Cyclic Hypersequent Calculi for Some Modal Logics with the Master Modality

Jan Rooduijn(✉)

University of Amsterdam, Amsterdam, Netherlands
j.m.w.rooduijn@uva.nl
https://staff.fnwi.uva.nl/j.m.w.rooduijn/

Abstract. At LICS 2013, O. Lahav introduced a technique to uniformly construct cut-free hypersequent calculi for basic modal logics characterised by frames satisfying so-called 'simple' first-order conditions. We investigate the generalisation of this technique to modal logics with the master modality (a.k.a. reflexive-transitive closure modality). The (co)inductive nature of this modality is accounted for through the use of non-well-founded proofs, which are made cyclic using focus-style annotations. We show that the peculiarities of hypersequents hinder the usual method of completeness via infinitary proof-search. Instead, we construct countermodels from maximally unprovable hypersequents. We show that this yields completeness for a small (yet infinite) subset of simple frame conditions.

Keywords: Hypersequent calculi · Modal logic · Master modality · Non-well-founded proofs · Cyclic proofs

1 Introduction

Cyclic and *non-well-founded* proofs have turned out to be highly effective in the proof theory of modal fixpoint logics. They have been applied to obtain proof-theoretic proofs of known results, such as the completeness of Kozen's axiomatisation of the modal μ-calculus [1], and Lyndon interpolation for Gödel-Löb logic [12]. Moreover, cyclic proof systems have been constructed for logics for which until then no proof system was known, *e.g.* Game Logic [7] and the hybrid μ-calculus [6]. The key advantage of cyclic proof systems over systems with explicit (co)induction rules, is that they enjoy a variant of the subformula property. Among other benefits, this makes them more suitable for proof search. Although cyclic proof systems have by now been devised for many modal fixpoint logics, little work has been done on constructing such systems in a uniform way. In particular, there is no general method to obtain cyclic proof systems for modal fixpoint logics characterised by various classes of frames. This paper attempts to take a first step in that direction.

This research has been made possible by a grant from the Dutch Research Council NWO, project nr. 617.001.857.

A. Das and S. Negri (Eds.): TABLEAUX 2021, LNAI 12842, pp. 354–370, 2021.
https://doi.org/10.1007/978-3-030-86059-2_21

Already without fixpoints, many modal logics call for a deviation from the standard sequent calculus. A typical example is the modal logic S5 (characterised by frames whose accessibility relation is an equivalence relation), for which obtaining a cut-free calculus in the standard sequent system is notoriously difficult. Several alternatives have been proposed, most of which equip ordinary sequents with extra structure, often echoing the Kripke semantics (for an overview we refer the reader to Chap. 4 of [8]). The alternative that arguably stays closest to Gentzen's original approach is that of *hypersequents*, which are nothing but finite disjunctions of sequents. Already with this minor modification, many more modal logics, including S5, can be given a sound and complete proof system. In [10], Ori Lahav presents a systematic method for constructing hypersequent calculi for any extension of one of the modal logics K, K4 or KB, characterised by frames satisfying any finite number of so-called *simple* frame conditions.

In this paper we adapt Lahav's method to uniformly obtain cyclic proof systems for a comparatively simple modal fixpoint logic: *unimodal logic with the master modality.* This language, denoted ML^*, augments the basic modal language with a modality $\boxast$, which is to be thought of as the reflexive-transitive closure of the basic modality $\Box$. For each finite set $\mathcal{C}$ of simple frame conditions we uniformly construct both an infinitary and a cyclic hypersequent calculus for ML^* interpreted on the class of $\mathcal{C}$-frames. In the cyclic systems, sequents are annotated using a focus mechanism originally due to Lange and Stirling (see *e.g.* [11]). All systems are proven to be sound, but completeness is only proven for a subset of the simple frame conditions which we shall call *equable*. While many simple frame conditions are not equable, there are infinitely many equable frame conditions, including: seriality, reflexivity, directedness and universality. As a corollary, we obtain decidability for each of these logics.

As for related work, a finitary analytic proof system for ML^* interpreted on the class of all frames is given in [5]. In [3], a cyclic proof system is presented for LTL and CTL, two modal fixpoint logics that are interpreted on restricted frames classes. In [9], a general method is given for constructing sound and complete Hilbert systems for ML^* interpreted on various frame classes, but this concerns non-analytic systems having both a cut-rule and an explicit induction rule. Another notable example of related work is [4], where, like here, cyclic proofs are combined with some calculus that extends the ordinary sequent calculus. However, they use labelled sequents rather than hypersequents and do not consider multiple logics at once.

In Sect. 2 we introduce the syntax and semantics of ML^* and define simple and equable frame conditions. In Sect. 3 we introduce our hypersequent calculi. Section 4 proves soundness for all calculi. Finally, in Sect. 5 completeness is proven for those calculi that contain only rules for equable frame classes.

2 Preliminaries

For the rest of this article, we fix a countable set of P of *propositional variables.*

Definition 1. *The syntax* ML^* *of* modal $*$-formulas *over* P *is generated by:*

$$\varphi ::= p \mid \bot \mid \varphi \to \varphi \mid \Box\varphi \mid \boxast\varphi$$

where $p \in \mathsf{P}$*.*

As usual, formulas will be interpreted in Kripke models. We will refer to modal $*$-formulas as just *formulas.*

Definition 2. *A* Kripke frame *is a pair* (S, R) *consisting of a set* S *of* states *together with an* accessibility relation $R \subseteq S \times S$*. A* Kripke model *is a triple* (S, R, V)*, where* (S, R) *is a* Kripke frame *and* $V : \mathsf{P} \to \mathcal{P}(S)$ *a valuation function.*

Formulas are interpreted in Kripke models in the usual way, with the following additional clause for $\boxast$:

$$\mathbb{S}, s \Vdash \boxast\psi \ :\Leftrightarrow \ \text{for all } t \in S \text{ such that } sR^*t\text{: } \mathbb{S}, t \Vdash \psi$$

where R^* is the reflexive-transitive closure of R. Whenever the intended the model $\mathbb{S}$ is clear from the context, we will simply write $s \Vdash \varphi$ instead of $\mathbb{S}, s \Vdash \varphi$.

Let L_1 be the first-order language with equality and a single relation symbol R. In contrast to ML^*, we let L_1 include the propositional connectives $\wedge$, $\vee$ and $\neg$. A *frame condition* then is nothing but an L_1-sentence. For Θ a set of frame conditions, a Kripke frame (S, R) is said to be a Θ*-frame* whenever, when regarded an L_1-structure, the frame (S, R) satisfies each sentence φ in Θ. A Kripke model will be called a Θ*-model* whenever its underlying frame is a Θ-frame.

The following definitions and proposition are taken from [10].

Definition 3. *A frame condition is called* n-simple *whenever it is of the form* $\forall s_1 \cdots s_n \exists u \varphi$*, where* φ *is built up using the connectives* $\vee$ *and* $\wedge$ *from atomic formulas of the form* $s_i R u$ *and* $s_i = u$ *with* $1 \leq i \leq n$*.*

Definition 4. *Given* $n \in \omega$*, an* abstract n-simple frame condition *is a finite non-empty set* C *consisting of pairs* $(C_R, C_=)$ *of subsets* $C_R, C_= \subseteq \{1, \dots, n\}$ *such that at least one of* C_R *and* $C_=$ *is non-empty.*

Definition 5. *The* interpretation *of some abstract* n*-simple frame condition* C *is the following first-order formula:*

$$\forall s_1 \cdots s_n \exists u \bigvee_{(C_R, C_=) \in C} \Big(\bigwedge_{i \in C_R} s_i R u \wedge \bigwedge_{j \in C_=} s_j = u \Big).$$

Using disjunctive normal forms, the following proposition is immediate.

