+1}$.

Proof. For $\alpha \in \Pi_1$ and $\beta \in \Sigma_1$, a cyclic proof of $\alpha \Rightarrow \beta$ is, by requirement, bud-free. ◀

▶ **Corollary 3.15.** *For $\alpha \in \Sigma_1$ and $\beta \in \Pi_1$, $\alpha \Rightarrow \beta$ is valid iff there exists a cyclic derivation of $\alpha \Rightarrow \beta$ of height at most $2^{|\mathsf{FL}(\alpha)|+|\mathsf{FL}(\beta)|+1}$.*

Proof. Similar to Corollary 3.14, but noting that for $\alpha \in \Sigma_1$ and $\beta \in \Pi_1$, every cyclic derivation is a cyclic proof. ◀

4 Lyndon Interpolation at the First Level

Via the simple cyclic proofs introduced in the previous section, Craig and Lyndon interpolation can be established for the first level of the modal μ-calculus, namely for sequents built from $\Sigma_1 \cup \Pi_1$. Lyndon interpolation is the property that if $\vdash \alpha \Rightarrow \beta$ then $\vdash \alpha \Rightarrow \iota$ and $\vdash \iota \Rightarrow \beta$ for some ι with $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$. The formula ι is called the *interpolant*. Recall that $\mathrm{sig}(\alpha)$ – the *signature* of α – records the polarity (negated or unnegated) of the constants occurring within α. In particular, for α closed $\mathrm{sig}(\alpha) = \mathsf{FL}(\alpha) \cap \mathrm{Lit}$.

The construction of interpolants directly generalises the proof-theoretic method for modal logic K presented in [9]. That argument recursively ascribes an interpolant to each sequent in a finite proof, starting at leaves and progressing towards the root through the modal and other logical rules. The desired Craig/Lyndon interpolant is the formula assigned to the root sequent.

The quantifier rules $\sigma\mathsf{L}$ and $\sigma\mathsf{R}$ do not affect the construction of interpolants as the signature is unaltered from premise and conclusion. That is, an interpolant chosen for the premise of such a rule is already an interpolant for the conclusion. As we shall see, it is the relaxation to *cyclic* proofs that presents the main hurdle to lifting interpolation from K to $\mathsf{K}\mu$.

4.1 Lyndon Candidates

Lifting interpolation from bud-free to cyclic proofs is not immediate. To visualise the difficultly, consider the cyclic proof of the sequent $b, \mu x.\Box(x \wedge a) \vee \Diamond(b \wedge x) \Rightarrow \nu x(\Box a \vee \Diamond x)$ from Figure 2, repeated in Figure 6. To directly apply the recursive construction from [9] one must first assign an 'interpolant' to the single bud $b, \alpha \Rightarrow \beta$ before building from this an interpolant for the root. Let z denote a hypothetical interpolant of the bud sequent. Following the standard construction leads to the formula $\Box a \vee \Diamond z$ as an interpolant for the sequent at the root (cf. left-hand derivation in Figure 6, ignoring the companion vertex). As bud and end-sequent are identical, it follows that each of z, $\Box a \vee \Diamond z$, $\Box a \vee \Diamond(\Box a \vee \Diamond z)$, etc. are all interpolants for $b, \alpha \Rightarrow \beta$. The dependency on z can be removed by finding a formula $\iota \equiv \Box a \vee \Diamond \iota$. That is, we can hope that an interpolant of the companion will be either $\mu z(\Box a \vee \Diamond z)$ or $\nu z(\Box a \vee \Diamond z)$ with the choice of quantifier determined by properties of the cyclic proof as a whole. For this specific example, the reader can confirm that both formulas are (Lyndon) interpolants, though this is not in general the case.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{a, \alpha \Rightarrow a}\,\mathsf{id}}{a \wedge \alpha \Rightarrow a}\,\bigwedge\mathsf{L}}{b, \Box(a \wedge \alpha) \Rightarrow \Box a, \Diamond\beta}\,\Box\mathsf{R} \qquad \dfrac{\dfrac{\dfrac{}{b, \alpha \Rightarrow \beta}\,\mathsf{b}}{b \wedge \alpha \Rightarrow \beta}\,\bigwedge\mathsf{L}}{b, \Diamond(b \wedge \alpha) \Rightarrow \Box a, \Diamond\beta}\,\Diamond\mathsf{L}}{b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \Box a, \Diamond\beta}\,\bigvee\mathsf{L}}{b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \Box a \vee \Diamond\beta}\,\bigvee\mathsf{R}}{b, \Box(a \wedge \alpha) \vee \Diamond(b \wedge \alpha) \Rightarrow \beta}\,\nu\mathsf{R}}{b, \alpha \Rightarrow \beta}\,\mu\mathsf{L}}{b, \alpha \Rightarrow \beta}\,\mathsf{c}
\qquad\qquad
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{a}\,\mathsf{id}}{a}}{\Box a}\,\Box\mathsf{R} \qquad \dfrac{\dfrac{\dfrac{}{z}\,\mathsf{b}}{z}}{\Diamond z}\,\Diamond\mathsf{L}}{\Box a \vee \Diamond z}\,\bigvee\mathsf{L}}{\Box a \vee \Diamond z}}{\Box a \vee \Diamond z}}{\Box a \vee \Diamond z}}{\sigma z(\Box a \vee \Diamond z)}\,\mathsf{c}$$

Figure 6 The cyclic proof from Figure 2 (left) and the formula associated to each vertex following Construction 1 for an arbitrary marking (right), with the root formula being the Lyndon candidate induced by this derivation. As before, $\alpha = \mu x(\Box(a \wedge x) \vee \Diamond(b \wedge x))$ and $\beta = \nu x(\Box a \vee \Diamond x)$.

The introduction of quantifiers into interpolants can be identified with companion vertices, meaning that the construction remains recursive in the given cyclic proof. The process begins at buds which are assigned variables so that buds with a common companion are given the same variable. Interpolants for other vertices are chosen to mirror the structure of the derivation, with quantifiers being introduced by companions, illustrated in Figure 6.

For reasons that will become clear, the choice of quantifier will be specified as an input to the construction of interpolants along with the cyclic derivation. Define a *marked (cyclic) derivation* to be a pair (D, σ) where D is a (cyclic) derivation and $\sigma\colon \mathrm{comp}(D) \to \{\mu, \nu\}$ is an assignment of a quantifier (either μ or ν) to each companion. The function σ will be referred to as a *marking* of D. The construction of a 'candidate' Lyndon interpolant from a marked cyclic derivation is formalised by the following.

▶ **Construction 1.** *Let (D, σ) be a marked cyclic derivation. Let $v \mapsto x_v\colon \mathrm{comp}(D) \to \mathrm{Var}$ be an injection of companion vertices in D into variables. In the following we write σ_v for $\sigma(v)$. If $v \in D$ has just one immediate successor, we denote it by v^+. Each $v \in D$ is associated a formula ι_v defined by recursion on D:*

Bud: *$\iota_v = x_{v^c}$, namely the variable associated to the bud's companion.*

Logical Rule:

id. *$\iota_v = l$ where l is the principal literal.*

idL. *$\iota_v = \bot$.*

idR. *$\iota_v = \top$.*

$\bigvee$L. *$\iota_v = \bigvee\{\iota_w \mid w \in \mathrm{Suc}_D(v)\}$.*

$\bigwedge$R. *$\iota_v = \bigwedge\{\iota_w \mid w \in \mathrm{Suc}_D(v)\}$.*

$\Diamond$L. *$\iota_v = \Diamond\iota_{v^+}$.*

$\Box$R. *$\iota_v = \Box\iota_{v^+}$.*

Any other logical rule. *$\iota_v = \iota_{v^+}$.*

Companion: *$\iota_v = \sigma_v x_v\, \iota_{v^+}$.*

The Lyndon candidate induced by (D, σ) *is the formula* $\iota_{(D,\sigma)} = \iota_r$ *where* r *is the root of* D*.*

We begin with some observations concerning the structure of Lyndon candidates. Let (D, σ) be any marked cyclic derivation.

▶ **Lemma 4.1.** $|\iota_{(D,\sigma)}| \leq h(D)$.

▶ **Lemma 4.2.** *If the endsequent of* D *is* $\Gamma \Rightarrow \Delta$ *then* $\mathrm{sig}(\iota_{(D,\sigma)}) \subset \mathrm{sig}(\Gamma) \cap \mathrm{sig}(\Delta)$*.*

▶ **Lemma 4.3.**
1. *If* $\mathrm{comp}(D) = \emptyset$ *then* $\iota_{(D,\sigma)}$ *is quantifier-free.*
2. *If* $\sigma(u) = \mu$ *for every* $u \in \mathrm{comp}(D)$ *then* $\iota \in \Sigma_1$*.*
3. *If* $\sigma(u) = \nu$ *for every* $u \in \mathrm{comp}(D)$ *then* $\iota \in \Pi_1$*.*

All three lemmas above are immediate consequences of Construction 1. The following generalisation of Lemma 4.3 will be useful in Section 6.

▶ **Lemma 4.4.**
1. *If* $\sigma(u) = \sigma(v)$ *whenever* $u \lhd_D v \in \mathrm{comp}(D)$*, then* $\iota_{(D,\sigma)}$ *is alternation-free.*
2. *If for every sequence* $u_0 \lhd_D \cdots \lhd_D u_k \in \mathrm{comp}(D)$ *we have* $|\{i < k \mid \sigma(u_i) \neq \sigma(u_{i+1})\}| \leq n$*, then* $\iota_{(D,\sigma)} \in \mathit{Clos}(\Sigma_{n+1} \cup \Pi_{n+1})$*.*

To validate that a candidate interpolant actually interpolates the intended sequent, it remains to build the witnessing proofs. This step also follows the proof-theoretic method presented in [9] and [17] whereby the derivation is modified by inserting in place of either the left or right cedent a subformula of the candidate interpolant. From a cyclic derivation of $\alpha \Rightarrow \beta$ and a candidate interpolant $\iota = \iota_{(D,\sigma)}$ derived from this derivation we can, without difficulty, obtain cyclic derivations for the sequents $\alpha \Rightarrow \iota$ and $\iota \Rightarrow \beta$.

Whether the constructed derivations are cyclic *proofs* depends on properties of the marking. Under the constant marking $u \mapsto \nu$, the sequent $\alpha \Rightarrow \iota$ always admits a cyclic proof, whereas provability of $\iota \Rightarrow \beta$ is trivialised by the opposite marking $u \mapsto \mu$. Finding a compromise that works for *both* sequents is the main hurdle to proving interpolation, though this problem manifests only for sequents beyond the first quantifier level.

