

Proof-theoretic aspects of modal logic with fixed points

Inauguraldissertation
der Philosophisch-naturwissenschaftlichen Fakultät
der Universität Bern

vorgelegt von
Mathis Kretz
von Kriens

Leiter der Arbeit:
Prof. Dr. G. Jäger
Institut für Informatik und angewandte Mathematik

Proof-theoretic aspects of modal logic with fixed points

Inauguraldissertation
der Philosophisch-naturwissenschaftlichen Fakultät
der Universität Bern

vorgelegt von
Mathis Kretz
von Kriens

Leiter der Arbeit:
Prof. Dr. G. Jäger
Institut für Informatik und angewandte Mathematik

Contents

Introduction	5
Outline	9
Related work	10
Prerequisites	11
Acknowledgements	11
1 Monotone operators	13
1.1 Basic definitions	13
1.2 Some standard results	14
2 Stratified modal fixed point logic	19
2.1 The language \mathcal{L}_{SFL}	19
2.2 The semantics of \mathcal{L}_{SFL}	22
2.3 Formula complexity	24
3 Some notable fragments of SFL	27
3.1 The Hilbert system H_{SFL}	27
3.2 Logic of common knowledge	29
3.3 Computational tree logic	31
3.4 Propositional dynamic logic	33
4 The infinitary system $\text{T}_{\text{SFL}}^\omega$	35
4.1 The system $\text{T}_{\text{SFL}}^\omega$	35
4.2 Saturated sequents	38
4.3 Completeness of $\text{T}_{\text{SFL}}^\omega$	42
4.4 Embedding H_{SFL} into $\text{T}_{\text{SFL}}^\omega$	45
5 Finitising $\text{T}_{\text{SFL}}^\omega$	49
5.1 The small model property and the system T_{SFL}	49
5.2 Soundness of T_{SFL}	51

6	Closure ordinals in SFL	53
6.1	Closure ordinals in $\mathcal{L}_{\text{SFL}}^1$	53
6.2	Closure ordinals in $\mathcal{L}_{\text{SFL}}^2$	59
7	The modal μ-calculus	67
7.1	The languages \mathcal{L}_μ and \mathcal{L}_μ^+	67
7.2	Semantics of \mathcal{L}_μ^+	71
7.3	Formula complexity	73
7.4	Semantics of \mathcal{L}_μ^+ under signatures	79
8	The infinitary systems \mathbb{T}_μ^ω and $\mathbb{T}_{\mu+}^\omega$	87
8.1	The systems \mathbb{T}_μ^ω and $\mathbb{T}_{\mu+}^\omega$	87
8.2	D -saturated sequents	90
8.3	Completeness of $\mathbb{T}_{\mu+}^\omega$	95
8.4	Finitising \mathbb{T}_μ^ω	98
8.5	A note on expressivity	99
	Concluding remarks	103
	Further work	104
	Bibliography	107
	Index	111

Introduction

Propositional modal logic, which in the scope of this thesis will be referred to merely as modal logic, arises from propositional logic by enriching the syntax with the two dual constructs $\Box A$ and $\Diamond A$. The intended meaning of the former is that A “is necessary” while the latter is usually read as A “being possible”. Even though the syntactic definition of (various systems of) modal logic dates at least as far back as the work of Gödel [17] a formal definition of the semantics in terms of so-called possible worlds is more recent and most commonly ascribed to Kripke [26]. In this semantics, formulae of the language of modal logic are evaluated in Kripke structures which are structures consisting of a set of possible worlds, a binary accessibility relation on this set and a valuation function. When evaluating a formula of modal logic in a Kripke structure we proceed as follows: an atomic proposition is satisfied at a world in a given Kripke structure, if the valuation function makes it true at this world. Furthermore, satisfaction of propositional formulae is defined canonically. This leaves the case of modal formulae for which the binary relation of the structure is consulted: a formula $\Box A$ is satisfied at a world in a given Kripke structure if A is satisfied at all worlds accessible from the current one. Dually, a formula $\Diamond A$ is satisfied at a world in a given Kripke structure if there exists a world accessible from the current one at which A is satisfied.

The specification of an abstract semantics in terms of Kripke structures has led to a more widespread interest in modal logic beyond philosophical questions concerning possibility and necessity. Replacing the accessibility relation by an indexed family of such relations, for instance, can lead to a formal framework suitable for reasoning about belief (or knowledge, depending on additional properties required of the accessibility relation) in a group of agents. In this context a formula $\Box_i A$ is interpreted to mean “agent i knows A ”, thus understanding knowledge as an agent’s inability to imagine a world in which A does not hold. A good overview of many different systems of modal epistemic logics is provided for example by Halpern and Moses [18]. From a different viewpoint, Kripke structures can be seen as state transition

graphs describing the behaviour of a system where the formula $\Box_i A$ is taken to mean that all executions of the action i in the current state lead to a state in which A holds. A classical study in this direction is that of Hennessy and Milner [19] where modal logic is used to characterise notions of process invariance. More directly mathematical applications of modal logic range from so-called provability logic as described for example by Artemov and Beklemishev [4] where a formula $\Box A$ is interpreted as the statement “ A is provable (in a suitable axiomatic system)” to non-wellfounded set theory where formulae of modal logic can be employed to characterise classes of sets generated by the antifoundation axiom as treated amongst others by Barwise and Moss [6].

Compared to other extensions of propositional logic like, for instance, first order predicate logic, modal logics have some decisive advantages. Among these are the decidability of the satisfaction problem (thus of answering the question whether a formula is satisfied in a Kripke structure at a certain world) and the finite model property (thus the property that if a formula is satisfiable then it is satisfied in a finite model, the size of which is a function of the size of the formula). Nevertheless, from the point of view of expressivity modal logics commonly suffer from a major drawback. Modal formulae can only be used to describe inherently bounded and finitary properties of Kripke structures. Thus, for example, when using modal logic to describe the knowledge of a group of agents, complex epistemic situations like common knowledge which require infinite iterations of the knowledge operator are not adequately expressible. Similarly, when interpreting a Kripke structure as a state transition graph of a system, modal logic cannot be used to capture behaviour which is in any way unbounded or refers to infinite runs of the system. Thus, for example, properties like safety (the system, running for an infinitely long time, never reaches a certain undesirable state), liveness (the system, running for an infinitely long time, reaches a certain desirable state eventually) or fairness (the system, running for an infinitely long time, always returns to a certain desirable state after finitely many steps) cannot be expressed as formulae of modal logic.

In order to overcome these limitations fixed point constructs may be introduced at the syntactic level. Central to this approach is the fact that if second order variables X, Y, Z, \dots are added to the modal language at hand, then every formula \mathcal{A} in which a given variable X appears only positively determines a monotone operator $F_{\mathcal{A}, X}^K$ on the domain of any Kripke structure K in which \mathcal{A} is evaluated. By standard results which will be treated in detail later the monotone operator $F_{\mathcal{A}, X}^K$ always has a least and a greatest fixed point for which special syntactic constructs are then added to the lan-

guage. With respect to the structure of the positive formula \mathcal{A} various levels of generality are conceivable, three of which will be studied in this thesis:

1. \mathcal{A} is an instance of a fixed formula scheme. As we will see later, this approach is taken in the case of the Logic of Common Knowledge as treated for example by Fagin, Halpern, Moses and Vardi [14] where epistemic modal logic is extended by an operator CA to express that “ A is common knowledge among all agents” and which turns out to be the greatest fixed point of the formula $\Box_1(A \wedge X) \wedge \dots \wedge \Box_m(A \wedge X)$ where $\{1, \dots, m\}$ is the set of all agents. Similarly, we will see that Computational Tree Logic (CTL) investigated for example by Clarke and Emerson [11] and Propositional Dynamic Logic (PDL) introduced by Fischer and Ladner [15, 16] can both be obtained from modal logic by adding fixed points generated by fixed formula schemes.
2. \mathcal{A} is a general formula, positive in X but fixed points which occur in \mathcal{A} do not depend on the denotation of \mathcal{A} itself. This will be the central principle in connection with stratified modal fixed point logic (SFL) for which we will provide sound and complete cut-free axiomatisations and which can be shown to contain all of the logics mentioned under 1.
3. \mathcal{A} is a general formula positive in X , possibly containing fixed points in which X is a free variable and which thus in turn depend on the denotation of \mathcal{A} . This is the situation which we will encounter when studying the propositional modal μ -calculus as introduced by Kozen [23] for which we will also provide sound and complete cut-free axiomatisations and which is in a sense the most general framework we shall consider.

Two main facts contribute greatly to the appeal of the extensions of modal logics by fixed points described above. Firstly, all three types represent true extensions in terms of expressivity. Thus, for example, viewing a Kripke structure as a state transition graph, the logic CTL is able to express properties like safety and liveness, while the full propositional μ -calculus can capture fairness. Similarly PDL can express unbounded properties like the fact that “all finite iterations of a program i result in a state in which A holds” which would not be expressible purely in terms of Hennessy Milner logic. The second appealing fact about modal logic with fixed points is that both decidability and the finite model property are preserved across all three levels of generality described above. On the most general level the finite model property and decidability have been established by Streett and Emerson [32] as well as by Kozen [24]. More particular proofs of the finite model

property for less general logics, along with respective upper bounds on the size of the finite models, have been provided by Fagin, Halpern, Moses and Vardi [14] for the case of the Logic of Common Knowledge and Fisher and Ladner [16] for PDL. Similarly, decidability of the model-checking problem (the problem of deciding whether a state transition graph satisfies a certain property given as a formula) for CTL, along with its complexity, has been established by Clarke, Emerson and Sistla [12].

For the purpose of exposition we will briefly focus on some further important results obtained for the most general setting of modal logic with fixed points, all of which are of semantical nature. The model checking problem for the modal μ -calculus has been addressed among others by Stirling and Walker [31], Winskel [36] as well as by Emerson, Jutla and Sistla [13]. With respect to expressivity Janin and Walukiewicz [22] have shown that the propositional modal μ -calculus corresponds to the bisimulation invariant fragment of monadic second order logic whereas Bradfield [9] and Lenzi [27] have independently shown that each level of the alternation hierarchy over the language of the μ -calculus is strictly more expressive than any levels below. A similar result was recently also obtained by Berwanger and Lenzi [8] with respect to the variable hierarchy over the same language.

As suggested by its title, the emphasis of this thesis lies on proof-theoretic aspects of modal logics with fixed points and thus on an approach which has received considerably less attention in the literature so far (for related work see below). Unlike the situation in model checking where the interest lies in semantically evaluating a given formula in one particular structure and often at one particular world, the proof-theoretic approach focuses on the problem of syntactically characterising the set of all valid formulae of a logic – that is to say formulae which are satisfied in any Kripke structure at any world – using a set of axioms and inference rules. Roughly structured into two parts, this thesis will introduce axiomatisations for modal logic with fixed points from two of the three levels of generality described above. Firstly, we will study axiomatisations for SFL and, secondly, axiomatisations for the propositional modal μ -calculus. In both cases this will mainly consist of proving the soundness and completeness of the respective axiomatisation, that is to say proving that any formula which is derivable syntactically is valid semantically and, conversely, that any semantically valid formula is indeed derivable syntactically. As will become apparent, proving the latter is heavily dependent on the generality of the fixed point principle in question.

All axiomatisations presented in this thesis are cut-free in the sense that they are complete without a rule allowing the deduction of a formula A once the formulae $B \vee A$ and $\neg B \vee A$ have been derived, even though such a rule

can be added without destroying the soundness property. The absence of a cut rule is an important feature of these axiomatisations with respect to potential applications, since it represents a first step in obtaining an effective procedure for deciding whether a given formula is valid by searching for a proof in a suitable axiomatisation. Once such a procedure is established the proof-theoretic approach could have many potential applications complementing existing model checking methods. To name but one example, when interpreting Kripke structures as state transition graphs, model checking is ideal for verifying systems for which the complete state space along with every possible state transition is known in advance. On the other hand, proof-theoretic methods could be used for checking properties against partial specifications of systems, given as finite sets of formulae. If $\{A_1, \dots, A_n\}$ is such a partial specification, checking a property B would amount to searching for a proof of the implication $(A_1 \wedge \dots \wedge A_n) \rightarrow B$. If a proof is found, this means that any system which satisfies all of the properties A_1, \dots, A_n , also satisfies B . If no proof is found, an ideal proof-search procedure would produce a countermodel to the implication in question.

Outline

As mentioned before, this thesis is implicitly structured into two parts. The first part is concerned with cut-free axiomatisations for the logic SFL, the second part is concerned with such axiomatisations for the full propositional modal μ -calculus. Following the introduction, Chapter 1 will provide the essentials on the theory of monotone operators and inductive definitions which will be used constantly throughout the rest of the thesis. Chapter 2 will introduce the logic SFL and Chapter 3 will discuss Logic of Common Knowledge, CTL and PDL as fragments of SFL. Proceeding towards establishing a cut-free axiomatisation of SFL, Chapter 4 will introduce an infinitary deductive system and show that this system is complete. The soundness proof is deferred to Chapter 5 where we transform the infinitary system into a finitary one which is at least as powerful in terms of the set of provable formulae. In Chapter 6 we conclude the first part of the thesis by investigating a direct application of the proof-theoretic methods devised so far, leading to results on closure ordinals for two fragments of SFL. The second part begins with Chapter 7 where the propositional modal μ -calculus is introduced along with some important notions used in subsequent proofs. Chapter 8 presents two infinitary deductive systems one of which is more practical from a technical point of view when proving completeness, the other is more streamlined, its completeness is implied by that of the first system and it also admits finitisa-

tion analogous to Chapter 5. The thesis ends with some concluding remarks and directions for further work.

The methods used for the propositional modal μ -calculus in the second part of the thesis are generalisations of those used in the first part. Nevertheless, it is instructive to read both parts in order to obtain a first idea of the structure of the argument in a simpler setting before moving on to see at which point the more general fixed point principle complicates matters.

Related work

Before embarking on the path of research outlined above we discuss some of the work which is more closely related to the subject of this thesis, reviewing a number of proof-theoretic approaches to modal logic with fixed points, mainly on the level of the propositional modal μ -calculus. In his initial study, Kozen [23] proposes an axiomatisation of the modal μ -calculus which he shows to be sound and complete for the so-called aconjunctive fragment. Furthermore, in his paper on the relation of the theory of wellquasiorders to the finite model property, Kozen [24] introduces an infinitary deduction rule similar to the ones used in this thesis and claims soundness and completeness for a system with this rule, however, making crucial use of a cut rule. The proof of completeness for Kozen's finitary axiomatisation with respect to the standard semantics and the full language of the propositional modal μ -calculus remained open for some time, although Walukiewicz [34] shows completeness for an interesting alternative system. Ambler, Kwiatkowska and Measor [2] prove completeness for both of Kozen's axiomatisations with respect to an alternative semantics in terms of an extension of modal duality theory. Kozen's original finitary axiomatisation is readdressed by Walukiewicz [35] who shows its completeness using some deep results from automata theory. An interesting direction of proof-theoretical research on the subject is also represented by Miculan [28] who studies natural deduction-style translations of Kozen's systems and their implementation in interactive theorem proving environments. More remotely, the approach taken by Andersen, Stirling and Winskel [3] as well as the follow-up work by Berezin and Gurov [7] also studies proof-systems for the propositional modal μ -calculus. However, the purpose of these systems is to derive local satisfaction statements of the form "A holds for process p " and not the global validity of a given formula.

The problem of obtaining cut-free axiomatisations has also been addressed for less general modal logics with fixed points. For Logic of Common Knowledge Alberucci and Jäger [1] obtain a partial cut-elimination result for a

finitary axiomatisation and a total cut-elimination result for an infinitary one. Using a different idea, also taken up in this thesis, Jäger, Kretz and Studer [21] are able to obtain a total semantic cut-elimination result for a finitary axiomatisation of Logic of Common Knowledge.

Prerequisites

The reader of this thesis is assumed to have some familiarity with the following subjects:

1. Basic set-theory, namely standard notions such as union, intersection, complementation, De Morgan's laws, set inclusion, power set, relations, functions and the notation $\{x \in X : P(x)\}$ for defining sets by collecting all elements of a set X which satisfy a property P .
2. The theory of cardinals, namely the notion of the cardinality $|X|$ of a set X and the least cardinal $|X|^+$ greater than the cardinality of X .
3. The theory of ordinals and wellorderings, namely notions such as transfinite induction along an arbitrary wellordering, ordinal arithmetic and the natural sum for ordinals.
4. Basic propositional logic, namely notions such as contraposition and tautology.

With regard to subjects which are not among those mentioned above, the thesis is intended to be self contained.

Acknowledgements

My foremost thanks are due to Gerhard Jäger for supervising this thesis, for giving me the chance to work in a stimulating scientific environment and for letting me profit from some of his great experience and intuition in the field of theoretical computer science and logic. The second person I would like to thank is Thomas Studer for taking so much interest in my work and for spending a lot of his own time to solve some of my problems. Many results presented in this thesis are based on joint work with the two people above. I am also grateful to the co-referee, Pierluigi Minari, for providing me with important comments on an earlier version of this thesis. Furthermore, the work presented in Chapter 6 was inspired by correspondence with Giacomo Lenzi and a discussion with Martin Lange. I would also like to thank all

members of the research group for theoretical computer science and logic in Bern for many interesting discussions during the traditional coffee breaks and lunch times and for generally making our offices an enjoyable place to work at. Moreover, I especially thank Thomas Strahm for entrusting me with the exercise lecture for his course on programming which was always a challenge and a refreshing change from my more theoretical work.

On a more private note, I would like to express my deepest gratitude towards my parents Marion and Andreas Kretz-Lenz as well as my sister Nicolette Kretz for the relentless love and support I receive from them. Furthermore, I thank all of my friends, old and new, who keep me going and who simply make everything worthwhile. Finally, my thanks go to Simone Artho for writing that postcard and for the summers, autumns, winters and springs that have passed ever since.

Bern, September 2006

Mathis Kretz

Chapter 1

Monotone operators

One of the central notions when studying modal logic with fixed points is that of a monotone operator. Implicitly, monotone operators are at the core of every inductive definition and are thus abundant in every-day mathematics. For example, whenever a set is defined by providing some basic elements along with certain closure conditions, the set can be viewed as having been obtained by starting with the empty set and iteratively applying a monotone operator which ensures the presence of the basic elements and implements the closure conditions until nothing new is added. We begin our account by reviewing those parts of the theory of monotone operators which are prerequisite to our work. For a more thorough introduction to the field the reader is referred to the standard literature, for example the textbooks by Moschovakis [29] or Barwise [5].

We proceed by first discussing the basic definitions and, secondly, by proving some results concerning fixed points and the closure behaviour of monotone operators. We will prove the results even though they are standard in order to acquire some of the mathematical flavour of the subject.

1.1 Basic definitions

An operator is a function from the subsets of a given set to the subsets of that same set and it is said to be monotone if it respects set inclusion.

Definition 1.1.1 (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 $F(B) \subset F(C)$.

The application of a given operator may be iterated, potentially transfinitely many times. There are two ways of defining such iterations which we will

consider. Firstly, we may start from the empty set and collect up all elements which are possibly added by repeated applications of the operator. When speaking informally, we will refer to this process as iteration from below. Secondly, we may start with the whole domain and throw away any elements which are possibly removed by a repeated application of the operator. This process we will informally refer to as iteration from above.

Definition 1.1.2 (Iterations). Let A be a set, F a monotone operator on A and α an ordinal. Define the sets I_F^α , I_F and $I_F^{<\alpha}$ as follows:

$$I_F^{<\alpha} = \bigcup_{\beta < \alpha} I_F^\beta \qquad I_F^\alpha = F(I_F^{<\alpha}) \qquad I_F = \bigcup_{\beta} I_F^\beta$$

Furthermore, we define the sets J_F^α , J_F and $J_F^{<\alpha}$ as follows:

$$J_F^{<\alpha} = \bigcap_{\beta < \alpha} J_F^\beta \qquad J_F^\alpha = F(J_F^{<\alpha}) \qquad J_F = \bigcap_{\beta} J_F^\beta$$

1.2 Some standard results

The sequence of sets determined by iterating a monotone operator transfinitely many times in the two ways described in Definition 1.1.2 always converges to a certain set and the number of iterations at which convergence is reached is bounded by the cardinality of the domain. Furthermore, the sets obtained by this iteration process are always fixed points of the operator in question, indeed they are least fixed points when iterating from below and greatest fixed points when iterating from above. In this sense, given a monotone operator F the iterations I_F^α and J_F^α can be viewed as being approximations of the least and greatest fixed points of F , respectively. As a consequence we also obtain the well-known result that each monotone operator has a least and a greatest fixed point, commonly attributed to Tarski and Knaster [33]. The first theorem which we prove takes care of the least fixed point case.

Theorem 1.2.1. *Let A be a set and F a monotone operator on A . Then the following statements hold:*

- (i) *If $\beta \leq \alpha$, then $I_F^{<\beta} \subset I_F^{<\alpha}$ and $I_F^\beta \subset I_F^\alpha$.*
- (ii) *There exists an ordinal κ such that $|\kappa| \leq |A|$ and $I_F = I_F^\kappa = I_F^{<\kappa}$.*
- (iii) *$F(I_F) = I_F$.*
- (iv) *$I_F = \bigcap \{B : F(B) \subset B\} = \bigcap \{B : F(B) = B\}$.*

Proof.

(i): Trivially we have $I_F^{<\beta} \subset I_F^{<\alpha}$, thus since F is monotone

$$I_F^\beta = F(I_F^{<\beta}) \subset F(I_F^{<\alpha}) = I_F^\alpha.$$

(ii): We first show that there exists some $\kappa < |A|^+$ for which

$$I_F^\kappa = I_F^{<\kappa}. \quad (1.1)$$

Assume the contrary. Then for all $\alpha < |A|^+$ we would have $I_F^{<\alpha} \subsetneq I_F^\alpha$. Thus for each $\alpha < |A|^+$ we could choose some $x_\alpha \in I_F^\alpha \setminus I_F^{<\alpha}$. Then the set $X := \{x_\alpha : \alpha \in |A|^+\}$ would contain $|A|^+$ many elements, but this is impossible since $X \subset A$! Thus we have shown (1.1).

In order to prove that $I_F^\kappa = I_F$ we show by induction on α that for every $\alpha \geq \kappa$ we have $I_F^\alpha = I_F^\kappa$. If $\kappa = 0 = \alpha$, then the claim trivially holds. Thus assume the claim holds for all $\beta < \alpha$ and furthermore $\alpha > \kappa$ since otherwise the claim is trivial. Then by induction hypothesis for all β such that $\kappa \leq \beta < \alpha$ we have $I_F^\beta = I_F^\kappa$. Since by (i) furthermore $I_F^\gamma \subset I_F^\kappa$ for all $\gamma < \kappa$, we obtain by (1.1)

$$\bigcup_{\beta < \alpha} I_F^\beta = I_F^\kappa = I_F^{<\kappa}.$$

This yields the equalities

$$I_F^\alpha = F\left(\bigcup_{\beta < \alpha} I_F^\beta\right) = F(I_F^{<\kappa}) = I_F^\kappa$$

and the claim is shown.

Therefore, by the claim and (1.1) we obtain

$$I_F = \bigcup_{\alpha} I_F^\alpha = \bigcup_{\alpha < \kappa} I_F^\alpha = I_F^{<\kappa} = I_F^\kappa$$

and (ii) is shown.

(iii): By (ii) we may reason as follows:

$$F(I_F) = F(I_F^\kappa) = F(I_F^{<\kappa}) = I_F^\kappa = I_F.$$

(iv): We begin with the first equality. Assume $B \subset A$ has the property that $F(B) \subset B$. We show by induction on α that for any ordinal α

$$I_F^\alpha \subset B. \quad (1.2)$$

Since $I_F^0 = \emptyset \subset B$ the claim holds for $\alpha = 0$. Therefore, assume the claim holds for all $\beta < \alpha$. By the induction hypothesis we have

$$I_F^{<\alpha} = \bigcup_{\beta < \alpha} I_F^\beta \subset B,$$

thus since F is monotone

$$I_F^\alpha = F(I_F^{<\alpha}) \subset F(B) \subset B.$$

This completes the induction and (1.2) holds.

Now by (1.2) we get

$$I_F = \bigcup_{\alpha} I_F^\alpha \subset B.$$

Since B was arbitrary we obtain

$$I_F \subset \bigcap \{B : F(B) \subset B\}. \quad (1.3)$$

By (iii) we also have $F(I_F) \subset I_F$, so $I_F \in \{B : F(B) \subset B\}$ and thus if $x \in \bigcap \{B : F(B) \subset B\}$, then also $x \in I_F$, yielding

$$I_F \supset \bigcap \{B : F(B) \subset B\}. \quad (1.4)$$

Now from (1.3) and (1.4) we obtain $I_F = \bigcap \{B : F(B) \subset B\}$. The second equality, namely

$$I_F = \bigcap \{B : F(B) = B\},$$

follows immediately by (iii) and thus we have shown (iv). □

A similar theorem can also be shown with respect to greatest fixed points. Its proof works by appropriately dualising the proof of Theorem 1.2.1.

Theorem 1.2.2. *Let A be a set and F a monotone operator on A . Then the following statements hold:*

- (i) If $\beta \leq \alpha$, then $J_F^{<\alpha} \subset J_F^{<\beta}$ and $J_F^\alpha \subset J_F^\beta$.
- (ii) There exists an ordinal κ such that $|\kappa| \leq |A|$ and $J_F = J_F^\kappa = J_F^{<\kappa}$.
- (iii) $F(J_F) = J_F$.
- (iv) $J_F = \bigcup\{B : B \subset F(B)\} = \bigcup\{B : B = F(B)\}$.

Finally, we shall also require the fact that for each monotone operator F we may construct a complement operator G in such a way, that the greatest and least fixed points of F are the complements of the least and greatest fixed points of G respectively. This fact, reflected in the next theorem, will be used at a later stage when defining syntactic negation in various modal languages which feature constructs for least and greatest fixed points in such a way that these constructs behave as duals.

Theorem 1.2.3. *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) $I_G = A \setminus J_F$.
- (iv) $J_G = A \setminus I_F$.

Proof. We first need to show that G is monotone. Let $B, C \subset A$ such that $B \subset C$. Then $A \setminus B \supset A \setminus C$ and by monotonicity of F we obtain that $F(A \setminus B) \supset F(A \setminus C)$. Therefore $A \setminus F(A \setminus B) \subset A \setminus F(A \setminus C)$, thus $G(B) \subset G(C)$ and so G is indeed monotone since B and C were arbitrary. We are left to show the statements (i) – (iv).

(i): Assume $F(B) = B$, then

$$G(A \setminus B) = A \setminus F(A \setminus (A \setminus B)) = A \setminus F(B) = A \setminus B,$$

thus the statement holds.

(ii): Assume $G(B) = B$, then $G(B) = A \setminus F(A \setminus B) = B$, therefore we have $F(A \setminus B) = A \setminus B$ and the claim is shown.

(iii): By Theorem 1.2.2 we have $J_F = \bigcup\{C : F(C) = C\}$, so

$$A \setminus J_F = A \setminus \bigcup\{C : F(C) = C\} = \bigcap\{A \setminus C : F(C) = C\}$$

By (i) and (ii) we obtain

$$\bigcap\{A \setminus C : F(C) = C\} = \bigcap\{D : G(D) = D\} = I_G$$

which proves the statement.

(iv): This statement follows by Theorem 1.2.1 and an dual argument to the one given for (iii).

□

Chapter 2

Stratified modal fixed point logic

Stratified modal fixed point logic, abbreviated as SFL, is based on multi-modal logic which is syntactically enriched by a single propositional variable X and its dual $\sim X$. The use of such a variable will enable the construction of arbitrary monotone operators, as long as they are defined using formulae which are syntactically positive in the variable X . To each such formula \mathcal{A} we will assign two constants $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$, the intended meaning of which will be the least and greatest fixed point of the operator defined by \mathcal{A} respectively. Stratification is achieved by allowing \mathcal{A} to contain other constants $P_{\mathcal{B}}$ or $Q_{\mathcal{B}}$, as long as these have been defined at an earlier stage. This mechanism restricts the expressive power of the language to nested but non-interleaving fixed points.

In this chapter, we will first set about defining the language of SFL, taking care that the stratification described above is achieved in a formal way. Then we proceed to establishing the semantics of our language in terms of Kripke structures. After that, we will introduce a way of measuring the complexity of formulae of the language of SFL in such a way that this measure has certain desirable properties.

2.1 The language \mathcal{L}_{SFL}

We define the language of SFL in a level-by-level fashion, starting with level 0, at which we allow only formulae of modal logic which possibly contain X or $\sim X$. For all formulae \mathcal{A} which are positive in X we then add constants $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$. At level 1 we allow formulae of modal logic possibly containing X or $\sim X$ as well as any constant $P_{\mathcal{B}}$ or $Q_{\mathcal{B}}$ from level 0. Again, we add new constants

P_A and Q_A as before and iterate this procedure inductively, obtaining levels 2, 3, 4 and so on. Each formula A of the language is thus 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 reflecting De Morgan's laws, the law of double negation and the law of fixed point duality stated in Theorem 1.2.3.

To facilitate the construction just described we first define the concept of the modal language generated by a set of symbols (and a set of indices for the modalities). Every level of the stratification will then be the modal language over the basic symbols along with the fixed point constants from the level below.

Definition 2.1.1 (Modal language). Let H be a set of symbols and M a set of indices. The modal language over H (with respect to M), denoted by \mathcal{L}_H^M , is the least set containing every element of H which is closed under the following conditions:

1. If A and B are in \mathcal{L}_H^M , then so are $(A \wedge B)$ and $(A \vee B)$.
2. If A is in \mathcal{L}_H^M and i is an index from M , then $\Box_i A$ and $\Diamond_i A$ are also in \mathcal{L}_H^M .

In case there is no danger of confusion, parentheses will be omitted. Furthermore, depending on whether M is a singleton set or not, we will speak of a mono-modal or a multi-modal language. In the mono-modal case we omit the indices when writing the symbols \Box and \Diamond .

Definition 2.1.2 (The language \mathcal{L}_{SFL} , level, length). Let

$$\Phi = \{p, \sim p, q, \sim q, r, \sim r, \dots\}$$

be a countable set of atomic propositions, $V = \{X, \sim X\}$ a set containing one variable and its negation, $T = \{\top, \perp\}$ a set containing symbols for truth and falsehood and M a set of indices.

1. Define $\mathcal{L}_{\text{SFL}}^0$ as the modal language over $\Phi \cup V \cup T$ with respect to M . Given a formula $A \in \mathcal{L}_{\text{SFL}}^0$, we inductively define $\neg A$ by:

$$\begin{aligned} \neg p &:= \sim p, & \neg X &:= \sim X, & \neg \sim p &:= p, & \neg \sim X &:= X, & \neg \top &:= \perp, \\ \neg \perp &:= \top, & \neg(B \wedge C) &:= \neg B \vee \neg C, & \neg(B \vee C) &:= \neg B \wedge \neg C, \\ \neg \Box_i B &:= \Diamond_i \neg B, & \neg \Diamond_i B &:= \Box_i \neg B. \end{aligned}$$

2. A formula $A \in \mathcal{L}_{\text{SFL}}^0$ is called \mathbf{X} -positive if $\sim\mathbf{X}$ does not occur in A . In the following, \mathbf{X} -positive formulae will be denoted by symbols $\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$ (possibly with primes and subscripts). Given formulae \mathcal{A} and B of $\mathcal{L}_{\text{SFL}}^0$ where \mathcal{A} is \mathbf{X} -positive we write $\mathcal{A}[B]$ for the formula which is obtained by replacing every occurrence of \mathbf{X} in \mathcal{A} by B . We define the dual $\overline{\mathcal{A}}$ of an \mathbf{X} -positive formula \mathcal{A} as $\overline{\mathcal{A}} = \neg(\mathcal{A}[\sim\mathbf{X}])$. Furthermore, define the sets

$$L_0 := \{P_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_{\text{SFL}}^0 \text{ } \mathbf{X}\text{-positive}\} \text{ and}$$

$$G_0 := \{Q_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_{\text{SFL}}^0 \text{ } \mathbf{X}\text{-positive}\}.$$

3. Define $\mathcal{L}_{\text{SFL}}^{k+1}$ as the modal language over $\Phi \cup \mathbf{V} \cup \mathbf{T} \cup L_k \cup G_k$ with respect to \mathbf{M} . Again a formula $\mathcal{A} \in \mathcal{L}_{\text{SFL}}^{k+1}$ is called \mathbf{X} -positive if $\sim\mathbf{X}$ does not occur in \mathcal{A} . For any formulae $A \in \mathcal{L}_{\text{SFL}}^{k+1}$ 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 \mathbf{X} -positive formulae $\mathcal{A} \in \mathcal{L}_{\text{SFL}}^{k+1}$ we also define $\overline{\mathcal{A}}$ as before. Similar to the base case we also define sets of fixed point constants

$$L_{k+1} := \{P_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_{\text{SFL}}^k \text{ } \mathbf{X}\text{-positive}\} \text{ and}$$

$$G_{k+1} := \{Q_{\mathcal{A}} : \mathcal{A} \in \mathcal{L}_{\text{SFL}}^k \text{ } \mathbf{X}\text{-positive}\}.$$

4. Define the language \mathcal{L}_{SFL} by setting $\mathcal{L}_{\text{SFL}} = \bigcup_{k \in \omega} \mathcal{L}_{\text{SFL}}^k$.
5. For every formula $A \in \mathcal{L}_{\text{SFL}}$ define $level(A)$ to be the least k so that $A \in \mathcal{L}_{\text{SFL}}^k$.
6. The length $|A|$ of an $A \in \mathcal{L}_{\text{SFL}}$ 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.

