An analytic proof system for common knowledge logic over S5

Jan Rooduijn [1]

ILLC, University of Amsterdam

Lukas Zenger [2]

Institute of Computer Science, University of Bern

Abstract

In this paper we present an analytic proof system for multi-modal logic with common-knowledge over S5 (called S5-CKL). The system is an annotated cyclic calculus manipulating two-sided Gentzen sequents and extending a known system for multi-modal S5. First a direct argument is used to show that the system is sound. Using a canonical model construction, we then show that the system is analytically complete. In particular, the use of the cut-rule is restricted to analytic cuts. Exploiting this analyticity, we then reduce the provability problem of a given sequent to the problem of solving a certain parity game. As a consequence we obtain an optimal decision procedure for proof search and thereby for the validity problem of S5-CKL.

Keywords: Common knowledge, S5 multi-modal logic, analytic proof systems, cyclic proofs, proof search games.

1 Introduction

Common knowledge is an important notion of group knowledge with applications ranging from philosophy to computer science. A proposition p is *common knowledge* in a group of agents if all agents know that p, all agents know that all agents know that p, and so on. In a formal setting, common knowledge is often studied using a logic called *common knowledge logic* (CKL, for short). The logic CKL is an extension of multi-agent modal logic by a fixed point operator $\boxtimes$ meant to express common knowledge. As such it is a fragment of the alternation-free modal μ-calculus [4]. CKL was introduced in 1990 by Halpern and Moses [12]. For an introduction to logics of common knowledge in general we refer the reader to [22].

[1] J.M.W. Rooduijn is supported by a grant from the Dutch Research Council NWO, project nr. 617.001.857.

[2] L. Zenger is supported by the Swiss National Science Foundation Grant 200021L_196176 Proof and Model Theory of Intuitionistic Temporal Logic.

Most commonly CKL is axiomatised using a Hilbert-style system for multimodal logic, extended by fixed point axioms expressing that $\boxplus$ is a greatest fixed point (see for example [22]). However, axiomatisiations of this kind suffer from the usual drawback of Hilbert-style systems: the presence of the *modus ponens* rule frustrates proof-theoretic analysis.

The common solution to this problem is to construct a sequent system with restricted applications of the *cut* rule, but this has been proven difficult for the logic CKL. When the base modal logic is required to satisfy certain frame conditions (motivated for example by the study of distributed systems), the difficulty often further increases. It is for example an open question whether there exists a finite and *analytic* sequent calculus for the logic CKL interpreted over S5-frames (this logic will be called S5-CKL). As usual, we call a calculus analytic whenever every valid sequent admits a finite proof containing only formulas in some sense relevant to the endsequent. In the context of modal fixed point logics one usually counts as relevant the formulas in the endsequent's *Fischer-Ladner closure* (originally defined in [10]). Notable sequent calculi for S5-CKL have been constructed for example by Alberucci & Jäger [5] and Hill & Poggiolesi [14], but although both are finite, neither are analytic. A description of these systems, as well as a comparison between them and the present paper, is postponed to Section 7 below.

More related work exists in the area of tableau-based decision procedures. In [3], Ajspur et al. give such a procedure for S5-CKL and several of its extensions. While their procedure is analytic, the fact that it requires multiple passes makes it unclear how one could extract an analogous sequent-style proof system. A single pass tableau-based decision procedure for CKL is given by Abate et. al in [1], but only for the interpretation of CKL over the class of all frames.

In the first part of this paper we give a positive answer to the question of whether S5-CKL admits a finite and analytic sequent calculus. To that end we present the cyclic sequent calculus $\mathsf{sCKL_f}$. Instead of an induction rule as in [5] and [14], $\mathsf{sCKL_f}$ uses cycles to characterise $\boxplus$ as a greatest fixed point. These cycles are handled by a focus mechanism, a technique originally proposed by Lange & Stirling in [16]. Roughly, a cyclic branch of some proof is deemed valid whenever some form of progress is made. The soundness argument then exploits the fact that this form of progress cannot be made infinitely often. In our case we show the soundness of $\mathsf{sCKL_f}$ by a minimal countermodel approach that is sometimes found in the literature (see *e.g.* [21]). For completeness we use a canonical model construction that is similar to the construction in [20]. Importantly, in the completeness proof the cut rule is only applied to formulas in the Fischer-Ladner closure of the endsequent. As every other rule of $\mathsf{sCKL_f}$ enjoys the subformula property (or its analogue for the Fischer-Ladner closure), we obtain that the calculus is analytic. The approach of using cyclic proof systems for common knowledge logic has also been been taken by Ricardo Wehbe in his PhD Thesis [23], but, like Abate et. al, he does not consider the restriction to S5-frames.

In the second part of the paper we show that $\mathsf{sCKL_f}$ is suitable for proof search. To that end we translate the calculus into a parity game which is played by a player called Prover, who tries to show that a given sequent is derivable in $\mathsf{sCKL_f}$, and a player called Refuter who tries to show the opposite. We then show that a sequent is derivable in $\mathsf{sCKL_f}$ if and only if Prover has a memoryless winning strategy in the corresponding game. Next, we use a result proven by Calude et al. in [8] to establish the existence of an efficient algorithm which computes the winner of each game. Finally, by combining these two results we establish the existence of an algorithm deciding whether a given sequent is derivable in $\mathsf{sCKL_f}$, which runs in optimal time.

The paper is structured as follows: in Section 2 we introduce the basic definitions and notations for S5-CKL, including its syntax and semantics. In Section 3 we define the sequent calculus $\mathsf{sCKL_f}$. Sections 4 and 5 are devoted to prove soundness and completeness, respectively. In Section 6 we define the aforementioned proof search game and establish that there exists and optimal proof search algorithm. Finally, in Section 7 we compare our work to related work of different authors and sketch some ideas for further research.

2 Basic definitions

The language of CKL consists of a finite set of atomic agents $\mathcal{A} = \{1, \dots n\}$ and a countably infinite set of atomic propositions P.

Definition 2.1 *Formulas* φ, ψ of CKL are inductively defined as follows:

$$\varphi, \psi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box_i \varphi \mid \boxplus \varphi$$

where $p \in \mathsf{P}$ and $i \in \mathcal{A}$. The set of CKL-formulas is denoted by Fm.

We will use the abbreviation $\Box\varphi :\equiv \bigwedge_{i \in \mathcal{A}} \Box_i \varphi$. Furthermore, we define $\Box^0\varphi := \varphi$ and $\Box^{n+1}\varphi := \Box\Box^n\varphi$. The expression $\Box_i^n\varphi$ is defined analogously.

Definition 2.2 An *epistemic Kripke model* is a tuple $\mathbb{S} = (S, \{R_i \mid i \in \mathcal{A}\}, V\}$ where

- S is a non-empty set;
- R_i is an equivalence relation on S for each $i \in \mathcal{A}$;
- V is a function $S \to \mathcal{P}(P)$.

Elements in S are called *states.* The binary relations R_i are called *transition relations* and V is called a *valuation.* Instead of $(s,t) \in R_i$ we usually write sR_it.

Formulas of CKL are evaluated in epistemic Kripke models as follows:

Definition 2.3 Let $\mathbb{S} = (S, \{R_i \mid i \in \mathcal{A}\}, V\}$ be an epistemic Kripke model and $s \in S$ be a state. The relation $\Vdash \subseteq S \times Fm$ is inductively defined by:

$$\begin{array}{lcl}
\mathbb{S}, s \Vdash p & \Leftrightarrow & p \in V(s) \\
\mathbb{S}, s \Vdash \neg\varphi & \Leftrightarrow & \mathbb{S}, s \nVdash \varphi \\
\mathbb{S}, s \Vdash \varphi \wedge \psi & \Leftrightarrow & \mathbb{S}, s \Vdash \varphi \text{ and } \mathbb{S}, s \Vdash \psi \\
\mathbb{S}, s \Vdash \Box_i\varphi & \Leftrightarrow & \text{for all } t \in S \text{ with } sR_it, \text{ it holds that } \mathbb{S}, t \Vdash \varphi \\
\mathbb{S}, s \Vdash \boxplus\varphi & \Leftrightarrow & \text{for all } t \in S \text{ with } sR^*t, \text{ it holds that } \mathbb{S}, t \Vdash \varphi
\end{array}$$

where R^* is the transitive closure of the relation $\bigcup_{i \in \mathcal{A}} R_i$.

