Cut–Free Axiomatizations for Stratified Modal Fixed Point Logic

Gerhard Jäger, Mathis Kretz*, and Thomas Studer

Institut für Informatik und angewandte Mathematik, Universität Bern,
Neubrückstrasse 10, CH–3012 Bern, Switzerland
{jaeger, kretz, tstuder}@iam.unibe.ch

Abstract. We present an infinitary and a finitary cut–free axiomatization for a fragment of the modal μ–calculus in which nesting of fixed points is restricted to non–interleaving occurrences. In this study we prove soundness and completeness of both axiomatizations. Completeness is established by constructing a canonical countermodel to any non–provable formula using an extension of the method of saturated sequents. Soundness of the finitary axiomatization is a consequence of the small model property and well–known results about monotone operators while completeness follows from the corresponding result for the infinitary case.

1 Introduction

Stratified modal fixed point logic, henceforth abbreviated as SFL, is an extension of standard multi–modal logic with syntactic constructs for representing least and greatest fixed points of stratified positive formulae. As such SFL is contained in Kozen's well known modal μ–calculus which – along with many of its fragments – is an important tool for the specification and verification of properties of programs. It has been studied extensively, for example, with respect to model–checking. Classical studies in this direction include Kozen [11], Streett and Emerson [14], Winskel [17] and Stirling and Walker [13].

Interesting efforts have also been undertaken to obtain sound and complete deductive systems for the modal μ–calculus. Completeness of the (trivially sound) axiomatization proposed by Kozen [11] was addressed by Walukiewicz [15], making use of some deep results from the theory of automata on infinite words. While certainly being interesting from a foundational point of view the proposed deductive systems have so far left one important question largely unaddressed, namely the possibility of obtaining cut–free axiomatizations. All previously proposed complete deductive systems for the modal μ–calculus crucially include a cut–rule and semantic or syntactic cut–elimination procedures were not known.

In this paper we present an infinitary and a closely related finitary cut–free axiomatization of SFL and show that these axiomatizations are sound and complete with respect to the standard Kripke semantics. The language of SFL captures –

* Research partly supported by the Swiss National Science Foundation

as we shall see later – many important logics such as PDL, CTL and the logic of common knowledge. It is stratified in the following sense: Consider a formula $\mu \mathsf{X}.\mathcal{A}[\mathsf{X}]$, where $\mathcal{A}[\mathsf{X}]$ is positive in the variable X. $\mathcal{A}[\mathsf{X}]$ may contain a subformula $\mu \mathsf{Y}.\mathcal{B}[\mathsf{Y}]$ and $\nu \mathsf{Y}.\mathcal{B}[\mathsf{Y}]$ only if X does not appear free in $\nu \mathsf{Y}.\mathcal{B}[\mathsf{Y}]$. This allows us to compute the meaning of $\mu \mathsf{Y}.\mathcal{B}[\mathsf{Y}]$ and $\nu \mathsf{Y}.\mathcal{B}[\mathsf{Y}]$ and then use it to determine the meaning of $\mu \mathsf{X}.\mathcal{A}[\mathsf{X}]$. Stratification guarantees that inner fixed points do not depend on the outer ones. Hence it is possible to determine the meaning of any formula by a simple induction on the levels of its fixed points and its complexity.

This is not possible when interleaving of fixed points is allowed. Consider the formula of the form $\mu \mathsf{X}.\nu \mathsf{Y}.\mathcal{A}[\mathsf{X}, \mathsf{Y}]$. Here the meaning of the inner fixed point $\nu \mathsf{Y}.\mathcal{A}[\mathsf{X}, \mathsf{Y}]$ depends on the value assigned to X by the interpretation of the outer fixed point $\mu \mathsf{X}.\nu \mathsf{Y}.\mathcal{A}[\mathsf{X}, \mathsf{Y}]$ which in turn depends on $\nu \mathsf{Y}.\mathcal{A}[\mathsf{X}, \mathsf{Y}]$. Hence interleaving has the effect that the meaning of nested fixed points cannot be determined one after another, but has to be treated in a more complicated way.

Turning to our two deductive systems for SFL we show (i) completeness of the infinitary version $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ and (ii) the soundness of the finitary system $\mathsf{T}_{\mathsf{SFL}}$. Furthermore, it will be obvious that everything provable in $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ is also provable in $\mathsf{T}_{\mathsf{SFL}}$ and, consequently, $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ as well as $\mathsf{T}_{\mathsf{SFL}}$ are sound and complete.

The completeness of $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ is proved by extending a method also used in Alberucci and Jäger [1]. For establishing the soundness of the finitary system $\mathsf{T}_{\mathsf{SFL}}$, we heavily rely on the small model property of the modal μ–calculus; see [14, 5, 9, 16]. Generalizing an idea of Jäger, Kretz and Studer [10], we can exploit the fact that with respect to provability only finitely many stages of greatest fixed points are relevant.

The focus of this paper is on proof-theoretic aspects of an important subsystem of the modal μ–calculus and the question whether there are cut–free formalizations. We do *not* claim that the systems presented in this article are a direct basis for implementations which perform better, for example, with respect to average case behavior. However, we think that it is important to know that finite cut–free calculi do exist so that it makes sense to search for cut–free frameworks whose structural proof theory is more developed than that of $\mathsf{T}_{\mathsf{SFL}}$. Furthermore, the proof systems studied in this paper allow direct and logic–oriented completeness proofs, a fact which is interesting in its own right.

2 Monotone Operators

To facilitate the introduction of SFL we will first briefly review some basic notions and results from the theory of monotone operators; for all details we refer, for example, to Moschovakis [12].

Definition 1 (Monotone Operators). *Let A be a set. A function F is called an operator on A if $F : \mathcal{P}(A) \longrightarrow \mathcal{P}(A)$. F is called monotone if for all $B, C \subset A$ such that $B \subset C$ we have $F(B) \subset F(C)$.*

An important notion which we will require in various forms is that of a (potentially transfinite) iteration of a given operator. There are two natural ways of defining iterations. Firstly, we may start from the empty set and collect up all elements which are possibly added by repeated applications of a given operator. Secondly, we may start with the whole domain and throw away any elements which are possibly removed by a repeated application of the operator. Both these possibilities are reflected in the next definition.

Definition 2 (Iterations). *Let A be a set, F a monotone operator on A and α an ordinal. Define the sets I_F^α, I_F, $I_F^{<\alpha}$ J_F^α, J_F and $J_F^{<\alpha}$ as follows:*

$$I_F^{<\alpha} = \bigcup_{\beta<\alpha} I_F^\beta, \qquad I_F^\alpha = F(I_F^{<\alpha}), \qquad I_F = \bigcup_\beta I_F^\beta.$$
$$J_F^{<\alpha} = \bigcap_{\beta<\alpha} J_F^\beta, \qquad J_F^\alpha = F(J_F^{<\alpha}), \qquad J_F = \bigcap_\beta J_F^\beta.$$

The next two theorems state that the sequences of sets determined by iterating a monotone operator transfinitely many times in the two ways described in Definition 2 always converge to the least and greatest fixed point of this operator and that the number of iterations at which convergence is reached is bounded by the cardinality of the domain.

Theorem 1 (Least Fixed Point). *Let A be a set and F a monotone operator on A. Then the following statements hold:*

(i) *For any ordinals α, β, if $\beta \leq \alpha$, then $I_F^\beta \subset I_F^\alpha$.*
(ii) *There exists an ordinal α such that $\alpha < |A|^+$ and $I_F = I_F^\alpha = I_F^{<\alpha}$.*
(iii) $F(I_F) = I_F = \bigcap\{B : F(B) \subset B\} = \bigcap\{B : F(B) = B\}$.

A proof of this theorem is given in Moschovakis [12]. This takes care of the least fixed point case. The analogous situation also holds for greatest fixed points and their approximations.

Theorem 2 (Greatest Fixed Point). *Let A be a set and F a monotone operator on A. Then the following statements hold:*

(i) *For any ordinals α, β, if $\beta \leq \alpha$, then $J_F^\alpha \subset J_F^\beta$.*
(ii) *There exists an ordinal α such that $\alpha < |A|^+$ and $J_F = J_F^\alpha = J_F^{<\alpha}$.*
(iii) $F(J_F) = J_F = \bigcup\{B : B \subset F(B)\} = \bigcup\{B : B = F(B)\}$.

Finally, for reasons of duality we shall also mention the following statement.