The next two results describe the transformation in detail. As mentioned, the idea is to replace in each sequent of a cyclic proof $P \vdash \alpha \Rightarrow \beta$ the left (resp. right) cedent by the candidate interpolant associated to that vertex. Taken literally, however, this will not yield a derivation because the interpolants associated to the premise and conclusion of a companion cannot be linked by a rule of the calculus. The former will be a formula ι_u with a free variable x_v whereas the latter will be a quantified formula $\iota_v = \sigma_v x_v \, \iota_u$. But as the reader has likely guessed, the 'real' interpolant associated to the premise of a companion should be the unfolding of ι_v, namely $\iota_u(\iota_v/x_v)$. The necessary substitutions are expressed as follows.

▶ **Definition 4.5** (Lyndon candidate at a vertex). *Let (D, σ) be a marked cyclic derivation with root $\Gamma \Rightarrow \Delta$ and let ι_v be the formula associated to v by Construction 1. Let $C_{\lhd v} = \{u \in \mathrm{comp}(D) \mid u \lhd_D v\}$ be the companion vertices strictly below $v \in D$ relative to the ordering $\lhd_D$. For each $v \in D$, the* Lyndon candidate at v *is the formula $\hat{\iota}_v$ defined as*

$$\hat{\iota}_v = \iota_v(\iota_{c_1}/x_{c_1}) \cdots (\iota_{c_n}/x_{c_n}) \quad \text{where } c_1 < \cdots < c_n < v \text{ enumerate the elements of } C_{\lhd v}.$$

The reader can confirm that $\hat{\iota}_r = \iota_r = \iota_{(D,\sigma)}$ where r is the root of D and that $\hat{\iota}_v$ is a closed formula for every v.

The next results describe the cyclic derivations obtained by combining a cyclic proof with the candidate interpolants.

▶ **Lemma 4.6.** *Let (D, σ) be a marked cyclic derivation with root $\alpha \Rightarrow \beta$ and let $\Gamma_v \Rightarrow \Delta_v$ denote the sequent labelling vertex $v \in D$. Let $\hat{\iota}_v$ be the Lyndon candidate at v. There exist cyclic derivations of*

1. *$\alpha \Rightarrow \iota_{(D,\sigma)}$ with buds $\{\Gamma_w \Rightarrow \hat{\iota}_w \mid w \in \mathrm{comp}(D)\}$.*
2. *$\iota_{(D,\sigma)} \Rightarrow \beta$ with buds $\{\hat{\iota}_w \Rightarrow \Delta_w \mid w \in \mathrm{comp}(D)\}$.*

Lemma 4.6 is a special case of the following refinement. Although the next result violates the assumption that sequents comprise closed formulas only, we take it as implicit that the variables in the statement are considered as fresh constants.

▶ **Proposition 4.7.** *Let (D, σ) be a cyclic derivation as per Lemma 4.6 and fix a vertex v. Let $C_{\lhd v} = \{w \in \mathrm{comp}(D) \mid w \lhd_D v\}$ and $C_{\geq v} = \{w \in \mathrm{comp}(D) \mid v \leq w\}$. There exist finite derivations of:*

1. *$\Gamma_v \Rightarrow \iota_v$ from initial sequents $\{\Gamma_w \Rightarrow x_w \mid w \in C_{\lhd v}\}$ and $\{\Gamma_w \Rightarrow \hat{\iota}_w \mid w \in C_{\geq v}\}$, wherein buds of the latter kind each have a companion.*
2. *$\iota_v \Rightarrow \Delta_v$ from initial sequents $\{x_w \Rightarrow \Delta_w \mid w \in C_{\lhd v}\}$ and $\{\hat{\iota}_w \Rightarrow \Delta_w \mid w \in C_{\geq v}\}$, wherein buds of the latter kind each have a companion.*

Proof. By induction on v. If v is a bud, then $v^c \in C_{\lhd v}$ and the claim is trivial. If v is an identity sequent, so are $\Gamma_v \Rightarrow \iota_v$ and $\iota_v \Rightarrow \Delta_v$. Suppose that $v \in \mathrm{comp}(D)$ and let v^+ be the immediate successor of v. We treat case 1. By the induction hypothesis, there is a finite derivation of $\Gamma_{v^+} \Rightarrow \iota_{v^+}$ from initial sequents

$$\{\Gamma_w \Rightarrow x_w \mid w \in C_{\lhd v^+}\} \cup \{\Gamma_w \Rightarrow \iota_w \mid w \in C_{\geq v^+}\}.$$

Note, $C_{\lhd v^+} = C_{\lhd v} \cup \{v\}$. Uniformly substituting $\iota_v = \sigma_v x_v \, \iota_{v^+}$ for x_v in this derivation yields a finite derivation of $\Gamma_v \Rightarrow \iota_{v^+}(\iota_v/x_v)$ from initial sequents $\{\Gamma_w \Rightarrow w \mid w \in C_{\lhd v}\}$ and $\{\Gamma_w \Rightarrow \iota_w \mid w \in C_{\geq v}\}$. Appending an instance of the rule $\sigma_v\mathsf{R}$ to the root of this derivation followed by the companion rule c yields a derivation of $\Gamma_v \Rightarrow \iota_v$ from the desired initial sequents. The remaining cases are straightforward. ◀

4.2 Lyndon Interpolation and Applications

Considering derivations with no buds, Lemmas 4.1 and 4.6 imply our first interpolation result for the modal μ-calculus:

▶ **Theorem 4.8** (Quantifier-free Lyndon Interpolation). *Suppose* $\vdash \alpha \Rightarrow \beta$ *with* $\alpha \in \Pi_1$ *and* $\beta \in \Sigma_1$. *There exists a quantifier-free formula* ι *such that*
1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$.
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.
3. $|\iota| \leq 2^{|\mathsf{FL}(\alpha)|+|\mathsf{FL}(\beta)|+1}$.

Proof. Let D be a cyclic proof witnessing $\vdash \alpha \Rightarrow \beta$. This proof, by virtue of the assumption on α and β, contains no buds and, hence, we may assume that it has no companions. Let $\iota = \iota_{(D,\sigma)}$ where σ is the empty marking. Lemma 4.2 confirms that ι is quantifier-free and satisfies the desired restriction on signatures. Lemma 4.6 describes cyclic derivations of the sequents $\alpha \Rightarrow \iota$ and $\iota \Rightarrow \beta$. As D contains no companions, the derivations are bud-free (finite) proofs. The bound is based on size of derivation is a consequence of Corollary 3.14 and Lemma 4.1. ◀

As the quantifier-free formulas are both Π_1 and Σ_1, Theorem 4.8 yields

▶ **Corollary 4.9.** *Modal logic* K *has the Lyndon interpolation property.*

Reading Theorem 4.8 as a separation result yields

▶ **Corollary 4.10** (Π_1-separation). *Let* $\alpha, \beta \in \Pi_1$. *If* $\langle\!|\alpha|\!\rangle \cap \langle\!|\beta|\!\rangle = \emptyset$ *then there exists a quantifier-free* ι *such that* $\langle\!|\alpha|\!\rangle \subset \langle\!|\iota|\!\rangle$ *and* $\langle\!|\iota|\!\rangle \cap \langle\!|\beta|\!\rangle = \emptyset$.

Proof. Let $\alpha, \beta \subset \Pi_1$ be such that $\langle\!|\alpha|\!\rangle \cap \langle\!|\beta|\!\rangle = \emptyset$. Then $\beta^\perp \in \Sigma_1$ and $\langle\!|\alpha|\!\rangle \subset \langle\!|\beta^\perp|\!\rangle$ by Corollary 2.12. By completeness, $\vdash \alpha \Rightarrow \beta^\perp$ and Theorem 4.8 yields a quantifier-free formula ι fulfilling the requirements. ◀

A special case of the previous corollary is:

▶ **Corollary 4.11** (Δ_1-collapse). $\Delta_1 = \Delta_0$. *That is, every* Δ_1*-formula is equivalent to a quantifier-free formula.*

Proof. By Corollary 4.10: Suppose $\alpha \in \Pi_1$ and $\alpha^\perp \equiv \beta \in \Pi_1$. Since $\langle\!|\alpha|\!\rangle \cap \langle\!|\alpha^\perp|\!\rangle = \emptyset$ there is a formula $\iota \in \Delta_0$ such that $\langle\!|\alpha|\!\rangle \subset \langle\!|\iota|\!\rangle \subset \langle\!|\alpha|\!\rangle$. ◀

The next results treat the other scenarios of interpolating sequents at the first level.

▶ **Theorem 4.12** (Σ_1-Lyndon Interpolation). *Suppose* $\vdash \alpha \Rightarrow \beta$ *with* $\alpha, \beta \in \Sigma_1$. *There exists a* Σ_1 *formula* ι *such that*
1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$.
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

Proof. Let (D, σ) be a marked cyclic proof of $\vdash \alpha \Rightarrow \beta$ where $\sigma_v = \mu$ for all companions. Let $\iota = \iota_{(D,\sigma)}$ be the candidate interpolant and let D_1 and D_2 be the cyclic derivations of $\alpha \Rightarrow \iota$ and $\iota \Rightarrow \beta$ respectively given by Lemma 4.6. We observe that these derivations are cyclic proofs. A bud $v \in \mathrm{bud}(D_1)$ will be labelled by a sequent $\Gamma_w \Rightarrow \hat{\iota}_w$ where $w \in \mathrm{comp}(D)$. As D is a proof but $\beta \in \Sigma_1$, Γ_w contains a genuine Σ_1 formula. So v is left successful. Similarly, every bud v of D_2 is left successful as, by construction, $\hat{\iota}_w$ is genuine Σ_1 for $w \in \mathrm{comp}(D)$. ◀

▶ **Corollary 4.13** (Π_1-Lyndon Interpolation). *Suppose $\vdash \alpha \Rightarrow \beta$ with $\alpha, \beta \in \Pi_1$. There exists a Π_1 formula ι such that*
1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$.
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

One configuration not covered by the above results is for sequents $\alpha \Rightarrow \beta$ where $\alpha \in \Sigma_1$ and $\beta \in \Pi_1$. From Theorem 3.13 we know that such a sequent is valid iff it admits a cyclic proof. The difficulty with this scenario, however, is that the constructed interpolant will not in general be Σ_1 or Π_1. This happens in the case of a derivation D with three companions $u < v_1, v_2$ where the left cedent at v_1 is empty and the right cedent at v_2 is empty. For the construction above to work, the chosen marking would need to send v_1 to ν and v_2 to μ, resulting in a candidate interpolant for u in neither Σ_1 nor Π_1. However, a simple modification of the construction achieves the aim.

