Algebraic Filtrations of the Modal μ-Calculus

Michael Benjamin Cromberge
University of the Witwatersrand, Johannesburg

Supervised by: Professor Clint Van Alten and
Professor Willem Conradie

26 August 2016

A Dissertation submitted to the Faculty of Science, University of the Witwatersrand, Johannesburg, in fulfilment of the requirements for the degree of Master of Science in Mathematics.

Declaration

I, the undersigned, hereby declare that the work contained in this thesis is my original work, and that any work done by others or by myself previously has been acknowledged and referenced accordingly. It has not been submitted before for any degree or examination at any other University.

Michael Benjamin Cromberge. Date: 26 August 2016

Abstract

In this thesis we analyse the issue of decidability for two modal logics which contain least binders. Towards this goal, we begin the work with a brief survey of modal logic, **PDL**, the modal μ-calculus and algebraic filtrations as exposited by Conradie et al. The first such modal logic we analyse is the fragment of the modal μ-calculus corresponding to **PDL**; the second logic is the equational theory of the class of ρ-algebras (motivated by the least root calculus of Pratt). We offer a new, algebraic, proof for the decidability of **PDL** by showing that **PDL** has the finite model property with respect to the class of dynamic algebras. We then show that the equational theory of the class of ρ-algebras has the finite model property with respect to the class of ρ-algebras; this is based on the proof of Pratt but differs in an important detail. The finite model property results for these two modal logics are achieved by an algebraic filtration method based on that of Conradie et al.

Acknowledgements

I would like to thank my supervisors Clint and Willem. Without your guidance I would not have been able to undertake this project; I am extremely fortunate to have worked with the both of you. Thank you for all the time you spent meeting with me at Wits and UJ and for your patience.

I would also like to thank my dad, my brothers Pete, Dave and Dan as well as Mel, for your many words of love and encouragement.

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at, are those of the author and are not necessarily to be attributed to the NRF. The financial assistance of the African Institute for Mathematical Sciences in Cape Town is also greatly appreciated.

In loving memory of my mom

Gaye Cromberge

1952-2007

Contents

1. Introduction

Given a formula φ and a deductive system Λ, a classical question is whether there exists an algorithm that can check in finite time if φ is a theorem of Λ i.e.

Is the Λ-provability problem, a decidable problem?

A deductive system Λ will be called decidable if the Λ-provability problem is decidable. The goal of this thesis is to algebraically address the issue of decidability for two modal logics involving least binders. The first is a fragment of the modal μ-calculus and the second is a similar fragment based on the least root modal language put forward by Pratt in [34].

To better understand these algebraic perspectives on decidability, we shall discuss how they were first proved and the background leading up to the results. We begin with an overview of modal logic. The historical information in this chapter is taken from [15], [1], [33], [4] and [34].

1.1 Modal logic

The phrase "in the past" is an example of a modality in every day language. We can prefix it to a statement p, thus giving a new statement with some bearing on when p is true. Other examples of modalities abound in every day discussions involving possibility (eg. it is necessary), epistemology (eg. it is known) and ethics (eg. it is immoral) to name but a few. Our interest in this work is in the mathematics of modalities: modal logic.

The language of modal logic uses $\wedge, \vee, \neg$ to denote conjunction, disjunction and negation respectively and $\bot, \top$ for truth values. The set of propositional variables is denoted by Φ and elements thereof by $p_1, p_2, ...$ etc. The symbol $\Diamond$ is used to denote a unary modality. There may be more than one modality and these may be of arity greater than 1; these are represented by symbols $\triangle$ in a modal type τ. Modal formulas will be formed over propositional variables using the connectives $\wedge, \vee, \neg$ and $\triangle$ from τ. The formula $\neg p \vee q$ will be given the shorthand $p \rightarrow q$. The symbol $\Box$ is shorthand for $\neg\Diamond\neg$ and is called the dual of $\Diamond$.

Modal logic is the extension of propositional logic that includes the study of statements involving modalities. Largely due to Boole [3] the study of modern propositional logic began in the setting of algebra. In a similar manner, the study of modal logic began in the setting of algebra with one of the earliest algebraic analyses performed by MacColl between 1880 and 1906 [29]. MacColl denoted statement "p is impossible" by p^{η}.

In the papers of MacColl [29], much was made of the meanings and properties of logical operations but there was not a rigorous notion of logical deductive systems, as is expected by modern standards. The first truly modal deductive systems were introduced by Lewis in 1932 when he defined the systems S_1 through to S_5 in [28]. Lewis intended to make a distinction between the ordinary and algebraic meanings of implication. Since implication is equivalently defined via disjunction (p *implies* q is equivalent to q *or not* p) Lewis defined disjunction in the former case to be the usual "or" connective (he called this extensional disjunction) and in the latter case he defined intensional disjunction to be "such that at least one of the disjoined propositions is necessarily true" [27].

In his systems Lewis thus introduced the $\Diamond$ symbol for possibility which he used to define intensional disjunction and in turn the notion of strict implication, as p strictly implies q if and only if $\neg\Diamond(p\wedge\neg q)$ [27]. The Lewis systems S_1–S_5 contained axioms involving strict implications and were closed under certain proof rules. In 1931 Gödel aimed to have a deductive system formalising the modality of provability [14]. He introduced the connective B, interpreting Bp as "p is provable" and he included in his deductive system the axiom $Bp \to p$. His definition of a logic has become standard [15]. A set of modal formulae Λ is a logic (in the basic modal language) if it contains all propositional tautologies, the axioms $\Box(p \to q) \to (\Box p \to \Box q)$, $\Diamond p \leftrightarrow \neg\Box\neg p$ and is closed under modus ponens, uniform substitution and the rule: if $\varphi \in \Lambda$ then $\Box\varphi \in \Lambda$. A formula φ is said to be a theorem of Λ if $\varphi \in \Lambda$. Logics are the standard deductive systems for modal logic and are geared towards producing modal logic validities.

Briefly steering from the history back towards our first example of a modality, we can define $\langle P\rangle\, p$ to mean "p was true at some past time" and $\langle F\rangle\, p$ to mean "p will be true at some future time". The duals of $\langle P\rangle\, p$ and $\langle F\rangle\, p$ are $[P]p = \neg\langle P\rangle\neg p$ and $[F]p = \neg\langle F\rangle\neg p$ and thus can be interpreted as "it always has been p" and "it is always going to be p" respectively. This is an example of a modal language with two modal symbols: $\langle P\rangle$ and $\langle F\rangle$. As seen from the discussion of the deductive systems of Lewis and Gödel, the axioms of a modal proof system are motivated by the intended meanings of the symbols involved. Thus if we insisted that our notion of time be dense we would consider as an axiom for a temporal logic: $\langle F\rangle\, p \to \langle F\rangle\langle F\rangle\, p$. If we were in the unfortunate situation that the future holds nothing new we would require the axiom $\langle F\rangle\, p \to \langle P\rangle\, p$.

Modal algebras supply an algebraic semantics for modal logic. A modal algebra $\mathbf{A}=(A, \vee, \wedge, \neg, 0, 1, f)$ is an algebra where $(A, \vee, \wedge, \neg, 0, 1)$ is a Boolean algebra and f is a unary operator on $\mathbf{A}$: $f(0) = 0$ and $f(x \vee y) = f(x) \vee f(y)$ for all $x, y \in A$. The operator f on $\mathbf{A}$ emulates the behaviour of a modality $\Diamond$ in a logic. A modal formula φ is a validity of a modal algebra $\mathbf{A}$ if every evaluation of φ in $\mathbf{A}$ equals 1. And φ is a validity of a class of modal algebras $\mathfrak{A}$ if φ is a validity of each member from $\mathfrak{A}$. In this way modal algebras offer an algebraic semantics for modal logic.

Given a structure $\mathcal{G} = (U, R)$ where R is a binary relation on the set U, the full complex algebra of $\mathcal{G}$ is the structure $\mathcal{G}^+ = (\mathcal{P}(U), m_R)$ where $\mathcal{P}(U)$ is the power set algebra of U and for any $U' \subseteq U$ we define $m_R(U') = \{u \in U \mid \text{there exists } u' \in U' \text{ such that } Ruu'\}$. It can be easily checked that m_R is an operator on $\mathcal{P}(U)$ and thus $\mathcal{G}^+$ is a modal algebra. In 1951 Jónsson and Tarski established the result that every modal algebra is isomorphic to a complex algebra (a subalgebra of a full complex algebra) [21].

In 1963 Kripke introduced his relational semantics of modal logic in his analysis of S_2, S_5 and two other logics [25]. The Kripke semantics of modal logic involve interpreting modal formulae to be true relative to a possible world in a Kripke structure. A Kripke structure or frame $\mathcal{F} = (W, R)$ consists of a non-empty set W (the possible worlds) and a binary relation R on W (the accessibility relation). A valuation $V : \Phi \to \mathcal{P}(W)$ is added to a frame to get a (Kripke) model $\mathcal{M} = (\mathcal{F}, V)$. Models are used to give semantics to modal languages: the valuation designates at which worlds propositional variables are satisfied and in turn where more complex modal formulae are satisfied. For example the formula $\Diamond p$ is satisfied at a world $w \in W$ if there is a world w' in W such that Rww' and w' satisfies p. A formula is a validity of a frame if it is satisfied at every point in the frame under every valuation. It can be checked that the set of validities of a frame, denoted Λ_F, is a logic. Since model valuations are the same as evaluations in a full complex algebra, it can be shown that a formula is validity of a frame $\mathcal{F}$ if and only if it is a validity of the full complex algebra $\mathcal{F}^+$.

Kripke's semantics was a major breakthrough in modal logic because it offered clear insights into how logics can be characterised. For example, using semantic tableaux Kripke showed that a formula is a theorem of S_5 if and only if it is true in all transitive and symmetric models [24]. The ideas of Kripke were present in some form in the works of Hintikka, Kanger, Prior and Montague (see Goldblatt's article [15] for a full account of their contributions).

Given that there are two descriptions of modal validities, it is natural to ask if they match up i.e. given a logic Λ and frame class F is it the case that $\Lambda = \Lambda_F$? The logic Λ is said to be sound with respect to F if every theorem of Λ is a validity of F. The logic Λ is said to be complete with respect to F if every validity of the class F can be derived in Λ (i.e. is a theorem of Λ). Proofs of soundness are usually straightforward whereas those concerning completeness can be far more difficult. The canonical model of a logic Λ is a natural tool in attempting a proof of completeness for Λ and was first used by Lemmon and Scott in the "Lemmon Notes" of 1966 [26]. The canonical model of Λ is constructed from the Λ-maximal consistent sets and is in fact a modal version of the Henkin construction used for the completeness of first order logic [17].

A logic Λ is characterised by a class of frames F if it is both sound and complete with respect to F. The first logics that were studied all turned out to be characterised by a class of frames [15]; in 1966 this led Lemmon to conjecture that this was the case for all logics [26]. However, in 1972 Thomason gave the first example of a logic lacking the frame characterisation property with a tense logic [39]. Furthermore in 1977–1979 work by Blok showed that the majority of logics are not characterised by any class of frames [2]. Blok used algebraic methods in his work, studying varieties rather than the logics themselves. These incompleteness results demonstrated the inadequacy of relational semantics of modal logic and led to a renewed interest in algebraic semantics.

Thomason was among those that spearheaded the renewed interest in the algebraic semantics in the modern era (1972–onwards). In particular he introduced general frames in his paper [39]. Thomason used his general frame semantics to show that, in contrast to the relational semantics, every logic is characterised by a class of modal algebras. Although his result shows completeness of any logic with respect to a class of abstract modal algebras, it was not necessarily a completeness result with respect to a class of full complex modal algebras. But the representation theorem of Jónsson and Tarski shows that every modal algebra is isomorphic to a complex algebra [21]. The Jónsson–Tarski theorem therefore gives a path towards relational completeness and it was during this era that the importance of their work to modal logic finally came to be realised.

Two other active areas of modal logic research during this time were correspondence theory and modal logic applied to computer science. Correspondence theory was initiated by van Benthem and focuses on the extent to which modal languages can describe frames and models (see [1]). The application of modal logic to computer science began mainly with Pratt and Pnueli and will be briefly discussed in Section 1.3.

1.2 Decidability of a logic

In this section we will discuss the decidability of a logic Λ in a little more detail. Recall that Λ is decidable if the Λ-provability problem is decidable i.e. there is an algorithm that can check in finite time whether or not $\varphi \in \Lambda$ for any modal formula φ.

In 1941 it was proved by McKinsey that S_2 and S_4 are decidable logics [31]. He did this by showing that the logic in question has the finite model property with respect to a class $\mathfrak{A}$ of algebras. This property states that a formula is a theorem of the logic if and only if it is a validity of $\mathfrak{A}_{fin}$ (where $\mathfrak{A}_{fin}$ denotes the finite members of $\mathfrak{A}$). McKinsey's method for proving the finite model property is the first example of an algebraic filtration [31].

It was demonstrated by Harrop in 1958 that any logic Λ which has a finite number of axioms and has the finite model property with respect to a class of algebras $\mathfrak{A}$, is a decidable logic [16]. A schematic version of his argument is as follows: construct an algorithm that makes use of the finite number of axioms of Λ to enumerate all the theorems of Λ; construct a second algorithm that enumerates all of the finite algebras from $\mathfrak{A}$. These two algorithms in parallel can check in finite time if φ is a theorem of Λ or not: if φ is a theorem of Λ then the first algorithm will halt; if φ is not a theorem of Λ then, as Λ has the finite model property with respect to $\mathfrak{A}$, there exists a finite algebra in $\mathfrak{A}$ on which φ is not a validity which will lead the second algorithm to halt.

A logic Λ has the finite model property with respect to a class of frames F if: a formula is a theorem of Λ if and only if it is a validity of F_{fin}. Note that by definition, if Λ is characterised by F and the validities of F and F_{fin} are the same then Λ has the finite model property with respect to F. The statement that the validities of F and F_{fin} are the same, is equivalent to saying that when a model from F (i.e. a model based on a frame in F) satisfies a modal formula then there is a finite model from F which also satisfies the formula. In 1966 Lemmon and Scott developed a technique for proving the finite model property in the Lemmon Notes [26]. Given a model $\mathcal{M}$ that satisfies a formula φ, they showed that one of the ways to attempt to prove there exists a finite model which also satisfies φ is by constructing a filtration of $\mathcal{M}$. A filtration of $\mathcal{M}$ is formed by first identifying two worlds in $\mathcal{M}$ to be equivalent if they satisfy the same formulae in $\text{sub}(\varphi)$ (where $\text{sub}(\varphi)$ denotes the subformulae of φ). A relation and valuation are then imposed on this finite quotient structure to form a finite model that preserves the satisfaction of formulae from $\text{sub}(\varphi)$. Note that Harrop's argument carries over to the frames as well and thus if a logic Λ has the finite model property with respect to a class of frames then Λ is decidable.

Although a logic which has finitely many axioms and the finite model property is decidable, work by Gabbay in 1972 showed that the converse is not true i.e. he showed that there are decidable logics which are finitely axiomatisable but lack the finite model property [12]. Another method for proving the decidability of a logic is by using automata: representations of machines which perform computations on a given input (such as a finite string) by passing through a sequence of states (see for example [35]).

1.3 Propositional dynamic logic (PDL)

Given a computer system, one wants to be able to verify whether the system satisfies certain requirements. An example of a requirement is: the system eventually responds to an instruction; this is referred to as a liveness property. Applying formal logic to test such requirements in computer science goes back to the 1960s with Floyd–Hoare logic [19]. For example Hoare introduced the construct $p\{a\}q$ which reads: if p holds before executing program a then q holds afterwards. This line of work was proof theoretic with logics representing the properties of programs and was aimed at testing the correctness of programs. In 1969 Pnueli introduced a semantic approach to questions of program correctness in that he turned assertions about programs into assertions of satisfiability or validity using models of temporal logic [30]. Apart from showing program correctness, formal logic applied to computer science can also be used for showing two programs to be equivalent as well as for comparing their expressive power (see [9]).

In 1976 Pratt introduced more expressive logics into the study of program correctness by building on the ideas of Hoare [32]. In particular Pratt introduced a dynamic logic where modalities represent the actions of programs. For example the formula $[a]p$ expresses the statement: after program a executes, p holds. Using this, Hoare's construct $p\{a\}q$ can be translated into dynamic logic as $p \to [a]q$. The dual of $[a]$ is $\langle a \rangle$ with $\langle a \rangle p$ expressing that there is an execution of a after which p holds. The program modalities from Pratt's dynamic logic [32] are built from atomic programs using dynamic constructs. These include $\pi_1; \pi_2$ expressing the program that executes program π_1 then program π_2, the program $\pi_1 \cup \pi_2$ that non-deterministically executes program π_1 or program π_2 and finally π^* which expresses the program that executes π a finite number of times. This language is referred to as propositional dynamic logic, PDL. We use **PDL** to denote the logic (deductive system) of PDL.

The language of PDL is interpreted on Kripke structures such that for each program modality there is a relation on the Kripke structure (if a program a is deterministic then the relation $R_{\langle a \rangle}$ is required to be a partial function). The relations in such a structure respect the dynamic constructs of the programs, for example $R_{\langle \pi_1; \pi_2 \rangle} = R_{\langle \pi_1 \rangle} \circ R_{\langle \pi_2 \rangle}$ where $\circ$ represents relation composition; $R_{\langle \pi_1 \cup \pi_2 \rangle} = R_{\langle \pi_1 \rangle} \cup R_{\langle \pi_2 \rangle}$ and $R_{\langle \pi^* \rangle} = R^*_{\langle \pi \rangle}$ where $R^*_{\langle \pi \rangle}$ represents the reflexive transitive closure of $R_{\langle \pi \rangle}$. Such a frame is called a regular frame.

A finite axiomatisation of PDL was given by Segerberg in 1977 [37]. This logic included axioms that modelled the behaviour of the program constructs; the least trivial of which is the axiom which resembles induction given as

$$p \to ([\pi^*](p \to [\pi]p) \to [\pi^*]p).$$

The first proof of completeness for **PDL** was given by Parikh in 1978 with later proofs given by Segerberg and Gabbay in 1982. Fischer and Ladner showed in 1979 that **PDL** has the finite model property with respect to the class of regular frames by using a modified filtration incorporating their Fischer–Ladner closure [11]. Loosely speaking their closure is needed to "educate" the filtration about the dynamic constructs of programs. Together with the fact that **PDL** has finitely many axioms, this shows that **PDL** is a decidable logic by Harrop's argument [16].

Dynamic algebras were introduced by Kozen and Pratt in 1979 in order to algebraically interpret

PDL; in the words of Pratt: "the class [of dynamic algebras] can be considered the algebraic home of induction" [33]. A dynamic algebra is a modal algebra with operators for each modality; the operators obey similar rules to those on a regular frame and the full complex algebra of a regular frame indeed gives a dynamic algebra. In this work we will give a new proof that **PDL** is decidable by using dynamic algebras (see Chapter 5).

There are also extensions of **PDL**: in 1979 Pratt introduced process logic which is **PDL** enriched with temporal operators. Towards showing two program systems to be equivalent, Hennessy and Milner introduced their logic in 1980 [18]. The Hennessy–Milner logic (**HML**) has unary modalities representing actions in computer a system but differs from **PDL** in that the action labels have no structure. **HML** is inadequate to express enduring properties in **PDL** such as $\langle \pi^* \rangle\, p$ since formulae in **HML** can only "see" up to the number of modal connectives [18]. In 1981 **HML** was enriched with higher arity modalities by Emerson and Clarke [6]. Among them is the until binary modality, **U**, which is interpreted over paths in a model (W, R, V): given $w \in W$, a path from w is a finite or infinite sequence of states $w, w_1, w_2, w_3...$ in W with Rww_1, Rw_1w_2 etc. The formula $\exists[\varphi \mathbf{U} \psi]$ is satisfied at w if there is a path (w_i) from w and a j in $\mathbb{N}$ such that φ is satisfied at all w_i with $0 \leq i < j$ and ψ is satisfied at w_j. In this way $\exists[\varphi \mathbf{U} \psi]$ expresses that there is a path from w on which φ is satisfied until ψ is satisfied. This enrichment of **HML** is called computational tree logic (**CTL**).