If $\mathbb{S}, s \Vdash \varphi$, then we say that φ is *true* or *holds* at state s of the epistemic Kripke model $\mathbb{S}$. A formula φ is *satisfiable* if there exists an epistemic Kripke model $\mathbb{S}$ and a state s such that $\mathbb{S}, s \Vdash \varphi$, and *unsatisfiable* if not satisfiable. A formula φ is *valid* if $\mathbb{S}, s \Vdash \varphi$ for every epistemic Kripke model $\mathbb{S}$ and every state s in $\mathbb{S}$, and *invalid* if not valid.

Definition 2.4 The *Fischer-Ladner closure* of a formula φ is the smallest set of formulas $Cl(\varphi)$ which contains φ and is closed under the following conditions:

- $\neg\psi \in Cl(\varphi)$ implies $\psi \in Cl(\varphi)$;
- $\psi_1 \wedge \psi_2 \in Cl(\varphi)$ implies $\psi_k \in Cl(\varphi)$ for each $k \in \{1, 2\}$;
- $\Box_i \psi \in Cl(\varphi)$ implies $\psi \in Cl(\varphi)$;
- $\boxplus\psi \in Cl(\varphi)$ implies $\psi \in Cl(\varphi)$ and $\{\Box_i \boxplus \psi \mid i \in \mathcal{A}\} \subseteq Cl(\varphi)$.

We will usually denote $Cl(\varphi)$ simply as the *closure* of φ. The definition of the closure of a formula is extended to the definition of the *closure* of a set of formulas A as follows:

$$Cl(A) := \bigcup \{Cl(\varphi) \mid \varphi \in A\}.$$

A set A of formulas is called *closed* whenever $Cl(A) = A$.

3 An annotated sequent system

An *annotated formula* is a pair (φ, a), usually written φ^a, where φ is a formula and a is either u (designating that the formula is *unfocussed*) or f (designating that the formula is *in focus*). A *sequent* is a pair (Γ, Δ), usually written $\Gamma \Rightarrow \Delta$, where Γ and Δ are finite sets of annotated formulas. In later proofs we will often denote sequents using the Greek letter σ, possibly with subscript.

We will only consider sequents whose formulas are annotated in a very specific way. Namely, if $\Gamma \Rightarrow \Delta$ is a sequent, we require that every formula in Γ is unfocussed, and at most one formula in Δ is in focus. Moreover, if Δ contains a formula that is in focus, then it must be of the form $\boxplus\varphi$ or of the form $\Box_i \boxplus \varphi$. To emphasise that we are only considering sequents of this restricted form, we will also refer to them as CKL-*sequents*.

For Γ a finite set of annotated formulas, we define the following abbreviations:

$$\Gamma^u := \{\varphi^u \mid \varphi^a \in \Gamma\}, \qquad \Box_i \Gamma := \{\Box_i \varphi^a \mid \varphi^a \in \Gamma\}$$
$$\Gamma^- := \{\varphi \mid \varphi^a \in \Gamma\}, \qquad \Box_i^{-1} \Gamma := \{\varphi^a \mid \Box_i \varphi^a \in \Gamma\}$$

The following definition presents our proof system. Note that basic modal part, *i.e.* the part without the rules $\boxplus_L$, $\boxplus_R$, U, F, is based on a standard system for the modal logic $\mathsf{S5}$. This system was originally presented by Ohnishi and Matsumoto in 1957 [19].

Definition 3.1 The sequent calculus $\mathsf{sCKL_f}$ manipulates CKL-sequents by the following axioms and rules.

$$\mathsf{id}\ \frac{}{\varphi^u \Rightarrow \varphi^a} \qquad \mathsf{w}_L\ \frac{\Gamma \Rightarrow \Delta}{\Gamma, \varphi^u \Rightarrow \Delta} \qquad \mathsf{w}_R\ \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \varphi^a, \Delta}$$

$$\neg_L\ \frac{\Gamma \Rightarrow \varphi^u, \Delta}{\Gamma, \neg\varphi^u \Rightarrow \Delta} \qquad \neg_R\ \frac{\Gamma, \varphi^u \Rightarrow \Delta}{\Gamma \Rightarrow \neg\varphi^u, \Delta}$$

$$\wedge_L\ \frac{\Gamma, \varphi^u, \psi^u \Rightarrow \Delta}{\Gamma, (\varphi \wedge \psi)^u \Rightarrow \Delta} \qquad \wedge_R\ \frac{\Gamma \Rightarrow \varphi^u, \Delta \qquad \Gamma \Rightarrow \psi^u, \Delta}{\Gamma \Rightarrow (\varphi \wedge \psi)^u, \Delta}$$

$$\boxast_L\ \frac{\Gamma, \varphi^u, \{\Box_i \boxast \varphi^u\}_{i=1}^n \Rightarrow \Delta}{\Gamma, \boxast\varphi^u \Rightarrow \Delta} \qquad \boxast_R\ \frac{\Gamma \Rightarrow \varphi^u, \Delta \qquad \{\Gamma \Rightarrow \Box_i \boxast \varphi^a, \Delta\}_{i=1}^n}{\Gamma \Rightarrow \boxast\varphi^a, \Delta}$$

$$\Box_{\mathsf{T}}\ \frac{\Gamma, \varphi^u \Rightarrow \Delta}{\Gamma, \Box_i \varphi^u \Rightarrow \Delta} \qquad \Box_{\mathsf{S5}}\ \frac{\Box_i \Gamma \Rightarrow \varphi^a, \Box_i \Delta}{\Box_i \Gamma \Rightarrow \Box_i \varphi^a, \Box_i \Delta}$$

$$\mathsf{U}\ \frac{\Gamma \Rightarrow \Delta^u}{\Gamma \Rightarrow \Delta} \qquad \mathsf{F}\ \frac{\Gamma \Rightarrow \varphi^f, \Delta^u}{\Gamma \Rightarrow \varphi^u, \Delta^u} \qquad \mathsf{cut}\ \frac{\Gamma \Rightarrow \varphi^u, \Delta \qquad \Gamma, \varphi^u \Rightarrow \Delta}{\Gamma \Rightarrow \Delta}$$

As usual, we will call the main formula introduced in the conclusion of a rule the *principal formula* of that rule. In the case of the rule $\Box_{\mathsf{S5}}$, this is the formula $\Box_i \varphi^a$. The rules U and cut have no principal formula.

Remark 3.2 To readers familiar with cyclic proof systems for modal fixed point logics, we offer the following motivation for the focus annotations. The purpose is to capture *traces*, in the sense of [18]. Because we are working in a very simple fragment of the modal μ-calculus, our ν-traces are relatively elementary. In particular, since we are in the alternation-free fragment, our ν-traces do not pass through μ-unfoldings. Moreover, our ν-traces do not pass trough disjunctions, or, in terms of two-sided sequents, they do not pass through conjunctions on the left-hand side of the sequent. It is because of this latter property, characteristic of the *completely additive* fragment of the μ-calculus (see [9]), that the traces do not *split*, whence it suffices to have at most one formula in focus.

The fact that our sequent calculus only manipulates CKL-sequents imposes restrictions on how its rules can be applied. This is illustrated by the following lemma.

Lemma 3.3 *In any application of the* $\boxast_R$*-rule such that the principal formula is in focus, the leftmost premiss has no formula in focus.*

Proof. Suppose that in the rule application

$$\boxast_R\ \frac{\Gamma \Rightarrow \varphi^u, \Delta \qquad \{\Gamma \Rightarrow \Box_i \boxast \varphi^f, \Delta\}_{i=1}^n}{\Gamma \Rightarrow \boxast\varphi^f, \Delta}$$

the leftmost premiss has a formula in focus. Then this formula must belong to Δ. Specifically, for the conclusion is a CKL-sequent, we have $\boxplus\varphi^f \in \Delta$. However, that means that every other premiss has two formulas in focus, contradicting the fact that they must also be CKL-sequents. □

A *derivation* in $\mathsf{sCKL_f}$ is a finite tree whose nodes are labelled by CKL-sequents and which is generated by the rules of the calculus $\mathsf{sCKL_f}$. Given a derivation π in $\mathsf{sCKL_f}$, an *upward path* ρ in π is a finite sequence of nodes $\rho = \rho(0), \rho(1), ..., \rho(n)$ of π such that for each $0 \leq i < n$ the node $\rho(i+1)$ is a child of the node $\rho(i)$. Observe that we do not require upward paths to start in the root of the derivation.

Definition 3.4 An upward path ρ in an $\mathsf{sCKL_f}$-derivation is said to be *successful* if the following holds:

(i) Every sequent $\Gamma \Rightarrow \Delta$ on the path ρ has a formula in focus, *i.e.* Δ contains an annotated formula of the form φ^f.

(ii) The path ρ passes through at least one application of $\boxplus_R$, where the principal formula is in focus.