Implication and equivalence for formulae of \mathcal{L}_{SFL} can be defined as abbreviations in the standard way. Furthermore, we shall also use a notation for formulae which are prefixed by a finite number of boxes or diamonds of a given index.

Definition 2.1.3 (Abbreviations). Given formulae A and B of \mathcal{L}_{SFL} and an index i from \mathbf{M} we define the following syntactic abbreviations:

$$A \rightarrow B := \neg A \vee B,$$

$$A \leftrightarrow B := (A \rightarrow B) \wedge (B \rightarrow A),$$

$$\square_i^1 A := \square_i A,$$

$$\square_i^{k+1} A := \square_i \square_i^k A,$$

$$\diamond_i^1 A := \diamond_i A,$$

$$\diamond_i^{k+1} A := \diamond_i \diamond_i^k A.$$

2.2 The semantics of \mathcal{L}_{SFL}

Formulae of the language \mathcal{L}_{SFL} are interpreted as standing for sets of worlds in Kripke structures. A Kripke structure is a tuple consisting of a set of possible worlds of the structure, a family of binary relations stating which worlds are accessible from which other ones and a valuation function which assigns a set of fulfilling worlds to each primitive proposition and also to the symbols \mathbf{X} and $\sim\mathbf{X}$. We now define this concept more formally and also introduce a notation for explicitly stating that the valuation function of a Kripke structure assigns a certain set of worlds to the variable \mathbf{X} . This will be convenient when speaking of operators determined by \mathbf{X} -positive formulae.

Definition 2.2.1 (Kripke structure). A Kripke structure for \mathcal{L}_{SFL} is a triple $\mathbf{K} = (S, R, \pi)$, where S is a non-empty set, $R : \mathbf{M} \rightarrow \mathcal{P}(S \times S)$ and $\pi : (\Phi \cup \mathbf{V}) \rightarrow \mathcal{P}(S)$ is a function such that $\pi(\sim\mathbf{X}) = S \setminus \pi(\mathbf{X})$ and for all $\mathbf{p} \in \Phi$ we have $\pi(\sim\mathbf{p}) = S \setminus \pi(\mathbf{p})$. The function R assigns an accessibility relation to each $i \in \mathbf{M}$ where we write R_i for the relation $R(i)$. In case \mathbf{M} is a singleton set (and thus \mathcal{L}_{SFL} is based on a mono-modal language), we treat R as a single relation $R \subset S \times S$. Furthermore, given a set $T \subset S$ we define the Kripke structure $\mathbf{K}[\mathbf{X} := T]$ as the triple (S, R, π') , where $\pi'(\mathbf{X}) = T$, $\pi'(\sim\mathbf{X}) = S \setminus T$ and $\pi'(P) = \pi(P)$ for all $P \in \Phi$.

Given a Kripke structure \mathbf{K} the denotations of the formulae of \mathcal{L}_{SFL} in \mathbf{K} are defined inductively on the levels of the language. Primitive propositions and the symbols \mathbf{X} and $\sim\mathbf{X}$ are treated by the valuation function and the boolean and modal constructs are interpreted as usual. The meaning of constants $P_{\mathcal{A}}$ and $Q_{\mathcal{A}}$ is defined to reflect the least and greatest fixed points of the monotone operator associated with the formula $\mathcal{A}[\mathbf{X}]$. This last fact is established in Theorem 2.2.3 which follows the definition.

Definition 2.2.2 (Denotation). Let $\mathbf{K} = (S, R, \pi)$ be a Kripke structure. For every $A \in \mathcal{L}_{\text{SFL}}^k$ we define the set $\|A\|_{\mathbf{K}} \subset S$ inductively as follows:

$$\begin{aligned} \|P\|_{\mathbf{K}} &:= \pi(P) \text{ for all } P \in \Phi \cup \mathbf{V}, & \|\top\|_{\mathbf{K}} &:= S, & \|\perp\|_{\mathbf{K}} &:= \emptyset, \\ \|B \wedge C\|_{\mathbf{K}} &:= \|B\|_{\mathbf{K}} \cap \|C\|_{\mathbf{K}}, & \|B \vee C\|_{\mathbf{K}} &:= \|B\|_{\mathbf{K}} \cup \|C\|_{\mathbf{K}}, \\ \|\Box_i B\|_{\mathbf{K}} &:= \{w \in S : v \in \|B\|_{\mathbf{K}} \text{ for all } v \text{ such that } wR_i v\}, \\ \|\Diamond_i B\|_{\mathbf{K}} &:= \{w \in S : v \in \|B\|_{\mathbf{K}} \text{ for some } v \text{ such that } wR_i v\}. \end{aligned}$$

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

$$\begin{aligned} \|P_{\mathcal{A}}\|_{\mathbf{K}} &:= \bigcap \{T \subset S : F_{\mathcal{A}}^{\mathbf{K}}(T) \subset T\} \text{ and} \\ \|Q_{\mathcal{A}}\|_{\mathbf{K}} &:= \bigcup \{T \subset S : T \subset F_{\mathcal{A}}^{\mathbf{K}}(T)\} \end{aligned}$$

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

Theorem 2.2.3. $F_{\mathcal{A}}^{\mathbf{K}}$ is monotone and $\|P_{\mathcal{A}}\|_{\mathbf{K}}$ and $\|Q_{\mathcal{A}}\|_{\mathbf{K}}$ are the least and greatest fixed points of $F_{\mathcal{A}}^{\mathbf{K}}$ respectively. That is if $\mathbf{K} = (S, R, \pi)$ is a Kripke structure, then

$$(i) \quad U \subset T \implies F_{\mathcal{A}}^{\mathbf{K}}(U) \subset F_{\mathcal{A}}^{\mathbf{K}}(T) \text{ for all } U, T \subset S,$$

$$(ii) \quad \|\mathcal{A}[P_{\mathcal{A}}]\|_{\mathbf{K}} = \|P_{\mathcal{A}}\|_{\mathbf{K}},$$

$$(iii) \quad F_{\mathcal{A}}^{\mathbf{K}}(T) = T \implies \|P_{\mathcal{A}}\|_{\mathbf{K}} \subset T \text{ for all } T \subset S,$$

$$(iv) \quad \|\mathcal{A}[Q_{\mathcal{A}}]\|_{\mathbf{K}} = \|Q_{\mathcal{A}}\|_{\mathbf{K}},$$

$$(v) \quad F_{\mathcal{A}}^{\mathbf{K}}(T) = T \implies T \subset \|Q_{\mathcal{A}}\|_{\mathbf{K}} \text{ for all } T \subset S.$$

Statement (i) is shown by a straightforward induction on the structure of \mathcal{A} . Statements (ii) and (iii) are consequences of (i) and Theorem 1.2.1, (iv) and (v) analogously follow from (i) and Theorem 1.2.2. Using the concept of the denotation $\|A\|_{\mathbf{K}}$ of a formula A in a Kripke structure \mathbf{K} the customary notion of satisfaction and satisfiability of a formula is defined. However, since our goal is to present a complete deduction system for stratified modal fixed point logic the dual notion of validity is in a sense more central in our context.

Definition 2.2.4 (Satisfaction and validity). Let $\mathbf{K} = (S, R, \pi)$ be a Kripke structure. We say a formula $A \in \mathcal{L}_{\text{SFL}}$ is satisfied in \mathbf{K} if $\|A\|_{\mathbf{K}} \neq \emptyset$ and valid in \mathbf{K} if $\|A\|_{\mathbf{K}} = S$. We say A is satisfiable if there exists a Kripke structure in which A is satisfied. Furthermore, we say that A is valid if it is valid in all Kripke structures.

As seen in Theorem 1.2.2, if a world is an element of a greatest fixed point, it must be an element of every iteration from above of the corresponding operator. It shall be established later that indeed when dealing with valid formulae finite such iterations suffice. We thus introduce a syntactic way of representing finite iterations from above as formulae of \mathcal{L}_{SFL} . For reasons of duality we also introduce a notation for finite iterations from below. On the other hand, a convenient semantic notation for arbitrary transfinite iterations both from below and above is also defined.

Definition 2.2.5 (Iterations). Let \mathcal{A} be an X -positive formula.

1. For every $k \in \omega$ define the formulae $Q_{\mathcal{A}}^k$ and $P_{\mathcal{A}}^k$ inductively as follows:

$$P_{\mathcal{A}}^0 := \perp, \quad P_{\mathcal{A}}^{k+1} := \mathcal{A}[P_{\mathcal{A}}^k], \quad Q_{\mathcal{A}}^0 := \top, \quad Q_{\mathcal{A}}^{k+1} := \mathcal{A}[Q_{\mathcal{A}}^k].$$

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

$$I_{\mathcal{A},\mathbf{K}}^{<\alpha} := I_{F_{\mathcal{A}}^{\mathbf{K}}}^{<\alpha}, \quad I_{\mathcal{A},\mathbf{K}}^\alpha := I_{F_{\mathcal{A}}^{\mathbf{K}}}^\alpha, \quad J_{\mathcal{A},\mathbf{K}}^{<\alpha} := J_{F_{\mathcal{A}}^{\mathbf{K}}}^{<\alpha}, \quad J_{\mathcal{A},\mathbf{K}}^\alpha := J_{F_{\mathcal{A}}^{\mathbf{K}}}^\alpha.$$

Remark 2.2.6. Two facts are evident from Definition 2.2.5 and will be used several times in the subsequent argument: Let \mathcal{A} be a formula of \mathcal{L}_{SFL} which is \mathbf{X} -positive. Then for all natural numbers k we have $\text{level}(Q_{\mathcal{A}}^k) < \text{level}(Q_{\mathcal{A}})$ and $\text{level}(P_{\mathcal{A}}^k) < \text{level}(P_{\mathcal{A}})$ and, furthermore, for all Kripke structures \mathbf{K} we have $\|P_{\mathcal{A}}^k\|_{\mathbf{K}} = I_{\mathcal{A},\mathbf{K}}^{<k}$ and $\|Q_{\mathcal{A}}^k\|_{\mathbf{K}} = J_{\mathcal{A},\mathbf{K}}^{<k}$.

2.3 Formula complexity

The language \mathcal{L}_{SFL} is built up in layers and at each layer a set of constants for certain formulae of the layer below is added. In view of the interpretation of these constants as fixed points, we will require a measure of formula complexity under which a constant at a higher level is strictly more complex than any formula in a level below. In particular, under this measure a greatest fixed point constant should always be more complex than all of its finite approximations.

Definition 2.3.1 (Rank). The rank $rk(A)$ of a formula $A \in \mathcal{L}_{\text{SFL}}$ is an ordinal defined inductively as follows:

1. If A is an element of $\Phi \cup \mathbf{V} \cup \mathbf{T}$, then $rk(A) := 0$.
2. If A is a fixed point constant $P_{\mathcal{A}}$ or $Q_{\mathcal{A}}$ and $\text{level}(A) = n$, then

$$rk(A) := \omega n.$$

3. If A is a formula $B \wedge C$ or $B \vee C$, then $rk(A) := \max(rk(B), rk(C)) + 1$.
4. If A is a formula $\Box_i B$ or $\Diamond_i B$ for some $i \in \mathbf{M}$, then $rk(A) := rk(B) + 1$.

The next two lemmata summarise the important properties required of the rank function. Since all proofs are routine, they will only be sketched.

Lemma 2.3.2. *For all formulae $A, B \in \mathcal{L}_{\text{SFL}}$ the following statements hold:*

- (i) if $\text{level}(A) = n$, then $\omega n \leq rk(A) < \omega n + \omega$.
- (ii) $rk(A) = rk(\neg A)$

(iii) $rk(A), rk(B) < rk(A \wedge B), rk(A \vee B)$

(iv) $rk(A) < rk(\Box_i A), rk(\Diamond_i A)$ for all $i \in M$.

Claims (i) and (ii) are shown by induction on $level(A)$ and the structure of A , noting for the fixed point cases that $level(\mathcal{A}) = level(\bar{\mathcal{A}})$ for all \mathbf{X} -positive formulae \mathcal{A} of \mathcal{L}_{SFL} . Claims (iii) and (iv) follow directly from Definition 2.3.1.

Lemma 2.3.3. *For all fixed point constants $Q_{\mathcal{A}}$ and all $k \in \omega$ we have*

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

The right hand inequality is a straightforward consequence of Remark 2.2.6 and claim (i) of Lemma 2.3.2, the left hand one is shown by induction on k . Due to Lemmata 2.3.2 and 2.3.3 we may carry out a good number of our subsequent definitions and proofs by induction on $rk(A)$. The first claim we can show in this manner states that our defined negation corresponds to complementation on the semantic level and thus behaves in the expected way.

Lemma 2.3.4. *For all $A \in \mathcal{L}_{\text{SFL}}$ and all Kripke structures $\mathcal{K} = (S, R, \pi)$ we have $\|\neg A\|_{\mathcal{K}} = S \setminus \|A\|_{\mathcal{K}}$.*

Chapter 3

Some notable fragments of SFL

In this chapter we introduce the Hilbert system H_{SFL} for stratified modal fixed point logic. We mention H_{SFL} because it is simple and easily accessible from an intuitive point of view. The system is basically Kozen's axiomatisation of the modal μ -calculus [23] adapted to our more restrictive setting. We then use the system H_{SFL} to gain some intuition about the expressive strength of SFL. To this end, we will discuss three notable fragments of SFL, namely Logic of Common Knowledge, CTL and PDL and their embeddings into H_{SFL} . Later, in Chapters 7 and 8, we will meet a prominent logic which in turn strictly contains SFL as a fragment, namely the modal μ -calculus.

3.1 The Hilbert system H_{SFL}

Before introducing H_{SFL} , we define the Hilbert style axiomatisation $H_{\mathbf{K}}^{\mathcal{L}}$ for standard modal logic \mathbf{K} over an arbitrary language \mathcal{L} closed under \Box_i and \Diamond_i . When defining Hilbert systems in general we proceed by providing axioms on the one hand and deduction rules on the other. An axiom represents a formula which may be used anywhere in a proof. A deduction rule expresses the fact that if all the premises have been deduced, the conclusion may also be deduced. Both axioms and deduction rules are given names, which are displayed in brackets to the right of them. Viewed strictly, these names would need to be parametrised by the formulae involved in the respective axiom or rule. However, for brevity we ignore such parameters since they are always clear from the context. Informally speaking (and an informal notion being sufficient for our purposes), a proof of a formula A in a Hilbert system is a sequence of formulae ending with A such that each element of this sequence is either an axiom or follows from previous elements by the application of a rule.

Definition 3.1.1 (The system $H_K^{\mathcal{L}}$). Assume M is a set of indices and \mathcal{L} an arbitrary language closed under \Box_i and \Diamond_i for all $i \in M$. The system $H_K^{\mathcal{L}}$ for modal logic (over \mathcal{L}) is defined by the following axioms and rules:

Logical axioms: For all propositional tautologies A of \mathcal{L} , all formulae B and C of \mathcal{L} and all indices i from M

$$A \quad (\text{TAUT}), \quad (\Box_i B \wedge \Box_i (B \rightarrow C)) \rightarrow \Box_i C \quad (\text{K}).$$

Logical rules: For all formulae B, C of \mathcal{L} and indices i from M

$$\frac{B \quad B \rightarrow C}{C} \quad (\text{MP}), \quad \frac{B}{\Box_i B} \quad (\text{NEC}).$$

The system H_{SFL} consists of $H_K^{\mathcal{L}_{\text{SFL}}}$ plus one axiom and one rule for the fixed point constants. Given an X -positive formula \mathcal{A} the additional axiom expresses the fact that $P_{\mathcal{A}}$ stands for a fixed point and the additional rule states that this fixed point is 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 3.1.2 (The system H_{SFL}). The system H_{SFL} is defined by adding the following axioms and rules to $H_K^{\mathcal{L}_{\text{SFL}}}$:

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

$$\mathcal{A}[P_{\mathcal{A}}] \rightarrow P_{\mathcal{A}} \quad (\text{CLO}).$$

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

$$\frac{\mathcal{A}[B] \rightarrow B}{P_{\mathcal{A}} \rightarrow B} \quad (\text{IND}).$$

The axioms and rules concerning the greatest fixed points are derivable in H_{SFL} . To see this we first need to show a technical lemma stating that variable substitution commutes with dualisation of X -positive formulae in the expected way.

Lemma 3.1.3. *For any formulae \mathcal{A} and B of \mathcal{L}_{SFL} where \mathcal{A} is X -positive we have*

$$\neg(\mathcal{A}[B]) = \overline{\mathcal{A}}[\neg B].$$

Proof. This claim is shown by a straightforward induction on the structure of \mathcal{A} . \square

Lemma 3.1.4. *The system H_{SFL} derives the formula $Q_{\mathcal{A}} \rightarrow \mathcal{A}[Q_{\mathcal{A}}]$ and the rule*

$$\frac{B \rightarrow \mathcal{A}[B]}{B \rightarrow Q_{\mathcal{A}}}.$$

Proof. For the first claim consider the contraposition of the closure axiom

$$\neg P_{\mathcal{B}} \rightarrow \neg \mathcal{B}[P_{\mathcal{B}}].$$

By Lemma 3.1.3 we obtain $Q_{\overline{\mathcal{B}}} \rightarrow \overline{\mathcal{B}}[Q_{\overline{\mathcal{B}}}]$. Now for any formula \mathcal{A} we always find a suitable formula \mathcal{B} such that $\overline{\mathcal{B}} = \mathcal{A}$, thus the claim is shown. For the second claim assume the implication $B \rightarrow \mathcal{A}[B]$ or rather its contraposition $\neg \mathcal{A}[B] \rightarrow \neg B$. Again using Lemma 3.1.3 we obtain $\overline{\mathcal{A}}[\neg B] \rightarrow \neg B$, so by applying the induction rule we arrive at $P_{\overline{\mathcal{A}}} \rightarrow \neg B$ and ultimately its contraposition $B \rightarrow Q_{\mathcal{A}}$ which proves the second claim. \square

3.2 Logic of common knowledge

Logic of Common Knowledge [14] 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 then 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 is formally speaking 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, we define the language of the Logic of Common Knowledge as a modal language with the addition of an operator C and its dual $\tilde{\mathsf{C}}$.

Definition 3.2.1 (Language \mathcal{L}_{C}). Let M be a finite set of indices. The language \mathcal{L}_{C} of Logic of Common Knowledge is defined as the least set containing every element of $\Phi \cup \mathsf{T}$ which is closed under the following conditions:

1. If A and B are in \mathcal{L}_{C} , then so are $(A \wedge B)$ and $(A \vee B)$.
2. If A is in \mathcal{L}_{C} and i is an index from M , then $\Box_i A$ and $\Diamond_i A$ are also in \mathcal{L}_{C} .
3. If A is in \mathcal{L}_{C} , then $\mathsf{C}A$ and $\tilde{\mathsf{C}}A$ are also in \mathcal{L}_{C} .

Furthermore, we define the abbreviation $EA := \bigwedge_{i \in \mathbf{M}} \Box_i A$ as well as its dual $\tilde{E}A := \bigvee_{i \in \mathbf{M}} \Diamond_i A$ and inductively set $E^0 A := \top$ and $E^{k+1} A := EE^k A$ as well as $\tilde{E}^0 A := \perp$ and $\tilde{E}^{k+1} A := \tilde{E}\tilde{E}^k A$.

The semantics of \mathcal{L}_C is designed to reflect the intuition that a formula of the form CA stands for an infinite conjunction over all finite iterations $E^k A$ of the “everybody knows”-operator.

Definition 3.2.2 (Semantics of \mathcal{L}_C). Let $\mathbf{K} = (S, R, \pi)$ be a (multi-modal) Kripke structure and A a formula of \mathcal{L}_C . We inductively define the denotation $\|A\|_{\mathbf{K}}$ of A in \mathbf{K} as usual if A is an atomic, propositional or modal formula and add the following two clauses for the common knowledge operator and its dual:

$$\begin{aligned} \|CA\|_{\mathbf{K}} &:= \bigcap_{k \in \omega} \|E^k A\|_{\mathbf{K}} \\ \|\tilde{C}A\|_{\mathbf{K}} &:= \bigcup_{k \in \omega} \|\tilde{E}^k A\|_{\mathbf{K}} \end{aligned}$$

The Logic of Common Knowledge can be axiomatised as $\mathbf{H}_{\mathbf{K}}^{\mathcal{L}_C}$ plus the following axioms and rules for the common knowledge operator C :

Closure axioms: For every formula A

$$CA \rightarrow (EA \wedge ECA).$$

Induction rules: For every formula A and B

$$\frac{B \rightarrow (EA \wedge EB)}{B \rightarrow CA}.$$

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 CA as $Q_{E(A^* \wedge X)}$ where A^* denotes the translation of A . By Lemma 3.1.4 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 \mathbf{H}_{SFL} . Since the rest of the axioms and rules of Logic of Common Knowledge are also a part of \mathbf{H}_{SFL} , this already takes care of our embedding result.

Theorem 3.2.3. *For all formulae A and B of Logic of Common Knowledge we have:*

1. \mathbf{H}_{SFL} derives $(CA \rightarrow (EA \wedge ECA))^*$.
2. If \mathbf{H}_{SFL} derives $(B \rightarrow (EA \wedge EB))^*$, then \mathbf{H}_{SFL} also derives $(B \rightarrow CA)^*$.

3.3 Computational tree logic

The next fragment of SFL we consider is Computational Tree Logic or CTL for short [11]. 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 behaviour which is in a sense unbounded like “in all runs extending from the current state A holds until B is the case”, denoted by $\forall(AUB)$ or “in some runs extending from the current state A holds until B is the case”, written as $\exists(AUB)$. The language \mathcal{L}_{CTL} of CTL is defined formally by adding the constructs $\forall(AUB)$ and $\exists(AUB)$ to a modal language.

Definition 3.3.1 (Language \mathcal{L}_{CTL}). The language \mathcal{L}_{CTL} of CTL is defined as the least set containing every element of $\Phi \cup \mathbb{T}$ which is closed under the following conditions:

1. If A and B are in \mathcal{L}_{CTL} , then so are $(A \wedge B)$ and $(A \vee B)$.
2. If A is in \mathcal{L}_{CTL} and i is an index from \mathbb{M} , then $\Box A$ and $\Diamond A$ are also in \mathcal{L}_{CTL} .
3. If A and B are in \mathcal{L}_{CTL} , then $\forall(AUB)$ and $\exists(AUB)$ are also in \mathcal{L}_{CTL} .

Interpreting the worlds in a Kripke structure as states of a system and the edges as transitions between such states, it makes sense to introduce the concept of a run of the system. A run is to be understood as one possible sequence of states which the system may assume as it is executed.

Definition 3.3.2 (Run). Let $\mathbf{K} = (S, R, \pi)$ be a (mono-modal) Kripke structure. For every world w in S an infinite sequence (w_0, w_1, w_2, \dots) of elements of S is called a run of \mathbf{K} extending from w if $w_0 = w$ and for every natural number i we have $w_i R w_{i+1}$ or w_i is a leaf with respect to R and $w_j = w_i$ for all natural numbers $j > i$.

Using the concept of a run we now assign a precise meaning to formulae of \mathcal{L}_{CTL} . As was to be expected, the construct $\forall(AUB)$ quantifies over all runs extending from a certain state while $\exists(AUB)$ corresponds to existential quantification.

Definition 3.3.3 (Semantics of \mathcal{L}_{CTL}). Let $\mathbf{K} = (S, R, \pi)$ be a (mono-modal) Kripke structure and A and B formulae of \mathcal{L}_{C} . We inductively define the denotation $\|A\|_{\mathbf{K}}$ of A in \mathbf{K} as usual if A is an atomic, propositional or

modal formula and add the following two clauses for the operators $\forall(AUB)$ and $\exists(AUB)$:

$$\begin{aligned} \|\forall(AUB)\|_{\mathbf{K}} &:= \{w \in S : \text{for every run } (w_0, w_1, w_2, \dots) \text{ of } \mathbf{K} \text{ extending from} \\ &\quad w \text{ there exists } i \text{ such that } w_i \in \|B\|_{\mathbf{K}} \text{ and} \\ &\quad w_j \in \|A\|_{\mathbf{K}} \text{ for all } j < i\} \\ \|\exists(AUB)\|_{\mathbf{K}} &:= \{w \in S : \text{there exists a run } (w_0, w_1, w_2, \dots) \text{ of } \mathbf{K} \text{ extending} \\ &\quad \text{from } w \text{ and an } i \text{ such that } w_i \in \|B\|_{\mathbf{K}} \text{ and} \\ &\quad w_j \in \|A\|_{\mathbf{K}} \text{ for all } j < i\} \end{aligned}$$

Axiomatically, we obtain CTL by extending $\mathbf{H}_{\mathbf{K}}^{\mathcal{L}^{\text{CTL}}}$ by the following axioms and rules governing the use of the $\forall(AUB)$ and $\exists(AUB)$ constructs:

Closure axioms: For every formula A and B

- (1) $B \vee (A \wedge \Box \forall(AUB)) \rightarrow \forall(AUB)$
- (2) $B \vee (A \wedge \Diamond \exists(AUB)) \rightarrow \exists(AUB)$

Induction rules: For all formulae A , B and C

- (3)
$$\frac{C \rightarrow (\neg B \wedge \Diamond C)}{C \rightarrow \neg \forall(AUB)}$$
- (4)
$$\frac{C \rightarrow (\neg B \wedge (A \rightarrow \Box C))}{C \rightarrow \neg \exists(AUB)}$$

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 $\forall(AUB)$ and $\exists(AUB)$. More precisely, we translate a formula of the form $\forall(AUB)$ into $P_{B^* \vee (A^* \wedge \Box X \wedge \Diamond \top)}$ and a formula of the form $\exists(AUB)$ into $P_{B^* \vee (A^* \wedge \Diamond X)}$ where in both cases A^* and B^* stand for the translations of A and B respectively. Using this translation and Definition 3.1.2, we again immediately get the following theorem which ensures that CTL is embeddable into \mathbf{H}_{SFL} .

Theorem 3.3.4. *For all formulae A , B and C of CTL we have:*

1. \mathbf{H}_{SFL} derives $[B \vee (A \wedge \Box \forall(AUB)) \rightarrow \forall(AUB)]^*$.
2. \mathbf{H}_{SFL} derives $[B \vee (A \wedge \Diamond \exists(AUB)) \rightarrow \exists(AUB)]^*$.
3. If \mathbf{H}_{SFL} derives $[C \rightarrow (\neg B \wedge \Diamond C)]^*$, then also $[C \rightarrow \neg \forall(AUB)]^*$.
4. If \mathbf{H}_{SFL} derives $[C \rightarrow (\neg B \wedge (A \rightarrow \Box C))]^*$, then also $[C \rightarrow \neg \exists(AUB)]^*$.

3.4 Propositional dynamic logic

The last fragment of SFL we mention is Propositional Dynamic Logic, abbreviated as PDL [15, 16]. 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 infinitary 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 .

As a matter of fact, the language we have just described informally does not correspond fully to that of standard PDL as found in the literature. In more customary accounts, the set M is provided with additional structure by closing under the program operators $;$ (for composition), \cup (for nondeterministic choice) and $*$ (for iteration) and furthermore by adding a mixed operator $?$ for testing whether a formula holds at a given world. Since we are concerned with fixed point extensions to modal logic, we will however consider only the $*$ -fragment of full PDL and, by a slight abuse of terminology, refer to this fragment merely as PDL. Accordingly, we define the language \mathcal{L}_{PDL} of PDL as a modal language, adding clauses for formulae of the form $\Box_i^* A$ and $\Diamond_i^* A$.

Definition 3.4.1 (Language \mathcal{L}_{PDL}). Let M be an infinite set of indices. The language \mathcal{L}_{PDL} of PDL is defined as the least set containing every element of $\Phi \cup \mathsf{T}$ which is closed under the following conditions:

1. If A and B are in \mathcal{L}_{PDL} , then so are $(A \wedge B)$ and $(A \vee B)$.
2. If A is in \mathcal{L}_{PDL} and i is an index from M , then $\Box_i A$ and $\Diamond_i A$ are also in \mathcal{L}_{PDL} .
3. If A is in \mathcal{L}_{PDL} , then $\Box_i^* A$ and $\Diamond_i^* A$ are also in \mathcal{L}_{PDL} .

The semantics of the language \mathcal{L}_{PDL} is set up in such a way that the \Box_i^* construct acts a quantification over all finite iterations of action i , whereas \Diamond_i^* amounts to an existential quantification.

Definition 3.4.2 (Semantics of \mathcal{L}_{PDL}). Let $K = (S, R, \pi)$ be a (multi-modal) Kripke structure and A a formula of \mathcal{L}_{PDL} . We inductively define the denotation $\|A\|_K$ of A in K as usual if A is an atomic, propositional or modal

formula and, for every i in \mathbf{M} , add the following two clauses for the universal and existential iteration operators:

$$\begin{aligned}\|\Box_i^* A\|_{\mathbf{K}} &:= \bigcap_{k \in \omega} \|\Box_i^k A\|_{\mathbf{K}} \\ \|\Diamond_i^* A\|_{\mathbf{K}} &:= \bigcup_{k \in \omega} \|\Diamond_i^k A\|_{\mathbf{K}}\end{aligned}$$

Again we may axiomatise PDL by taking $\mathbf{H}_{\mathbf{K}}^{\mathcal{L}^{\text{PDL}}}$ and extending it by the following axioms and rules for the \Box_i^* and \Diamond_i^* operators:

Closure axioms: For every formula A

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

Induction rules: For all formulae A, B

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

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 $Q_{A^* \wedge \Box_i X}$ and $\Diamond_i^* A$ as $P_{A^* \vee \Diamond_i X}$ where A^* stands for the translation of A . Using this translation, Definition 3.1.2 and Lemma 3.1.4 we prove the following embedding theorem.

Theorem 3.4.3. *For all formulae A and B of PDL we have:*

1. \mathbf{H}_{SFL} derives $[\Box_i^* A \rightarrow (A \wedge \Box_i \Box_i^* A)]^*$.
2. \mathbf{H}_{SFL} derives $[(A \vee \Diamond_i \Diamond_i^* A) \rightarrow \Diamond_i^* A]^*$.
3. If \mathbf{H}_{SFL} derives $[B \rightarrow (A \wedge \Box_i B)]^*$, then \mathbf{H}_{SFL} derives $[B \rightarrow \Box_i^* A]^*$.
4. If \mathbf{H}_{SFL} derives $[(A \vee \Diamond_i B) \rightarrow B]^*$, then \mathbf{H}_{SFL} derives $[\Diamond_i^* A \rightarrow B]^*$.

Chapter 4

The infinitary system T_{SFL}^ω

In this chapter we introduce the syntactic calculus T_{SFL}^ω for deriving valid formulae of stratified modal fixed point logic. In order to focus on more important aspects of the argument, we shall restrict the remainder of our study of **SFL** to the mono-modal case, that is to say the case where M is a singleton set. A generalisation of the subsequent arguments to the full multi-modal case is an easy exercise and, furthermore, a fully general version of these arguments is carried out in Chapter 8 for an extension of **SFL**. Thus the generality of our account is not affected in an essential way by the above mentioned restriction.

The first section of this chapter will introduce the cut-free infinitary Tait style deductive system T_{SFL}^ω . In the second section we will introduce the central technical notion of saturated sequents on which the completeness proof for T_{SFL}^ω will be based. While soundness is postponed to a later chapter, the third section will address the completeness proof itself. In the fourth section we will investigate an important relationship between the Hilbert style system H_{SFL} introduced in Chapter 3 and the new system T_{SFL}^ω by studying in particular how the induction rule (IND) can be derived in T_{SFL}^ω using an added cut rule.

4.1 The system T_{SFL}^ω

The calculus T_{SFL}^ω is designed in Tait-style, that is to say, it can be used to derive finite sets of formulae of \mathcal{L}_{SFL} . It is infinitary in the sense that in order to apply the rule for a greatest fixed point constant Q_A we must derive infinitely many premises. As a consequence of this, proofs in T_{SFL}^ω can become infinite in depth. The infinitary nature of the greatest fixed point rule is in accord with the fact that in order for a greatest fixed point formula to be satisfied, all of its approximations must be satisfied. However, since

stratified modal fixed point logic enjoys the so-called finite model property, in order to assert the validity of a greatest fixed point it will turn out to be sufficient to check the validity of all its finite approximations only. This is indeed slightly surprising in view of the fact that (as we shall also see in Chapter 6) there are fixed point formulae of $\mathcal{L}_{\mathsf{SFL}}$ which, on certain Kripke structures, do not close at any finite approximation. We can resolve this apparent contradiction by noting that the rule for greatest fixed points which we are about to introduce preserves validity but not necessarily satisfaction in one particular Kripke structure. As one of its main features we note that our calculus will not include a cut rule although such a rule can be added for convenience. That is to say we have soundness for the calculus including the cut rule and completeness for its cut free fragment. This amounts to a semantic cut elimination result.

As mentioned above we want to introduce a calculus for deriving finite sets of formulae. For this purpose we first introduce some convenient notation for dealing with such sets. In this context a finite set of formulae is always read as the disjunction of all its elements.