Theorem 3 (Fixed Point Duality). *Let F be a monotone operator on a set A. Then the operator $G(X) := A \setminus F(A \setminus X)$ is monotone and the following properties hold for all $B \subset A$:*

(i) *If $F(B) = B$, then $G(A \setminus B) = A \setminus B$.*
(ii) *If $G(B) = B$, then $F(A \setminus B) = A \setminus B$.*
(iii) *$I_G = A \setminus J_F$ and $J_G = A \setminus I_F$.*

3 Stratified Modal Fixed Point Logic

Starting off from the usual language of multi-modal logic, we introduce the language of SFL in a level-by-level fashion, beginning at level 0, at which we allow only formulae of modal logic possibly containing a fresh propositional variable X and its negation $\sim\mathsf{X}$. For all formulae $\mathcal{A}$ which are positive in X we then add constants $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$ for the least and greatest fixed point of the associated operator. At level 1 we allow formulae of modal logic possibly containing X and $\sim\mathsf{X}$ as well as any constant $P_{\mathcal{A}}$ or $Q_{\mathcal{A}}$ from level 0. Again, for those formulae $\mathcal{B}$ of level 1, which are positive in X we add new constants $P_{\mathcal{B}}$ and $Q_{\mathcal{B}}$ as before and iterate this procedure inductively, thus obtaining levels $2, 3, 4$ and so on. Each formula A of the language is therefore assigned a level in a natural way, namely the minimal level of this construction at which A appears. At each level of the language we also define the negation $\neg A$ for an arbitrary formula A using De Morgan's laws, the law of double negation and the law of fixed point duality stated in Theorem 3.

Definition 3 (Language, Level, Length). *Let $\Phi = \{\mathsf{p}, \sim\mathsf{p}, \mathsf{q}, \sim\mathsf{q}, \mathsf{r}, \sim\mathsf{r}, \ldots\}$ be a countable set of atomic propositions, $\mathsf{V} = \{\mathsf{X}, \sim\mathsf{X}\}$ a set containing one variable and its negation, $\mathsf{T} = \{\top, \bot\}$ a set containing symbols for truth and falsehood and M a set of indices.*

1. *Define $\mathcal{L}_0$ as being the least superset of $\Phi \cup \mathsf{V} \cup \mathsf{T}$ closed under $\wedge, \vee, \Box_i, \Diamond_i$ where $i \in \mathsf{M}$. Given a formula $A \in \mathcal{L}_0$, we inductively define $\neg A$ by:*

$$\neg\mathsf{p} := \sim\mathsf{p}, \quad \neg\mathsf{X} := \sim\mathsf{X}, \quad \neg\sim\mathsf{p} := \mathsf{p}, \quad \neg\sim\mathsf{X} := \mathsf{X}, \quad \neg\top := \bot, \quad \neg\bot := \top,$$
$$\neg(B \wedge C) := \neg B \vee \neg C, \quad \neg(B \vee C) := \neg B \wedge \neg C,$$
$$\neg\Box_i B := \Diamond_i \neg B, \quad \neg\Diamond_i B := \Box_i \neg B.$$

2. *A formula $A \in \mathcal{L}_0$ is called X-positive if $\sim\mathsf{X}$ does not occur in A. In the following, X-positive formulae will be denoted by symbols $\mathcal{A}, \mathcal{B}, \mathcal{C}, \ldots$ (possibly with primes and subscripts). Given formulae $\mathcal{A}$ and B of $\mathcal{L}_0$ where $\mathcal{A}$ is X-positive we write $\mathcal{A}[B]$ for the formula which is obtained by replacing every occurrence of X in $\mathcal{A}$ by B. We define the dual $\overline{\mathcal{A}}$ of an X-positive formula $\mathcal{A}$ as $\overline{\mathcal{A}} = \neg(\mathcal{A}[\sim\mathsf{X}])$. Furthermore, define the sets*

$$L_0 := \{P_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_0 \ \mathsf{X}\text{-positive}\} \text{ and } G_0 := \{Q_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_0 \ \mathsf{X}\text{-positive}\}.$$

3. *Define $\mathcal{L}_{k+1}$ as the least superset of $\Phi \cup \mathsf{V} \cup \mathsf{T} \cup L_k \cup G_k$ closed under the symbols $\wedge, \vee, \Box_i, \Diamond_i$ where $i \in \mathsf{M}$. Again a formula $\mathcal{A} \in \mathcal{L}_{k+1}$ is called X-positive if $\sim\mathsf{X}$ does not occur in $\mathcal{A}$. For any formulae $A \in \mathcal{L}_{k+1}$ define $\neg A$ as before and adding the clauses $\neg P_{\mathcal{A}} := Q_{\overline{\mathcal{A}}}$ and $\neg Q_{\mathcal{A}} := P_{\overline{\mathcal{A}}}$. For X-positive formulae $\mathcal{A} \in \mathcal{L}_{k+1}$ we also define $\overline{\mathcal{A}}$ as before. Similar to the base case we also define sets of fixed point constants*

$$L_{k+1} := \{P_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_k \ \mathsf{X}\text{-positive}\} \text{ and } G_{k+1} := \{Q_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_k \ \mathsf{X}\text{-positive}\}.$$

4. *Define the language $\mathcal{L}$ by setting $\mathcal{L} = \bigcup_{k \in \omega} \mathcal{L}_k$.*

5. *For every formula $A \in \mathcal{L}$ define* $\mathrm{level}(A)$ *to be the least k so that $A \in \mathcal{L}_k$.*
6. *The length $|A|$ of an $A \in \mathcal{L}$ is simply the number of symbols occurring in A with the proviso that fixed point constants $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$ count (recursively) as $|\mathcal{A}| + 1$ many symbols.*

Turning to the semantics of $\mathcal{L}$, a Kripke structure, as usual, is a triple $\mathsf{K} = (S, R, \pi)$ where S is a non–empty set, R a mapping from M to $\mathcal{P}(S \times S)$ and π a mapping from $\Phi \cup \mathsf{V}$ to $\mathcal{P}(S)$; we also require that $\pi({\sim}\mathsf{X}) = S \setminus \pi(\mathsf{X})$ and $\pi({\sim}\mathsf{p}) = S \setminus \pi(\mathsf{p})$ for all $\mathsf{p} \in \Phi$. The function R assigns an accessibility relation to each $i \in \mathsf{M}$ where we write R_i for the relation $R(i)$. Furthermore, given a set $T \subset S$ we define $\mathsf{K}[\mathsf{X} := T]$ to be the the Kripke structure (S, R, π'), where π' maps X to T, ${\sim}\mathsf{X}$ to $S \setminus T$ and otherwise agrees with π.

Given a Kripke structure K we now define the denotations of the formulae of $\mathcal{L}$ by main induction on their levels and side induction on their lengths.

Definition 4 (Denotation). *Let $\mathsf{K} = (S, R, \pi)$ be a Kripke structure. For every $A \in \mathcal{L}_k$ we define the set $\|A\|_{\mathsf{K}} \subset S$ inductively as follows:*

$$\|P\|_{\mathsf{K}} := \pi(P) \text{ for all } P \in \Phi \cup \mathsf{V}, \quad \|\top\|_{\mathsf{K}} := S, \quad \|\bot\|_{\mathsf{K}} := \emptyset,$$
$$\|B \wedge C\|_{\mathsf{K}} := \|B\|_{\mathsf{K}} \cap \|C\|_{\mathsf{K}}, \quad \|B \vee C\|_{\mathsf{K}} := \|B\|_{\mathsf{K}} \cup \|C\|_{\mathsf{K}},$$
$$\|\Box_i B\|_{\mathsf{K}} := \{w \in S : v \in \|B\|_{\mathsf{K}} \text{ for all } v \text{ such that } wR_iv\},$$
$$\|\Diamond_i B\|_{\mathsf{K}} := \{w \in S : v \in \|B\|_{\mathsf{K}} \text{ for some } v \text{ such that } wR_iv\}.$$

For every $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$ we define

$$\|P_{\mathcal{A}}\|_{\mathsf{K}} := \bigcap\{T \subset S : T \supset F^{\mathsf{K}}_{\mathcal{A}}(T)\} \text{ and } \|Q_{\mathcal{A}}\|_{\mathsf{K}} := \bigcup\{T \subset S : T \subset F^{\mathsf{K}}_{\mathcal{A}}(T)\}$$

where $F^{\mathsf{K}}_{\mathcal{A}}$ is the operator on S given by $F^{\mathsf{K}}_{\mathcal{A}}(T) := \|\mathcal{A}\|_{\mathsf{K}[\mathsf{X}:=T]}$ for every subset T of S.

Obviously by the main induction hypothesis the operator $F^{\mathsf{K}}_{\mathcal{A}}$ is defined and monotone and $\|P_{\mathcal{A}}\|_{\mathsf{K}}$ and $\|Q_{\mathcal{A}}\|_{\mathsf{K}}$ are the least and greatest fixed points of $F^{\mathsf{K}}_{\mathcal{A}}$ respectively.

Definition 5 (Satisfaction and Validity). *Let $\mathsf{K} = (S, R, \pi)$ be a Kripke structure. We say a formula $A \in \mathcal{L}$ is satisfied in K if $\|A\|_{\mathsf{K}} \neq \emptyset$ and valid in K if $\|A\|_{\mathsf{K}} = S$. We say A is satisfiable if there exists a Kripke structures in which A is satisfied. Furthermore, we say that A is valid if it is valid in all Kripke structures.*

The next definition introduces formulae describing the finite iterations of monotone operators, both from below and above. We also define the stages of such inductive definitions over a given Kripke structure.