Proposition 1. *Any n-simple frame condition is equivalent to the interpretation of some abstract n-simple frame condition.*

In the following, we use the general term *(abstract) simple frame condition* to encapsulate every (abstract) frame condition that is n-simple for some $n \in \omega$. For the sake simplicity we will sometimes blur the distinction between an abstract frame condition C and its interpretation. In particular, for $\mathcal{C}$ a finite set of abstract simple frame conditions and Θ the set of their interpretations, we often use the terms $\mathcal{C}$-model and $\mathcal{C}$-frame where we mean Θ-model and Θ-frame.

We close this section by defining the subclasses of the class of simple frame conditions, to which we will restrict most of our attention for the rest of this paper.

Definition 6. *An abstract n-simple frame condition C is called:*

- equality-free *if $C_= = \emptyset$ for all $(C_R, C_=) \in C$;*
- disjunction-free *if C is a singleton;*
- equable *if for some $U \subseteq \{1, \ldots, n\}$, we have $U = C_=$ for all $(C_R, C_=) \in C$.*

Clearly if C is equality-free or disjunction-free, then it is equable. It turns out that the converse is also true (up to logical equivalence). The verification of this fact is left to the reader. Some examples of equable frame conditions are *reflexivity*, given by $C = \{\langle\{1\}, \{1\}\rangle\}$, and *k-bounded top width*, which is given by $C = \{\langle\{i, j\}, \emptyset\rangle : 1 \leq i < j \leq k\}$ for any $k \geq 2$. An example of a simple frame condition which is *not* equable is $C = \{\langle\{1\}, \{2\}\rangle, \langle\{2\}, \{1\}\rangle\}$, which in [10] is called *linearity*. For more examples of simple frame conditions, we refer the reader to the aforementioned article.

3 Infinitary and Cyclic Hypersequent Calculi

In this section we introduce families of infinitary and cyclic hypersequent calculi for ML^* interpreted on classes of Θ-models, where Θ is an arbitrary set of simple frame conditions.

3.1 Hypersequents and Pre-proofs

Definition 7. *A* sequent *is an ordered pair (Γ, Δ) of finite sets of formulas, usually written $\Gamma \Rightarrow \Delta$. A* hypersequent *is a finite set $\{\sigma_0, \ldots, \sigma_n\}$ of sequents, usually written $\sigma_0 \mid \cdots \mid \sigma_n$.*

We adopt the convention of using shorthand notation for singleton formulas and sequents. For instance, we write $\Gamma, \varphi \Rightarrow \psi, \Delta$ where we mean $\{\varphi\} \cup \Gamma \Rightarrow \{\psi\} \cup \Delta$, and the hypersequent $H \cup \{\sigma\}$ may be written as $H \mid \sigma$.

(Hyper)sequents are interpreted in Kripke models as follows.

Definition 8. *Let $\mathbb{S}$ be a Kripke model. Then:*

- *A sequent $\Gamma \Rightarrow \Delta$ is said to be* satisfied *at a state s of $\mathbb{S}$ whenever:*
  If $s \Vdash \varphi$ for all $\varphi \in \Gamma$, then $s \Vdash \psi$ for some $\psi \in \Delta$.
- *A sequent is* valid *in $\mathbb{S}$ if it is satisfied at every state of $\mathbb{S}$.*
- *A hypersequent H is* valid *in $\mathbb{S}$ if there is a $\sigma \in H$ which is valid in $\mathbb{S}$.*

A hypersequent valid in all $\mathcal{C}$-models will be called *$\mathcal{C}$-valid*.

The following hypersequent calculus is an expansion by two additional fixpoint rules of the system HK given in [10] for basic modal logic.

Definition 9. *The hypersequent calculus* HK^* *has the following axioms and rules.*

$$\mathsf{id}\ \frac{}{\varphi \Rightarrow \varphi} \qquad\qquad \bot\ \frac{}{\bot \Rightarrow}$$

$$\mathsf{iw}_L\ \frac{H \mid \Gamma \Rightarrow \Delta}{H \mid \Gamma, \varphi \Rightarrow \Delta} \qquad \mathsf{iw}_R\ \frac{H \mid \Gamma \Rightarrow \Delta}{H \mid \Gamma \Rightarrow \varphi, \Delta} \qquad \mathsf{ew}\ \frac{H}{H \mid \Gamma \Rightarrow \Delta}$$

$$\to_L\ \frac{H \mid \Gamma, \psi \Rightarrow \Delta \qquad H \mid \Gamma \Rightarrow \varphi, \Delta}{H \mid \Gamma, \varphi \to \psi \Rightarrow \Delta} \qquad \to_R\ \frac{H \mid \Gamma, \varphi \Rightarrow \psi, \Delta}{H \mid \Gamma \Rightarrow \varphi \to \psi, \Delta}$$

$$\Box\ \frac{H \mid \Gamma \Rightarrow \varphi}{H \mid \Box\Gamma \Rightarrow \Box\varphi} \qquad \mathsf{cut}\ \frac{H \mid \Gamma_1, \varphi \Rightarrow \Delta_1 \qquad H \mid \Gamma_2 \Rightarrow \varphi, \Delta_2}{H \mid \Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2}$$

$$\boxast_L\ \frac{H \mid \Gamma, \varphi, \Box\boxast\varphi \Rightarrow \Delta}{H \mid \Gamma, \boxast\varphi \Rightarrow \Delta} \qquad \boxast_R\ \frac{H \mid \Gamma \Rightarrow \varphi, \Delta \qquad H \mid \Gamma \Rightarrow \Box\boxast\varphi, \Delta}{H \mid \Gamma \Rightarrow \boxast\varphi, \Delta}$$

Following [10], we augment HK^* with rules corresponding to certain simple frame conditions.

Definition 10. *Let C be an abstract n-simple frame condition. The rule* $\mathsf{r}_C^{\mathsf{HK}^*}$ induced *by C is defined as follows:*

$$\mathsf{r}_C^{\mathsf{HK}^*}\ \frac{\{H \mid \bigcup_{i \in C_R} \Gamma_i', \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j : (C_R, C_=) \in C\}}{H \mid \Box\Gamma_1', \Gamma_1 \Rightarrow \Delta_1 \mid \cdots \mid \Box\Gamma_n', \Gamma_n \Rightarrow \Delta_n}$$

Given a finite set $\mathcal{C}$ of abstract simple frame conditions, we let $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ be the system HK^* augmented with the rules $\mathsf{r}_C^{\mathsf{HK}^*}$ for each $C \in \mathcal{C}$.