Definition 4.1.1 (Sequents). A sequent (of $\mathcal{L}_{\mathsf{SFL}}$) is a finite set of formulae of $\mathcal{L}_{\mathsf{SFL}}$. 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 \perp 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 $\mathsf{T}_{\mathsf{SFL}}^\omega$. 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 $\mathsf{T}_{\mathsf{SFL}}^\omega$. As in the case of $\mathsf{H}_{\mathsf{SFL}}$ a name is indicated in brackets to the right of each rule and again parametrisation of these names with the formulae and sequents involved is omitted.

Definition 4.1.2 (The system $\mathsf{T}_{\mathsf{SFL}}^\omega$). The system $\mathsf{T}_{\mathsf{SFL}}^\omega$ is defined by the following inference rules:

Axioms: For all sequents Γ of $\mathcal{L}_{\mathsf{SFL}}$ and \mathfrak{p} in Φ

$$\frac{}{\Gamma, \mathfrak{p}, \sim \mathfrak{p}} \text{ (ID1)}, \quad \frac{}{\Gamma, \mathsf{X}, \sim \mathsf{X}} \text{ (ID2)}, \quad \frac{}{\Gamma, \top} \text{ (ID3)}.$$

Propositional rules: For all sequents Γ and formulae A and B of \mathcal{L}_{SFL}

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

Modal rules: For all sequents Γ and Σ and formulae A of \mathcal{L}_{SFL}

$$\frac{\Gamma, A}{\diamond\Gamma, \square A, \Sigma} \quad (\square).$$

Fixed point rules: For all sequents Γ and all X -positive formulae \mathcal{A} of \mathcal{L}_{SFL}

$$\frac{\Gamma, \mathcal{A}[P_{\mathcal{A}}]}{\Gamma, P_{\mathcal{A}}} \quad (P), \quad \frac{\Gamma, Q_{\mathcal{A}}^k \quad \text{for all } k \in \omega}{\Gamma, Q_{\mathcal{A}}} \quad (Q^\omega).$$

We may also add the following cut rule for every sequent Γ and formula A of \mathcal{L}_{SFL}

$$\frac{\Gamma, A \quad \Gamma, \neg A}{\Gamma} \quad (\text{cut})$$

in which case the resulting system will be referred to as $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut})$.

We may now state formally what it means for a sequent Γ to be provable in one of the systems just introduced. In particular, we are also interested in measuring the length of a possible proof of Γ . Since proofs in $\mathsf{T}_{\text{SFL}}^\omega$ and $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut})$ may be infinite in length, but are wellfounded objects, ordinals turn out to be the ideal tool for such a measurement.

Definition 4.1.3 (Provability). Assume Γ is a sequent of \mathcal{L}_{SFL} and α an ordinal. We define the provability of Γ in $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut})$ in α many steps, denoted by $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^\alpha \Gamma$, by induction as follows:

1. If Γ is obtained by one of the axioms of $\mathsf{T}_{\text{SFL}}^\omega$, then $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^\beta \Gamma$ holds for all ordinals β .
2. If Γ is obtained by one of the propositional, modal, fixed point or cut rules where Γ_i are the premises of the respective rule, $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^{\beta_i} \Gamma_i$ holds for all of these premises and β is an ordinal such that $\beta_i < \beta$ for all β_i , then $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^\beta \Gamma$.

The notion of provability of Γ in $\mathsf{T}_{\text{SFL}}^\omega$ (without cut) in α many steps, denoted by $\mathsf{T}_{\text{SFL}}^\omega \vdash^\alpha \Gamma$, is obtained by not including the cut rule in the above induction. Furthermore, we say Γ is provable and write $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash \Gamma$ or $\mathsf{T}_{\text{SFL}}^\omega \vdash \Gamma$ if there exists an ordinal β such that Γ is provable in the respective system in β many steps. Finally, we write $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \Gamma$ or $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \not\vdash \Gamma$ if Γ is not provable in the respective system.

With respect to provability, both $\mathsf{T}_{\text{SFL}}^\omega$ and $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut})$ have the important property of weakening. That is, whenever a sequent Γ is provable, any sequent which extends Γ is also provable with the same length. This can be shown by a straightforward induction on the length of the proof of Γ .

Lemma 4.1.4 (Weakening). *For all sequents Γ and Δ of \mathcal{L}_{SFL} and all ordinals β we have*

1. *If $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^\beta \Gamma$ and $\Gamma \subset \Delta$, then $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash^\beta \Delta$.*
2. *If $\mathsf{T}_{\text{SFL}}^\omega \vdash^\beta \Gamma$ and $\Gamma \subset \Delta$, then $\mathsf{T}_{\text{SFL}}^\omega \vdash^\beta \Delta$.*

It is not immediately obvious that $\mathsf{T}_{\text{SFL}}^\omega$ and $\mathsf{T}_{\text{SFL}}^\omega + (\text{cut})$ are sound. Problems might occur in connection with the infinitary rule (Q^ω) whose premises are exactly the finite stages of greatest fixed points, whereas in arbitrary Kripke structures transfinite stages cannot be ruled out. At a later point we will, however, prove the soundness of a system T_{SFL} which contains $\mathsf{T}_{\text{SFL}}^\omega$, thus dismissing such concerns.

4.2 Saturated sequents

Completeness of $\mathsf{T}_{\text{SFL}}^\omega$ with respect to the specified semantics can be shown using an extension of the method of saturated sequents used by Alberucci and Jäger in [1]. A sequent is saturated if it is in a sense maximally non-provable. The first major step in proving the completeness of $\mathsf{T}_{\text{SFL}}^\omega$ 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. For technical reasons which will shortly become apparent it is useful to define the slightly finer grained notions of k -presaturation and k -saturation of which the notion of saturation turns out to be a special case.

Definition 4.2.1 (Saturated sequents). Let k be a natural number.

1. A sequent $\Gamma \subset \mathcal{L}_{\text{SFL}}$ is called k -presaturated if all of the following properties hold:
 - (i) $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \Gamma$
 - (ii) For all formulae $A \wedge B$ with $\text{level}(A \wedge B) \geq k$, if $A \wedge B \in \Gamma$, then $A \in \Gamma$ or $B \in \Gamma$.

- (iii) For all formulae $A \vee B$ with $level(A \vee B) \geq k$, if $A \vee B \in \Gamma$, then $A \in \Gamma$ and $B \in \Gamma$.
- (iv) For all formulae $Q_{\mathcal{A}}$ with $level(Q_{\mathcal{A}}) \geq k$, if $Q_{\mathcal{A}} \in \Gamma$, then $Q_{\mathcal{A}}^n \in \Gamma$ 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 $level(P_{\mathcal{A}}) \geq k$, 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 4.2.1, 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. Conditions (ii) to (iv) are relatively unproblematic in this respect since the formulae on the right hand side of the implications are always of strictly lower complexity than those on the left hand side. What makes matters more complicated is condition (v), since satisfying this condition 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 stop Γ from being k -presaturated. The second definition introduces the set $csub(A)$ of critical subformulae of formula A of \mathcal{L}_{SFL} . The intended meaning of $csub(A)$ is being the set of all subformulae of A which could be considered during the process of saturation. Note that formulae of the form $\Box B$ and $\Diamond B$ are treated as atomic in this context, since there are no closure conditions for them in Definition 4.2.1.

Definition 4.2.2 (k -deficiency set). Let Γ be a sequent with $\top_{\text{SFL}}^\omega \not\vdash \Gamma$ and k a natural number. Define the k -deficiency set $ds_k(\Gamma)$ 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 4.2.1.

Definition 4.2.3 (Critical subformulae). For any $A \in \mathcal{L}_{\text{SFL}}$ define the set $csub(A)$ of critical subformulae of A inductively as follows:

1. If A is an element of $\Phi \cup \mathcal{V} \cup \mathcal{T}$, a fixed point constant $P_{\mathcal{A}}$, a formula $\Diamond B$ or a formula $\Box B$, then define $csub(A)$ as the set $\{A\}$.

2. If A is a fixed point constant $Q_{\mathcal{A}}$, then $\text{csub}(A) := \{Q_{\mathcal{A}}\} \cup \{Q_{\mathcal{A}}^n : n \in \omega\}$.
3. If A is a formula $B \vee C$ or $B \wedge C$, then $\text{csub}(A) := \{A\} \cup \text{csub}(B) \cup \text{csub}(C)$.

For any sequent $\Gamma = \{A_1, \dots, A_n\}$ set $\text{csub}(\Gamma) := \text{csub}(A_1) \cup \dots \cup \text{csub}(A_n)$.

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 4.2.1, 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 4.2.1 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 formalising the saturation argument just described. For this purpose, we show three crucial lemmata. The first one confirms the intuition that every non-provable sequent Γ is k -saturated with respect to some sufficiently large number k . The second lemma proves the fact that every $(k + 1)$ -saturated sequent can be extended to a k -presaturated one. This fact is an important ingredient to the proof of the third lemma which states that given a $(k + 1)$ -saturated sequent, we may extend it to a k -saturated sequent. Using the first and iterating the third lemma thus allows us to arrive at a 0-saturated and thus saturated extension of any given sequent.

Lemma 4.2.4. *For every sequent Γ with $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \Gamma$ there exists a natural number k such that Γ is k -saturated.*

Proof. By Definition 4.2.1 it is clear that the claim holds if we take k as the maximum level of all formulae in Γ plus 1. \square

Lemma 4.2.5. *Suppose that k is a natural number and Γ a $(k + 1)$ -saturated sequent. Then there exists a k -presaturated sequent Δ so that*

$$\Gamma \subset \Delta \text{ and } (\Delta \setminus \Gamma) \subset \text{csub}(ds_k(\Gamma)).$$

Proof. The proof of this lemma is routine and we will thus only give a short sketch. An almost identical claim is shown by Alberucci and Jäger [1] in full detail. The idea is to start from the $(k + 1)$ -saturated sequent provided by the

assumption and iteratively satisfy conditions (ii) to (iv) of Definition 4.2.1 with respect to those formulae of level k which violate one of these conditions. To show the claim, we essentially use the fact that each step in the iterative process preserves non-provability, as well as the set inclusions proposed by the lemma and only produces new formulae of strictly lower complexity than the one just treated. The strict decrease in rank guarantees termination of the procedure after the addition of only finitely many formulae. \square

Lemma 4.2.6. *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 4.2.5 to obtain a k -presaturated sequent Δ_0 so that

$$\Gamma \subset \Delta_0 \text{ and } (\Delta_0 \setminus \Gamma) \subset \text{csub}(ds_k(\Gamma)).$$

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

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

Now, in turn, Δ_1 need not be k -presaturated since we have no guarantee that

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

is empty. However, again by using Lemma 4.2.5 we can extend Δ_1 to a k -presaturated sequent Δ with the properties

$$\Delta_1 \subset \Delta \text{ and } (\Delta \setminus \Delta_1) \subset \text{csub}(ds_k(\Delta_1)).$$

Thus all elements of $(\Delta \setminus \Delta_1)$ belong to the set

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

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 4.2.4 and 4.2.6 now yields the result that any non-provable sequent may be extended to a saturated one. This thus takes care of the first part of our completeness proof for $\mathbb{T}_{\text{SFL}}^\omega$.

Lemma 4.2.7. *For every sequent Γ which is not provable in $\mathbb{T}_{\text{SFL}}^\omega$ there exists a saturated sequent Δ such that $\Gamma \subset \Delta$.*

Proof. Assume $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \Gamma$. Then by Lemma 4.2.4 we know that Γ is k -saturated for a suitably chosen k . Moreover, according to Lemma 4.2.6 there are sequents $\Gamma_{k-1}, \dots, \Gamma_1, \Gamma_0$ so that each Γ_i is i -saturated for $0 \leq i \leq k-1$ and

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

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

4.3 Completeness of $\mathsf{T}_{\text{SFL}}^\omega$

Based on the collection of all saturated sequents we now define a Kripke structure K_{sat} which will turn out to be a suitable countermodel for any non-provable formula of \mathcal{L}_{SFL} . The worlds of K_{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 \mathfrak{p} and $\sim \mathfrak{p}$ are elements of Γ at once. The same is also true for the symbols X and $\sim \mathsf{X}$.

Definition 4.3.1 (Canonical countermodel). Define the triple K_{sat} as follows:

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

It is easily verified that K_{sat} is a Kripke structure in the sense of Definition 2.2.1. Given a formula $A \in \mathcal{L}_{\text{SFL}}$ and a set $T \subset S_{\text{sat}}$ we will write $\|A\|_{\text{sat}}$ for $\|A\|_{\mathsf{K}_{\text{sat}}}$ and $\|A\|_{\text{sat}[\mathsf{X}:=T]}$ for $\|A\|_{\mathsf{K}_{\text{sat}}[\mathsf{X}:=T]}$.

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

Lemma 4.3.2. *Assume $\Gamma \subset \mathcal{L}_{\text{SFL}}$ is a saturated sequent.*

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

Proof. To show (i), consider that since Γ is saturated, we have $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \Gamma$ and thus also $\mathsf{T}_{\text{SFL}}^\omega \not\vdash \{B \in \mathcal{L}_{\text{SFL}} : \diamond B \in \Gamma\}, A$. Thus, by Lemma 4.2.7, there exists a saturated sequent Δ such that $\{B \in \mathcal{L}_{\text{SFL}} : \diamond B \in \Gamma\}, A \subset \Delta$ and thus $A \in \Delta$ and $\Gamma R_{\text{sat}} \Delta$. Part (ii) of the claim is obvious by Definition 4.3.1. \square

The proofs of the lemmata leading up to the completeness theorem formally proceed as an induction along the levels of the language \mathcal{L}_{SFL} . 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 K_{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 K_{sat} .

Definition 4.3.3 (k -adequacy). Let k be a natural number. We call the Kripke structure K_{sat} k -adequate if for all saturated sequents Γ and all formulae A of \mathcal{L} we have

$$\text{level}(A) \leq k \text{ and } A \in \Gamma \implies \Gamma \notin \|A\|_{\text{sat}}.$$

The naive approach to showing that if a formula A is in a saturated sequent Γ , then K_{sat} is a countermodel for A at world Γ would be an induction on the complexity $rk(A)$. Hence the argument would run along the following lines: A is in Γ , therefore by saturation of Γ some suitable formulae B_i are also in some saturated Δ and $rk(B_i) < rk(A)$. Thus, by the induction hypothesis, K_{sat} is a countermodel for B_i at world Δ and therefore by the structure of A we conclude that K_{sat} is a countermodel for A at Γ . However, this argument breaks down in one crucial case, namely when A is a constant P_A . In this case saturation of Γ implies that $\mathcal{A}[P_A]$ is an element of Γ but in general $rk(P_A) \leq rk(\mathcal{A}[P_A])$. For this reason, we must treat the case of constants P_A separately in the shape of the next lemma.

Lemma 4.3.4. *Suppose that K_{sat} is k -adequate for the natural number k and let P_A be a fixed point constant of level $k + 1$. Then we have for all X -positive formulae \mathcal{B} such that $\text{level}(\mathcal{B}) \leq k$, all saturated sequents Γ and all ordinals α that*

$$\mathcal{B}[P_A] \in \Gamma \implies \Gamma \notin \|\mathcal{B}\|_{\text{sat}[X := I_{\mathcal{A}, \mathsf{K}_{\text{sat}}}^{\leq \alpha}]}$$

Proof. We prove this claim by main induction on α and side induction on $rk(\mathcal{B})$. 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 4.3.2 and the hypothesis of the side induction. We are thus left with

the fixed point and variable cases: In case $\mathcal{B} = P_C$ or $\mathcal{B} = Q_C$ we have $\mathcal{B}[P_A] = \mathcal{B}$ and since $\text{level}(\mathcal{B}) \leq k$ and \mathbf{K}_{sat} is k -adequate, we get

$$\Gamma \notin \|\mathcal{B}\|_{\text{sat}} = \|\mathcal{B}\|_{\text{sat}[\mathbf{X}:=I_{\mathcal{A},\mathbf{K}_{\text{sat}}}^{\leq\alpha}]}$$

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

$$\Gamma \notin \|\mathcal{A}\|_{\text{sat}[\mathbf{X}:=I_{\mathcal{A},\mathbf{K}_{\text{sat}}}^{\leq\beta}]}$$

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

$$\Gamma \notin \|\mathbf{X}\|_{\text{sat}[\mathbf{X}:=I_{\mathcal{A},\mathbf{K}_{\text{sat}}}^{\leq\beta}]}$$

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

$$\Gamma \notin \|\mathbf{X}\|_{\text{sat}[\mathbf{X}:=I_{\mathcal{A},\mathbf{K}_{\text{sat}}}^{\leq\alpha}]},$$

and the claim is shown. \square

The semantics of the least fixed point constants along with the facts mentioned in Theorem 1.2.1 allow us to obtain the following corollary to Lemma 4.3.4, using the property that if no approximation of a least fixed point holds at a world then the least fixed point itself cannot hold at that world either.

Corollary 4.3.5. *Suppose that \mathbf{K}_{sat} is k -adequate for the natural number k and let P_A be a fixed point constant of level $k + 1$. Then we have for all saturated sequents Γ that*

$$P_A \in \Gamma \implies \Gamma \notin \|P_A\|_{\text{sat}}.$$

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

Lemma 4.3.6. *Suppose that \mathbf{K}_{sat} is k -adequate for the natural number k . Then for all formulae A of \mathcal{L}_{SFL} and all saturated sequents Γ we have*

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

Proof. We show this lemma by induction on $rk(A)$. The atomic, variable and truth value symbol cases are trivial. The propositional and modal cases follow by induction hypothesis, the latter using Lemma 4.3.2. 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 $\text{level}(Q_{\mathcal{B}}) \leq k + 1$. Then, by saturation of Γ , we have $Q_{\mathcal{B}}^l \in \Gamma$ for some $l \in \omega$ and $\text{level}(Q_{\mathcal{B}}^l) < \text{level}(Q_{\mathcal{B}}) \leq k + 1$. Thus, by induction hypothesis, $\Gamma \notin \|Q_{\mathcal{B}}^k\|_{\text{sat}}$ and thus also $\Gamma \notin \|Q_{\mathcal{B}}\|_{\text{sat}}$. For the second case assume that $A = P_{\mathcal{B}}$, $P_{\mathcal{B}} \in \Gamma$ and $\text{level}(P_{\mathcal{B}}) \leq k + 1$. In the case where $\text{level}(P_{\mathcal{B}}) \leq k$ it follows that $\Gamma \notin \|P_{\mathcal{B}}\|_{\text{sat}}$ by k -adequacy of \mathbf{K}_{sat} , in the case where $\text{level}(P_{\mathcal{B}}) = k + 1$ the claim follows by Corollary 4.3.5. \square

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 \mathbf{K}_{sat} at the world Γ . Thus, together with Lemma 4.2.7, we obtain the following construction of a canonical countermodel to any non-provable formula A : expand the non-provable sequent $\{A\}$ to a saturated sequent Δ . Since A is in Δ we are finished. We use this argument to prove the final completeness result.

Lemma 4.3.7. *For any natural number k , the Kripke structure \mathbf{K}_{sat} is k -adequate; i.e. for all formulae $B \in \mathcal{L}_{\text{SFL}}$ and all saturated sequents $\Gamma \subset \mathcal{L}_{\text{SFL}}$ we have*

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

Proof. This lemma is shown by an easy induction on k using Lemma 4.3.6 in the induction step. \square

Theorem 4.3.8 (Completeness of T_{SFL}^ω). *The system T_{SFL}^ω is complete, that is for all formulae $A \in \mathcal{L}_{\text{SFL}}$ if A is valid, then $T_{\text{SFL}}^\omega \vdash A$.*

Proof. We show the claim by contraposition. Assume $T_{\text{SFL}}^\omega \not\vdash A$. Then, by Lemma 4.2.7, there exists a saturated set $\Gamma \subset \mathcal{L}_{\text{SFL}}$ such that $A \in \Gamma$. Therefore by Lemma 4.3.7 we have $\Gamma \notin \|A\|_{\text{sat}}$ and thus A cannot be valid. Hence we have shown the completeness of T_{SFL}^ω . \square

4.4 Embedding H_{SFL} into T_{SFL}^ω

From the completeness of T_{SFL}^ω which was shown in Section 4.3 we can immediately deduce that H_{SFL} is contained in T_{SFL}^ω . Nevertheless, it is instructive to carry through an embedding of the closure axioms and induction rules of H_{SFL} into $T_{\text{SFL}}^\omega + (\text{cut})$. In the latter case we will see how induction in H_{SFL} corresponds to the use of the infinitary rule (Q^ω) along with a series of cuts. Before we may proceed, a series of auxiliary facts need to be established. The first is that duality extends to syntactic iterations of least and greatest fixed points.

Lemma 4.4.1. *For every natural number k we have $\neg P_{\mathcal{A}}^k = Q_{\mathcal{A}}^k$.*

Proof. The proof of this claim is a straightforward induction on k using Lemma 3.1.3 for the induction step. \square

A second fact we will require several times treats iterations of an X -positive formula \mathcal{A} on a sequent $\neg B, C$ which can be taken to stand for the implication $B \rightarrow C$. The lemma ensures that if $B \rightarrow C$ is derivable in $\mathsf{T}_{\text{SFL}}^\omega$, then so is $\mathcal{A}[B] \rightarrow \mathcal{A}[C]$ and $\mathcal{A}[\mathcal{A}[B]] \rightarrow \mathcal{A}[\mathcal{A}[C]]$ and so on.

Lemma 4.4.2. *For all formulae B, C and \mathcal{A} of \mathcal{L}_{SFL} where \mathcal{A} is X -positive we have*

$$\mathsf{T}_{\text{SFL}}^\omega \vdash \neg B, C \implies \mathsf{T}_{\text{SFL}}^\omega \vdash \neg \mathcal{A}[B], \mathcal{A}[C]$$

Proof. The claim is shown by induction on $rk(\mathcal{A})$. The base cases are trivial, the boolean and modal cases follow directly from the induction hypothesis and Lemma 4.1.4. This leaves the fixed point cases of which we will consider $\mathcal{A} = P_{\mathcal{B}}$. The case of $\mathcal{A} = Q_{\mathcal{B}}$ then follows by appropriately dualising the argument. Assume thus that $\mathcal{A} = P_{\mathcal{B}}$. It will be sufficient to show that for any natural number k

$$\mathsf{T}_{\text{SFL}}^\omega \vdash Q_{\mathcal{B}}^k, P_{\mathcal{B}} \tag{4.1}$$

since then by (Q^ω) we obtain $\mathsf{T}_{\text{SFL}}^\omega \vdash Q_{\mathcal{B}}, P_{\mathcal{B}}$ and thus $\mathsf{T}_{\text{SFL}}^\omega \vdash \neg \mathcal{P}_{\mathcal{B}}[B], \mathcal{P}_{\mathcal{B}}[C]$. We show (4.1) by induction on k , noticing that the base case is trivial since $Q_{\mathcal{B}}^0 = \top$. Thus assume that the claim holds for k and therefore in particular $\mathsf{T}_{\text{SFL}}^\omega \vdash Q_{\mathcal{B}}^k, P_{\mathcal{B}}$. By Lemma 4.4.1 we get $\mathsf{T}_{\text{SFL}}^\omega \vdash \neg P_{\mathcal{B}}^k, P_{\mathcal{B}}$. Now we have $rk(\mathcal{B}) < rk(P_{\mathcal{B}}) = rk(\mathcal{A})$ so the hypothesis of the outer induction and an application of the rule (P) yields $\mathsf{T}_{\text{SFL}}^\omega \vdash \neg \mathcal{B}[P_{\mathcal{B}}^k], P_{\mathcal{B}}$ and therefore we obtain $\mathsf{T}_{\text{SFL}}^\omega \vdash \neg P_{\mathcal{B}}^{k+1}, P_{\mathcal{B}}$. Thus, again by Lemma 4.4.1, we have $\mathsf{T}_{\text{SFL}}^\omega \vdash Q_{\mathcal{B}}^{k+1}, P_{\mathcal{B}}$ and (4.1) is shown. \square

A third fact which we will require is that the principle of excluded middle extends to arbitrary formulae of \mathcal{L}_{SFL} . This needs to be shown since the system $\mathsf{T}_{\text{SFL}}^\omega$ a priori only postulates excluded middle at the level of primitive propositions and the variable X in the shape of the axioms (ID1) and (ID2) respectively.

Lemma 4.4.3. *For all formulae A of \mathcal{L}_{SFL} we have $\mathsf{T}_{\text{SFL}}^\omega \vdash A, \neg A$.*

Proof. In order to prove this claim, we note that there always exists an X -positive formula \mathcal{A} such that the formula A can be written either as $\mathcal{A}[\mathfrak{p}]$ or $\mathcal{A}[\sim \mathfrak{p}]$ for some suitable atomic proposition \mathfrak{p} or $\sim \mathfrak{p}$ or otherwise as $\mathcal{A}[\mathsf{X}]$, $\mathcal{A}[\sim \mathsf{X}]$, $\mathcal{A}[\top]$ or $\mathcal{A}[\perp]$. Thus since by the rules (ID1), (ID2) and (ID3) we have $\mathsf{T}_{\text{SFL}}^\omega \vdash \mathfrak{p}, \sim \mathfrak{p}$, $\mathsf{T}_{\text{SFL}}^\omega \vdash \mathsf{X}, \sim \mathsf{X}$ and $\mathsf{T}_{\text{SFL}}^\omega \vdash \top, \perp$, the claim follows in each case by an application of Lemma 4.4.2. \square

We are now ready to show that the closure axiom of H_{SFL} is provable in the system T_{SFL}^ω . Moving from Hilbert-style to a Tait-style framework, this means showing the derivability of the sequent $\{\neg\mathcal{A}[P_{\mathcal{A}}], P_{\mathcal{A}}\}$.

Theorem 4.4.4. *For all formulae \mathcal{A} of \mathcal{L}_{SFL} where \mathcal{A} is X -positive we have*

$$T_{\text{SFL}}^\omega \vdash \neg\mathcal{A}[P_{\mathcal{A}}], P_{\mathcal{A}}$$

Proof. By Lemma 4.4.3 we have $T_{\text{SFL}}^\omega \vdash \neg\mathcal{A}[P_{\mathcal{A}}], \mathcal{A}[P_{\mathcal{A}}]$. Thus, by rule (P), we obtain $T_{\text{SFL}}^\omega \vdash \neg\mathcal{A}[P_{\mathcal{A}}], P_{\mathcal{A}}$ and so the claim holds. \square

We accomplish our goal by showing that the induction rule of H_{SFL} is derivable in $T_{\text{SFL}}^\omega + (\text{cut})$. Again translating between the two frameworks, this means showing that the provability of the sequent $\{\neg\mathcal{A}[B], B\}$ implies the provability of $\{\neg P_{\mathcal{A}}, B\}$. It is here that the (cut) rule is employed in a crucial way in order to bring the argument's assumption into play.

Theorem 4.4.5. *For all formulae B and \mathcal{A} of \mathcal{L}_{SFL} where \mathcal{A} is X -positive we have*

$$T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg\mathcal{A}[B], B \implies T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg P_{\mathcal{A}}, B$$

Proof. We assume that $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg\mathcal{A}[B], B$. It suffices to show by induction on k that

$$T_{\text{SFL}}^\omega + (\text{cut}) \vdash Q_{\overline{\mathcal{A}}}^k, B \tag{4.2}$$

for all natural numbers k . The base case is trivial since $Q_{\overline{\mathcal{A}}}^0 = \top$. We thus assume that the claim holds for k and aim to show it for $k+1$. By hypothesis $T_{\text{SFL}}^\omega + (\text{cut}) \vdash Q_{\overline{\mathcal{A}}}^k, B$, thus by Lemma 4.4.1 we obtain $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg P_{\mathcal{A}}^k, B$. By Lemma 4.4.2 we get $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg\mathcal{A}[P_{\mathcal{A}}^k], \mathcal{A}[B]$. Now, using Lemma 3.1.3 and again Lemma 4.4.1, we conclude $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \overline{\mathcal{A}}[Q_{\overline{\mathcal{A}}}^k], \mathcal{A}[B]$. The rule (cut) with the assumption yields $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \overline{\mathcal{A}}[Q_{\overline{\mathcal{A}}}^k], B$ and therefore $T_{\text{SFL}}^\omega + (\text{cut}) \vdash Q_{\overline{\mathcal{A}}}^{k+1}, B$ which proves (4.2). Now, from (4.2) and the rule (Q^ω) it follows that $T_{\text{SFL}}^\omega + (\text{cut}) \vdash Q_{\overline{\mathcal{A}}}, B$ and thus $T_{\text{SFL}}^\omega + (\text{cut}) \vdash \neg P_{\mathcal{A}}, B$ which concludes the proof. \square

Chapter 5

Finitising T_{SFL}^ω

As we have seen in Chapter 4, owing to the finite model property, the rule (Q^ω) relies only on the set of all finite iterations of a given greatest fixed point as premises. We will now use the so-called small model property of SFL to reduce the number of premises of the greatest fixed point rule down to a single premise. In doing so, we proceed in a similar way to the finitisation for Logic of Common Knowledge, obtained by the authors in [21]. This procedure will result in a truly finitary system in which all proofs are finite in length. It will turn out that completeness of the finitary system is implied by the completeness of T_{SFL}^ω , thus the remaining task is to prove soundness which then boils down to verifying the modified rule for greatest fixed points.

In the first section of this rather short chapter we will review the finite model property for SFL and use it to state the finitary cut-free system T_{SFL} . We also note an important proof-theoretical relationship between the infinitary system T_{SFL}^ω and T_{SFL} . In the second section we show the soundness of T_{SFL} , along with the soundness of T_{SFL}^ω which we obtain as a corollary.

5.1 The small model property and the system T_{SFL}

Before we address the soundness proof, we state the small model property in its customary form.

Remark 5.1.1 (Small model property). There exists an exponential function $f : \omega \rightarrow \omega$ such that for every formula $A \in \mathcal{L}_{\text{SFL}}$ if A is satisfiable, then there exists a Kripke structure $\mathbf{K} = (S, R, \pi)$ with $|S| < f(|A|)$ which satisfies A .

Similar to the case of the finite model property, 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 5.1.1 can be reconstructed from results presented in [32] or [10]. The exact shape of f or indeed a minimal candidate with respect to our framework shall not concern us here. We will however show how we may use f to bound the number of premises of the greatest fixed point rules.

Definition 5.1.2 (The system T_{SFL}). The system T_{SFL} is defined by replacing the rule (Q^ω) in the system $\mathsf{T}_{\text{SFL}}^\omega$ by the rule

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

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

The function f guaranteed to exist by Remark 5.1.1 is used to bound the number of iterations $Q_{\mathcal{A}}^k$ for which we need to check the derivability of $\Gamma, Q_{\mathcal{A}}^k$ before applying the rule (Q) 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) needs to be weakened by an arbitrary sequent Σ explicitly.

The next theorem establishes a proof-theoretical relationship between $\mathsf{T}_{\text{SFL}}^\omega$ and T_{SFL} and makes it clear that completeness of T_{SFL} is implied by completeness of $\mathsf{T}_{\text{SFL}}^\omega$: the rule (Q) has essentially the same shape as the rule (Q^ω) but relies on only one of the premises. Thus whenever (Q^ω) is applicable, then so is (Q) and therefore the system T_{SFL} proves at least the same formulae as $\mathsf{T}_{\text{SFL}}^\omega$.

Theorem 5.1.3. *For all finite sets Γ of \mathcal{L}_{SFL} and all ordinals α we have that*

$$\mathsf{T}_{\text{SFL}}^\omega \vdash^\alpha \Gamma \implies \mathsf{T}_{\text{SFL}} \vdash \Gamma.$$

Proof. This assertion is shown by induction on α . By Definition 5.1.2 the only non-trivial case to consider is that Γ is the conclusion of an application of the rule (Q^ω) of $\mathsf{T}_{\text{SFL}}^\omega$. Then there exist a finite set Δ , a formula $Q_{\mathcal{A}}$ and ordinals $\alpha_1, \alpha_2, \dots$ such that $\Gamma = \Delta, Q_{\mathcal{A}}$ and furthermore

$$\mathsf{T}_{\text{SFL}}^\omega \vdash^{\alpha_m} \Delta, Q_{\mathcal{A}}^m \quad \text{and} \quad \alpha_m < \alpha$$

for all natural numbers $m > 0$. Now, by induction hypothesis, we obtain

$$\mathsf{T}_{\text{SFL}} \vdash \Delta, Q_{\mathcal{A}}^m$$

for all natural numbers $m > 0$, thus in particular for all natural numbers m such that $1 \leq m \leq f(|\bigvee(\Delta, Q_{\mathcal{A}})|)$. Hence, by rule (Q) of T_{SFL} , we may conclude $\mathsf{T}_{\text{SFL}} \vdash \Delta, Q_{\mathcal{A}}$ which completes the proof. \square

Corollary 5.1.4 (Completeness of T_{SFL}). *The system T_{SFL} is complete. That is, if A is a valid formula of \mathcal{L}_{SFL} , then A is provable in T_{SFL} .*

Proof. Assume A is valid, then A is provable in $\mathsf{T}_{\text{SFL}}^\omega$ by Theorem 4.3.8 and so by Theorem 5.1.3 also in T_{SFL} . \square

5.2 Soundness of T_{SFL}

Before we proceed to show the soundness of T_{SFL} , we investigate the impact of the small model property on the closure behaviour of monotone operators in our framework. The first fact we note about the small model property is a reformulation in terms of validity instead of satisfiability and is essentially just the contraposition of Remark 5.1.1.