1.4 The modal μ-calculus

The modal μ-calculus is an enrichment of modal logic with least fixed point binders/quantifiers. The modal μ-calculus provides a powerful and elegant framework for program semantics as well as recursion. The notion of adding fixed point binders to modal logic first appeared in a paper by Scott and De Bakker in 1969 [36]. Other authors who developed the theory from 1970–1981 include Park, De Roever, Emerson and Pratt (see [15]). Hennessy and Milner used the largest fixed point in their definition of an "observational equivalence relation" between programs [18].

In 1983 Kozen introduced the most widely used version of the modal μ-calculus: L_μ [22]. In addition to using the usual binary connectives, negation and unary modal connectives, the formulae of L_μ are formed using a least fixed point binder μ. The fixed point formulae can represent assertions of safety or liveness regarding a computer system. Informally a liveness property refers to the assertion that "something good will eventually happen" and a safety property refers to the assertion "nothing bad will ever happen".

Towards interpreting fixed point formulae in a model $\mathcal{M} = (\mathcal{F}, V)$, the valuation of the formula

$\varphi(X)$ gives a different set as X varies. In this way the valuation of $\varphi(X)$ induces an increasing function on $\mathcal{P}(W)$ (increasing by a syntactic restriction on $\varphi(X)$). The formula $\mu X.\varphi(X)$ is then interpreted as the least fixed point of this function, guaranteed to exist by the theorem of Knaster and Tarski proved in 1930.

Kozen gave a finite axiomatisation for L_μ in [22] but was only able to show completeness for a proper fragment of L_μ. The problem of showing that the full L_μ is complete under Kozen's finite axiomatisation was only solved in 1995 by Walukiewicz [40]. In 1984 Kozen and Parikh demonstrated that L_μ is decidable via automata [23]. As demonstrated by Streett [38] the decidability of the modal μ-calculus cannot be shown using a filtration style technique due to L_μ formulae such as $\mu X.[a]X$.

The modal μ-calculus strictly subsumes many temporal logics, including PDL and CTL mentioned previously [38]. We shall show algebraically that **PDL** is decidable which will in turn serve to demonstrate that the fragment of the modal μ-calculus corresponding to **PDL**, denoted L_μ^*, is decidable. Alternation of fixed points makes the μ-calculus highly expressive and much work has been done on classifying on which level modal languages fall in the so-called "μ-calculus alternation hierarchy" [4]. Whether the model checking of formulae containing alternating fixed points necessarily takes exponential time is still an open problem [13].

Before Kozen introduced L_μ in 1983, Pratt examined a similar modal language, L_ρ, with least root binders in 1981 [34]. Unlike L_μ, Pratt's least root language is amenable to filtration arguments. This is achieved by imposing strong syntactic conditions on the minuends. The algebraic semantics for Pratt's fragment are given via $\mathfrak{R}$, the class of ρ-algebras. There is an overlap between the languages L_ρ and L_μ but, as shall be seen in Chapter 5, L_ρ is not a subset of L_μ.

1.5 Thesis contribution and outline

As can be seen by the argument of Harrop, if one has a finitely axiomatisable logic which has the finite model property with respect to a class of models, then the logic will be decidable [16]. The objective of this work is to use this argument in an attempt to prove the decidability of two modal logics with least binders by showing that they have the finite model property with respect to certain classes of modal algebras.

In particular, the first such logic is **PDL** and is identified with the fragment L_μ^* of the modal μ-calculus. **PDL** is a finitely axiomatisable logic and thus we will show that **PDL** is decidable by proving that it has the finite model property with respect to the class of dynamic algebras. In this way we offer a new proof for the decidability of **PDL**.

The second logic we will look at is $\Lambda_{\mathfrak{R}}$, the equational theory of $\mathfrak{R}$, where $\mathfrak{R}$ is the class of ρ-algebras i.e. the algebraic class which interprets Pratt's least root calculus [34]. We will only show that $\Lambda_{\mathfrak{R}}$ has the finite model property with respect to $\mathfrak{R}$. Pratt provided a sketch of the proof in Theorem 3 of [34]; based on his sketch it is not clear that the constructed finite ρ-algebra falsifies least root formulae. Hence we will provide a proof based on his sketch but filling in this important step. Both of these finite model property results will be proved by an algebraic filtration method based on that of Conradie et al. [7].

The layout for the remainder of the thesis is as follows. In Chapter 2 the necessary background for relational and algebraic modal logic will be covered. PDL and the modal μ-calculus will be formally introduced in Chapter 3. The more recent results in the development of modal logic begin in Chapter 4 where we review the algebraic filtrations of Conradie et al. [7]. The main contribution of this thesis lies in Chapter 5; there we give a new proof for the decidability of **PDL** (see Theorem 5.1.11 and Theorem 5.1.15) and then we review the result of Pratt (Theorem 3 of [34]) that $\Lambda_{\mathfrak{R}}$ has the finite model property with respect to $\mathfrak{R}$ (Theorem 5.2.22). In particular for Pratt's result, we fill in an important step which appears to be missing from his proof. Chapter 6 is the concluding chapter where we review the work as well as discuss possible future research directions.

2. Modal Logic

The results in this chapter are taken from the sources [1], [8], [5] and [21]. For each result, the proof is provided when we feel it offers some insight. When a proof is provided, we fill out more details than those found in the source material.

In this chapter we explore various classical results in modal logic that are of interest in studying our decidability question. We shall begin with an overview of the relational semantics of modal logic in Section 2.1 as well as modal deductive systems in Section 2.2. We shall then discuss soundness and completeness in Section 2.3 and the finite model property in Section 2.4. We then discuss the previous topics again but through the lens of modal algebras. In particular, the algebraic semantics of modal logic is looked at in Section 2.5 and then completeness with respect to algebras is discussed in Section 2.6. Finally the duality between frames and algebras is briefly discussed in Section 2.7.

2.1 The relational semantics of modal logic

2.1.1 Definition. A *type* $\mathfrak{t}$ is a pair (O, ζ) where O is a non-empty set of *function symbols* and $\zeta : O \to \mathbb{N}$. For $f \in O$, the number $\zeta(f)$ is the *arity* of f.

2.1.2 Remark. When the symbols of O are used to represent modalities we call the type a *modal type* and denote it by τ. If τ consists only of a single modal symbol of arity 1, then τ is referred to as the *basic* modal type with the modal symbol denoted by the diamond $\Diamond$ i.e. $\tau = (\{\Diamond\}, \{(\Diamond, 1)\})$.

Given a modal type $\tau = (O, \zeta)$, we will usually denote the situation $\triangle \in O$ by $\triangle \in \tau$.

2.1.3 Definition. Let Φ be a non-empty set of propositional variables and $p \in \Phi$. A τ-*modal formula* φ is defined recursively as $\varphi = \bot \mid p \mid \neg\varphi \mid \varphi \vee \psi \mid \triangle(\varphi_1, ..., \varphi_{\zeta(\triangle)})$ where $\triangle \in \tau$. The set of all such formulae is denoted by $Form(\tau, \Phi)$.

2.1.4 Remark. Some formulae in $Form(\tau, \Phi)$ are assigned the following shorthand:

formula	**shorthand formula**
$\neg\varphi \vee \psi$	$\varphi \to \psi$
$\neg\bot$	$\top$
$\neg(\neg\varphi \vee \neg\psi)$	$\varphi \wedge \psi$
$\neg\triangle(\neg\varphi_1, ..., \neg\varphi_{\zeta(\triangle)})$	$\triangledown(\varphi_1, ..., \varphi_{\zeta(\triangle)})$

When τ is the basic modal type, the dual of $\Diamond$, $\neg\Diamond\neg$, is denoted by $\Box$.

We will now introduce objects with just enough structure to determine a modal formula's truth. In particular we will introduce frames and then models.

2.1.5 Definition. A τ-*frame* is a pair $\mathcal{F} = (W, R_\triangle)_{\triangle\in\tau}$ where W is a non-empty set and $R_\triangle \subseteq W^{\zeta(\triangle)+1}$ for each $\triangle \in \tau$; these are the *accessibility relations* of $\mathcal{F}$.

2.1.6 Example. Let τ be a modal type with two modal symbols of arity 1: $\tau = \{\langle a\rangle, \langle b\rangle\}$. An example of a τ-frame is $\mathcal{F} = (\{w_1, w_2, w_3, w_4, w_5\}, R_{\langle a\rangle}, R_{\langle b\rangle})$ where $R_{\langle a\rangle} = \{(w_1, w_2), (w_1, w_3), (w_1, w_4), (w_2, w_3), (w_2, w_4), (w_3, w_4)\}$ and $R_{\langle b\rangle} = \{(w_1, w_2), (w_1, w_3), (w_1, w_5), (w_2, w_3), (w_2, w_5), (w_3, w_5)\}$. The frame $\mathcal{F}$ is represented by the following diagram where elements of relations are given by arrows; the dashed arrows are from $R_{\langle b\rangle}$. It can be seen that both relations are transitive (i.e. $R_{\langle a\rangle}u_1u_2$ and $R_{\langle a\rangle}u_2u_3$ imply $R_{\langle a\rangle}u_1u_3$ for all $u_1, u_2, u_3 \in W$, and similarly for $R_{\langle b\rangle}$).

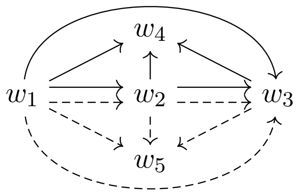

2.1.7 Definition. Let $\mathcal{F} = (W, R_\triangle)_{\triangle\in\tau}$ be a τ-frame, let $\triangle \in \tau$ and let $w \in W$. A $R_\triangle$-*path from* w is a sequence of states (w_i) such that $w_0 = w$ and $R_\triangle w_i w_{i+1}$ for all $0 \le i < n$ (if $|(w_i)| = n$) or $R_\triangle w_i w_{i+1}$ for all i (if (w_i) infinite).

2.1.8 Example. In the frame from Example 2.1.6, $w_2w_3w_4$ is a $R_{\langle a\rangle}$-path from w_2.

2.1.9 Definition. A τ-*model* is a pair $\mathcal{M} = (\mathcal{F}, V)$ where $\mathcal{F} = (W, R_\triangle)_{\triangle\in\tau}$ is a τ-frame and V is a *valuation*: $V : \Phi \to \mathcal{P}(W)$. The τ-model $\mathcal{M}$ is said to be *based on the* τ-*frame* $\mathcal{F}$.

2.1.10 Example. Let τ be the basic modal type and $\Phi = \{p, q\}$. Then $\mathcal{M} = (\{w_1, w_2, w_3\}, \{(w_1, w_1), (w_1, w_2), (w_3, w_3)\ (w_3, w_1), (w_3, w_2)\}, \{(p, \{w_1\}), (q, \varnothing)\})$ is a τ-model. $\mathcal{M}$ can be represented by the following diagram.

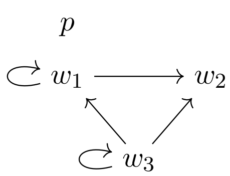

For a τ-model $\mathcal{M} = (W, R_\triangle, V)_{\triangle\in\tau}$ and a propositional variable p, the set $V(p)$ determines which points in the model satisfy p; the semantics of more complex formulae are determined recursively over those of the propositional variables. The accessibility relation is needed for the semantics of non-Boolean modal formulae.

2.1.11 Definition. Let $\varphi, \varphi_1, ..., \varphi_n, \psi \in Form(\tau, \Phi)$, let $\mathcal{M} = (W, R_\triangle, V)_{\triangle \in \tau}$ be a τ-model with $w \in W$. The formulae in $Form(\tau, \Phi)$ are interpreted in the following way:

$$\begin{array}{rcl} \mathcal{M}, w \Vdash p \in \Phi & \text{if} & w \in V(p) \\ \mathcal{M}, w \Vdash \neg\varphi & \text{if} & \mathcal{M}, w \nVdash \varphi \\ \mathcal{M}, w \Vdash \varphi \vee \psi & \text{if} & \mathcal{M}, w \Vdash \varphi \text{ or } \mathcal{M}, w \Vdash \psi \\ \mathcal{M}, w \Vdash \triangle(\varphi_1, ..., \varphi_{\zeta(\triangle)}) & \text{if} & \text{there exists a } (w_1, ..., w_{\zeta(\triangle)}) \in W^{\zeta(\triangle)} \text{ such that} \\ & & R_\triangle w w_1 ... w_{\zeta(\triangle)} \text{ and } \mathcal{M}, w_i \Vdash \varphi_i, 1 \leq i \leq \zeta(\triangle). \end{array}$$

It is never the case that $\mathcal{M}, w \Vdash \bot$. The formula φ is *satisfied in* $\mathcal{M}$ *at* w if $\mathcal{M}, w \Vdash \varphi$. The formula φ is *satisfiable in* $\mathcal{M}$ if there exists a point in $\mathcal{M}$ at which φ is satisfied.

2.1.12 Example. Recall the model $\mathcal{M}$ from Example 2.1.10. $\mathcal{M}, w_1 \Vdash \Diamond p$; $\mathcal{M}, w_3 \Vdash \Diamond p$ and vacuously $\mathcal{M}, w_2 \Vdash \Box p \wedge \Box q$.

2.1.13 Remark. For a τ-model $\mathcal{M} = (W, R_\triangle, V)_{\triangle \in \tau}$ with $w \in W$, we will sometimes denote $\mathcal{M}, w \Vdash \varphi$ simply by $w \Vdash \varphi$ when $\mathcal{M}$ is understood from the context.

2.1.14 Remark. The formula $\Diamond p$ is satisfied at a point w if p is satisfied at a point accessible from w. Likewise $\Box p = \neg\Diamond\neg p$ is satisfied at w if it is not the case that there is an accessible point from w where p is not satisfied or equivalently, if w_1 is accessible from w then p is satisfied at w_1. This motivates the interpretations of $\Diamond p$ and $\Box p$ as *possibly* p and *necessarily* p respectively.

2.1.15 Remark. Given that we have defined the semantics of any τ-modal formula φ, we can now extend our valuation V to a $V^* : Form(\tau, \Phi) \to \mathcal{P}(W)$ by $V^*(\varphi) = \{w \in W \mid w \Vdash \varphi\}$. In this way $w \Vdash \varphi$ iff $w \in V^*(\varphi)$ and it can easily be checked against Definition 2.1.11, with τ as the basic modal type, that for all $\varphi, \psi \in Form(\tau, \Phi)$:

$$\begin{array}{rcl} V(\bot) & = & \varnothing \\ V^*(p) & = & V(p) \\ V^*(\neg\varphi) & = & -V^*(\varphi) \\ V^*(\varphi \vee \psi) & = & V^*(\varphi) \cup V^*(\psi) \\ V^*(\Diamond\varphi) & = & m_R(V^*(\varphi)) \end{array}$$

where $-V^*(\varphi) = W \setminus V^*(\varphi)$ and $m_R : \mathcal{P}(W) \to \mathcal{P}(W)$ is defined by $m_R(S) = \{w \in W \mid (\exists w' \in S)(Rww')\}$. Note that $V^*(\Box\varphi) = l_R(V^*(\varphi))$ where $l_R : \mathcal{P}(W) \to \mathcal{P}(W)$ is defined by $l_R(S) = \{w \in W \mid (\forall w' \in W)(Rww' \implies w' \in S)\}$.

2.1.16 Lemma. Let $\mathcal{M} = (W, R_\triangle, V)_{\triangle\in\tau}$ be a τ-model and let $w \in W$. The following are equivalent for all $\varphi, \psi \in Form(\tau, \Phi)$:

(i) $\mathcal{M}, w \Vdash \varphi \to \psi$;

(ii) $\mathcal{M}, w \Vdash \varphi$ implies $\mathcal{M}, w \Vdash \psi$.

Proof. $(i) \implies (ii)$ Let $\mathcal{M}, w \Vdash \varphi \to \psi$ and let $\mathcal{M}, w \Vdash \varphi$. We will show that $\mathcal{M}, w \Vdash \psi$. Since $\mathcal{M}, w \Vdash \varphi \to \psi$ we have by Definition 2.1.11 that $\mathcal{M}, w \nVdash \varphi$ or $\mathcal{M}, w \Vdash \psi$. If $\mathcal{M}, w \nVdash \varphi$ we would have a contradiction since we have $\mathcal{M}, w \Vdash \varphi$ by assumption. Thus we must have $\mathcal{M}, w \Vdash \psi$.
$(ii) \implies (i)$ Assume the conditional: $\mathcal{M}, w \Vdash \varphi$ implies $\mathcal{M}, w \Vdash \psi$. We will show that $\mathcal{M}, w \Vdash \varphi \to \psi$, in particular, $\mathcal{M}, w \nVdash \varphi$ or $\mathcal{M}, w \Vdash \psi$. If $\mathcal{M}, w \nVdash \varphi$ we are done. Suppose $\mathcal{M}, w \Vdash \varphi$. Therefore by the assumed conditional $\mathcal{M}, w \Vdash \psi$ and we are done. □

2.1.17 Definition. Let $\varphi \in Form(\tau, \Phi)$ and let $\mathcal{F} = (W, R_\triangle)_{\triangle\in\tau}$ be a τ-frame. The formula φ *is a validity on* $\mathcal{F}$ if $(\mathcal{F}, V), w \Vdash \varphi$ for all valuations V on $\mathcal{F}$ and all points w in $\mathcal{F}$. This is denoted by $\mathcal{F} \Vdash \varphi$.

2.1.18 Example. Recall the frame $\mathcal{F} = (W, R_{\langle a\rangle}, R_{\langle b\rangle})$ from Example 2.1.6 which had both relations $R_{\langle a\rangle}$ and $R_{\langle b\rangle}$ transitive. Let $\varphi = \langle a\rangle\langle a\rangle p \to \langle a\rangle p$, we shall show $\mathcal{F} \Vdash \varphi$. Let $w \in W$ and let $V : \Phi \to \mathcal{P}(W)$. We need to show $(\mathcal{F}, V), w \Vdash \langle a\rangle\langle a\rangle p \to \langle a\rangle p$. Suppose $w \Vdash \langle a\rangle\langle a\rangle p$; therefore there exists a $w_1 \in W$ such that $R_{\langle a\rangle} w w_1$ and $w_1 \Vdash \langle a\rangle p$. Thus again there exists a $w_2 \in W$ such that $R_{\langle a\rangle} w_1 w_2$ and $w_2 \Vdash p$. Now since $R_{\langle a\rangle}$ is transitive, $R_{\langle a\rangle} w w_2$ and thus $w \Vdash \langle a\rangle p$. Therefore $w \Vdash \varphi$ by Lemma 2.1.16 and thus $\mathcal{F} \Vdash \varphi$. Similarly it can be shown that $\mathcal{F} \Vdash \langle b\rangle\langle b\rangle p \to \langle b\rangle p$.

2.1.19 Example. Let $(4) = \Diamond\Diamond p \to \Diamond p$. We will show that for a frame of the basic modal type $\mathcal{F} = (W, R)$ that $\mathcal{F} \Vdash (4)$ implies R is transitive. Let $\mathcal{F} \Vdash (4)$ and let $w_1, w_2, w_3 \in W$ such that Rw_1w_2, Rw_2w_3. We want to show that Rw_1w_3. Consider the valuation $V : \Phi \to \mathcal{P}(W)$ by $V(p) = \{w_3\}$; this gives $(\mathcal{F}, V), w_2 \Vdash \Diamond p$ and $(\mathcal{F}, V), w_1 \Vdash \Diamond\Diamond p$. But $w_1 \Vdash (4)$ since $\mathcal{F} \Vdash (4)$; this all together gives by Lemma 2.1.16 that $w_1 \Vdash \Diamond p$. Therefore there exists a point $w' \in W$ such that Rw_1w' and $w' \Vdash p$. But $V(p) = \{w_3\}$, therefore $w' = w_3$ and Rw_1w_3. Thus we have shown that R is transitive.