Definition 6 (Iterations). *Let $\mathcal{A}$ be an X-positive formula.*

1. *For every $k \in \omega$ define the formulae $Q^k_{\mathcal{A}}$ and $P^k_{\mathcal{A}}$ inductively as follows:*

$$P^0_{\mathcal{A}} := \bot, \quad P^{k+1}_{\mathcal{A}} := \mathcal{A}[P^k_{\mathcal{A}}], \qquad Q^0_{\mathcal{A}} := \top, \quad Q^{k+1}_{\mathcal{A}} := \mathcal{A}[Q^k_{\mathcal{A}}].$$

2. *Let $\mathsf{K} = (S, R, \pi)$ be a Kripke structure. For every ordinal α define the subsets $I^{<\alpha}_{\mathcal{A},\mathsf{K}}$, $I^{\alpha}_{\mathcal{A},\mathsf{K}}$ $J^{<\alpha}_{\mathcal{A},\mathsf{K}}$ and $J^{\alpha}_{\mathcal{A},\mathsf{K}}$ of S as follows:*

$$I^{<\alpha}_{\mathcal{A},\mathsf{K}} := I^{<\alpha}_{F^{\mathsf{K}}_{\mathcal{A}}}, \qquad I^{\alpha}_{\mathcal{A},\mathsf{K}} := I^{\alpha}_{F^{\mathsf{K}}_{\mathcal{A}}}, \qquad J^{<\alpha}_{\mathcal{A},\mathsf{K}} := J^{<\alpha}_{F^{\mathsf{K}}_{\mathcal{A}}}, \qquad J^{\alpha}_{\mathcal{A},\mathsf{K}} := J^{\alpha}_{F^{\mathsf{K}}_{\mathcal{A}}}.$$

Remark 1. We note two facts, which are evident from Definition 6 and which will be used several times in the subsequent argument. Let $\mathcal{A}$ be a formula of $\mathcal{L}$ which is X-positive. Then for all natural numbers k we have $\mathrm{level}(Q^k_{\mathcal{A}}) < \mathrm{level}(Q_{\mathcal{A}})$ and $\mathrm{level}(P^k_{\mathcal{A}}) < \mathrm{level}(P_{\mathcal{A}})$ and, furthermore, for all Kripke structures K we have $\|P^k_{\mathcal{A}}\|_{\mathsf{K}} = I^{<k}_{\mathcal{A},\mathsf{K}}$ and $\|Q^k_{\mathcal{A}}\|_{\mathsf{K}} = J^{<k}_{\mathcal{A},\mathsf{K}}$.

Our language $\mathcal{L}$ is built up in layers and at each layer we add a set of fixed point constants for certain formulae of the layer below. The following definition is so that the rank of a formula at a higher level is strictly greater than the rank of any formula in a level below.

Definition 7 (Rank). *The rank $\mathrm{rk}(A)$ of a formula $A \in \mathcal{L}$ is an ordinal defined inductively as follows:*

1. *If A is an element of $\Phi \cup \mathsf{V} \cup \mathsf{T}$, then $\mathrm{rk}(A) := 0$.*
2. *If A is a fixed point constant $P_{\mathcal{A}}$ or $Q_{\mathcal{A}}$ and $\mathrm{level}(A) = n$, then $\mathrm{rk}(A) := \omega n$.*
3. *If A is a formula $B \wedge C$ or $B \vee C$, then $\mathrm{rk}(A) := \max(\mathrm{rk}(B), \mathrm{rk}(C)) + 1$.*
4. *If A is a formula $\Box_i B$ or $\Diamond_i B$ for some $i \in \mathsf{M}$, then $\mathrm{rk}(A) := \mathrm{rk}(B) + 1$.*

The following lemma summarizes some important properties of the rank function. Its proof is routine and left to the reader.

Lemma 1. *For all formulae $A, B \in \mathcal{L}$ the following statements hold:*

(i) $\omega \cdot \mathrm{level}(A) \leq \mathrm{rk}(A) < (\omega \cdot \mathrm{level}(A)) + \omega$.
(ii) $\mathrm{rk}(A) = \mathrm{rk}(\neg A)$.
(iii) $\mathrm{rk}(A), \mathrm{rk}(B) < \mathrm{rk}(A \wedge B), \mathrm{rk}(A \vee B)$.
(iv) $\mathrm{rk}(A) < \mathrm{rk}(\Box_i A), \mathrm{rk}(\Diamond_i A)$ *for all $i \in \mathsf{M}$.*

The right hand inequality of the following Lemma is a straightforward consequence of Remark 1 and Lemma 1, the left hand one is an induction on k.

Lemma 2. *For all fixed point constants $Q_{\mathcal{A}}$ and all $k \in \omega$ we have*

$$\mathrm{rk}(Q^k_{\mathcal{A}}) \leq \mathrm{rk}(Q^{k+1}_{\mathcal{A}}) < \mathrm{rk}(Q_{\mathcal{A}}).$$

The preceding considerations show that the rank of a formula A as given in Definition 7 allow us to prove certain claims – for example the next one – by a simple induction on $\mathrm{rk}(A)$.

Lemma 3. *For all formulae A of $\mathcal{L}$ and all Kripke structures $\mathsf{K} = (S, R, \pi)$ we have $\|\neg A\|_{\mathsf{K}} = S \setminus \|A\|_{\mathsf{K}}$.*

4 A Hilbert System for SFL

In this section we briefly introduce the Hilbert system $\mathsf{H}_{\mathsf{SFL}}$ for stratified modal fixed point logic. We mention $\mathsf{H}_{\mathsf{SFL}}$ because it is simple and easily accessible from an intuitive point of view. The system is basically Kozen's axiomatization of the modal μ–calculus [11] adapted to our more restrictive setting. It consists of the standard axioms and rules for multi–modal logic plus axioms and rules for the fixed point constants. Given X-positive formulae $\mathcal{A}$ the additional axioms express the fact that the constants $P_{\mathcal{A}}$ stand for fixed points and the additional induction rules state that these fixed points are minimal. Indeed, as we are about to see we only require additional axioms and rules for least fixed points. Their counterparts with respect to greatest fixed points may be derived due to the syntactic duality of least and greatest fixed points in our language.

Definition 8 (The System $\mathsf{H}_{\mathsf{SFL}}$). *The system $\mathsf{H}_{\mathsf{SFL}}$ is defined by adding the following axioms and rules to any suitable Hilbert–style axiomatization for a multi–modal formulation of the logic K:*

Closure axioms *For every X-positive formula $\mathcal{A}$:*

$$\mathcal{A}[P_{\mathcal{A}}] \to P_{\mathcal{A}}$$

Induction rules *For every X-positive formula $\mathcal{A}$ and every formula B:*

$$\frac{\mathcal{A}[B] \to B}{P_{\mathcal{A}} \to B}$$

It is easily checked that for greatest fixed points the duals of the closure axioms and induction rules are derivable in $\mathsf{H}_{\mathsf{SFL}}$. Basically all we need is the fixed point duality formulated in Theorem 3.

Lemma 4. *The system $\mathsf{H}_{\mathsf{SFL}}$ derives the formula $Q_{\mathcal{A}} \to \mathcal{A}[Q_{\mathcal{A}}]$ and the rule:*

$$\frac{B \to \mathcal{A}[B]}{B \to Q_{\mathcal{A}}}$$

We may use the system $\mathsf{H}_{\mathsf{SFL}}$ to gain some intuition about the expressive strength of SFL. To this end we will briefly discuss three notable fragments of SFL, namely logic of common knowledge, computational tree logic CTL and propositional dynamic logic PDL. Logic of common knowledge [6] is a multi–modal logic which can be used to talk about certain epistemic situations among a group of agents. The index set M is interpreted as standing for a finite set of agents and a modal formula $\Box_i A$ is taken to mean that agent i knows the statement A. If A is known by all agents in M, we write $\mathsf{E}A$ which, formally speaking, is just an abbreviation for the conjunction $\bigwedge_{i \in \mathsf{M}} \Box_i A$. If A is common knowledge among all agents – all agents know A and all agents know that all agents know A and so on ad infinitum – then we write $\mathsf{C}A$. More formally, the logic of common knowledge can be axiomatized as a multi–modal version of K plus the following axioms and rules to treat the common knowledge operator C:

Closure axioms For every formula A:

$$\mathsf{C}A \to \mathsf{E}(A \wedge \mathsf{C}A)$$

Induction rules For every formula A and B:

$$\frac{B \to \mathsf{E}(A \wedge B)}{B \to \mathsf{C}A}$$

It can easily be seen that the logic of common knowledge is a fragment of SFL by defining a syntactic embedding from the former into the latter: we translate atomic, propositional and modal formulae as themselves and a formula $\mathsf{C}A$ as $Q_{\mathsf{E}(A^* \wedge \mathsf{X})}$ where A^* denotes the translation of A. By Lemma 4 it is clear that we then obtain the following theorem stating that the translations of the axioms and rules for common knowledge are derivable in $\mathsf{H}_{\mathsf{SFL}}$. Since the rest of the axioms and rules of logic of common knowledge are also a part of $\mathsf{H}_{\mathsf{SFL}}$, this already takes care of the embedding result.