In any application of some rule of $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$, the sequents outside of the context H are called *active*. Furthermore, the *active* formulas of an active sequent are those that occur outside of Γ and Δ. All other formulas and sequents are called *inactive*. Note that due to the fact the (hyper)sequents are sets, the contexts H might also contain active sequents (and likewise Γ and Δ might contain active formulas). In the case of $\mathsf{r}_C^{\mathsf{HK}^*}$, the i-th active sequent in the conclusion is said to have *index* i and the premiss corresponding to $(C_R, C_=) \in C$ is said to have *index* $(C_R, C_=)$. Here the fact that hypersequents are sets means that one sequent might have multiple indices.

For the rest of this paper we assume that $\mathcal{C}$ is an arbitrary finite set of simple frame conditions, unless specified otherwise.

Definition 11. *An* $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$-pre-proof *is a (possibly infinite) derivation in* $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$.

For any $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$-pre-proof π with root H, we say that π is a $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$-pre-proof *of* H.

This derivation system has a property akin to the subformula property.

Definition 12. *The* closure *of a set Φ of formulas is the least $\Psi \supseteq \Phi$ such that:*

(i) If $\varphi \to \psi \in \Psi$, then $\varphi, \psi \in \Psi$;
(ii) If $\Box\varphi \in \Psi$, then $\varphi \in \Psi$;
(iii) If $\boxast\varphi \in \Psi$, then $\varphi, \Box\boxast\varphi \in \Psi$.

We write $Cl(\Phi)$ for the closure of Φ. It is easy to see that Cl is a closure operator and that the closure of any finite set of formulas is finite. The following lemma can be verified by direct inspection of the rules.

Lemma 1. *Let π be a cut-free $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$-pre-proof of H. Any formula occurring in π belongs to the closure of the set of formulas occurring in H.*

3.2 Infinitary Proofs with Trace Condition

It is not hard to show that the system $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ need not be sound with respect to all Kripke models based on a $\mathcal{C}$-frame. In fact, already when $\mathcal{C}$ is empty there are infinite pre-proofs of invalid hypersequents. We therefore need a way to recognize valid infinite proofs. The technical treatment in this section takes inspiration from [4], which in turn follows [2].

We use $\Box^n\varphi$ as a shorthand for the formula φ preceded by n instances of $\Box$.

Definition 13. *A formula φ is said to be a* trace formula *if it is of the form $\Box^i\boxast\psi$ for $i \in \{0, 1\}$. If $i = 1$, we say that φ is* unfolded.

Definition 14. *A* trace value *is either the* empty trace value *ϵ, or a pair (φ, σ), where σ is a sequent and φ a trace formula in the right-hand side of σ.*

If τ is the empty trace value or $\tau = (\varphi, \sigma)$ such that the sequent σ belongs to some hypersequent H, then τ is said to be a trace value *for* H.

Definition 15. *Let (H, H') be a pair consisting of the conclusion and a premiss, respectively, of an application of a some rule r of $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ and let τ and τ' be trace values for H and H'. The pair (τ, τ') is called a* trace pair *for (H, H') if one of τ and τ' is the empty trace value, or one of the following conditions holds for $\tau = (\varphi, \sigma)$ and $\tau = (\varphi', \sigma')$:*

1. *σ' is an inactive sequent equal to σ and $\varphi = \varphi'$.*
2. *σ and σ' are active sequents and one of the following holds:*
   (a) r is among $\mathsf{iw}_L, \mathsf{iw}_R, \mathsf{cut}, \to_L, \to_R, \boxast_L, \boxast_R$ and $\varphi' = \varphi$.
   (b) $\mathsf{r} \in \{\mathsf{r}_C^{\mathsf{HK}^*} \mid C \in \mathcal{C}\}$, the index of σ is in $C_=$, where $(C_R, C_=)$ is the index of H', and $\varphi = \varphi'$.
   (c) r is $\Box$ and $\Box\varphi' = \varphi$.
   (d) r is $\boxast_R$, φ is active, H' is the right-hand premiss, and $\varphi' = \Box\varphi$.

When (τ, τ') is a trace pair by virtue of item 2(d), it will be called an unfolding.

Remark 1. There are several subtleties involved with Definition 15:

– Consider the following inference, where every trace value is marked.

$$\boxast_R \; \dfrac{p \Rightarrow \overbrace{\Box\boxast p}^{\tau_3} \mid p \Rightarrow \overbrace{\boxast p}^{\tau_4} \mid p \Rightarrow p \qquad p \Rightarrow \overbrace{\boxast p}^{\tau_5} \mid p \Rightarrow \overbrace{\Box\boxast p}^{\tau_6}}{p \Rightarrow \underbrace{\Box\boxast p}_{\tau_1} \mid p \Rightarrow \underbrace{\boxast p}_{\tau_2}}$$

The trace value τ_1 does *not* form a trace pair with τ_6, because the sequent belonging to τ_6 is active, whereas the one belonging to τ_1 is not. In contrast, since the sequent of τ_5 is inactive, the pair (τ_2, τ_5) is a trace value, even though τ_2 is active. The other trace pairs are (τ_1, τ_3), (τ_2, τ_4) and (τ_2, τ_6) is a trace pair, with the latter being an unfolding.

– In case (2)(c), the shape of the rule $\Box$ forces σ' to be of the form $\Gamma \Rightarrow \varphi'$, where σ is of the form $\Box\Gamma \Rightarrow \Box\varphi'$.

Definition 16. *A* trace *is a sequence of trace pairs. A trace is called* good *if it contains finitely many empty trace values and infinitely many unfoldings.*

Definition 17. *A path $(H_i)_{i\in I}$ in some proof is said to be* covered *by a trace $(\tau_i)_{i\in I}$ if (τ_i, τ_{i+1}) is a trace pair for (H_i, H_{i+1}) for each $i \in I$ such that $i+1 \in I$.*

Definition 18. *An* $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$-proof *is an $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$-pre-proof of which every infinite branch is covered by a good trace.*

A hypersequent H will be called $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$-provable if there is an $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$-proof whose root is labelled by H.

3.3 Cyclic Proofs

In this section we assume that $\mathcal{C}$ is a finite set of equable frame conditions.

Definition 19. *An* annotated hypersequent *is a hypersequent H together with a trace value τ for H. We call τ an* annotation, *say that H is* annotated *by τ and write $\tau \vdash H$.*

In proof trees, we often simplify notation by, instead of writing $\tau \vdash$, putting the formula designated by τ between square brackets. This formula is then said to be *in focus*. When τ is empty, we signify this by putting no formula between brackets.

Definition 20. *The derivation system $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ is obtained from $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ by making the following adaptations:*

1. *The basic judgments are annotated hypersequents.*
2. *If H is derivable from $H_1, \ldots, H_n$ by some rule r of $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$, then $\tau \vdash H$ is derivable from $\tau_1 \vdash H_1, \ldots, \tau_n \vdash H_n$ by r in $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ if and only if the pair (τ, τ_i) is a trace pair for (H, H_i).*

3. The following structural rule, called focus change*, is added:*

$$\mathsf{fc}\,\frac{\tau \vdash H}{\tau' \vdash H}$$

Here τ and τ' may be any two trace values for H.

Although the rules of the derivation system $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ are given in an indirect fashion, it is clearly decidable whether some given inference is a valid rule application.

Whenever some leaf l of some derivation in $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ is the conclusion of an application of id or $\bot$, we say that l is an *axiomatic* leaf.