In light of Remark 2.1.15 the definition of a validity (Definition 2.1.17) could be reformulated to: $\mathcal{F} \Vdash \varphi$ iff $V^*(\varphi) = W$ for all valuations V. Indeed this is used in the algebraic semantics of modal logic (see Definition 2.5.10).

2.1.20 Definition. Let $\varphi \in Form(\tau, \Phi)$ and let F be a *τ-frame class* (F is a class of τ-frames). The formula φ *is a validity on* F if $\mathcal{F} \Vdash \varphi$ for all $\mathcal{F} \in F$. This is denoted by $F \Vdash \varphi$.

We now look at the notion of local semantic consequence.

2.1.21 Definition. Let $\Gamma \subseteq Form(\tau, \Phi)$ and let $\mathcal{M} = (W, R_\triangle, V)_{\triangle \in \tau}$ a τ-model. We say Γ is *satisfiable in* $\mathcal{M}$ if there is a $w \in W$ such that $\varphi \in \Gamma$ implies $\mathcal{M}, w \Vdash \varphi$. This is denoted by $\mathcal{M}, w \Vdash \Gamma$.

2.1.22 Definition. Let $\Gamma \subseteq Form(\tau, \Phi)$, let $\varphi \in Form(\tau, \Phi)$ and let F be a τ-frame class. The formula φ is a *local semantic consequence of* Γ *over* F if $\mathcal{M}, w \Vdash \Gamma$ implies $\mathcal{M}, w \Vdash \varphi$ for all models $\mathcal{M}$ based on frames in F and all points w in $\mathcal{M}$. This is denoted by $\Gamma \Vdash_F \varphi$.

2.1.23 Example. Let τ be the basic modal type, let $\Gamma = \{\Diamond p, \Box\Diamond p\}$ and let F be the class of all τ-frames. We shall show $\Gamma \Vdash_{\mathrm{F}} \Diamond\Diamond p$. Consider a τ-model $\mathcal{M}$ and $w \in W$ such that $\mathcal{M}, w \Vdash \Gamma$. We need to show $w \Vdash \Diamond\Diamond p$. Since $w \Vdash \Diamond p$ there exists $w_1 \in W$ such that Rww_1 and $w_1 \Vdash p$. Since Rww_1 and $w \Vdash \Box\Diamond p$ we have that $w_1 \Vdash \Diamond p$ and thus $w \Vdash \Diamond\Diamond p$.

2.2 Deductive systems of modal logic

In this section we take for granted the notions of a *propositional tautology* and *lattice*. The relevant material can be found in [10] and [8], respectively.

We have seen how to generate well formed modal formulae as well as how to formally handle their semantics. We shall introduce modal logic proof systems known as logics, the axioms of which are motivated towards generating the validities on a frame class. For the sake of readability, most of the definitions and theorems will now be given with respect to the basic modal type and can be suitably generalized to an arbitrary modal type. When working with τ as the basic modal type, we suppress τ in definitions and results, for example $Form(\tau, \Phi)$ will be $Form(\Phi)$. We shall also work with a countable set of propositional variables Φ.

2.2.1 Definition. A *logic* is a set of modal formulae Λ that contains all propositional tautologies together with the following axioms:

(K) $\quad \Box(p \to q) \to (\Box p \to \Box q)$
$(Dual)$ $\quad \Diamond p \leftrightarrow \neg\Box\neg p.$

Λ is closed under modus ponens (m.p.), uniform substitution (u.s.) and *generalization (gen)*: if $\varphi \in \Lambda$ then $\Box\varphi \in \Lambda$. A formula φ is a *theorem of* Λ if $\varphi \in \Lambda$; this is denoted by $\vdash_\Lambda \varphi$.

2.2.2 Remark. Let F be a frame class. It can be seen using the model satisfaction (Definition 2.1.11) as well as the definition of frame validity (Definition 2.1.17), that the set Λ_F of F validities ($\Lambda_F = \{\varphi \in Form(\Phi) \mid F \Vdash \varphi\}$) is a logic.

The next theorem shows that logics can be defined using the diamond; we state it without proof.

2.2.3 Theorem. *Let $\Lambda \subseteq Form(\Phi)$ such that Λ contains all propositional tautologies and is closed under modus ponens and uniform substitution. The following are equivalent:*

(*i*) Λ *is a logic.*

(*ii*) $\vdash_\Lambda \Diamond\bot \leftrightarrow \bot$;
$\vdash_\Lambda \Diamond(p \vee q) \leftrightarrow (\Diamond p \vee \Diamond q)$;
if $\vdash_\Lambda p \to q$ *then* $\vdash_\Lambda \Diamond p \to \Diamond q$.

2.2.4 Remark. We will refer to the closure rule, if $\vdash_\Lambda (p \to q)$ then $\vdash_\Lambda \Diamond p \to \Diamond q$, as $\Diamond$-*generalization* ($\Diamond$-*gen*).

We now want to show that for any set of modal formulae, one can find a smallest logic which contains it. Towards this result we now define complete lattices.

2.2.5 Definition. A lattice $\mathbf{L} = (L, \vee, \wedge)$ is a *complete lattice* if $\bigwedge C$ exists for all $C \subseteq L$.

2.2.6 Lemma. Let the class of all logics be denoted N. Then $\mathbf{N} = (N, \cup, \cap)$ is a complete lattice with $\bigwedge C = \bigcap C$ for all $C \subseteq N$.

Proof. Let $C \subseteq N$. The fact that $\bigcap C$ contains all the logic axioms and is closed under modus ponens, uniform substitution and generalization is immediate. Thus $\bigcap C \in N$. Dually $\bigvee C = \bigcap C^u$, where $C^u = \{\Lambda \in N \mid (\forall \Psi \in C)(\Psi \subseteq \Lambda)\}$. □

2.2.7 Corollary. Let $\Gamma \subseteq Form(\Phi)$. There exists a smallest logic containing Γ, denoted $K \oplus \Gamma$. In particular, if Ψ is a logic such that $\Gamma \subseteq \Psi$ then $K \oplus \Gamma \subseteq \Psi$.

Proof. Consider the following set of logics: $C = \{\Psi \in N \mid \Gamma \subseteq \Psi\}$. We have that C is non-empty as $Form(\Phi) \in C$. Define $K \oplus \Gamma$ as $\bigcap C$. The set $K \oplus \Gamma \in N$ since $\mathbf{N}$ is a complete lattice (Lemma 2.2.6). Also $\Gamma \subseteq K \oplus \Gamma$ since $\Gamma \subseteq \Psi$ for all $\Psi \in C$ by definition of C. Finally if $\Psi \in N$ such that $\Gamma \subseteq \Psi$ then $\Psi \in C$ (by definition of C) and $K \oplus \Gamma \subseteq \Psi$ since $K \oplus \Gamma = \bigcap C$. Thus we have shown that $K \oplus \Gamma$ is the smallest logic containing Γ. □

2.2.8 Remark. The $K\oplus$ mapping (i.e. $\Gamma \mapsto K \oplus \Gamma$ for $\Gamma \subseteq Form(\Phi)$) is an order-preserving mapping: if $\Gamma_1 \subseteq \Gamma_2$ then $K \oplus \Gamma_1 \subseteq K \oplus \Gamma_2$. To see this suppose $\Gamma_1 \subseteq \Gamma_2$. Then $\Gamma_1 \subseteq \Gamma_2 \subseteq K \oplus \Gamma_2$ (by Corollary 2.2.7) and since $K \oplus \Gamma_1$ is the smallest logic containing Γ_1 (Corollary 2.2.7) we have $K \oplus \Gamma_1 \subseteq K \oplus \Gamma_2$. Therefore the logic K, which is defined as $K \oplus \varnothing$, is the bottom of $\mathbf{N}$ since for any logic Ψ we have $\varnothing \subseteq \Psi$ and $K \oplus \Psi = \Psi$ (Corollary 2.2.7). The logic $Form(\Phi)$ is the top of $\mathbf{N}$.

2.2.9 Definition. A logic Λ is *finitely axiomatisable* if there exists a $\Gamma \subseteq_{fin} Form(\Phi)$ such that $\Lambda = K \oplus \Gamma$.

We now look at the notion of Λ-deduction.

2.2.10 Definition. Let Λ be a logic, let $\Gamma \subseteq Form(\Phi)$ and let $\varphi \in Form(\Phi)$. The formula φ is Λ-*deducible from* Γ if $\vdash_\Lambda \varphi$ or there exists $\psi_1, ..., \psi_n \in \Gamma$ such that $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \varphi$. This is denoted by $\Gamma \vdash_\Lambda \varphi$.

2.2.11 Example. Recall $\Gamma = \{\Diamond p, \Box\Diamond p\}$ from Example 2.1.23. We shall show $\Gamma \vdash_K \Diamond\Diamond p$. Also recall the $\Diamond$-generalization ($\Diamond$-gen) rule from Remark 2.2.4.

(1)	$\vdash_K p \to \top$	tautology
(2)	$\vdash_K \top \to (\neg p \vee p)$	tautology
(3)	$\vdash_K \top \to (\neg\Diamond p \vee \Diamond p)$	u.s.: 2
(4)	$\vdash_K p \to (\neg\Diamond p \vee \Diamond p)$	prop. logic: 1,3
(5)	$\vdash_K \Diamond p \to \Diamond(\neg\Diamond p \vee \Diamond p)$	$\Diamond$-gen: 4
(6)	$\vdash_K \neg\Diamond p \vee \Diamond(\neg\Diamond p \vee \Diamond p)$	$\to$ defn.: 5
(7)	$\vdash_K \Diamond(p \vee q) \leftrightarrow (\Diamond p \vee \Diamond q)$	thm 2.2.3
(8)	$\vdash_K \Diamond(\neg\Diamond p \vee \Diamond p) \leftrightarrow (\Diamond\neg\Diamond p \vee \Diamond\Diamond p)$	u.s.: 7
(9)	$\vdash_K (p \leftrightarrow q) \to ((r \vee p) \leftrightarrow (r \vee q))$	tautology
(10)	$\vdash_K (\Diamond(\neg\Diamond p \vee \Diamond p) \leftrightarrow (\Diamond\neg\Diamond p \vee \Diamond\Diamond p)) \to$	
	$((\neg\Diamond p \vee \Diamond(\neg\Diamond p \vee \Diamond p)) \leftrightarrow (\neg\Diamond p \vee (\Diamond\neg\Diamond p \vee \Diamond\Diamond p)))$	u.s.: 9
(11)	$\vdash_K (\neg\Diamond p \vee \Diamond(\neg\Diamond p \vee \Diamond p)) \leftrightarrow (\neg\Diamond p \vee (\Diamond\neg\Diamond p \vee \Diamond\Diamond p))$	m.p.: 8,10
(12)	$\vdash_K \neg\Diamond p \vee \Diamond\neg\Diamond p \vee \Diamond\Diamond p$	m.p.: 6,11
(13)	$\vdash_K \neg\neg p \leftrightarrow p$	tautology
(14)	$\vdash_K \neg\Diamond p \vee \neg\neg\Diamond\neg\Diamond p \vee \Diamond\Diamond p$	u.s.: 13; then
		u.s. and m.p.: 9,12
(15)	$\vdash_K \neg\Diamond p \vee \neg\Box\Diamond p \vee \Diamond\Diamond p$	$\Box$ defn.: 14
(16)	$\vdash_K \neg\neg(\neg p \vee \neg q) \leftrightarrow (\neg p \vee \neg q)$	u.s.: 13
(17)	$\vdash_K \neg(p \wedge q) \leftrightarrow (\neg p \vee \neg q)$	$\wedge$ defn.: 16
(18)	$\vdash_K (\neg\Diamond p \vee \neg\Box\Diamond p) \leftrightarrow \neg(\Diamond p \wedge \Box\Diamond p)$	u.s.: 17
(19)	$\vdash_K (\neg\Diamond p \vee \neg\Box\Diamond p \vee \Diamond\Diamond p) \leftrightarrow (\neg(\Diamond p \wedge \Box\Diamond p) \vee \Diamond\Diamond p)$	u.s. and m.p.: 9,18
(20)	$\vdash_K \neg(\Diamond p \wedge \Box\Diamond p) \vee \Diamond\Diamond p$	m.p.: 15,19
(21)	$\vdash_K (\Diamond p \wedge \Box\Diamond p) \to \Diamond\Diamond p$	$\to$ defn.: 20

Note that in lines 14 and 19 we skipped a few steps; these are similar to the steps from lines 9–12. Also note that at some points we have taken the associativity and commutativity of $\vee$ for granted; these properties can be shown by a uniform substitution on a propositional tautology. These omissions were to prevent the derivation from becoming longer.

It is no mere coincidence that we have $\Gamma \vdash_K \Diamond\Diamond p$ and $\Gamma \Vdash_{\mathsf{F}} \Diamond\Diamond p$ (Example 2.1.23, where F is the class of all frames). We shall see in the next section (Theorem 2.3.18) that K is strongly complete with respect to F i.e. $\Sigma \Vdash_F \varphi$ implies $\Sigma \vdash_K \varphi$ for all $\Sigma \subseteq Form(\Phi)$ and all $\varphi \in Form(\Phi)$.

2.2.12 Definition. Let $\Gamma \subseteq Form(\Phi)$ and let Λ be a logic. Γ is said to be Λ-*inconsistent* if $\Gamma \vdash_\Lambda \bot$; otherwise Γ is said to be Λ-*consistent*.

2.2.13 Remark. Let Λ be a logic, let F be a non-empty class of frames such that $\Lambda \subseteq \Lambda_F$ (recall from Remark 2.2.2 that $\Lambda_F = \{\varphi \in Form(\Phi) \mid F \Vdash \varphi\}$) and let $\mathcal{M}$ be a model based on a frame $\mathcal{F} = (W, R)$ in F with $w \in W$. We shall show that the set $\Gamma = \{\varphi \in Form(\Phi) \mid \mathcal{M}, w \Vdash \varphi\}$ is a Λ-consistent set. Suppose Γ is Λ-inconsistent; thus either $\vdash_\Lambda \bot$ or there exists $\psi_1, ..., \psi_n \in \Gamma$ such that $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \bot$. The former case causes a contradiction since $\Lambda \subseteq \Lambda_F$ gives that $\bot$ is a validity on the non-empty frame class F. For the latter case we also derive a contradiction and thus show Γ to be Λ-consistent:

(1)	$\vdash_\Lambda$	$(\psi_1 \wedge ... \wedge \psi_n) \to \bot$	
(2)	$\vdash_\Lambda$	$(p \to q) \to (\neg q \to \neg p)$	tautology
(3)	$\vdash_\Lambda$	$((\psi_1 \wedge ... \wedge \psi_n) \to \bot) \to (\top \to \neg(\psi_1 \wedge ... \wedge \psi_n))$	u.s.: 2
(4)	$\vdash_\Lambda$	$\top \to \neg(\psi_1 \wedge ... \wedge \psi_n)$	m.p.: 1,3
(5)	$\vdash_\Lambda$	$\neg\neg p \leftrightarrow p$	tautology
(6)	$\vdash_\Lambda$	$\neg\neg(\neg p_1 \vee ... \vee \neg p_n) \leftrightarrow (\neg p_1 \vee ... \vee \neg p_n)$	u.s.: 5
(7)	$\vdash_\Lambda$	$\neg(p_1 \wedge ... \wedge p_n) \leftrightarrow (\neg p_1 \vee ... \vee \neg p_n)$	$\wedge$ defn.: 6
(8)	$\vdash_\Lambda$	$\neg(\psi_1 \wedge ... \wedge \psi_n) \leftrightarrow (\neg\psi_1 \vee ... \vee \neg\psi_n)$	u.s.: 7
(9)	$\vdash_\Lambda$	$\top \to (\neg\psi_1 \vee ... \vee \neg\psi_n)$	prop. logic: 4,8
(10)	$\vdash_\Lambda$	$\top$	tautology
(11)	$\vdash_\Lambda$	$\neg\psi_1 \vee ... \vee \neg\psi_n$	m.p.: 9,10.

Consider the statement $\vdash_\Lambda \neg\psi_1 \vee ... \vee \neg\psi_n$ from line 11 of the above derivation. Since $\Lambda \subseteq \Lambda_F$ we thus have that $w \Vdash \neg\psi_1 \vee ... \vee \neg\psi_n$ and therefore $w \nVdash \psi_i$ for some $i \in \{1, .., n\}$. However $\psi_i \in \Gamma = \{\varphi \in Form(\Phi) \mid \mathcal{M}, w \Vdash \varphi\}$ which gives $w \Vdash \psi_i$, a contradiction. Therefore Γ is Λ-consistent.

2.2.14 Theorem. *Let $\Gamma \subseteq Form(\Phi)$ and let Λ be a logic. The following are equivalent:*

(*i*) Γ *is Λ-inconsistent;*
(*ii*) *there exists $\psi \in Form(\Phi)$ such that $\Gamma \vdash_\Lambda \psi \wedge \neg\psi$;*
(*iii*) $\Gamma \vdash_\Lambda \varphi$ *for all $\varphi \in Form(\Phi)$.*

Proof. (*i*) $\implies$ (*ii*) Let $\Gamma \vdash_\Lambda \bot$. Therefore either $\vdash_\Lambda \bot$ or $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \bot$ for $\psi_1, ..., \psi_n \in \Gamma$. If $\vdash_\Lambda \bot$ then (*ii*) is true since we have the tautology $\vdash_\Lambda \bot \to (p \wedge \neg p)$ which shows that for any $\varphi \in Form(\Phi)$ we will have $\vdash_\Lambda \varphi \wedge \neg\varphi$ and in particular the result $\Gamma \vdash_\Lambda \varphi \wedge \neg\varphi$ will hold. Suppose alternatively that $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \bot$ for $\psi_1, ..., \psi_n \in \Gamma$. Let $\chi = \psi_1 \wedge ... \wedge \psi_n$, we thus have:

(1)	$\vdash_\Lambda$	$\chi \to \bot$	
(2)	$\vdash_\Lambda$	$\neg\chi$	see derivation in Remark 2.2.13
(3)	$\vdash_\Lambda$	$p \to (q \to (p \wedge q))$	tautology
(4)	$\vdash_\Lambda$	$\neg\chi \to (\chi \to (\neg\chi \wedge \chi))$	u.s.: 3
(5)	$\vdash_\Lambda$	$\chi \to (\neg\chi \wedge \chi)$	m.p.: 2,4
(6)	$\vdash_\Lambda$	$(\psi_1 \wedge ... \wedge \psi_n) \to (\neg\chi \wedge \chi)$	χ defn.: 5.

We thus have that χ satisfies the condition from (ii).
$(ii) \implies (iii)$ Let $\Gamma \vdash_\Lambda \psi \wedge \neg\psi$ for some $\psi \in Form(\Phi)$ and let $\varphi \in Form(\Phi)$. We will show that $\Gamma \vdash_\Lambda \varphi$. If $\vdash_\Lambda \psi \wedge \neg\psi$ then $\vdash_\Lambda \bot$ and the result is immediate (see lines 4 and 5 below). Suppose $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to (\psi \wedge \neg\psi)$ where $\{\psi_1, ..., \psi_n\} \subseteq \Gamma$:

(1)	$\vdash_\Lambda$	$(\psi_1 \wedge ... \wedge \psi_n) \to (\psi \wedge \neg\psi)$	
(2)	$\vdash_\Lambda$	$(p \wedge \neg p) \to \bot$	tautology
(3)	$\vdash_\Lambda$	$(\psi \wedge \neg\psi) \to \bot$	u.s.: 2
(4)	$\vdash_\Lambda$	$\bot \to p$	tautology
(5)	$\vdash_\Lambda$	$\bot \to \varphi$	u.s.: 4
(6)	$\vdash_\Lambda$	$(\psi_1 \wedge ... \wedge \psi_n) \to \varphi$	prop. logic: 1,3,5.