Theorem 4. *For all formulae A and B of logic of common knowledge we have:*

1. $\mathsf{H}_{\mathsf{SFL}}$ *derives* $(\mathsf{C}A \to \mathsf{E}(A \wedge \mathsf{C}A))^*$.
2. *If* $\mathsf{H}_{\mathsf{SFL}}$ *derives* $(B \to \mathsf{E}(A \wedge B))^*$, *then* $\mathsf{H}_{\mathsf{SFL}}$ *also derives* $(B \to \mathsf{C}A)^*$.

The next fragment of SFL we consider is computational tree logic or CTL for short [4]. CTL is based on mono–modal logic and may be used to talk about the set of all possible runs of a system. Using this logic we may express such properties as "in all runs extending from the current state A holds in the next state", written as $\Box A$, or "in some runs extending from the current state A holds in the next state", written as $\Diamond A$. More importantly we may also express behavior which is in a sense unbounded like "in all runs extending from the current state A holds until B is the case", denoted by $\forall(A \mathsf{U} B)$ or "in some runs extending from the current state A holds until B is the case", written as $\exists(A \mathsf{U} B)$. Axiomatically we obtain CTL by extending a suitable axiomatization for K by the following axioms and rules governing the use of the $\forall(A \mathsf{U} B)$ and $\exists(A \mathsf{U} B)$ constructs:

Closure axioms For every formula A and B

(1) $B \vee (A \wedge \Box\forall(A \mathsf{U} B)) \to \forall(A \mathsf{U} B)$

(2) $B \vee (A \wedge \Diamond\exists(A \mathsf{U} B)) \to \exists(A \mathsf{U} B)$

Induction rules For all formulae A, B and C

(3) $\dfrac{C \to (\neg B \wedge \Diamond C)}{C \to \neg\forall(A \mathsf{U} B)}$

(4) $\dfrac{C \to (\neg B \wedge (A \to \Box C))}{C \to \neg\exists(A \mathsf{U} B)}$

An embedding of CTL into SFL can be obtained again by translating atomic, propositional and modal constructs as themselves and using in this case least fixed point constants to translate formulae of the form $\forall(A \mathsf{U} B)$ and $\exists(A \mathsf{U} B)$. More precisely we translate a formula of the form $\forall(A \mathsf{U} B)$ into $P_{B^* \vee (A^* \wedge \Box \mathsf{X} \wedge \Diamond \top)}$ and a formula of the form $\exists(A \mathsf{U} B)$ into $P_{B^* \vee (A^* \wedge \Diamond \mathsf{X})}$ where in both cases A^* and B^* stand for the translations of A and B respectively. Using this translation and Definition 8 we again immediately get the following theorem which ensures that CTL is embeddable into $\mathsf{H}_{\mathsf{SFL}}$.

Theorem 5. *For all formulae A, B and C of* CTL *we have:*

1. $\mathsf{H_{SFL}}$ *derives* $[B \vee (A \wedge \Box\forall(A\mathsf{U}B)) \to \forall(A\mathsf{U}B)]^*$.
2. $\mathsf{H_{SFL}}$ *derives* $[B \vee (A \wedge \Diamond\exists(A\mathsf{U}B)) \to \exists(A\mathsf{U}B)]^*$.
3. *If* $\mathsf{H_{SFL}}$ *derives* $[C \to (\neg B \wedge \Diamond C)]^*$, *then also* $[C \to \neg\forall(A\mathsf{U}B)]^*$.
4. *If* $\mathsf{H_{SFL}}$ *derives* $[C \to (\neg B \wedge (A \to \Box C))]^*$, *then also* $[C \to \neg\exists(A\mathsf{U}B)]^*$.

The last fragment of SFL which we mention is propositional dynamic logic, abbreviated as PDL [7, 8]. It is once again a multi–modal logic this time featuring an infinite set M of indices. The logic is primarily used for reasoning about programs in the following sense: a formula $\Box_i A$ is interpreted as the statement "whenever an i action is executed in the current state, we terminate in a state which satisfies A". Consequently, $\Diamond_i A$ is taken to mean "in the current state it is possible to execute an i action and terminate in a state which satisfies A". Similar to our two previous examples PDL also features constructs for expressing "unbounded" properties. $\Box_i^* A$ is used to state that for any finite iteration of the action i we end up in a state satisfying A. Dually, $\Diamond_i^* A$ states that there exists a finite iteration of i actions after which we end up in a state satisfying A. Again we may axiomatize PDL by taking suitable axioms for multi–modal K and extending them by the following axioms and rules for the $\Box_i^*$ and $\Diamond_i^*$ operators:

Closure axioms For every formula A

(1) $\Box_i^* A \to (A \wedge \Box_i \Box_i^* A)$
(2) $(A \vee \Diamond_i \Diamond_i^* A) \to \Diamond_i^* A$

Induction rules For all formulae A, B

(3) $$\frac{B \to (A \wedge \Box_i B)}{B \to \Box_i^* A}$$
(4) $$\frac{(A \vee \Diamond_i B) \to B}{\Diamond_i^* A \to B}$$

The above axioms and rules suggest a translation of a formula $\Box_i^* A$ as a greatest and $\Diamond_i^* A$ as a least fixed point. More precisely, in order to embed PDL into SFL we translate atomic, propositional and modal formulae as themselves, $\Box_i^* A$ as $Q_{A^* \wedge \Box_i \mathsf{x}}$ and $\Diamond_i^* A$ as $P_{A^* \vee \Diamond_i \mathsf{x}}$ where A^* stands for the translation of A. Using this translation, Definition 8 and Lemma 4 we prove the following embedding theorem.

Theorem 6. *For all formulae A and B of* PDL *we have:*

1. $\mathsf{H_{SFL}}$ *derives* $[\Box_i^* A \to (A \wedge \Box_i \Box_i^* A)]^*$.
2. $\mathsf{H_{SFL}}$ *derives* $[(A \vee \Diamond_i \Diamond_i^* A) \to \Diamond_i^* A]^*$.
3. *If* $\mathsf{H_{SFL}}$ *derives* $[B \to (A \wedge \Box_i B)]^*$, *then* $\mathsf{H_{SFL}}$ *derives* $[B \to \Box_i^* A]^*$.
4. *If* $\mathsf{H_{SFL}}$ *derives* $[(A \vee \Diamond_i B) \to B]^*$, *then* $\mathsf{H_{SFL}}$ *derives* $[\Diamond_i^* A \to B]^*$.

5 The Infinitary Calculus $\mathsf{T}^\omega_{\mathsf{SFL}}$

In the following we introduce the system $\mathsf{T}^\omega_{\mathsf{SFL}}$ for deriving valid formulae of stratified modal fixed point logic. For ease of presentation we shall restrict the remainder of our account to the case where we are dealing with just one modality, that is to say where M is a singleton set. A generalization of the subsequent arguments to the full multi–modal case is an easy exercise. Consequently, we drop the indices when writing the symbols $\Box$ and $\Diamond$. Furthermore, given a Kripke structure $\mathsf{K} = (S, R, \pi)$ we henceforth treat R as a single relation $R \subset S \times S$.

The calculus $\mathsf{T}^\omega_{\mathsf{SFL}}$ is designed in Tait–style, that is to say, using it we derive finite sets of formulae of $\mathcal{L}$. It is infinitary in the sense that in order to apply the rule for a greatest fixed point constant $Q_{\mathcal{A}}$ we must derive infinitely many premises. As a consequence of this, proofs in $\mathsf{T}^\omega_{\mathsf{SFL}}$ can become infinite in length. Our calculus will not include a cut rule although such a rule could be added for convenience.

Definition 9 (Sequents). *A sequent is a finite set of formulae of $\mathcal{L}$. Henceforth, unless otherwise stated capital Greek letters $\Gamma, \Delta, \Sigma, \dots$ (possibly with primes and subscripts) shall be used to denote sequents. Given a formula A we write Γ, A for $\Gamma \cup \{A\}$. By $\bigvee \Gamma$ we denote the formula $((\dots(A_1 \vee A_2) \vee \dots) \vee A_n)$ if Γ is the set $\{A_1, \dots, A_n\}$ or the formula $\bot$ if Γ is empty. Furthermore, by $\Diamond\Gamma$ we denote the sequent obtained by prefixing each formula in Γ by $\Diamond$.*

We are now in a position to state the rules of $\mathsf{T}^\omega_{\mathsf{SFL}}$. These rules are to be read in the following way: if all sequents displayed in the premise have already been derived, then the sequent displayed in the conclusion may also be derived. In this sense, rules with an empty premise correspond to sequents which may always be derived, that is to say they are axioms of $\mathsf{T}^\omega_{\mathsf{SFL}}$.