▶ **Theorem 4.14.** *Suppose $\vdash \alpha \Rightarrow \beta$ with $\alpha \in \Sigma_1$ and $\beta \in \Pi_1$. There exists a Σ_1 formula ι such that*
1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$.
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

Proof. Let ι be the Lyndon candidate induced by D with the trivial marking $\sigma\colon u \mapsto \mu$ but with the alteration that a vertex with empty left cedent is assigned the candidate $\iota_v = \top$. ◀

5 Uniform Interpolation at the First Level

The method of constructing Lyndon interpolants can be adapted to generate *uniform* interpolants based on observations made by Pitts [32] (see also [17] and [18]). The idea is to explore, via proof-search, all ways in which a given formula α can influence a proof of a conditional sequent of the form $\alpha \Rightarrow \beta$ and encode this information as a single formula of L_μ. The evolution of left cedents in cyclic derivations is captured precisely by the template calculus of Section 3.1. This section demonstrates the encoding of a template T for a formula $\alpha \in \Sigma_1 \cup \Pi_1$ as a formula $\iota \in \Pi_1$ which uniformly interpolates α against all Π_1 formulas. The general uniform interpolation property is addressed in Section 6.

The template calculus, recall, is the sequent calculus of structural and left logical rules of $\mathsf{K}\mu$ with the modal rule $\Diamond\mathsf{L}$ replaced by a branching form:

$$\frac{\gamma_1, \Delta \Rightarrow \quad \cdots \quad \gamma_n, \Delta \Rightarrow \qquad \Delta \Rightarrow}{\Diamond\gamma_1, \ldots, \Diamond\gamma_n, \Box\Delta, L \Rightarrow}\ \Diamond^*\mathsf{L}$$

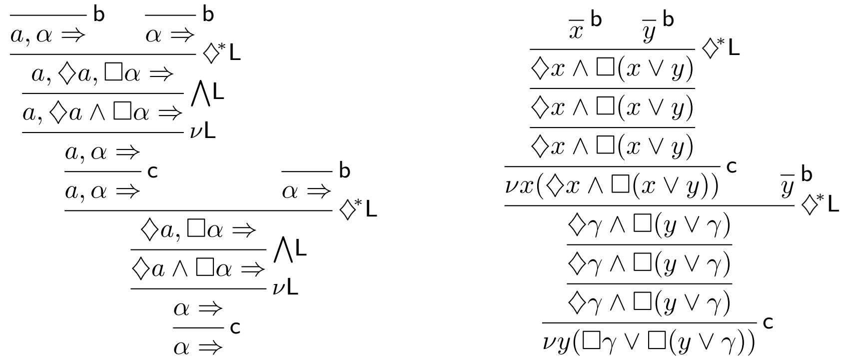

Figure 7 An α-template where $\alpha = \nu x(\Diamond a \wedge \Box x)$ and the uniform interpolant relative to Π_1 formulas and the empty signature; $\gamma = \nu x(\Diamond x \wedge \Box(x \vee y))$.

and subject to the constraint that L is a consistent set of literals. The formulas $\Diamond\gamma_1, \ldots, \Diamond\gamma_n$ are considered the principal formulas of the rule and the corresponding premises $\gamma_i, \Delta \Rightarrow$ are the *major* premises. The final premise $\Delta \Rightarrow$ with no active formula is the *minor* premise of the rule. A cyclic derivation in the template calculus is called a *template*. We recall also the concept of a *pruning* of a template T, which is any subtree of T given by removing all but one premise of each instance of $\Diamond^*\mathsf{L}$ in T and any vertex above a removed premise.

Unlike our use of the template calculus in Section 3, the minor premise plays an important role in what follows. We view a template for a cedent Γ as a record of the (potential) dependence on Γ of an arbitrary cyclic proof with endsequent $\Gamma \Rightarrow \Delta$. In the simplest scenario a cyclic proof D of $\Gamma \Rightarrow \Delta$ can be seen as carving out a pruning of the fixed template T. Instances of $\Diamond\mathsf{L}$ in D pick out a major premise of $\Diamond^*\mathsf{L}$ and the rule $\Box\mathsf{R}$ corresponds to the minor premise. In this light, the proof-search argument of Proposition 3.11 corresponds to the specific case $\Delta = \emptyset$.

Templates can vary in two essential ways. Obviously, two templates may employ different instances of logical rules to derive the same sequent. They could, however, always agree on the choice and instance of logical rules yet differ in respect to which vertices/sequents are buds or which, of potentially many, repetitions are allocated as companions. In other words, two templates could have identical unravellings yet differ as finite derivations.

Changing the order of non-modal rules within a template will never change the instances of $\Diamond^*\mathsf{L}$. As the structure of the logical rules does not affect the success condition, we can readily assume that any two templates for the same cedent agree modulo the selection of buds and companions. In other words, we can implicitly assume that all templates unravel to the same infinite derivation in the template calculus.

A similar observation holds for cyclic derivations in $\mathsf{K}\mu$. The order of non-modal rules is irrelevant, and we can assume that instances of a modal rule always have a maximally decomposed conclusion (each cedent having the form $L, \Box\Gamma, \Diamond\Delta$ for $L \subset \mathrm{Lit}$ consistent) and maximal choice of side formulas, though the choice of principal formula is permitted to vary between instances. These assumptions allow a cyclic derivation in $\mathsf{K}\mu$ to be related to a pruning of the template along the lines of Proposition 3.11. The statement of the connection relies on one further operation on derivations. For a detailed account, we refer the reader to [4]; a different approach to the same problem is presented in [5].

The *(left) projection* of a $\mathsf{K}\mu$-derivation D is the tree D_{L} in which the right cedent at every vertex has been removed along with any vertex that is the conclusion of a non-modal right rule. We tacitly assume that the root of D is a left rule (otherwise D_{L} may not be a tree), but such an assumption is harmless.

▶ **Assumption 1** (Canonicity of derivations)**.** *For any Γ-template T and bud-free derivation D with endsequent $\Gamma' \Rightarrow \Delta$ where $\Gamma' \subseteq \Gamma$, D_{L} is a pruning of the unravelling of T modulo the removal of some non-modal and structural rules from the latter.*

5.1 Uniform Candidates

The construction of a candidate uniform interpolant from a template follows the same procedure as that of Lyndon interpolants, with the only alteration in the treatment of the branching modal rule. A *marked template* is a pair (T, σ) where T is a template and σ, called the *marking*, is a function from the set of companions of T into $\{\mu, \nu\}$.

▶ **Construction 2.** *Fix a signature S. Let (T, σ) be a marked cyclic derivation and let $v \mapsto x_v \colon \mathrm{comp}(D) \to \mathrm{Var}$ be an injection of companion vertices in D into variables. We write σ_v for $\sigma(v)$. Let Φ_v be the left cedent at vertex $v \in T$. Each vertex $v \in T$ is associated a formula ι_v defined by recursion on T, with the first applicable case taking precedence:*

Empty sequent: *If $\Phi_v = \emptyset$ then $\iota_v = \top$.*

Bud: *$\iota_v = x_{v^{\mathsf{c}}}$, namely the variable associated to the bud's companion.*

Identity: *$\iota_v = \bot$*

Modal Rule: *Let $v_0, \ldots, v_n$ enumerate the immediate successors of v with the minor premise last. Define*

$$\iota_v = \bigwedge\Gamma \text{ where } \Gamma = (\Phi_v \cap S) \cup \left\{\Diamond\iota_{v_i} \;\middle|\; i < n\right\} \cup \left\{\Box(\iota_0 \vee \cdots \vee \iota_n)\right\}.$$

Logical Rule $\bigvee$L: *$\iota_v = \bigvee\{\iota_w \mid w \in \mathrm{Suc}_T(v)\}$.*

Any other logical rule: *$\iota_v = \iota_{v^+}$.*

Companion: *$\iota_v = \sigma_v x_v\, \iota_{v^+}$.*

The uniform candidate induced by (T, σ, S) *is the formula $\iota_{(T,\sigma,S)} = \iota_r$ where r is the root of T.* ⌟

▶ **Example 5.1.** Figure 7 presents an example of Construction 2 to a template for the Π_1 formula $\nu x(\Diamond a \wedge \Box x)$. ⌟

The verification that certain templates induce uniform interpolants boils down to properties of the marking of T. Given a marked α-template (T, σ), a cyclic derivation of $\alpha \Rightarrow \iota_{(T,\sigma,S)}$ is readily available, given by inserting into T a substituted form of the formula ι_v as the right cedent at each vertex v. The derivation will be a cyclic proof provided that sufficient companions of T are assigned the quantifier ν (as these induce right successful buds). Furthermore, the template induces a transformation on cyclic derivation of sequents $\alpha \Rightarrow \beta$ into derivations of $\iota_{(T,\sigma,S)} \Rightarrow \beta$ by a similar process, subject to constraints on the signature of β. That the transformation preserves the proofhood condition depends on sufficiently many companions being marked μ (inducing left successful buds). Uniform interpolation relies on finding templates whose marking satisfy these competing conditions.

As in the case of Lyndon candidates, uniform candidates mirror the structure of the derivation – the template – from which they are constructed. If the marking assigns ν to all companions, for instance, the constructed candidate will be Π_1. Indeed, Lemmas 4.1 to 4.4 carry over to the realm of uniform candidates without alteration.

The substitutional instances of the candidate formulas are defined in the same way as for candidate Lyndon interpolants. Let (T, σ) be a marked template. The notation Φ_u is generally used for the left cedent at vertex $u \in T$.

▶ **Definition 5.2** (Uniform candidate at a vertex). *Let ι_v be the formula associated to vertex v by Construction 2. Let $C_{\lhd v} = \{w \in \mathrm{comp}(T) \mid w \lhd_T v\}$ be the companions strictly below $v \in T$ relative to the relation $\lhd_T$. For each $v \in T$, the* uniform candidate at v *is the formula*

$$\hat{\iota}_v = \iota_v(\iota_{c_1}/x_{c_1}) \cdots (\iota_{c_n}/x_{c_n}) \quad \textit{where } c_1 < \cdots < c_n < v \textit{ enumerate the elements of } C_{\lhd v}.$$

The reader can confirm that $\hat{\iota}_r = \iota_r = \iota_{(T,\sigma,S)}$ where r is the root of D and that $\hat{\iota}_v$ is a closed formula for every v.

We begin with correspondents to the Interpolation Lemma 4.6.