Lemma 5.2.1. *Let $A \in \mathcal{L}_{\text{SFL}}$. If A is valid in all Kripke structures K with $|\mathsf{K}| \leq f(|A|)$, then A is valid.*

Proof. Assume A is not valid. Then $\neg A$ is satisfiable and thus $\neg A$ is satisfied in some Kripke structure K with $|\mathsf{K}| \leq f(|\neg A|)$ by Remark 5.1.1. Since $|\neg A| = |A|$ there thus exists a Kripke structure L with $|\mathsf{L}| \leq f(|A|)$ such that A is not valid in L , namely $\mathsf{L} = \mathsf{K}$, and thus the claim holds. \square

The second fact we require in order to show the soundness of T_{SFL} is slightly more involved. It 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 fact we use to prove this claim is of course that the greatest fixed point of a monotone operator is always reached at the latest at the cardinality of the underlying structure.

Lemma 5.2.2. *If the formula $\bigvee(\Gamma, Q_{\mathcal{A}}^k)$ 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 $\bigvee(\Gamma, Q_{\mathcal{A}}^k)$ is valid it is also valid in K . Since $|\mathsf{K}| \leq k$, Theorem 1.2.2 guarantees that the formula $Q_{\mathcal{A}}^k \leftrightarrow Q_{\mathcal{A}}$ is valid in K . Thus $\bigvee(\Gamma, Q_{\mathcal{A}})$ is valid in K and the claim is shown. \square

We are now ready to state and prove the soundness of the system T_{SFL} . Since our calculus is based on standard rules for modal logic, the only really interesting part is the one concerning the rules (P) and (Q) . In both cases we need to exploit facts about monotone operators. In the first case this is the fixed point property, in the second in a sense maximality. As a corollary to the soundness of T_{SFL} we then also obtain that of $\mathsf{T}_{\text{SFL}}^\omega$.

Theorem 5.2.3 (Soundness of T_{SFL}). *The system T_{SFL} is sound, that is for all sequents $\Gamma \subset \mathcal{L}_{\text{SFL}}$ if $\mathsf{T}_{\text{SFL}} \vdash \Gamma$, then the formula $\bigvee \Gamma$ is valid.*

Proof. We must show that all axioms of T_{SFL} are valid and that all rules of T_{SFL} preserve validity. However, in view of the definition of T_{SFL} , we merely need to check the rules (P) and (Q).

(P): Assume the formula $\bigvee(\Gamma, \mathcal{A}[P_{\mathcal{A}}])$ representing the premise of (P) is valid. By Theorem 2.2.3 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) is valid and thus the rule is sound.

(Q): Assume thus that the formula $\bigvee(\Gamma, Q_{\mathcal{A}}^k)$ where $k = f(|\bigvee(\Gamma, Q_{\mathcal{A}})|)$ representing the premise of (Q) is valid. Therefore, by Lemma 5.2.2, the formula $\bigvee(\Gamma, Q_{\mathcal{A}})$ is valid in all Kripke structures K such that $|\mathsf{K}| \leq k$. Hence, by Lemma 5.2.1, the formula $\bigvee(\Gamma, Q_{\mathcal{A}})$ representing the conclusion of (Q) is valid and thus the rule is sound.

This concludes the soundness proof for the system T_{SFL} . □

Corollary 5.2.4 (Soundness of $\mathsf{T}_{\text{SFL}}^\omega$). *The system $\mathsf{T}_{\text{SFL}}^\omega$ is sound. That is, if a formula A of \mathcal{L}_{SFL} is provable in $\mathsf{T}_{\text{SFL}}^\omega$, then A is valid.*

Proof. If A is provable in $\mathsf{T}_{\text{SFL}}^\omega$, then it is also provable in T_{SFL} by Theorem 5.1.3. Thus, by Theorem 5.2.3, we obtain that A is valid. □

Chapter 6

Closure ordinals in SFL

By Theorem 1.2.1 we know that for any monotone operator F and any set S we find a minimal ordinal β such that the least fixed point of F is reached after β many iterations on S from below. Theorem 1.2.2 guarantees the existence of such an ordinal with respect to greatest fixed points and iteration of F from above. With this in mind, we will refer to the ordinal β as the closure ordinal of the least or greatest fixed point of F on S .

A well-known result (see for example [14]) about Logic of Common Knowledge, as introduced in Section 3.2, states that for any formula A and any Kripke structure \mathbf{K} we have $\|CA\|_{\mathbf{K}} = J_{E(X \wedge A), \mathbf{K}}^{<\omega} = \bigcap_{k \in \omega} \|Q_{E(X \wedge A)}^k\|_{\mathbf{K}}$, that is to say common knowledge of a formula is always attained after less than ω many iterations of the operator $F_{E(X \wedge A)}^{\mathbf{K}}$ at the latest. A similar result can be shown for least fixed points of monotone operators induced by Σ_1^0 formulae of an appropriate relational language. These turn out to also close at the ordinal ω on standard models as shown for example in Hinman [20]. In view of such results, it makes sense to ask whether similar upper bounds on the closure ordinals of fixed points can be established for SFL and in particular whether the fact that a fixed point of \mathcal{L}_{SFL} is provable or, equivalently, valid enables the establishment of such upper bounds.

This chapter is organised in two parts: the first section deals with a very simple fragment of \mathcal{L}_{SFL} in which closure ordinals of valid fixed points turn out to have a strict upper bound. The second section investigates a slightly more complex fragment, for which such an upper bound does not exist.

6.1 Closure ordinals in $\mathcal{L}_{\text{SFL}}^1$

Valid fixed points in the fragment $\mathcal{L}_{\text{SFL}}^1$ without nesting all close at the ordinal ω . This is established in the current section using the completeness result of

Section 4.3 and mostly proof-theoretic arguments. We first define the notion of a closure ordinal of a set of formulae of \mathcal{L}_{SFL} more formally. Our definitions needs to take into account that such an ordinal may not exist for a given set.

Definition 6.1.1 (Closure ordinal). Theorems 1.2.1, 1.2.2 and 2.2.3 guarantee that for every X -positive formula \mathcal{A} and every Kripke structure K there exists a least ordinal α_{K} such that the α_{K} -th iteration of $F_{\mathcal{A}}^{\mathsf{K}}$ from below equals the set $\|P_{\mathcal{A}}\|_{\mathsf{K}}$ and a least ordinal β_{K} such that the β_{K} -th approximation of $F_{\mathcal{A}}^{\mathsf{K}}$ from above equals the set $\|Q_{\mathcal{A}}\|_{\mathsf{K}}$. We say that a least fixed point $P_{\mathcal{A}}$ has closure ordinal γ if the collection of all ordinals α_{K} where K is a Kripke structure has γ as its least upper bound. Dually, we say that a greatest fixed point $Q_{\mathcal{A}}$ has closure ordinal γ if the collection of all ordinals β_{K} where K is a Kripke structure has γ as its least upper bound. Furthermore, we say a set \mathcal{G} of fixed point constants of \mathcal{L}_{SFL} has closure ordinal α if each $A \in \mathcal{G}$ has a closure ordinal and these closure ordinals have α as their least upper bound.

Following Definition 6.1.1, a proof that a set \mathcal{G} of formulae has closure ordinal α consists of two parts: firstly, we need to show that α is an upper bound for the closure ordinals of all elements of \mathcal{G} and, secondly, that there is no smaller upper bound than α . Accordingly, we first show that ω is an upper bound for the closure ordinals of all valid formulae of $\mathcal{L}_{\text{SFL}}^1$. In order to treat the least fixed point case, a restriction of the language \mathcal{L}_{SFL} to formulae without greatest fixed points may be used. The system $\mathsf{T}_{\text{SFL}}^{\omega}$ then turns out to have some useful properties with respect to this restriction of the language.

Definition 6.1.2 (The language $\mathcal{L}_{\text{SFL}}^-$). Let $\mathcal{L}_{\text{SFL}}^-$ be the set of all formulae of \mathcal{L}_{SFL} which do not contain greatest fixed point constants $Q_{\mathcal{A}}$.

Remark 6.1.3. Obviously, $\mathsf{T}_{\text{SFL}}^{\omega}$ is also complete for all formulae of $\mathcal{L}_{\text{SFL}}^-$. That is, for any formula A of $\mathcal{L}_{\text{SFL}}^-$ we have that if A is valid, then $\mathsf{T}_{\text{SFL}}^{\omega} \vdash A$. Furthermore, by the structure of the rules of $\mathsf{T}_{\text{SFL}}^{\omega}$ we note that a proof of a valid formula A of $\mathcal{L}_{\text{SFL}}^-$ may not contain applications of the rule (Q^{ω}) , since otherwise A would have to contain greatest fixed point constants, which is not allowed. It follows that such a proof is always finite in length.

An important ingredient to showing the mentioned result is the fact that if a least fixed point of $\mathcal{L}_{\text{SFL}}^-$ is provable in $\mathsf{T}_{\text{SFL}}^{\omega}$ in k many steps, then the k -th approximation of that fixed point is also provable. Indeed, we show a slightly more general version of this fact.

Lemma 6.1.4. *Let Γ be a sequent of formulae of $\mathcal{L}_{\text{SFL}}^-$, \mathcal{B} an X -positive formula of $\mathcal{L}_{\text{SFL}}^-$ and $P_{\mathcal{A}} \in \mathcal{L}_{\text{SFL}}^-$. Then we have*

$$\mathsf{T}_{\text{SFL}}^{\omega} \vdash^p \Gamma, \mathcal{B}[P_{\mathcal{A}}] \implies \mathsf{T}_{\text{SFL}}^{\omega} \vdash^p \Gamma, \mathcal{B}[P_{\mathcal{A}}^l] \text{ for all } l \geq p$$

Proof. We show the claim by induction on p . The base case of $p = 0$ is trivial, since in this case $\Gamma, \mathcal{B}[P_{\mathcal{A}}]$ is axiomatic and thus so is $\Gamma, \mathcal{B}[P_{\mathcal{A}}^l]$ for any l . Therefore, we may assume that the claim holds for any $q < p$ and that $\mathbb{T}_{\text{SFL}}^\omega \stackrel{p}{\vdash} \Gamma, \mathcal{B}[P_{\mathcal{A}}]$. We then distinguish two cases: either $\mathcal{B}[P_{\mathcal{A}}]$ was the distinguished formula of the last inference used in the proof of $\Gamma, \mathcal{B}[P_{\mathcal{A}}]$ or not.

Case 1: $\mathcal{B}[P_{\mathcal{A}}]$ was not the distinguished formula. We make a further case distinction on the rule used for the last inference. In each case the claim follows in a straightforward manner using the induction hypothesis noting that the case of the rule (Q^ω) cannot happen since $\Gamma \subset \mathcal{L}_{\text{SFL}}^-$.

Case 2: $\mathcal{B}[P_{\mathcal{A}}]$ was the distinguished formula. We again make a further case distinction on the rule used for the last inference. The rules (ID1), (ID2) and (ID3) are trivial, the rules (\vee) , (\wedge) and (\Box) are straightforward applications of the induction hypothesis and the rule (Q^ω) cannot have been used by assumption. We are thus left to consider the case of the rule (P) . In this case either $\mathcal{B} = \mathbf{X}$ or $\mathcal{B} = P_{\mathcal{C}}$.

Case 2.1: $\mathcal{B} = \mathbf{X}$. Then we have $\mathbb{T}_{\text{SFL}}^\omega \stackrel{p}{\vdash} \Gamma, P_{\mathcal{A}}$ and thus by the premise of the rule (P) and Lemma 4.1.4 we may assume that $\mathbb{T}_{\text{SFL}}^\omega \stackrel{q}{\vdash} \Gamma, P_{\mathcal{A}}, \mathcal{A}[P_{\mathcal{A}}]$ for some $q < p$. Therefore, by induction hypothesis $\mathbb{T}_{\text{SFL}}^\omega \stackrel{q}{\vdash} \Gamma, P_{\mathcal{A}}^l, \mathcal{A}[P_{\mathcal{A}}]$ for all $l \geq q$ and again by induction hypothesis ($\mathbb{T}_{\text{SFL}}^\omega \stackrel{q}{\vdash} \Gamma, P_{\mathcal{A}}^l, \mathcal{A}[P_{\mathcal{A}}^m]$ for all $m \geq q$) for all $l \geq q$. Thus, in particular, also $\mathbb{T}_{\text{SFL}}^\omega \stackrel{q}{\vdash} \Gamma, P_{\mathcal{A}}^{l+1}, \mathcal{A}[P_{\mathcal{A}}^l]$ for all $l \geq q$ and so $\mathbb{T}_{\text{SFL}}^\omega \stackrel{p}{\vdash} \Gamma, \mathcal{A}[P_{\mathcal{A}}^l]$ for all $l \geq p > q$.

Case 2.2: $\mathcal{B} = P_{\mathcal{C}}$. Then, trivially, $\mathcal{B}[P_{\mathcal{A}}] = \mathcal{B} = \mathcal{B}[P_{\mathcal{A}}^l]$ for all l , thus the claim holds. □

To treat the greatest fixed point case an inversion lemma can be shown for the rule (Q^ω) . It can be formulated over arbitrary formulae of \mathcal{L}_{SFL} and does not require a restriction to $\mathcal{L}_{\text{SFL}}^-$.

Lemma 6.1.5. *For all sequents Γ and formulae $Q_{\mathcal{A}}$ of \mathcal{L}_{SFL} we have*

$$\mathbb{T}_{\text{SFL}}^\omega \stackrel{\alpha}{\vdash} \Gamma, Q_{\mathcal{A}} \implies \mathbb{T}_{\text{SFL}}^\omega \stackrel{\alpha}{\vdash} \Gamma, Q_{\mathcal{A}}^k \text{ for all } k \in \omega.$$

Proof. We prove the claim by induction on α . For $\alpha = 0$ it is trivial since in that case $\Gamma, Q_{\mathcal{A}}$ is axiomatic and thus so is $\Gamma, Q_{\mathcal{A}}^k$. We assume thus that the claim holds for all $\beta < \alpha$ and distinguish two cases: either $Q_{\mathcal{A}}$ is the distinguished formula of the last inference in the proof of $\Gamma, Q_{\mathcal{A}}$ or $Q_{\mathcal{A}}$ is not the distinguished formula.

Case 1: $Q_{\mathcal{A}}$ is the distinguished formula. Then the rule used for the last inference was (Q^ω) and thus by the premise of this rule and Lemma 4.1.4 we may assume that $\mathbb{T}_{\text{SFL}}^\omega \vdash^{\beta_k} \Gamma, Q_{\mathcal{A}}, Q_{\mathcal{A}}^k$ where $\beta_k < \alpha$ for all $k \in \omega$. By induction hypothesis we have $\mathbb{T}_{\text{SFL}}^\omega \vdash^{\beta_k} \Gamma, Q_{\mathcal{A}}^l, Q_{\mathcal{A}}^k$ for all k and l in ω so we especially have $\mathbb{T}_{\text{SFL}}^\omega \vdash^{\beta_k} \Gamma, Q_{\mathcal{A}}^k$ and thus $\mathbb{T}_{\text{SFL}}^\omega \vdash^\alpha \Gamma, Q_{\mathcal{A}}^k$ for all $k \in \omega$.

Case 2: $Q_{\mathcal{A}}$ is not the distinguished formula. In this case we make a further case distinction on the rule used for the last inference. In the case of each rule the claim then follows by applying the induction hypothesis to the premises and applying the respective rule again.

□

Lemmata 6.1.4 and 6.1.5 essentially imply that ω is an upper bound for the closure ordinals of all valid formulae of $\mathcal{L}_{\text{SFL}}^1$. We now need to prove that ω is indeed the least upper bound. To this end, we show that for each natural number n we find a valid fixed point which does not close before n .

Definition 6.1.6. Let n be a natural number. Define \mathcal{A}_n to be the X -positive formula $\Box^n \mathbf{p} \vee \Diamond(\sim \mathbf{p} \vee \mathsf{X})$. Obviously $\mathcal{A}_n \in \mathcal{L}_{\text{SFL}}^0$ for any n .

Lemma 6.1.7. For all natural numbers $n > 0$ the formula $P_{\mathcal{A}_n}^n$ is valid.

Proof. We prove that for all $n > 0$ we have $\mathbb{T}_{\text{SFL}}^\omega \vdash P_{\mathcal{A}_n}^n$. The lemma then follows from the soundness of the system $\mathbb{T}_{\text{SFL}}^\omega$. To arrive at this goal, we first show that for any natural number m we have $\mathbb{T}_{\text{SFL}}^\omega \vdash \Box^m \mathbf{p}, \sim \mathbf{p}, P_{\mathcal{A}_n}^m$. We do this by induction on m . In the case of $m = 0$ the claim is trivial since $\mathbf{p}, \sim \mathbf{p}, P_{\mathcal{A}_n}^0$ is axiomatic. Thus assume the claim holds for m and consider the following derivation in $\mathbb{T}_{\text{SFL}}^\omega$:

$$\frac{\frac{\Box^m \mathbf{p}, \sim \mathbf{p}, P_{\mathcal{A}_n}^m}{\Box^m \mathbf{p}, \sim \mathbf{p} \vee P_{\mathcal{A}_n}^m} (\vee)}{\Box^{m+1} \mathbf{p}, \Diamond(\sim \mathbf{p} \vee P_{\mathcal{A}_n}^m), \sim \mathbf{p}, \Box^n \mathbf{p}} (\Box)}{\Box^{m+1} \mathbf{p}, \sim \mathbf{p}, P_{\mathcal{A}_n}^{m+1}} (\vee)$$

Therefore, the claim also holds for $m + 1$ and is thus shown for all m . By the claim and the fact that $n > 0$ we have $\mathbb{T}_{\text{SFL}}^\omega \vdash \Box^{n-1} \mathbf{p}, \sim \mathbf{p}, P_{\mathcal{A}_n}^{n-1}$. Now consider the following derivation in $\mathbb{T}_{\text{SFL}}^\omega$:

$$\frac{\frac{\frac{\Box^{n-1}\mathbf{p}, \sim\mathbf{p}, P_{\mathcal{A}_n}^{n-1}}{\Box^{n-1}\mathbf{p}, \sim\mathbf{p} \vee P_{\mathcal{A}_n}^{n-1}} (\vee)}{\Box^n\mathbf{p}, \Diamond(\sim\mathbf{p} \vee P_{\mathcal{A}_n}^{n-1})} (\Box)}{P_{\mathcal{A}_n}^n} (\vee)$$

Thus the lemma is shown. \square

For each natural number $m < n$ we may construct a Kripke structure in which the formula $P_{\mathcal{A}_m}^n$ is not valid. This ensures closure of $P_{\mathcal{A}_n}$ after exactly n iterations.

Definition 6.1.8. Let n be a natural number. Define the Kripke structure \mathbf{K}_n as

$$\mathbf{K}_n = (\omega, R, \pi_n),$$

where we define $R := \{(n, n+1) : n \in \omega\}$, $\pi_n(\mathbf{p}) := \{m \in \omega : m < n\}$, $\pi_n(\sim\mathbf{p}) := \omega \setminus \pi_n(\mathbf{p})$, $\pi_n(\mathbf{q}) := \pi_n(\mathbf{X}) := \emptyset$ for all other $\mathbf{q} \in \Phi$ and furthermore $\pi_n(\sim\mathbf{q}) := \pi_n(\sim\mathbf{X}) := S$ for all $\sim\mathbf{q} \in \Phi$.

The next lemma establishes a crucial relation between two Kripke structures \mathbf{K}_m and \mathbf{K}_{m+1} for any natural number m . It states that under certain conditions the number $n+1$, seen as a world of \mathbf{K}_{m+1} , behaves like the number n , seen as a world of \mathbf{K}_m .

Lemma 6.1.9. Let \mathcal{A} be an \mathbf{X} -positive formula of \mathcal{L}_{SFL} and T, S be subsets of ω such that if $n+1 \in T$, then $n \in S$. Then the following three statements hold for all natural numbers n and m :

$$(i) \quad n+1 \in \|\mathcal{A}\|_{\mathbf{K}_{m+1}[\mathbf{X}:=T]} \implies n \in \|\mathcal{A}\|_{\mathbf{K}_m[\mathbf{X}:=S]},$$

$$(ii) \quad n+1 \in I_{\mathcal{A}, \mathbf{K}_{m+1}}^{<\beta} \implies n \in I_{\mathcal{A}, \mathbf{K}_m}^{<\beta},$$

$$(iii) \quad n+1 \in J_{\mathcal{A}, \mathbf{K}_{m+1}}^{<\beta} \implies n \in J_{\mathcal{A}, \mathbf{K}_m}^{<\beta}.$$

Proof. We show these claims simultaneously by induction on $rk(\mathcal{A})$ and side induction on β . If \mathcal{A} is the atomic proposition \mathbf{p} , then (i) follows by Definition 6.1.8 and (ii) and (iii) follow by hypothesis of the side induction and (i). In the case where \mathcal{A} is a truth value symbol or an atomic proposition distinct from \mathbf{p} all three claims are trivial. If \mathcal{A} is the variable \mathbf{X} , then (i) follows by our assumption about the sets T and S and (ii) and (iii) follow by hypothesis of the side induction. In the cases where \mathcal{A} is either a formula $\mathcal{B} \wedge \mathcal{C}$ or $\mathcal{B} \vee \mathcal{C}$ claim (i) follows by hypothesis of the main induction, whereas claims (ii) and (iii) follow by hypothesis of the side induction and (i). If \mathcal{A} is a formula $\Box\mathcal{B}$,

then Definition 6.1.8 and the main induction hypothesis lead to the following chain of implications showing (i):

$$\begin{aligned}
n + 1 \in \|\Box\mathcal{B}\|_{\kappa_{m+1}[X:=T]} &\implies \\
n + 2 \in \|\mathcal{B}\|_{\kappa_{m+1}[X:=T]} &\implies \\
n + 1 \in \|\mathcal{B}\|_{\kappa_m[X:=S]} &\implies \\
n \in \|\Box\mathcal{B}\|_{\kappa_m[X:=S]}. &
\end{aligned}$$

Claims (ii) and (iii) again follow by hypothesis of the side induction and (i). For the case where \mathcal{A} is a formula $\Diamond\mathcal{B}$ we proceed analogously. This leaves us with the fixed point cases, of which we will treat the one where \mathcal{A} is a least fixed point constant $P_{\mathcal{B}}$. To show (i), we assume that we have $n + 1 \in \|P_{\mathcal{B}}\|_{\kappa_{m+1}[X:=T]} = \|P_{\mathcal{B}}\|_{\kappa_{m+1}}$. This means that for some ordinal β we have $n + 1 \in I_{\mathcal{B}, \kappa_{m+1}}^{<\beta}$ which by the main induction hypothesis using (ii) yields $n \in I_{\mathcal{B}, \kappa_m}^{<\beta}$. From this we conclude $n \in \|P_{\mathcal{B}}\|_{\kappa_m} = \|P_{\mathcal{B}}\|_{\kappa_{m+1}[X:=S]}$ and thus (i) is shown. For (ii) and (iii) we again use the hypothesis of the side induction and (i). We conclude the proof by noting that the greatest fixed point case works absolutely analogously, using (iii) and the main induction hypothesis to prove (i). \square

Remark 6.1.10. Let m and n be natural numbers such that $m \leq n$. Then from Definition 6.1.8 it follows that $0 \notin \|\Box^n \mathbf{p}\|_{\kappa_m}$. Furthermore, applying Lemma 6.1.9 with $T := S := \emptyset$ yields for any closed formula $A \in \mathcal{L}_{\text{SFL}}$ the implication $0 \in \|\Diamond A\|_{\kappa_{m+1}} \implies 0 \in \|A\|_{\kappa_m}$.

Lemma 6.1.11. *Let n be a natural number. For all $m \leq n$ and all $k < m$ we have $0 \notin \|P_{\mathcal{A}_n}^k\|_{\kappa_m}$.*

Proof. We show the claim by induction on m . In the base case of $m = 0$ there is nothing to show, since there cannot be any $k < m$. Assume thus the claim holds for m and let $k < m + 1$. If $k = 0$, then $P_{\mathcal{A}_n}^k = \perp$ thus the claim is trivially true. Otherwise, assume $k = k' + 1 > 0$ and for a contradiction $0 \in \|P_{\mathcal{A}_n}^k\|_{\kappa_{m+1}}$. Then we have $0 \in \|\Box^n \mathbf{p} \vee \Diamond(\sim \mathbf{p} \vee P_{\mathcal{A}_n}^{k'})\|_{\kappa_{m+1}}$ and thus also $0 \in \|\Box^n \mathbf{p} \vee \Diamond \sim \mathbf{p} \vee \Diamond P_{\mathcal{A}_n}^{k'}\|_{\kappa_{m+1}}$. Since by assumption $m + 1 \leq n$ we have by Remark 6.1.10 that $0 \notin \|\Box^n \mathbf{p}\|_{\kappa_{m+1}}$, therefore $0 \in \|\Diamond \sim \mathbf{p} \vee \Diamond P_{\mathcal{A}_n}^{k'}\|_{\kappa_{m+1}}$. However, since $0 < k < m + 1$ we must have $m + 1 \geq 2$ and thus we know that $0 \notin \|\Diamond \sim \mathbf{p}\|_{\kappa_{m+1}}$. This means that $0 \in \|\Diamond P_{\mathcal{A}_n}^{k'}\|_{\kappa_{m+1}}$ and again by Remark 6.1.10 we get $0 \in \|P_{\mathcal{A}_n}^{k'}\|_{\kappa_m}$. Since $k = k' + 1 < m + 1$ we also have $k' < m$ and thus we have a contradiction to the induction hypothesis. Therefore, we cannot have any $k < m + 1$ for which $0 \in \|P_{\mathcal{A}_n}^k\|_{\kappa_{m+1}}$ holds and so the claim is shown. \square

We now combine the facts obtained in Lemmata 6.1.4, 6.1.5, 6.1.7 and 6.1.11 to end up with the desired result.

Theorem 6.1.12. *The valid fixed points of $\mathcal{L}_{\text{SFL}}^1$ have closure ordinal ω .*

Proof. We must show that ω is the least upper bound of the closure ordinals of all fixed points $P_{\mathcal{A}}$ or $Q_{\mathcal{A}}$ which are valid. To see that ω is an upper bound, first consider an arbitrary constant $Q_{\mathcal{A}}$ which is valid. By Remark 6.1.3 we must have $\mathbb{T}_{\text{SFL}}^{\omega} \vdash Q_{\mathcal{A}}$. Now due to Lemma 6.1.5 we have that $\mathbb{T}_{\text{SFL}}^{\omega} \vdash Q_{\mathcal{A}}^k$ for all natural numbers k . By soundness of $\mathbb{T}_{\text{SFL}}^{\omega}$ this means that $Q_{\mathcal{A}}^k$ is valid for each k , thus the denotation of $Q_{\mathcal{A}}^k$ remains the same across all k in all Kripke structures, thus $Q_{\mathcal{A}}$ has a closure ordinal of 0. We are thus left to consider an arbitrary constant $P_{\mathcal{A}}$ which is valid. Again, by Remark 6.1.3 we may assume $\mathbb{T}_{\text{SFL}}^{\omega} \vdash^k P_{\mathcal{A}}$ for some $k \in \omega$. Now since $P_{\mathcal{A}} \in \mathcal{L}_{\text{SFL}}^1$ we may use Lemma 6.1.4 to conclude that $\mathbb{T}_{\text{SFL}}^{\omega} \vdash P_{\mathcal{A}}^k$. Thus by soundness of $\mathbb{T}_{\text{SFL}}^{\omega}$ the formula $P_{\mathcal{A}}^k$ is valid meaning that the fixed point $P_{\mathcal{A}}$ reaches closure after at most k many iterations on any Kripke structure. Thus we have established that ω is an upper bound.

To show leastness, it suffices to argue that for each natural number n , the fixed point $P_{\mathcal{A}_n}$ is valid and has closure ordinal n . By Lemma 6.1.7 and the completeness of $\mathbb{T}_{\text{SFL}}^{\omega}$ the formula $P_{\mathcal{A}_n}^n$ is valid, thus $P_{\mathcal{A}_n}$ must also be valid. By Lemma 6.1.11 the denotation of each iteration $P_{\mathcal{A}_n}^k$ where $k < n$ on \mathbb{K}_n does not already contain all of \mathbb{K}_n , thus closure of $P_{\mathcal{A}_n}$ on \mathbb{K}_n is reached only after n many steps. Therefore, ω is the least upper bound of the closure ordinals of all valid fixed points of $\mathcal{L}_{\text{SFL}}^1$. \square

6.2 Closure ordinals in $\mathcal{L}_{\text{SFL}}^2$

There is no closure ordinal for the set of valid fixed points in the fragment $\mathcal{L}_{\text{SFL}}^2$ which allows at most one nesting of fixed points. This will be the result of the current section. To obtain this result, we find for each ordinal α a valid fixed point of $\mathcal{L}_{\text{SFL}}^2$ which closes at exactly the α -th stage when evaluated on α viewed as a Kripke structure.

Definition 6.2.1. Let α be an ordinal. Define the Kripke structure \mathbb{K}_{α} as

$$\mathbb{K}_{\alpha} = (\alpha + 1, >, \pi),$$

where $\pi(P) = \emptyset$ for all $P \in \Phi \cup \mathbb{V}$.

When dealing with ordinals as Kripke structures an important notion is that of a structure \mathbb{L} being a substructure of a structure \mathbb{K} . This shall be the

case if \mathbf{K} contains at least all of the worlds of \mathbf{L} , the accessibility relation and valuation function of \mathbf{L} are the appropriate restrictions of those of \mathbf{K} and the relation R never points from a world in \mathbf{L} back to one which is in \mathbf{K} but not in \mathbf{L} .

Definition 6.2.2 (Substructure). Let $\mathbf{K} = (S, R, \pi)$ and $\mathbf{L} = (S', R', \pi')$ be Kripke structures. We say that \mathbf{L} is a substructure of \mathbf{K} if all of the following conditions hold

1. $S' \subset S$,
2. $R' = R \upharpoonright (S' \times S')$,
3. $\pi'(P) = \pi(P) \cap S'$ for all $P \in \Phi \cup \mathbf{V}$,
4. For all w in S' and all v in S we have $wRv \implies v \in S'$.

Remark 6.2.3. By Definition 6.2.1 it is clear that for all ordinals α and β if $\beta < \alpha$, then \mathbf{K}_β is a substructure of \mathbf{K}_α .

If \mathbf{L} is a substructure of \mathbf{K} , then the denotation of a formula in \mathbf{L} is a subset of the denotation of the same formula in \mathbf{K} . The same goes for all approximations of fixed points from below and above. This fact will be used several times when comparing denotations of formulae on different ordinal numbers.

Lemma 6.2.4. Let $\mathbf{K} = (S, R, \pi)$ and $\mathbf{L} = (S', R', \pi')$ be Kripke structures such that \mathbf{L} is a substructure of \mathbf{K} . Furthermore, let $T \subset S$ and $T' \subset S'$ such that $T' \subset T$. Then for all \mathbf{X} -positive formulae \mathcal{A} of \mathcal{L}_{SFL} and all ordinals α

$$(i) \ \|\mathcal{A}\|_{\mathbf{L}[X:=T']} \subset \|\mathcal{A}\|_{\mathbf{K}[X:=T]}$$

$$(ii) \ I_{\mathcal{A}, \mathbf{L}}^\alpha \subset I_{\mathcal{A}, \mathbf{K}}^\alpha$$

$$(iii) \ J_{\mathcal{A}, \mathbf{L}}^\alpha \subset J_{\mathcal{A}, \mathbf{K}}^\alpha$$

Proof. We prove all three claims simultaneously by induction on $rk(\mathcal{A})$ and side induction on α . If \mathcal{A} is an atomic proposition or a truth value symbol, all claims are trivial. If \mathcal{A} is \mathbf{X} , then (i) follows by the assumption that $T' \subset T$, (ii) and (iii) follow by hypothesis of the side induction. In case \mathcal{A} is a formula $\mathcal{A}_1 \wedge \mathcal{A}_2$ or $\mathcal{A}_1 \vee \mathcal{A}_2$, claim (i) follows by hypothesis of the main induction, claims (ii) and (iii) by hypothesis of the side induction and (i). For the modal cases we will consider $\mathcal{A} = \Box\mathcal{B}$. $\mathcal{A} = \Diamond\mathcal{B}$ then follows by

a dual argument. By the fact that \mathbf{L} is a substructure of \mathbf{K} , the induction hypothesis, and some basic set-theoretic reasoning we have

$$\begin{aligned} \|\Box\mathcal{B}\|_{\mathbf{L}[X:=T']} &= \{w \in S' : v \in \|\mathcal{B}\|_{\mathbf{L}[X:=T']} \text{ for all } v \text{ such that } wR'v\} \\ &= \{w \in S' : v \in \|\mathcal{B}\|_{\mathbf{L}[X:=T']} \text{ for all } v \text{ such that } wRv\} \\ &\subset \{w \in S : v \in \|\mathcal{B}\|_{\mathbf{L}[X:=T']} \text{ for all } v \text{ such that } wRv\} \\ &\subset \{w \in S : v \in \|\mathcal{B}\|_{\mathbf{K}[X:=T]} \text{ for all } v \text{ such that } wRv\} \\ &= \|\Box\mathcal{B}\|_{\mathbf{K}[X:=T]} \end{aligned}$$

which proves claim (i). For claims (ii) and (iii) we again use the hypothesis of the side induction and (i). We are thus left to consider the cases where \mathcal{A} is a fixed point constant. Assume that $\mathcal{A} = P_{\mathcal{B}}$. We first aim to show (i) and thus assume that $w \in \|P_{\mathcal{B}}\|_{\mathbf{L}[X:=T']} = \|P_{\mathcal{B}}\|_{\mathbf{L}}$. Therefore, we have $w \in I_{\mathcal{B},\mathbf{L}}^{\beta}$ for some ordinal β which means $w \in \|\mathcal{B}\|_{\mathbf{L}[X:=I_{\mathcal{B},\mathbf{L}}^{<\beta}]}$. Now $rk(\mathcal{B}) < rk(P_{\mathcal{B}})$, so by hypothesis of the main induction we have

$$I_{\mathcal{B},\mathbf{L}}^{<\beta} \subset I_{\mathcal{B},\mathbf{K}}^{<\beta} \text{ and } \|\mathcal{B}\|_{\mathbf{L}[X:=I_{\mathcal{B},\mathbf{L}}^{<\beta}]} \subset \|\mathcal{B}\|_{\mathbf{K}[X:=I_{\mathcal{B},\mathbf{K}}^{<\beta}]}$$

Thus also $w \in I_{\mathcal{B},\mathbf{K}}^{\beta}$ and so $w \in \|P_{\mathcal{B}}\|_{\mathbf{K}} = \|P_{\mathcal{B}}\|_{\mathbf{K}[X:=T]}$. Claims (ii) and (iii) again follow by (i) and the hypothesis of the side induction. This leaves the case of $\mathcal{A} = Q_{\mathcal{B}}$, which is treated analogously to the case of the least fixed point constants. \square

The accessibility relation in a Kripke structure \mathbf{K}_{α} runs from larger to smaller ordinals only. This means that evaluating certain modal formulae at some world $\beta < \alpha$ in \mathbf{K}_{α} is equivalent to evaluating the same formula at β in \mathbf{K}_{β} . A specific case where this fact holds is shown next.