Observe that on a successful path there are no applications of the focus rules U and F. Given a derivation π in $\mathsf{sCKL_f}$ and a leaf l in π we call l an *axiomatic leaf* if l is the conclusion of an application of the rule id. If l is not axiomatic, we call it a *non-axiomatic leaf.* A *repetition* in π is a pair of nodes $\langle u, v\rangle$ such that u is a proper ancestor of v and both u and v are labelled by the same sequent. A repetition $\langle u, v\rangle$ is called *successful* if the path from u to v is successful.

Definition 3.5 A derivation π in $\mathsf{sCKL_f}$ is a *proof* if every leaf l of π is either axiomatic or there exists a node $c(l)$ such that $\langle c(l), l\rangle$ is a successful repetition.

If π is a proof with root sequent $\Gamma \Rightarrow \Delta$, then π is said to be an $\mathsf{sCKL_f}$-proof *of* $\Gamma \Rightarrow \Delta$ and $\Gamma \Rightarrow \Delta$ is called $\mathsf{sCKL_f}$-*provable.*

4 Soundness

A sequent $\Gamma \Rightarrow \Delta$ is said to be *satisfied* at a state s of an epistemic Kripke model $\mathbb{S}$ - denoted by $\mathbb{S}, s \Vdash \Gamma \Rightarrow \Delta$ - if it holds that either $\mathbb{S}, s \nVdash \varphi$ for some $\varphi^u \in \Gamma$, or $\mathbb{S}, s \Vdash \psi$ for some $\psi^a \in \Delta$. A sequent is called *valid* if it is satisfied at every state of every epistemic Kripke model. Note that the focus annotations play no meaningful role in the above definitions.

The proof of the following lemma, which states that every rule of $\mathsf{sCKL_f}$ is individually sound, is standard and therefore omitted.

Lemma 4.1 *Let*

$$\mathsf{r}\,\frac{\sigma_1 \quad \cdots \quad \sigma_n}{\sigma}$$

be a rule application of $\mathsf{sCKL_f}$. *If* σ *is invalid, then so is one of the premisses.*

In order to prove that the system $\mathsf{sCKL_f}$ is sound as a whole, we will first prove a strengthening of Lemma 4.1 which takes the annotations into account.

Let σ be a sequent that has a formula in focus, *i.e.* for $j \in \{0,1\}$ the right-hand side Δ of σ contains a formula of the form $\Box_i^j \boxplus \psi^f$. We denote by $\sigma(n)$ the sequent obtained by adding the formula $\Box_i^j \Box^n \psi^u$ to Δ (recall the definition of $\Box^n$ below Definition 2.1). For any invalid sequent σ that has a formula in focus, we define:

$$\mu(\sigma) := \min\{n \in \omega : \sigma(n) \text{ is invalid}\}.$$

Observe that the function μ is well-defined.

Lemma 4.2 *Let*

$$\mathsf{r}\,\frac{\sigma_1 \quad \cdots \quad \sigma_n}{\sigma}$$

be any rule application of $\mathsf{sCKL_f}$*. If* σ *is invalid, then there is an invalid premiss* σ_i *such that, if* σ *and* σ_i *both have a formula in focus, then:*

$$\mu(\sigma_i) \leq \mu(\sigma), \tag{1}$$

and, if moreover $\mathsf{r} = \boxplus_R$ *and the principal formula is in focus, then the inequality (1) is strict.*

Proof. Note that the statement becomes vacuous if r is id, since id derives only valid sequents. Moreover, if either σ, or all of the σ_i, have no formula in focus, then the statement reduces to Lemma 4.1. This covers the cases where r is among $\{\mathsf{U}, \mathsf{F}\}$, because those rules require either the conclusion or the sole premiss to have no formula in focus.

Now suppose that $\mathsf{r} \notin \{\mathsf{U}, \mathsf{F}, \mathsf{id}\}$ and both σ and at least one of the σ_i have a formula in focus. We first consider the case where the formula that is in focus in σ is *not* the principal formula of the rule application (this includes the case where $\mathsf{r} = \mathsf{cut}$). Direct inspection of the rules shows that in this case every premiss σ_i has a formula in focus and that

$$\mathsf{r}\,\frac{\sigma_1(\mu(\sigma)) \quad \cdots \quad \sigma_n(\mu(\sigma))}{\sigma(\mu(\sigma))}$$

is a valid rule application. The required result then follows from Lemma 4.1.

The only cases left are those in which the principal formula is in focus, which can only be the case if $\mathsf{r} \in \{\mathsf{w}_R, \Box_{\mathsf{S5}}, \boxplus_R\}$. The case $\mathsf{r} = \mathsf{w}_R$ is immediate, as the sole premiss will have no formula in focus whenever the principal formula is in focus. We treat the two remaining cases separately.

▷ $\Box_{\mathsf{S5}}$: Then σ is of the form:

$$\Box_i \Gamma \Rightarrow \Box_i \boxplus \psi^f, \Box_i \Delta.$$

Let $n := \mu(\sigma)$. By the definition of μ, there is an epistemic Kripke model $\mathbb{S}$, and a state s of $\mathbb{S}$ such that $\mathbb{S}, s \nVdash \sigma(n)$. In particular, it holds that

$$\mathbb{S}, s \nVdash \Box_i \Box^n \psi.$$

It follows that there is a state t in $\mathbb{S}$ such that sR_it and $\mathbb{S}, t \nVdash \Box^n\psi$. Clearly this also means that $\mathbb{S}, t \nVdash \boxplus\psi$. We claim that, in fact,

$$\mathbb{S}, t \nVdash \Box_i\Gamma \Rightarrow \boxplus\psi^f, \Box^n\psi^u, \Box_i\Delta,$$

which gives the required result.

By the fact that R_i is transitive, it holds for all φ such that $\mathbb{S}, s \Vdash \Box_i\varphi$, that $\mathbb{S}, t \Vdash \Box_i\varphi$. It follows that $\mathbb{S}, t \Vdash \Box_i\varphi$ for each $\Box_i\varphi^u \in \Box_i\Gamma$. Moreover, suppose that $\Box_i\psi^a \in \Box_i\Delta$. Then $\mathbb{S}, s \nVdash \Box_i\psi$. Thus there is a state r in $\mathbb{S}$ such that sR_ir and $\mathbb{S}, r \nVdash \psi$. By symmetry and transitivity, we get tR_is, whence $\mathbb{S}, t \nVdash \Box_i\psi$, as required.

▷ $\boxplus_R$: As in the previous case, let $n := \mu(\sigma)$ and let $\mathbb{S}, s$ be such that $\mathbb{S}, s \nVdash \sigma(n)$. Then $\mathbb{S}, s \nVdash \boxplus\varphi$, where $\boxplus\varphi^f$ is the principal formula.

If $n = 0$, then $\mathbb{S}, s \nVdash \varphi$ and thus the leftmost premiss is invalid and forms a witness to the statement, as it has no formula in focus. If $n > 0$, then $\mathbb{S}, s \nVdash \Box_i\Box^{n-1}\varphi$, for some $i \in \mathcal{A}$. This means that there is an invalid premiss σ_k with $\mu(\sigma_k) = n - 1$, as required. □

We are now ready to prove the soundness theorem.

Theorem 4.3 *If there is an* $\mathsf{sCKL_f}$*-proof with root* σ*, then* σ *is valid.*

Proof. Suppose, towards a contradiction, that an invalid sequent σ is the root of some $\mathsf{sCKL_f}$-proof π. Repeatedly applying Lemma 4.2, we obtain an upward path

$$\rho = \sigma_0, \sigma_1, \ldots, \sigma_n$$

through π such that $\sigma_0 = \sigma$ and σ_n labels a leaf of π. Since σ_n is invalid by construction, this leaf cannot be axiomatic. Therefore, there there must be some $k < n$ such that $\langle\sigma_k, \sigma_n\rangle$ is a successful repeat. Observe that this implies that $\sigma_k = \sigma_n$. However, by the fact that we constructed this path using Lemma 4.2, it holds that $\mu(\sigma_k) < \mu(\sigma_n)$, a contradiction. □

5 Completeness

Let Σ be a finite and closed set of formulas and let $\Gamma \Rightarrow \Delta$ be a sequent. We say that $\Gamma \Rightarrow \Delta$ is a Σ*-sequent* if $\Gamma^- \cup \Delta^- \subseteq \Sigma$. A sequent $\Gamma \Rightarrow \Delta$ will be called Σ*-provable* whenever there is an $\mathsf{sCKL_f}$-proof of $\Gamma \Rightarrow \Delta$ that contains only Σ-sequents. Finally, we say of a Σ-sequent $\Gamma \Rightarrow \Delta$ that it is *saturated* whenever it is Σ-unprovable and $\Gamma^- \cup \Delta^- = \Sigma$.