▶ **Lemma 5.3** (First Template Lemma). *Let (T, σ) be a marked template for Γ and let Φ_u be the left cedent at $v \in T$. For every signature S there is a cyclic derivation of $\Gamma \Rightarrow \iota_{(T,\sigma,S)}$ in* $\mathsf{K}\mu + \mathsf{w}$ *with buds labelled by sequents $\Phi_u \Rightarrow \hat{\iota}_u$ for $u \in \mathrm{bud}(T)$.*

Proof. The cyclic derivation of $\Gamma \Rightarrow \iota_{(T,\sigma,S)}$ is constructed by recursion through the template. Let T' be the template T in which the sequent at each vertex $v \in T$ is replaced by $\Phi_v \Rightarrow \hat{\iota}_v$. The root of T' is labelled $\Gamma \Rightarrow \iota_{(T,\sigma,S)}$. To turn T' into a cyclic derivation, we maintain the same bud-companion association and merely insert sufficient right rules between the vertices of T' and replace occurrences of $\Diamond^*\mathsf{L}$ by finite derivations. Thus, a bud/companion $\Phi_b \Rightarrow$ in T becomes a bud/companion $\Phi_b \Rightarrow \hat{\iota}_b$ in T' and an identity sequent $\Phi_v \Rightarrow$ becomes $\Phi_v \Rightarrow \bot$. If $\Phi_v = \emptyset$, then v becomes the provable sequent $\Phi_v \Rightarrow \top$ and the vertices above v are discarded. The transformation in the case of logical rules is straightforward and omitted. This leaves only the modal rule $\Diamond^*\mathsf{L}$ and companions to consider. Suppose the modal rule occurs at a vertex v. Let $v_0, \ldots, v_n$ enumerate the immediate successors of v with the minor premise last. Let $L = \Phi_v \cap S$. The construction defines

$$\hat{\iota}_v = \bigwedge\Gamma \text{ where } \Gamma = L \cup \left\{\Diamond\hat{\iota}_{v_i} \;\middle|\; i < n\right\} \cup \left\{\Box(\hat{\iota}_{v_0} \vee \cdots \vee \hat{\iota}_{v_n})\right\}.$$

The vertex v is expanded into the derivation

$$\dfrac{\dfrac{}{\Phi_v \Rightarrow l}\,\mathsf{id} \quad \text{for } l \in L \qquad \dfrac{\Phi_{v_i} \Rightarrow \hat{\iota}_{v_i}}{\Phi_v \Rightarrow \Diamond\hat{\iota}_{v_i}}\,\Diamond\mathsf{L} \quad \text{for } i < n \qquad \dfrac{\dfrac{\Phi_{v_n} \Rightarrow \hat{\iota}_{v_n}}{\Phi_{v_n} \Rightarrow \hat{\iota}_{v_0} \vee \cdots \vee \hat{\iota}_{v_n}}\,\mathsf{w},\bigvee\mathsf{R}}{\Phi_v \Rightarrow \Box(\hat{\iota}_{v_0} \vee \cdots \vee \hat{\iota}_{v_n})}\,\Box\mathsf{R}}{\Phi_v \Rightarrow \hat{\iota}_v}\,\bigwedge\mathsf{R}$$

The remaining case is the companion. Then $\hat{\iota}_v = \sigma_v x_v\, \beta$ for some formula β, and $\hat{\iota}_{v^+} = \beta(\hat{\iota}_v/x_v)$. Inserting an instance of $\sigma_v\mathsf{R}$ above the companion suffices. As bud and companions are associated the same sequent, a cyclic derivation has been constructed. ◀

▶ **Lemma 5.4** (Second Template Lemma)**.** *For any signature S and formula Δ with $\mathrm{sig}(\Delta) \subset S$, given a cyclic derivation D of $\alpha \Rightarrow \Delta$ there exists a cyclic derivation of $\iota_{(T,\sigma,S)} \Rightarrow \Delta$ whose buds (and companions) are sequents of the form $\hat{\iota}_u \Rightarrow \Theta$ such that $\Phi_u \Rightarrow \Theta$ is a bud of D.*

Proof. This argument closely follows that of the Lyndon Interpolation Lemma 4.6. The main complication is the treatment of modal rules.

Let D be a cyclic derivation of $\alpha \Rightarrow \Delta$ and consider its unravelling to a bud-free derivation which we name D_ω. Let $\Phi_u \Rightarrow$ be the sequent at $u \in T$ and $\Gamma_u \Rightarrow \Delta_u$ the sequent at $u \in D_\omega$. The idea is to replace the left cedent in each vertex of D_ω by an appropriately chosen candidate and close off the infinite branches of D_ω to obtain a cyclic derivation. Assumption 1 ensures that the left projection of D_ω aligns with those occurring in the template.

We begin by associating to each vertex $u \in D_\omega$ a set $V_u \subset T$ of template vertices which identifies the candidate interpolants $I_u = \{\hat{\iota}_u \mid u \in V_u\}$ that can act in the place of the left cedent of u. The assignment of template vertices to D_ω is the naive choice determined by compatibility of left cedents: $V_u = \{v \in T \mid \Gamma_u \subset \Phi_v\}$. We now show how to build a bud-free derivation of $\hat{\iota}_v \Rightarrow \Delta_u$ for each $v \in V_u$. Let $S = \{\hat{\iota}_v \Rightarrow \Delta_u \mid v \in V_u\}$ be the collection of such sequents. The construction proceeds co-recursively in the sense that we associate to each sequent $\hat{\iota}_v \Rightarrow \Delta_u$ from S a finite derivation $D_{u,v}$ whose buds are sequents in S. Patching together sufficiently many such derivations will yield a cyclic derivation of the desired form.

Let $v \in V_u$. There are four distinct cases to consider based on the rule at $u \in D_\omega$.

Non-modal right rule. These rules do not alter the left cedent and we simply maintain the right rule employed. If, for example, the rule at u is $\bigwedge\mathsf{R}$ with principal formula $\bigwedge\{\delta_1, \ldots, \delta_n\} \in \Delta_u$, then we have $v \in V_{u'}$ for each successor $u' \in \mathrm{Suc}_D(u)$ and $D_{u,v}$ is the following finite derivation for appropriate choice of Δ.

$$\left.\dfrac{\dfrac{}{\hat{\iota}_v \Rightarrow \delta_1, \Delta}\,\mathsf{b} \qquad \cdots \qquad \dfrac{}{\hat{\iota}_v \Rightarrow \delta_n, \Delta}\,\mathsf{b}}{\hat{\iota}_v \Rightarrow \Delta_u}\,\bigvee\mathsf{L}\right\} = D_{u,v}$$

Non-modal left rule. Let $\gamma \in \Gamma_u$ be the principal formula at u. If this is also the principal formula at $v \in T$ then $D_{u,v}$ is formed analogously to the previous case as the single left rule with corresponding buds. For instance, in the case that $\gamma = \bigvee\{\gamma_1, \ldots, \gamma_n\}$ then $\hat{\iota}_v = \bigvee\{\hat{\iota}_w \mid w \in \mathrm{Suc}_T(v)\}$ and

$$\left.\dfrac{\dfrac{}{\hat{\iota}_{v_1} \Rightarrow \Delta_u}\,\mathsf{b} \qquad \cdots \qquad \dfrac{}{\hat{\iota}_{v_n} \Rightarrow \Delta_u}\,\mathsf{b}}{\hat{\iota}_v \Rightarrow \Delta_u}\,\textstyle\bigwedge\mathsf{R}\right\} = D_{u,v}$$

The buds each correspond to a sequent in S. However, it may be that γ is not principal at v. First, we may assume that v is not a bud of T as the companion is labelled by the same sequent. Assumption 1 ensures that Γ_u is a subset of each cedent above v until a point that γ is principal. Moreover, only (left) logical rules occur in this portion of the template. So we let $D_{u,v}$ be this finite derivation mimicking the rules that alter the candidate interpolants ending at the point that γ is decomposed.

Identity sequent. The left or right identity rules are covered in the previous cases. So assume that the cedents Δ_u and Φ_v contain a literal l in common. As $\Delta_u \subset \mathsf{FL}(\Delta)$, it follows that $l \in S$. By the previous case we can assume that v is a modal rule, so $\hat{\iota}_u = \bigwedge\Gamma$ with $l \in \Gamma$ and $\hat{\iota}_u \Rightarrow \Delta_u$ has a straightforward (finite) proof.

The final case is that u is a modal rule. As above, we may assume that v is also the conclusion of a modal rule in T, whence

$$\hat{\iota}_v = \textstyle\bigwedge\Gamma \text{ where } \Gamma = L \cup \left\{\Diamond\iota_{v_i} \;\middle|\; i < n\right\} \cup \left\{\Box(\iota_{v_0} \vee \cdots \vee \iota_{v_n})\right\}$$

where $v_0, \ldots, v_n$ enumerate the successors of v (in T) with v_n the minor premise and L is some set of literals. If the rule at u is the left modal rule ($\Diamond\mathsf{L}$), this rule is carried over to $D_{u,v}$ with the premise u^+ being matched to the major premise v_i such that $\Gamma_{u^+} \subset \Phi_{v_i}$. If, however, the rule at u is the right modal rule ($\Box\mathsf{R}$), then we appeal to the fact that $\Gamma_{u^+} \subset \Phi_{v_i}$ for all $i \leq n$ to form the disjunction $\bigvee\{\hat{\iota}_{v_0}, \ldots, \hat{\iota}_{v_n}\}$ before introducing the modal operator. The two cases each induce one of the following derivations for $D_{u,v}$:

$$\dfrac{\dfrac{\dfrac{}{\hat{\iota}_{v_i} \Rightarrow \Delta_{u^+}}\,\mathsf{b}}{\ldots, \Diamond\hat{\iota}_{v_i}, \ldots \Rightarrow \Delta_u}\,\Diamond\mathsf{L}}{\hat{\iota}_v \Rightarrow \Delta_u}\,\textstyle\bigwedge\mathsf{L} \qquad \dfrac{\dfrac{\dfrac{\dfrac{}{\hat{\iota}_{v_0} \Rightarrow \Delta_{u^+}}\,\mathsf{b} \quad \cdots \quad \dfrac{}{\hat{\iota}_{v_n} \Rightarrow \Delta_{u^+}}\,\mathsf{b}}{\hat{\iota}_{v_0} \vee \cdots \vee \hat{\iota}_{v_n} \Rightarrow \Delta_{u^+}}\,\textstyle\bigvee\mathsf{L}}{\ldots, \Box(\hat{\iota}_{v_0} \vee \cdots \vee \hat{\iota}_{v_n}) \Rightarrow \Delta_u}\,\Box\mathsf{R}}{\hat{\iota}_v \Rightarrow \Delta_u}\,\textstyle\bigwedge\mathsf{L}$$

With all cases of $D_{u,v}$ defined, let E be the derivation with root $\iota_{(T,S)} \Rightarrow \Delta$ given by repeatedly pasting at every bud the next applicable partial derivation $D_{u,v}$. This derivation has preserved all companion rules from D_ω. Cutting off each branch at the first repeated companion completes the construction. ◀

5.2 Uniform Interpolation and Applications

From the two template lemmas, we obtain our first uniform interpolation result.