Definition 10 (The System $\mathsf{T}^\omega_{\mathsf{SFL}}$). *The system $\mathsf{T}^\omega_{\mathsf{SFL}}$ is defined by the following set of inference rules:*

$$\frac{}{\Gamma, \mathsf{p}, {\sim}\mathsf{p}}\ (ID1) \qquad \frac{}{\Gamma, \mathsf{X}, {\sim}\mathsf{X}}\ (ID2) \qquad \frac{}{\Gamma, \top}\ (ID3)$$

$$\frac{\Gamma, A, B}{\Gamma, A \vee B}\ (\vee) \qquad \frac{\Gamma, A \quad \Gamma, B}{\Gamma, A \wedge B}\ (\wedge)$$

$$\frac{\Gamma, A}{\Diamond\Gamma, \Box A, \Sigma}\ (\Box)$$

$$\frac{\Gamma, \mathcal{A}[P_{\mathcal{A}}]}{\Gamma, P_{\mathcal{A}}}\ (P_{\mathcal{A}}) \qquad \frac{\Gamma, Q^k_{\mathcal{A}} \quad \textit{for all } k \in \omega}{\Gamma, Q_{\mathcal{A}}}\ (Q^\omega_{\mathcal{A}})$$

Provability of a sequent Γ in $\mathsf{T}^\omega_{\mathsf{SFL}}$ is defined as usual and denoted by $\mathsf{T}^\omega_{\mathsf{SFL}} \vdash \Gamma$. It is obvious from the formulation of the system that $\mathsf{T}^\omega_{\mathsf{SFL}}$ satisfies weakening in the sense that if Γ is provable and $\Gamma \subset \Delta$, then Δ is also provable.

However, it is not immediately obvious that $\mathsf{T}^\omega_{\mathsf{SFL}}$ is sound. Problems might occur in connection with the infinitary rule $(Q^\omega_{\mathcal{A}})$ whose premises are exactly the finite stages of greatest fixed points, whereas in arbitrary Kripke structures transfinite stages cannot be ruled out. However, in Section 7 we will prove the soundness of a system $\mathsf{T}_{\mathsf{SFL}}$ which contains $\mathsf{T}^\omega_{\mathsf{SFL}}$. Thus we do not elaborate on the soundness proof for $\mathsf{T}^\omega_{\mathsf{SFL}}$ here.

From the completeness of $\mathsf{T}^\omega_{\mathsf{SFL}}$ which will be shown in Section 6 we can immediately deduce that the Hilbert system $\mathsf{H}_{\mathsf{SFL}}$ is contained in $\mathsf{T}^\omega_{\mathsf{SFL}}$. More formally and moving from Hilbert–style to a Tait–style framework, completeness implies the following Theorem

Theorem 7. *For all formulae $\mathcal{A}$ and B of $\mathcal{L}$ where $\mathcal{A}$ is* X*-positive we have:*

(i) $\mathsf{T}^\omega_{\mathsf{SFL}} \vdash \neg\mathcal{A}[P_{\mathcal{A}}], P_{\mathcal{A}}$.
(ii) $\mathsf{T}^\omega_{\mathsf{SFL}} \vdash \neg\mathcal{A}[B], B \implies \mathsf{T}^\omega_{\mathsf{SFL}} \vdash \neg P_{\mathcal{A}}, B$.

6 A Completeness Proof for $\mathsf{T}^\omega_{\mathsf{SFL}}$

We now aim to prove completeness for the system $\mathsf{T}^\omega_{\mathsf{SFL}}$, that is the claim that any valid formula of $\mathcal{L}$ must be provable in $\mathsf{T}^\omega_{\mathsf{SFL}}$. In order to arrive at this result, we will in fact show its contrapositive. That is we will show that given a non–provable formula A, we may always construct a Kripke structure in which A is not satisfied. Such a formula A can thus never be valid. To show completeness we use an extension of the method of saturated sequents. A sequent is saturated if it is in a sense maximally non–provable. The first major step in proving the completeness of $\mathsf{T}^\omega_{\mathsf{SFL}}$ is thus to show that any non–provable sequent can be expanded to a saturated sequent. The second step is then to show that from the set of all saturated sequents we may construct a suitable countermodel for a non–provable formula. We thus first need to state what saturated sequents are.

Definition 11 (Saturated Sequents). *Let k be a natural number.*

1. *A sequent $\Gamma \subset \mathcal{L}$ is called k–presaturated if all of the following properties hold:*
   (i) $\mathsf{T}^\omega_{\mathsf{SFL}} \nvdash \Gamma$
   (ii) For all formulae $A \wedge B$ with $\mathrm{level}(A\wedge B) \geq k$*, if $A \wedge B \in \Gamma$, then $A \in \Gamma$ or $B \in \Gamma$.*
   (iii) For all formulae $A \vee B$ with $\mathrm{level}(A\vee B) \geq k$*, if $A \vee B \in \Gamma$, then $A \in \Gamma$ and $B \in \Gamma$.*
   (iv) For all formulae $Q_{\mathcal{A}}$ with $\mathrm{level}(Q_{\mathcal{A}}) \geq k$*, if $Q_{\mathcal{A}} \in \Gamma$, then $Q^n_{\mathcal{A}} \in \Gamma$ for some $n \in \omega$.*
2. *Γ is called k–saturated if it is k–presaturated and in addition the following property holds*

(v) For all formulae $P_{\mathcal{A}}$ with $\mathrm{level}(P_{\mathcal{A}}) \geq k$*, if* $P_{\mathcal{A}} \in \Gamma$*, then* $\mathcal{A}[P_{\mathcal{A}}] \in \Gamma$*.*

In general Γ is simply called saturated if it is 0–saturated. To show that any non–provable sequent Γ may be expanded to a saturated one we will use the strategy of choosing a formula in Γ which violates one of the conditions (ii) to (v) of Definition 11, adding suitable formulae to make the respective condition satisfied, then iterating this with a new formula which violates the conditions and so on. The essential step in the proof is then to show that this procedure converges after finitely many steps, thus yielding a saturated sequent. What makes matters complicated is condition (v) since satisfying this condition – in contrast to conditions (ii) to (iv) – means increasing formula complexity. In order to tackle these complications we need to make two technical definitions. The first one introduces a notation for the set of all those formulae in a sequent Γ which prevent Γ from being k–presaturated. The second definition introduces the set $\mathrm{csub}(A)$ of critical subformulae of formula A of $\mathcal{L}$. The intended meaning of $\mathrm{csub}(A)$ is being the set of all subformulae of A which could be considered during the process of saturation.

Definition 12 (k–Deficiency Set). *Let Γ be a sequent such that $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash \Gamma$ and k a natural number. Define the k–deficiency set $\mathrm{ds}_k(\Gamma)$ of Γ as the empty set if Γ is k–presaturated and otherwise as the set of all elements of Γ of level k which violate one of the conditions (ii) – (iv) of Definition 11.*

Definition 13 (Critical Subformulae). *For any $A \in \mathcal{L}$ define the set* $\mathrm{csub}(A)$ *of critical subformulae of A inductively as follows:*

1. *If A is an element of $\Phi \cup \mathsf{V} \cup \mathsf{T}$, a fixed point constant $P_{\mathcal{A}}$, a formula $\Diamond B$ or a formula $\Box B$, then define* $\mathrm{csub}(A)$ *as the set* $\{A\}$*.*
2. *If A is a fixed point constant $Q_{\mathcal{A}}$, then* $\mathrm{csub}(A) := \{Q_{\mathcal{A}}\} \cup \{Q^n_{\mathcal{A}} : n \in \omega\}$*.*
3. *If A is a formula $B \vee C$ or $B \wedge C$, then* $\mathrm{csub}(A) := \{A\} \cup \mathrm{csub}(B) \cup \mathrm{csub}(C)$*.*

For any sequent $\Gamma = \{A_1, \ldots, A_n\}$ set $\mathrm{csub}(\Gamma) := \mathrm{csub}(A_1) \cup \ldots \cup \mathrm{csub}(A_n)$*.*

In order to saturate an arbitrary non–provable sequent Γ we proceed as follows: Start with Γ and find the least number k for which Γ is k–saturated. If $k = 0$, then Γ is saturated and we are done. Thus assume $k = l + 1$ for some number l. Now we iteratively satisfy conditions (ii) to (iv) of Definition 11, producing in a finite number of steps a sequent Γ' which is l–presaturated, but not necessarily l–saturated. To achieve the latter, we add $\mathcal{A}[P_{\mathcal{A}}]$ to Γ' for each fixed point constant of level l and thus obtain a sequent Γ''. Unfortunately Γ'' need not necessarily be l–presaturated any longer but we notice that it is still $l + 1$–saturated. Thus we again iteratively satisfy conditions (ii) to (iv) of Definition 11 where after we arrive at an l–saturated sequent Γ'''. Repeating this procedure yields a strictly decreasing sequence of saturation numbers and thus after finitely many steps of adding formulae we reach a sequent which is saturated. We will now set about formalizing the saturation argument just described in the shape of the next three lemmata.