By the presence of the cut rule, the following lemma is immediate.

Lemma 5.1 *For any* Σ*-unprovable* Σ*-sequent* $\Gamma \Rightarrow \Delta$*, there is a saturated* Σ*-sequent* $\overline{\Gamma} \Rightarrow \overline{\Delta}$ *such that* $\Gamma \subseteq \overline{\Gamma}$ *and* $\Delta \subseteq \overline{\Delta}$*.*

Similarly, the following follows by the presence of the rule T.

Lemma 5.2 *If* $\Gamma \Rightarrow \Delta$ *is saturated and* $\Box_i\varphi \in \Gamma^-$*, then* $\varphi \in \Gamma^-$*.*

The following is a standard definition for the canonical model of $\mathsf{S5\text{-}CKL}$ (although the canonical model is usually defined with respect to a Hilbert-style proof system).

Definition 5.3 Let Σ be a non-empty, finite and closed set of formulas. The *canonical model* $\mathbb{S}^\Sigma$ of Σ is given by:

$$S^\Sigma := \{\Gamma^- \mid \Gamma \Rightarrow \Delta \text{ is a saturated } \Sigma\text{-sequent}\}$$
$$AR_i^\Sigma B :\Leftrightarrow \Box_i\Box_i^{-1}A = \Box_i\Box_i^{-1}B$$
$$V^\Sigma(A) := \{p \in \mathsf{P} \mid p \in A\}$$

It is immediate to verify that for every non-empty, finite and closed set Σ, its canonical model $\mathbb{S}^\Sigma$ is an epistemic Kripke model.

We are now ready to prove the Truth Lemma.

Lemma 5.4 (Truth Lemma) *For every $\varphi \in \Sigma$: $\mathbb{S}^\Sigma, A \Vdash \varphi$ if and only if $\varphi \in A$.*

Proof. We prove this by induction on φ. We only treat the cases $\varphi = \Box_i\psi$ and $\varphi = \boxplus\psi$. The other cases are standard.

▷ $\varphi = \Box_i\psi$.
In case $\varphi \in A$, we must show that for every B with $AR_i^\Sigma B$, it holds that $\mathbb{S}^\Sigma, B \Vdash \psi$. By the induction hypothesis it suffices to show that $\psi \in B$. First note that, by the definition of R_i^Σ, we have $\Box_i\psi \in B$. Lemma 5.2 then gives $\psi \in B$.

Now suppose that $\varphi \notin A$. By definition there is a saturated Σ-sequent $\Gamma \Rightarrow \Delta$ such that $\Gamma^- = A$. Note that by saturation, we have $\varphi^a \in \Delta$ for some $a \in \{u, f\}$. Let $\Delta_0 = \Delta \setminus \{\varphi^a\}$. We claim that the Σ-sequent

$$\Box_i\Box_i^{-1}\Gamma \Rightarrow \psi^a, \Box_i\Box_i^{-1}\Delta_0 \tag{2}$$

is Σ-unprovable. Indeed, consider the inference

$$\Box_{\mathsf{S5}}\ \frac{\Box_i\Box_i^{-1}\Gamma \Rightarrow \psi^a, \Box_i\Box_i^{-1}\Delta_0}{\Box_i\Box_i^{-1}\Gamma \Rightarrow \Box_i\psi^a, \Box_i\Box_i^{-1}\Delta_0}$$

If the premiss were Σ-provable, then so would be the conclusion. However, this cannot be the case because $\Gamma \Rightarrow \Delta$ can be obtained from the conclusion by a series of weakenings.

By Lemma 5.1, there is a saturated Σ-sequent $\overline{\Gamma} \Rightarrow \overline{\Delta}$ extending the sequent depicted in (2). Define the set $B := (\overline{\Gamma})^-$. Since $\psi \notin B$, the induction hypothesis gives $\mathbb{S}^\Sigma, B \nVdash \psi$. Finally, we claim that $AR_i^\Sigma B$ and thus $\mathbb{S}^\Sigma, A \nVdash \Box_i\psi$. Indeed, we clearly have $\Box_i\Box_i^{-1}A \subseteq \Box_i\Box_i^{-1}B$. For the other direction, first note that $\Box_i\psi \notin B$, for otherwise we would have $\psi \in B$. It follows for any $\Box_i\chi \in \Sigma$ that $\Box_i\chi \in B$ entails $\Box_i\chi \notin \Delta^-$, whence $\Box_i\chi \in A$, as required.

▷ $\varphi = \boxplus\psi$. As before, we will first consider the case where $\boxplus\psi \in A$. Let $\Gamma \Rightarrow \Delta$ be a saturated Σ-sequent such that $A = \Gamma^-$. Then it holds by saturation that $\psi^u \in \Gamma$ and $\Box_i\boxplus\psi^u \in \Gamma$ for all $1 \leq i \leq n$. By the induction hypothesis, we have $\mathbb{S}^\Sigma, A \Vdash \psi$. Moreover, for every $B \in S^\Sigma$ such that $AR_i^\Sigma B$, it holds by the definition of R_i^Σ and Lemma 5.2 that $\boxplus\psi \in B$. By repeating this argument we obtain that $\mathbb{S}^\Sigma, A \Vdash \boxplus\psi$.

In case $\varphi \notin A$, we again consider a saturated Σ-sequent $\Gamma \Rightarrow \Delta$ such that $A = \Gamma^-$. By the presence of the rules U and F, we may assume without loss of generality that $\varphi^f \in \Delta$. Now suppose, towards a contradiction, that $\mathbb{S}^\Sigma, A \Vdash \varphi$.

For every $A(R^\Sigma)^*B$ it holds that $\mathbb{S}^\Sigma, B \Vdash \psi$. In particular, it follows that $\mathbb{S}^\Sigma, A \Vdash \psi$, whence, by the induction hypothesis, we have $\psi^u \in \Gamma$.

As in the previous case, we take $\Delta_0 = \Delta \setminus \{\varphi^f\}$. Consider the following derivation:

$$\boxplus_R\ \dfrac{\begin{matrix}\pi \\ \Gamma \Rightarrow \psi^u, \Delta_0\end{matrix} \qquad \begin{matrix}\pi_1 \\ \Gamma \Rightarrow \Box_1\boxplus\psi^f, \Delta_0\end{matrix} \qquad \cdots \qquad \begin{matrix}\pi_n \\ \Gamma \Rightarrow \Box_n\boxplus\psi^f, \Delta_0\end{matrix}}{\Gamma \Rightarrow \boxplus\psi^f, \Delta_0}$$

where π is a series of weakenings followed by an application of id to derive $\psi^u \Rightarrow \psi^u$, and each π_i is constructed as follows:

$$\begin{array}{rc}
 & \pi' \qquad \pi'_1 \qquad \dots \qquad \pi'_n \\
\boxplus_R & \underline{\sigma' \qquad \sigma'_1 \qquad \cdots \qquad \sigma'_n} \\
\Box_{\mathsf{S5}} & \underline{\Box_i\Box_i^{-1}\Gamma \Rightarrow \boxplus\psi^f, \Box_i\Box_i^{-1}\Delta_0} \\
\mathsf{w}_L & \underline{\Box_i\Box_i^{-1}\Gamma \Rightarrow \Box_i\boxplus\psi^f, \Box_i\Box_i^{-1}\Delta_0} \\
 & \vdots \\
\mathsf{w}_L & \underline{\qquad\qquad\qquad} \\
\mathsf{w}_R & \underline{\Gamma \Rightarrow \Box_i\boxplus\psi^f, \Box_i\Box_i^{-1}\Delta_0} \\
 & \vdots \\
\mathsf{w}_R & \underline{\qquad\qquad\qquad} \\
 & \Gamma \Rightarrow \Box_i\boxplus\psi^f, \Delta_0
\end{array}$$

In the above derivation the sequent σ' is given by

$$\sigma' = \Box_i\Box_i^{-1}\Gamma \Rightarrow \psi^u, \Box_i\Box_i^{-1}\Delta_0$$

and the derivation π' is obtained from the Σ-provability of the sequent $\Box_i\Box_i^{-1}\Gamma \Rightarrow \psi^u, \Box_i\Box_i^{-1}\Delta_0$. Indeed, if it were not Σ-provable, then by applying Lemma 5.1 and the induction hypothesis, we would obtain a state B such that $\mathbb{S}^\Sigma, B \nVdash \psi$. By the same argument as in the previous case, we would moreover have $AR_i^\Sigma B$, contradicting the assumption that $A \Vdash \boxplus\psi$

Furthermore, each sequent σ'_k in the derivation π_i is given by

$$\sigma'_k = \Box_i\Box_i^{-1}\Gamma \Rightarrow \Box_k\boxplus\psi^f, \Box_i\Box_i^{-1}\Delta_0$$

and each derivation π'_k is constructed by repeatedly applying cut to add formulas from Σ until every leaf is either saturated or Σ-provable. To the leaves that are Σ-provable we append their respective proofs, and to the saturated sequents we apply the same process as we have now applied to $\Gamma \Rightarrow \Box_i \boxplus \psi^f, \Delta_0$. Crucially, for every such saturated leaf $\Gamma' \Rightarrow \Box_k \boxplus \psi^f, \Delta'$ we have that $AR_i^{\Sigma} C$, where $C := (\Gamma')^-$. Our assumption $\mathbb{S}^{\Sigma}, A \Vdash \boxplus \psi$ therefore entails that $\mathbb{S}^{\Sigma}, C \Vdash \boxplus \psi$, allowing us to repeat the same argument.