Lemma 6.2.5. *For all ordinals α , β and γ where $\beta < \alpha$ we have:*

$$(i) \beta \in J_{\Box X, \mathbf{K}_{\alpha}}^{\gamma} \implies \beta \in J_{\Box X, \mathbf{K}_{\beta}}^{\gamma}$$

$$(ii) \beta \in I_{\Box X \vee Q_{\Box X}, \mathbf{K}_{\alpha}}^{\gamma} \implies \beta \in I_{\Box X \vee Q_{\Box X}, \mathbf{K}_{\beta}}^{\gamma}$$

Proof. Both claims are shown by separate inductions on β .

(i): Assume $\beta \in J_{\Box X, \mathbf{K}_{\alpha}}^{\gamma}$, meaning that

$$\beta \in \|\Box X\|_{\mathbf{K}_{\alpha}[X:=J_{\Box X, \mathbf{K}_{\alpha}}^{<\gamma}]} = \{\delta \in \alpha + 1 : \delta' \in J_{\Box X, \mathbf{K}_{\alpha}}^{<\gamma} \text{ for some } \delta' < \delta\}.$$

We may thus chose a $\delta' < \beta$ such that $\delta' \in J_{\Box X, \mathbf{K}_{\alpha}}^{<\gamma}$. Furthermore, let $\gamma' < \gamma$ be arbitrary. Then $\delta' \in J_{\Box X, \mathbf{K}_{\alpha}}^{\gamma'}$ and thus by induction hypothesis

$\delta' \in J_{\diamond X, K_{\delta'}}^{\gamma'}$. By Lemma 6.2.4 we also have $\delta' \in J_{\diamond X, K_{\beta}}^{\gamma'}$ and since γ' was arbitrary $\delta' \in J_{\diamond X, K_{\beta}}^{<\gamma}$. Therefore,

$$\beta \in \{\delta \in \alpha + 1 : \delta' \in J_{\diamond X, K_{\beta}}^{<\gamma} \text{ for some } \delta' < \delta\} = J_{\diamond X, K_{\beta}}^{\gamma}.$$

(ii): Assume $\beta \in I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{\gamma}$. This means that

$$\beta \in \|\square X \vee Q_{\diamond X}\|_{K_{\alpha}[X:=I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{<\gamma}]},$$

so we must distinguish two cases.

Case 1: $\beta \in \|Q_{\diamond X}\|_{K_{\alpha}[X:=I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{<\gamma}]}$. Then $\beta \in \|Q_{\diamond X}\|_{K_{\alpha}}$, so $\beta \in J_{\diamond X, K_{\alpha}}^{\xi}$ for all ordinals ξ . Now, by (i) we also have $\beta \in J_{\diamond X, K_{\beta}}^{\xi}$ for all ordinals ξ , so $\beta \in \|Q_{\diamond X}\|_{K_{\beta}} = \|Q_{\diamond X}\|_{K_{\beta}[X:=I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{<\gamma}]}$.

Case 2: $\beta \in \|\square X\|_{K_{\alpha}[X:=I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{<\gamma}]}$. Then we have

$$\beta \in \{\delta \in \alpha + 1 : \delta' \in \bigcup_{\xi < \gamma} I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{\xi} \text{ for all } \delta' < \delta\}. \quad (6.1)$$

Now if $\beta = 0$, then K_{β} consists of just one world, namely 0. In this case we trivially have $\beta \in \|\square X\|_{K_{\beta}[X:=I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{<\gamma}]}$ for any γ .

If $\beta > 0$ consider some $\delta' < \beta$. By (6.1) we have $\delta' \in I_{\square X \vee Q_{\diamond X}, K_{\alpha}}^{\xi}$ for some $\xi < \gamma$. Thus by induction hypothesis $\delta' \in I_{\square X \vee Q_{\diamond X}, K_{\delta'}}^{\xi}$ and by Lemma 6.2.4 we get $\delta' \in I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{\xi}$ and therefore $\delta' \in I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{<\gamma}$. Thus

$$\beta \in \{\delta \in \beta + 1 : \delta' \in I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{<\gamma} \text{ for all } \delta' < \delta\},$$

$$\text{so } \beta \in \|\square X\|_{K_{\beta}[X:=I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{<\gamma}]}$$

This means that in both cases $\beta \in I_{\square X \vee Q_{\diamond X}, K_{\beta}}^{\gamma}$ and so the claim is shown. \square

Every ordinal α , viewed as a world of the Kripke structure K_{α} , eventually enters the fixed point $P_{\square X \vee Q_{\diamond X}}$ but not before the α -th approximation and thus the fixed point does not close before α . This is shown next.

Lemma 6.2.6. *For all ordinals α we have*

(i) $\alpha \notin J_{\diamond X, K_\alpha}^\alpha$

(ii) $\alpha \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^\alpha$

(iii) $\alpha \notin I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\alpha}$

Proof. All claims are shown by separate inductions on α .

(i): We have

$$J_{\diamond X, K_\alpha}^\alpha = \|\diamond X\|_{K_\alpha[X:=J_{\diamond X, K_\alpha}^{<\alpha}]} = \{\gamma \in \alpha + 1 : \delta \in J_{\diamond X, K_\alpha}^{<\alpha} \text{ for some } \delta < \gamma\}.$$

Now if $\alpha = 0$, then the claim holds because K_0 consists of only the world 0. Thus we assume that $\alpha > 0$ and $\alpha \in J_{\diamond X, K_\alpha}^\alpha$, $\delta' < \alpha$ and $\delta' \in J_{\diamond X, K_\alpha}^{<\alpha}$. In this case $\delta' \in J_{\diamond X, K_\alpha}^\beta$ for all $\beta < \alpha$. Since $\delta' < \alpha$, we also have $\delta' \in J_{\diamond X, K_\alpha}^{\delta'}$ and so by Lemma 6.2.5 we obtain $\delta' \in J_{\diamond X, K_{\delta'}}^{\delta'}$ which is a contradiction to the induction hypothesis. Thus we must have $\alpha \notin J_{\diamond X, K_\alpha}^\alpha$.

(ii): We have

$$\begin{aligned} I_{\square X \vee Q_{\diamond X}, K_\alpha}^\alpha &\supset \|\square X\|_{K_\alpha[X:=I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\alpha}]} \\ &= \{\gamma \in \alpha + 1 : \delta \in \bigcup_{\beta < \alpha} I_{\square X \vee Q_{\diamond X}, K_\alpha}^\beta \text{ for all } \delta < \gamma\}. \end{aligned}$$

Assume we have $\alpha \notin I_{\square X \vee Q_{\diamond X}, K_\alpha}^\alpha$. Then by the above there exists a $\delta' < \alpha$ such that $\delta' \notin \bigcup_{\beta < \alpha} I_{\square X \vee Q_{\diamond X}, K_\alpha}^\beta$. But by induction hypothesis we have $\delta' \in I_{\square X \vee Q_{\diamond X}, K_{\delta'}}^{\delta'}$ and thus by Remark 6.2.3 and Lemma 6.2.4 we have $\delta' \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^{\delta'}$ meaning that $\delta' \in \bigcup_{\beta < \alpha} I_{\square X \vee Q_{\diamond X}, K_\alpha}^\beta$ which is a contradiction. Thus we must have $\alpha \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^\alpha$.

(iii): In the case where $\alpha = 0$ the claim is trivial. Therefore, assume that $\alpha > 0$ and $\alpha \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\alpha}$. Thus $\alpha \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^\beta$ for some $\beta < \alpha$. We distinguish two cases as to whether $\alpha \in \|\square X\|_{K_\alpha[X:=I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\beta}]}$ or not. In the first case we have $\alpha \in \|\square X\|_{K_\alpha}$ and thus $\alpha \in J_{\diamond X, K_\alpha}^{<\gamma}$ for all ordinals γ . Therefore, we also have $\alpha \in J_{\diamond X, K_\alpha}^\alpha$ which contradicts (i). In the case where $\alpha \notin \|\square X\|_{K_\alpha[X:=I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\beta}]}$ it must hold that $\alpha \in \|\square X\|_{K_\alpha[X:=I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\beta}]}$ meaning that

$$\alpha \in \{\delta \in \alpha + 1 : \delta' \in I_{\square X \vee Q_{\diamond X}, K_\alpha}^{<\beta} \text{ for all } \delta' < \delta\}.$$

Thus for all $\delta' < \alpha$ we have that $\delta' \in I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\alpha}^{< \beta}$ and so $\beta \in I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\alpha}^{< \beta}$. Therefore, $\beta \in I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\alpha}^\gamma$ for some $\gamma < \beta < \alpha$ and so by Lemma 6.2.5 we obtain $\beta \in I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\beta}^\gamma$ which means that $\beta \in I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\beta}^{< \beta}$ contradicting the induction hypothesis. Thus we must have $\alpha \notin I_{\square X \vee Q_{\diamond X}, \mathbf{K}_\alpha}^{< \alpha}$. \square

Combining the previous lemmata and the following easy observation we can state and prove the main result of this section.

Lemma 6.2.7. *For all X -positive formulae A , closed formulae B and Kripke structures \mathbf{K} we have $\|P_A \vee B\|_{\mathbf{K}} \subset \|P_{A \vee B}\|_{\mathbf{K}}$.*

Proof. Assume $\mathbf{K} = (S, R, \pi)$ is an arbitrary Kripke structure. We first show the following claim:

$$\begin{aligned} K &:= \bigcap \{T \subset S : T \supset F_A^{\mathbf{K}[X:=T]}\} \cup \|B\|_{\mathbf{K}} \\ &\subset \bigcap \{T \subset S : T \supset F_A^{\mathbf{K}[X:=T]} \cup \|B\|_{\mathbf{K}}\} =: L \end{aligned} \quad (6.2)$$

Assume thus that $w \in K$. Then the following cases can hold: either we have $w \in \|B\|_{\mathbf{K}}$ or $w \in \bigcap \{T \subset S : T \supset F_A^{\mathbf{K}[X:=T]}\}$ (or both). Now assume that T is an arbitrary set such that $T \supset F_A^{\mathbf{K}[X:=T]} \cup \|B\|_{\mathbf{K}}$. In the first case we trivially have $w \in T$ and thus $w \in L$. In the second case we also have $T \supset F_A^{\mathbf{K}[X:=T]}$, thus by assumption $w \in T$, so $w \in L$ and, therefore, we have shown (6.2). Since $\|P_A \vee B\|_{\mathbf{K}} = K$ and, furthermore, by the fact that B is a closed formula

$$L = \bigcap \{T \subset S : T \supset F_A^{\mathbf{K}[X:=T]} \cup \|B\|_{\mathbf{K}[X:=T]}\} = \|P_{A \vee B}\|_{\mathbf{K}}$$

the lemma follows by (6.2). \square

Theorem 6.2.8. *The valid fixed points of $\mathcal{L}_{\text{SFL}}^2$ have no closure ordinal.*

Proof. Obviously, the formula $P_{\square X} \vee \neg P_{\square X} = P_{\square X} \vee Q_{\diamond X}$ is valid. Thus by Lemma 6.2.7 the $\mathcal{L}_{\text{SFL}}^2$ fixed point $P_{\square X \vee Q_{\diamond X}}$ is valid. However, by Lemma 6.2.6 for every ordinal α there exists a Kripke structure, namely \mathbf{K}_α , on which $P_{\square X \vee Q_{\diamond X}}$ closes at an ordinal greater than α . Thus $P_{\square X \vee Q_{\diamond X}}$ has no closure ordinal and so the claim is shown. \square

Theorem 6.2.8 implies, furthermore, that the set of all fixed points of $\mathcal{L}_{\text{SFL}}^2$, without restriction to the valid ones does not have a closure ordinal either. Indeed, the set of all fixed points of $\mathcal{L}_{\text{SFL}}^1$ does not have a closure ordinal either. This can easily be seen by proving a lemma similar to 6.2.6 with

respect to the least fixed point P_{\Box} of $\mathcal{L}_{\text{SFL}}^1$. This non-valid fixed point again closes at arbitrarily large ordinals depending on the Kripke structure it is evaluated in. Thus in the case of fixed points of $\mathcal{L}_{\text{SFL}}^1$ validity implies closure, whereas this is not true for $\mathcal{L}_{\text{SFL}}^2$.

Chapter 7

The modal μ -calculus

We now shift our attention to the so called modal μ -calculus (or μ -calculus for short) introduced by Kozen in [23], a modal fixed point logic which contains but greatly surpasses SFL in terms of expressivity. Whereas fixed points in SFL are constructed in a restricted way so that interdependencies cannot occur, these restrictions are abolished in the case of the μ -calculus. More precisely, an arbitrary fixed point formula in the language of the μ -calculus, say for example the least fixed point $(\mu X)\mathcal{A}$ may contain free variables, say Y . Such a free variable may be bound by further fixed point quantifiers, say $(\nu Y)(\mu X)\mathcal{A}$, yielding an interleaved fixed point. In our example the denotation of $(\nu Y)(\mu X)\mathcal{A}$ then obviously depends on that of $(\mu X)\mathcal{A}$ which in turn again depends on the denotation of $(\nu Y)(\mu X)\mathcal{A}$ for the interpretation of the variable Y . It will turn out that this more general modal fixed point logic still admits a cut-free axiomatisation, a completeness argument for which will, however, be considerably more intricate.

This chapter will start by defining the language of the μ -calculus in its standard form as well as a syntactic extension of this language needed for technical reasons later. The second section will introduce the semantics in terms of Kripke structures of the languages discussed. In the third section we investigate the use of finite sequences of ordinals to measure the complexity of formulae and ultimately wellorder them. Finite sequences of ordinals are also central to the fourth section of this chapter where they are employed to define an alternative semantics which will play a crucial role later.

7.1 The languages \mathcal{L}_μ and \mathcal{L}_μ^+

The language of the μ -calculus has less structure than that of SFL and is thus easier to define. Starting from a countable set of atomic propositions

and a countable set of distinct variable symbols, we build the formulae of the language by closing under boolean and modal operators as well as under the fixed point quantifiers (μX) for least fixed points and (νX) for greatest fixed points, as long as X is a variable which appears only positively in the formula to be quantified. Simultaneously, we also define the free variables of a formula to be the set of all those variables which are not in the scope of any fixed point quantifier. For technical reasons an extension of the pure language of the μ -calculus will also be needed. This extended language also contains quantifiers ($\nu^k X$) for every natural number $k > 0$, which will be used to represent finite approximations of greatest fixed points, again assuming that X is positive variable.

Definition 7.1.1 (Language \mathcal{L}_μ , free variables). Let

$$\Phi = \{p, \sim p, q, \sim q, r, \sim r, \dots\}$$

be a countable set of atomic propositions,

$$\mathbf{V} = \{X, \sim X, Y, \sim Y, Z, \sim Z, \dots\}$$

a set containing countably many variables and their negations, $\mathbf{T} = \{\top, \perp\}$ a set containing symbols for truth and falsehood and \mathbf{M} a set of indices. Define the formulae of the language \mathcal{L}_μ as well as the set fv of free variables of each formula inductively as follows:

1. If P is an element of $\Phi \cup \mathbf{V} \cup \mathbf{T}$, then P is a formula of \mathcal{L}_μ . Furthermore, if $P = X$ or $P = \sim X$ for some X or $\sim X$ from \mathbf{V} , then $fv(P) := \{X\}$, otherwise $fv(P) := \emptyset$.
2. If A and B are formulae of \mathcal{L}_μ , then so are $(A \wedge B)$ and $(A \vee B)$. Furthermore, we define $fv((A \wedge B)) := fv((A \vee B)) := fv(A) \cup fv(B)$.
3. If A is a formula of \mathcal{L}_μ and $i \in \mathbf{M}$, then $\Box_i A$ and $\Diamond_i A$ are also formulae of \mathcal{L}_μ . Furthermore, let $fv(\Box_i A) := fv(\Diamond_i A) := fv(A)$.
4. If A is a formula of \mathcal{L}_μ and the negated variable $\sim X$ does not occur in A , then $(\mu X)A$ and $(\nu X)A$ are also formulae of \mathcal{L}_μ . Furthermore, we define $fv((\mu X)A) := fv((\nu X)A) := fv(A) \setminus \{X\}$.

In case there is no danger of confusion, we will omit parentheses in formulae. If the negated variable $\sim X$ does not occur in a formula A of \mathcal{L}_μ , we say that A is X -positive or alternatively positive in X . Formulae which are positive in a certain variable determined by the context will henceforth be denoted by letters $\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots$. A variable Y which occurs in A but is not an element of $fv(A)$ is called a bound variable of A . Furthermore, we will call A closed, if $fv(A) = \emptyset$.

Definition 7.1.2 (Language \mathcal{L}_μ^+). The formulae of the extended language \mathcal{L}_μ^+ (and their free variables) are defined by adding the following clause to the induction of Definition 7.1.1

5. If A is a formula of \mathcal{L}_μ^+ and the negated variable $\sim X$ does not occur in A , then for every natural number $k > 0$ $(\nu^k X)A$ is also a formula of \mathcal{L}_μ^+ . Furthermore, we define $fv((\nu^k X)A) := fv(A) \setminus \{X\}$.

We define X -positive and closed formulae as well as bound variables of \mathcal{L}_μ^+ , analogously to those of \mathcal{L}_μ . Given a formula B of \mathcal{L}_μ^+ , we define B^- as the formula obtained from B by first replacing all subexpressions of the form $(\nu^k X)C$ by $(\nu X)C$ and then all free variables by \top . Clearly, B^- is a formula of \mathcal{L}_μ .

Similar to the case of SFL, a very important syntactic operation is that of the substitution of positive variables with formulae. Since positive formulae in a sense stand for monotone operators, the operation of substitution can be seen as standing for operator application. Given a fixed point formula $(\mu X)\mathcal{A}$, a notable special case of positive substitution is the fixed point unfolding $\mathcal{A}[(\mu X)\mathcal{A}/X]$ where all occurrences of X in \mathcal{A} are replaced by $(\mu X)\mathcal{A}$. We adopt a convenient shorthand for a generalisation of this case, omitting the redundant mention of the positive variable X . The same is done for greatest fixed point formulae and their finite approximations.

Definition 7.1.3 (Positive substitution). If \mathcal{A} and B are formulae of \mathcal{L}_μ^+ and \mathcal{A} is positive in the variable X , then $\mathcal{A}[B/X]$ is defined as the formula obtained by simultaneously replacing every occurrence of X in \mathcal{A} by B , renaming bound variables of \mathcal{A} if they coincide with free variables of B . If \mathcal{B} is an X -positive formula of \mathcal{L}_μ^+ , then we will abbreviate $\mathcal{B}[(\mu X)\mathcal{C}/X]$ and $\mathcal{B}[(\nu X)\mathcal{C}/X]$ to $\mathcal{B}[(\mu X)\mathcal{C}]$ and $\mathcal{B}[(\nu X)\mathcal{C}]$ respectively. Similarly, we will abbreviate $\mathcal{B}[(\nu^k X)\mathcal{C}/X]$ to $\mathcal{B}[(\nu^k X)\mathcal{C}]$.

For formulae of the language \mathcal{L}_μ we define negation as usual reflecting De Morgan's laws, the law of double negation and the law of fixed point duality as asserted in Theorem 1.2.3. Negation is not defined for the language \mathcal{L}_μ^+ since we have not included duals for formulae of the form $(\nu^k X)\mathcal{A}$ where k is a natural number greater than 0. It will turn out that syntactic negation is not needed for formulae of this form.

Definition 7.1.4 (Negation). The negation $\neg A$ of a formula A of \mathcal{L}_μ is defined inductively as follows:

1. $\neg p := \sim p$ and $\neg \sim p := p$ for all $p, \sim p \in \Phi$.

2. $\neg X := \sim X$ and $\neg \sim X := X$ for all $X, \sim X \in \mathcal{V}$.
3. $\neg \top := \perp$ and $\neg \perp := \top$.
4. If B and C are formulae of \mathcal{L}_μ , then $\neg(B \wedge C) := \neg B \vee \neg C$ and $\neg(B \vee C) := \neg B \wedge \neg C$.
5. If B is a formula of \mathcal{L}_μ , then $\neg \square_i B := \diamond_i \neg B$ and $\neg \diamond_i B := \square_i \neg B$ for all $i \in \mathcal{M}$.
6. If \mathcal{A} is an X -positive formula of \mathcal{L}_μ , then $\neg(\mu X)\mathcal{A} := (\nu X)\overline{\mathcal{A}}$ and $\neg(\nu X)\mathcal{A} := (\mu X)\overline{\mathcal{A}}$ where $\overline{\mathcal{A}}$ is the formula $\neg(\mathcal{A}[\sim X/X])$.

We observe that the definitions of $\neg(\mu X)\mathcal{A} = (\nu X)\overline{\mathcal{A}}$ and $\neg(\nu X)\mathcal{A} = (\mu X)\overline{\mathcal{A}}$ are well formed since the formula $\overline{\mathcal{A}}$ is indeed X -positive.

While finite approximations of greatest fixed points are explicitly included in the language \mathcal{L}_μ^+ these are not featured in the smaller language \mathcal{L}_μ . Nevertheless, for any X -positive formula \mathcal{A} a set of formulae of \mathcal{L}_μ can be defined to play the same role.

Definition 7.1.5 (Finite approximations). Let \mathcal{A} be an X -positive formula of \mathcal{L}_μ . We define the finite approximations of $(\nu X)\mathcal{A}$ as formulae of \mathcal{L}_μ inductively for each natural number $n > 0$:

$$(\nu X)^1 \mathcal{A} := \mathcal{A}[\top/X] \quad \text{and} \quad (\nu X)^{n+1} \mathcal{A} := \mathcal{A}[(\nu X)^n \mathcal{A}]$$

We can now set up a translation of formulae of the language \mathcal{L}_μ^+ into formulae of the language \mathcal{L}_μ by using the abbreviations introduced in Definition 7.1.5 to translate greatest fixed point formulae. Later, when defining deductive systems for formulae of the two languages \mathcal{L}_μ^+ and \mathcal{L}_μ , we will see that this translation also behaves adequately with respect to provability.

Definition 7.1.6 (Translation). The translation A° of a formula A of \mathcal{L}_μ^+ is defined inductively on the structure of A :

1. If $A \in \Phi \cup \mathcal{V} \cup \top$, then $A^\circ := A$.
2. If A is of the form $B \wedge C$, then $A^\circ := B^\circ \wedge C^\circ$. If A is of the form $B \vee C$, then $A^\circ := B^\circ \vee C^\circ$.
3. If A is of the form $\square_i B$, then $A^\circ := \square_i B^\circ$. If A is of the form $\diamond_i B$, then $A^\circ := \diamond_i B^\circ$.
4. If A is of the form $(\mu X)\mathcal{A}$, then $A^\circ := (\mu X)\mathcal{A}^\circ$. If A is of the form $(\nu X)\mathcal{A}$, then $A^\circ := (\nu X)\mathcal{A}^\circ$.

5. If A is of the form $(\nu^n X)A$ for some natural number $n > 0$, then $A^\circ := (\nu X)^n A^\circ$.

The definition is extended to sequents of \mathcal{L}_μ^+ : for $\Gamma = \{A_1, \dots, A_m\}$ we set $\Gamma^\circ := \{A_1^\circ, \dots, A_m^\circ\}$.

The next lemma summarises two important facts, both of which are straightforward to prove. On the one hand, as already mentioned, the translation introduced in Definition 7.1.6 maps the formulae of \mathcal{L}_μ^+ into those of \mathcal{L}_μ . On the other hand, nothing happens when translating a formula of \mathcal{L}_μ .

Lemma 7.1.7. *If A is a formula of \mathcal{L}_μ^+ , then A° is a formula of \mathcal{L}_μ . Moreover, if B is a formula of \mathcal{L}_μ , then $B^\circ = B$.*

7.2 Semantics of \mathcal{L}_μ^+

The standard semantics for formulae of \mathcal{L}_μ^+ are again given in terms of Kripke structures. Even though the definition of a Kripke structure for \mathcal{L}_μ^+ is very similar to the one given with respect to \mathcal{L}_{SFL} in Chapter 2 we repeat it here for ease of reference and because it needs a slight adjustment to cater for the presence of countably many distinct variables in \mathcal{L}_μ^+ .

Definition 7.2.1 (Kripke structure). A Kripke structure for \mathcal{L}_μ^+ is a triple $\mathbb{K} = (S, R, \pi)$, where S is a non-empty set, $R : \mathbb{M} \rightarrow \mathcal{P}(S \times S)$ and where $\pi : (\Phi \cup \mathbb{V}) \rightarrow \mathcal{P}(S)$ is a function such that $\pi(\sim X) = S \setminus \pi(X)$ for all $\sim X \in \mathbb{V}$ and $\pi(\sim \mathbf{p}) = S \setminus \pi(\mathbf{p})$ for all $\sim \mathbf{p} \in \Phi$. The function R assigns an accessibility relation to each $i \in \mathbb{M}$ where we write R_i for the relation $R(i)$. Furthermore, given a set $T \subset S$ and a variable $X \in \mathbb{V}$ we define the Kripke structure $\mathbb{K}[X := T]$ as the triple (S, R, π') , where $\pi'(X) = T$, $\pi'(\sim X) = S \setminus T$ and $\pi'(P) = \pi(P)$ for all other $P \in \Phi \cup \mathbb{V}$.

We are now ready to assign a meaning to the formulae of \mathcal{L}_μ^+ in terms of Kripke structures. This is achieved in a straightforward way by induction on the structure of formulae, with a side induction on all natural numbers greater than 0 to treat finite greatest fixed point approximations.

Definition 7.2.2 (Denotation). Let $\mathbb{K} = (S, R, \pi)$ be a Kripke structure. For every $A \in \mathcal{L}_\mu^+$ we define the set $\|A\|_{\mathbb{K}} \subset S$ inductively as follows:

$$\begin{aligned} \|P\|_{\mathbb{K}} &:= \pi(P) \text{ for all } P \in \Phi \cup \mathbb{V}, & \|\top\|_{\mathbb{K}} &:= S, & \|\perp\|_{\mathbb{K}} &:= \emptyset, \\ \|B \wedge C\|_{\mathbb{K}} &:= \|B\|_{\mathbb{K}} \cap \|C\|_{\mathbb{K}}, & \|B \vee C\|_{\mathbb{K}} &:= \|B\|_{\mathbb{K}} \cup \|C\|_{\mathbb{K}}, \\ \|\Box_i B\|_{\mathbb{K}} &:= \{w \in S : v \in \|B\|_{\mathbb{K}} \text{ for all } v \text{ such that } wR_i v\}, \\ \|\Diamond_i B\|_{\mathbb{K}} &:= \{w \in S : v \in \|B\|_{\mathbb{K}} \text{ for some } v \text{ such that } wR_i v\}. \end{aligned}$$

For every formula $(\mu X)\mathcal{A}$ and $(\nu X)\mathcal{A}$ we define

$$\begin{aligned} \|(\mu X)\mathcal{A}\|_{\mathbf{K}} &:= \bigcap \{T \subset S : T \supset F_{\mathcal{A},X}^{\mathbf{K}}(T)\} \text{ and} \\ \|(\nu X)\mathcal{A}\|_{\mathbf{K}} &:= \bigcup \{T \subset S : T \subset F_{\mathcal{A},X}^{\mathbf{K}}(T)\} \end{aligned}$$

where $F_{\mathcal{A},X}^{\mathbf{K}}$ is the operator on S given by $F_{\mathcal{A},X}^{\mathbf{K}}(T) := \|\mathcal{A}\|_{\mathbf{K}[X:=T]}$ for every subset T of S . Furthermore, if \mathcal{A} is an X -positive formula, then we define $\|(\nu^k X)\mathcal{A}\|_{\mathbf{K}}$ for every $k > 0$ by side induction on k as follows:

$$\begin{aligned} \|(\nu^1 X)\mathcal{A}\|_{\mathbf{K}} &:= \|\mathcal{A}[T/X]\|_{\mathbf{K}} \\ \|(\nu^{n+1} X)\mathcal{A}\|_{\mathbf{K}} &:= \|\mathcal{A}[(\nu^n X)\mathcal{A}]\|_{\mathbf{K}}. \end{aligned}$$

In analogy to Theorem 2.2.3, we must show that the manner in which the denotation of formulae of \mathcal{L}_{μ}^+ is defined actually guarantees that formulae of the form $(\mu X)\mathcal{A}$ and $(\nu X)\mathcal{A}$ are interpreted as least and greatest fixed points.

Theorem 7.2.3. *If \mathcal{A} is an X -positive formula of \mathcal{L}_{μ}^+ , then $F_{\mathcal{A},X}^{\mathbf{K}}$ is monotone and $\|(\mu X)\mathcal{A}\|_{\mathbf{K}}$ and $\|(\nu X)\mathcal{A}\|_{\mathbf{K}}$ are the least and greatest fixed points of $F_{\mathcal{A},X}^{\mathbf{K}}$ respectively. That is if $\mathbf{K} = (S, R, \pi)$ is a Kripke structure, then*

- (i) $U \subset T \implies F_{\mathcal{A},X}^{\mathbf{K}}(U) \subset F_{\mathcal{A},X}^{\mathbf{K}}(T)$ for all $U, T \subset S$
- (ii) $\|\mathcal{A}[(\mu X)\mathcal{A}]\|_{\mathbf{K}} = \|(\mu X)\mathcal{A}\|_{\mathbf{K}}$
- (iii) $F_{\mathcal{A},X}^{\mathbf{K}}(T) = T \implies T \supset \|(\mu X)\mathcal{A}\|_{\mathbf{K}}$ for all $T \subset S$
- (iv) $\|\mathcal{A}[(\nu X)\mathcal{A}]\|_{\mathbf{K}} = \|(\nu X)\mathcal{A}\|_{\mathbf{K}}$
- (v) $F_{\mathcal{A},X}^{\mathbf{K}}(T) = T \implies T \subset \|(\nu X)\mathcal{A}\|_{\mathbf{K}}$ for all $T \subset S$.

The only assertion which requires some work is (i). Assertions (ii) to (v) follow from (i) along with Theorem 1.2.1 for the least fixed point case and Theorem 1.2.2 for the greatest fixed point case. The proof of assertion (i) proceeds by main induction on the structure of the formula \mathcal{A} , in each case showing the claim for $(\nu^k X)\mathcal{A}$ by side induction on k .

For the sake of easy reference, we next state the definition of the notions of satisfaction and validity for formulae of \mathcal{L}_{μ}^+ . Since we will be dealing with sound and complete axiomatisations of the μ -calculus, validity will remain the central semantic notion in our context.

Definition 7.2.4 (Satisfaction and validity). Let $\mathbf{K} = (S, R, \pi)$ be a Kripke structure. We say a formula $A \in \mathcal{L}_{\mu}^+$ is satisfied in \mathbf{K} if $\|A\|_{\mathbf{K}} \neq \emptyset$ and valid in \mathbf{K} if $\|A\|_{\mathbf{K}} = S$. We say A is satisfiable if there exists a Kripke structure in which A is satisfied. Furthermore, we say that A is valid if it is valid in all Kripke structures.

In order to deal with arbitrary transfinite iterations semantically, we introduce a similar notation as was already used in the case of SFL. When specifying an iteration we are required to specify four parameters: an ordinal for the number of iterations, a variable X , a formula \mathcal{A} which is positive in X and a Kripke structure K on which the iteration is evaluated. This leads to a notation which is rather bulky at first sight but will prove to be quite intuitive when it is used in subsequent arguments.