Lemma 5. *For every sequent Γ with $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash \Gamma$ there exists a natural number k such that Γ is k–saturated.*

Proof. By Definition 11 it is clear that the claim holds if we take k as the maximum level of all formulae in Γ plus 1. $\square$

Lemma 6. *Suppose that k is a natural number and Γ a $(k+1)$–saturated sequent. Then there exists a k–presaturated sequent Δ so that*

$$\Gamma \subset \Delta \text{ and } (\Delta \setminus \Gamma) \subset \mathrm{csub}(\mathrm{ds}_k(\Gamma)).$$

Proof. The proof of this lemma is routine and will thus not be carried out. An almost identical claim is shown by Alberucci and Jäger [1] in full detail. $\square$

Lemma 7. *Suppose that k is a natural number and Γ a $(k+1)$–saturated sequent. Then there exists a k–saturated sequent Δ such that $\Gamma \subset \Delta$.*

Proof. In a first step we apply Lemma 6 in order to obtain a k–presaturated sequent Δ_0 so that $\Gamma \subset \Delta_0$ and $(\Delta_0 \setminus \Gamma) \subset \mathrm{csub}(\mathrm{ds}_k(\Gamma))$. This sequent need not necessarily be k–saturated as condition (v) could be violated for some fixed point constants $P_{\mathcal{A}}$ of level k. To rectify this problem, our next step is to define the sequent $\Delta_1 := \Delta_0 \cup \{\mathcal{A}[P_{\mathcal{A}}] : \mathrm{level}(P_{\mathcal{A}}) = k \text{ and } P_{\mathcal{A}} \in \Delta_0\}$. Now, in turn, Δ_1 need not be k–presaturated since we have no guarantee that

$$\mathrm{ds}_k(\Delta_1) = \mathrm{ds}_k(\{\mathcal{A}[P_{\mathcal{A}}] : \mathrm{level}(P_{\mathcal{A}}) = k \text{ and } P_{\mathcal{A}} \in \Delta_0\})$$

is empty. However, again by using Lemma 6 we can extend Δ_1 to a k–presaturated sequent Δ with the properties $\Delta_1 \subset \Delta$ and $(\Delta \setminus \Delta_1) \subset \mathrm{csub}(\mathrm{ds}_k(\Delta_1))$. Thus all elements of $(\Delta \setminus \Delta_1)$ belong to the set

$$\mathrm{csub}(\{\mathcal{A}[P_{\mathcal{A}}] : \mathrm{level}(P_{\mathcal{A}}) = k \text{ and } P_{\mathcal{A}} \in \Delta_0\})$$

with the consequence that all fixed point constants of level k which are elements of Δ are already elements of Δ_0. Hence Δ is k–saturated and the lemma is shown. $\square$

Combining Lemmata 5 and 7 now takes care of the first part of our completeness proof for $\mathsf{T}^{\omega}_{\mathsf{SFL}}$.

Lemma 8. *For every sequent Γ which is not provable in $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ there exists a saturated sequent Δ such that $\Gamma \subset \Delta$.*

Proof. Assume $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash \Gamma$. Then by Lemma 5 we know that Γ is k–saturated for a suitably chosen k. Moreover, according to Lemma 7 there are sequents $\Gamma_{k-1}, \ldots, \Gamma_1, \Gamma_0$ so that each Γ_i is i–saturated for $0 \leq i \leq k-1$ and

$$\Gamma \subset \Gamma_{k-1} \subset \ldots \subset \Gamma_1 \subset \Gamma_0.$$

To conclude the proof simply set $\Delta := \Gamma_0$. $\square$

Based on the collection of all saturated sequents we now define a Kripke structure $\mathsf{K}_{\mathrm{sat}}$ which will turn out to be a suitable countermodel for any non–provable formula of $\mathcal{L}$. The worlds of $\mathsf{K}_{\mathrm{sat}}$ are just the saturated sequents themselves. Accessibility is defined to treat formulae of the form $\Box B$ and $\Diamond B$ correctly and the valuation function makes a primitive proposition true in any world which does not contain it. This is possible since the non–provability of any saturated sequent Γ guarantees that not both p and $\sim\mathsf{p}$ are elements of Γ at once. The same is also true for the symbols X and $\sim\mathsf{X}$.

Definition 14 (Canonical Countermodel). *Set* $\mathsf{K}_{\mathrm{sat}} := (S_{\mathrm{sat}}, R_{\mathrm{sat}}, \pi_{\mathrm{sat}})$ *and*

$$\begin{aligned} S_{\mathrm{sat}} &:= \{\Gamma \subset \mathcal{L} : \Gamma \textit{ saturated}\}, \\ R_{\mathrm{sat}} &:= \{(\Gamma, \Delta) \in S_{\mathrm{sat}} \times S_{\mathrm{sat}} : \{B \in \mathcal{L} : \Diamond B \in \Gamma\} \subset \Delta\}, \\ \pi_{\mathrm{sat}}(P) &:= \{\Gamma \in S_{\mathrm{sat}} : P \notin \Gamma\} \textit{ for } P \in \Phi \cup \mathsf{V}. \end{aligned}$$

It is easily verified that $\mathsf{K}_{\mathrm{sat}}$ *is a Kripke structure in our sense. Given a formula* $A \in \mathcal{L}$ *and a set* $T \subset S_{\mathrm{sat}}$ *we will write* $\|A\|_{\mathrm{sat}}$ *for* $\|A\|_{\mathsf{K}_{\mathrm{sat}}}$ *and* $\|A\|_{\mathrm{sat}[\mathsf{X}:=T]}$ *for* $\|A\|_{\mathsf{K}_{\mathrm{sat}}[\mathsf{X}:=T]}$.

The completeness proof for $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ will effectively be finished once we have shown for an arbitrary formula A of $\mathcal{L}$ that if A is an element of a saturated set Γ, then $\mathsf{K}_{\mathrm{sat}}$ is a countermodel for A at the world Γ. The next lemma shows that with respect to this property the construction of $\mathsf{K}_{\mathrm{sat}}$ consistently treats formulae of the form $\Box B$ and $\Diamond B$.

Lemma 9. *Assume* $\Gamma \subset \mathcal{L}$ *is a saturated sequent.*

(i) If $\Box A \in \Gamma$*, then there exists a sequent* Δ *such that* $\Gamma R_{\mathrm{sat}} \Delta$ *and* $A \in \Delta$*.*
(ii) If $\Diamond A \in \Gamma$*, then* $A \in \Delta$ *for all sequents* Δ *such that* $\Gamma R_{\mathrm{sat}} \Delta$*.*

Proof. To show (i) consider that since Γ is saturated, we have $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash \Gamma$ and thus also $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash \{B \in \mathcal{L} : \Diamond B \in \Gamma\}, A$. Thus by Lemma 8 there exists a saturated sequent Δ such that $\{B \in \mathcal{L} : \Diamond B \in \Gamma\}, A \subset \Delta$ and thus $A \in \Delta$ and $\Gamma R_{\mathrm{sat}} \Delta$. Part (ii) of the claim is obvious by Definition 14. $\Box$

The proofs of the lemmata leading up to the completeness theorem formally proceed as an induction along the levels of the language $\mathcal{L}$. That is to say we prove that for all formulae A of level 0 if A is an element of a saturated sequent Γ, then $\mathsf{K}_{\mathrm{sat}}$ is a countermodel for A at Γ. Then we show that if this claim holds for all formulae at level k, then it also holds for all formulae at level $k+1$. For this purpose we introduce the notion of k–adequacy of the Kripke structure $\mathsf{K}_{\mathrm{sat}}$.

Definition 15 (k–Adequacy). *Let k be a natural number. We call the Kripke structure* $\mathsf{K}_{\mathrm{sat}}$ *k–adequate if for all saturated sequents Γ and all formulae A of $\mathcal{L}$ we have*

$$\mathrm{level}(A) \le k \textit{ and } A \in \Gamma \implies \Gamma \notin \|A\|_{\mathrm{sat}}.$$

The naive approach to showing that if a formula A is in a saturated sequent Γ, then $\mathsf{K}_{\mathrm{sat}}$ is a countermodel for A at world Γ would be an induction on $\mathrm{rk}(A)$. However, this argument breaks down in the case where A is a constant $P_{\mathcal{A}}$. For this reason we must treat the case of constants $P_{\mathcal{A}}$ separately in the shape of the next lemma.