▶ **Theorem 5.5** (Π_1 Uniform Interpolation). *For every signature S and formula $\alpha \in \Pi_1 \cup \Sigma_1$ there exists $\iota \in \Pi_1$ satisfying*

1. $\mathrm{sig}(\iota) \subset S \cap \mathrm{sig}(\alpha)$ *and* $|\iota| \leq 2^{|\mathsf{FL}(\alpha)|+1}$.
2. $\vdash \alpha \Rightarrow \iota$.
3. *for all $\beta \in \Pi_1$ with $\mathrm{sig}(\beta) \subset S$, if $\vdash \alpha \Rightarrow \beta$ then $\vdash \iota \Rightarrow \beta$.*

As we are targeting a Π_1 interpolant it is natural to consider the candidate $\iota_{(T,\sigma,S)}$ for an α-template in which every companion is marked ν. This choice guarantees condition 2 by Lemma 5.3 as right cedents in the generated cyclic derivation always comprise Π_1 formulas. This candidate, however, will not necessarily satisfy the third condition. It will fail, for example, if $\alpha = \mu x \Diamond x$: The natural template for α induces the Π_1 candidate $\iota = \nu x(\Diamond x \wedge \Box(\alpha \vee \top))$, which is equivalent to $\nu x \Diamond x$. Letting $\beta = \bot$, we have $\vdash \alpha \Rightarrow \beta$ but not $\vdash \iota \Rightarrow \beta$ as the latter sequent is not valid. The construction fails in this case because the sequent α is unsatisfiable, yet ι is satisfiable. Selecting $\bot$ for ι, indeed, provides a trivial uniform interpolant.

The solution is to 'correct' the candidate interpolant by pre-emptively setting the subformula assigned to $v \in T$ to be $\bot$ if $\Phi_v \Rightarrow$ is provable. Recall that provability can be checked from within the template alone by Theorem 3.11. We did not incorporate this step into Construction 2 because it is only necessary for constraining the uniform candidate to a particular quantifier class.

Proof. Fix a signature S and formula $\alpha \in \Pi_1 \cup \Sigma_1$. Let T be an α-template in which buds are the first repeat on each path and let $\sigma\colon \mathrm{comp}(T) \to \{\nu\}$ be the trivial marking assigning ν to all companions. By construction, $\iota_{(T,\sigma,S)} \in \Pi_1$ and $|\iota_{(T,\sigma,S)}| \leq 2^{|\mathsf{FL}(\alpha)|+1}$. Let Φ_u be the left cedent of $u \in T$ and $\hat{\iota}_u$ be the uniform candidate assigned to $u \in T$. We choose ι to be the modification of $\iota_{(T,\sigma,S)}$ given by replacing any subformula ι_u such that $\Phi_u \Rightarrow$ is provable by the formula $\bot$. This change can be carried out by inspection of the template alone (cf. Theorem 3.11). Clearly, $\iota \in \Pi_1$ and $\mathrm{sig}(\iota) \subset S \cap \mathrm{sig}(\alpha)$ and $|\iota| \leq |\iota_{(T,\sigma,S)}|$.

The First Template Lemma provides a cyclic derivation of $\alpha \Rightarrow \iota_{(T,\sigma,S)}$ in which all buds are sequents of the form $\Phi_u \Rightarrow \hat{\iota}_u$ for $u \in \mathrm{bud}(T)$. As each such formula $\hat{\iota}_u$ is (genuine) Π_1 all buds are right successful. The corresponding cyclic derivation of $\alpha \Rightarrow \iota$ is also a proof because if a bud is (now) labelled by $\Phi_u \Rightarrow \bot$ it is because $\Phi_u \Rightarrow$ is provable. Thus, $\vdash \alpha \Rightarrow \iota$.

For the final condition, fix $\beta \in \Pi_1$ and suppose D is a cyclic proof of $\alpha \Rightarrow \beta$. Consider the cyclic derivation of $\iota_{(T,\sigma,S)} \Rightarrow \beta$ described by the Second Template Lemma. A corresponding cyclic derivation D' of $\iota \Rightarrow \beta$ is readily obtained by simply closing every sequent with $\bot$ in the left cedent as an initial sequent. We claim that D' is a cyclic proof. Let $u \in \mathrm{bud}(D')$ be a bud of D' which is not right successful. As β is Π_1, the right cedent of u must be empty. Let $v \leq u$ be the first vertex on the path to u with empty right cedent. This vertex corresponds to a vertex $v' \in D$ of the form $\Gamma' \Rightarrow$. As D is a proof, the definition of ι means that the sequent at v in D' happens to be $\bot \Rightarrow$. ◀

Because the definition of signature takes into account polarities of variables, Theorem 5.5 actually establishes uniform *Lyndon* interpolation for the class of Σ_1 formulas:

▶ **Corollary 5.6.** *The Π_1 fragment of the modal μ-calculus has the right uniform Lyndon interpolation property. By duality, the Σ_1 fragment has the left uniform Lyndon interpolation property.*

▶ **Corollary 5.7.** *Every $\alpha \in \Sigma_1$ can be computably associated a formula $\alpha_\Pi \in \Pi_1$ with $|\alpha_\Pi| \leq 2^{|\mathsf{FL}(\alpha)|+1}$ such that $\alpha \in \Delta_1$ iff $\alpha \equiv \alpha_\Pi$.*

Proof. Fix $\alpha \in \Sigma_1$ and let $S = \text{Lit}$ be the universal signature. Set α_Π to be the uniform interpolant for α given by Theorem 5.5. In particular, $\langle\!|\alpha|\!\rangle \subset \langle\!|\alpha_\Pi|\!\rangle$ by soundness. Since $\alpha_\Pi \in \Pi_1$, the right to left direction is immediate. So suppose $\alpha \equiv \beta$ for $\beta \in \Pi_1$. Then $\vdash \alpha \Rightarrow \beta$ by completeness and $\vdash \alpha_\Pi \Rightarrow \beta$ by Theorem 5.5. So $\alpha \equiv \alpha_\Pi$. ◀

In the previous section we confirmed, via the construction of Lyndon interpolants, that every Δ_1 formula is equivalent to a quantifier-free formula (Corollary 4.11). Combining that observation with the characterisation of Δ_1 above yields a decidability result first observed by Martin Otto [31].

▶ **Corollary 5.8** (Otto's Theorem for Σ_1). *The problem of determining whether an arbitrary Σ_1-formula is equivalent to a quantifier-free formula is decidable.*

We will obtain the full version of Otto's Theorem, lifting the restriction to Σ_1, as a corollary of interpolation for formulas of arbitrary complexity, treated in the next section.

6 Cyclic Proofs and Interpolation

We present a general notion of cyclic proof that is sound and complete for all sequents in L_μ. The point of departure from the account in Definition 3 is that notion of proof can longer be expressed by a local condition on bud/companion pairs. Rather, proofhood in the general case is a *global* property determined by a condition on the *infinite* paths through a cyclic derivation.

The next definition expresses the idea of following the (potentially infinite) decomposition of a formula through a cyclic derivation.

▶ **Definition 6.1.** *Let D be a cyclic derivation. A* left *(resp.* right*)* occurrence *in D is a pair $f = (\alpha, v)$ where v is a vertex of D and α is a formula in the left (right) cedent at v. A* left trace *in D is an infinite sequence $t = (\alpha_i, v_i)_{i<\omega}$ of left occurrences such that for all $i < \omega$,*

1. *$(v_i)_i$ is a path through D,*
2. *if α_i is principal at v then α_{i+1} is an active formula at v_{i+1},*
3. *if α_i is not principal and v_i is not a modal rule then $\alpha_{i+1} = \alpha_i$,*
4. *if α_i is not principal and v_i is a modal rule then $\Box\alpha_{i+1} = \alpha_i$.*

A sequence of right occurrences fulfilling the same criteria but with $\Diamond$ in place of $\Box$ in the final clause is called a right trace*. The path $(v_i)_i$ is called the* carrier *of t.*

A simple argument, appealing to (weak) König's Lemma, shows that every path through a cyclic derivation is the carrier of at least one (left or right) trace. A cyclic derivation, recall, must contain a modal rule between every bud-companion pair. As a consequence, the formulas along a trace never stabilise:

▶ **Lemma 6.2.** *Let $t = (\alpha_i, v_i)_{i<\omega}$ be a left or right trace. The sequence of formulas $(\alpha_i)_i$ is a progressing thread.*

We can now isolate the cyclic proofs.

▶ **Definition 6.3** (Good trace). *Let t be a (left or right) trace through a cyclic derivation and let τ be the progressing thread carried by t. We call t* good *if it is a left trace and τ is μ or a right trace and τ is ν.*

▶ **Definition 6.4** (Cyclic proof). *A cyclic derivation D is a* cyclic proof *iff every infinite path through P is the carrier of a good trace.*

Definition 6.4 agrees with the previous notion of cyclic proofs for endsequents in the first quantifier level: An infinite path through a cyclic proof according to Definition 3.5 must be such that either every vertex along it is left-successful, or every vertex is right-successful. Weak König's Lemma implies that the path carries either a left trace of genuine Σ_1 formulas or a right trace of genuine Π_1 formulas. Both examples are good traces. In the general setting, checking whether a cyclic derivation fulfils the condition of Definition 6.4 is decidable and, for many logics including the modal μ-calculus, is PSPACE-complete [6].

Soundness and completeness of cyclic proofs is established by the same arguments as before observing that the connection between paths through derivations and strategies extends to one between traces and plays. Soundness in particular is a simple refinement of the proof of Proposition 3.9.

▶ **Lemma 6.5** (Duality). $\vdash \Gamma \Rightarrow \Delta$ *iff* $\vdash\ \Rightarrow \Delta, \Gamma^\perp$, *iff* $\Gamma, \Delta^\perp \Rightarrow$.

▶ **Theorem 6.6** (Soundness). $\Gamma \Rightarrow \Delta$ *is valid if* $\vdash \Gamma \Rightarrow \Delta$.

Proof. Suppose D is a cyclic derivation of $\Gamma \Rightarrow$ and S is a winning Verifier strategy in $G(\bigwedge\Gamma, F)$. The proof of Proposition 3.9 determines an infinite path $P = (u_i)_i$ through D such that every left trace carried by P can be associated a play consistent with S that carries the same thread ignoring consecutive repetitions. As S is winning for Verifier, P does not carry a good trace, hence D is not a cyclic proof. ◀

Completeness is not so straightforward. Although the definition of a template applies in the general setting, without an analogue of the *unsuccessful* vertices a lifting of Proposition 3.11 is not immediately available. We return to completeness at the end of Section 6.1.

6.1 A Normal Form for Cyclic Derivations

In order to ensure the candidate constructions from previous sections work at higher levels of the quantifier hierarchy, it is necessary to impose constraints on the structure of cyclic proofs. This is accomplished through a normal-form theorem for cyclic derivations that shows that every cyclic proof can be represented as a marked derivation where the marking encapsulates the existence of good threads. The use of marked derivations in this way first appeared in [20] (conceived of as a 'colouring' of companions). Our use of them is more abstract and motivated by properties of cyclic systems known as *reset systems* developed in [1, 4, 21, 36, 25]. A normal form theorem similar to the one below is implicit in the syntactic proofs of Lyndon and uniform interpolation in [3, 4].