Definition 7.2.5. Let \mathcal{A} be an X -positive formula of \mathcal{L}_μ^+ and $K = (S, R, \pi)$ be a Kripke structure for \mathcal{L}_μ^+ . For every ordinal α define the subsets $I_{\mathcal{A},X,K}^{<\alpha}$, $I_{\mathcal{A},X,K}^\alpha$, $J_{\mathcal{A},X,K}^{<\alpha}$ and $J_{\mathcal{A},X,K}^\alpha$ of S as follows:

$$I_{\mathcal{A},X,K}^{<\alpha} := I_{F_{\mathcal{A},X}^K}^{<\alpha}, \quad I_{\mathcal{A},X,K}^\alpha := I_{F_{\mathcal{A},X}^K}^\alpha, \quad J_{\mathcal{A},X,K}^{<\alpha} := J_{F_{\mathcal{A},X}^K}^{<\alpha}, \quad J_{\mathcal{A},X,K}^\alpha := J_{F_{\mathcal{A},X}^K}^\alpha.$$

7.3 Formula complexity

The language \mathcal{L}_μ^+ , unlike \mathcal{L}_{SFL} , is not structured hierarchically. Thus we may not rely on any such structure when defining a measure for the complexity of formulae of \mathcal{L}_μ^+ under which a greatest fixed point formula $(\nu X)\mathcal{A}$ is strictly greater than any finite approximation formula $(\nu^{k+1}X)\mathcal{A}$ and also, as we shall require, strictly greater than $\mathcal{A}[(\nu^k X)\mathcal{A}]$. In order to nevertheless succeed in the construction of a measure with the properties just mentioned, two notions will be of central importance: firstly, finite sequences of ordinals and, secondly, the lexicographic ordering on these sequences.

Definition 7.3.1 (Signatures). A signature σ is a finite sequence of ordinals $\langle \alpha_1, \dots, \alpha_n \rangle$ the length $lh(\sigma)$ of which is n and the i -th component $(\sigma)_i$ of which is the ordinal α_i . The empty signature is denoted by $\langle \rangle$ and we set $lh(\langle \rangle) := 0$. We will henceforth denote signatures by boldface Greek letters σ, τ, \dots . Assume that σ and τ are arbitrary signatures. We define two ordering relations on signatures:

1. τ is lexicographically smaller than σ , in symbols $\tau <_{lex} \sigma$, if and only if τ is a proper initial segment of σ or there exists an i such that $(\tau)_i < (\sigma)_i$ and for all $j < i$ we have $(\tau)_j = (\sigma)_j$. The reflexive closure of $<_{lex}$ shall be denoted by \leq_{lex} .
2. τ is componentwise smaller than σ , in symbols $\tau \trianglelefteq \sigma$, if and only if $lh(\tau) \leq lh(\sigma)$ and $(\tau)_i \leq (\sigma)_i$ for all $1 \leq i \leq lh(\tau)$.

Furthermore, for arbitrary signatures $\sigma = \langle \alpha_1, \dots, \alpha_n \rangle$ and $\tau = \langle \beta_1, \dots, \beta_m \rangle$ we define two operations:

1. The concatenation of σ and τ , in symbols $\sigma * \tau$, is defined by setting $\sigma * \langle \rangle := \langle \rangle * \sigma := \sigma$ and $\sigma * \tau := \langle \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_m \rangle$ for the case where τ is not empty.
2. The maximisation of σ and τ , in symbols $\sigma \sqcup \tau$, is defined by setting $\sigma \sqcup \langle \rangle := \langle \rangle \sqcup \sigma := \sigma$ and

$$\sigma \sqcup \tau := \begin{cases} \langle \max(\alpha_1, \beta_1), \dots, \max(\alpha_m, \beta_m), \alpha_{m+1}, \dots, \alpha_n \rangle & \text{if } m \leq n \\ \langle \max(\alpha_1, \beta_1), \dots, \max(\alpha_n, \beta_n), \beta_{n+1}, \dots, \beta_m \rangle & \text{otherwise} \end{cases}$$

for the case where τ is not empty.

Remark 7.3.2. It is easily seen that the relation \sqsubseteq is transitive. Furthermore, we note that the relation $<_{lex}$ is a wellordering on every set of signatures with length bounded by some natural number, though not a wellordering in general.

Making a first use of signatures, we define the rank of a formula of \mathcal{L}_μ^+ , assigning each formula a signature containing all iteration numbers k of subformulae of the form $(\nu^k X)\mathcal{A}$ in the appropriate order. Subformulae in the form of greatest fixed points $(\nu X)\mathcal{A}$ are assigned the ordinal ω in order to majorise all finite approximations.

Definition 7.3.3 (Rank, length). The rank $rk(A)$ of a formula A of \mathcal{L}_μ^+ is inductively defined as follows:

1. $rk(P) := \langle 0 \rangle$ for all $P \in \Phi \cup \mathbb{V} \cup \mathbb{T}$.
2. $rk(B \wedge C) := rk(B \vee C) := (rk(B) \sqcup rk(C)) * \langle 0 \rangle$
3. $rk(\Box_i B) := rk(\Diamond_i B) := rk(B) * \langle 0 \rangle$
4. $rk((\mu X)\mathcal{A}) := rk(\mathcal{A}) * \langle 0 \rangle$
5. $rk((\nu X)\mathcal{A}) := rk(\mathcal{A}) * \langle \omega \rangle$
6. $rk((\nu^n X)\mathcal{A}) := rk(\mathcal{A}) * \langle n \rangle$

The length $lh(A)$ of A is defined as $lh(rk(A))$.

Formulae of \mathcal{L}_μ^+ can be ordered according to their ranks using the relations $<_{lex}$ and \sqsubseteq . Later, it will turn out that $<_{lex}$ is indeed a wellordering on a particularly important set of formulae used in our arguments. Let us, however, first focus on some of the more basic properties of the ordering relations $<_{lex}$ and \sqsubseteq , all of which are direct consequences of Definitions 7.3.1 and 7.3.3.

Lemma 7.3.4. *For all formulae A, B and \mathcal{A} of \mathcal{L}_μ^+ where \mathcal{A} is \mathbf{X} -positive and all natural numbers $n > 0$ we have:*

1. $rk(A), rk(B) <_{lex} rk(A \vee B) = rk(A \wedge B)$
2. $lh(A), lh(B) < lh(A \vee B) = lh(A \wedge B)$
3. $rk(B) <_{lex} rk(\Box_i B) = rk(\Diamond_i B)$
4. $lh(B) < lh(\Box_i B) = lh(\Diamond_i B)$
5. $rk(\mathcal{A}) = rk(\mathcal{A}[\perp/\mathbf{X}]) = rk(\mathcal{A}[\top/\mathbf{X}])$
6. $lh(\mathcal{A}) = lh(\mathcal{A}[\perp/\mathbf{X}]) = lh(\mathcal{A}[\top/\mathbf{X}])$
7. $rk(\mathcal{A}) <_{lex} rk((\mu\mathbf{X})\mathcal{A}), rk((\nu\mathbf{X})\mathcal{A}), rk((\nu^n\mathbf{X})\mathcal{A})$
8. $lh(\mathcal{A}) < lh((\nu\mathbf{X})\mathcal{A}) = lh((\nu^n\mathbf{X})\mathcal{A})$

Next, we show a technical lemma which will be used to ensure that the rank function behaves as required with respect to greatest fixed point approximations. The lemma states that if a formula B is substituted for a variable \mathbf{X} in an \mathbf{X} -positive formula \mathcal{A} which has a componentwise smaller rank than B , then the rank of B is an initial segment of the rank of the resulting formula.

Lemma 7.3.5. *Let \mathcal{A} be an \mathbf{X} -positive formula of \mathcal{L}_μ^+ such that \mathbf{X} occurs in \mathcal{A} and suppose furthermore that B is a formula of \mathcal{L}_μ^+ such that $rk(\mathcal{A}) \preceq rk(B)$. Then there exists a signature σ so that $rk(\mathcal{A}[B/\mathbf{X}]) = rk(B) * \sigma$.*

Proof. We show this lemma by induction on $lh(\mathcal{A})$ and distinguish the following cases:

1. $lh(\mathcal{A}) = 1$. Then \mathcal{A} is the variable \mathbf{X} and the claim is trivial.
2. We are assuming the induction hypothesis and \mathcal{A} is the formula $\mathcal{B} \vee \mathcal{C}$. Then we have

$$rk(\mathcal{B}), rk(\mathcal{C}) \preceq rk(\mathcal{A}) \preceq rk(B). \quad (7.1)$$

If \mathbf{X} occurs in both \mathcal{B} and \mathcal{C} , then we can apply the induction hypothesis twice to obtain

$$rk(\mathcal{B}[B/\mathbf{X}]) = rk(B) * \sigma_1 \quad \text{and} \quad rk(\mathcal{C}[B/\mathbf{X}]) = rk(B) * \sigma_2 \quad (7.2)$$

for suitable signatures σ_1 and σ_2 . Clearly, this implies

$$rk(\mathcal{A}[B/\mathbf{X}]) = ((rk(B) * \sigma_1) \sqcup (rk(B) * \sigma_2)) * \langle 0 \rangle \quad (7.3)$$

and thus the claim holds for $\sigma = (\sigma_1 \sqcup \sigma_2) * \langle 0 \rangle$.

If X occurs in only one of the formulae \mathcal{B} and \mathcal{C} , say in \mathcal{B} , then the induction hypothesis yields

$$rk(\mathcal{B}[B/X]) = rk(B) * \sigma_1 \quad (7.4)$$

for some suitable σ_1 . We also know that $\mathcal{C}[B/X] = \mathcal{C}$ and deduce from (7.1) that

$$rk(\mathcal{C}[B/X]) \leq rk(B). \quad (7.5)$$

From (7.4) and (7.5) we conclude that

$$rk(\mathcal{A}[B/X]) = ((rk(B) * \sigma_1) \sqcup rk(\mathcal{C}[B/X])) * \langle 0 \rangle = rk(B) * \sigma \quad (7.6)$$

where σ is the signature $\sigma_1 * \langle 0 \rangle$.

3. We are assuming the induction hypothesis and \mathcal{A} is the formula $\mathcal{B} \wedge \mathcal{C}$. Then we proceed analogously to the previous case.
4. We are assuming the induction hypothesis and \mathcal{A} is a formula of a form not covered so far. Then the claim is immediate by induction hypothesis.

□

We now address some important issues concerning the rank of greatest fixed point approximations. The most notable of these is the fact that the rank of a formula $\mathcal{A}[(\nu^n X)\mathcal{A}]$ is lexicographically smaller than that of $(\nu^{n+1} X)\mathcal{A}$ for all natural numbers $n > 0$. To prove this we will require Lemma 7.3.5.

Theorem 7.3.6. *For all formulae $(\nu X)\mathcal{A}$ of \mathcal{L}_μ^+ and all natural numbers $n > 0$ we have*

1. $rk(\mathcal{A}[\top/X]) <_{lex} rk((\nu^1 X)\mathcal{A})$
2. $rk(\mathcal{A}[(\nu^n X)\mathcal{A}]) <_{lex} rk((\nu^{n+1} X)\mathcal{A})$
3. $rk((\nu^n X)\mathcal{A}) <_{lex} rk((\nu X)\mathcal{A})$

Proof. The first and third assertions are immediate consequences of Definition 7.3.3. In order to prove the second assertion we make a case distinction as to whether $X \in fv(\mathcal{A})$ or not. If X is not free in \mathcal{A} , then $\mathcal{A}[(\nu^n X)\mathcal{A}] = \mathcal{A}$ and hence

$$rk(\mathcal{A}[(\nu^n X)\mathcal{A}]) <_{lex} rk(\mathcal{A}) * \langle n + 1 \rangle = rk((\nu^{n+1} X)\mathcal{A}).$$

We are thus left to treat the case where $X \in fv(\mathcal{A})$. In view of Definition 7.3.3 we know that $rk(\mathcal{A}) \leq rk((\nu^n X)\mathcal{A})$. Hence Lemma 7.3.5 and again Definition 7.3.3 yield

$$rk(\mathcal{A}[(\nu^n X)\mathcal{A}]) = rk((\nu^n X)\mathcal{A}) * \sigma = (rk(\mathcal{A}) * \langle n \rangle) * \sigma$$

for some suitable signature σ . Together with the fact that

$$rk((\nu^{n+1} X)\mathcal{A}) = rk(\mathcal{A}) * \langle n + 1 \rangle$$

we get $rk(\mathcal{A}[(\nu^n X)\mathcal{A}]) <_{lex} rk((\nu^{n+1} X)\mathcal{A})$ which concludes the proof. \square

So far, we have seen that the rank function has all desirable properties with respect to the ordering relation $<_{lex}$ except for one particular property, namely that of being a wellordering. Indeed, by Remark 7.3.2 we know that $<_{lex}$ cannot be a wellordering on all formulae of \mathcal{L}_μ^+ , as the lengths of all of these formulae are obviously not bounded by any natural number. We will now introduce the strong closure of a formula D of \mathcal{L}_μ which, as it turns out, contains only formulae of bounded length and is thus wellordered by $<_{lex}$. For this reason, the strong closure will play a central role in the arguments to follow.

The strong closure is related to a closure set which is better known in the context of modal logic with fixed points, namely the so-called Fischer-Ladner closure of a formula. In order to show that the strong closure of a formula contains only formulae of bounded length, we will rely on the fact that the Fischer-Ladner closure of that same formula is finite. This can easily be shown by adapting the finiteness proof found in [16] to the current formal framework.

Definition 7.3.7 (Fischer-Ladner closure). Let D be a closed formula of \mathcal{L}_μ . The Fischer-Ladner closure $\mathbb{FL}(D)$ of D is defined inductively as follows:

1. $D \in \mathbb{FL}(D)$
2. If $A \wedge B \in \mathbb{FL}(D)$ or $A \vee B \in \mathbb{FL}(D)$, then $A \in \mathbb{FL}(D)$ and $B \in \mathbb{FL}(D)$.
3. If $\Box_i A \in \mathbb{FL}(D)$ or $\Diamond_i A \in \mathbb{FL}(D)$, then $A \in \mathbb{FL}(D)$.
4. If $(\mu X)\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A} \in \mathbb{FL}(D)$ and $\mathcal{A}[(\mu X)\mathcal{A}] \in \mathbb{FL}(D)$.
5. If $(\nu X)\mathcal{A} \in \mathbb{FL}(D)$, then $\mathcal{A} \in \mathbb{FL}(D)$ and $\mathcal{A}[(\nu X)\mathcal{A}] \in \mathbb{FL}(D)$.

Lemma 7.3.8. *The cardinality of $\mathbb{FL}(D)$ of a formula D of \mathcal{L}_μ is linear in $lh(D)$, thus in particular $\mathbb{FL}(D)$ is a finite set.*

The defining clauses for the strong closure of a formula are identical to those of the Fischer-Ladner closure in what concerns propositional, modal and least fixed point formulae. Greatest fixed point formulae, however, are treated differently, introducing infinitely many finite approximations into the closure.

Definition 7.3.9 (Strong closure). Let D be a closed formula of \mathcal{L}_μ . The strong closure $\mathbb{S}\mathbb{C}(D)$ of D is defined inductively as follows:

1. $D \in \mathbb{S}\mathbb{C}(D)$
2. If $A \wedge B \in \mathbb{S}\mathbb{C}(D)$ or $A \vee B \in \mathbb{S}\mathbb{C}(D)$, then $A \in \mathbb{S}\mathbb{C}(D)$ and $B \in \mathbb{S}\mathbb{C}(D)$.
3. If $\Box_i A \in \mathbb{S}\mathbb{C}(D)$ or $\Diamond_i A \in \mathbb{S}\mathbb{C}(D)$, then $A \in \mathbb{S}\mathbb{C}(D)$.
4. If $(\mu X)\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$, then $\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$ and $\mathcal{A}[(\mu X)\mathcal{A}] \in \mathbb{S}\mathbb{C}(D)$.
5. If $(\nu X)\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$, then $\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$ and for every natural number $n > 0$ also $(\nu^n X)\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$.
6. If $(\nu^1 X)\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$, then $\mathcal{A}[\top/X] \in \mathbb{S}\mathbb{C}(D)$.
7. If n is a natural number greater than 0 and $(\nu^{n+1} X)\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$, then $\mathcal{A}[(\nu^n X)\mathcal{A}] \in \mathbb{S}\mathbb{C}(D)$.
8. If \mathcal{A} is X -positive and $\mathcal{A} \in \mathbb{S}\mathbb{C}(D)$, then for every variable Y also $\mathcal{A}[Y/X] \in \mathbb{S}\mathbb{C}(D)$.

From Definition 7.3.9 it is clear that the strong closure of a formula is in general an infinite set. Nevertheless, it will turn out that the set only contains formulae of finitely many different lengths, a fact which is sufficient to guarantee the wellfoundedness property we are looking for. We require two rather obvious properties of the syntactic operation introduced in Definition 7.1.2 which maps a formula A of \mathcal{L}_μ^+ into a formula A^- by omitting all fixed point iteration numbers. The first property is that the said operation does not decrease the length of A or the rank of A with respect to the ordering \preceq . This is shown by a simple induction on the structure of A . The second property establishes the crucial relationship between the strong closure and the Fischer-Ladner closure of a formula D and is easily shown by induction on the build-up the set $\mathbb{S}\mathbb{C}(D)$.

Lemma 7.3.10. *For all formulae A of \mathcal{L}_μ^+ we have*

$$rk(A) \preceq rk(A^-) \quad \text{and} \quad lh(A) = lh(A^-)$$

Lemma 7.3.11. *Let D be a closed formula of \mathcal{L}_μ . Then for all formulae A of \mathcal{L}_μ^+ we have*

$$A \in \mathbb{SC}(D) \implies A^- \in \mathbb{FL}(D).$$

Using these two properties, we may now prove that the elements of the strong closure of a formula D have only finitely many different lengths and thus arrive at the claim that the set $\mathbb{SC}(D)$ is wellordered by $<_{lex}$.

Lemma 7.3.12. *The set $\{lh(A) : A \in \mathbb{SC}(D)\}$ where D is a closed formula of \mathcal{L}_μ is finite.*

Proof. By Lemmata 7.3.10 and 7.3.11 we have the following set inclusions

$$\begin{aligned} \{lh(A) : A \in \mathbb{SC}(D)\} &= \{lh(A^-) : A \in \mathbb{SC}(D)\} \\ &\subset \{lh(A^-) : A^- \in \mathbb{FL}(D)\} \subset \{lh(A) : A \in \mathbb{FL}(D)\} \end{aligned}$$

By Lemma 7.3.8 the claim now follows. \square

Theorem 7.3.13. *Let D be a closed formula of \mathcal{L}_μ . The restriction of the ordering relation $<_{lex}$ to the set $\{rk(A) : A \in \mathbb{SC}(D)\}$ is a wellordering.*

Proof. By Lemma 7.3.12 the formulae in $\mathbb{SC}(D)$ have only finitely many different lengths. Thus the lengths of their ranks are bounded by some natural number and so by Remark 7.3.2 the relation $<_{lex}$ wellorders these ranks. \square

7.4 Semantics of \mathcal{L}_μ^+ under signatures

Until now we have used signatures to define a syntactic measure for formulae which has certain desirable properties with respect to unfolding greatest fixed points. We now make use of signatures a second time in order to define a semantical measure of formulae which behaves nicely with respect to unfolding least fixed points. More precisely, we are looking for a semantical way of measuring formulae, which decreases when moving from a formula of the form $(\mu X)\mathcal{A}$ to its unfolding $\mathcal{A}[(\mu X)\mathcal{A}]$. Our approach follows that of Streett and Emerson [32] and works by assigning a formula A its signed denotation $\|A\|_{\mathbf{K}}^\sigma$ in a Kripke structure \mathbf{K} where $\sigma = \langle \sigma_1, \dots, \sigma_m \rangle$ is a signature of appropriate length. The signed denotation of A semantically assigns each subformula of A of the form $(\mu X)\mathcal{A}$ its σ_i -th approximation, where i is the number of nested μ -quantifiers occurring in $(\mu X)\mathcal{A}$. We thus first define a function for measuring the degree of μ -nesting of a formula and collect some useful properties of this function in a subsequent lemma.

Definition 7.4.1 (μ -height). The μ -height $h_\mu(A)$ of a formula A of \mathcal{L}_μ^+ is defined by induction on the structure of A .

1. If $A \in \Phi \cup \mathbf{V} \cup \mathbf{T}$, then $h_\mu(A) = 0$.
2. If A is of the form $B \wedge C$ or $B \vee C$, then

$$h_\mu(A) := \max(h_\mu(B), h_\mu(C)).$$

3. If A is of the form $\Box_i B$ or $\Diamond_i B$, then

$$h_\mu(A) := h_\mu(B).$$

4. If A is of the form $(\mu X)\mathcal{B}$, then

$$h_\mu(A) := h_\mu(\mathcal{B}) + 1$$

5. If A is of the form $(\nu X)\mathcal{A}$ or $(\nu^n X)\mathcal{A}$ for any $n > 0$, then

$$h_\mu(A) := h_\mu(\mathcal{B}).$$

Lemma 7.4.2.

1. For all formulae A of \mathcal{L}_μ^+ we have $h_\mu(A) = h_\mu(A^-)$ and $h_\mu(A) < lh(A)$.
2. If D is a formula of \mathcal{L}_μ , then the set $\{h_\mu(A) : A \in \mathbb{S}\mathbb{C}(D)\}$ is finite.

Proof. The first assertion follows from Definition 7.4.1 and the second is an immediate consequence of the first and Lemma 7.3.12. \square

By Lemma 7.4.2 we know that the set of all μ -heights of formulae of $\mathbb{S}\mathbb{C}(D)$ has a least upper bound which we call the μ -bound of D .

Definition 7.4.3 (μ -bound). The μ -bound $b_\mu(D)$ of a closed formula D of \mathcal{L}_μ is the least natural number n such that $h_\mu(A) \leq n$ for all $A \in \mathbb{S}\mathbb{C}(D)$.

We are now ready to define the signed denotation by induction on $<_{lex}$. In view of Remark 7.3.2 and Theorem 7.3.13 this notion is defined for formulae occurring in the strong closure $\mathbb{S}\mathbb{C}(D)$ of a formula D of \mathcal{L}_μ only. It turns out that working in such a restricted set of formulae is nevertheless sufficient for the purpose of our argument. Simultaneously to the signed denotation, we also define signed approximations of least and greatest fixed points.

Definition 7.4.4 (Signed denotation). Let $\mathbf{K} = (S, R, \pi)$ be a Kripke structure, D a formula of \mathcal{L}_μ and $n = b_\mu(D)$. For every $A \in \mathbb{S}\mathbb{C}(D)$ and signature $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ we define the set $\|A\|_{\mathbf{K}}^\sigma \subset S$ by induction on $<_{lex}$ as follows:

1. For atomic, propositional and modal formulae we proceed as before.

$$\begin{aligned} \|P\|_{\mathbf{K}}^\sigma &:= \pi(P) \text{ for all } P \in \Phi \cup \mathbb{V}, & \|\top\|_{\mathbf{K}}^\sigma &:= S, & \|\perp\|_{\mathbf{K}}^\sigma &:= \emptyset, \\ \|B \wedge C\|_{\mathbf{K}}^\sigma &:= \|B\|_{\mathbf{K}}^\sigma \cap \|C\|_{\mathbf{K}}^\sigma, & \|B \vee C\|_{\mathbf{K}}^\sigma &:= \|B\|_{\mathbf{K}}^\sigma \cup \|C\|_{\mathbf{K}}^\sigma, \\ \|\Box_i B\|_{\mathbf{K}}^\sigma &:= \{w \in S : v \in \|B\|_{\mathbf{K}}^\sigma \text{ for all } v \text{ such that } wR_i v\}, \\ \|\Diamond_i B\|_{\mathbf{K}}^\sigma &:= \{w \in S : v \in \|B\|_{\mathbf{K}}^\sigma \text{ for some } v \text{ such that } wR_i v\}. \end{aligned}$$

2. For each \mathbf{X} -positive formula \mathcal{A} we note that the operator $F_{\mathcal{A}, \mathbf{X}}^{\mathbf{K}, \sigma}$ defined as $F_{\mathcal{A}, \mathbf{X}}^{\mathbf{K}, \sigma}(T) := \|\mathcal{A}\|_{\mathbf{K}[\mathbf{X}:=T]}^\sigma$ for every $T \subset S$ is monotone. For every ordinal α we thus set

$$I_{\mathcal{A}, \mathbf{X}, \mathbf{K}, \sigma}^{<\alpha} := I_{F_{\mathcal{A}, \mathbf{X}}^{\mathbf{K}, \sigma}}^{<\alpha}, \quad J_{\mathcal{A}, \mathbf{X}, \mathbf{K}, \sigma}^{<\alpha} := J_{F_{\mathcal{A}, \mathbf{X}}^{\mathbf{K}, \sigma}}^{<\alpha}.$$

3. For least fixed point formulae $(\mu\mathbf{X})\mathcal{A}$ we define

$$\|(\mu\mathbf{X})\mathcal{A}\|_{\mathbf{K}}^\sigma := I_{\mathcal{A}, \mathbf{X}, \mathbf{K}, \sigma}^{<\sigma_m}$$

where $m = h_\mu((\mu\mathbf{X})\mathcal{A})$.

4. For greatest fixed point formulae and their finite approximations we set

$$\begin{aligned} \|(\nu^1\mathbf{X})\mathcal{A}\|_{\mathbf{K}}^\sigma &:= \|\mathcal{A}[\top/\mathbf{X}]\|_{\mathbf{K}}^\sigma, \\ \|(\nu^{n+1}\mathbf{X})\mathcal{A}\|_{\mathbf{K}}^\sigma &:= \|\mathcal{A}[(\nu^n\mathbf{X})\mathcal{A}]\|_{\mathbf{K}}^\sigma, \\ \|(\nu\mathbf{X})\mathcal{A}\|_{\mathbf{K}}^\sigma &:= \bigcap_{i < \omega} \|(\nu^i\mathbf{X})\mathcal{A}\|_{\mathbf{K}}^\sigma. \end{aligned}$$

We continue by investigating some important properties of the signed denotation. To begin with, we note that the signed denotation $\|\mathcal{A}\|_{\mathbf{K}}^\sigma$ of an \mathbf{X} -positive formula \mathcal{A} does not behave in a compositional way with respect to positive substitution. That is to say, the set $\|\mathcal{A}[B/\mathbf{X}]\|_{\mathbf{K}}^\sigma$ is in general different from the set $\|\mathcal{A}\|_{\mathbf{K}[\mathbf{X}:=S]}^\sigma$ where $S = \|B\|_{\mathbf{K}}^\sigma$, since the μ -height of $\mathcal{A}[B/\mathbf{X}]$ is generally greater than that of \mathcal{A} . However, in the special case where all components of σ are identical syntactic and semantic substitution do commute in the way described above.

Lemma 7.4.5. *Let D be a closed formula of \mathcal{L}_μ , B some formula of \mathcal{L}_μ^+ and \mathcal{A} an \mathbf{X} -positive formula of $\mathbb{S}\mathbb{C}(D)$ such that $\mathcal{A}[B/\mathbf{X}]$ also belongs to $\mathbb{S}\mathbb{C}(D)$. For all ordinals σ , signatures $\boldsymbol{\sigma} = \langle \sigma, \dots, \sigma \rangle$ of length $b_\mu(D)$, Kripke structures $\mathbb{K} = (S, R, \pi)$ and subsets T of S we have*

$$S = \|B\|_{\mathbb{K}}^{\boldsymbol{\sigma}} \implies \|\mathcal{A}[B/\mathbf{X}]\|_{\mathbb{K}}^{\boldsymbol{\sigma}} = \|\mathcal{A}\|_{\mathbb{K}[X:=S]}^{\boldsymbol{\sigma}}.$$

This lemma is shown by a straightforward induction on the $<_{lex}$ ordering. The next thing we note is that in general $\|A\|_{\mathbb{K}}^{\boldsymbol{\sigma}}$ is a proper subset of the real denotation $\|A\|_{\mathbb{K}}$ since when considering the signed denotation of a formula, a least fixed point subformula $(\mu\mathbf{X})\mathcal{A}$ is interpreted by an approximation $I_{\mathcal{A},\mathbf{X},\mathbb{K},\boldsymbol{\sigma}}^{<\alpha}$, thus in general by a subset of the fixed point itself. Similarly, greatest fixed point subformulae are interpreted by the intersection over the signed denotations of all of their finite approximations, thus in general by a superset of the fixed point itself. The next lemma asserts that the notion of signed denotation, nevertheless, coincides with that of the plain denotation, if a sufficiently large signature is chosen. This is mainly due to the fact that the number of iterations required for a monotone operator to reach its least fixed point is essentially bounded by the cardinality of the underlying Kripke structure (see Theorem 1.2.1).

Lemma 7.4.6. *Let us assume that*

- (1) D is a closed formula of \mathcal{L}_μ ,
- (2) $\mathbb{K} = (S, R, \pi)$ is a Kripke structure,
- (3) κ is the least cardinal greater than the cardinality of S ,
- (4) $\boldsymbol{\kappa}$ is the signature $\langle \kappa, \dots, \kappa \rangle$ of length $b_\mu(D)$.

Then for all formulae A of $\mathbb{S}\mathbb{C}(D)$

$$\|A\|_{\mathbb{K}} \subset \|A\|_{\mathbb{K}}^{\boldsymbol{\kappa}}.$$

Proof. We proceed by induction on $rk(A)$ and distinguish the following cases:

1. If $A \in \Phi \cup \mathbf{V} \cup \mathbf{T}$, then the assertion is obvious.
2. If A is a disjunction, a conjunction, a modal formula or a formula $(\nu^n\mathbf{X})\mathcal{A}$ for some natural number $n > 0$, then the assertion follows directly from the induction hypothesis and Definitions 7.2.2 and 7.4.4.

3. If A is a formula of the form $(\mu X)\mathcal{A}$ or $(\nu X)\mathcal{A}$, then we first note that by induction hypothesis $F_{\mathcal{A},X}^K(T) \subset F_{\mathcal{A},X}^{K,\kappa}(T)$ for all subsets T of S . Therefore, we have $I_{\mathcal{A},X,K} \subset I_{\mathcal{A},X,K,\kappa}$ and $J_{\mathcal{A},X,K} \subset J_{\mathcal{A},X,K,\kappa}$ and so

$$\|(\mu X)\mathcal{A}\|_K = I_{\mathcal{A},X,K} \subset I_{\mathcal{A},X,K,\kappa} = I_{\mathcal{A},X,K,\kappa}^{<\kappa} = \|(\mu X)\mathcal{A}\|_K^\kappa \quad (7.7)$$

$$\|(\nu X)\mathcal{A}\|_K = J_{\mathcal{A},X,K} \subset J_{\mathcal{A},X,K,\kappa} \subset J_{\mathcal{A},X,K,\kappa}^{<\omega} = \|(\nu X)\mathcal{A}\|_K^\kappa \quad (7.8)$$

easily follow from Definitions 7.2.2 and 7.4.4. For the last equality in (7.8) we furthermore use Lemma 7.4.5. This completes the proof. \square

Signed denotations have a very important property which we address next. Considering a least fixed point formula $(\mu X)\mathcal{A}$ and its unfolding $\mathcal{A}[(\mu X)\mathcal{A}]$, we will see that given a signature σ there exists a lexicographically smaller signature τ such that $\|(\mu X)\mathcal{A}\|_K^\sigma \subset \|\mathcal{A}[(\mu X)\mathcal{A}]\|_K^\tau$. This property will be of crucial use to us when considering signed denotations in an induction along signatures of some bounded length. For technical reasons we prove a slightly more general assertion.

Lemma 7.4.7. *Let us assume that*

- (1) D is a closed formula of \mathcal{L}_μ and \mathcal{A} and \mathcal{B} are X -positive formulae of $\mathbb{S}\mathbb{C}(D)$,
- (2) $(\mu X)\mathcal{A}$ and $\mathcal{B}[(\mu X)\mathcal{A}]$ belong to $\mathbb{S}\mathbb{C}(D)$,
- (3) $h_\mu((\mu X)\mathcal{A}) = m + 1$ and $h_\mu(\mathcal{B}) \leq m$,
- (4) $K = (S, R, \pi)$ is a Kripke structure,
- (5) κ is the least cardinal greater than the cardinality of S ,
- (6) σ is a signature $\langle \sigma_1, \dots, \sigma_n \rangle$ of length $n = b_\mu(D)$,
- (7) τ is the signature $\langle \sigma_1, \dots, \sigma_m, \alpha, \kappa, \dots, \kappa \rangle$ of length n .

Then we have

$$\|\mathcal{B}\|_{K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^\sigma \subset \|\mathcal{B}[(\mu X)\mathcal{A}]\|_K^\tau.$$

Proof. The assertion is shown by induction on $rk(\mathcal{B})$ making the following case distinction:

1. If U does not occur in \mathcal{B} , then trivially

$$\|\mathcal{B}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{\sigma} = \|\mathcal{B}\|_{\mathbb{K}}^{\sigma} \quad \text{and} \quad \|\mathcal{B}\|_{\mathbb{K}}^{\tau} = \|\mathcal{B}[(\mu X)\mathcal{A}]\|_{\mathbb{K}}^{\tau}$$

From the fact that $h_{\mu}(\mathcal{B}) \leq m$ we further obtain $\|\mathcal{B}\|_{\mathbb{K}}^{\sigma} = \|\mathcal{B}\|_{\mathbb{K}}^{\tau}$ which completes the discussion of this case.