Lemma 10. *Suppose that $\mathsf{K}_{\mathrm{sat}}$ is k–adequate for the natural number k and let $P_{\mathcal{A}}$ be a fixed point constant of level $k+1$. Then we have for all X-positive formulae $\mathcal{B}$ such that $\mathrm{level}(\mathcal{B}) \leq k$, all saturated sequents Γ and all ordinals α that*

$$\mathcal{B}[P_{\mathcal{A}}] \in \Gamma \implies \Gamma \notin \|\mathcal{B}\|_{\mathrm{sat}[\mathsf{X}:=I^{<\alpha}_{\mathcal{A},\mathsf{K}_{\mathrm{sat}}}]}.$$

Proof. We prove this claim by main induction on α and side induction on $\mathrm{rk}(\mathcal{B})$. The atomic and truth value symbol cases are trivial, the propositional cases follow by hypothesis of the side induction and the modal cases by Lemma 9 and the hypothesis of the side induction. We are thus left with the fixed point and variable cases: In case $\mathcal{B} = P_{\mathcal{C}}$ or $\mathcal{B} = Q_{\mathcal{C}}$ we have $\mathcal{B}[P_{\mathcal{A}}] = \mathcal{B}$ and since $\mathrm{level}(\mathcal{B}) \leq k$ and $\mathsf{K}_{\mathrm{sat}}$ is k–adequate, we get

$$\Gamma \notin \|\mathcal{B}\|_{\mathrm{sat}} = \|\mathcal{B}\|_{\mathrm{sat}[\mathsf{X}:=I^{<\alpha}_{\mathcal{A},\mathsf{K}_{\mathrm{sat}}}]}.$$

In the case where $\mathcal{B} = \mathsf{X}$ from $\mathcal{B}[P_{\mathcal{A}}] \in \Gamma$, i.e. $P_{\mathcal{A}} \in \Gamma$, we immediately obtain $\mathcal{A}[P_{\mathcal{A}}] \in \Gamma$. Now we apply the hypothesis of the main induction and conclude that we have

$$\Gamma \notin \|\mathcal{A}\|_{\mathrm{sat}[\mathsf{X}:=I^{<\beta}_{\mathcal{A},\mathsf{K}_{\mathrm{sat}}}]}$$

for all $\beta < \alpha$. Semantic reasoning yields

$$\Gamma \notin \|\mathsf{X}\|_{\mathrm{sat}[\mathsf{X}:=I^{\beta}_{\mathcal{A},\mathsf{K}_{\mathrm{sat}}}]}$$

for all $\beta < \alpha$. Consequently we have

$$\Gamma \notin \|\mathsf{X}\|_{\mathrm{sat}[\mathsf{X}:=I^{<\alpha}_{\mathcal{A},\mathsf{K}_{\mathrm{sat}}}]},$$

and the claim is shown. $\square$

Corollary 1. *Suppose that $\mathsf{K}_{\mathrm{sat}}$ is k–adequate for the natural number k and let $P_{\mathcal{A}}$ be a fixed point constant of level $k+1$. Then we have for all saturated sequents Γ that*

$$P_{\mathcal{A}} \in \Gamma \implies \Gamma \notin \|P_{\mathcal{A}}\|_{\mathrm{sat}}.$$

The next lemma takes care of the induction step, showing that if the Kripke structure $\mathsf{K}_{\mathrm{sat}}$ is k–adequate, then it is also $(k+1)$–adequate. This property will shortly lead us to the statement that indeed $\mathsf{K}_{\mathrm{sat}}$ is k–adequate for all natural numbers k.

Lemma 11. *Suppose that $\mathsf{K}_{\mathrm{sat}}$ is k–adequate for the natural number k. Then for all formulae A of $\mathcal{L}$ and all saturated sequents Γ we have*

$$A \in \Gamma \textit{ and } \mathrm{level}(A) \leq k+1 \implies \Gamma \notin \|A\|_{\mathrm{sat}}.$$

Proof. We show this lemma by induction on $\mathrm{rk}(A)$. The atomic, variable and truth value symbol cases are trivial. The propositional and modal cases follow by induction hypothesis, the latter using Lemma 9. We are thus left with the fixed point cases: For the first case assume that $A = Q_{\mathcal{B}}$, $Q_{\mathcal{B}} \in \Gamma$ and that $\mathrm{level}(Q_{\mathcal{B}}) \leq k+1$. Then by saturation of Γ we have $Q^l_{\mathcal{B}} \in \Gamma$ for some $l \in \omega$ and $\mathrm{level}(Q^l_{\mathcal{B}}) < \mathrm{level}(Q_{\mathcal{B}}) \leq k+1$. Thus by induction hypothesis $\Gamma \notin \|Q^k_{\mathcal{B}}\|_{\mathrm{sat}}$ and thus also $\Gamma \notin \|Q_{\mathcal{B}}\|_{\mathrm{sat}}$. For the second case assume that $A = P_{\mathcal{B}}$, $P_{\mathcal{B}} \in \Gamma$ and $\mathrm{level}(P_{\mathcal{B}}) \leq k+1$. In the case where $\mathrm{level}(P_{\mathcal{B}}) \leq k$ it follows that $\Gamma \notin \|P_{\mathcal{B}}\|_{\mathrm{sat}}$ by k–adequacy of $\mathsf{K}_{\mathrm{sat}}$, in the case where $\mathrm{level}(P_{\mathcal{B}}) = k+1$ the claim follows by Corollary 1. □

We are now ready to prove the crucial lemma establishing the fact that any formula contained in a saturated sequent Γ is not satisfied in $\mathsf{K}_{\mathrm{sat}}$ at Γ.

Lemma 12. *For any natural number k, the Kripke structure $\mathsf{K}_{\mathrm{sat}}$ is k–adequate; i.e. for all formulae $B \in \mathcal{L}$ and all saturated sequents $\Gamma \subset \mathcal{L}$ we have*

$$B \in \Gamma \implies \Gamma \notin \|B\|_{\mathrm{sat}}.$$

Proof. This lemma is shown by a straightforward induction on k using Lemma 11 in the induction step. □

Theorem 8 (Completeness of $\mathsf{T}^{\omega}_{\mathsf{SFL}}$). *The system $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ is complete, that is for all formulae $A \in \mathcal{L}$ if A is valid, then $\mathsf{T}^{\omega}_{\mathsf{SFL}} \vdash A$.*

Proof. We show the claim by contraposition. Assume $\mathsf{T}^{\omega}_{\mathsf{SFL}} \nvdash A$. Then by Lemma 8 there exists a saturated set $\Gamma \subset \mathcal{L}$ such that $A \in \Gamma$. Therefore by Lemma 12 we have $\Gamma \notin \|A\|_{\mathrm{sat}}$ and thus A cannot be valid. Hence we have shown the completeness of $\mathsf{T}^{\omega}_{\mathsf{SFL}}$. □

7 Finitizing $\mathsf{T}^{\omega}_{\mathsf{SFL}}$

We will now use the so–called small model property of SFL to reduce the number of premises of the greatest fixed point rule $(Q^{\omega}_{\mathcal{A}})$ down to a single premise. Doing so will result in a truly finitary system in which all proofs are finite in length. Before we address the soundness proof of $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ we state the small model property in its customary form.

Remark 2 (Small model property). There exists an exponential function f on the natural numbers such that for every formula $A \in \mathcal{L}$ if A is satisfiable, then there exists a Kripke structure $\mathsf{K} = (S, R, \pi)$ with $|S| < f(|A|)$ which satisfies A.

The small model property holds for stratified modal fixed point logic since it holds for the modal μ–calculus. A candidate for the exponential function f mentioned in Remark 2 can be reconstructed from results presented in [9] or [16]. The exact shape of f or indeed a minimal candidate with respect to our framework shall not concern us in the current study.

Definition 16 (The System $\mathsf{T}_{\mathsf{SFL}}$). *The system $\mathsf{T}_{\mathsf{SFL}}$ is defined by replacing the rule $(Q^{\omega}_{\mathcal{A}})$ in the system $\mathsf{T}^{\omega}_{\mathsf{SFL}}$ by the rule*

$$\frac{\Gamma, Q^{k}_{\mathcal{A}}}{\Gamma, Q_{\mathcal{A}}, \Sigma} \quad (Q_{\mathcal{A}})$$

where $k = f(|\bigvee(\Gamma, Q_{\mathcal{A}})|)$.

The function f guaranteed to exists by Remark 2 is used to bound the number of iterations $Q^{k}_{\mathcal{A}}$ for which we need to check the derivability of $\Gamma, Q^{k}_{\mathcal{A}}$ before applying the rule $(Q_{\mathcal{A}})$ to conclude $\Gamma, Q_{\mathcal{A}}$. Indeed f supplies us with the only such iteration we need to check explicitly, as the subsequent argument will show. In order to obtain weakening the conclusion of the rule $(Q_{\mathcal{A}})$ needs to weakened by an arbitrary sequent Σ explicitly.