$(iii) \implies (i)$ is trivial. □

2.3 Completeness with respect to frame classes

It is natural to ask whether our syntactic and semantic descriptions of validities match up i.e. given a logic Λ is there a frame class F such that $\Lambda = \Lambda_F$?

2.3.1 Definition. Let Λ be a logic and let F be a frame class. The logic Λ is *sound with respect to* F if $\Lambda \subseteq \Lambda_F$.

2.3.2 Remark. Let $\Gamma \subseteq Form(\Phi)$ and let F be a frame class. Suppose we want to show $K \oplus \Gamma \subseteq \Lambda_F$. Since $K \oplus \Gamma$ is the smallest logic containing Γ (Corollary 2.2.7), it suffices to show that Λ_F is a logic containing Γ. But by Remark 2.2.2 we already have that Λ_F is a logic. Thus to prove $K \oplus \Gamma \subseteq \Lambda_F$ it suffices to only show $\Gamma \subseteq \Lambda_F$.

2.3.3 Example. Let $D = \Box p \to \Diamond p$ and define the class of right unbounded frames as $ruFrm = \{(W, R) \mid (\forall w \in W)(\exists w_1 \in W)(Rww_1)\}$. We shall show $K \oplus \{D\} \subseteq \Lambda_{ruFrm}$. From the previous remark all we need to check is that $ruFrm \Vdash D$. To this end let $\mathcal{F} = (W, R) \in ruFrm$, let $w \in W$

and let $V : \Phi \to \mathcal{P}(W)$. Suppose $(\mathcal{F}, V), w \Vdash \Box p$. Since $\mathcal{F} \in ruFrm$ there exists $w_1 \in W$ such that Rww_1. Thus $w_1 \Vdash p$ and thus $w \Vdash \Diamond p$. Therefore $w \Vdash D$ and $ruFrm \Vdash D$.

2.3.4 Example. Let F be a frame class. Since $K = K \oplus \varnothing$ is the bottom of the logic lattice **N** (Remark 2.2.8) and Λ_F is a logic (Remark 2.2.2), we have that K is sound with respect F.

2.3.5 Definition. Let Λ be a logic and let F be a frame class. The logic Λ is *strongly complete with respect to* F if $\Gamma \Vdash_F \varphi$ implies $\Gamma \vdash_\Lambda \varphi$ for all $\varphi \in Form(\Phi)$ and all $\Gamma \subseteq Form(\Phi)$. When $F \Vdash \varphi$ implies $\vdash_\Lambda \varphi$ for all φ in $Form(\Phi)$, the logic Λ is said to be *weakly complete* or simply *complete with respect to* F.

2.3.6 Remark. It can be checked that strong completeness implies weak completeness.

Now given a logic Λ and a frame class F, we can rephrase the statement $\Lambda = \Lambda_F$ as: Λ is sound and complete with respect to F; the logic Λ is said to be *characterised by* F. Demonstrating $\Lambda \subseteq \Lambda_F$ comes down to showing that the Λ axioms are validities on F (Remark 2.3.2). However, showing that Λ is complete with respect to F is in general not as straightforward as demonstrating soundness. The following simple theorem will translate the problem of showing strong-completeness to a model satisfaction problem.

2.3.7 Theorem. *Let Λ be a logic and let F be a frame class. Λ is strongly complete with respect to F iff every Λ-consistent set of formulae Γ is satisfiable in a model $\mathcal{M}$ based on a frame $\mathcal{F}$ in F.*

Proof. $\implies$ We prove the contrapositive. Let Γ be a Λ-consistent set of formulae that is not satisfiable in any model based on a frame in F. Since Γ is Λ-consistent there exists a $\varphi \in Form(\Phi)$ such that $\Gamma \nvdash_\Lambda \varphi$ (Theorem 2.2.14). Recall from Definition 2.1.22 that $\Gamma \Vdash_F \varphi$ if $\mathcal{M}, w \Vdash \Gamma$ implies $\mathcal{M}, w \Vdash \varphi$ for all models $\mathcal{M}$ based on a frame in F and all points w in $\mathcal{M}$. Since Γ is not satisfiable in any model based on a frame in F we then vacuously have $\Gamma \Vdash_F \varphi$. Thus Λ is not strongly complete with respect to F since we have $\Gamma \Vdash_F \varphi$ and $\Gamma \nvdash_\Lambda \varphi$.
$\impliedby$ We prove the contrapositive. Suppose Λ is not strongly complete with respect to F. Then there exists a $\Gamma \cup \{\varphi\} \subseteq Form(\Phi)$ such that $\Gamma \Vdash_F \varphi$ and $\Gamma \nvdash_\Lambda \varphi$. The set $\Gamma \cup \{\neg\varphi\}$ is Λ-consistent. To see this suppose towards a contradiction that $\Gamma \cup \{\neg\varphi\}$ is Λ-inconsistent. Therefore $\Gamma \cup \{\neg\varphi\} \vdash_\Lambda \varphi$ (Theorem 2.2.14) and thus either $\vdash_\Lambda \varphi$ or $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \varphi$ for $\psi_1, ..., \psi_n \in \Gamma \cup \{\neg\varphi\}$. The former case ($\vdash_\Lambda \varphi$) leads to a contradiction since $\Gamma \nvdash_\Lambda \varphi$. In the latter case there are two scenarios: either $\psi_1, ..., \psi_n \subseteq \Gamma$ or $\psi_i = \neg\varphi$ for some $i \in \{1, ..., n\}$. If $\psi_1, ..., \psi_n \subseteq \Gamma$ then this contradicts $\Gamma \nvdash_\Lambda \varphi$. Suppose then that $\psi_i = \neg\varphi$ for some $i \in \{1, ..., n\}$ then we have $\vdash_\Lambda (\psi_1 \wedge ... \wedge \neg\varphi \wedge ... \wedge \psi_n) \to \varphi$ i.e. $\vdash_\Lambda (\chi \wedge \neg\varphi) \to \varphi$ where $\chi = \bigwedge_{k \neq i, k=1}^{n} \psi_k$. By definition of $\to$ we then have $\vdash_\Lambda \neg(\chi \wedge \neg\varphi) \vee \varphi$; it can be shown that this implies $\vdash_\Lambda \neg\chi \vee \varphi \vee \varphi$ which in turn implies $\vdash_\Lambda \chi \to \varphi$, a contradiction since $\Gamma \nvdash_\Lambda \varphi$. Thus $\Gamma \cup \{\neg\varphi\}$ is Λ-consistent. The Λ-consistent set of formulae $\Gamma \cup \{\neg\varphi\}$ is not satisfiable in any model based on a frame in F since if it were, we would arrive at a contradiction as $\Gamma \Vdash_F \varphi$. $\square$

Motivated by this theorem we begin the construction of a model on which all of the Λ-consistent sets of formulae will be satisfied, a canonical model [26].

2.3.8 Definition. Let Λ be a logic and let $\Gamma \subseteq Form(\Phi)$. The set Γ is a Λ-*maximal consistent set* (Λ-MCS) if Γ is a Λ-consistent set and $\Gamma \subseteq \Sigma$ implies $\Gamma = \Sigma$ for all Λ-consistent sets Σ.

2.3.9 Example. Recall Remark 2.2.13: Λ is a logic, F is a class of frames such that $\Lambda \subseteq \Lambda_F$, $\mathcal{M}$ is a model based on a frame $\mathcal{F} = (W, R)$ in F with $w \in W$ and $\Gamma = \{\varphi \mid w \Vdash \varphi\}$. We shall show Γ is a Λ-MCS. Γ has already been shown to be Λ-consistent. In order to show Γ is a Λ-MCS we will show that any set of formulae that properly contains Γ must be Λ-inconsistent. Let $\Sigma \subseteq Form(\Phi)$ such that $\Gamma \subset \Sigma$ and let $\varphi \in \Sigma \setminus \Gamma$. We have $w \nVdash \varphi$ and thus $w \Vdash \neg\varphi$. Thus $\neg\varphi \in \Gamma$. But $\Gamma \subset \Sigma$ thus $\neg\varphi \in \Sigma$. Therefore $\Sigma \vdash_\Lambda \bot$ since $\vdash_\Lambda (\varphi \wedge \neg\varphi) \to \bot$. Thus Σ is Λ-inconsistent.

The following result shows that any Λ-consistent set can be extended to a Λ-MCS.

2.3.10 Lemma. *(Lindenbaum's Lemma)* Let Λ be a logic and let Σ be a Λ-consistent set. There exists a Λ-MCS Σ^+ such that $\Sigma \subseteq \Sigma^+$.

Proof. Consider $\Sigma^+ = \bigcup_{n=0}^{\infty} \Sigma_n$ where Σ_0 is defined as Σ and Σ_n is defined inductively as:

$$\Sigma_{n+1} = \begin{cases} \Sigma_n \cup \{\varphi_n\} & \text{if this is } \Lambda\text{-consistent} \\ \Sigma_n \cup \{\neg\varphi_n\} & \text{otherwise} \end{cases}$$

where $\varphi_1, \varphi_2, ...$ is an enumeration of $Form(\Phi)$ (this is possible since we assume a countable Φ). We have that Σ_n is Λ-consistent for all n. The proof comes down to the fact that if Σ_n is Λ-consistent and $\Sigma_n \cup \{\varphi_n\}$ is Λ-inconsistent then $\Sigma_n \cup \{\neg\varphi_n\}$ is Λ-consistent. We shall now look at why this is true. Since $\Sigma_n \cup \{\varphi_n\}$ is Λ-inconsistent we have $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_l) \to \bot$ for $\psi_1, ..., \psi_l \in \Sigma_n \cup \{\varphi_n\}$. There exists a $j \in \{1, ..., l\}$ such that $\psi_j = \varphi_n$ (otherwise Σ_n would be Λ-inconsistent). It can be shown that this leads to $\vdash_\Lambda \bigwedge_{i \neq j, i=1}^{l} \psi_i \to \neg\varphi_n$ i.e. $\Sigma_n \vdash_\Lambda \neg\varphi_n$. Suppose towards a contradiction that $\Sigma_n \cup \{\neg\varphi_n\}$ is Λ-inconsistent i.e. $\vdash_\Lambda (\sigma_1 \wedge ... \wedge \sigma_m) \to \bot$ where $\sigma_1, ..., \sigma_m \in \Sigma_n \cup \{\neg\varphi_n\}$. By the same reasoning as before we have $\sigma_k = \neg\varphi_n$ for some $k \in \{1, ..., m\}$ and thus $\vdash_\Lambda \bigwedge_{i \neq k, i=1}^{m} \sigma_i \to \varphi_n$ i.e. $\Sigma_n \vdash_\Lambda \varphi_n$. But we also have $\Sigma_n \vdash_\Lambda \neg\varphi_n$ i.e. Σ_n is Λ-inconsistent (Theorem 2.2.14), a contradiction.

Thus given Σ_n is Λ-consistent for all n, we have Σ^+ is Λ-consistent. To see this suppose Σ^+ is Λ-inconsistent i.e. $\vdash_\Lambda (\psi_1 \wedge ... \wedge \psi_n) \to \bot$ for $\psi_1, ..., \psi_n \in \Sigma^+$. By definition of Σ^+ we have for all $i \in \{1, ..., n\}$ that $\psi_i \in \Sigma_{\mathrm{f}(i)}$ for some $\mathrm{f} : \{1, ..., n\} \to \mathbb{N}$. Since $\Sigma_n \subseteq \Sigma_{n+1}$ for all n, we have $\psi_1, ..., \psi_n \in \Sigma_{\max_{1 \leq i \leq n} \mathrm{f}(i)}$ and thus $\bot$ is Λ-deducible from $\Sigma_{\max_{1 \leq i \leq n} \mathrm{f}(i)}$. This causes a contradiction since $\Sigma_{\max_{1 \leq i \leq n} \mathrm{f}(i)}$ is Λ-consistent.

To see Σ^+ is a Λ-MCS consider $\Sigma^+ \subset \Gamma \subseteq Form(\Phi)$ with $\chi \in \Gamma \setminus \Sigma^+$ and suppose the formula χ is equal to φ_i in the enumeration of $Form(\Phi)$. By the definition of Σ_{i+1} we have:

$$\Sigma_{i+1} = \begin{cases} \Sigma_i \cup \{\varphi_i\} & \text{if this is } \Lambda\text{-consistent} \\ \Sigma_i \cup \{\neg\varphi_i\} & \text{otherwise.} \end{cases}$$

The set $\Sigma_i \cup \{\varphi_i\}$ is Λ-inconsistent else otherwise we would have $\Sigma_{i+1} = \Sigma_i \cup \{\varphi_i\}$ which would contradict the fact that $\varphi_i \notin \Sigma^+$. Thus $\Sigma_{i+1} = \Sigma_i \cup \{\neg\varphi_i\}$ and therefore $\neg\varphi_i \in \Sigma^+ \subset \Gamma$. From this we deduce that $\Gamma \vdash_\Lambda \varphi_i \wedge \neg\varphi_i$ i.e. Γ is Λ-inconsistent. □

2.3.11 Theorem. *Let Λ be a logic and let Σ be a Λ-MCS. The following hold:*

- (i) *for all $\varphi \in Form(\Phi)$ either $\varphi \in \Sigma$ or $\neg\varphi \in \Sigma$;*
- (ii) *Σ is closed under modus ponens;*
- (iii) $\Lambda \subseteq \Sigma$;
- (iv) *for all $\varphi, \psi \in Form(\Phi)$ we have $\varphi \vee \psi \in \Sigma$ iff $\varphi \in \Sigma$ or $\psi \in \Sigma$.*

Proof. (i) Let $\varphi \in Form(\Phi)$. Consider the set $\Sigma \cup \{\varphi\}$. If $\Sigma \cup \{\varphi\}$ is Λ-consistent then, since Σ is a Λ-MCS, we would have $\Sigma = \Sigma \cup \{\varphi\}$ and thus $\varphi \in \Sigma$. If $\Sigma \cup \{\varphi\}$ is Λ-inconsistent then $\Sigma \cup \{\neg\varphi\}$ is Λ-consistent (see proof of Lemma 2.3.10) and similarly $\neg\varphi \in \Sigma$.
(ii) Let $\varphi, \psi \in Form(\Phi)$ and let $\varphi, \varphi \to \psi \in \Sigma$. If $\psi \in \Sigma$ then we are done. Suppose towards a contradiction that $\psi \notin \Sigma$. Thus by (i) we have $\neg\psi \in \Sigma$. Now, as Λ is a logic (and consequently contains all propositional tautologies) we know that $\vdash_\Lambda (p \wedge \neg p) \to \bot$ which leads to the fact that $\vdash_\Lambda ((\varphi \wedge \neg\psi) \wedge \neg(\varphi \wedge \neg\psi)) \to \bot$ under uniform substitution. It is also straightforward to show that $\vdash_\Lambda \neg(\varphi \wedge \neg\psi) \leftrightarrow (\varphi \to \psi)$. Thus we have both that $\varphi, \neg\psi, \varphi \to \psi \in \Sigma$ and $\vdash_\Lambda ((\varphi \wedge \neg\psi) \wedge (\varphi \to \psi)) \to \bot$ which is a contradiction to Σ being Λ-consistent.
(iii) Let $\vdash_\Lambda \varphi$. If $\varphi \in \Sigma$ then we are done. Suppose towards a contradiction that $\varphi \notin \Sigma$. By statement (i) we then have $\neg\varphi \in \Sigma$. As Λ is a logic we know that $\vdash_\Lambda p \leftrightarrow (p \vee (p \wedge \neg p))$ which leads to $\vdash_\Lambda \varphi \leftrightarrow (\varphi \vee (\varphi \wedge \neg\varphi))$ under uniform substitution. Since $\vdash_\Lambda \varphi$ it follows that $\vdash_\Lambda \varphi \vee (\varphi \wedge \neg\varphi)$. It is straightforward to demonstrate that $\vdash_\Lambda (\varphi \vee (\varphi \wedge \neg\varphi)) \leftrightarrow (\neg\varphi \to (\varphi \wedge \neg\varphi))$. Therefore we have both that $\neg\varphi \in \Sigma$ and $\vdash_\Lambda \neg\varphi \to (\varphi \wedge \neg\varphi)$ which, by Theorem 2.2.14, implies that Σ is Λ-inconsistent, a contradiction.
(iv) Let $\varphi, \psi \in Form(\Phi)$. Suppose $\psi \vee \varphi \notin \Sigma$. We will show $\neg\varphi \in \Sigma$ and $\neg\psi \in \Sigma$. By statement (i) we have $\neg(\psi \vee \varphi) \in \Sigma$. But $\vdash_\Lambda \neg(\psi \vee \varphi) \leftrightarrow (\neg\varphi \wedge \neg\psi)$; due to the fact that $\Lambda \subseteq \Sigma$ and Σ is closed under modus ponens (statements (ii), (iii)) we have that $(\neg\varphi \wedge \neg\psi) \in \Sigma$. Since $\vdash_\Lambda (\neg\varphi \wedge \neg\psi) \to \neg\varphi$ and $\vdash_\Lambda (\neg\varphi \wedge \neg\psi) \to \neg\psi$ we have $\neg\varphi \in \Sigma$ and $\neg\psi \in \Sigma$ again by statements (ii) and (iii). The reverse implication follows by similar reasoning. □

2.3.12 Definition. Let Λ be a logic. The *canonical model for* Λ is $\mathcal{M}^\Lambda = (W^\Lambda, R^\Lambda, V^\Lambda)$ where $W^\Lambda = \{\Sigma \subseteq Form(\Phi) \mid \Sigma \text{ is a } \Lambda\text{-}MCS\}$, $V^\Lambda(p) = \{\Sigma \in W^\Lambda \mid p \in \Sigma\}$ and $R^\Lambda\Gamma_1\Gamma_2$ iff $\varphi \in \Gamma_2$ implies $\Diamond\varphi \in \Gamma_1$ for all $\varphi \in Form(\Phi)$. The *canonical frame for* Λ is $\mathcal{F}^\Lambda = (W^\Lambda, R^\Lambda)$.

2.3.13 Remark. V^Λ and R^Λ are referred to as the *canonical valuation* and *canonical relation* respectively.

We will not state the full proofs of the following two lemmas. Their complete proofs can be found in [1].

2.3.14 Lemma. Let Λ be a logic; $R^\Lambda\Gamma_1\Gamma_2$ iff $\Box\varphi \in \Gamma_1$ implies $\varphi \in \Gamma_2$ for all $\varphi \in Form(\Phi)$.

Proof. The proof uses the definition of R^Λ (Definition 2.3.12) and the properties of Λ-MCS from Theorem 2.3.11. □

2.3.15 Lemma. Let Λ be a logic, let $\Gamma \in W^\Lambda$ and let $\varphi \in Form(\Phi)$. If $\Diamond\varphi \in \Gamma$ then there exists $\Sigma^+ \in W^\Lambda$ such that $R^\Lambda\Gamma\Sigma^+$ and $\varphi \in \Sigma^+$.

Proof. Let $\Diamond\varphi \in \Gamma$. Consider the set $\Sigma = \{\varphi\} \cup \{\psi \mid \Box\psi \in \Gamma\}$. The fact that $\Diamond\varphi \in \Gamma$ can be used to check that Σ is Λ-consistent. From Lindenbaum's Lemma (Lemma 2.3.10) there exists $\Sigma^+ \in W^\Lambda$ such that $\Sigma \subseteq \Sigma^+$. We have by the construction of Σ that $\varphi \in \Sigma^+$ and $R^\Lambda\Gamma\Sigma^+$ (Lemma 2.3.14). □

The preceding results seem to suggest that in a canonical model a formula is satisfied at a Λ-MCS if and only if it is an element of the Λ-MCS. We will now give a formal proof of this. As with many results that we explore, the theorem is proved by induction on the length of formulae.