Definition 21. *An* $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof *is a finite derivation π in $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ together with a back edge map f assigning to each non-axiomatic leaf l of π a node $f(l)$ such that:*

- *$f(l)$ is a proper ancestor of l, labelled by the same annotated hypersequent.*
- *For each step $\langle \tau \vdash H, \tau' \vdash H' \rangle$ on the path between $f(l)$ and l, it holds that τ' is not empty and the surrounding rule application is not* fc*.*
- *For some step $\langle \tau \vdash H, \tau' \vdash H' \rangle$ on the path between $f(l)$ and l it holds that (τ, τ') is an unfolding.*

An (unannotated) hypersequent H will be called $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$*-provable* if there is an $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof with root $\tau \vdash H$, where τ may be any annotation. Note that, by the availability of fc, this is equivalent to there being an $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof whose root is annotated by the empty trace value.

Definition 22. *Let (T, f) be a finite tree with back edges. The* one-step dependency order $\preceq_1$ *on* $\mathsf{ran}(f)$ *is given by:*

$$u \preceq_1 v :\Leftrightarrow u \textit{ lies on the path between } v \textit{ and } v' \textit{ for some } v' \in f^{-1}(v).$$

The dependency order $\preceq$ *on* $\mathsf{ran}(f)$ *is defined as the transitive closure of* $\preceq_1$.

For α a sequence, we let $\mathrm{Inf}(\alpha)$ denote the set of elements occurring infinitely often in α. The proof of the following lemma is omitted to conserve space.

Lemma 2. *For any infinite path α through some finite tree with back edges (T, f), the set* $\mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$ *has a $\preceq$-greatest element.*

Proposition 2. *If H is* $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$*-provable, then H is* $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$*-provable.*

Proof. Let (π, f) be an $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof with root $\tau \vdash H$. We let π_0 be the $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$-proof obtained by unravelling (π, f) and removing all annotations and applications of fc. It suffices to show that every infinite branch γ of π_0 is covered by a good trace. To that end, note that any such γ corresponds to an infinite path ρ through (π, f). Let $\preceq$ be the dependency order on $\mathsf{ran}(f)$ given in Definition 22. For any two $u, v \in \mathsf{ran}(f)$ such that $u \preceq v$, it holds that the focus rule is not applied on the path from v to u, because this path is an initial segment of the path from v to a leaf l with $f(l) = v$. By Lemma 2, the set

$\text{Inf}(\rho) \cap \mathsf{ran}(f)$ must contain a $\preceq$-greatest element u_0. It follows that from some point in ρ every node has a formula in focus and the focus rule is not applied. Since, moreover, the node u_0 is visited infinitely often, an unfolding happens infinitely often on the trace corresponding to the formulas in focus on this tail of ρ. Therefore, the infinite branch γ is covered by a good trace, as required. □

4 Soundness

This section is devoted to proving the following soundness theorem. Again, our treatment is based on [4].

Theorem 1. *Let $\mathcal{C}$ be a finite set of abstract simple frame conditions. If a hypersequent is $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_{\mathcal{C}}$-provable, then it is valid in every $\mathcal{C}$-model.*

Definition 23. *Let H be a hypersequent and $\mathbb{S}$ a Kripke model. A* countermodel state assignment (cmsa) *of H in $\mathbb{S}$ is a function $\alpha : H \to S$ assigning to each sequent σ of H a state $\alpha(\sigma)$ of $\mathbb{S}$ in which σ is not satisfied.*

Clearly for every model $\mathbb{S}$ in which H is invalid, there is a cmsa of H in $\mathbb{S}$.

Definition 24. *Let α be a cmsa of H in $\mathbb{S}$ and let $\tau := (\Box^i \boxast \psi, \sigma)$ be a non-empty trace value in H. The* weight *of τ with respect to α is given by*

$$\mu_\alpha(\tau) := \min\{n \in \omega : \mathbb{S}, \alpha(\sigma) \nVdash \Box^i \Box^n \psi\}.$$

Note that the minimum taken in the above definition always exists by the fact that α is assumed to be a cmsa.

Lemma 3. *Let H be the conclusion of an application of some rule r of $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ with premisses $H_1, \ldots, H_n$ and let $\mathbb{S}$ be a $\mathcal{C}$-model. For every cmsa α of H in $\mathbb{S}$, there is a premiss H_k and a cmsa α_k of H_k in $\mathbb{S}$ such that for every trace pair (τ, τ_k) for (H, H_k) consisting of non-empty trace values, it holds that*

$$\mu_{\alpha_k}(\tau_k) \leq \mu_\alpha(\tau),$$

and if (τ, τ_k) is an unfolding, then this inequality is strict.

Proof. For the choice of H_k and α_k we make a case distinction on the rule r of $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ that is applied. We first define α_k only on the active sequent of H_k (if it exists). Because of space issues, we only treat three cases, leaving the others to the reader.

- $\mathsf{r} = \Box$. There is a single premiss H_1 and there are two active sequents $\sigma \in H$ and $\sigma_1 \in H_1$. Moreover, the sequent σ is of the form $\Box\Gamma \Rightarrow \Box\varphi$. Since α is a cmsa, there must be some state s_1 for which it holds that $\alpha(\sigma) R s_1$ and $s_1 \nVdash \varphi$. If φ is not of the form $\boxast\psi$, we pick any such s_1 and put $\alpha_1(\sigma_1) := s_1$. If, on the other hand, the formula φ is of the form $\boxast\psi$, then we need to take a bit more care in picking the successor s_1 of $\alpha(\sigma)$. By definition, it holds that $\alpha(\sigma) \nVdash \Box\Box^{\mu_\alpha(\sigma, \Box\varphi)}\psi$. Thus it has a successor s_1 such that $s_1 \nVdash \Box^{\mu_\alpha(\sigma, \Box\varphi)}\psi$. We set $\alpha_1(\sigma_1) := s_1$.

- $\mathsf{r} = \mathsf{r}_C^{\mathsf{HK}^*}$ for some n-simple $C \in \mathcal{C}$. Let $\sigma_1, \ldots, \sigma_n$ be the active sequents in H. By the fact that $\mathbb{S}$ is a $\mathcal{C}$-model, there must be some state s of $\mathbb{S}$ and pair $(C_R, C_=) \in C$ such that for all $i \in C_R$ it holds that $\alpha(\sigma_i)Rs$ and for all $j \in C_=$ that $\alpha(\sigma_j) = s$. As premiss we pick the H_k corresponding to this $(C_R, C_=) \in C$ and we set: $\alpha_k(\bigcup_{i \in C_R} \Gamma_i', \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j) := s$.
- $\mathsf{r} = \circledast_R$. Then H has an active sequent σ of the form $\Gamma \Rightarrow \circledast\varphi, \Delta$, and there are two premisses H_1 and H_2, with as active sequent respectively $\sigma_1 = \Gamma \Rightarrow \varphi, \Delta$ and $\sigma_2 = \Gamma \Rightarrow \Box\circledast\varphi, \Delta$. If $\alpha(\sigma) \not\Vdash \varphi$, we pick H_1 and set $\alpha_1(\sigma_1) := \alpha(\sigma)$. If, on the other hand, we have $\alpha(\sigma) \not\Vdash \Box\circledast\varphi$, then we pick H_2 and set $\alpha_2(\sigma_2) := \alpha(\sigma)$.