The structure of the rule $(Q_{\mathcal{A}})$ makes it clear that completeness of $\mathsf{T}_{\mathsf{SFL}}$ is implied by completeness of $\mathsf{T}^{\omega}_{\mathsf{SFL}}$. With completeness of $\mathsf{T}_{\mathsf{SFL}}$ taken care of the only task that remains is to show soundness. Before we proceed to show the soundness of $\mathsf{T}_{\mathsf{SFL}}$ we establish some consequences of the small model property. The first fact we note is essentially just the contraposition of Remark 2.

Lemma 13. *Let $A \in \mathcal{L}$. If A is valid in all Kripke structures K with $|\mathsf{K}| \leq f(|A|)$, then A is valid.*

Proof. The proof of this claim is trivial by the fact that $|\neg A| = |A|$. □

The second fact we require in order to show the soundness of $\mathsf{T}_{\mathsf{SFL}}$ states that in order to check the validity of greatest fixed point in all Kripke structures with at most k worlds, we only need to check the validity of its k-th approximation. The essential ingredient to proving this claim is the boundedness of closure ordinals established in Theorem 2.

Lemma 14. *If the formula $\bigvee(\Gamma, Q^{k}_{\mathcal{A}})$ is valid, then the formula $\bigvee(\Gamma, Q_{\mathcal{A}})$ is valid in all Kripke structures K with $|\mathsf{K}| \leq k$.*

Proof. Let $k \in \omega$ be arbitrary and K be a Kripke structure with $|\mathsf{K}| \leq k$. Since the formula $\bigvee(\Gamma, Q^{k}_{\mathcal{A}})$ is valid it is also valid in K. Since $|\mathsf{K}| \leq k$, Theorem 2 guarantees that the formula $Q^{k}_{\mathcal{A}} \leftrightarrow Q_{\mathcal{A}}$ is valid in K. Thus $\bigvee(\Gamma, Q_{\mathcal{A}})$ is valid in K and the claim is shown. □

We are now ready to state and prove the soundness of the system $\mathsf{T}_{\mathsf{SFL}}$.

Theorem 9 (Soundness of $\mathsf{T}_{\mathsf{SFL}}$). *The system $\mathsf{T}_{\mathsf{SFL}}$ is sound, that is for all sequents $\Gamma \subset \mathcal{L}$ if $\mathsf{T}_{\mathsf{SFL}} \vdash \Gamma$, then the formula $\bigvee \Gamma$ is valid.*

Proof. We must show that all axioms of $\mathsf{T}_{\mathsf{SFL}}$ are valid and that all rules of $\mathsf{T}_{\mathsf{SFL}}$ preserve validity. However, in view of the definition of $\mathsf{T}_{\mathsf{SFL}}$ we merely need to check the rules $(P_{\mathcal{A}})$ and $(Q_{\mathcal{A}})$. For the rule $(P_{\mathcal{A}})$ assume the formula $\bigvee(\Gamma, \mathcal{A}[P_{\mathcal{A}}])$ representing the premise of $(P_{\mathcal{A}})$ is valid. By the semantics of $\mathcal{L}$ the formula $\mathcal{A}[P_{\mathcal{A}}] \leftrightarrow P_{\mathcal{A}}$ is also valid and hence the formula $\bigvee(\Gamma, P_{\mathcal{A}})$ representing the conclusion of $(P_{\mathcal{A}})$ is valid and thus the rule is sound. For the rule $(Q_{\mathcal{A}})$ assume that the formula $\bigvee(\Gamma, Q^k_{\mathcal{A}})$ where $k = f(|\bigvee(\Gamma, Q_{\mathcal{A}})|)$ representing the premise of $(Q_{\mathcal{A}})$ is valid. Therefore by Lemma 14 the formula $\bigvee(\Gamma, Q_{\mathcal{A}})$ is valid in all Kripke structures K with $|\mathsf{K}| \leq k$. Therefore by Lemma 13 the formula $\bigvee(\Gamma, Q_{\mathcal{A}})$ representing the conclusion of $(Q_{\mathcal{A}})$ is valid and thus the rule is sound. This concludes the soundness proof for the system $\mathsf{T}_{\mathsf{SFL}}$. $\square$

8 Concluding Remarks

The approach to stratification presented in this paper can be easily generalized to a form of stratification which permits, at each level, the introduction of simultaneously defined least and greatest fixed points. To give a simple example, extend our language $\mathcal{L}_0$ so that it comprises countably many variables $\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3, \ldots$ and their negations $\sim\mathsf{X}_1, \sim\mathsf{X}_2, \sim\mathsf{X}_3, \ldots$. Now assume that we are given a system of, say, three formulas

$$\mathcal{A}_1[\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3], \quad \mathcal{A}_2[\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3], \quad \mathcal{A}_3[\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3],$$

all positive in $\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3$. Then we add three new constants $P_{\mathcal{A}_1}$, $P_{\mathcal{A}_2}$ and $P_{\mathcal{A}_3}$ as well as three constants $Q_{\mathcal{A}_1}$, $Q_{\mathcal{A}_2}$ and $Q_{\mathcal{A}_3}$ to represent the least and greatest fixed points simultaneously generated by these three formulas.

In the Hilbert system for stratified simultaneous modal fixed point logic the corresponding closure axioms read as follows:

$$\begin{aligned}
\mathcal{A}_1[P_{\mathcal{A}_1}, P_{\mathcal{A}_2}, P_{\mathcal{A}_3}] &\to P_{\mathcal{A}_1}, \\
\mathcal{A}_2[P_{\mathcal{A}_1}, P_{\mathcal{A}_2}, P_{\mathcal{A}_3}] &\to P_{\mathcal{A}_2}, \\
\mathcal{A}_3[P_{\mathcal{A}_1}, P_{\mathcal{A}_2}, P_{\mathcal{A}_3}] &\to P_{\mathcal{A}_3}.
\end{aligned}$$

Furthermore, for all formulas B_1, B_2, B_3 we have as induction rules:

$$\frac{(\mathcal{A}_1[B_1, B_2, B_3] \to B_1) \wedge (\mathcal{A}_2[B_1, B_2, B_3] \to B_2) \wedge (\mathcal{A}_3[B_1, B_2, B_3] \to B_3)}{P_{\mathcal{A}_i} \to B_i}$$

where $i \in \{1, 2, 3\}$. This can be done for arbitrary (finite) systems of positive formulas and analogously carried through for all levels of stratification.

Even though the languages thus obtained and the corresponding systems are strictly more expressive than the ones studied in this article (cf. e.g. [2, 3]), all our methods remain applicable in this more general framework.

References

1. Luca Alberucci and Gerhard Jäger. About cut–elimination for logics of common knowledge. *Annals of Pure and Applied Logic*, 133:73–99, 2005.

2. Dietmar Berwanger, Erich Grädel, and Giacomo Lenzi. On the variable hierarchy of the modal μ–calculus. In *CSL '02: Proceedings of the 16th International Workshop and 11th Annual Conference of the EACSL on Computer Science Logic*, pages 352–366. Springer, 2002.
3. Dietmar Berwanger and Giacomo Lenzi. The variable hierarchy of the modal μ–calculus is strict. In *STACS 2005, Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science*, LNCS. Springer, 2005. (to appear).
4. Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Workshop on Logics of Programs*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.
5. E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu–calculus and determinacy (extended abstract). In *Proceedings of the 32th Annual IEEE Symposium on Foundations of Computer Science*, pages 3368–3377. IEEE Computer Society Press, 1991.
6. Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
7. Michael J. Fischer and Richard E. Ladner. Propositional modal logic of programs. In *Proc. 9th ACM Annual Symposium on Theory of Computing*, pages 286 – 294, 1977.
8. Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computing and System Science*, 18(2):194 – 211, 1979.
9. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
10. Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut–free common knowledge. submitted, 2005.
11. Dexter Kozen. Results on the propositional modal μ–calculus. *Theoretical Computer Science*, 27:333–354, 1983.
12. Yannis Moschovakis. *Elementary Induction on Abstract Structures*. Studies in Logic and the Foundations of Mathematics. North-Holland, 1974.
13. Colin Stirling and David Walker. Local model–checking in the modal mu–calculus. *Theoretical Computer Science*, 89:161–177, 1991.
14. Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional modal mu-calculus. *Information and Computation*, 81(3):249–264, 1989.
15. Igor Walukiewicz. Completeness of Kozen's axiomatization of the propositional μ–calculus. *Information and Computation*, 157:142–182, 2000.
16. Thomas Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bull. Soc. Math. Belg.*, 8(2), 2001.
17. Glynn Winskel. A note on model checking the modal ν–calculus. *Theoretical Computer Science*, 83:157–167, 1991.