2.3.16 Theorem. *Let Λ be a logic and let Γ be a Λ-MCS. For all $\psi \in Form(\Phi)$: $\mathcal{M}^\Lambda, \Gamma \Vdash \psi$ iff $\psi \in \Gamma$.*

Proof. We want to show for all $\psi \in Form(\Phi)$ that:

$$\mathcal{M}^\Lambda, \Gamma \Vdash \psi \text{ iff } \psi \in \Gamma. \tag{2.3.1}$$

Base case: Suppose $\mathrm{k}(\psi) = 0$[1]; then $\psi = p \in \Phi$ or $\psi = \bot$. In the former case $\mathcal{M}^\Lambda, \Gamma \Vdash p$ iff $\Gamma \in V^\Lambda(p)$ iff $p \in \Gamma$ by definition of the canonical valuation (Definition 2.3.12). In the latter case: since $\mathcal{M}^\Lambda, \Gamma \nVdash \bot$ and $\bot \notin \Gamma$ (as Γ is Λ-consistent) we have $\mathcal{M}^\Lambda, \Gamma \Vdash \bot$ iff $\bot \in \Gamma$.

Inductive hypothesis: Assume statement 2.3.1 is true for all $\psi \in Form(\Phi)$ such that $\mathrm{k}(\psi) < n$.

[1] $\mathrm{k}(\psi)$ returns the number of connectives in ψ.

Inductive step: Let $\mathrm{k}(\psi) = n$. Thus for $\varphi, \varphi_1, \varphi_2 \in Form(\Phi)$ we have either:

$$\psi = \begin{cases} \varphi_1 \vee \varphi_2 & (i) \\ \neg\varphi & (ii) \\ \Diamond\varphi & (iii). \end{cases}$$

(i) $\mathcal{M}^\Lambda, \Gamma \Vdash \varphi_1 \vee \varphi_2$ iff $\mathcal{M}^\Lambda, \Gamma \Vdash \varphi_1$ or $\mathcal{M}^\Lambda, \Gamma \Vdash \varphi_2$. From the inductive hypothesis $\Gamma \Vdash \varphi_1$ or $\Gamma \Vdash \varphi_2$ iff $\varphi_1 \in \Gamma$ or $\varphi_2 \in \Gamma$. From part (iv) of Theorem 2.3.11 $\varphi_1 \in \Gamma$ or $\varphi_2 \in \Gamma$ iff $\varphi_1 \vee \varphi_2 \in \Gamma$.
(ii) $\mathcal{M}^\Lambda, \Gamma \Vdash \neg\varphi$ iff $\mathcal{M}^\Lambda, \Gamma \nVdash \varphi$. From the inductive hypothesis $\Gamma \nVdash \varphi$ iff $\varphi \notin \Gamma$. Part (i) of Theorem 2.3.11 gives $\varphi \notin \Gamma$ iff $\neg\varphi \in \Gamma$ ($\neg\varphi \in \Gamma$ implies $\varphi \notin \Gamma$ since otherwise Γ would be Λ-inconsistent).
(iii) $\Longrightarrow$ Let $\mathcal{M}^\Lambda, \Gamma \Vdash \Diamond\varphi$. Then there exists $\Gamma' \in W^\Lambda$ such that $R^\Lambda\Gamma\Gamma'$ and $\mathcal{M}^\Lambda, \Gamma' \Vdash \varphi$. By the inductive hypothesis $\varphi \in \Gamma'$. But since $R^\Lambda\Gamma\Gamma'$ this gives $\Diamond\varphi \in \Gamma$.
$\Longleftarrow$ Let $\Diamond\varphi \in \Gamma$. By Lemma 2.3.15 there exists $\Gamma' \in W^\Lambda$ such that $R^\Lambda\Gamma\Gamma'$ and $\varphi \in \Gamma'$. By the inductive hypothesis we have $\Gamma' \Vdash \varphi$. Therefore $\Gamma \Vdash \Diamond\varphi$. □

2.3.17 Remark. The case when ψ is $\top$ or $\Box\varphi$ falls under case (ii) since by definition $\top = \neg\bot$ and $\Box\varphi = \neg\Diamond\neg\varphi$. When $\psi = \varphi_1 \wedge \varphi_2$, this falls under cases (i) and (ii) since $\varphi_1 \wedge \varphi_2 = \neg(\neg\varphi_1 \vee \neg\varphi_2)$.

2.3.18 Theorem. *(Canonical Model Theorem) Let Λ be a logic; then Λ is strongly complete with respect to $\{\mathcal{F}^\Lambda\}$.*

Proof. By Theorem 2.3.7 Λ is strongly complete with respect to $\{\mathcal{F}^\Lambda\}$ if and only if every Λ-consistent set is satisfiable in a model based on $\mathcal{F}^\Lambda$. Let Γ be a Λ-consistent set. Then by Lindenbaum's lemma (Lemma 2.3.10) there exists a Λ-MCS, Γ^+, extending Γ. Therefore by Theorem 2.3.16 $\mathcal{M}^\Lambda, \Gamma^+ \Vdash \Gamma$. □

The canonical model theorem shows that for a logic Λ and frame class F, if the canonical frame $\mathcal{F}^\Lambda$ is in F then Λ is strongly complete with respect F. With this we can see K is strongly complete with respect to the class of all basic frames F since $\mathcal{F}^K$ is a basic frame. This offers a pain-free alternative proof to Example 2.2.11. Here are another two worked examples that are amenable to the canonical method.

2.3.19 Example. Recall Example 2.3.3: $D = \Box p \to \Diamond p$ and $ruFrm = \{(W, R) \mid (\forall w \in W)(\exists w' \in W)(Rww')\}$. We shall show $K \oplus \{D\}$ is strongly complete with respect to $ruFrm$. From the above results, all we need to check is $\mathcal{F}^{K\oplus\{D\}} \in ruFrm$. Let $\Gamma \in W^{K\oplus\{D\}}$ and consider the set $S = \{\Gamma_1 \in W^{K\oplus\{D\}} \mid R^{K\oplus\{D\}}\Gamma\Gamma_1\}$. If $S \neq \varnothing$ then we are done. Suppose $S = \varnothing$ then vacuously $\Gamma \Vdash \Box p$. But $\Gamma \Vdash D$ (since $K \oplus \{D\} \subseteq \Gamma$ (Theorem 2.3.11) and $D \in \Gamma$ implies $\Gamma \Vdash D$ (Theorem 2.3.16)). Thus $\Gamma \Vdash \Diamond p$, which gives a contradiction to S being empty. Therefore together with the soundness result of Example 2.3.3 we now have $K \oplus \{D\} = \Lambda_{ruFrm}$.

2.3.20 Example. Let $1.1 = \Diamond p \to \Box p$ and define the class of partial function frames as $pfFrm = \{(W, R) \mid (\forall a, b, c \in W)(Rab \;\&\; Rac \implies b = c)\}$. We will show that $K \oplus \{1.1\}$ is strongly complete with respect to $pfFrm$. Let $R^{K\oplus\{1.1\}}\Gamma_1\Gamma_2$, $R^{K\oplus\{1.1\}}\Gamma_1\Gamma_3$ and $\varphi \in \Gamma_2$. We shall show $\varphi \in \Gamma_3$. We have by definition of $R^{K\oplus\{1.1\}}$ (see Definition 2.3.12) that $\Diamond\varphi \in \Gamma_1$ and thus $\Gamma_1 \Vdash \Diamond\varphi$ (Theorem 2.3.16). But $\Diamond\varphi \to \Box\varphi \in \Gamma_1$ ($\vdash_{K\oplus\{1.1\}} \Diamond\varphi \to \Box\varphi$ under uniform substitution of 1.1 and $K \oplus \{1.1\} \subseteq \Gamma_1$ by Theorem 2.3.11). Therefore $\Gamma_1 \Vdash \Diamond\varphi \to \Box\varphi$ (Theorem 2.3.16) and thus $\Gamma_1 \Vdash \Box\varphi$. This shows that $\Gamma_3 \Vdash \varphi$ and thus $\varphi \in \Gamma_3$ (Theorem 2.3.16). Therefore $\Gamma_2 \subseteq \Gamma_3$ and it can be similarly demonstrated that $\Gamma_3 \subseteq \Gamma_2$. Thus $\Gamma_2 = \Gamma_3$ and $\mathcal{F}^{K\oplus\{1.1\}} \in pfFrm$ giving $K \oplus \{1.1\}$ is strongly complete with respect to $pfFrm$. It can easily be checked that $1.1 \in \Lambda_{pfFrm}$ and thus $K \oplus \{1.1\}$ is sound with respect to $pfFrm$ (see Remark 2.3.2 for why this suffices for a soundness proof). We thus have $K \oplus \{1.1\} = \Lambda_{pfFrm}$.

The previous two examples have offered two logics that are sound and strongly complete with respect to a frame class. However there exist many logics lacking this property, for example $K \oplus \{L\}$ where L is the Löb axiom: $\Box(\Box p \to p) \to \Box p$. To see this suppose $K \oplus \{L\}$ is sound and strongly complete with respect to a frame class F. Thus every $K \oplus \{L\}$-consistent set of formulae is satisfiable in a model based on a frame in F (Theorem 2.3.7). It can be shown that there is a $K \oplus \{L\}$-consistent set that can only be satisfied in a model with infinite paths. But L forces only finite paths on any frame $\mathcal{F}$ for which $\mathcal{F} \Vdash L$, a contradiction. Thus $K \oplus \{L\}$ is not sound and strongly complete with respect to any frame class.

2.4 The finite model property

As noted in Section 1.2 a finitely axiomatisable logic (Definition 2.2.9) which has the finite model property is decidable [16].

2.4.1 Definition. Let Λ be a logic and let F be a frame class. The logic Λ is said to have the *finite model property (fmp) with respect to* F if the following property holds: $\vdash_\Lambda \varphi$ if and only if $F_{fin} \Vdash \varphi$ for all $\varphi \in Form(\Phi)$.

2.4.2 Remark. If Λ is sound and complete with respect to a frame class F such that $F \Vdash \varphi$ if and only if $F_{fin} \Vdash \varphi$ for all $\varphi \in Form(\Phi)$, then by definition Λ has the finite model property with respect to F.

Since, in the previous section, we already looked at methods of proving soundness and completeness, for the proof of the finite model property we will focus on how to show that that the validities of F and F_{fin} are the same i.e. $F_{fin} \Vdash \varphi$ implies $F \Vdash \varphi$ for all $\varphi \in Form(\Phi)$[2]. In particular we will look

[2] By definition we already have $F \Vdash \varphi$ implies $F_{fin} \Vdash \varphi$ for all $\varphi \in Form(\Phi)$.

at the contrapositive of $F_{fin} \Vdash \varphi$ implies $F \Vdash \varphi$, which states that if there is a model $\mathcal{M}$ based on a frame (W, R) in F with $w \in W$ such that $\mathcal{M}, w \Vdash \neg\varphi$ then there exists a model $\mathcal{M}'$ based on a frame (W', R') in F_{fin} with $w' \in W'$ such that $\mathcal{M}', w' \Vdash \neg\varphi$. We will now discuss a technique for proving this which involves constructing a finite model by identifying points. In particular we will look at a filtration of a model [26].

2.4.3 Definition. Let $\mathcal{M} = (W, R, V)$ and let $\Sigma \subseteq Form(\Phi)$. Define the relation $\sim_\Sigma$ on $\mathcal{M}$ as $\sim_\Sigma = \{(w, w_1) \in W^2 \mid (\forall\varphi \in \Sigma)(w \Vdash \varphi \iff w_1 \Vdash \varphi)\}$.

2.4.4 Remark. It is routine to check that $\sim_\Sigma$ is an equivalence relation on W.

2.4.5 Remark. The equivalence class of w under $\sim_\Sigma$ (the set $\{w' \in W \mid w \sim_\Sigma w'\}$) is denoted by $|w|_\Sigma$ or simply $|w|$ when Σ is inferred from the context.

2.4.6 Definition. Let $\Sigma \subseteq Form(\Phi)$ be a subformula closed set (Σ contains all of its subformulae) and let $\mathcal{M} = (W, R, V)$. A *filtration of* $\mathcal{M}$ *through* Σ is a model $\mathcal{M}^\Sigma = (W^\Sigma, R^\Sigma, V^\Sigma)$ where $W^\Sigma = W/{\sim_\Sigma}$, $V^\Sigma(p) = \{|w| \in W^\Sigma \mid w \in V(p)\}$ for $p \in \Sigma$, and R^Σ is defined to have the following properties:

(*i*) Rw_1w_2 implies $R^\Sigma|w_1||w_2|$;
(*ii*) $R^\Sigma|w_1||w_2|$ and $w_2 \Vdash \varphi$ imply $w_1 \Vdash \Diamond\varphi$ for all $\Diamond\varphi \in \Sigma$.

The filtration $\mathcal{M}^\Sigma$ is a model based on the frame $\mathcal{F}^\Sigma = (W^\Sigma, R^\Sigma)$.

2.4.7 Example. Given a subformula closed set Σ and model $\mathcal{M} = (W, R, V)$, it can easily be checked that the relation $R^{\Sigma,s}$ on W^Σ given by $R^{\Sigma,s}|w_1||w_2|$ iff there exists $u_1 \in |w_1|$ and $u_2 \in |w_2|$ such that Ru_1u_2, fulfils conditions (*i*) and (*ii*) of a filtration relation. Given any filtration relation R^Σ on W^Σ, it is immediate that $R^{\Sigma,s} \subseteq R^\Sigma$ by the definition of $R^{\Sigma,s}$ and property (*i*) of R^Σ. The largest filtration relation (relative to set inclusion) on W^Σ is $R^{\Sigma,l}|w_1||w_2|$ iff $w_2 \Vdash \varphi$ implies $w_1 \Vdash \Diamond\varphi$ for all $\Diamond\varphi \in \Sigma$.

2.4.8 Example. Consider the model $\mathcal{M}$ as shown:

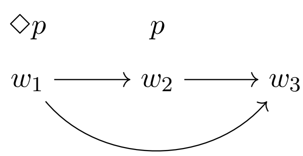

If we filtrate $\mathcal{M}$ through the subformula closed set $\Sigma = \{p, \Diamond p\}$ using the largest filtration $R^{\Sigma,l}$ then we get $\mathcal{M}^\Sigma$ as:

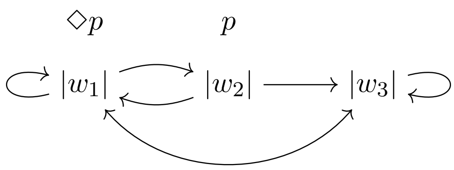

Most of the arrows are added vacuously, for instance: $R^{\Sigma,l}|w_2||w_1|$ since $w_1 \nVdash p$ and $\Diamond p$ is the only diamond formula in Σ, thus for all $\Diamond\varphi \in \Sigma$ we have $w_1 \Vdash \varphi$ implies $w_2 \Vdash \Diamond\varphi$. Notice how the filtration does not preserve the transitivity of $\mathcal{M}$; the filtration relation $R^{\Sigma,t}|w_1||w_2|$ iff $w_2 \Vdash \varphi \vee \Diamond\varphi$ implies $w_1 \Vdash \Diamond\varphi$ for all $\Diamond\varphi \in \Sigma$, will preserve the transitivity of $\mathcal{M}$.

Towards demonstrating that the filtration construction will prove the finite model property, we must first show that a filtration through a finite subformula closed set is finite.

2.4.9 Lemma. Given a model $\mathcal{M} = (W, R, V)$ and a subformula closed set Σ with $|\Sigma| = n$, we have $|W^\Sigma| \leq 2^n$.

Proof. Consider the map $m : W^\Sigma \to \mathcal{P}(\Sigma)$ by $m(|w|) = \{\varphi \in \Sigma \mid \mathcal{M}, w \Vdash \varphi\}$. It remains to check that m is a well-defined injective map and thus $|W^\Sigma| \leq |\mathcal{P}(\Sigma)| = 2^n$. The map m is a well-defined injective map since: $|w_1| = |w_2|$ iff $w_1 \sim_\Sigma w_2$ iff $\{\varphi \in \Sigma \mid \mathcal{M}, w_1 \Vdash \varphi\} = \{\varphi \in \Sigma \mid \mathcal{M}, w_2 \Vdash \varphi\}$ iff $m(|w_1|) = m(|w_2|)$. $\square$

We still have not checked whether the formulae in Σ that are satisfiable on $\mathcal{M}$ are also satisfiable on $\mathcal{M}^\Sigma$. The Definition 2.4.6 of R^Σ will suffice to prove this in an induction on the length of formulae; this is also where the subformula condition for Σ is necessary. We will state the result without proof (see Theorem 2.39 of [1]).

2.4.10 Theorem. *Let Σ be a subformula closed set and let $\mathcal{M} = (W, R, V)$ be a model. For all $w \in W$ and all $\varphi \in \Sigma$: $\mathcal{M}, w \Vdash \varphi$ iff $\mathcal{M}^\Sigma, |w| \Vdash \varphi$.*

Therefore suppose we have $\mathcal{M}$ based on a frame in a given frame class F with $\mathcal{M}, w \Vdash \varphi$. If we filtrate $\mathcal{M}$ through $\text{sub}(\varphi) = \{\psi \mid \psi \text{ subformula of } \varphi\}$ we will have a finite model $\mathcal{M}^\Sigma$ (since $\text{sub}(\varphi)$ is finite; Theorem 2.4.9) with $\mathcal{M}^\Sigma, |w| \Vdash \varphi$ (since $\varphi \in \text{sub}(\varphi)$; Theorem 2.4.10). All that remains to show the validities of F and F_{fin} are the same is to check that $\mathcal{M}^\Sigma$ is based on a frame in F, in particular, $\mathcal{F}^\Sigma \in F$.

2.4.11 Example. Define the class of Euclidean frames E=$\{(W, R) \mid (\forall w_1, w_2, w_3 \in W)(Rw_1w_2 \;\&\; Rw_1w_3 \implies Rw_2w_3)\}$. The following diagram shows a member of E.

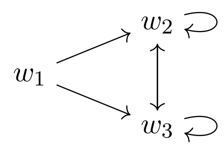

We shall show that the validities of E and E_{fin} are the same. Let $\mathcal{M}$ be a model based on a $\mathcal{F} = (W, R)$ in E with $\mathcal{M}, w \Vdash \varphi$. Consider the filtration of $\mathcal{M}$ through $\Sigma = \text{sub}(\varphi)$ by the relation

$$\begin{aligned} R^\Sigma|w_1||w_2| \quad \text{iff} \quad & w_2 \Vdash \psi \text{ implies } w_1 \Vdash \Diamond\psi \text{ for all } \Diamond\psi \in \Sigma; \\ & w_1 \Vdash \Diamond\psi \text{ implies } w_2 \Vdash \Diamond\psi \text{ for all } \Diamond\psi \in \Sigma; \\ & w_2 \Vdash \Diamond\Diamond\psi \text{ implies } w_2 \Vdash \Diamond\psi \text{ for all } \Diamond\Diamond\psi \in \Sigma. \end{aligned}$$

The relation R^Σ satisfies conditions (i) and (ii) of Definition 2.4.6. In particular to see that the relation satisfies condition (i), consider $w_1 \in W$ and $w_2 \in W$ such that Rw_1w_2. We will show that the three implications above are satisfied and hence that $R^\Sigma|w_1||w_2|$. By definition, if $w_2 \Vdash \psi$ then $w_1 \Vdash \Diamond\psi$ for all $\Diamond\psi \in \Sigma$. For the second implication suppose $w_1 \Vdash \Diamond\psi$ for some $\Diamond\psi \in \Sigma$. Therefore there exists $w \in W$ such that Rw_1w and $w \Vdash \psi$. Since Rw_1w_2, Rw_1w and $\mathcal{F} \in E$ we have Rw_2w and hence $w_2 \Vdash \Diamond\psi$. For the third implication suppose $w_2 \Vdash \Diamond\Diamond\psi$ for some $\Diamond\Diamond\psi \in \Sigma$. Therefore there exists $w \in W$ and $w' \in W$ such that Rw_2w, Rww' and $w' \Vdash \psi$. Since Rw_1w_2 and $\mathcal{F} \in E$, we know that Rw_2w_2; together with the fact that Rw_2w and $\mathcal{F} \in E$, this implies that Rww_2. Thus we have Rww_2, Rww' and $\mathcal{F} \in E$ which together imply that Rw_2w' and hence $w_2 \Vdash \Diamond\psi$. The following diagram depicts the situation (where the extra relations follow as $\mathcal{F}$ is Euclidean).