To complete the definition of α_k, for each inactive sequent $\sigma_k \in H_k$, we put $\alpha_k(\sigma_k) := \alpha(\sigma_k)$. We leave it to the reader to verify that in each case α_k is a cmsa of H_k in $\mathbb{S}$.

It remains to verify the condition on trace pairs (τ, τ_k) for (H, H_k). First note that if τ_k is inactive, then τ and τ_k must have the same underlying sequent. By definition, it follows that $\mu_{\alpha_k}(\tau_k) = \mu_\alpha(\tau)$.

For trace pairs between active sequents, we only cover the case $\mathsf{r} = \circledast_R$, leaving the other cases to the reader. Suppose that $\langle(\varphi, \sigma), (\varphi_k, \sigma_k)\rangle$ is a trace pair for (H, H_k) such that both σ and σ_k are active sequents in an application of $\circledast_R$. By the definition of α_k given above, it holds that $\alpha(\sigma) = \alpha_k(\sigma_k)$. If $\varphi_k = \varphi$, then clearly $\mu_\alpha(\varphi, \sigma) = \mu_{\alpha_k}(\varphi_k, \sigma_k)$. If, on the other hand, the trace pair is an unfolding, then φ is active and H_k is the right-hand premiss. It follows there is some ψ such that $\varphi = \circledast\psi$ and $\varphi_k = \Box\circledast\psi$. Therefore we have that $\mu_{\alpha_k}(\varphi_k, \sigma_k) = \mu_\alpha(\varphi, \sigma) - 1 < \mu_\alpha(\varphi, \sigma)$, as required. □

Proof of Theorem 1. Suppose, towards a contradiction, that some $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_\mathcal{C}$-provable hypersequent H is $\mathcal{C}$-invalid. Then there is a cmsa α of H in some model $\mathbb{S}$. Repeatedly applying Lemma 3, we obtain a branch $H = H_0, H_1, H_2 \ldots$ in the proof of H, with for each H_i a cmsa α_i of H_i in $\mathbb{S}$. Note that this branch must be infinite, for otherwise the final H_i is an axiom, contradicting the fact that it has a cmsa. Moreover, by the condition of infinite branches it contains a good trace $\overline{\tau}$ which from some point, say, from the hypersequent H_k, contains no empty trace values. By construction, we have $\mu_{\alpha_k}(\tau_0) \leq \mu_{\alpha_{k+1}}(\tau_1) \leq \mu_{\alpha_{k+2}}(\tau_2) \leq \ldots$ and, since infinitely many unfoldings occur on $\overline{\tau}$, this inequality is strict infinitely often. Clearly we have reached the desired contradiction.

Question 1. Suppose we weaken Condition 2 of Definition 15 to allow σ to be inactive, provided that it is equal to σ'. The pair (τ_1, τ_6) of Remark 1 then becomes a trace pair. Is the system $\mathsf{HK}^*_{\mathsf{inf}} + \mathsf{R}_\mathcal{C}$ still sound for any finite $\mathcal{C}$?

5 Completeness

In this section we prove cut-free completeness for $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_\mathcal{C}$, where $\mathcal{C}$ is any finite set of equable frame conditions. Our method is an adaptation of the one in [10]. We close the section with a brief explanation for why the more common method of completeness via infinitary proof search is hard to apply to our hypersequent calculi.

5.1 Completeness of $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ for Equable $\mathcal{C}$

This subsection will be devoted to proving the following theorem, which is the main theorem of this paper.

Theorem 2. *Let $\mathcal{C}$ be a finite set of equable frame conditions. If a hypersequent is valid in every $\mathcal{C}$-model, then it has a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof.*

We will prove this theorem by constructing a countermodel for each unprovable hypersequent. For Γ a set of formulas, we define $\Box^{-1}\Gamma := \{\varphi \mid \Box\varphi \in \Gamma\}$.

Definition 25. *Let H be a hypersequent. The* canonical model $\mathbb{S}^H$ *for H is the model (S, R, V) given by:*

- $S := H$.
- $\Gamma_1 \Rightarrow \Delta_1 R \Gamma_2 \Rightarrow \Delta_2 :\Leftrightarrow \Box^{-1}\Gamma_1 \subseteq \Gamma_2$.
- $V(p) := \{\Gamma \Rightarrow \Delta \mid p \in \Gamma\}$.

The key property of canonical models is that, for certain unprovable hypersequents H, they satisfy a *Truth Lemma*, with the consequence that H is invalid in the canonical model $\mathbb{S}^H$ of H. The bulk of this subsection concerns constructing such unprovable hypersequents and establishing the Truth Lemma.

Definition 26. *Let Σ be a finite closed set of formulas. An (annotated) (hyper)sequent is said to be a Σ-*(annotated) (hyper)sequent *if it contains only formulas from Σ.*

For the rest of this section we assume an arbitrary finite closed set of formulas Σ. First, we want our unprovable hypersequent to satisfy the following saturation properties.

Definition 27. *A sequent $\Gamma \Rightarrow \Delta$ is said to be* propositionally saturated *if the following closure conditions hold:*

(i) $\bot \notin \Gamma$.
(ii) $\Gamma \cap \Delta = \emptyset$.
(iii) *If $\varphi_1 \to \varphi_2 \in \Gamma$, then $\varphi_2 \in \Gamma$ or $\varphi_1 \in \Delta$.*
(iv) *If $\varphi_1 \to \varphi_2 \in \Delta$, then $\varphi_1 \in \Gamma$ and $\varphi_2 \in \Delta$.*
(v) *If $\boxast\varphi \in \Gamma$, then $\varphi \in \Gamma$ and $\Box\boxast\varphi \in \Gamma$.*
(vi) *If $\boxast\varphi \in \Delta$, then $\varphi \in \Delta$ or $\Box\boxast\varphi \in \Delta$.*

A hypersequent is propositionally saturated *whenever each of its sequents is.*

Definition 28. *Let $\mathcal{C}$ be a finite set of abstract simple frame conditions. A hypersequent H is said to be $\mathcal{C}$-*presaturated *if $\mathbb{S}^H$ is a $\mathcal{C}$-model. If, moreover, the hypersequent H is propositionally saturated, it is be said to be $\mathcal{C}$-*saturated.

An annotated hypersequent will be called *$\mathcal{C}$-(pre)saturated* whenever the underlying hypersequent is.

Definition 29. *Let $\Gamma_1 \Rightarrow \Delta_1$ and $\Gamma_2 \Rightarrow \Delta_2$ be sequents and let H_1, H_2 be hypersequents. We define:*

- $\Gamma_1 \Rightarrow \Delta_1 \sqsubseteq \Gamma_2 \Rightarrow \Delta_2$ *if* $\Gamma_1 \subseteq \Gamma_2$ *and* $\Delta_1 \subseteq \Delta_2$.
- $H_1 \sqsubseteq H_2$ *if for all* $\sigma_1 \in H_1$, *there is some* $\sigma_2 \in H_2$ *such that* $\sigma_1 \sqsubseteq \sigma_2$.

If two (hyper)sequents are related by $\sqsubseteq$, *we say that the former is* encompassed *by the latter.*

The following definition and lemmas are based on the notion of a *propositional retract* in [5].