By the pidgeonhole principle, at some point one of these saturated leaves must be identical to a saturated leaf reached earlier in the construction. Note that in this case the upward path from the earlier saturated leaf to the later one is successful. We then terminate the construction of this branch. Since every branch is terminated at some point, we end up with a Σ-proof of $\Gamma \Rightarrow \Delta$, a contradiction. □

The following theorem can now be proven by a standard argument.

Theorem 5.5 *If a Σ-sequent σ is valid, then it is Σ-provable in* $\mathsf{sCKL_f}$.

6 A proof search game

Deciding whether a CKL-formula is valid is known to be EXPTIME-complete [13]. In this section we are going to show that our proof system admits optimal proof search, *i.e.* there exists an algorithm that decides whether a formula is provable in $\mathsf{sCKL_f}$ in time exponential in the size of the formula. To that end we are going to take a game-theoretic perspective on our proof system. For each sequent σ we will define a parity game $\mathcal{G}_\sigma$ which is played by two players called *Prover* and *Refuter*. Roughly, the goal of Prover is to show that σ has a proof in $\mathsf{sCKL_f}$ while Refuter tries to show the opposite. We will show that a sequent σ has a $\mathsf{sCKL_f}$-proof if and only if Prover has a memoryless winning strategy in the game $\mathcal{G}_\sigma$. In order to obtain our complexity-theoretic result, we will then refer to an algorithm given in [8] which decides for each game $\mathcal{G}_\sigma$ whether Prover has a memoryless winning strategy in time polynomial in the size of the input i.e. the size of $\mathcal{G}_\sigma$. Finally, as the size of $\mathcal{G}_\sigma$ is exponential in the size of σ, we obtain an optimal proof search result.

Throughout this section Σ always denotes a finite and closed set of formulas. For the basic definitions around parity games we refer the reader to [11]. The game $\mathcal{G}_\sigma$ has two types of positions: CKL-sequents and rule instances.

Definition 6.1 A *rule instance* (or *instance* for short) is a triple $(\sigma, \mathsf{r}, \langle \sigma_1, \ldots, \sigma_n \rangle)$ such that

$$\mathsf{r}\,\frac{\sigma_1 \ \cdots \ \sigma_n}{\sigma}$$

is a valid rule application in $\mathsf{sCKL_f}$.

If i is a rule instance, we write $\mathsf{conc}(i)$ for the first element of i, i.e. for the conclusion of i. A Σ-*instance* is a rule instance involving only Σ-sequents. For a Σ-sequent σ the game $\mathcal{G}_\sigma$ is defined as follows:

Definition 6.2 Let σ be a Σ-sequent. The *proof search game* $\mathcal{G}_\sigma$ associated to σ takes positions in $S \cup I$, where S is the set of Σ-sequents and I is the set of Σ-instances. The ownership function and admissible moves are as described in the following table:

Position	Owner	Admissible moves
σ	Prover	$\{i \in I \mid \mathsf{conc}(i) = \sigma\}$
$(\sigma, \mathsf{r}, \langle \sigma_1, \ldots, \sigma_n \rangle)$	Refuter	$\{\sigma_i \mid 1 \leq i \leq n\}$

The positions are given the following priorities:

(i) Every position of the form $\Gamma \Rightarrow \Delta^u$ has priority 3;

(ii) Every position of the form $(\sigma, \boxplus_R, \langle \sigma_0, \ldots, \sigma_n \rangle)$ where the principal formula is in focus has priority 2;

(iii) Every other position has priority 1.

A position is called a *dead end* if its owner has no admissible moves in this position available. A *match* in $\mathcal{G}_\sigma$ is a sequence of positions starting in σ, such that any two consecutive positions are related by an admissible move. A match is either finite and ends in a dead end or infinite. The *winning conditions* for a match are as follows: Prover wins every finite match in which the dead end belongs to Refuter and she wins every infinite match in which the highest priority encountered infinitely often is even. Refuter wins every finite match in which the dead end belongs to Prover and he wins every infinite match in which the highest priority visited infinitely often is odd.

Observe that the only positions that are dead ends are rule instances of id. Therefore Prover wins every finite match. In the following we are going to use standard terminology in the theory of parity games such as *strategy*, *memoryless strategy*, *winning strategy* and *strategy tree*. For definitions of these concepts we again refer the reader to [11].

Proposition 6.3 *Let σ be a Σ-sequent. If σ is Σ-provable, then Prover has a memoryless winning strategy in $\mathcal{G}_\sigma$.*

The basic proof idea is to read off a winning strategy for Prover from a Σ-proof π for σ. This can be done by identifying each play in $\mathcal{G}_\sigma$ with a finite or infinite path through π. As Prover can always choose which rule to apply to a given sequent and every play starts in σ (which labels the root of π), we can ensure that every possible play in $\mathcal{G}_\sigma$ - when Prover uses this strategy - corresponds to some path through π. The fact that π is a proof then guarantees that Prover wins every play. Finally, since in each parity game exactly one player has a memoryless winning strategy [11], the existence of a winning strategy for Prover implies the existence of a memoryless winning strategy for her. In order to present a detailed proof, we require some preliminary work.

Recall that in a proof π there exists for every non-axiomatic leaf l a node $c(l)$ such that $\langle c(l), l \rangle$ is a successful repetition. As there might exist several candidates for the node $c(l)$, we fix for each non-axiomatic leaf l one candidate $c(l)$ which we call the *companion* of l.

Definition 6.4 Let π be a $\mathsf{sCKL_f}$-proof. A *path ρ through* π is a (possibly infinite) sequence $\rho = \rho(0), \rho(1), \ldots$ of nodes in π which satisfies the following properties:

(i) $\rho(0)$ is the root of π.

(ii) If $\rho(i)$ is an axiomatic leaf, then ρ is finite and ends at $\rho(i)$.

(iii) If $\rho(i)$ is an non-axiomatic leaf l with companion $c(l)$, then $\rho(i+1)$ is a child node of $c(l)$.

(iv) Otherwise, $\rho(i+1)$ is a child node of $\rho(i)$.

If condition (iii) applies, then we say that ρ *passes through the non-axiomatic leaf* l. Observe that paths never pass through axiomatic leaves. Notice that there is a difference between paths as defined here and *upward paths* as defined is section 3. Namely, an upward path is always finite, has to end at latest at a leaf of π (axiomatic or non-axiomatic) and might start at any node. In particular, an upward path cannot pass through a non-axiomatic leaf and continue at its companion. A path as defined here has to start at the root and in case it reaches a non-axiomatic leaf, it has to continue at its companion. Moreover, a finite path can only end at an axiomatic leaf. Notice that a finite path might still pass through some non-axiomatic leaves first before eventually reaching an axiomatic leaf. Furthermore, notice that paths are defined with respect to $\mathsf{sCKL_f}$-proofs and not with respect to arbitrary derivations like upward paths. The reason for this is simply to avoid the case that some non-axiomatic leaf is not part of a successful repetition.

The following lemma states a first basic result about infinite paths through $\mathsf{sCKL_f}$-proofs. The proof of the lemma follows immediately from the fact that $\mathsf{sCKL_f}$-proofs are finite.

Lemma 6.5 *Suppose π is an $\mathsf{sCKL_f}$-proof and ρ is an infinite path through π. Then there exists a non-axiomatic leaf l through which ρ passes infinitely often.*

Given sets X, Y and a function $f : X \longrightarrow Y$ we denote by $dom(f)$ the *domain* of f and by $ran(f)$ the *range* of f.

Definition 6.6 A *finite tree with back edges* is a pair (T, f) consisting of a finite tree T and a (partial) function $f : T \longrightarrow T$, such that every $u \in dom(f)$ is a leaf of T and the node $f(u)$ is a proper ancestor of u.