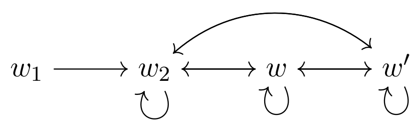

Note that the definition of R^Σ also ensures that $\mathcal{F}^\Sigma \in E$.

2.5 Algebraic semantics of modal logic

In this section we take for granted the notions of *subalgebra, homomorphism* and *Boolean algebra*. See [5] for the relevant material.

Up to now we have been concerned with the relational semantics of modal logic. However another interpretation of modal logic is through algebra. This is not foreign as the truth table semantics of propositional calculus are nothing more than algebraic semantics under the $\mathbf{2} = \{0, 1\}$ algebra.

We begin with the general definition of an algebra: algebras of type $\mathfrak{t}$ (recall Definition 2.1.1).

2.5.1 Definition. Let $\mathfrak{t} = (O, \zeta)$ be a type. An *algebra of type* $\mathfrak{t}$ is a structure $\mathbf{A} = (A, g_{\mathbf{A}})_{g \in \mathfrak{t}}$ where A is a non-empty set, called the *carrier set*, and $(g_{\mathbf{A}})_{g \in \mathfrak{t}}$ is a family of functions on A, the *fundamental operations of* $\mathbf{A}$, with $arity(g_{\mathbf{A}}) = \zeta(g)$ for all $g \in \mathfrak{t}$.

2.5.2 Example. Let *Bool* be the Boolean type and let τ be the basic modal type. The *term algebra* of type $Bool \cup \tau$[3] over Φ, denoted $\mathbf{T}_\tau(\Phi)$, is defined with carrier set $Form(\tau, \Phi)$ and operations $\vee_{\mathbf{T}_\tau(\Phi)}(\varphi, \psi) = \varphi \vee \psi$, $\wedge_{\mathbf{T}_\tau(\Phi)}(\varphi, \psi) = \varphi \wedge \psi$, $\neg_{\mathbf{T}_\tau(\Phi)}(\varphi) = \neg\varphi$, $\diamond_{\mathbf{T}_\tau(\Phi)}(\varphi) = \diamond\varphi$; the nullary operations $0_{\mathbf{T}_\tau(\Phi)}$ and $1_{\mathbf{T}_\tau(\Phi)}$ are identified with the formulae $\perp$ and $\top$ respectively. For a term $\varphi \in T_\tau(\Phi)$ the set $st(\varphi)$ denotes the set of subterms of φ. The carrier of the term algebra will be denoted by $T_\tau(\Phi)$ or $Form(\tau, \Phi)$.

2.5.3 Remark. For a modal type τ and algebra $\mathbf{A}$ of type $Bool \cup \tau$, we will often write g for the fundamental operation $g_{\mathbf{A}}$ for all $g \in Bool$. For all $\triangle \in \tau$ the fundamental operation $\triangle_{\mathbf{A}}$ will be denoted by $f_\triangle$ or just f when τ is the basic modal type.

For the sake of readability the results for the remainder of this section (as well as those of Sections 2.6 and 2.7) will be presented with τ as the basic modal type (in this circumstance the term algebra will be denoted by $\mathbf{T}(\Phi)$). Motivated by the characteristics of the diamond (Theorem 2.2.3) we look at the notion of an operator on an algebra.

2.5.4 Definition. Let τ be the basic modal type and let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f)$ be an algebra of type $Bool \cup \tau$. The function f is said to be a *normal operator on* $\mathbf{A}$ (or simply *operator on* $\mathbf{A}$) if it preserves joins and the bottom i.e. for all $x, y \in A$:

$$\begin{array}{lrcl} (i) & f(x \vee y) & = & f(x) \vee f(y); \\ (ii) & f(0) & = & 0. \end{array}$$

2.5.5 Definition. Let τ be the basic modal type. A *modal algebra* (or *Boolean algebra with operators*) is an algebra $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f)$ of type $Bool \cup \tau$ such that $(A, \vee, \wedge, \neg, 0, 1)$ is a Boolean algebra and f is an operator on $\mathbf{A}$.

2.5.6 Example. Let $\mathcal{F} = (W, R)$ be a frame and recall the mapping $m_R : \mathcal{P}(W) \to \mathcal{P}(W)$ defined by $m_R(S) = \{w \in W \mid (\exists w' \in S)(Rww')\}$. It can easily be checked m_R is an operator on the power set Boolean algebra $(\mathcal{P}(W), \bigcup, \bigcap, -, \varnothing, W)$ where $-$ refers to set complementation.

2.5.7 Definition. Let $\mathcal{F} = (W, R)$ be a frame. The *full complex algebra of* $\mathcal{F}$ is the modal algebra $\mathcal{F}^+ = (\mathcal{P}(W), \bigcup, \bigcap, -, \varnothing, W, m_R)$. A proper sub-algebra of a full complex algebra is called a *complex algebra*.

[3] $Bool \cup \tau$ is defined as $(\{\vee, \wedge, \neg, \diamond, 0, 1\}, \{(\vee, 2), (\wedge, 2), (\neg, 1), (\diamond, 1), (0, 0), (1, 0)\})$.

Modal algebras are good candidates for modal logic semantics since they have a Boolean algebra structure and an operator for the Boolean and modal formula semantics respectively. It is natural for algebraic semantics to be defined with respect to equations. For $\sigma, \psi \in T(\Phi)$ we write $\sigma \approx \psi$ for an equation since equality depends on the particular evaluation of variables in σ, ψ in a particular algebra. The following theorem will be used to define what it means for an equation to be a validity on a modal algebra.

2.5.8 Theorem. *Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f)$ be a modal algebra and $h : \Phi \to A$. The mapping h has a unique homomorphic extension h^* to $\mathbf{T}(\Phi)$. The following commutative diagram describes the situation (where $e_\Phi : \Phi \to T(\Phi)$ is defined by $e_\Phi(p) = p$).*

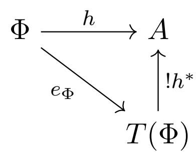

Proof. Define h^* as:

$$h^*(\varphi) = \begin{cases} 0 & \varphi = \bot \\ h(p) & \varphi = p \in \Phi \\ \neg h^*(\varphi_1) & \varphi = \neg\varphi_1 \\ h^*(\varphi_1) \vee h^*(\varphi_2) & \varphi = \varphi_1 \vee \varphi_2 \\ f(h^*(\varphi_1)) & \varphi = \Diamond\varphi_1. \end{cases}$$

From the definition it is immediate that h^* is a homomorphism that extends h. Towards uniqueness suppose $j : T(\Phi) \to A$ is an arbitrary homomorphism extending h i.e. $j \circ e_\Phi = h$. Since $h^* \circ e_\Phi = h$ we have that j and h^* agree on Φ. But homomorphisms are only defined up to generating subsets and by the definition of $\mathbf{T}(\Phi)$ (see Example 2.5.2) we have that Φ generates $\mathbf{T}(\Phi)$, thus $j = h^*$. □

2.5.9 Remark. In light of this theorem the statement that a logic is closed under uniform substitution (from Definition 2.2.1) can be stated rigorously as: for all $g : \Phi \to T(\Phi)$ and all $\varphi \in T(\Phi)$, $\vdash_\Lambda \varphi$ implies $\vdash_\Lambda g^*(\varphi)$.

2.5.10 Definition. Let $\sigma, \psi \in T(\Phi)$ and let $\mathbf{A}$ be a modal algebra. The equation $\sigma \approx \psi$ is *valid on* $\mathbf{A}$ if $g^*(\sigma) = g^*(\psi)$ for every $\mathbf{A}$-*assignment* $g : \Phi \to A$. This is denoted by $\mathbf{A} \models \sigma \approx \psi$. The formula σ is a *validity on* $\mathbf{A}$ if $\mathbf{A} \models \sigma^\approx$ where $\sigma^\approx = \sigma \approx \top$.

2.5.11 Definition. Let $\mathfrak{A}$ be a class of modal algebras and let $\sigma, \psi \in T(\Phi)$. The equation $\sigma \approx \psi$ is *valid on* $\mathfrak{A}$ if $\mathbf{A} \models \sigma \approx \psi$ for every $\mathbf{A} \in \mathfrak{A}$. This is denoted $\mathfrak{A} \models \sigma \approx \psi$. The formula σ is a *validity on* $\mathfrak{A}$ if $\sigma^\approx$ is valid on $\mathfrak{A}$.

2.5.12 Example. Recall Example 2.3.3: $D = \Box p \to \Diamond p$ and $ruFrm = \{(W, R) \mid (\forall w \in W)(\exists w' \in W)(Rww')\}$. Let $\mathfrak{A}$ be the class of full complex algebras $\mathcal{F}^+$ where $\mathcal{F} \in ruFrm$. We will show $\mathfrak{A} \models D^{\approx}$. Let $\mathcal{F}^+ = (W, R)^+ \in \mathfrak{A}$ and $g : \Phi \to \mathcal{P}(W)$. We need to show $g^*(D) = W$ as $g^*(\top) = W$. Now $g^*(D) = g^*(\Box p \to \Diamond p) = g^*(\neg\neg\Diamond\neg p \vee \Diamond p)$ by definition of $\to$ and $\Box$. Since g^* is a homomorphism we have $g^*(\neg\neg\Diamond\neg p \vee \Diamond p) = - - m_R(-g^*(p)) \cup m_R(g^*(p))$ where $-$ is set complementation. Then $- - m_R(-g^*(p)) \cup m_R(g^*(p)) = m_R(-g^*(p)) \cup m_R(g^*(p))$ since $\neg\neg x = x$ on a Boolean algebra. Then $m_R(-g^*(p)) \cup m_R(g^*(p)) = m_R(-g^*(p) \cup g^*(p)) = m_R(W)$ since m_R is an operator on $\mathcal{P}(W)$. But $m_R(W) = W$ since $\mathcal{F} \in ruFrm$. Thus $g^*(D) = W$ and therefore $\mathfrak{A} \models D^{\approx}$.

In general we have $\mathcal{F} \Vdash \sigma \leftrightarrow \psi$ iff $\mathcal{F}^+ \models \sigma \approx \psi$. This comes down to the fact that the notions of $\mathcal{P}(W)$-assignment and valuation are the same and thus have equal homomorphic extensions to $\mathbf{T}(\Phi)$ (see Remark 2.1.15). In this way complex algebras precisely capture frame validity.

2.5.13 Definition. Let $\mathfrak{A}$ be a class of modal algebras. The *equational theory of* $\mathfrak{A}$ is the set of formulae $\{\sigma \in T(\Phi) \mid \mathfrak{A} \models \sigma^{\approx}\}$ and is denoted by $\Lambda_{\mathfrak{A}}$.

2.5.14 Remark. It can be demonstrated that for any class of modal algebras $\mathfrak{A}$ the set $\Lambda_{\mathfrak{A}}$ is a logic.

2.6 Completeness with respect to modal algebras

In this section we take for granted the notion of *congruence*. See [5] for the relevant material.

2.6.1 Definition. Let Λ be a logic and let $\mathfrak{A}$ be a class of modal algebras. The logic Λ is said to be *sound with respect to* $\mathfrak{A}$ if $\vdash_\Lambda \sigma \leftrightarrow \psi$ implies $\mathfrak{A} \models \sigma \approx \psi$ for all $\sigma, \psi \in T(\Phi)$. Λ is said to be *complete with respect to* $\mathfrak{A}$ if $\mathfrak{A} \models \sigma \approx \psi$ implies $\vdash_\Lambda \sigma \leftrightarrow \psi$ for all $\sigma, \psi \in T(\Phi)$.

2.6.2 Remark. The proof that a logic is sound with respect to a class of modal algebras is straightforward. It involves checking that the axioms of the logic are validities on the class of algebras as well as checking that the validities of the class of algebras are closed under modus ponens, uniform substitution and generalisation.

We shall now look at completeness of a logic with respect to a class of algebras. As we shall see, the advantage of algebraic semantics is that, unlike with frames, every logic is complete with respect to some class of algebras (Theorem 2.6.9). In particular we will show completeness with respect to Lindenbaum–Tarski algebras: modal algebras formed by factoring equivalent formulae (modulo a given logic) out of the term algebra. In Section 2.7 we will see that Lindenbaum–Tarski algebras are the algebraic analogues of canonical models.

2.6.3 Definition. Let Λ be a logic. The binary relation $\equiv_\Lambda$ on $T(\Phi)$ is defined as $\{(\sigma, \psi) \in T(\Phi)^2 \mid \vdash_\Lambda \sigma \leftrightarrow \psi\}$.

2.6.4 Lemma. Let Λ be a logic. The relation $\equiv_\Lambda$ is a congruence on $\mathbf{T}(\Phi)$.

Proof. $\equiv_\Lambda$ **is an equivalence relation**: Uniform substitution and modus ponens with the following propositional tautologies serves to demonstrate that $\equiv_\Lambda$ satisfies the desired properties of an equivalence relation.

$\vdash_\Lambda$	$p \leftrightarrow p$	reflexivity
$\vdash_\Lambda$	$(p \leftrightarrow q) \to (q \leftrightarrow p)$	symmetry
$\vdash_\Lambda$	$((p \leftrightarrow q) \wedge (q \leftrightarrow r)) \to (p \leftrightarrow r)$	transitivity

$\equiv_\Lambda$ **is a congruence**: Uniform substitution and modus ponens with the following tautologies will show that $\equiv_\Lambda$ preserves the relevant operation on $\mathbf{T}(\Phi)$.

$\vdash_\Lambda (p \leftrightarrow q) \to ((p \vee r) \leftrightarrow (q \vee r))$	$\equiv_\Lambda$ preserves $\vee$
$\vdash_\Lambda (p \leftrightarrow q) \to ((p \wedge r) \leftrightarrow (q \wedge r))$	$\equiv_\Lambda$ preserves $\wedge$
$\vdash_\Lambda (p \leftrightarrow q) \leftrightarrow (\neg p \leftrightarrow \neg q)$	$\equiv_\Lambda$ preserves $\neg$

It can also be seen that $\varphi \equiv_\Lambda \varphi_1$ implies $\Diamond\varphi \equiv_\Lambda \Diamond\varphi_1$ by uniform substitution and modus ponens with the Λ formula $(p \to q) \to (\Diamond p \to \Diamond q)$ (Theorem 2.2.3). $\square$

2.6.5 Remark. Given a logic Λ, the equivalence class of φ under $\equiv_\Lambda$ (the set $\{\psi \in T(\Phi) \mid \psi \equiv_\Lambda \varphi\}$) is denoted by $[\varphi]_\Lambda$ or simply $[\varphi]$.

2.6.6 Definition. Let Λ be a logic. The *Lindenbaum–Tarski algebra over* Λ is defined as $\mathcal{L}_\Lambda(\Phi) = (T(\Phi)/\equiv_\Lambda, \vee, \wedge, \neg, 0, 1, f)$ with the operations $(\vee, \wedge, \neg, 0, 1, f)$ defined as the canonical quotient operations:

$$\begin{aligned} 0 &= [\bot] \\ 1 &= [\top] \\ \neg[\varphi] &= [\neg\varphi] \\ [\varphi] \vee [\psi] &= [\varphi \vee \psi] \\ [\varphi] \wedge [\psi] &= [\varphi \wedge \psi] \\ f[\varphi] &= [\Diamond\varphi]. \end{aligned}$$

2.6.7 Remark. The fact that the operations on $\mathcal{L}_\Lambda(\Phi)$ are well defined is due precisely to the fact that $\equiv_\Lambda$ respects the algebraic operations on $\mathbf{T}(\Phi)$ (Lemma 2.6.4).

2.6.8 Lemma. Let Λ be a logic. The algebra $\mathcal{L}_\Lambda(\Phi)$ is a modal algebra.

Proof. The fact that the Lindenbaum–Tarski operations $(\vee, \wedge, \neg, 0, 1)$ obey the Boolean algebra axioms can be shown using the appropriate propositional tautologies in Λ. The function f on $\mathcal{L}_\Lambda(\Phi)$ is an operator on $\mathcal{L}_\Lambda(\Phi)$ since $\vdash_\Lambda \Diamond(p \vee q) \leftrightarrow (\Diamond p \vee \Diamond q)$ and $\vdash_\Lambda \Diamond\bot \leftrightarrow \bot$ (Theorem 2.2.3). □

2.6.9 Theorem. *Let Λ be a logic. Then Λ is complete with respect to $\{\mathcal{L}_\Lambda(\Phi)\}$.*

Proof. Let $\sigma, \psi \in T(\Phi)$ and suppose $\nvdash_\Lambda \sigma \leftrightarrow \psi$. We want to show $\{\mathcal{L}_\Lambda(\Phi)\} \not\models \sigma \approx \psi$. In particular we will show there exists a $g : \Phi \to \mathcal{L}_\Lambda(\Phi)$ such that $g^*(\sigma) \neq g^*(\psi)$. Consider the $g : \Phi \to \mathcal{L}_\Lambda(\Phi)$ defined by $g(p) = [p]$. The fundamental operations on $\mathcal{L}_\Lambda(\Phi)$ facilitate a straightforward induction on the length of formulae to show that $g^*(\varphi) = [\varphi]$ for all $\varphi \in T(\Phi)$. Therefore $g^*(\sigma) = [\sigma]$ and $g^*(\psi) = [\psi]$. Since $\nvdash_\Lambda \sigma \leftrightarrow \psi$ it follows that $g^*(\sigma) \neq g^*(\psi)$. Therefore Λ is complete with respect to $\{\mathcal{L}_\Lambda(\Phi)\}$. □

2.6.10 Remark. To prove that a logic Λ is complete with respect to a class of modal algebras $\mathfrak{A}$, this theorem shows it suffices to demonstrate $\mathcal{L}_\Lambda(\Phi) \in \mathfrak{A}$.

2.7 The duality between frames and algebras

Up to now we have looked at aspects of modal validities through lenses that are apparently different: frames and modal algebras. It is natural to ask whether there is some relationship between the two approaches. In this section we will discuss how a representation theorem due to Jónsson and Tarski holds for modal algebras and complex algebras, analogous to the Stone representation of every Boolean algebra as a set algebra [21]. This is a useful result because in the same way that set algebras model propositional logic, complex algebras model modal logic (see Example 2.5.12 and paragraph following it).

In the same way that taking the full complex algebra of a frame (see Definition 2.5.7) takes us from frames to modal algebras, we will now discuss a method of moving from modal algebras to frames.

2.7.1 Definition. Let $\mathbf{B} = (B, \vee, \wedge, \neg, 0, 1)$ be a Boolean algebra and let $C \subseteq B$. The set C is a *filter of* $\mathbf{B}$ if for all $x, y \in B$:

(*i*) $x, y \in C$ implies $x \wedge y \in C$;
(*ii*) $x \in C$ and $x \leq y$ implies $y \in C$;
(*iii*) $1 \in C$.

A filter C of $\mathbf{B}$ is called a *proper filter* if $0 \notin C$. A maximal proper filter (relative to set inclusion) of $\mathbf{B}$ is called an *ultrafilter of* $\mathbf{B}$. The set of ultrafilters of $\mathbf{B}$ is denoted Uf$\mathbf{B}$.

We state the following theorem about ultrafilters without proof.