Definition 30. *Let $\mathcal{C}$ be a finite set of simple frame conditions. A* retract *of an annotated Σ-hypersequent $(\varphi, \sigma) \vdash H \mid \sigma$ is a finite set $\mathcal{H}$ consisting of annotated Σ-hypersequents of the form $(\varphi, \sigma') \vdash H \mid \sigma'$ with $\sigma \sqsubseteq \sigma'$, such that $(\varphi, \sigma) \vdash H \mid \sigma$ is derivable from $\mathcal{H}$ in $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ without using the rules $\Box$,* fc, *and* cut. *Moreover, the retract $\mathcal{H}$ is said to be $\mathcal{C}$-*saturated *if for every $(\varphi, \sigma') \vdash H' \in \mathcal{H}$ such that $\sigma' \sqsubseteq H$ it holds that H' is $\mathcal{C}$-saturated.*

The following crucial lemma is the only part of the completeness proof where we rely on the restriction to equable frame conditions.

Lemma 4. *Let $\mathcal{C}$ be a finite set of equable frame conditions and H a $\mathcal{C}$-saturated hypersequent. Then any annotated Σ-hypersequent $(\varphi, \sigma) \vdash H \mid \sigma$ has a $\mathcal{C}$-saturated retract.*

Proof. We say that a retract $\mathcal{H}$ is *$\mathcal{C}$-presaturated (propositionally saturated)* if for every $(\varphi, \sigma') \vdash H' \in \mathcal{H}$ such that $\sigma' \sqsubseteq H$ it holds that H' is $\mathcal{C}$-presaturated (propositionally saturated). The proof rests on the following two claims.

1. Any annotated Σ-hypersequent has a propositionally saturated retract.
2. For $\mathcal{C}$ a finite set of equable frame conditions, and H a $\mathcal{C}$-presaturated hypersequent, any annotated Σ-hypersequent $(\varphi, \sigma) \vdash H \mid \sigma$ has a $\mathcal{C}$-presaturated retract.

The proof of Claim 1 is analogous to the proof of Lemma 6.1 in [5]. For Claim 2, we argue by induction on the number of Σ-formulas not occurring in the sequent σ. If $\sigma \not\sqsubseteq H$ or if $H \mid \sigma$ is already $\mathcal{C}$-presaturated, then we simply set $\mathcal{H} := \{H \mid \sigma\}$.

Now suppose, towards a contradiction, that $\sigma \sqsubseteq H$ and $H \mid \sigma$ is not $\mathcal{C}$-presaturated, that is:

$$\text{There is an } n\text{-simple } C \in \mathcal{C} \text{ and a list } (\Gamma_k \Rightarrow \Delta_k)_{1 \le k \le n} \text{ of sequents in } H \cup \{\sigma\} \text{ such that for every } \Gamma \Rightarrow \Delta \in H \cup \{\sigma\} \text{ and } (C_R, C_=) \in C \text{ there is an } i \in C_R \text{ s.t. } \Box^{-1}\Gamma_i \not\subseteq \Gamma \text{ or a } j \in C_= \text{ s.t. } \Gamma_j \Rightarrow \Delta_j \neq \Gamma \Rightarrow \Delta. \quad (1)$$

For the rest of this proof we fix a condition $C \in \mathcal{C}$ and a list $(\Gamma_k \Rightarrow \Delta_k)_{1 \le k \le n}$ that witness (1).

Since $\sigma \sqsubseteq H$, there is a sequent $\overline{\sigma} \in H$ such that $\sigma \sqsubseteq \overline{\sigma}$. Let $(\overline{\Gamma_k} \Rightarrow \overline{\Delta_k})_{1 \le k \le n}$ be the list obtained by replacing in $(\Gamma_k \Rightarrow \Delta_k)_{1 \le k \le n}$ each occurrence of σ by $\overline{\sigma}$.

By the $\mathcal{C}$-presaturation of H, there must be some $(C_R, C_=) \in C$ and $\overline{\Gamma} \Rightarrow \overline{\Delta} \in H$ such that it holds for each $i \in C_R$ that $\Box^{-1}\overline{\Gamma_i} \subseteq \overline{\Gamma}$ and for each $j \in C_=$ that $\overline{\Gamma_j} \Rightarrow \overline{\Delta_j} = \overline{\Gamma} \Rightarrow \overline{\Delta}$.

It follows for every $i \in C_R$ that $\Box^{-1}\Gamma_i \subseteq \Box^{-1}\overline{\Gamma_i} \subseteq \overline{\Gamma}$. Thus, by the fact that $H \mid \sigma$ is not $\mathcal{C}$-presaturated, there must be some $k \in C_=$ such that $\Gamma_k \Rightarrow \Delta_k \neq \overline{\Gamma_k} \Rightarrow \overline{\Delta_k}$. By construction this can only be the case if $\Gamma_k \Rightarrow \Delta_k = \sigma$.

Now consider the following inference.

$$\mathsf{r}_C^{\mathsf{HK}^*} \; \frac{\{H \mid \bigcup_{i \in C_R} \Box^{-1}\Gamma_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow [\varphi], \bigcup_{j \in C_=} \Delta_j : (C_R, C_=) \in C\}}{H \mid \Gamma_k \Rightarrow [\varphi], \Delta_k}$$

Observe that the right-hand side of each premiss contains φ. The reason is that φ belongs to the right-hand side of $\sigma = \Gamma_k \Rightarrow \Delta_k$ and, by equability, $k \in C_=$ for every $(C_R, C_=) \in C$. We claim that for any $(C_R, C_=) \in C$, the Σ-sequent

$$\sigma_R := \bigcup_{i \in C_R} \Box^{-1}\Gamma_i \cup \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j$$

is such that $\sigma \sqsubset \sigma_R$, whence contains strictly less Σ-formulas than σ. Note that no information is lost by indexing σ_R solely by R, since, by equability, for each $(C_R^1, C_=^1), (C_R^2, C_=^2) \in C$ it holds that $C_=^1 = C_=^2$. Since $k \in C_=$, we already have $\sigma \sqsubseteq \sigma_R$. Now suppose, towards a contradiction, that $\sigma = \sigma_R$. Then by (1), there must be some $j \in C_=$ such that $\Gamma_j \Rightarrow \Delta_j \neq \sigma$. It follows that

$$\begin{aligned} \Gamma_j \Rightarrow \Delta_j &= \overline{\Gamma_j} \Rightarrow \overline{\Delta_j} && \text{(because } \Gamma_j \Rightarrow \Delta_j \neq \sigma) \\ &= \overline{\Gamma_k} \Rightarrow \overline{\Gamma_k} && \text{(because } j, k \in C_=) \\ &= \overline{\sigma}. \end{aligned}$$

But then $\overline{\sigma} \sqsubseteq \sigma$, so $\sigma = \overline{\sigma}$ and $H \mid \sigma = H$, contradicting the assumption that $H \mid \sigma$ is not $\mathcal{C}$-presaturated.

Finally, the induction hypothesis gives, for each $(C_R, C_=) \in C$, a suitable retract $\mathcal{H}_R$ of $(\varphi, \sigma_R) \vdash H \mid \sigma_R$. We put:

$$\mathcal{H} := \bigcup_{(C_R, C_=) \in C} \mathcal{H}_R,$$

which finishes the proof of Claim 2.