2. If \mathcal{B} is the formula X , then since $h_{\mu}(\mathcal{A}) \leq m$ we have $F_{\mathcal{A},X}^{K,\sigma}(T) = F_{\mathcal{A},X}^{K,\tau}(T)$ for all $T \subset S$ and so

$$\|\mathcal{B}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{\sigma} = I_{\mathcal{A},X,K,\sigma}^{<\alpha} = I_{\mathcal{A},X,K,\tau}^{<\alpha} = \|(\mu X)\mathcal{A}\|_{\mathbb{K}}^{\tau} = \|\mathcal{B}[(\mu X)\mathcal{A}]\|_{\mathbb{K}}^{\tau}.$$

3. If \mathcal{B} is a conjunction, a disjunction, a modal formula, a formula $(\nu X)\mathcal{A}$ or a formula $(\nu^k X)\mathcal{A}$ for some natural number $n > 0$, then the assertion follows directly from the induction hypothesis.
4. If \mathcal{B} is a formula of the form $(\mu Y)\mathcal{C}$, then we first claim that

$$I_{\mathcal{C},Y,K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{<\xi} \subset I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\xi} \quad (7.9)$$

for all ordinals ξ . This claim is shown by side induction on ξ . If we have $w \in I_{\mathcal{C},Y,K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{<\xi}$, then

$$\begin{aligned} w &\in F_{\mathcal{C},Y}^{K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}],\sigma}(I_{\mathcal{C},Y,K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{<\zeta}) \\ &= \|\mathcal{C}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}][Y:=I_{\mathcal{C},Y,K[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}]}^{<\zeta}]}^{\sigma} \end{aligned}$$

for some $\zeta < \xi$. Now by hypothesis of the side induction

$$w \in \|\mathcal{C}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{<\alpha}][Y:=I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\zeta}]}^{\sigma}.$$

We know that the Kripke structures

$$\mathbb{K}[X := I_{\mathcal{A},X,K,\sigma}^{<\alpha}][Y := I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\zeta}]$$

and

$$\mathbb{K}[Y := I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\zeta}][X := I_{\mathcal{A},X,K,\sigma}^{<\alpha}]$$

are identical. Since $rk(\mathcal{C}) <_{lex} rk(\mathcal{B})$ we may thus apply the hypothesis of the main induction and infer that

$$\begin{aligned} w &\in \|\mathcal{C}[(\mu X)\mathcal{A}]\|_{\mathbb{K}[Y:=I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\zeta}]}^{\tau} \\ &= F_{\mathcal{C}[(\mu X)\mathcal{A}],Y}^{K,\tau}(I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\zeta}) \subset I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{<\xi} \end{aligned}$$

and thus (7.9) is established. From assumption (3) we obtain that $h_\mu(\mathcal{B}) = k$ for some $k \leq m$ and thus that

$$\|\mathcal{B}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{\leq\alpha}]}^\sigma = I_{\mathcal{C},Y,\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{\leq\alpha}],\sigma}^{\leq\sigma_k}$$

Combining this with (7.9) and assumption (5), we obtain

$$\|\mathcal{B}\|_{\mathbb{K}[X:=I_{\mathcal{A},X,K,\sigma}^{\leq\alpha}]}^\sigma \subset I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{\leq\sigma_k} \subset I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{\leq\kappa} \quad (7.10)$$

Since X really occurs in \mathcal{B} we have $m + 1 < h_\mu(\mathcal{B}(\mu X)\mathcal{A})$ and therefore

$$\|\mathcal{B}[(\mu X)\mathcal{A}]\|_{\mathbb{K}}^\tau = \|(\mu Y)\mathcal{C}[(\mu X)\mathcal{A}]\|_{\mathbb{K}}^\tau = I_{\mathcal{C}[(\mu X)\mathcal{A}],Y,K,\tau}^{\leq\kappa} \quad (7.11)$$

By (7.10) and (7.11) the assertion follows for this case.

Thus we have shown the assertion in all cases and the proof is done. \square

Chapter 8

The infinitary systems T_μ^ω and $\mathsf{T}_{\mu+}^\omega$

We now introduce two infinitary, cut-free deductive systems for the modal μ -calculus and show that both are sound and complete for formulae of the language \mathcal{L}_μ . After formally presenting the two systems in the first section of this chapter, the second section will examine the notion of D -saturated sequents which will play an important role in showing that the two systems are complete. Following the completeness proof in the third section, we will show in the fourth section how the method used in Chapter 5 can be adapted to provide a finitary cut-free system also in the case of the μ -calculus. The fifth section will elaborate on the relationship between SFL and the μ -calculus, showing that the former may be embedded into the latter but not vice versa.

8.1 The systems T_μ^ω and $\mathsf{T}_{\mu+}^\omega$

The system $\mathsf{T}_{\mu+}^\omega$ is designed to provide a notion of provability for formulae of the extended language \mathcal{L}_μ^+ making sure that the explicit finite iterations of greatest fixed point formulae are axiomatised correctly. The system T_μ^ω – which we are ultimately interested in – differs from $\mathsf{T}_{\mu+}^\omega$ merely by the omission of these rules. Since for formulae of the language \mathcal{L}_μ provability in the extended system $\mathsf{T}_{\mu+}^\omega$ implies provability in the restricted system T_μ^ω a completeness proof for the former will also imply completeness of the latter with respect to the language \mathcal{L}_μ .

Both $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω are given in Tait style. In this context, we will adopt the same terminology as in Section 4.1, thus speaking of sequents of \mathcal{L}_μ or \mathcal{L}_μ^+

and adhering to the same notational conventions as introduced in Definition 4.1.1. For the statement of the rules of $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω we adopt the same naming conventions as before: a name is given in brackets to the right of every rule, ignoring the fact that each such name would need to be parametrised by the sequents and formulae involved in the rule. From a superficial point of view the systems $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω look very similar to $\mathsf{T}_{\text{SFL}}^\omega$ presented for SFL. However, the far richer syntax of the modal μ -calculus will lead to a much more intricate completeness argument.

Definition 8.1.1 (The system $\mathsf{T}_{\mu+}^\omega$). The system $\mathsf{T}_{\mu+}^\omega$ is defined by the following inference rules:

Axioms: For all sequents Γ of \mathcal{L}_μ^+ and \mathfrak{p} in Φ

$$\frac{}{\Gamma, \mathfrak{p}, \sim \mathfrak{p}} \text{ (ID1)}, \quad \frac{}{\Gamma, \mathsf{X}, \sim \mathsf{X}} \text{ (ID2)}, \quad \frac{}{\Gamma, \top} \text{ (ID3)}.$$

Propositional rules: For all sequents Γ and formulae A and B of \mathcal{L}_μ^+

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

Modal rules: For all sequents Γ and Σ and formulae A of \mathcal{L}_μ^+ and all indices i from M

$$\frac{\Gamma, A}{\diamond_i \Gamma, \square_i A, \Sigma} \text{ (}\square\text{)}$$

Approximation rules: For all sequents Γ and X -positive formulae \mathcal{A} of \mathcal{L}_μ^+ and all natural numbers $k > 0$

$$\frac{\Gamma, \mathcal{A}[\top/\mathsf{X}]}{\Gamma, (\nu^1 \mathsf{X})\mathcal{A}} \text{ (}\nu.1\text{)} \quad \frac{\Gamma, \mathcal{A}[(\nu^k \mathsf{X})\mathcal{A}]}{\Gamma, (\nu^{k+1} \mathsf{X})\mathcal{A}} \text{ (}\nu.k+1\text{)}$$

Fixed point rules: For all sequents Γ and X -positive formulae \mathcal{A} of \mathcal{L}_μ^+

$$\frac{\Gamma, \mathcal{A}[(\mu \mathsf{X})\mathcal{A}]}{\Gamma, (\mu \mathsf{X})\mathcal{A}} \text{ (}\mu\text{)} \quad \frac{\Gamma, (\nu^k \mathsf{X})\mathcal{A} \text{ for all } k \in \omega}{\Gamma, (\nu \mathsf{X})\mathcal{A}} \text{ (}\nu.\omega\text{)}$$

Definition 8.1.2 (The system T_μ^ω). The system T_μ^ω is defined by the following inference rules:

Axioms: For all sequents Γ of \mathcal{L}_μ and \mathfrak{p} in Φ

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

Propositional rules: For all sequents Γ and formulae A and B of \mathcal{L}_μ

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

Modal rules: For all sequents Γ and Σ and formulae A of \mathcal{L}_μ and all indices i from M

$$\frac{\Gamma, A}{\diamond_i \Gamma, \square_i A, \Sigma} \quad (\square)$$

Fixed point rules: For all sequents Γ and X -positive formulae \mathcal{A} of \mathcal{L}_μ

$$\frac{\Gamma, \mathcal{A}[(\mu\mathsf{X})\mathcal{A}]}{\Gamma, (\mu\mathsf{X})\mathcal{A}} \quad (\mu) \qquad \frac{\Gamma, (\nu\mathsf{X})^k \mathcal{A} \quad \text{for all } k \in \omega}{\Gamma, (\nu\mathsf{X})\mathcal{A}} \quad (\nu)$$

Both $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω can be extended by the usual cut rule in an way analogous to Definition 4.1.2. We will, however, refrain from doing so, since none of the subsequent arguments depend on such a rule. Indeed, both systems will turn out to be complete without cut. Provability of a sequent Γ in $\mathsf{T}_{\mu+}^\omega$ or T_μ^ω is defined as usual and we again obtain a weakening lemma.

Definition 8.1.3 (Provability). Assume Γ is a sequent of \mathcal{L}_μ^+ and α an ordinal. We define the provability of Γ in $\mathsf{T}_{\mu+}^\omega$ in α many steps, denoted by $\mathsf{T}_{\mu+}^\omega \vdash^\alpha \Gamma$, by induction as follows:

1. If Γ is obtained by one of the axioms of $\mathsf{T}_{\mu+}^\omega$, then $\mathsf{T}_{\mu+}^\omega \vdash^\beta \Gamma$ holds for all ordinals β .
2. If Γ is obtained by one of the propositional, modal, approximation or fixed point rules where Γ_i are the premises of the respective rule, $\mathsf{T}_{\mu+}^\omega \vdash^{\beta_i} \Gamma_i$ holds for all of these premises and β is an ordinal such that $\beta_i < \beta$ for all β_i , then $\mathsf{T}_{\mu+}^\omega \vdash^\beta \Gamma$.

The notion of provability of a sequent Γ of \mathcal{L}_μ in T_μ^ω in α many steps, denoted by $\mathsf{T}_\mu^\omega \vdash^\alpha \Gamma$, is obtained by an induction analogous to the above one not including the approximation rules. Furthermore, we say a suitable Γ is provable and write $\mathsf{T}_{\mu+}^\omega \vdash \Gamma$ or $\mathsf{T}_\mu^\omega \vdash \Gamma$ if there exists an ordinal β such that Γ is provable in the respective system in β many steps. Finally, we write $\mathsf{T}_{\mu+}^\omega \not\vdash \Gamma$ or $\mathsf{T}_\mu^\omega \not\vdash \Gamma$ if Γ is not provable in the respective system.

Lemma 8.1.4 (Weakening). For all sequents Γ and Δ of \mathcal{L}_μ^+ , sequents Σ and Λ of \mathcal{L}_μ and ordinals β we have

1. If $\mathsf{T}_{\mu+}^\omega \vdash^\beta \Gamma$ and $\Gamma \subset \Delta$, then $\mathsf{T}_{\mu+}^\omega \vdash^\beta \Delta$.

2. If $\mathsf{T}_\mu^\omega \mid^\beta \Sigma$ and $\Sigma \subset \Lambda$, then $\mathsf{T}_\mu^\omega \mid^\beta \Lambda$.

The proof of the weakening lemma is a straightforward induction on the proof length β . Perhaps more importantly, provability in the two systems is related in a very crucial way. Whenever a sequent of \mathcal{L}_μ^+ is provable in $\mathsf{T}_{\mu+}^\omega$, then its standard translation to a sequent of \mathcal{L}_μ is provable T_μ^ω .

Theorem 8.1.5. *For all sequents Γ of \mathcal{L}_μ^+ we have*

$$\mathsf{T}_{\mu+}^\omega \vdash \Gamma \implies \mathsf{T}_\mu^\omega \vdash \Gamma^\circ$$

Proof. The claim is shown by induction on the proof of Γ in $\mathsf{T}_{\mu+}^\omega$. The only non-trivial cases are dealt with by observing that all applications of the rules $(\nu.1)$ and $(\nu.k+1)$ trivialise in view of the translation $^\circ$ and that $(\nu.\omega)$ goes over into (ν) . \square

As in the case of the infinitary system $\mathsf{T}_{\text{SFL}}^\omega$, the soundness of $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω is not obvious. Again, the greatest fixed point rules $(\nu.\omega)$ and (ν) look too strong at first glance, requiring only the finite iterations of a monotone operator to be valid in order to arrive at the conclusion that the greatest fixed point itself is also valid. However, similarly to the situation of SFL , we will later prove the soundness of a finitary system which contains T_μ^ω , yielding its soundness together with that of $\mathsf{T}_{\mu+}^\omega$.

8.2 D -saturated sequents

The technique used to prove the completeness of $\mathsf{T}_{\mu+}^\omega$ will again be that of saturated sequents as used by Alberucci and Jäger in [1]. However, the first step of showing that any sequent not provable in $\mathsf{T}_{\mu+}^\omega$ has a saturated sequent extending it will prove more challenging since our argument may not rely on any sort of level structure as in the case of SFL . Furthermore, in order to have the tool of induction on formula rank at our disposition we must ensure that all formulae involved in the argument occur in the strong closure of the distinguished non-provable formula to which a countermodel is constructed. We therefore introduce a notion of saturation which is parameterised by a particular closed formula D of \mathcal{L}_μ .

Definition 8.2.1 (D -saturated sequent). Let D be a closed formula of \mathcal{L}_μ . A sequent Γ of $\text{SC}(D)$ is called D -saturated (with respect to $\mathsf{T}_{\mu+}^\omega$) if all of the following conditions are satisfied:

(S.1) $\mathsf{T}_{\mu+}^\omega \not\vdash \Gamma$.

(S.2) For all formulae A and B of \mathcal{L}_μ^+ we have

$$\begin{aligned} A \vee B \in \Gamma &\implies A \in \Gamma \quad \text{and} \quad B \in \Gamma, \\ A \wedge B \in \Gamma &\implies A \in \Gamma \quad \text{or} \quad B \in \Gamma. \end{aligned}$$

(S.3) For all X -positive formulae \mathcal{A} of \mathcal{L}_μ^+ and all natural numbers $n > 0$ we have

$$\begin{aligned} (\mu X)\mathcal{A} \in \Gamma &\implies \mathcal{A}[(\mu X)\mathcal{A}] \in \Gamma, \\ (\nu X)\mathcal{A} \in \Gamma &\implies (\nu^i X)\mathcal{A} \in \Gamma \text{ for some natural number } i > 0, \\ (\nu^{n+1}X)\mathcal{A} \in \Gamma &\implies \mathcal{A}[(\nu^n X)\mathcal{A}] \in \Gamma, \\ (\nu^1 X)\mathcal{A} \in \Gamma &\implies \mathcal{A}[\top/X] \in \Gamma. \end{aligned}$$

We now set about showing that given a closed formula D of \mathcal{L}_μ any non-provable sequent consisting only of formulae from $\mathbb{S}\mathbb{C}(D)$ may be extended to a D -saturated sequent which also only contains formulae from $\mathbb{S}\mathbb{C}(D)$. As before, starting from a non-provable sequent we chose an iterative approach, repeatedly selecting a formula which violates one of the conditions (S.2) or (S.3) and adding suitable formulae to the sequent in order to make the respective condition satisfied. Seeing that this process becomes stable after a finite number of iterations then finishes the proof. Similar to the case of Lemma 4.2.7, problems arise when we encounter a least fixed point formula, say of the form $(\mu X)\mathcal{A}$ which violates condition (S.3) and for which we must thus add $\mathcal{A}[(\mu X)\mathcal{A}]$. Since this latter formula may itself violate one of the saturation conditions and in general has a greater rank than $(\mu X)\mathcal{A}$ the overall rank of violating formulae does not decrease during this step and termination is not guaranteed. Whereas in the case of **SFL** unfolding a least fixed point formula produces violating fixed point formulae from a lower level only, such an argument cannot be applied to the case of the μ -calculus, due to the presence of interleaved fixed points. Instead, we use the approach of a modified rank function, keeping a history of least fixed point formulae which have already been considered and ignoring these. Thus, for example, before treating a violating formula $(\mu X)\mathcal{A}$ the rank assigns to it a non-zero measure, since the formula is not yet in the history. When $(\mu X)\mathcal{A}$ is treated, it is added to the rank's history – formally speaking just a set of least fixed point formulae passed as a parameter to the rank function – and the fixed point is unfolded to $\mathcal{A}[(\mu X)\mathcal{A}]$. However, when measuring this unfolding the subformula $(\mu X)\mathcal{A}$ which is now in the history, is ignored by the rank function

thus leading to a decrease in rank. The main technical complication of this approach lies in ensuring that a least fixed point formula is not ignored too early in the whole process. This, however, can be taken care of by checking a suitable assertion after every step of the iteration.

Lemma 8.2.2. *Let D be a closed formula of \mathcal{L}_μ . For every sequent Γ of $\mathsf{SC}(D)$ which is not provable in $\mathsf{T}_{\mu+}^\omega$ there exists a sequent Δ of $\mathsf{SC}(D)$ which is D -saturated and $\Gamma \subset \Delta$.*

Proof. We begin by fixing an arbitrary enumeration F_1, F_2, F_3, \dots of the formulae in $\mathsf{SC}(D)$ calling the least i such that $A = F_i$ the index of A . Furthermore, we need to introduce some auxiliary notation:

Let N be a subset of $\mathsf{SC}(D)$. The N -rank $rk(N, A)$ of a formula $A \in \mathsf{SC}(D)$ is defined by an induction identical to that of Definition 7.3.3, adding the clause

$$A \in N \implies rk(N, A) := \langle 0 \rangle.$$

Analogously, we write $lh(N, A)$ for $lh(rk(N, A))$. Clearly, for all subsets N, N_1, N_2 of $\mathsf{SC}(D)$ and all $A \in \mathsf{SC}(D)$ this modified rank has the following properties:

$$lh(N, A) \leq lh(A) \tag{8.1}$$

$$A \in N \implies rk(N, \mathcal{B}[A/X]) = rk(N, \mathcal{B}[\top/X]) \tag{8.2}$$

$$N_1 \subset N_2 \implies rk(N_2, A) \leq_{lex} rk(N_1, A) \tag{8.3}$$

From (8.1) and Lemma 7.3.12 we obtain a strengthening of Lemma 7.3.13: the restriction of $<_{lex}$ to the set $\{rk(N, A) : A \in \mathsf{SC}(D) \text{ and } N \subset \mathsf{SC}(D)\}$ is also a wellordering. Given a subset N of $\mathsf{SC}(D)$ and a formula $A \in \mathsf{SC}(D)$ we therefore write $ot(N, A)$ for the order type of $rk(N, A)$ with respect to this wellordering.

Starting with our sequent Γ which is not provable in $\mathsf{T}_{\mu+}^\omega$, we inductively define a sequence of pairs $(\Gamma_0, M_0), (\Gamma_1, M_1), (\Gamma_2, M_2), \dots$ where for each natural number i the sequent Γ_i is not provable in $\mathsf{T}_{\mu+}^\omega$ and M_i is a set of formulae of $\mathsf{SC}(D)$.

1. $\Gamma_0 := \Gamma$ and $M_0 := \emptyset$.
2. If Γ_n is D -saturated, then $\Gamma_{n+1} := \Gamma_n$ and $M_{n+1} := M_n$.
3. If Γ_n is not D -saturated, we choose the formula A with least index which violates one of the conditions (S.2) or (S.3). Γ_{n+1} and M_{n+1} are now determined by case distinction on the form of A .

3.1. If A is of the form $B \vee C$, then we define

$$\Gamma_{n+1} := \Gamma_n \cup \{B, C\} \quad \text{and} \quad M_{n+1} := M_n.$$

3.2. If A is of the form $B \wedge C$, then, since Γ_n is not provable in $\mathbb{T}_{\mu+}^\omega$, we know that

$$\mathbb{T}_{\mu+}^\omega \not\vdash \Gamma_n, B \quad \text{or} \quad \mathbb{T}_{\mu+}^\omega \not\vdash \Gamma_n, C.$$

Accordingly, we define

$$\Gamma_{n+1} := \begin{cases} \Gamma_n \cup \{B\} & \text{if } \mathbb{T}_{\mu+}^\omega \not\vdash \Gamma_n, B \\ \Gamma_n \cup \{C\} & \text{otherwise} \end{cases} \quad \text{and} \quad M_{n+1} := M_n.$$

3.3. If A is of the form $(\mu X)\mathcal{B}$, then we define

$$\Gamma_{n+1} := \Gamma_n \cup \{\mathcal{B}[(\mu X)\mathcal{B}]\} \quad \text{and} \quad M_{n+1} := M_n \cup \{(\mu X)\mathcal{B}\}.$$

3.4. If A is of the form $(\nu X)\mathcal{B}$, then, since Γ_n is not provable in $\mathbb{T}_{\mu+}^\omega$, we know that

$$\mathbb{T}_{\mu+}^\omega \not\vdash (\nu^i X)\mathcal{B}$$

for some natural number i greater than 0. We choose the least such i and define

$$\Gamma_{n+1} := \Gamma_n \cup \{(\nu^i X)\mathcal{B}\} \quad \text{and} \quad M_{n+1} := M_n.$$

3.5. If A is of the form $(\nu^{i+1}X)\mathcal{B}$ for some number $i > 0$, then we define

$$\Gamma_{n+1} := \Gamma_n \cup \{\mathcal{B}[(\nu^i X)\mathcal{B}]\} \quad \text{and} \quad M_{n+1} := M_n.$$

3.6. If A is of the form $(\nu^1 X)\mathcal{B}$, then we define

$$\Gamma_{n+1} := \Gamma_n \cup \{\mathcal{B}[\top/X]\} \quad \text{and} \quad M_{n+1} := M_n.$$

It is easily verified during the above case distinction that for all natural numbers n we have

$$\mathbb{T}_{\mu+}^\omega \not\vdash \Gamma_n, \tag{8.4}$$

$$\Gamma \subset \Gamma_n \subset \Gamma_{n+1} \quad \text{and} \quad M_n \subset M_{n+1}, \tag{8.5}$$

$$(\mu X)\mathcal{B} \in M_n \implies \mathcal{B}[(\mu X)\mathcal{B}] \in \Gamma_n. \tag{8.6}$$

Next, we show two further crucial properties of the sequence just constructed.

(i) If a formula $(\mu\mathbf{X})\mathcal{B}$ belongs to M_{n+1} but not to M_n , then

$$ot(M_{n+1}, \mathcal{B}[(\mu\mathbf{X})\mathcal{B}]) < ot(M_n, (\mu\mathbf{X})\mathcal{B}).$$

(ii) If the formula picked at step $n + 1$ in the inductive definition above violates one of the conditions (2) or (3) and is of the form $(\mu\mathbf{X})\mathcal{B}$, then

$$ot(M_{n+1}, \mathcal{B}[(\mu\mathbf{X})\mathcal{B}]) < ot(M_n, (\mu\mathbf{X})\mathcal{B}).$$

To prove (i), assume that $(\mu\mathbf{X})\mathcal{B}$ is an element of $M_{n+1} \setminus M_n$. By (8.2) and (8.3) this implies

$$rk(M_{n+1}, \mathcal{B}[(\mu\mathbf{X})\mathcal{B}]) = rk(M_{n+1}, \mathcal{B}[\top/\mathbf{X}]) <_{lex} rk(M_n, \mathcal{B}[\top/\mathbf{X}]). \quad (8.7)$$

However, since $(\mu\mathbf{X})\mathcal{B} \notin M_n$ we also have

$$rk(M_n, \mathcal{B}[\top/\mathbf{X}]) <_{lex} rk(M_n, \mathcal{B}[\top/\mathbf{X}]) * \langle 0 \rangle = rk(M_n, (\mu\mathbf{X})\mathcal{B}). \quad (8.8)$$

Assertions (8.7) and (8.8) imply $rk(M_{n+1}, \mathcal{B}[(\mu\mathbf{X})\mathcal{B}]) <_{lex} rk(M_n, (\mu\mathbf{X})\mathcal{B})$ and so (i) is shown. Property (i) follows immediately from (i) together with (8.6). In a next step, we assign to all sequents Π of $\mathsf{SC}(D)$ which are not provable in $\mathsf{T}_{\mu+}^\omega$ and all finite subsets N of $\mathsf{SC}(D)$ a deficiency number $dn(N, A)$ as follows:

(D.1) If Π is saturated, then $dn(N, \Pi) := 0$.

(D.2) Otherwise,

$$dn(N, \Pi) := \omega^{ot(N, A_1)} \# \dots \# \omega^{ot(N, A_m)}$$

where A_1, \dots, A_m are the formulae in Π which violate one of the conditions (2) or (3) and $\#$ denotes the natural sum operation on ordinals as introduced by Schütte [30].

To conclude the proof, we note that (ii) together with (8.3) yields

$$\Gamma_n \text{ is not } D\text{-saturated} \implies dn(M_{n+1}, \Gamma_{n+1}) < dn(M_n, \Gamma_n).$$

Since, therefore, the deficiency number strictly decreases with every step of the sequence $(\Gamma_0, M_0), (\Gamma_1, M_1), (\Gamma_2, M_2), \dots$ where Γ_i is not saturated, there must exist an m such that $dn(M_m, \Gamma_m) = 0$, meaning that Γ_m is saturated and thus a candidate for Δ . \square

Based on the collection of all D -saturated sequents we now construct a canonical countermodel much in the same way as we have done in Definition 4.3.1. This construction is dependent on the formula D , the role of which will later be played by the non-provable formula used in the completeness argument.

Definition 8.2.3 (Canonical countermodel). Let D be a closed formula of \mathcal{L}_μ . Define the triple $\mathsf{K}_D = (S_D, R_D, \pi_D)$ as follows, where $i \in \mathsf{M}$:

$$\begin{aligned} S_D &:= \{\Gamma \subset \mathsf{SC}(D) : \Gamma \text{ } D\text{-saturated}\}, \\ R_D(i) &:= \{(\Gamma, \Delta) \in S_D \times S_D : \{B \in \mathsf{SC}(D) : \diamond_i B \in \Gamma\} \subset \Delta\}, \\ \pi_D(P) &:= \{\Gamma \in S_D : P \notin \Gamma\} \text{ for } P \in \Phi \cup \mathsf{V}. \end{aligned}$$

It is easily verified that K_D is a Kripke structure in the sense of Definition 2.2.1. Given a formula $A \in \mathsf{SC}(D)$ and a set $T \subset S_D$ we will write $\|A\|_D$ for $\|A\|_{\mathsf{K}_D}$ and $\|A\|_{D[x:=T]}$ for $\|A\|_{\mathsf{K}_D[x:=T]}$.

8.3 Completeness of $\mathsf{T}_{\mu+}^\omega$

Using the canonical countermodel construction the completeness of the system $\mathsf{T}_{\mu+}^\omega$ can be shown. The general pattern which we follow is that of showing that a formula A not provable in $\mathsf{T}_{\mu+}^\omega$ is not valid in the canonical countermodel and thus not valid in the more general sense. The canonical countermodel is built in such a way that if a saturated sequent Γ contains A , then A is not satisfied at Γ . We will require most of the machinery prepared so far in order to show this but we first note that the canonical countermodel construction treats modalities in a suitable way.

Lemma 8.3.1. *Assume D is a closed formula of \mathcal{L}_μ , i is an index from M , $\Box_i A$ and $\Diamond_i A$ are formulae in $\mathsf{SC}(D)$ and Γ is a D -saturated sequent. Then we have the following implications:*

1. *If $\Box_i A \in \Gamma$, then there exists some D -saturated sequent Δ such that $(\Gamma, \Delta) \in R_D(i)$ and $A \in \Delta$.*
2. *If $\Diamond_i A \in \Gamma$, then $A \in \Delta$ for all D -saturated sequents Δ such that $(\Gamma, \Delta) \in R_D(i)$.*

Proof.

Claim 1: Since Γ is D -saturated, it is not provable in $\mathsf{T}_{\mu+}^\omega$. By the rule (\Box) of $\mathsf{T}_{\mu+}^\omega$ we infer that

$$\mathsf{T}_{\mu+}^\omega \not\vdash \{C : \Diamond_i C \in \Gamma\}.$$

From Lemma 8.2.2 we obtain a D -saturated sequent Σ with the properties

$$\{C : \diamond_i C \in \Gamma\} \subset \Sigma, \quad (8.9)$$

$$B \in \Sigma. \quad (8.10)$$

By (8.9) we have $(\Gamma, \Sigma) \in R_D(i)$. Thus the claim is shown by setting $\Delta := \Sigma$.

Claim 2: Assume Δ is an arbitrary D -saturated sequent which satisfies $(\Gamma, \Delta) \in R_D(i)$. Therefore, $\{C : \diamond_i C \in \Gamma\} \subset \Delta$ and $B \in \Delta$. Thus the claim is shown. □

Employing both an induction on signatures of bounded length and an induction on formula rank, we arrive at a signed version of the truth lemma which thanks to Lemma 7.4.6 implies the truth lemma without signatures.

Lemma 8.3.2. *Let D be some closed formula of \mathcal{L}_μ and let $n := b_\mu(D)$. Then for all signatures σ of length less than or equal to n , all closed formulae A of $\mathsf{SC}(D)$ and all D -saturated sequents Γ we have*

$$A \in \Gamma \implies \Gamma \notin \|A\|_D^\sigma$$

Proof. We show this lemma by main induction on all signatures σ of length less than or equal to n and side induction on $rk(A)$. In doing so we distinguish the following cases:

1. If $A \in \Phi \cup \mathsf{T}$, then the assertion is trivially verified.
2. If A is a disjunction, a conjunction, a formula $(\nu^k \mathsf{X})\mathcal{A}$ for some natural number $k > 0$ or a formula $(\nu \mathsf{X})\mathcal{A}$, then the assertion follows directly by Lemma 7.3.4, Theorem 7.3.6 and the hypothesis of the side induction.
3. If A is a modal formula, then the assertion follows from Lemma 8.3.1 along with the side induction hypothesis.
4. If A is a formula of the form $(\mu \mathsf{X})\mathcal{A}$, then, since $A \in \Gamma$ and Γ is D -saturated, we also have

$$\mathcal{A}[(\mu \mathsf{X})\mathcal{A}] \in \Gamma. \quad (8.11)$$

We may assume that $h_\mu((\mu\mathsf{X})\mathcal{A}) = m + 1$ with $h_\mu(\mathcal{A}) \leq m$. Furthermore, let us assume that $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ for suitable ordinals $\sigma_1, \dots, \sigma_n$. According to Definition 7.4.4 we thus have

$$\|A\|_D^\sigma = \|(\mu\mathsf{X})\mathcal{A}\|_D^\sigma = I_{\mathcal{A}, \mathsf{X}, \mathsf{K}_D, \sigma}^{<\sigma_{m+1}}. \quad (8.12)$$

In order to establish our assertion, we assume that

$$\Gamma \in \|A\|_D^\sigma \quad (8.13)$$

and aim to arrive at a contradiction. In view of (8.12), there exists an ordinal $\alpha < \sigma_{m+1}$ so that

$$\Gamma \in \|\mathcal{A}\|_{D[\mathsf{X}:=I_{\mathcal{A}, \mathsf{X}, \mathsf{K}_D, \sigma}^{<\alpha}]}^\sigma. \quad (8.14)$$

We choose κ to be the least cardinal greater than the cardinality of K_D and define the signature $\tau := \langle \sigma_1, \dots, \sigma_m, \alpha, \kappa, \dots, \kappa \rangle$. By Lemma 7.4.7 we get

$$\Gamma \in \|\mathcal{A}[(\mu\mathsf{X})\mathcal{A}]\|_D^\tau. \quad (8.15)$$

On the other hand $\tau <_{lex} \sigma$ thus (8.11), together with the main induction hypothesis, yields

$$\Gamma \notin \|\mathcal{A}[(\mu\mathsf{X})\mathcal{A}]\|_D^\tau. \quad (8.16)$$

Since (8.15) and (8.16) present us with a contradiction, our assumption (8.13) must have been false. Hence Γ cannot be an element of $\|(\mu\mathsf{X})\mathcal{A}\|_D^\sigma$ and the assertion is shown in this case.

Thus we have treated all possible cases and our proof is complete. \square

Theorem 8.3.3. *Let D be a closed formula of \mathcal{L}_μ and A a closed formula of $\mathsf{SC}(D)$. Then for all D -saturated sequents Γ of $\mathsf{SC}(D)$ we have*

$$A \in \Gamma \implies \Gamma \notin \|A\|_D.$$

Proof. As before, we define κ to be the least cardinal greater than the cardinality of K_D and $\boldsymbol{\kappa}$ to be the signature $\langle \kappa, \dots, \kappa \rangle$ of length $b_\mu(D)$. Now, from $A \in \Gamma$ it follows that $\Gamma \notin \|A\|_D^{\boldsymbol{\kappa}}$ by Lemma 8.3.2. Thus, applying Lemma 7.4.6, we may deduce $\Gamma \notin \|A\|_D$ concluding the argument. \square

The truth lemma now immediately implies completeness of the systems $\mathsf{T}_{\mu+}^\omega$ and T_μ^ω which was the main goal of this chapter. While $\mathsf{T}_{\mu+}^\omega$ is technically more suitable in the completeness proof the system T_μ^ω is definitely more elegant as a final product.