2.7.2 Theorem. *Let* $\mathbf{B} = (B, \vee, \wedge, \neg, 0, 1)$ *be a Boolean algebra and let* $C \subseteq B$ *be a proper filter of* $\mathbf{B}$*.* C *is an ultrafilter of* $\mathbf{B}$ *iff* $y \in C$ *or* $\neg y \in C$ *for every* $y \in B$*.*

2.7.3 Remark. An *ultrafilter over a set* W is an ultrafilter of the power set algebra $\mathcal{P}(W)$. Every ultrafilter over an infinite set W is uncountable. To see this consider an ultrafilter C over an infinite set W. For $S \in C$, it is guaranteed that $W \setminus S \in \mathcal{P}(W) \setminus C$ else otherwise we would have $\varnothing \in C$ which is a contradiction to C being a proper filter. Thus consider the mapping $t : C \to \mathcal{P}(W) \setminus C$ by $t(S) = W \setminus S$. It is easy to check that t is injective; the fact that t is surjective can be seen from Theorem 2.7.2. Thus t is a bijection. Since W is infinite, $\mathcal{P}(W)$ is uncountable by Cantor's theorem. The ultrafilter C cannot be countable as this would imply (by the bijectivity of t) that $\mathcal{P}(W) \setminus C$ is countable which would lead to $C \cup (\mathcal{P}(W) \setminus C) = \mathcal{P}(W)$ being countable.

We state the following theorem without proof.

2.7.4 Theorem. *Let* $\mathbf{B} = (B, \vee, \wedge, \neg, 0, 1)$ *be a Boolean algebra and let* $C \subseteq B$ *be an ultrafilter of* $\mathbf{B}$*. For all* $x, y \in B$*,* $x \vee y \in C$ *iff* $x \in C$ *or* $y \in C$*.*

2.7.5 Definition. Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f)$ be a modal algebra. The *ultrafilter frame of* $\mathbf{A}$ is the frame $\mathbf{A}_+ = (\text{Uf}\mathbf{A}, Q_f)$ where $Q_f = \{(C_1, C_2) \in \text{Uf}\mathbf{A}^2 \mid (\forall x \in C_2)(f(x) \in C_1)\}$.

Taking the ultrafilter frame of a modal algebra offers a construction that leads from modal algebras to frames. The following diagram depicts the situation between the two settings of frames and modal algebras.

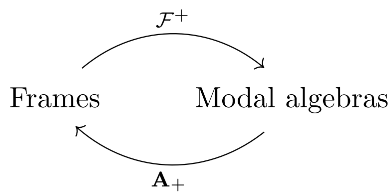

One would know that $(\ )_+$ and $(\ )^+$ are the 'right' constructions to move between the two settings if one could recover $\mathbf{A}$ from $\mathbf{A}_+{}^+$ and $\mathcal{F}$ from $\mathcal{F}^+{}_+$ where $\mathbf{A}$ is a modal algebra and $\mathcal{F}$ is a frame. The Jónsson–Tarski theorem shows this is indeed the case for algebras by stating that a modal algebra $\mathbf{A}$ is embedded in its *canonical algebra*: $\mathbf{A}_+{}^+$. Furthermore it shows that every modal algebra is isomorphic to a complex algebra, algebras that capture frame validity. This embedding of a modal algebra into its canonical algebra has links to topology, in particular, Stone spaces (see [8]). We will state the Jónsson–Tarski theorem [21] without proof (see Theorem 5.43 of [1]).

2.7.6 Theorem. *(Jónsson–Tarski Theorem) Let* $\mathbf{A}$ *be a modal algebra and consider* ${A_+}^+$*, the carrier set of the algebra* ${\mathbf{A}_+}^+$*. The map* $r : A \to {A_+}^+$ *defined by* $r(x) = \{C \in Uf\mathbf{A} \mid x \in C\}$ *is an injective homomorphism.*

It can also be shown that $\mathcal{F} = (W, R)$ is embedded in ${\mathcal{F}^+}_+ = (\mathrm{Uf}\,\mathcal{P}(W), Q_{m_R})$ by the map $w \mapsto l(w) = \{C \in \mathrm{Uf}\,\mathcal{P}(W) \mid \{w\} \in C\}$. Notice the strong similarity of the Λ-MCS properties (Theorem 2.3.11) and those of ultrafilters. The following theorem firmly establishes the connection between the two.

2.7.7 Theorem. *Let* Λ *be a logic. Then* $\mathcal{L}_\Lambda(\Phi)_+ \cong \mathcal{F}^\Lambda$*.*

3. PDL and the Modal μ-Calculus

3.1 Propositional dynamic logic (PDL)

The results in this section are taken from the sources [1] and [11]. For each result, the proof is provided when we feel it offers some insight. When a proof is provided, we fill out more details than those found in the source material.

PDL is a branch of modal logic where modalities represent the actions of programs. In this section we will first look at the PDL type which contains more than one (possibility infinitely many) unary modalities (Definition 3.1.1). Next we discuss regular frames which offer a relational interpretation for PDL formulae (Definition 3.1.3). We then give the Segerberg axiomatisation of PDL (Definition 3.1.8) and finally we present a brief outline of the result of Fischer and Ladner which states that **PDL** is decidable (Corollary 3.1.13).

Most of the results and definitions up to this point have been with respect to the basic modal type and will be assumed to be suitably adjusted when used in the context of a multi-modal unary type (see [1]).

3.1.1 Definition. The set of *program symbols* Π is defined recursively over a set of *atomic program symbols* Π_0 as $\pi = a \mid \pi_1;\pi_2 \mid \pi_1 \cup \pi_2 \mid \pi^*$ where $a \in \Pi_0$. The *PDL type* (denoted τ^*) is a modal type consisting of unary modal symbols indexed by program symbols in Π (modal symbols are written as $\langle\pi\rangle$ for $\pi \in \Pi$). The *basic program type* (denoted τ_0) is a modal type consisting of unary modal symbols indexed by atomic program symbols in Π_0.

3.1.2 Remark. The formulae of PDL are the elements in the set $Form(\tau^*, \Phi)$.

Intuitively $\pi_1;\pi_2$ represents the program which executes program π_1 then program π_2, the symbol $\pi_1 \cup \pi_2$ represents the program which executes program π_1 or program π_2. Finally, π^* represents the program that executes program π a finite (possibly zero) number of times. According to these intended meanings, PDL formulae are interpreted on regular frames.

3.1.3 Definition. A τ^*-frame $\mathcal{F} = (W, R_{\langle\pi\rangle})_{\pi\in\Pi}$ is *regular* if for all $\pi, \pi_1, \pi_2 \in \Pi$:

$$\begin{array}{lrcl} (i) & R_{\langle\pi_1;\pi_2\rangle} & = & R_{\langle\pi_1\rangle} \circ R_{\langle\pi_2\rangle}; \\ (ii) & R_{\langle\pi_1\cup\pi_2\rangle} & = & R_{\langle\pi_1\rangle} \cup R_{\langle\pi_2\rangle}; \\ (iii) & R_{\langle\pi^*\rangle} & = & R^*_{\langle\pi\rangle} \end{array}$$

where $R_{\langle\pi_1\rangle} \circ R_{\langle\pi_2\rangle}$ is relation composition of $R_{\langle\pi_1\rangle}$ and $R_{\langle\pi_2\rangle}$, $R^*_{\langle\pi\rangle}$ is the transitive reflexive closure of $R_{\langle\pi\rangle}$. The class of regular frames is denoted $RFrm$. A τ^*-model is *regular* if it is based on a regular τ^*-frame.

3.1.4 Example. Let $\Pi_0 = \{a\}$; then $\tau^* = \{\langle a\rangle, \langle a^*\rangle, \langle a; a^*\rangle, ...\}$. The diagram below shows a regular τ^*-model but only displaying the relations $R_{\langle a\rangle}$ (dashed) and $R_{\langle a^*\rangle}$.

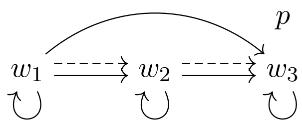

It can be seen $w_1 \Vdash \langle a^*\rangle p$, $w_2 \Vdash \langle a \cup a^*\rangle p$ and, although the relation is not shown, $w_1 \Vdash \langle a; a\rangle p$.

3.1.5 Remark. In Chapter 5 we will introduce the class of dynamic algebras $\mathfrak{D}$ which provide an algebraic semantics of PDL.

Intuitively, the set of worlds W of a regular frame can be viewed as the possible execution states of a computer; a relation $R_{\langle\pi\rangle}$ relates state w_1 to state w_2 if before program π executes the computer is in state w_1 and afterwards, it is in state w_2. In this way the formula $\langle\pi\rangle\varphi$ represents the statement that φ holds after a possible execution of π and $[\pi]\varphi = \neg\langle\pi\rangle\neg\varphi$ expresses that φ holds after every execution of π.

We now look at the following result for the valuation of PDL formulae. This theorem will be used to show that PDL is a fragment of the modal μ-calculus (Theorem 3.2.20).

3.1.6 Theorem. *Let $\langle\pi_1; \pi_2\rangle\varphi, \langle\pi_1 \cup \pi_2\rangle\varphi, \langle\pi^*\rangle\varphi \in T_{\tau^*}(\Phi)$ and let $\mathcal{M} = (W, R_{\langle\pi\rangle}, V)_{\pi\in\Pi}$ be a regular τ^*-model. Then*

$$\begin{aligned} V^*(\langle\pi_1;\pi_2\rangle\varphi) &= m_{R_{\langle\pi_1\rangle}}(m_{R_{\langle\pi_2\rangle}}(V^*(\varphi))) \\ V^*(\langle\pi_1\cup\pi_2\rangle\varphi) &= m_{R_{\langle\pi_1\rangle}}(V^*(\varphi)) \cup m_{R_{\langle\pi_2\rangle}}(V^*(\varphi)) \\ V^*(\langle\pi^*\rangle\varphi) &= \bigcup_{n=0}^{\infty} m^n_{R_{\langle\pi\rangle}}(V^*(\varphi)) \end{aligned}$$

where $m^0_{R_{\langle\pi\rangle}}(V^*(\varphi)) = V^*(\varphi)$ and $m^n_{R_{\langle\pi\rangle}}(V^*(\varphi)) = (m_{R_{\langle\pi\rangle}} \circ ... \circ m_{R_{\langle\pi\rangle}})(V^*(\varphi))$ nested n times.

Proof. Since $\mathcal{M}$ is a regular model we have

$$\begin{aligned} V^*(\langle\pi_1;\pi_2\rangle\varphi) &= m_{R_{\langle\pi_1;\pi_2\rangle}}(V^*(\varphi)) \\ &= m_{R_{\langle\pi_1\rangle}\circ R_{\langle\pi_2\rangle}}(V^*(\varphi)). \end{aligned}$$

Furthermore,

$$\begin{aligned} w \in m_{R_{\langle \pi_1 \rangle} \circ R_{\langle \pi_2 \rangle}}(V^*(\varphi)) &\iff \text{there exists a } w' \in V^*(\varphi) \text{ such that } R_{\langle \pi_1 \rangle} \circ R_{\langle \pi_2 \rangle} ww' \\ &\iff \text{there exists a } w' \in V^*(\varphi) \text{ and } z \in W \text{ such that } R_{\langle \pi_1 \rangle} wz \text{ and } \\ &\qquad R_{\langle \pi_2 \rangle} zw' \\ &\iff w \in m_{R_{\langle \pi_1 \rangle}}(m_{R_{\langle \pi_2 \rangle}}(V^*(\varphi))). \end{aligned}$$

Similarly for the second result, since $\mathcal{M}$ is a regular model, we have

$$\begin{aligned} V^*(\langle \pi_1 \cup \pi_2 \rangle \varphi) &= m_{R_{\langle \pi_1 \cup \pi_2 \rangle}}(V^*(\varphi)) \\ &= m_{R_{\langle \pi_1 \rangle} \cup R_{\langle \pi_2 \rangle}}(V^*(\varphi)). \end{aligned}$$

The fact that $m_{R_{\langle \pi_1 \rangle} \cup R_{\langle \pi_2 \rangle}}(V^*(\varphi)) = m_{R_{\langle \pi_1 \rangle}}(V^*(\varphi)) \cup m_{R_{\langle \pi_2 \rangle}}(V^*(\varphi))$ follows directly from the definition of $R_{\langle \pi_1 \rangle} \cup R_{\langle \pi_2 \rangle}$. For the final result we have, again by the regularity assumption on $\mathcal{M}$, that

$$\begin{aligned} V^*(\langle \pi^* \rangle \varphi) &= m_{R_{\langle \pi^* \rangle}}(V^*(\varphi)) \\ &= m_{R^*_{\langle \pi \rangle}}(V^*(\varphi)). \end{aligned}$$

Furthermore, since $R^*_{\langle \pi \rangle}$ is the transitive closure of $R_{\langle \pi \rangle}$, we have

$$\begin{aligned} w \in m_{R^*_{\langle \pi \rangle}}(V^*(\varphi)) &\iff \text{there exists a } w' \in V^*(\varphi) \text{ such that } R^*_{\langle \pi \rangle} ww' \\ &\iff \text{there exists a } w' \in V^*(\varphi) \text{ and a subset } \{w_0, ..., w_n\} \text{ of } W \text{ such that } \\ &\qquad w_0 = w, w_n = w' \text{ and } R_{\langle \pi \rangle} w_i w_{i+1} \text{ for all } 0 \leq i < n \\ &\iff w \in \bigcup_{n=0}^{\infty} m^n_{R_{\langle \pi \rangle}}(V^*(\varphi)). \end{aligned}$$

□

3.1.7 Remark. This lemma shows us that regular frames offer the intended semantics for PDL. For instance we now know $w \Vdash \langle \pi^* \rangle p$ iff $w \in m^n_{R_{\langle \pi \rangle}}(V(p))$ for some n i.e. p is satisfied on some finite $R_{\langle \pi \rangle}$-path from w (see Definition 2.1.7). Also the fact that the valuation of the formula $\langle \pi^* \rangle p$ involves an infinite disjunction demonstrates that PDL is strictly more expressive than a modal language equipped only with atomic unary modalities (such as Hennessy Milner logic discussed in Chapter 1).

The following is the Segerberg axiomatisation of PDL [37].

3.1.8 Definition. The logic **PDL** is the smallest logic containing the $T_{\tau^*}(\Phi)$ formulae:

(i)	$[\pi](p \to q) \to ([\pi]p \to [\pi]q)$
(ii)	$\langle\pi\rangle\, p \leftrightarrow \neg[\pi]\neg p$
(iii)	$\langle\pi_1;\pi_2\rangle\, p \leftrightarrow \langle\pi_1\rangle\,\langle\pi_2\rangle\, p$
(iv)	$\langle\pi_1 \cup \pi_2\rangle\, p \leftrightarrow (\langle\pi_1\rangle\, p \vee \langle\pi_2\rangle\, p)$
(v)	$\langle\pi^*\rangle\, p \leftrightarrow (p \vee \langle\pi\rangle\,\langle\pi^*\rangle\, p)$
(vi)	$[\pi^*](p \to [\pi]p) \to (p \to [\pi^*]p)$.

PDL can be shown to be decidable via a relational argument. In particular the result follows if **PDL** has the finite model property with respect to $RFrm$. The fact that **PDL** has the finite model property can be proven (in part) by using filtrations. However this process is complicated by the inductive structure of the modalities and one cannot always expect the filtration of a regular model through any subformula closed set to be regular. The Fischer–Ladner closure of a set of formulae solves this problem [11].

3.1.9 Definition. A set of formulae $\Sigma \subseteq T_{\tau^*}(\Phi)$ is *Fischer–Ladner closed* if

(i)	$\langle\pi_1;\pi_2\rangle\,\varphi \in \Sigma$ implies $\langle\pi_1\rangle\,\langle\pi_2\rangle\,\varphi \in \Sigma$;
(ii)	$\langle\pi_1 \vee \pi_2\rangle\,\varphi \in \Sigma$ implies $\langle\pi_1\rangle\,\varphi \vee \langle\pi_2\rangle\,\varphi \in \Sigma$;
(iii)	$\langle\pi^*\rangle\,\varphi \in \Sigma$ implies $\langle\pi\rangle\,\langle\pi^*\rangle\,\varphi \in \Sigma$;
(iv)	Σ is subformula closed.

3.1.10 Lemma. Let $\Sigma \subseteq T_{\tau^*}(\Phi)$. There exists a smallest *Fischer–Ladner* closed set containing Σ, denoted $FL(\Sigma)$. Furthermore if Σ is finite then $FL(\Sigma)$ is finite.

Proof. See Theorem 3.2 of [11]. □

3.1.11 Theorem. *$RFrm \Vdash \varphi$ iff $RFrm_{fin} \Vdash \varphi$ for all $\varphi \in T_{\tau^*}(\Phi)$.*

Proof. The proof involves taking a filtration of a model through a finite Fischer–Ladner closed set (see Theorem 3.2 of [11]). □

3.1.12 Theorem. ***PDL** is sound and complete with respect to $RFrm$.*

Proof. Soundness is straightforward. By Theorem 2.3.18 the completeness result follows by showing $\mathcal{F}^{\mathbf{PDL}} \in RFrm$. □

3.1.13 Corollary. **PDL** is decidable.

Proof. By Theorem 3.1.11 and Theorem 3.1.12 **PDL** has the finite model property with respect to $RFrm$. Therefore, together with the fact that **PDL** is finitely axiomatisable, we have that **PDL** is decidable by Harrop [16]. □

In Chapter 5 (Section 5.1) we give an algebraic proof that **PDL** is decidable.

3.2 The modal μ-calculus

The results in this section are taken from the sources [8], [4] and [38]. For each result, the proof is provided when we feel it offers some insight. When a proof is provided, we fill out more details than those found in the source material.

We shall now look at the modal μ-calculus. The modal μ-calculus (the language of which we denote by L_μ) adds a least fixed point binder, μ, to the modal language of basic programs. We will first consider the syntax and semantics of L_μ and then we will look at how L_μ can be used to express assertions regarding program correctness (Examples 3.2.13 and 3.2.15). We will then look at L_μ^*: a fragment of L_μ which expresses all PDL formulae (Theorem 3.2.19). We then finally review the result of Streett [38] which demonstrates both that L_μ is more expressive than PDL and why filtration techniques do not suffice for L_μ (Theorem 3.2.21).

3.2.1 Definition. The set Var denotes a set of *variables* disjoint from Φ. The *formulae of* L_μ are defined recursively as $\varphi = \bot \mid p \mid X \mid \neg\varphi_1 \mid \varphi_1 \vee \varphi_2 \mid \langle a\rangle\, \varphi_1 \mid \mu X.\varphi(X)$ where $p \in \Phi$, $X \in Var$, $a \in \Pi_0$ and $\varphi(X)$ refers to an L_μ formula possibly containing occurrences of the variable X; furthermore $\varphi(X)$ is *syntactically order-preserving in* X i.e. every X in $\varphi(X)$ occurs under an even number of negations. The *dual of* $\mu X.\varphi(X)$ is defined as $\nu X.\varphi(X) = \neg\mu X.\neg\varphi(\neg X)$. The collection of L_μ formulae will be denoted by $L_\mu(\Phi, Var)$.

3.2.2 Example. The string of symbols $\mu X_1.(\langle a\rangle\, X_1 \vee \mu X_2.(X_2 \wedge \langle b\rangle\, X_1))$ is an L_μ formula whilst the string of symbols $\mu X.(p \vee \langle a\rangle\, \neg X)$ is not an L_μ formula as $p \vee \langle a\rangle\, \neg X$ is not syntactically order-preserving in X.

Towards interpreting L_μ formulae we now discuss the least fixed point theorem of Knaster and Tarski [8].

3.2.3 Definition. Let $\mathbf{L} = (L, \vee, \wedge)$ be a lattice. A mapping $t : L \to L$ is *order-preserving* if $x \leq y$ implies $t(x) \leq t(y)$ for all $x, y \in L$. A *fixed point of* t is an element $x \in L$ such that $t(x) = x$; a *pre-fixed point of* t (*post-fixed point of* t) is an element $x \in L$ such that $t(x) \leq x$ ($x \leq t(x)$).