Let P be a path through a cyclic derivation. By $\mathrm{Inf}(P)$ we denote the set of vertices that appear infinitely often along P:

$$\mathrm{Inf}(P) = \{u \in D \mid u = P(i) \text{ for infinitely many } i\}.$$

The set $\mathrm{Inf}(P)$ is non-empty and has a least element, which must be a companion. The least vertex is denoted $\min(P)$.

▶ **Theorem 6.7** (Normal form theorem). *For every cyclic derivation in $\mathsf{K}\mu$ or the template calculus there exists a cyclic derivation D in the same calculus with the same unravelling which admits a marking σ such that any path P through D carries a good left trace iff $\sigma(\min(P)) = \mu$.*

The theorem is due to Janin and Walukiewicz who proved the result for a variation of our sequent calculus $\mathsf{K}\mu$ [20, Lemma 4.5]. Their proof appeals to standard results in automata theory (see, e.g., [19]). From a fixed cyclic derivation (or template) D one obtains, via McNaughton's theorem [29], a deterministic parity stream automaton $\mathcal{A}$ over the alphabet of sequents and rules which accepts an infinite sequence in the alphabet iff the sequences enumerates a path through D and the path carries a good left trace. Running $\mathcal{A}$ against the unravelling of D gives an assignment of states to vertices in D_ω which can be refolded into cyclic derivation with the desired property.

A completely syntactic proof of the normal form theorem can be given by elaborating the sequent calculus for $\mathsf{K}\mu$. Fix a countably infinite set of symbols, called names. Each sequent in D and formula within the sequent is annotated by a finite sequence of names recording the quantifiers encountered in the finite trace from the root to the formula. Aside from the logical rules inherited from D a new structural rule – the *reset* rule – is added which, read bottom-up, discards superfluous information from all annotations in the sequent. Sequent calculi incorporating annotations and a reset mechanism are called *reset systems*. The reset rule is the main component which ensures that the unravelling of D can be annotated by employing only finitely many names and finitely many distinct annotated sequents. As the annotation mechanism is deterministic, the annotated form of D_ω can be realised as the unravelling of a cyclic derivation in the reset calculus which validates Theorem 6.7. A reset calculus for the modal μ-calculus was introduced in [21] for satisfiability and dualised to validity in [36]. The connection to normal forms was realised in [2, Thm IV.10].

The automata- and proof-theoretic approaches to the normal form theorem are closely related. The first reset calculus for the modal μ-calculus was inspired by the structural operations underpinning the Safra construction for determinising stream automata [21]. Alternative reset systems arise from considering different determinisation algorithms for Büchi and parity automata [15]. A general method of associating reset systems to abstract notion of cyclic derivation is explored in [25].

▶ **Definition 6.8.** *A marked derivation which satisfies the property in Theorem 6.7 is said to be in* normal form *or, simply,* normal.

In a normal cyclic derivation, the existence of good left traces on a path is determined by the marking of the lowest companion that occurs infinitely often. If a companion v is marked by μ then *every* path that visits v infinitely often and companions below v only finitely often is the carrier of a good left trace. If v is marked by ν then *no* such path carries a good left trace.

The normal form property does not state anything about *right* traces. Under the knowledge that the derivation is a proof a symmetric condition arises:

▶ **Corollary 6.9.** *Suppose D is a normal cyclic proof. Then any path P through D carries a good right trace iff $\sigma(\min(P)) = \nu$.*

For endsequents from $\Sigma_1 \cup \Pi_1$, the normal form theorem is vacuous. Given a derivation of a sequent $\Gamma \Rightarrow \Delta$ in this fragment, mark a companion u as μ iff the left cedent at u contains a genuine Σ_1 formula. Then (D, σ) is in normal form. Normal derivations at the first level have the further property that the marking is constant along path segments: if $u \lhd_D v$ for $u, v \in \mathrm{comp}(D)$ then $\sigma(u) = \sigma(v)$. Lemma 4.4 shows that this property alone is sufficient to ensure that the Lyndon candidate constructed from D is in the class $\mathit{Clos}(\Sigma_1 \cup \Pi_1)$.

The normal form theorem for templates is sufficient to prove completeness, starting with a lifting of Proposition 3.11:

▶ **Proposition 6.10.** *Let T be a normal template for Γ. If $\Gamma \Rightarrow$ is valid then T admits a pruning to a cyclic proof.*

Proof. Let $\Gamma \Rightarrow$ be any sequent and (T, σ) a template for Γ in normal form. Say that a bud of T is *successful* iff its companion is marked μ, and a non-bud vertex is *successful* if it is either a modal rule and a major premise is successful, or a non-modal rule and all premises are successful. Let U be the set of unsuccessful vertices.

We assume that $\Gamma \Rightarrow$ does not admit a cyclic proof. In particular, no cyclic derivation can be witnessed as a subtree of the successful vertices in T as, by the definition of normal form, any such derivation would be a proof. Hence, U contains the root of T. From U we recursively construct a frame $\langle F, r\rangle$ and Verifier strategy S in $G(\bigwedge\Gamma, F)$ with the property that every infinite play consistent with S corresponds to a left trace through some path of U. By the choice of U, all left traces carry ν-threads, so S is winning for Verifier. ◀

Completeness and decidability is a corollary of Theorem 6.7 and Proposition 6.10:

▶ **Theorem 6.11** (Completeness). *If* $\Gamma \Rightarrow \Delta$ *is valid, then* $\vdash \Gamma \Rightarrow \Delta$.

As the construction of normal forms is computable, Proposition 6.10 implies

▶ **Corollary 6.12.** *Provability in* $\mathsf{K}\mu$ *is decidable.*

6.2 Lyndon Interpolation for the Full Language

The normal form theorem provides a direct path to Lyndon interpolants:

▶ **Theorem 6.13** (Lyndon Interpolation for μ-calculus). *Suppose* $\vdash \alpha \Rightarrow \beta$ *for* $\alpha, \beta \in \mathsf{L}_\mu$. *There exists a formula* ι *such that*

1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$,
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

Proof. Fix a marked normal cyclic proof (D, σ) and let ι be the induced Lyndon candidate following Construction 1. Condition 1 of the theorem holds by definition. We outline the argument for $\vdash \alpha \Rightarrow \iota$, with the other half being symmetric. Let $\Gamma_u \Rightarrow \Delta_u$ be the sequent labelling $u \in D$.

By Lemma 4.6, $\vdash \alpha \Rightarrow \iota$ admits a cyclic derivation D' where buds (and companions) are the sequents $\Gamma_w \Rightarrow \hat{\iota}_w$ for $w \in \mathrm{comp}(D)$. The proof of Proposition 4.7 presents D' as, essentially, the cyclic derivation D with the right cedent at $w \in D$ replaced by the formula $\hat{\iota}_w$. Any path P through D' corresponds to a path through D that carries exactly the same left traces. Let C be the companions of D visited infinitely often by that path. Furthermore, P carries exactly one right trace, whose infinitely occurring quantified formulas are precisely $\{\hat{\iota}_w \mid w \in C\}$. The shortest such formula is, by design, $\hat{\iota}_u$ for $u = \min P$. If P does not carry a good left trace, then $\sigma(u) = \nu$ and $\hat{\iota}_u = \nu x_u \zeta$ for some ζ. But then P carries a good right trace. ◀

We can establish a refinement of the result in case of the second quantifier level. For this we require an observation:

▶ **Proposition 6.14.** *Let* $\Gamma \subset \Pi_2$ *and* $\Delta \subset \Sigma_2$. *If* (D, σ) *is a normal proof of* $\Gamma \Rightarrow \Delta$ *then for all* $u, v \in \mathrm{comp}(D)$, *if* $u \lhd_D v$ *then* $\sigma(u) = \sigma(v)$.

Proof. Let $u, v \in \mathrm{bud}(D)$ such that $u^{\mathsf{c}} < v^{\mathsf{c}} < u$. Suppose $\sigma(u^{\mathsf{c}}) = \nu$. Let k exceed the number of formulas in the left cedent at v and consider the path P which alternates between the bud u and $K + 1$ cycles at v:

$$P := u^{\mathsf{c}} \ldots \underbrace{vv^{\mathsf{c}} \ldots vv^{\mathsf{c}}}_{k} \ldots uu^{\mathsf{c}} \ldots \underbrace{vv^{\mathsf{c}} \ldots vv^{\mathsf{c}}}_{k} \ldots uu^{\mathsf{c}} \ldots$$

By design, $\min P = u^{\mathsf{c}}$. Since $\sigma(u^{\mathsf{c}}) = \nu$, P carries a good right trace. From some point on, this trace never encounters the quantifier μ as the main connective. Also, by the choice of k, the trace must repeat a formula at v before the path returns next to u. As a consequence, the path which cycles between v^{c} and v only also carries a good right trace and $\sigma(v^{\mathsf{c}}) = \nu$ by normal form. The case of $\sigma(u^{\mathsf{c}}) = \mu$ is symmetric. ◀

The proof above breaks if $\Gamma \not\subset \Pi_2$: the good (right) trace carried by P could, for example, encounter a formula $\nu x\, \alpha$ in each path segment $u^{\mathsf{c}} \dots v$ and strictly longer formulas $\mu y\, \beta$ within the 'local' cycles $v^{\mathsf{c}} \dots v$. In that case the marking would satisfy $\sigma(u) = \nu$ but $\sigma(v) = \mu$.

▶ **Theorem 6.15** (Lyndon interpolation at the second level). *Suppose that $\vdash \alpha \Rightarrow \beta$ for $\alpha \in \Pi_2$ and $\beta \in \Sigma_2$. There exists an alternation-free formula ι such that*

1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$,
2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

Proof. By Corollary 6.9 we may assume a normal proof (D, σ) of $\alpha \Rightarrow \beta$. The associated Lyndon candidate $\iota_{(D,\sigma)}$ is, by the proof of Theorem 6.13, a Lyndon interpolant for $\alpha \Rightarrow \beta$. Lemma 4.4 and Proposition 6.14 show that $\iota_{(D,\sigma)} \in \mathit{Clos}(\Pi_1 \cap \Sigma_1)$. ◀

Interpolation for the second level has two interesting corollaries. The first result was shown by Marti and Venema [28] via a novel annotated cyclic calculus for alternation-free fragment.

▶ **Corollary 6.16.** *The alternation-free fragment of L_μ admits Lyndon interpolation.*

The second corollary is a modal rendering of Rabin's definability theorem for weak monadic second order logic [33], originally transferred to the μ-calculus by Arnold and Niwiński [7].