Theorem 8.3.4 (Completeness of $\mathsf{T}_{\mu+}^\omega$). *For all closed formulae A of \mathcal{L}_μ we have that if A is valid, then $\mathsf{T}_{\mu+}^\omega \vdash A$.*

Proof. We show the contrapositive of the asserted implication and thus assume that A is not provable in $\mathsf{T}_{\mu+}^\omega$. Then by Lemma 8.2.2 there exists an A -saturated sequent Γ of $\mathsf{SC}(A)$ such that $A \in \Gamma$. Applying Theorem 8.3.3, we conclude that $\Gamma \notin \llbracket A \rrbracket_A$, meaning that A cannot be valid. This concludes the proof. \square

Corollary 8.3.5 (Completeness of T_μ^ω). *For all closed formulae A of \mathcal{L}_μ we have that if A is valid, then $\mathsf{T}_\mu^\omega \vdash A$.*

Proof. This assertion follows from Theorem 8.3.4 by Lemma 7.1.7 and Theorem 8.1.5. \square

8.4 Finitising T_μ^ω

As mentioned in the context of SFL and as shown for example in [32] or [10], the μ -calculus enjoys the small model property. For this reason, we may use the same technique as presented in Chapter 5 to obtain a truly finitary deductive system for the μ -calculus, which is cut-free like its infinitary counterpart. To avoid excessive repetition, we will confine ourselves to mentioning the important definition and result without going through all of the intermediate steps again, all of which could be achieved in an absolutely analogous way to their counterparts in Chapter 5. Let us, nevertheless, remind ourselves of the exact formulation of the small model property in the context of the μ -calculus.

Remark 8.4.1. There exists an exponential function $f : \omega \rightarrow \omega$ such that for every formula $A \in \mathcal{L}_\mu$ if A is satisfiable, then there exists a Kripke structure $\mathsf{K} = (S, R, \pi)$ with $|S| < f(|A|)$ which satisfies A .

Employing the bounding function f from Remark 8.4.1, we may reduce the number of premises needed in the greatest fixed point rule down to one and obtain the deductive system T_μ .

Definition 8.4.2 (The system T_μ). The system T_μ is defined by replacing the rule (ν) in the system T_μ^ω by the rule

$$\frac{\Gamma, (\nu X)^k \mathcal{A}}{\Gamma, (\nu X) \mathcal{A}, \Sigma} \quad (\nu. < \omega)$$

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

As in the case of SFL, there is a decisive proof-theoretical relationship between the infinitary system \mathbb{T}_μ^ω and the finitary \mathbb{T}_μ .

Theorem 8.4.3. *For all finite sets Γ of \mathcal{L}_μ and all ordinals α we have that*

$$\mathbb{T}_\mu^\omega \vdash^\alpha \Gamma \implies \mathbb{T}_\mu \vdash \Gamma.$$

The proof of this theorem is analogous to the one given for Theorem 5.1.3. As an immediate consequence we obtain the completeness of the system \mathbb{T}_μ .

Corollary 8.4.4 (Completeness of \mathbb{T}_μ). *For all closed formulae A of \mathcal{L}_μ we have that if A is valid, then $\mathbb{T}_\mu \vdash A$.*

Without further complications, adaptations of Lemmata 5.2.1 and 5.2.2 to the current setting can be shown. This leads us to the soundness of \mathbb{T}_μ .

Theorem 8.4.5 (Soundness of \mathbb{T}_μ). *The system \mathbb{T}_μ is sound, that is for all finite sets $\Gamma \subset \mathcal{L}_\mu$ if $\mathbb{T}_\mu \vdash \Gamma$, then the formula $\bigvee \Gamma$ is valid.*

Finally, this soundness result implies the soundness of both \mathbb{T}_μ^ω and $\mathbb{T}_{\mu+}^\omega$: For the former soundness follows by Theorem 8.4.3. The soundness of the latter is obtained by applying Theorem 8.1.5 again.

Corollary 8.4.6 (Soundness of \mathbb{T}_μ^ω and $\mathbb{T}_{\mu+}^\omega$). *The systems \mathbb{T}_μ^ω and $\mathbb{T}_{\mu+}^\omega$ are sound. That is, if a formula A of \mathcal{L}_μ is provable in \mathbb{T}_μ^ω or $\mathbb{T}_{\mu+}^\omega$, then A is valid.*

8.5 A note on expressivity

In Chapter 7 we mentioned that the μ -calculus contains SFL as a fragment. We can now make this intuition more formal by providing a straightforward translation of formulae of \mathcal{L}_{SFL} to formulae of \mathcal{L}_μ and stating the expected embedding result.

Definition 8.5.1 (Translation). Let A be a formula of \mathcal{L}_{SFL} . We define the translation A^* of A inductively on $rk(A)$.

1. If $A \in \Phi \cup \mathbb{V} \cup \mathbb{T}$, then $A^* := A$.
2. If A is a formula $B \wedge C$, then $A^* := B^* \wedge C^*$. Similarly, if A is a formula $B \vee C$, then $A^* := B^* \vee C^*$.
3. If A is a formula $\Box_i B$ where i is an index from \mathbb{M} , then $A^* := \Box_i B^*$. Similarly, if A is a formula $\Diamond_i B$, then $A^* := \Diamond_i B^*$.

4. If A is a constant P_B , then $A^* := (\mu X)\mathcal{B}^*$. Similarly, if A is a formula Q_B , then $A^* := (\nu X)\mathcal{B}^*$.

If $\Gamma = \{A_1, \dots, A_n\}$ is a sequent of \mathcal{L}_{SFL} , then we set $\Gamma^* := \{A_1^*, \dots, A_n^*\}$.

Before addressing the actual embedding theorem, we mention an important technical lemma which states that the translation given in Definition 8.5.1 commutes with positive substitution. The lemma can be shown by a straightforward induction on formula rank in the sense of Definition 2.3.1.

Lemma 8.5.2. *For all formulae \mathcal{A} and B of \mathcal{L}_{SFL} where \mathcal{A} is X -positive we have*

$$(\mathcal{A}[B])^* = \mathcal{A}^*[B^*/X].$$

Theorem 8.5.3. *For all sequents Γ of \mathcal{L}_{SFL} we have that if $\mathsf{T}_{\text{SFL}} \vdash \Gamma$ or $\mathsf{T}_{\text{SFL}}^\omega \vdash \Gamma$, then $\mathsf{T}_\mu \vdash \Gamma^*$, $\mathsf{T}_\mu^\omega \vdash \Gamma^*$ and $\mathsf{T}_{\mu+}^\omega \vdash \Gamma^*$.*

Proof. We will consider the case where we assume that $\mathsf{T}_{\text{SFL}}^\omega \vdash \Gamma$ and show that $\mathsf{T}_\mu^\omega \vdash \Gamma^*$. The other cases then all follow by completeness of $\mathsf{T}_{\text{SFL}}^\omega$ and soundness of T_{SFL} (Theorems 4.3.8 and 5.2.3) as well as by soundness of T_μ^ω and completeness of $\mathsf{T}_{\mu+}^\omega$ and T_μ (Corollary 8.4.6, Theorem 8.3.4 and Corollary 8.3.5). The proof proceeds by induction on the derivation of Γ in $\mathsf{T}_{\text{SFL}}^\omega$ where in view of Definition 8.5.1 the only non-trivial cases are those in which the last rule used to derive Γ was either (P) or (Q).

We first consider the case of the rule (P). Then $\Gamma = \Delta, P_A$ for some suitable sequent Δ and X -positive formula \mathcal{A} of \mathcal{L}_{SFL} and by assumption of the rule we have $\mathsf{T}_{\text{SFL}}^\omega \vdash \Delta, \mathcal{A}[P_A]$. By induction hypothesis obtain $\mathsf{T}_\mu^\omega \vdash \Delta^*, (\mathcal{A}[P_A])^*$ and Lemma 8.5.2 yields $\mathsf{T}_\mu^\omega \vdash \Delta^*, \mathcal{A}^*[(\mu X)\mathcal{A}^*]$. Therefore, using the rule (μ) of T_μ^ω we obtain $\mathsf{T}_\mu^\omega \vdash \Gamma^*$.

Thus we are left to consider the rule (Q). In this case we have $\Gamma = \Delta, Q_A$ for some suitable sequent Δ and X -positive formula \mathcal{A} of \mathcal{L}_{SFL} and by assumption of the rule $\mathsf{T}_{\text{SFL}}^\omega \vdash \Delta, Q_A^k$ for all natural numbers k . Then, by the induction hypothesis, we have $\mathsf{T}_\mu^\omega \vdash \Delta^*, (Q_A^k)^*$ for all natural numbers k . Furthermore, by Lemma 8.5.2 and a straightforward induction on m we get the syntactic equality $(Q_A^m)^* = (\nu X)^m \mathcal{A}^*$ for every natural number m . Thus $\mathsf{T}_\mu^\omega \vdash \Delta^*, (\nu X)^k \mathcal{A}^*$ for every natural number k and so by the rule (ν) we obtain $\mathsf{T}_\mu^\omega \vdash \Gamma^*$. \square

So far, we have shown that SFL is contained in the μ -calculus. However, we can do better by identifying a very conspicuous fragment of the μ -calculus to which SFL corresponds exactly, namely the fragment which consists of all formulae of \mathcal{L}_μ which are built using only the variable X .

Definition 8.5.4 (X-fragment of \mathcal{L}_μ). The X-fragment of \mathcal{L}_μ , denoted by \mathcal{L}_μ^X is the set of all formulae of \mathcal{L}_μ in which only the variable X (or no variable at all) occurs.

The translation defined in Definition 8.5.1 turns out to provide a one-to-one correspondence between the formulae of \mathcal{L}_{SFL} and those of \mathcal{L}_μ^X . This insight will finally ensure that SFL and the X-fragment of the μ -calculus can be viewed as being two different syntactic formulations of the same logic.

Lemma 8.5.5. *For every sequent Δ which consists only of formulae of \mathcal{L}_μ^X there exists a sequent Γ of \mathcal{L}_{SFL} such that $\Gamma^* = \Delta$.*

Proof. The assertion immediately follows if we can prove it for a single formula A of \mathcal{L}_μ^X . We proceed by induction on the structure of A and notice that all of the cases are trivial by Definition 8.5.1 except for the fixed point cases. Assume thus that A is of the form $(\mu X)\mathcal{A}$ for some X-positive formula \mathcal{A} of \mathcal{L}_μ^X . By induction hypothesis, there exists a formula \mathcal{B} of \mathcal{L}_{SFL} such that $\mathcal{B}^* = \mathcal{A}$. It is also clear from Definition 8.5.1 that \mathcal{B} must be X-positive, thus $P_{\mathcal{B}}$ is a formula of \mathcal{L}_{SFL} and $P_{\mathcal{B}}^* = (\mu X)\mathcal{B}^* = (\mu X)\mathcal{A} = A$. The greatest fixed point case is obtained in an analogous way. \square

Theorem 8.5.6. *SFL is equivalent to the μ -calculus restricted to formulae of \mathcal{L}_μ^X . That is, the following two statements hold:*

1. *For every sequent Γ of \mathcal{L}_{SFL} there exists a sequent Δ of \mathcal{L}_μ^X such that $T_\mu^\omega \vdash \bigvee \Gamma \leftrightarrow \bigvee \Delta$.*
2. *For every sequent Γ of \mathcal{L}_μ^X there exists a sequent Δ of \mathcal{L}_{SFL} such that $T_\mu^\omega \vdash \bigvee \Gamma \leftrightarrow \bigvee \Delta$.*

Proof. The first statement follows directly by Theorem 8.5.3, the second by Lemma 8.5.5. \square

An important consequence of Theorem 8.5.6 is that SFL is properly contained in the μ -calculus, that is to say that there are formulae of the latter which are not expressible in the former. This follows directly from the fact that the variable hierarchy of the μ -calculus is strict, as shown by Berwanger and Lenzi [8].

Corollary 8.5.7. *SFL is a proper fragment of the μ -calculus.*

Concluding remarks

This thesis started by introducing the logic **SFL** and investigating three notable fragments thereof. We then proceeded to give an infinitary cut-free axiomatisation $\mathsf{T}_{\mathsf{SFL}}^\omega$ of **SFL** and prove its completeness by showing that for any non-provable formula we may construct a countermodel, using the method of saturated sequents. With the help of the small model property of **SFL**, the infinitary system $\mathsf{T}_{\mathsf{SFL}}^\omega$ was subsequently turned into a finitary system $\mathsf{T}_{\mathsf{SFL}}$ for which we showed soundness. Due to the structure of the rules of $\mathsf{T}_{\mathsf{SFL}}^\omega$ and $\mathsf{T}_{\mathsf{SFL}}$, the soundness of $\mathsf{T}_{\mathsf{SFL}}^\omega$ is implied by that of $\mathsf{T}_{\mathsf{SFL}}$ and, conversely, the completeness of $\mathsf{T}_{\mathsf{SFL}}$ is implied by that of $\mathsf{T}_{\mathsf{SFL}}^\omega$. The part on **SFL** was concluded by investigating closure ordinals for valid fixed points of certain fragments of the language $\mathcal{L}_{\mathsf{SFL}}$. Valid fixed points from $\mathcal{L}_{\mathsf{SFL}}^1$ turned out to have closure ordinal ω , whereas the set of such fixed points from $\mathcal{L}_{\mathsf{SFL}}^2$ turned out to have no closure ordinal at all.

For the case of the μ -calculus the thesis followed essentially the same pattern as for **SFL**. We first introduced the two cut-free infinitary systems $\mathsf{T}_{\mu^+}^\omega$ and T_μ^ω , the first of which features an explicit axiomatisation of finite approximations of greatest fixed points and lent itself more conveniently to a completeness proof. In order to prove completeness of $\mathsf{T}_{\mu^+}^\omega$, the methods used for $\mathsf{T}_{\mathsf{SFL}}^\omega$ needed to be significantly generalised. The complexity of formulae of \mathcal{L}_μ^+ was measured in terms of signatures, which were shown to be wellordered on the strong closure $\mathsf{SC}(D)$ of a formula D of \mathcal{L}_μ^+ . This had the effect that notions like saturation needed to be parametrised by this formula D and extra care needed to be taken, that all formulae used in subsequent arguments were contained in $\mathsf{SC}(D)$. Completeness of the system T_μ^ω , which does not feature explicit finite iterations, is implied via a simple translation by that of $\mathsf{T}_{\mu^+}^\omega$. Unlike the method used for completeness, the argument used to obtain the finitisation T_μ of T_μ^ω needed no generalisation and was virtually identical to the one used for $\mathsf{T}_{\mathsf{SFL}}$. Again, the soundness of the infinitary systems is implied by the soundness of T_μ and the completeness of T_μ^ω implies the completeness of T_μ . Our study ended with an argument showing that **SFL** is a proper fragment of the μ -calculus.

Further work

While providing valuable proofs of concept as to the possibility of cut-free axiomatisations of SFL and the μ -calculus, the work presented in this thesis is far from being definitive. There are two main issues which could be addressed as topics of further research:

1. Obtaining the small model property syntactically instead of using it from outside.
2. Turning the cut-free axiomatisations presented in this thesis into proof-search procedures.

Issue 1 could be addressed by providing more careful canonical countermodel constructions than the ones given to prove Theorems 4.3.8 and 8.3.4, in such a way that the size of the countermodel for a non-provable formula A turns out to be finite and bounded by a function of the length of A . Once such a construction is achieved, we could reason as follows (considering the example of the μ -calculus): if A is a satisfiable formula of \mathcal{L}_μ , then $\neg A$ is not valid, thus by soundness of $\mathsf{T}_{\mu+}^\omega$ we have $\mathsf{T}_{\mu+}^\omega \not\vdash \neg A$ which by our new countermodel construction would mean that there exists a Kripke structure $\mathsf{K} = (S, R, \pi)$ of size less than $f(|\neg A|)$ for some suitable function f and a world w in S such that $w \notin \|\neg A\|_{\mathsf{K}}$. This last fact would yield that $w \in \|A\|_{\mathsf{K}}$ guaranteeing that the satisfiable formula A has a Kripke structure of size less than $f(|A|)$ in which it is satisfied. Hence the small model property would be shown, even though the weaker and more straightforward finite model property would still be needed to show the soundness of $\mathsf{T}_{\mu+}^\omega$.

A possible way to achieve the construction of a canonical countermodel as described above could be to start with the non-provable formula A and gradually decompose it, building up a Kripke structure based on sets similar to the saturated sequents in the proofs of Lemmata 4.2.7 and 8.2.2. Unlike in the saturation argument, where the collection of all of the infinitely many saturated sequents is used, a more careful argument could add new sets of formulae only as they are needed, that is to say, whenever a formula of the form $\Box_i B$ or $\Diamond_i B$ is decomposed. The challenging part of such an argument would undoubtedly consist of showing that the decomposition process terminates after finitely many steps and that the size of the resulting model is bounded by a suitable function of the length of A .

Taking this approach a step further would then bring us closer to solving issue 2. Instead of decomposing non-provable formulae only, the method could be adjusted to work for arbitrary formulae A , yielding a countermodel

in case A is not provable and a proof of A otherwise. Such a deduction chain technique has already been studied by Kretz and Studer [25] for Logic of Common Knowledge. The major challenge of this extended approach is to show that the decomposition process is actually correct and considers all relevant parts of the formula A , since in this case decomposition cannot be guided by additional information like, for example, the non-provability of certain components of a complex non-provable formula.

Bibliography

- [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] Simon Ambler, Marta Kwiatkowska, and Nicholas Measor. Duality and the completeness of the modal μ -calculus. *Theoretical Computer Science*, 151:3–27, 1995.
- [3] Henrik R. Andersen, Colin Stirling, and Glynn Winskel. A compositional proof system for the modal μ -calculus. In *Proceedings of the IEEE Symposium on Logic in Computer Science LICS'94*, pages 144–153, 1994.
- [4] Sergei Artemov and Lev Beklemishev. Provability logic. In Dov Gabbay and Franz Guenther, editors, *Handbook of Philosophical Logic*, volume 13, pages 189–360. Springer, 2004.
- [5] Jon Barwise. *Admissible Sets and Structures*. Perspectives in Mathematical Logic. Springer, 1975.
- [6] Jon Barwise and Lawrence Moss. *Vicious Circles*, volume 60 of *CSLI Lecture Notes*. CSLI Publications, 1996.
- [7] Sergey Berezin and Dilian Gurov. A compositional proof system for the modal μ -calculus and CCS. Technical Report CMU-CS-97-105, Carnegie Mellon University, 1997.
- [8] Dietmar Berwanger and Giacomo Lenzi. The variable hierarchy of the μ -calculus is strict. In *STACS 2005, Proceedings of the 22nd Symposium on Theoretical Aspects of Computer Science*, volume 3404 of *LNCS*, pages 97–109. Springer-Verlag, 2005.
- [9] Julian Bradfield. The modal μ -calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195:133–153, 1998.

- [10] Julian Bradfield and Colin Stirling. Modal logics and mu-calculi: an introduction. In Jan Bergstra, Alban Ponse, and Scott Smolka, editors, *Handbook of Process Algebra*, pages 293–330. Elsevier, 2001.
- [11] Edmund M. Clarke and E. Allen Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proceedings of the Workshop on Logics of Programs*, volume 131 of *LNCS*, pages 52–71. Springer, 1981.
- [12] Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.
- [13] E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model checking for the μ -calculus and its fragments. *Theoretical Computer Science*, 258:491–522, 2001.
- [14] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [15] Michael J. Fischer and Richard E. Ladner. Propositional modal logic of programs. In *Proceedings of the 9th ACM Annual Symposium on Theory of Computing*, pages 286 – 294, 1977.
- [16] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computing and System Science*, 18(2):194 – 211, 1979.
- [17] Kurt Gödel. Eine Interpretation des intuitionistischen Aussagenkalküls. In *Ergebnisse eines mathematischen Kolloquiums 4*, pages 34–40, 1933.
- [18] Joseph Y. Halpern and Yoram Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:311–379, 1992.
- [19] Matthew Hennessy and Robin Milner. Algebraic laws for indeterminism and concurrency. *Journal of the ACM*, 32:137–162, 1985.
- [20] Peter G. Hinman. *Recursion-Theoretic Hierarchies*. Perspectives in Mathematical Logic. Springer, 1978.
- [21] Gerhard Jäger, Mathis Kretz, and Thomas Studer. Cut-free common knowledge. *Journal of Applied Logic*, 2006. To appear.

- [22] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *Proceedings of the 7th Conference on Concurrency Theory CONCUR'96*, volume 1119 of *LNCS*, pages 263–277. Springer, 1996.
- [23] Dexter Kozen. Results on the propositional modal μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [24] Dexter Kozen. A finite model theorem for the propositional μ -calculus. *Studia Logica*, 47(3):233–241, 1988.
- [25] Mathis Kretz and Thomas Studer. Deduction chains for common knowledge. *Journal of Applied Logic*, 4:331–357, 2006.
- [26] Saul Kripke. Semantic analysis of modal logic *I*, normal propositional calculi. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.
- [27] Giacomo Lenzi. A hierarchy theorem for the μ -calculus. In *Proceedings of the 23rd International Colloquium on Automata, Languages and Programming ICALP96*, volume 1099 of *LNCS*, pages 87–97. Springer, 1996.
- [28] Marino Miculan. On the formalization of the modal μ -calculus in the calculus of inductive constructions. *Information and Computation*, 164:199–231, 2001.
- [29] Yannis Moschovakis. *Elementary Induction on Abstract Structures*. Studies in Logic and the Foundations of Mathematics. North-Holland, 1974.
- [30] Kurt Schütte. *Proof Theory*. Springer, 1977.
- [31] Colin Stirling and David Walker. Local model checking the modal mu-calculus. *Theoretical Computer Science*, 89:161–177, 1991.
- [32] Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional modal mu-calculus. *Information and Computation*, 81:249–264, 1989.
- [33] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.

- [34] Igor Walukiewicz. A complete deductive system for the μ -calculus. In *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science*, pages 136–147. IEEE Computer Science Press, 1993.
- [35] Igor Walukiewicz. Completeness of Kozen’s axiomatization of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000.
- [36] Glynn Winskel. A note on model checking the modal ν -calculus. *Theoretical Computer Science*, 83:157–167, 1991.

Index

- $|A|$, *see* length of formulae of \mathcal{L}_{SFL}
- $\neg A$, *see* negation
- $\|A\|_{D[X:=T]}$, *see* canonical countermodel for \mathcal{L}_μ
- $\|A\|_D$, *see* canonical countermodel for \mathcal{L}_μ
- \overline{A} , *see* dual X -positive formula
- $A \leftrightarrow B$, 21
- $A \rightarrow B$, 21
- $\|A\|_K$, *see* denotation
- $\|A\|_{\text{sat}}$, *see* canonical countermodel for \mathcal{L}_{SFL}
- $\|A\|_{\text{sat}[X:=T]}$, *see* canonical countermodel for \mathcal{L}_{SFL}
- $\|A\|_K^\sigma$, *see* signed denotation
- \mathcal{A}_n , 56
- A° , 70
- A^* , 99
- $\mathcal{A}[B/X]$, *see* positive substitution
- B^- , 69
- $b_\mu(D)$, *see* μ -bound
- $\mathcal{B}[(\mu X)\mathcal{C}]$, *see* $\mathcal{A}[B/X]$
- $\mathcal{B}[(\nu^k X)\mathcal{C}]$, *see* $\mathcal{A}[B/X]$
- $\mathcal{B}[(\nu X)\mathcal{C}]$, *see* $\mathcal{A}[B/X]$
- bound variables
 - of \mathcal{L}_μ , 68
 - of \mathcal{L}_μ^+ , 69
- $\Box_k^i A$, 21
- canonical countermodel
 - for \mathcal{L}_μ , 95
 - for \mathcal{L}_{SFL} , 42
- closed formula
 - of \mathcal{L}_μ , 68
 - of \mathcal{L}_μ^+ , 69
- closure ordinal, 54
- componentwise ordering, 73
- Computational Tree Logic, 31
- $\sigma * \tau$, *see* concatenation
- concatenation, 74
- critical subformulae, 39
- $\text{csub}(A)$, *see* critical subformulae
- $\text{csub}(\Gamma)$, 40
- CTL, *see* Computational Tree Logic
- D -saturated, 90
- denotation
 - for \mathcal{L}_C , 30
 - for \mathcal{L}_{CTL} , 31
 - for \mathcal{L}_μ^+ , 71
 - for \mathcal{L}_{PDL} , 33
 - for \mathcal{L}_{SFL} , 22
- $\Diamond_k^i A$, 21
- $\text{ds}_k(\Gamma)$, *see* k -deficiency set
- dual X -positive formula
 - for \mathcal{L}_μ , 70
 - for \mathcal{L}_{SFL} , 21
- EA, 30
- $E^k A$, 30
- $\tilde{E}A$, 30
- $\tilde{E}^k A$, 30
- F_A^K , 23
- $F_{A,X}^K$, 72
- $F_{A,X}^{K,\sigma}$, 81

- finitary Tait system
 for \mathcal{L}_μ , 98
 for \mathcal{L}_{SFL} , 50
 Fischer-Ladner closure, 77
 $\text{FL}(D)$, *see* Fischer-Ladner closure
 free variables
 of \mathcal{L}_μ , 68
 of \mathcal{L}_μ^+ , 69
 $fv(A)$, *see* free variables
- $\diamond\Gamma$, 36
 $\bigvee\Gamma$, 36
 Γ° , 71
 Γ^* , 100
 Γ, A , 36
- $\text{H}_K^\mathcal{L}$, *see* Hilbert system for modal logic
 $h_\mu(A)$, *see* μ -height
 H_{SFL} , *see* Hilbert system for SFL
 Hilbert system
 for SFL, 28
 for modal logic, 28
- $I_{A,K}^\alpha$, 24
 $I_{A,K}^{<\alpha}$, 24
 $I_{A,X,K}^\alpha$, 73
 $I_{A,X,K}^{<\alpha}$, 73
 $I_{A,X,K,\sigma}^{<\alpha}$, 81
 I_F , 14
 $I_F^{<\alpha}$, 14
 I_F^α , 14
- infinitary Tait system
 for \mathcal{L}_μ , 88
 for \mathcal{L}_μ^+ , 88
 for \mathcal{L}_{SFL} , 36
- $J_{A,K}^\alpha$, 24
 $J_{A,K}^{<\alpha}$, 24
 $J_{A,X,K}^{<\alpha}$, 73
 $J_{A,X,K}^\alpha$, 73
 $J_{A,X,K,\sigma}^{<\alpha}$, 81
- J_F , 14
 $J_F^{<\alpha}$, 14
 J_F^α , 14
- k -adequate, 43
 K_α , 59
 K_D , *see* canonical countermodel for \mathcal{L}_μ
 k -deficiency set, 39
 K_n , 57
 k -presaturated, 38
 K_{sat} , 42
 k -saturated, 39
 $K[X := T]$, *see* variable assignment
 Kripke structure
 for \mathcal{L}_μ^+ , 71
 for \mathcal{L}_{SFL} , 22
 run of a, 31
- \mathcal{L}_C , *see* Language of Logic of Common Knowledge
 \mathcal{L}_{CTL} , *see* language of CTL
 \mathcal{L}_H^M , *see* modal language
 \mathcal{L}_μ , *see* language of the μ -calculus
 \mathcal{L}_μ^+ , *see* language of the extended μ -calculus
 \mathcal{L}_μ^X , *see* X-fragment of \mathcal{L}_μ
 \mathcal{L}_{PDL} , *see* language of PDL
 \mathcal{L}_{SFL} , *see* language of SFL
 $\mathcal{L}_{\text{SFL}}^i$, 20
 $\mathcal{L}_{\text{SFL}}^-$, 54
- language
 of CTL, 31
 of PDL, 33
 of SFL, 21
 of Logic of Common Knowledge, 29
 of the μ -calculus, 68
 of the extended μ -calculus, 69
- length
 of a signature, 73

- of formulae of \mathcal{L}_μ^+ , 74
- of formulae of \mathcal{L}_{SFL} , 21
- $level(A)$, 21
- lexicographical ordering, 73
- $lh(A)$, *see* length of formulae of \mathcal{L}_μ^+
- $lh(\sigma)$, *see* length of a signature
- Logic of Common Knowledge, 29
- maximisation, 74
- $\sigma \sqcup \tau$, *see* maximisation
- modal language, 20
- μ -bound, 80
- μ -height, 80
- negation
 - for formulae of \mathcal{L}_μ , 69
 - for formulae of \mathcal{L}_{SFL} , 20
- $(\nu X)^k \mathcal{A}$, 70
- operator, 13
 - monotone, 13
- $P_{\mathcal{A}}^k$, 23
- PDL, *see* Propositional Dynamic Logic
- π_D , *see* canonical countermodel for \mathcal{L}_μ
- π_n , 57
- π_{sat} , 42
- positive substitution, 69
- Propositional Dynamic Logic, 33
- provability
 - in \mathbb{T}_μ^ω , 89
 - in $\mathbb{T}_{\mu+}^\omega$, 89
 - in $\mathbb{T}_{\text{SFL}}^\omega$, 37
 - in $\mathbb{T}_{\text{SFL}}^\omega + (\text{cut})$, 37
- $Q_{\mathcal{A}}^k$, 23
- R_D , *see* canonical countermodel for \mathcal{L}_μ
- R_{sat} , 42
- rank
 - of formulae of \mathcal{L}_μ^+ , 74
 - of formulae of \mathcal{L}_{SFL} , 24
- $rk(A)$, *see* rank
- S_D , *see* canonical countermodel for \mathcal{L}_μ
- S_{sat} , 42
- satisfaction
 - for \mathcal{L}_μ^+ , 72
 - for \mathcal{L}_{SFL} , 23
- saturated, 39
- sequent
 - of \mathcal{L}_μ , 87
 - of \mathcal{L}_μ^+ , 88
 - of \mathcal{L}_{SFL} , 36
- $(\sigma)_i$, 73
- signatures, 73
- signed denotation, 81
- $\text{SC}(D)$, *see* strong closure
- strong closure, 78
- substructure, 60
- \mathbb{T}_μ , *see* finitary Tait system for \mathcal{L}_μ
- \mathbb{T}_μ^ω , *see* infinitary Tait system for \mathcal{L}_μ
- $\mathbb{T}_\mu^\omega \not\vdash \Gamma$, *see* provability in \mathbb{T}_μ^ω
- $\mathbb{T}_\mu^\omega \vdash \Gamma$, *see* provability in \mathbb{T}_μ^ω
- $\mathbb{T}_\mu^\omega \vdash^\alpha \Gamma$, *see* provability in \mathbb{T}_μ^ω
- $\mathbb{T}_{\mu+}^\omega$, *see* infinitary Tait system for \mathcal{L}_μ^+
- $\mathbb{T}_{\mu+}^\omega \not\vdash \Gamma$, *see* provability in $\mathbb{T}_{\mu+}^\omega$
- $\mathbb{T}_{\mu+}^\omega \vdash \Gamma$, *see* provability in $\mathbb{T}_{\mu+}^\omega$
- $\mathbb{T}_{\mu+}^\omega \vdash^\alpha \Gamma$, *see* provability in $\mathbb{T}_{\mu+}^\omega$
- \mathbb{T}_{SFL} , *see* finitary Tait system for \mathcal{L}_{SFL}
- $\mathbb{T}_{\text{SFL}}^\omega + (\text{cut}) \vdash \Gamma$, *see* provability in $\mathbb{T}_{\text{SFL}}^\omega + (\text{cut})$
- $\mathbb{T}_{\text{SFL}}^\omega$, *see* infinitary Tait system for \mathcal{L}_{SFL}
- $\mathbb{T}_{\text{SFL}}^\omega + (\text{cut})$, *see* infinitary Tait system for \mathcal{L}_{SFL}

$T_{\text{SFL}}^\omega + (\text{cut}) \not\vdash \Gamma$, *see* provability in

$T_{\text{SFL}}^\omega + (\text{cut}) \stackrel{\alpha}{\vdash} \Gamma$, *see* provability
in $T_{\text{SFL}}^\omega + (\text{cut})$

$T_{\text{SFL}}^\omega \not\vdash \Gamma$, *see* provability in T_{SFL}^ω

$T_{\text{SFL}}^\omega \vdash \Gamma$, *see* provability in T_{SFL}^ω

$T_{\text{SFL}}^\omega \stackrel{\alpha}{\vdash} \Gamma$, *see* provability in T_{SFL}^ω

validity

for \mathcal{L}_μ^+ , 72

for \mathcal{L}_{SFL} , 23

variable assignment

for \mathcal{L}_μ^+ , 71

for \mathcal{L}_{SFL} , 22

variable substitution $\mathcal{A}[B]$, 21

X-fragment of \mathcal{L}_μ , 101

X-positive formula

of \mathcal{L}_μ , 68

of \mathcal{L}_μ^+ , 69

of \mathcal{L}_{SFL} , 21

$<_{lex}$, *see* lexicographical ordering

\leq_{lex} , *see* lexicographical ordering

\sqsubseteq , *see* componentwise ordering