The main statement of the lemma can now be proven from claims 1 and 2 by a straightforward induction. □

Definition 31. *Let $\mathcal{C}$ be some finite set of abstract simple frame conditions. A Σ-hypersequent H is called* $\mathcal{C}$-maximal *if the following hold:*

(i) There is no cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof of H.
(ii) H is $\mathcal{C}$-saturated.
(iii) H is $\subseteq$-maximal as a Σ-hypersequent satisfying both (i) and (ii).
(iv) For every Σ-sequent σ:

Either $\sigma \sqsubseteq H$ or there is a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof of $H \mid \sigma$.

Because of space limitations the proof of the following lemma is only sketched.

Lemma 5. *Let $\mathcal{C}$ be some finite set of abstract simple frame conditions. Then any hypersequent H which has no cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_\mathcal{C}$-proof, can be $\sqsubseteq$-extended to be $\mathcal{C}$-maximal.*

Proof (sketch). In the same way as one can prove Lemma 2 of [10], it can be shown that there is a Σ-hypersequent H_0 such that $H \sqsubseteq H_0$ and H_0 satisfies conditions (i) and (iv) of $\mathcal{C}$-maximality. Using similar arguments as in the proof of Theorem 3 of [10], it can then be shown that H_0 also satisfies condition (ii). Finally, taking a $\subseteq$-maximal extension of H_0 with respect to conditions (i) and (ii) breaks neither condition (iv) nor the encompassing of H. $\square$

We will prove our Truth Lemma for the canonical models of $\mathcal{C}$-maximal hypersequents. We first prove the following existence lemma.

Lemma 6. *For $\mathcal{C}$ a finite set of equable simple frame conditions, let H be a $\mathcal{C}$-maximal Σ-hypersequent, and let $\mathbb{S}$ be its canonical model. Then for every sequent $\sigma := \Gamma \Rightarrow \Delta \in H$ the following hold:*

(i) For all $\Box\psi \in \Sigma$:
If $\Box\psi \in \Delta$, then there is $\sigma' := \Gamma' \Rightarrow \Delta' \in H$ such that $\sigma R \sigma'$ and $\psi \in \Delta'$.
(ii) For all $\circledast\psi \in \Sigma$:
If $\circledast\psi \in \Delta$, then there is $\sigma' := \Gamma' \Rightarrow \Delta' \in H$ such that $\sigma R^* \sigma'$ and $\psi \in \Delta'$.

Proof. We leave the proof of item (i) to the reader. For item (ii), define

$$\mathcal{S} := \{\Gamma' \Rightarrow \Delta' \in H : \sigma R^* \Gamma' \Rightarrow \Delta' \text{ and, } \psi \in \Delta' \text{ or } \Box\circledast\psi \in \Delta'\}.$$

We must show that $\mathcal{S}$ contains a sequent $\Gamma' \Rightarrow \Delta'$ with $\psi \in \Delta'$. Assume that this is not the case. We will reach a contradiction by constructing a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_\mathcal{C}$-proof (π, f) of H.

Since $\sigma \in \mathcal{S}$, we have $\Box\circledast\psi \in \Delta$ by our assumption. We begin the construction of (π, f) as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\overset{(\pi_1, f_1)}{H \mid \Box^{-1}\Gamma \Rightarrow \psi} \qquad \overset{\pi_2}{H \mid \Box^{-1}\Gamma \Rightarrow [\Box\circledast\psi]}}{H \mid \Box^{-1}\Gamma \Rightarrow [\circledast\psi]}\ \circledast_R}{H \mid \Box\Box^{-1}\Gamma \Rightarrow [\Box\circledast\psi]}\ \Box}{\vdots}\ \mathsf{iw}_L}{H \mid \Gamma \Rightarrow [\Box\circledast\psi]}\ \mathsf{iw}_L}{\vdots}\ \mathsf{iw}_R}{H \mid \Gamma \Rightarrow \Delta, [\Box\circledast\psi]}\ \mathsf{iw}_R$$

The cut-free $\mathsf{HK}^* + \mathsf{R}_\mathcal{C}$ proof (π_1, f_1) is obtained by the $\mathcal{C}$-maximality of H and the fact that $\Box^{-1}\Gamma \Rightarrow \psi \not\sqsubseteq H$. The latter must be the case, for otherwise there would be a $\Gamma_1 \Rightarrow \Delta_1 \in H$ such that $\Box^{-1}\Gamma \Rightarrow \psi \sqsubseteq \Gamma_1 \Rightarrow \Delta_1$. But that would mean that $\Gamma_1 \Rightarrow \Delta_1 \in \mathcal{S}$ with $\psi \in \Delta_1$, which we assumed not to be the case.

We invoke Lemma 4 to obtain a retract $\mathcal{H}$ of $H \mid \Box^{-1}\Gamma \Rightarrow \Box\boxtimes\psi$ and let π_2 be the derivation of this hypersequent from $\mathcal{H}$. By construction, every annotated hypersequent in $\mathcal{H}$ is of the form $(\Box\boxtimes\psi, \sigma') \vdash H \mid \sigma'$ where $\sigma' \sqsupseteq \Box^{-1}\Gamma \Rightarrow \Box\boxtimes\psi$. Furthermore, the sequent σ' is such that $\sigma' \not\sqsubseteq H$ or $H \mid \sigma'$ is $\mathcal{C}$-saturated. By the $\mathcal{C}$-maximality of H this means that either $H \mid \sigma'$ has a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof, or $\sigma' \in H$.

To every leaf of π_2 that has a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof, we append that proof. Observe that any other leaf is of the form $(\Box\boxtimes\psi, \sigma') \vdash H \mid \sigma'$ for some $\sigma' \in \mathcal{S}$. To each such leaf we recursively apply the above procedure. By the finiteness of $\mathcal{S}$, every branch created in this way must at some point encounter the same annotated hypersequent $(\Box\boxtimes\psi, \sigma') \vdash H \mid \sigma'$ twice, for some in $\sigma' \in \mathcal{S}$. Whenever that happens, we add a back edge from the second encounter to the first and terminate the procedure for this branch. Notice that between target of the newly added back edge and its source the focus rule is not applied, there is always a formula in focus, and at least one unfolding happens on the induced trace.

After finitely many steps this procedure terminates for every branch and we obtain a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof of H, giving the desired contradiction. □

The following Truth Lemma is now proven using a straightforward induction, which we leave to the reader.

Lemma 7. *Let $\mathbb{S}^H$ be the canonical model for some $\mathcal{C}$-maximal Σ-hypersequent H. Then for all $\sigma := \Gamma \Rightarrow \Delta \in S^H$ and $\varphi \in \Sigma$ the following hold:*

(a) If $\varphi \in \Gamma$, then $\mathbb{S}^H, s \Vdash \varphi$.
(b) If $\varphi \in \Delta$, then $\mathbb{S}^H, s \nVdash \varphi$.

Proof of Theorem 2. We argue by contraposition. Suppose H has no cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof. Let Σ be a finite closed set such that H is a Σ-sequent. By Lemma 5, there is a $\mathcal{C}$-maximal Σ-hypersequent H_0 encompassing H.