▶ **Corollary 6.17.** *Every Δ_2 formula is equivalent to an alternation-free formula in the same signature. In particular, $\Delta_2 = (\Pi_2 \cap \Sigma_2)^{\equiv}$.*

We conclude our voyage into Lyndon interpolation with two results on the higher ambiguous classes. Both results are due to Arnold and Santocanale [34]. The first is a negative result:

▶ **Theorem 6.18.** *For every $n > 2$, $(\Pi_n \cup \Sigma_n)^{\equiv} \subsetneq \Delta_n$.*

The proof of Theorem 6.18 is beyond the scope of this exposition. The second result from [34] is a refinement of Theorem 6.15 that can be lifted to higher levels (see Corollary 6.21 below). The theorem appeals to an important syntactic fragment of L_μ.

▶ **Definition 6.19** (Disjunctive formula). *A formula α is* disjunctive *if for every conjunction $\bigwedge\Gamma \in \mathsf{FL}(\alpha)$ there exists $L \subset \mathrm{Lit}$ and $\Delta \subset \Theta \subset \mathsf{FL}(\alpha)$ such that $\Gamma = L \cup \Diamond\Delta \cup \Box\bigvee\Theta$. The set of disjunctive formulas is denoted $\mathbb{D}$.*

In disjunctive formulas, the use of conjunction is restricted to specific arrangements of modal statements. The unary modalities can be expressed disjunctively as $\Diamond\alpha \equiv \bigwedge\{\Diamond\alpha, \Diamond\top, \Box\bigvee\{\alpha, \top\}\}$ and $\Box\alpha \equiv \Box\bigvee\{\alpha\}$. The notion was introduced by Janin and Walukiewicz [20] who showed that every formula in the modal μ-calculus is equivalent to a disjunctive formula. Definition 6.19 is more permissive than the standard presentation which is obtained by requiring $\Theta = \Delta$. Our formulation has better proof-theoretic properties [5] and can be reduced to the Janin–Walukiewicz definition by noting that, for $\Delta \subset \Theta$ and $\Delta' = \Delta \cup \{\bigvee(\Theta \setminus \Delta)\}$,

$$\bigwedge\Diamond\Delta \wedge \Box\bigvee\Theta \equiv \left(\bigwedge\Diamond\Delta \wedge \Box\bigvee\Delta\right) \vee \left(\Diamond\Delta' \wedge \Box\bigvee\Delta'\right).$$

We have already seen examples of disjunctive formulas: the uniform candidates in Construction 2. We explore the consequences of that observation in the next section.

▶ **Theorem 6.20** (Lyndon Interpolation for Disjunctive Sequents). *Suppose* $\vdash \alpha \Rightarrow \beta$ *where* $\alpha, \beta^\perp \in \Pi_{n+1} \cap \mathbb{D}$. *There exists* $\iota \in \mathit{Clos}(\Pi_n \cup \Sigma_n)$ *such that*

1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap \mathrm{sig}(\beta)$,

2. $\vdash \alpha \Rightarrow \iota$ *and* $\vdash \iota \Rightarrow \beta$.

Proof sketch. The structure of disjunctive formulas is such that a cyclic proof of $\alpha \Rightarrow \beta$ with α and $\beta^\perp$ both disjunctive can be transformed into one in which every bud contains at most one formula in each cedent. We may further assume that the cyclic proof is normal and utilises conservative instances of rules only, in the sense of the proof of Proposition 3.11.

From the observations above and the assumption that $\alpha \in \Pi_{n+1}$ and $\beta \in \Sigma_{n+1}$ we deduce that in every chain of companions $u_0 \lhd u_1 \lhd \cdots \lhd u_k$ the sequence of markings $\sigma(u_0), \ldots, \sigma(u_k)$ has fewer than n alternations, whence the Lyndon interpolant induced by the cyclic proof is in $\mathit{Clos}(\Pi_n \cup \Sigma_n)$. ◀

▶ **Corollary 6.21** (Collapse of Disjunctive-Δ_n). *Suppose* $\alpha \in \mathbb{D} \cap \Pi_{n+1}$. *If* $\alpha^\perp$ *is equivalent to a formula in* $\mathbb{D} \cap \Pi_{n+1}$, *then* $\alpha \in \mathit{Clos}(\Sigma_n \cup \Pi_n)^{\equiv}$.

6.3 Uniform Interpolation for the Full Language

The normal form theorem for templates (Theorem 6.7) also offers a means to prove uniform interpolation for the full language of the modal μ-calculus:

▶ **Theorem 6.22** (Uniform Lyndon Interpolation). *For every signature S and formula $\alpha \in \mathsf{L}_\mu$ there exists a formula ι such that*

1. $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap S$.

2. $\vdash \alpha \Rightarrow \iota$.

3. *For all* β, *if* $\mathrm{sig}(\beta) \subset S$ *and* $\vdash \alpha \Rightarrow \beta$, *then* $\vdash \iota \Rightarrow \beta$.

Proof. Fix S and α, and an α-template T. By Theorem 6.7 we may assume that (T,σ) is in normal form. Let $\iota = \iota_{(T,\sigma,S)}$ be the induced uniform candidate with $\hat{\iota}_u$ the candidate associated to $u \in T$. By design, $\mathrm{sig}(\iota) \subset \mathrm{sig}(\alpha) \cap S$. Lemmas 5.3 and 5.4 describe the cyclic derivations that witness conditions 2 and 3 and it suffices to confirm that these derivations are proofs. Concerning 2, recall that a cyclic derivation D_T of $\alpha \Rightarrow \iota$ is obtained by installing $\hat{\iota}_u$ as the right cedent of $u \in T$ and interspersing the derivation with right rules as needed. Let P be any path in this derivation and with $u = \min P$ the minimal companion visited infinitely often. Suppose that P does not carry a good left trace. So, $\sigma(u) = \nu$. But then $\hat{\iota}_u = \nu x_u \zeta$ for some ζ and, by design, $\hat{\iota}_u$ is the shortest formula along the single right trace carried by P. Hence, D_T is a cyclic proof.

The argument towards 3 is similar. Let $\mathrm{sig}(\beta) \subset S \cap \mathrm{sig}(\alpha)$. The proof of Lemma 5.4 describes the construction of a cyclic derivation D' with endsequent $\iota \Rightarrow \beta$ from a cyclic proof D of $\alpha \Rightarrow \beta$. Consider a path P through D'. The construction of D' was such that P is associated a path P_D through D and a path P_T through T such that (i) P and P_D carry the same right traces, (ii) P carries a good left trace if P_T does, and (iii) P_T carries every left trace of P_D. As D is a cyclic proof, P must carry a good trace. ◀

The candidate uniform interpolants obtained from Construction 2 are all disjunctive formulas (cf. Definition 6.19). An immediate corollary of our uniform interpolation theorem is the celebrated *disjunctive normal form theorem* for the modal μ-calculus, established in [20]:

▶ **Corollary 6.23** (Disjunctive normal form theorem)**.** *Every formula of the modal μ-calculus has an equivalent disjunctive formula in the same signature.*

Proof. Apply Theorem 6.22 for $S = \mathrm{sig}(\alpha)$. Then $\vdash \alpha \Rightarrow \iota$ (condition 2) and $\vdash \iota \Rightarrow \alpha$ (by condition 3). ◀

Corollary 5.7 gave a characterisation of the Σ_1 formulas that are equivalent to a Π_1 formula and, by extension, the Σ_1 formulas equivalent to quantifier-free formulas. The template approach to uniform interpolants lifts the characterisation to the whole calculus. The first step is a characterisation of Π_1-equivalence due to Küsters and Wilkie [24]. As a second corollary we obtain Otto's theorem on the decidability of quantifier-free equivalence [31].

▶ **Theorem 6.24** (Π_1-characterisation)**.** *For every $\alpha \in \mathsf{L}_\mu$ there is an effectively computed $\alpha_\Pi \in \Pi_1$ such that $\alpha \in \Pi_1^{\equiv}$ iff $\alpha \equiv \alpha_\Pi$.*

Proof. Let $S = \mathrm{Lit}$ be the universal signature. We carry out the proof of Theorem 6.22 with modification to candidates used in Theorem 5.5, from which the claim is an immediate consequence. ◀

▶ **Corollary 6.25** (Π_1-decidability)**.** *The problem of determining whether an arbitrary formula of the modal μ-calculus is equivalent to a Π_1 formula is decidable.*

Decidability of quantifier-free definability now extends to the full language:

▶ **Corollary 6.26** (Otto's theorem). *The problem of determining whether an arbitrary formula of the modal μ-calculus is equivalent to a quantifier-free formula is decidable.*

Proof. Define $\alpha_\Sigma = \neg(\neg\alpha)_\Pi \in \Sigma_1$. By the results above, $\langle\!|\alpha_\Sigma|\!\rangle \subset \langle\!|\alpha|\!\rangle \subset \langle\!|\alpha_\Pi|\!\rangle$ and $\langle\!|\alpha|\!\rangle$ is definable by a quantifier-free formula iff $\vdash \alpha_\Pi \Rightarrow \alpha_\Sigma$. ◀

The final result we present is an analogue to the automata-theoretic result stating that non-deterministic Büchi tree automata accept the same languages as alternating Büchi automata. From the perspective of interpolation, the result shows that the Lyndon interpolation theorems for the first and second quantifiers (Theorems 4.8 and 6.15) reduce to our interpolation theorem for disjunctive formulas (Theorem 6.20).

▶ **Theorem 6.27.** *Every Π_1 and Π_2 formula is equivalent to a Π_1 and Π_2 disjunctive formula respectively.*

Proof. We show the Π_2 case. Fix $\alpha \in \Pi_2$ and let (T, σ) be a normal α-template. Let ι be the uniform candidate for α with respect to $\mathrm{sig}(\alpha)$. In particular, $\iota \equiv \alpha$. As in Proposition 6.14, we have that for all $u \lhd v \in \mathrm{comp}(T)$, if $\sigma(u) = \mu$ then $\sigma(v) = \mu$. Thus, $\iota \in \Pi_2 \cap \mathbb{D}$. ◀

7 Discussion

From any cyclic derivation of a sequent $\alpha \Rightarrow \beta$, the cyclic Maehara method yields an explicit formula ι in the signature common to both α and β, and cyclic derivations of $\alpha \Rightarrow \iota$ and $\iota \Rightarrow \beta$. If the initial derivation happens to be a cyclic proof and in normal form then the resulting derivations are guaranteed to be cyclic proofs and ι a Lyndon interpolant of $\alpha \Rightarrow \beta$.