Observe that cyclic proofs can be considered as finite trees with back edges, that satisfy the property that if $u \in dom(f)$, then u is a non-axiomatic leaf and $f(u)$ is the companion of u. The definition of path through a cyclic proof is generalised to path through a finite tree with back edges in the obvious way.

Definition 6.7 Let (T, f) be a finite tree with back edges. Define the *one-step dependency order* $\preceq_1$ on $ran(f)$ as follows:

$$u \preceq_1 v :\Leftrightarrow u \text{ occurs on the upward path from } v \text{ to } v' \text{ for some } v' \in f^{-1}(v)$$

Define the *dependency order* $\preceq$ on $ran(f)$ as the transitive closure of $\preceq_1$.

Observe that the dependency order $\preceq$ is reflexive and transitive. Furthermore it is also anti-symmetric and so $\preceq$ is a partial order on $ran(f)$. Observe that $u \preceq v$ implies that there exists an upward path from v to u. Let (T, f)

be a finite tree with back edges and let ρ be an infinite path through (T, f). Denote by $Inf(\rho)$ the set of nodes of T that occur infinitely often in ρ.

Lemma 6.8 *Let (T, f) be a finite tree with back edges and let ρ be an infinite path through (T, f). Then the set $Inf(\rho) \cap ran(f)$ has a $\preceq$-greatest element.*

Proof. Observe that the set $Inf(\rho) \cap ran(f)$ is finite since T is a finite tree. Furthermore, observe that $Inf(\rho) \cap ran(f)$ is non-empty, as ρ must pass through some leaf $l \in dom(f)$ infinitely often (see Lemma 6.5). It therefore suffices to prove that all $\preceq$-maximal elements in $Inf(\rho) \cap ran(f)$ are identical. To that end let $u \in Inf(\rho) \cap ran(f)$ be a $\preceq$-maximal element. We claim that all nodes in $Inf(\rho)$ belong to T_u, where T_u denotes the subtree of T which is rooted at u. Notice that it suffices to prove the claim for each $f(l)$ for which $l \in Inf(\rho) \cap dom(f) \cap T_u$. Therefore let l be an arbitrary such leaf. Then there exists an upward path ρ_l starting at the root of T and ending in l, such that both $f(l)$ and u occur on ρ_l. Moreover, since $l \in Inf(\rho)$ we have that $f(l) \in Inf(\rho) \cap ran(f)$. As u is a $\preceq$-maximal element of $Inf(\rho) \cap ran(f)$ it follows that $u \not\prec f(l)$. Therefore $f(l)$ belongs to T_u. Since u was chosen arbitrarily, it follows that $Inf(\rho) \cap ran(f)$ has a $\preceq$-greatest element. □

Corollary 6.9 *Let π be a $\mathsf{sCKL_f}$-proof and ρ be an infinite path through π. Let $l_1, \ldots, l_k$ be the non-axiomatic leaves through which ρ passes infinitely often. There exists $1 \leq i \leq k$ such that there exists an upward path from $c(l_i)$ to $c(l_j)$ for each $1 \leq j \leq k$.*

Recall that a match m in a proof search game is a finite or infinite sequence of positions $m(0), m(1), m(2), \ldots$. Observe that each even position $m(2i)$ is owned by Prover and each odd position $m(2i + 1)$ by Refuter. For an initial segment $m(0), \ldots m(i)$ of m we say that its *length* is i.

Definition 6.10 Let σ be a Σ-sequent and let π be a Σ-proof of σ. Let ρ be an infinite path through π and let m be an infinite match in $\mathcal{G}_\sigma$.

- An initial segment $m(0), \ldots, m(i)$ of m *corresponds* to an initial segment $\rho(0), \ldots, \rho(j)$ of ρ if $i = 2j$ and for each $0 \leq k \leq i$ with $k = 2l$ it holds that $m(k)$ is the sequent that labels $\rho(l)$.
- The match m and the path ρ are called *corresponding* if every initial segment of m with even length corresponds to an initial segment of ρ.

Proof. [of Proposition 6.3] Suppose that σ is Σ-provable. So there exists a proof π of σ in which every occurring sequent is a Σ-sequent. We denote the root of π by r_π. We simultaneously define a strategy $\mathcal{S}$ for Prover in the game $\mathcal{G}_\sigma$ and show how to map each initial segment of a match of even length in which Prover uses $\mathcal{S}$ onto an initial segment of a path through π. The strategy $\mathcal{S}$ is a function which maps initial segments of matches $\langle m(0), \ldots, m(2i)\rangle$ of even length onto rule instances. Therefore strategy $\mathcal{S}$ uses a *memory.*

For the base case observe that each match in $\mathcal{G}_\sigma$ begins in σ. Therefore $\langle m(0)\rangle$ for $m(0) = \sigma$ is an initial segment of every match in $\mathcal{G}_\sigma$. Similarly, every path through π starts in r_π which is labelled by σ. Thus $\langle\rho(0)\rangle$ for $\rho(0) = r_\pi$ is an initial segment of every path through π. Observe that $\langle m(0)\rangle$ corresponds to $\langle\rho(0)\rangle$.

For the inductive step suppose that we have already mapped the initial segment

$$m_i = \langle m(0), \ldots, m(2i)\rangle$$

of a match onto the initial segment

$$\rho_i = \langle\rho(0), \ldots, \rho(i)\rangle$$

of a path, such that m_i corresponds to ρ_i, where $i \geq 0$. Let $j \in I$ be the Σ-instance

$$j = (m(2i), \mathsf{r}, \langle\sigma'_1, \ldots, \sigma'_k\rangle),$$

which generated $\rho(i)$ in π when read top down. Then define

$$\mathcal{S}(m_i) = j.$$

Now suppose that Refuter extends the match by choosing premiss σ'_l. Then let $m(2i+1) = j$ and let $m(2i+2) = \sigma'_l$ and extend the initial segment m_i to

$$m_{i+1} = \langle m(0), \ldots, m(2i), m(2i+1), m(2i+2)\rangle$$

Furthermore, let $\rho(i+1)$ be the child of $\rho(i)$ which is labelled by σ'_l and extend ρ_i to

$$\rho(i+1) = \langle\rho(0), \ldots, \rho(i), \rho(i+1)\rangle$$

Observe that $m(i+1)$ corresponds to $\rho(i+1)$.

Finally, in order to turn $\mathcal{S}$ into a total function which maps *every* initial segment of a match with even length onto a rule instance (and not just those that correspond to initial segments of paths), we add the following clause. Fix a Σ-instance $j' \in I$. For any initial segment $m'_i = \langle m(0)', \ldots, m(2i)'\rangle$ of a match which is not covered in the above construction we define

$$\mathcal{S}(m'_i) = j'$$

Observe that $\mathcal{S}$ is a well-defined strategy for Prover, which has the property that if m is a match of $\mathcal{G}_\sigma$ in which Prover uses strategy $\mathcal{S}$, then there exists by construction a path ρ through π such that every initial segment of m of even length corresponds to some initial segment of ρ. Therefore m corresponds to ρ.

We show that $\mathcal{S}$ is a winning strategy for Prover. To that end let m be a match in $\mathcal{G}_\sigma$ in which Prover uses strategy $\mathcal{S}$ and let ρ be the path through π which corresponds to m. In case m is finite Prover wins by default and we have nothing to show. So suppose m is infinite. Then ρ is also infinite. By Lemma 6.5 and Corollary 6.9 there exists a non-axiomatic leaf l_0 with companion $c(l_0)$ such that the following holds:

(i) ρ passes through l_0 infinitely often.

(ii) If $l_0, l_1, ..., l_k$ are the non-axiomatic leaves through which ρ passes infinitely often, then there is an upward path from $c(l_0)$ to $c(l_i)$ for each $1 \leq i \leq k$.

Consider the subtree π_{l_0} of π rooted at $c(l_0)$. Observe that for each upward path from $c(l_0)$ to one of the non-axiomatic leaves l_i there is always a formula in focus (for a proof of this claim, see for example Proposition 2 in [20]). Next, observe that ρ contains a final segment in which each of $l_0, l_1, ..., l_k$ occur infinitely often and no other non-axiomatic leaf occurs. Therefore on this final segment ρ only passes through nodes of π_{l_0} which occur on the upward paths between $c(l_0)$ and l_i and so there is a formula in focus in each step on that final segment. Hence, the match m passes, after finitely many moves, only through positions with priority 1 or 2. As π is a proof, there exists a rule application of $\boxplus_R$ between $c(l_0)$ and l_0 where the principal formula is in focus. This means that, since ρ passes infinitely often through l_0, priority 2 is encountered infinitely often. Hence, the highest priority encountered infinitely often is even and Prover wins the match. We conclude that $\mathcal{S}$ is a winning strategy for Prover. Finally, since in a given parity game exactly one of the two players has a memoryless winning strategy [11], the existence of a winning strategy for Prover implies the existence of a *memoryless* winning strategy for her. □

Let us now consider the converse direction of Proposition 6.3.