We claim that the canonical model $\mathbb{S}^{H_0}$ for H_0 is a countermodel to H. Indeed, let $\sigma \in H$, then there is $\sigma_0 := \Gamma_0 \Rightarrow \Delta_0 \in H_0$ such that $\sigma \sqsubseteq \sigma_0$. By Lemma 7, we have for each $\varphi \in \Gamma_0$ that $\mathbb{S}^{H_0}, s_0 \Vdash \varphi$ and for each $\psi \in \Delta_0$ that $\mathbb{S}^{H_0}, s_0 \nVdash \psi$. Thus σ_0 is not valid in $\mathbb{S}^{H_0}$, and the same holds for σ. Since σ was taken arbitrarily, the hypersequent H is not valid in $\mathbb{S}^{H_0}$.

Finally, the result follows that fact that, by $\mathcal{C}$-saturation, the model $\mathbb{S}^{H_0}$ is a $\mathcal{C}$-model. □

Question 2. For which other finite sets $\mathcal{C}$ of simple (not necessarily equable) frame conditions is $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$ cut-free complete?

Corollary 1. *For $\mathcal{C}$ a finite set of equable frame conditions, the logic obtained by interpreting ML^* on the class of $\mathcal{C}$-frames is decidable.*

Proof. Let H be an arbitrary hypersequent. Then H is a Σ-hypersequent for some finite closed set Σ. If H is invalid in some $\mathcal{C}$-model, then, by Theorem 1, it follows that H cannot have a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof. By the same reasoning as in the proof of Theorem 2 we obtain a model $\mathbb{S}^{H_0}$ in which H is not valid. Since the size of this model is by construction bounded by the size of Σ, we can decide the $\mathcal{C}$-validity of H by checking its validity in finitely many models.

Question 3. Is the size of the smallest cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof of some $\mathcal{C}$-valid Σ-hypersequent also bounded by the size of Σ? We conjecture that this question can be answered positively by showing that for every cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof, there is a cut-free $\mathsf{HK}^*_{\mathsf{circ}} + \mathsf{R}_{\mathcal{C}}$-proof of the same hypersequent, with the property that every branch contains at most one annotated hypersequent twice (in which case these two occurrences are connected by a back edge) and no annotated hypersequent more than twice.

5.2 Completeness via (Infinitary) Proof Search

A standard method for proving completeness of non-well-founded proof systems is via infinitary proof search. Roughly, the idea is to find some proof-search strategy such that a countermodel can be extracted from a failed attempt, *i.e.* an attempt that does not yield a proof. Then, since soundness entails that the proof-search must fail for any invalid hypersequent, completeness follows. This is also the method used to prove completeness in [4].

In this subsection we briefly sketch a complication that arises when one tries to apply this method to our hypersequent calculi. Because this already occurs in the case of $\mathsf{HK}^*_{\mathsf{inf}}$ (without additional rules for frame conditions), we restrict our attention to this system.

Suppose we obtain a pre-proof π from the failure of an application of some proof-search strategy for $\mathsf{HK}^*_{\mathsf{inf}}$ to the hypersequent H. The problem arises in the case that π is infinite. In this case we wish to use the fact that π has a branch β which is *not* covered by a good trace, in order to extract a countermodel. One might for example try to take the canonical model $\mathbb{S}^H$ of some hypersequent H that occurs infinitely often on β. To prove an analogue of part (ii) of Lemma 6 for $\mathbb{S}^H$, we would have to show that any $\boxplus\psi \in \Delta$ for some $\Gamma \Rightarrow \Delta$ in H is not unfolded infinitely often on β. The proof-search strategy would then ensure that at some point in the branch β the rule $\boxplus_R$ is applied to $\boxplus\psi$ in $\Gamma \Rightarrow \Delta$ and the branch β continues through the premiss on the left-hand side. This would give us a state $\Gamma' \Rightarrow \Delta'$ in $\mathbb{S}^H$ such that $\Gamma \Rightarrow \Delta R^* \Gamma' \Rightarrow \Delta'$ and $\psi \in \Delta'$. The problem is that we cannot guarantee that $\boxplus\psi$ is not unfolded infinitely often, because we might repeatedly lose its trace due to that trace being overtaken by some other active sequent (cf. Remark 1 and Question 1).

6 Conclusion

In this paper we have constructed sound and complete infinitary and cyclic proof systems for ML^* interpreted on any frame class characterised by a finite number of equable frame conditions.

In future work we wish to extend this to non-equable frame conditions. We conjecture that there are cases in which the single focus-style annotations are not sufficient, and one must turn a more complex annotating mechanism.

We would also like to extend this work to more expressive fragments of the modal μ-calculus, such as polymodal logic with the master modality, PDL, the alternation-free modal μ-calculus, or even the modal μ-calculus itself.

Another avenue for further research is to see whether our hypersequent calculi can be used to establish Craig interpolation for their respective logics.

Finally, we wish to combine non-well-founded proof theory with other enrichments of ordinary Gentzen sequents, such as nested sequents. It would be interesting to better understand which of such systems combine well with non-well-founded proof theory and why.

References

1. Afshari, B., Leigh, G.E.: Cut-free completeness for modal mu-calculus. In: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science. LICS 2017 (2017)
2. Brotherston, J., Simpson, A.: Sequent calculi for induction and infinite descent. J. Log. Comput. **21**(6), 1177–1216 (2011)
3. Brünnler, K., Lange, M.: Cut-free sequent systems for temporal logic. J. Log. Algebr. Program. **76**(2), 216–225 (2008)
4. Docherty, S., Rowe, R.N.S.: A non-wellfounded, labelled proof system for propositional dynamic logic. In: Cerrito, S., Popescu, A. (eds.) TABLEAUX 2019. LNCS (LNAI), vol. 11714, pp. 335–352. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29026-9_19
5. Doczkal, C., Smolka, G.: Constructive completeness for modal logic with transitive closure. In: Hawblitzel, C., Miller, D. (eds.) CPP 2012. LNCS, vol. 7679, pp. 224–239. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35308-6_18
6. Enqvist, S.: A circular proof system for the hybrid mu-calculus. arXiv preprint arXiv:2001.04971 (2020)
7. Enqvist, S., Hansen, H.H., Kupke, C., Marti, J., Venema, Y.: Completeness for game logic. In: 2019 34th Annual ACM/IEEE Symposium on Logic in computer Science (LICS), pp. 1–13. IEEE (2019)
8. Indrzejczak, A.: Sequents and Trees: An Introduction to the Theory and Applications of Propositional Sequent Calculi. Springer Nature, Basingstoke (2020)
9. Kikot, S., Shapirovsky, I., Zolin, E.: Completeness of logics with the transitive closure modality and related logics. arXiv preprint arXiv:2011.02205 (2020)
10. Lahav, O.: From frame properties to hypersequent rules in modal logics. In: 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science, pp. 408–417. IEEE (2013)
11. Lange, M., Stirling, C.: Focus games for satisfiability and completeness of temporal logic. In: Proceedings 16th Annual IEEE Symposium on Logic in Computer Science, pp. 357–365. IEEE (2001)
12. Shamkanov, D.S.: Circular proofs for the Gödel-Löb provability logic. Math. Notes **96**(3), 575–585 (2014)