The structure of the interpolating formula is completely determined by the initial cyclic derivation of $\alpha \Rightarrow \beta$. In selected cases, sufficient bounds on the complexity of the interpolant can be given. If α and β both sit in the first non-trivial level of the quantifier hierarchy, a candidate interpolant can always be selected from the same level (Theorems 4.8 through 4.14). If α is Π_2 and β is Σ_2, the candidate formed from a normal proof will be alternation-free, i.e., in $\Pi_2 \cap \Sigma_2$ (Theorem 6.15). In particular, the alternation-free fragment of L_μ has the Lyndon interpolation property (Corollary 6.16) and is the largest level of the quantifier hierarchy with this property (Theorem 6.18; [34]). Similar observations, both positive and negative, can be made for the designation of uniform interpolants (e.g. Theorems 5.5 and 6.27).

The question naturally arises whether the syntactic method can be applied also to fragments of the μ-calculus corresponding to specific fixed point logics. In such settings the quantifier complexity of the induced interpolant is less important than the issue of whether the constructed interpolant can be expressed by (or adjusted to be) a formula in the chosen fixed point logic, as the latter is clearly required to conclude the Lyndon/uniform interpolation property. Gödel–Löb provability logic GL is one example in which the gap between construction and logic can be bridged. Shamkanov's syntactic proof of Lyndon interpolation [35] stems from the realisation of Gödel–Löb logic as a cyclic proof system over the modal logic of transitive frames, K4. As the Maehara method supports the K4 modality, every GL-provable sequent admits a Lyndon interpolant in the μ-calculus over K4. Lyndon interpolation for GL is a corollary of the de Jong–Sambin–Bernardi fixed point theorem which establishes that the μ-calculus collapses over GL.

There are, however, many examples of logics that cannot internalise the structure of their cyclic proofs. Well-known fixed point logics that lack the Craig interpolation property include linear-time temporal logic, computation-tree logic [27] and the logic of common knowledge [37]. Each of these logics can be associated sound and complete analytic sequent calculi that support the cyclic Maehara construction of candidate interpolants. Lyndon interpolants for valid sequents are definable in the linear-time μ-calculus, multi-modal μ-calculus and μ-calculus over S5 respectively, but not always in the initial fixed point logic itself. This phenomenon was explored in detail in [23]. A deeper analysis of the cyclic Maehara method can potentially contribute to characterising the necessary complexity of interpolants and algorithms for deciding their existence within specific fragments.

Acknowledgments

Both authors were supported by the Knut and Alice Wallenberg Foundation [2020.0199]. The first author was additionally supported by the Swedish Research Council [2022-01685] and the Dutch Research Council [OCENW.M20.048]. The authors would like to thank the referees for their suggestions, Sam van Gool and his students at ENS Saclay and Cárolos Laméris for eliminating many errors from this work.

References

1 Bahareh Afshari and Graham E. Leigh. Finitary proof systems for Kozen's μ. Oberwolfach preprints (owp) issn 1864-7596, Mathematisches Forschungsinstitut Oberwolfach, 2016.

2 Bahareh Afshari and Graham E. Leigh. Cut-free completeness for modal mu-calculus. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12. IEEE Computer Society, 2017.

3 Bahareh Afshari and Graham E. Leigh. Lyndon interpolation for modal μ-calculus. In Aybüke Özgün and Yulia Zinova, editors, *Language, Logic, and Computation*, pages 197–213, Cham, 2022. Springer International Publishing. `doi:10.1007/978-3-030-98479-3_10`.

4 Bahareh Afshari, Graham E. Leigh, and Guillermo Menéndez Turata. Uniform interpolation from cyclic proofs: The case of modal mu-calculus. In Anupam Das and Sara Negri, editors, *Automated Reasoning with Analytic Tableaux and Related Methods. TABLEAUX 2021*, volume 12842 of *Lecture Notes in Computer Science*, Cham, 2021. Springer. `doi:10.1007/978-3-030-86059-2_20`.

5 Bahareh Afshari, Graham E. Leigh, and Guillermo Menèndez Turata. Demystifying μ. *Fundamenta Informaticae*, 194(2):12773, August 2025. `doi:10.46298/fi.12773`.

6 Bahareh Afshari and Dominik Wehr. Abstract cyclic proofs. *Mathematical Structures in Computer Science*, 34(7):552–577, August 2024. `doi:10.1017/S0960129524000070`.

7 André Arnold and Damian Niwiński. Fixed point characterization of Büchi automata on infinite trees. *J. Inf. Process. Cybern.*, 26(8-9):451–459, January 1991.

8 André Arnold and Damian Niwiński. *Rudiments of μ-calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2001.

9 Nick Bezhanishvili, Balder ten Cate, and Rosalie Iemhoff. Six proofs of interpolation for the modal logic K. In ten Cate et al. [12], chapter 4. Preprint: https://arxiv.org/abs/2510.16398. doi:10.5334/bdg.d.

10 Julian Bradfield and Colin Stirling. 12 Modal mu-calculi. In *Studies in Logic and Practical Reasoning*, volume 3, pages 721–756. Elsevier, 2007. doi:10.1016/S1570-2464(07)80015-2.

11 Balder ten Cate and Jesse Comer. Interpolation in first-order logic. In ten Cate et al. [12], chapter 2. Preprint: https://arxiv.org/abs/2510.03822. doi:10.5334/bdg.b.

12 Balder ten Cate, Jean Christoph Jung, Patrick Koopmann, Christoph Wernhard, and Frank Wolter, editors. *Theory and Applications of Craig Interpolation*. Ubiquity Press, 2026. To appear; preprints accessible from https://cibd.bitbucket.io/taci/. doi:10.5334/bdg.

13 Giovanna D'Agostino and Marco Hollenberg. Logical questions concerning the μ-calculus: Interpolation. Lyndon and Łos-Tarski. *The Journal of Symbolic Logic*, 65(1):310–332, 2000.

14 Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ-calculus. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Dough Tygar, Moshe Y. Vardi, Gerhard Weikum, S. Arun-Kumar, and Naveen Garg, editors, *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science*, volume 4337 of *Lecture Notes in Computer Science*, pages 273–284. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. doi:10.1007/11944836_26.

15 Maurice Dekker, Johannes Kloibhofer, Johannes Marti, and Yde Venema. Proof systems for the modal μ-calculus obtained by determinizing automata. In Revantha Ramanayake and Josef Urban, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 242–259, Cham, 2023. Springer Nature Switzerland. doi:10.1007/978-3-031-43513-3_14.

16 Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, April 1979. Publisher: Elsevier BV. doi:10.1016/0022-0000(79)90046-1.

17 Iris van der Giessen, Raheleh Jalali, and Roman Kuznets. Interpolation in proof theory. In ten Cate et al. [12], chapter 5. To appear; preprints accessible from https://cibd.bitbucket.io/taci/. doi:10.5334/bdg.e.

18 Sam van Gool. Uniform interpolation. In ten Cate et al. [12], chapter 9. Preprint: https://arxiv.org/abs/2512.15391. doi:10.5334/bdg.i.

19 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002. doi:10.1007/3-540-36387-4.

20 David Janin and Igor Walukiewicz. Automata for the modal μ-calculus and related results. In Jiří Wiedermann and Petr Hájek, editors, *Mathematical Foundations of Computer Science 1995: 20th International Symposium, MFCS '95 Prague, Czech Republic, August 28–September 1, 1995 Proceedings*, pages 552–562, Berlin, Heidelberg, 1995. Springer Berlin Heidelberg.

21 Natthapong Jungteerapanich. A tableau system for the modal μ-calculus. In Martin Giese and Arild Waaler, editors, *Automated Reasoning with Analytic Tableaux and Related Methods, 18th International Conference, TABLEAUX 2009, Oslo, Norway, July 6-10, 2009. Proceedings*, volume 5607 of *Lecture Notes in Computer Science*, pages 220–234. Springer, 2009.

22 Dexter Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

23 Agi Kurucz, Frank Wolter, and Michael Zakharyaschev. From interpolating formulas to separating languages and back again. In ten Cate et al. [12], chapter 10. Preprint: https://arxiv.org/abs/2508.12805. `doi:10.5334/bdg.j`.

24 Ralf Küsters and Thomas Wilke. Deciding the first level of the μ-calculus alternation hierarchy. In Manindra Agrawal and Anil Seth, editors, *Proceedings of the 22nd Conference Kanpur on Foundations of Software Technology and Theoretical Computer Science*. Springer-Verlag Berlin Heidelberg, December 2002.

25 Graham E. Leigh and Dominik Wehr. From GTC to Reset: Generating reset proof systems from cyclic proof systems. *Annals of Pure and Applied Logic*, 175(10):103485, December 2024. `doi:10.1016/j.apal.2024.103485`.

26 Shôji Maehara. Craig-interpolation theorem. *Sugaku*, 12(4):235–237, 1961.

27 Larisa Maksimova. Temporal logics of “the next” do not have the Beth property. *Journal of Applied Non-Classical Logics*, 1(1):73–76, January 1991. `doi:10.1080/11663081.1991.10510772`.

28 Johannes Marti and Yde Venema. Focus-style proof systems and interpolation for the alternation-free μ-calculus, April 2021. arXiv:2103.01671 [cs]. `doi:10.48550/arXiv.2103.01671`.

29 Robert McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9(5):521–530, October 1966. `doi:10.1016/S0019-9958(66)80013-X`.

30 Damian Niwinski and Igor Walukiewicz. Games for the mu-calculus. *Theoretical Computer Science*, 163(1&2):99–116, 1996.

31 Martin Otto. Eliminating recursion in the μ-calculus. In Gerhard Goos, Juris Hartmanis, Jan Van Leeuwen, Christoph Meinel, and Sophie Tison, editors, *STACS 99*, volume 1563, pages 531–540. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999. Series Title: Lecture Notes in Computer Science. `doi:10.1007/3-540-49116-3_50`.

32 Andrew M. Pitts. On an interpretation of second order quantification in first order intuitionistic propositional logic. *J. Symb. Log.*, 57(1):33–52, 1992.

33 Michael O. Rabin. Weakly definable relations and special automata. In *Studies in Logic and the Foundations of Mathematics*, volume 59, pages 1–23. Elsevier, 1970. `doi:10.1016/S0049-237X(08)71929-3`.

34 Luigi Santocanale and André Arnold. Ambiguous classes in μ-calculi hierarchies. *Theoretical Computer Science*, 333(1-2):265–296, March 2005. `doi:10.1016/j.tcs.2004.10.024`.

35 Daniyar Shamkanov. Circular proofs for the Gödel-Löb provability logic. *Math. Notes*, 96:575–585, 2014.

36 Colin Stirling. A tableau proof system with names for modal mu-calculus. In Andrei Voronkov and Margarita V. Korovina, editors, *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, volume 42 of *EPiC Series in Computing*, pages 306–318. EasyChair, 2014.

37 Thomas Studer. Common knowledge does not have the Beth property. *Information Processing Letters*, 109(12):611–614, May 2009. doi:10.1016/j.ipl.2009.02.011.