Proposition 6.11 *Let σ be a Σ-sequent. If Prover has a memoryless winning strategy in $\mathcal{G}_\sigma$, then σ is Σ-provable.*

Proof. Suppose that Prover has a memoryless winning strategy in $\mathcal{G}_\sigma$. Let T be the corresponding strategy tree. Let π be the finite subtree of T obtained by pruning every infinite branch of T at the first repetition. Observe that the root of π is labelled by σ and π is generated by rules of $\mathsf{sCKL_f}$. Therefore π is a derivation. In order to show that π is indeed a $\mathsf{sCKL_f}$-proof, let l be an arbitrary leaf of π. If l is also a leaf of T, then l is axiomatic as T is the strategy tree of a winning strategy for Prover. Otherwise, l was generated by pruning an infinite branch of T. In that case there exists a node $c(l)$, such that $\langle c(l), l\rangle$ is a repetition. Therefore it remains to show $\langle c(l), l\rangle$ is successful. Suppose towards a contradiction that $\langle c(l), l\rangle$ is not successful. Then on the upward path ρ from $c(l)$ to l either some sequent does not have a formula in focus or ρ does not pass through an application of $\boxplus_R$ where the principal formula is in focus. This implies that there is a position occurring in ρ which has priority 3, or every position in ρ has priority 1. Since T is the strategy tree of a *memoryless* strategy, there exists an infinite branch in T that has a final segment which is an infinite concatenation $\rho \cdot \rho \cdot \rho \cdots$. On this path the highest priority occurring infinitely often is odd, contradicting the assumption that T is the strategy tree of a winning strategy for Prover. Therefore each repetition is successful and so we conclude that π is a $\mathsf{sCKL_f}$-proof of σ. □

Observe that the above constructed proof is *uniform* in the following sense:

Definition 6.12 A Σ-proof π is *uniform* if there exists a function $f : S \longrightarrow I$ such that whenever a sequent $\sigma \in S$ occurs in π, it occurs as the conclusion of the rule instance $f(\sigma)$.

Observe that in a uniform proof the first repetition in each branch is successful. We conclude:

Theorem 6.13 *The following are equivalent for any sequent σ:*

(i) *σ is $\mathsf{sCKL_f}$-provable*

(ii) *Prover has a memoryless winning strategy in $\mathcal{G}_\sigma$.*

(iii) *σ has a uniform $\mathsf{sCKL_f}$-proof.*

Let us now turn towards complexity. The *size* $c(\varphi)$ of a formula φ is the number of subformulas of φ. The *size* $c(\Gamma \Rightarrow \Delta)$ of a sequent $\Gamma \Rightarrow \Delta$ is defined as follows:

$$c(\Gamma \Rightarrow \Delta) := \sum_{\varphi \in \Gamma^- \uplus \Delta^-} c(\varphi)$$

where $\uplus$ denotes the disjoint union. Observe that if Σ is the closure of $\Gamma^- \cup \Delta^-$, then $|\Sigma|$ is linear in $c(\Gamma \Rightarrow \Delta)$.

Lemma 6.14 *Given a Σ-sequent σ, the number of positions in $\mathcal{G}_\sigma$ is polynomially bounded by $|\mathcal{P}(\Sigma)|$.*

Proof. Observe that each unannotated Σ-sequent is an ordered pair of subsets of Σ. Therefore there are $|\mathcal{P}(\Sigma)|^2$ many unannotated Σ-sequents. By taking the focus annotations into account we obtain at most $|\mathcal{P}(\Sigma)|^3$ many Σ-sequents. Hence $|S| \leq |\mathcal{P}(\Sigma)|^3$. Next, observe that for each Σ-sequent σ' and each rule r, there are at most $|\Sigma|$ many ways of applying r to σ'. We therefore obtain the following upper bound:

$$|I| \leq 14 \cdot |\mathcal{P}(\Sigma)|^3 \cdot |\Sigma| \leq 14 \cdot |\mathcal{P}(\Sigma)|^4$$

Together, the set of positions of $\mathcal{G}_\sigma$ is bounded by $14 \cdot |\mathcal{P}(\Sigma)|^4 + |\mathcal{P}(\Sigma)|^3$ which is polynomial in $|\mathcal{P}(\Sigma)|$. □

In order to get a polynomial bound for deciding the winner of a given proof search game we can now make use of one of the many existing algorithms for solving parity games. For instance the following result by Calude et al. [8].

Theorem 6.15 ([8, Theorem 2.9]) *There is an algorithm which finds the winner of a parity game in time $\mathcal{O}(n^{log(m)+6})$ for a parity game with n positions and priorities in $\{1, 2, ..., m\}$. Furthermore, the algorithm can compute a memoryless winning strategy for the winner in time $\mathcal{O}(n^{log(m)+7} \cdot log(n))$.*

Let σ be a sequent and let Σ be its closure. By Lemma 6.14 the number n of positions in $\mathcal{G}_\sigma$ is polynomial in the size of $|\mathcal{P}(\Sigma)|$. Since the number of different priorities in our games is constant, Theorem 6.15 implies that there is an algorithm deciding the winner of $\mathcal{G}_\sigma$ in time polynomial in $|\mathcal{P}(\Sigma)|$ and so exponential in $c(\sigma)$. By the same argument the above mentioned algorithm also computes a memoryless winning strategy in time exponential in $c(\sigma)$.

Corollary 6.16 *For any* CKL*-sequent* σ*, there is an algorithm deciding whether* σ *is provable that runs in time exponential in* $c(\sigma)$*.*

7 Related research and future work

In this section we discuss the relation of the present paper with earlier research. In passing we also propose some directions for further research.

7.1 Explicit induction

In [5], Alberucci & Jäger present another proof system for $\mathsf{S5}$-CKL. In contrast to our system $\mathsf{sCKL_f}$, their system $\mathsf{S5}_n(\mathsf{C})$ does not allow cyclic proofs. Rather, it uses an *explicit induction* rule, which can be thought of as the Gentzen-style translation of the well-known *induction axiom*: $\boxplus(p \to \Box p) \to (p \to \boxplus p)$.

Like us, Alberucci & Jäger obtain a partial cut-elimination result. However, given some endsequent $\Gamma \Rightarrow \Delta$, they do not manage to restrict cuts to $Cl(\Gamma \cup \Delta)$, but only to a larger set that they call the *disjunctive-conjunctive* closure of $Cl(\Gamma \cup \Delta)$. One could try to sharpen their cut-elimination result by translating our $\mathsf{sCKL_f}$-proofs into $\mathsf{S5}_n(\mathsf{C})$-proofs. Such a translation from cyclic proofs into inductive proofs occurs more often in the literature (see *e.g.* [16] and [2]).

A first problem with this approach is caused by the fact that our language is not sufficiently expressive to capture the *strengthened induction* rule from [2]. Indeed, an adaptation of the rule $\mathsf{ind_s}$ to our system would have to take the following form:

$$\mathsf{ind_s}\ \frac{\Gamma \Rightarrow \varphi, \Delta \qquad \{\Gamma \Rightarrow \Box_i \nu x.\Box(\overline{\Gamma \Rightarrow \Delta} \vee x) \wedge \varphi\}_{i=1}^{n}}{\Gamma \Rightarrow \boxplus\varphi, \Delta}$$

where $\overline{\Gamma \Rightarrow \Delta}$ is defined as

$$\bigwedge_{\varphi \in \Gamma} \varphi \wedge \bigwedge_{\psi \in \Delta} \neg\psi$$

But the μ-calculus formulas in the right premisses are not expressible in the language of CKL. This problem is circumvented by Brünnler & Lange in [7] by augmenting the language with *annotations*. However, if we were to extend our language analogously, it would be unclear how to translate sequents from this extended language into the ordinary sequents of the system $\mathsf{S5}_n(\mathsf{C})$.

One could resort to a reformulation of $\mathsf{S5}_n(\mathsf{C})$ in terms of the augmented language, but then a second problem arises. Namely, to make the translation one has to show that the strengthened induction rule is derivable in $\mathsf{S5}_n(\mathsf{C})$. But for this one needs to apply cut to formulas of the form $\overline{\Gamma \Rightarrow \Delta}$, where Γ and Δ only contain formulas in the ordinary language. Interestingly, this seems to give a partial cut-elimination result very similar to the one obtained by Alberucci & Jäger. This suggests that inductive proofs for $\mathsf{S5}$-CKL might require strictly more complex cut formulas than cyclic proofs. We leave it as future work to investigate this conjecture.

Finally, we wish to point out that the sequent given in [5, p.10] as example of a sequent that does not admit a cut-free proof in $\mathsf{S5}_n(\mathsf{C})$, namely the sequent $\Box_1(p \wedge \boxplus q) \wedge \Box_2(q \wedge \boxplus p) \Rightarrow \boxplus(p \vee q)$, does have a cut-free proof in our system.

7.2 Generalised sequent calculi

Another proposal for a proof system for S5-CKL is made by Hill & Poggiolesi in [14]. Their system HS5C does not manipulate ordinary sequents, but a generalised form of sequents called *indexed hypersequents*. Indexed hypersequents are akin to the more well-known formalisms of nested sequents (see *e.g.* [6]) or labelled sequents (see *e.g.* [17]) .

Like us, Hill & Poggiolesi start with a non-analytic system, and then use semantic methods to show that proofs can be restricted to a certain shape. While this restriction does not result in fully analytic proofs, Hill & Poggiolesi claim that it nevertheless makes their system suitable for proof search. No complexity bound for proof search is given.

Using a syntactic cut-elimination procedure, Hill & Poggiolesi show that the system HS5C is cut-free. In particular, its non-analyticity does not arise from unavoidable applications of the cut-rule. Rather, the system HS5C features a non-analytic explicit induction rule similar to that of the system $\mathsf{S5}_n(\mathsf{C})$ discussed in the previous subsection. It is this induction rule whose non-analyticity cannot be avoided.

We consider it a very interesting avenue for further research to see if the explicit induction rule in HS5C can be replaced by cyclic proofs. The main obstacle to this approach is the fact that even an analytic proof of a given endsequent may still contain infinitely many distinct labels. In terms of our approach, this would obstruct the appeal to the pidgeonhole principle in the proof of the Truth Lemma.

7.3 Focus games

In [15], Lange introduces a two-player game for checking satisfiability of Converse-PDL (CPDL, for short). This is an extension of Propositional Dynamic Logic (or, PDL) introduced by Fischer and Ladner in [10]. The logic CPDL introduces the possibility to reason about the backwards application of programs. The satisfiability game introduced in [15] is played by a player $\exists$ whose goal is to show that a given formula of CPDL is satisfiable and a player $\forall$ whose goal is to show that the formula is unsatisfiable. As in our case, the game uses a focus mechanism to capture traces. Furthermore, as our proof system, the satisfiability game is cyclic, *i.e.* plays are finite and end when either a leaf is encountered or a certain condition for cycles (similar to our condition for successful repetitions) is met. For details about the game we refer the reader to [15]. Lange claims that the satisfiability game is sound and complete, *i.e.* player $\exists$ has a winning strategy in the game for a sequent Φ if and only if Φ is satisfiable. Unfortunately, however, the authors of the present paper discovered a counterexample to the soundness of Lange's system. Namely, consider the sequent $\Phi = \{[a]p, \langle a\rangle[\bar{a}]\langle a\rangle\neg p\}$, where a is an atomic program, p an atomic proposition and $\bar{a}$ is the converse program of a. It is readily checked that Φ is unsatisfiable in any Kripke model. However, player $\exists$ has a winning strategy in the satisfiability game for Φ. This can be seen by observing that no matter in what order rules are applied to Φ, neither winning condition 1 nor 2 of $\forall$ can ever be met. For winning condition 2, this is immediate, as Φ does not contain a formula of the form $\langle \alpha^* \rangle \varphi$. For winning condition 1, it suffices to notice that no play starting in Φ can ever reach a sequent containing the formulas p and $\neg p$, as by eliminating a modal operator in the left formula, one has to introduce a modal operator in the right formula and vice versa. Therefore, no play can ever be won by $\forall$, implying that $\exists$ has a winning strategy.

Due to the close connection between CPDL and S5-CKL, we conjecture that the method of the present paper could be adapted to CPDL. If true, one could extract a focus game for validity for CPDL in the same way as described in the previous section. Such a game could then be dualised into a focus game for satisfiability, thus fixing the original system proposed in [15]. The realisation of this task is left for future research.

Acknowledgements

The authors would like to thank Johannes Marti and the anonymous reviewers for helpful comments on a final draft of this paper.

References

[1] Abate, P., R. Goré and F. Widmann, *Cut-free single-pass tableaux for the logic of common knowledge*, in: *In Workshop on Agents and Deduction at TABLEAUX*, 2007.

[2] Afshari, B. and G. E. Leigh, *Cut-free completeness for modal mu-calculus*, in: *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, IEEE, 2017, pp. 1–12.

[3] Ajspur, M., V. Goranko and D. Shkatov, *Tableau-based decision procedure for the multiagent epistemic logic with all coalitional operators for common and distributed knowledge*, Log. J. IGPL **21** (2013), pp. 407–437.

[4] Alberucci, L., "The Modal μ-Calculus and the Logic of Common Knowledge," Ph.D. thesis, Universität Bern, Institut für Informatik und angewandte Mathematik (2002).

[5] Alberucci, L. and G. Jäger, *About cut elimination for logics of common knowledge*, Annals of Pure and Applied Logic **133** (2005), pp. 73–99.

[6] Brünnler, K., "Nested sequents," habilitation, University of Bern (2010).

[7] Brünnler, K. and M. Lange, *Cut-free sequent systems for temporal logic*, The Journal of Logic and Algebraic Programming **76** (2008), pp. 216–225.

[8] Calude, C. S., S. Jain, B. Khoussainov, W. Li and F. Stephan, *Deciding parity games in quasipolynomial time*, in: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017 (2017), p. 252–263.

[9] Carreiro, F. and Y. Venema, *PDL inside the μ-calculus: a syntactic and an automata-theoretic characterization*, Advances in Modal Logic **10** (2014), pp. 74–93.

[10] Fischer, M. J. and R. E. Ladner, *Propositional dynamic logic of regular programs*, J. Comput. Syst. Sci. **18** (1979), pp. 194–211.

[11] Grädel, E., W. Thomas and T. Wilke, editors, "Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]," Lecture Notes in Computer Science **2500**, Springer, 2002.

[12] Halpern, J. Y. and Y. Moses, *Knowledge and common knowledge in a distributed environment*, J. ACM **37** (1990), p. 549–587.

[13] Halpern, J. Y. and Y. Moses, *A guide to completeness and complexity for modal logics of knowledge and belief*, Artificial intelligence **54** (1992), pp. 319–379.

[14] Hill, B. and F. Poggiolesi, *Common knowledge: a finitary calculus with a syntactic cut-elimination procedure*, Logique et Analyse (2015), pp. 279–306.

[15] Lange, M., *Satisfiability and completeness of converse-PDL replayed*, in: A. Günter, R. Kruse and B. Neumann, editors, *KI 2003: Advances in Artificial Intelligence, 26th Annual German Conference on AI, KI 2003, Hamburg, Germany, September 15-18, 2003, Proceedings*, Lecture Notes in Computer Science **2821** (2003), pp. 79–92.

[16] Lange, M. and C. Stirling, *Focus games for satisfiability and completeness of temporal logic*, in: *Proceedings 16th Annual IEEE Symposium on Logic in Computer Science*, IEEE, 2001, pp. 357–365.

[17] Negri, S., *Proof theory for modal logic*, Philosophy Compass **6** (2011), pp. 523–538.

[18] Niwiński, D. and I. Walukiewicz, *Games for the μ-calculus*, Theoretical Computer Science **163** (1996), pp. 99–116.

[19] Ohnishi, M. and K. Matsumoto, *Gentzen method in modal calculi*, Osaka Mathematical Journal **9** (1957), pp. 113–130.

[20] Rooduijn, J. M. W., *Cyclic hypersequent calculi for some modal logics with the master modality*, in: A. Das and S. Negri, editors, *Automated Reasoning with Analytic Tableaux and Related Methods* (2021), pp. 354–370.

[21] Stirling, C., *A tableau proof system with names for modal mu-calculus*, in: A. Voronkov and M. V. Korovina, editors, *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, EPiC Series in Computing **42**, EasyChair, 2014 pp. 306–318.

[22] van Ditmarsch, H., W. van der Hoek and B. Kooi, "Dynamic Epistemic Logic," Springer Publishing Company, Incorporated, 2007, 1st edition.

[23] Wehbe, R., "Annotated Systems for Common Knowledge," Ph.D. thesis, Universität Bern, Institut für Informatik und angewandte Mathematik (2010).