3.2.4 Theorem. *(Knaster Tarski Fixed Point Theorem) Let* $\mathbf{L} = (L, \vee, \wedge)$ *be a complete lattice and let* $t : L \to L$ *be an order-preserving function. Then the least fixed point of* t *exists and is* $\bigwedge\{x \in L \mid t(x) \leq x\}$*.*

Proof. The set $C = \{x \in L \mid t(x) \leq x\}$ is non-empty since $t(\bigvee L) \leq \bigvee L$ (as $\bigvee L$ is of course the top of $\mathbf{L}$). It needs to be shown that $z = \bigwedge C$ is a fixed point of t i.e. $t(z) = z$. We will firstly show $t(z) \leq z$. Since $z = \bigwedge C$ it suffices to show $t(z)$ is a lower bound of C: let $x \in C$; then $z \leq x$ by definition of z; but t is order-preserving thus $t(z) \leq t(x)$; but $t(x) \leq x$ since $x \in C$; thus $t(z) \leq x$ which shows $t(z)$ is a lower bound of C and thus $t(z) \leq z$. The other inequality $z \leq t(z)$ follows from the first: we have $t(z) \leq z$ and thus $t(t(z)) \leq t(z)$ i.e. $t(z) \in C$; therefore $z \leq t(z)$ by definition of z. Thus $t(z) = z$ i.e. z is a fixed point of t. We need to show z is the least fixed point of t. Suppose z_1 is a fixed point of t i.e. $t(z_1) = z_1$; thus $z_1 \in C$ which gives $z \leq z_1$ by definition of z. □

3.2.5 Remark. Using this result we can deduce, in a complete Boolean algebra $\mathbf{L}$, that the greatest fixed point of t exists and is $\bigvee\{x \in L \mid x \leq t(x)\}$. To see this let $y = \bigvee\{x \in L \mid x \leq t(x)\}$ and let $t' : L \to L$ be defined as $t'(x) = \neg t(\neg x)$ (note that it is order-preserving).

$$\begin{aligned}
\neg y &= \bigwedge\{\neg x \in L \mid x \leq t(x)\} \\
&= \bigwedge\{x \in L \mid \neg x \leq t(\neg x)\} \\
&= \bigwedge\{x \in L \mid \neg t(\neg x) \leq x\} \\
&= \bigwedge\{x \in L \mid t'(x) \leq x\}.
\end{aligned}$$

Thus, from Theorem 3.2.4, we know that $\neg y$ is the least fixed point of t'. From this we know that y is a fixed point of t: $t'(\neg y) = \neg y$ iff $\neg t(\neg\neg y) = \neg y$ iff $\neg t(y) = \neg y$ iff $t(y) = y$. Furthermore y is the greatest fixed point of t: let $z_1 \in L$ such that $t(z_1) = z_1$; we have $\neg t(\neg\neg z_1) = \neg z_1$ i.e. $t'(\neg z_1) = \neg z_1$ and therefore $\neg y \leq \neg z_1$ (since $\neg y$ is the least fixed point of t'). Thus $y \geq z_1$.

3.2.6 Remark. From these results we also have that $\bigwedge\{x \in L \mid t(x) \leq x\}$ is the least pre-fixed point of t and $\bigvee\{x \in L \mid x \leq t(x)\}$ is the greatest post-fixed point of t.

When it comes to the semantics of L_μ formulae it is immediate that formulae involving the fixed point binder μ are difficult to formally interpret. The Knaster–Tarski theorem (Theorem 3.2.4) will be used to interpret such formulae. Specifically for a τ_0-frame $(W, R_{\langle a \rangle})_{a \in \Pi_0}$ and L_μ formula $\mu X.\varphi(X)$, the interpretation of $\varphi(X)$ varies as the interpretation of X varies. In this way $\varphi(X)$ induces an order-preserving function on $\mathcal{P}(W)$ (order-preserving by the syntactic condition on $\varphi(X)$). The interpretation of $\mu X.\varphi(X)$ will be the least fixed point of this function on the complete lattice $\mathcal{P}(W)$ (guaranteed to exist by Theorem 3.2.4).

3.2.7 Definition. Let $\mathbf{A}$ be a τ_0 modal algebra, let $X \in Var$, let $\overline{X} = (X_1, ..., X_n) \in Var^n$, let $z \in A$, let $\overline{z} = (z_1, ..., z_n) \in A^n$ and let $V : \Phi \cup Var \to A$. Then an *X-variant of V* is defined as $V_{X \to z} : \Phi \cup Var \to A$ where:

$$V_{X \to z}(p) = \begin{cases} z & \text{if } p = X \\ V(p) & \text{otherwise.} \end{cases}$$

A *$\overline{X}$-variant of V* is defined as $V_{\overline{X} \to \overline{z}} : \Phi \cup Var \to A$ where $V_{\overline{X} \to \overline{z}} = V_{X_1 \to z_1, ..., X_n \to z_n}$.

$V_{\overline{X} \to \overline{z}}$ will also be referred to as a *vector variant of V*. The maps V and $V_{X\to}$ will form the base case in the following recursive definition for the semantics of L_μ formulae.

3.2.8 Definition. Let $\psi \in L_\mu(\Phi, Var)$, let $(W, R_{\langle a \rangle})_{a \in \Pi_0}$ be a τ_0-frame and let $V : \Phi \cup Var \to \mathcal{P}(W)$. The *set of states of W satisfying ψ* is given by $V^*(\psi)$ where $V^* : L_\mu(\Phi, Var) \to \mathcal{P}(W)$ is *the extension of V to $L_\mu(\Phi, Var)$* and is defined as:

$$V^*(\psi) = \begin{cases} \varnothing & \psi = \bot \\ V(p) & \psi = p \in \Phi \cup Var \\ W \setminus V^*(\psi_1) & \psi = \neg \psi_1 \\ V^*(\psi_1) \cup V^*(\psi_2) & \psi = \psi_1 \vee \psi_2 \\ m_{R_{\langle a \rangle}}(V^*(\psi_1)) & \psi = \langle a \rangle \psi_1 \\ \bigcap \{S \in \mathcal{P}(W) \mid V^*_{X \to S}(\varphi(X)) \subseteq S\} & \psi = \mu X.\varphi(X) \end{cases}$$

where $V^*_{X \to S}$ is the extension of $V_{X \to S}$ to $L_\mu(\Phi, Var)$.

3.2.9 Remark. $V^*_{X\to}(\varphi(X)) : \mathcal{P}(W) \to \mathcal{P}(W)$ is defined by $V^*_{X\to}(\varphi(X))(S) = V^*_{X\to S}(\varphi(X))$. The function $V^*_{X\to}(\varphi(X))$ will be denoted by φ^* when there is no danger of ambiguity. Since the function $V^*_{X\to}(\varphi(X))$ is order-preserving (by the syntactic condition on $\varphi(X)$) and $(\mathcal{P}(W), \bigcup, \bigcap)$ is a complete lattice we have by Theorem 3.2.4 that the interpretation of $\mu X.\varphi(X)$ is the least fixed point of the function $V^*_{X\to}(\varphi(X))$. The L_μ formula $\nu X.\varphi(X) = \neg \mu X.\neg\varphi(\neg X)$ is interpreted as

$$\begin{aligned} V^*(\neg\mu X.\neg\varphi(\neg X)) &= W \setminus \bigcap \{S \in \mathcal{P}(W) \mid V^*_{X\to S}(\neg\varphi(\neg X)) \subseteq S\} \\ &= \bigcup \{W \setminus S \in \mathcal{P}(W) \mid W \setminus V^*_{X\to S}(\varphi(\neg X)) \subseteq S\} \\ &= \bigcup \{W \setminus S \in \mathcal{P}(W) \mid W \setminus V^*_{X\to W\setminus S}(\varphi(X)) \subseteq S\} \\ &= \bigcup \{W \setminus S \in \mathcal{P}(W) \mid W \setminus S \subseteq V^*_{X\to W\setminus S}(\varphi(X))\} \\ &= \bigcup \{S \in \mathcal{P}(W) \mid S \subseteq V^*_{X\to S}(\varphi(X))\} \end{aligned}$$

i.e. $\nu X.\varphi(X)$ is interpreted as the greatest fixed point of $V^*_{X\to}(\varphi(X))$ (see Remark 3.2.5).

3.2.10 Example. Let $\mathcal{M} = (W, R_{\langle a\rangle})_{a\in\Pi_0}$ be a τ_0-frame, let $V : \Phi \cup Var \to \mathcal{P}(W)$ and let $\varphi(X) = X \vee \neg p$. Recall from the previous remark that we usually denote the map $V^*_{X\to}(\varphi(X))$ by φ^* when there is no danger of ambiguity. Therefore we have $\varphi^*(S) = S \cup (W \setminus V(p))$ for $S \in \mathcal{P}(W)$. We have $V^*(\mu X.\varphi(X)) = \bigcap\{S \in \mathcal{P}(W) \mid S \cup (W \setminus V(p)) \subseteq S\} = \bigcap\{S \in \mathcal{P}(W) \mid W \setminus V(p) \subseteq S\} = W \setminus V(p)$. It can be seen that the greatest fixed point of φ^* is W.

We have a way of relationally interpreting L_μ formulae with binders. However it is not always as straightforward as Example 3.2.10 to find the least fixed point of a formula on a model (for instance consider the formula $\varphi(X) = p \wedge [a]X$). Towards having a way to easily approximate the fixed point, there is a technique called *ordinal unfolding*: on a complete lattice $\mathbf{L}$ with an order-preserving function $t : L \to L$ we have $0 \leq t(0)$ which gives $t(0) \leq t^2(0)$ since t is order-preserving, which gives $t^2(0) \leq t^3(0)$ and so on. In this way $\mathbf{L}$ has the increasing chain $0 \leq t(0) \leq t^2(0) \leq \ldots \leq t^\lambda(0) \leq t^{\lambda+1}(0)\ldots$ where λ is a limit ordinal and $t^\lambda(0)$ is recursively defined as $\bigvee_{\gamma<\lambda} t^\gamma(0)$ and $t^{\lambda+1}(0)$ is defined as $t(t^\lambda(0))$. We state the following result regarding the join of this chain and the least fixed point of t (guaranteed to exist by Knaster–Tarski).

3.2.11 Lemma. Let $\mathbf{L}$ be a complete lattice and $t : L \to L$ be an order-preserving function. The element $\bigvee_{\gamma<\kappa} t^\gamma(0)$ (where $\kappa = |L|$) is equal to the least fixed point of t.

3.2.12 Remark. Similarly the meet of the descending chain $1 \geq t(1) \geq t^2(1) \geq \ldots \geq t^\lambda(1) \geq \ldots$ is the greatest fixed point of t.

We shall now look at two examples of how L_μ can be used to express assertions of safety and liveness in program correctness.

3.2.13 Example. Let $\Pi_0 = \{a\}$. Suppose we are interested in the safety condition: "$\langle a\rangle\, p$ is always satisfied along every $R_{\langle a\rangle}$-path" (p could represent the statement "memory is not full"). The safety assertion can be encoded by the L_μ formula $\nu X.(\langle a\rangle\, p \wedge [a]X)$. Let $\varphi(X) = \langle a\rangle\, p \wedge [a]X$ and consider the τ_0-model $\mathcal{F} = (W, R_{\langle a\rangle})$ with $V : \Phi \cup Var \to \mathcal{P}(W)$ given as:

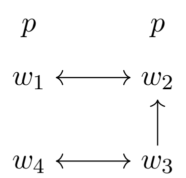

We want to see whether every state in W satisfies the safety assertion, in particular is $V^*(\nu X.\varphi(X)) = W$? We find $V^*(\nu X.\varphi(X))$ by using ordinal unfolding on φ^* in $\mathcal{P}(W)$: $\varphi^*(S) = m_{R_{\langle a\rangle}}(V(p)) \cap l_{R_{\langle a\rangle}}(S)$[1] thus giving rise to the following decreasing chain in $\mathcal{P}(W)$:

[1]Recall from Remark 2.1.15 that $l_{R_{\langle a\rangle}}(S) = \{w \in W \mid (\forall w' \in W)(R_{\langle a\rangle}ww' \implies w' \in S)\}$.

$$\begin{aligned}
\varphi^*(W) &= \{w_1, w_2, w_3\} \cap W \\
&= \{w_1, w_2, w_3\} \\
\varphi^{*2}(W) &= \{w_1, w_2, w_3\} \cap l_{R_{\langle a \rangle}}(\{w_1, w_2, w_3\}) \\
&= \{w_1, w_2, w_3\} \cap \{w_1, w_2, w_4\} \\
&= \{w_1, w_2\} \\
\varphi^{*3}(W) &= \{w_1, w_2, w_3\} \cap l_{R_{\langle a \rangle}}(\{w_1, w_2\}) \\
&= \{w_1, w_2, w_3\} \cap \{w_1, w_2\} \\
&= \{w_1, w_2\}.
\end{aligned}$$

This shows that $\bigcap_{n<|\mathcal{P}(W)|} \varphi^{*n}(W) = \{w_1, w_2\}$ and thus, by Remark 3.2.12, $V^*(\nu X.\varphi(X)) = \{w_1, w_2\}$. The points w_3 and w_4 fail the safety condition.

3.2.14 Remark. This safety property can be expressed by the PDL formula $[a^*]\langle a \rangle p$.

3.2.15 Example. Let $\Pi_0 = \{a, b\}$. Consider the liveness property: "there is a path to a state where b can be executed". This can be encoded by the formula $\mu X.(\langle b \rangle \top \vee \langle a \rangle X \vee \langle b \rangle X)$. Let $\varphi(X) = \langle b \rangle \top \vee \langle a \rangle X \vee \langle b \rangle X$ and consider the τ_0-frame $(W, R_{\langle a \rangle}, R_{\langle b \rangle})$ and $V : \Phi \cup Var \to \mathcal{P}(W)$ given by (where $R_{\langle b \rangle}$ elements are dashed):

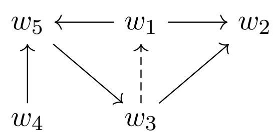

Towards finding $V^*(\mu X.\varphi(X))$:

$$\begin{aligned}
\varphi^*(\varnothing) &= m_{R_{\langle b \rangle}}(W) \cup m_{R_{\langle a \rangle}}(\varnothing) \cup m_{R_{\langle b \rangle}}(\varnothing) \\
&= \{w_3\} \\
\varphi^{*2}(\varnothing) &= \{w_3\} \cup m_{R_{\langle a \rangle}}(\{w_3\}) \cup m_{R_{\langle b \rangle}}(\{w_3\}) \\
&= \{w_3, w_5\} \\
\varphi^{*3}(\varnothing) &= \{w_3\} \cup m_{R_{\langle a \rangle}}(\{w_3, w_5\}) \cup m_{R_{\langle b \rangle}}(\{w_3, w_5\}) \\
&= \{w_1, w_3, w_4, w_5\} \\
\varphi^{*4}(\varnothing) &= \{w_3\} \cup m_{R_{\langle a \rangle}}(\{w_1, w_3, w_4, w_5\}) \cup m_{R_{\langle b \rangle}}(\{w_1, w_3, w_4, w_5\}) \\
&= \{w_3\} \cup \{w_1, w_3, w_4, w_5\} \cup \{w_3\} \\
&= \{w_1, w_3, w_4, w_5\}.
\end{aligned}$$

Therefore $\bigcup_{n<|\mathcal{P}(W)|} \varphi^{*n}(\varnothing) = \{w_1, w_3, w_4, w_5\}$ and thus by Lemma 3.2.11 $V^*(\mu X.\varphi(X)) = \{w_1, w_3, w_4, w_5\}$. The state w_2 is the only state to fail the liveness condition.

3.2.16 Remark. This liveness property of "there is a path to a state where b can be executed" can also be expressed by the PDL formula $\langle (a \vee b)^* \rangle \langle b \rangle \top$.

We have seen how L_μ can express formulae of PDL; we will now exactly isolate the fragment of L_μ corresponding to PDL, denoted L_μ^*.

3.2.17 Definition. The formulae of L_μ^* are defined recursively as $\varphi = \bot \mid p \mid X \mid \neg\varphi_1 \mid \varphi_1 \vee \varphi_2 \mid \langle a \rangle \varphi_1 \mid \mu X.(\varphi \vee \langle a \rangle X)$ where $p \in \Phi$, $X \in Var$ and $a \in \Pi_0$. The collection of formulae of L_μ^* is denoted $L_\mu^*(\Phi, Var)$.

Towards showing that PDL formulae are equivalent to L_μ^* formulae we define the notion of two such formulae being $\mathcal{M}$-equivalent for a model $\mathcal{M}$.

3.2.18 Definition. Let $\varphi \in T_{\tau^*}(\Phi)$, let $\psi \in L_\mu^*(\Phi, Var)$ and let $\mathcal{M} = (\mathcal{F}, V)$ be a τ^*-model. The formulae φ, ψ are *$\mathcal{M}$-equivalent* if $V^*(\varphi) = V^*(\psi)$. We denote this situation by $\varphi \sim_{\mathcal{M}} \psi$.

We now see the full translation from PDL formulae to L_μ^* formulae.

3.2.19 Theorem. *Let $\mathcal{M}$ be a regular τ^*-model and consider $l : T_{\tau^*}(\Phi) \cup Var \to L_\mu(\Phi, Var)$ by*

$$l(\varphi) = \begin{cases} X & \varphi = X \\ p & \varphi = p \in \Phi \\ \neg l(\varphi_1) & \varphi = \neg\varphi_1 \\ l(\varphi_1) \vee l(\varphi_2) & \varphi = \varphi_1 \vee \varphi_2 \\ \langle a \rangle l(\varphi_1) & \varphi = \langle a \rangle \varphi_1 \\ l(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi_1) & \varphi = \langle \pi_1; \pi_2 \rangle \varphi_1 \\ l(\langle \pi_1 \rangle \varphi_1) \vee l(\langle \pi_2 \rangle \varphi_1) & \varphi = \langle \pi_1 \cup \pi_2 \rangle \varphi_1 \\ \mu X.(l(\varphi_1) \vee l(\langle \pi \rangle X)) & \varphi = \langle \pi^* \rangle \varphi_1. \end{cases}$$

Then $\varphi \sim_{\mathcal{M}} l(\varphi)$ for all $\varphi \in T_{\tau^*}(\Phi)$.

Proving this full correspondence of PDL with L_μ^* involves a tedious double induction on the length of formulae and programs. We thus prove a simpler version of the theorem below.

3.2.20 Theorem. *Let $\mathcal{M} = (W, R_{\langle \pi \rangle}, V)_{\pi \in \Pi}$ be a regular τ^*-model, let $a, b \in \Pi_0$ and let $p \in \Phi$. Then it follows*

$$\begin{aligned} \langle a; b \rangle p &\sim_{\mathcal{M}} l(\langle a; b \rangle p) = \langle a \rangle \langle b \rangle p \\ \langle a \cup b \rangle p &\sim_{\mathcal{M}} l(\langle a \cup b \rangle p) = \langle a \rangle p \vee \langle b \rangle p \\ \langle a^* \rangle p &\sim_{\mathcal{M}} l(\langle a^* \rangle p) = \mu X.(p \vee \langle a \rangle X). \end{aligned}$$

Proof. From Theorem 3.1.6 we have

$$\begin{aligned} V^*(\langle a;b\rangle\, p) &= m_{R_{\langle a\rangle}}(m_{R_{\langle b\rangle}}(V(p))) \\ &= V^*(\langle a\rangle\, \langle b\rangle\, p) \end{aligned}$$

$$\begin{aligned} V^*(\langle a\cup b\rangle\, p) &= m_{R_{\langle a\rangle}}(V(p)) \cup m_{R_{\langle b\rangle}}(V(p)) \\ &= V^*(\langle a\rangle\, p \vee \langle b\rangle\, p). \end{aligned}$$

Let $\varphi(X) = p \vee \langle a\rangle\, X$. We now need to show that $\langle a^*\rangle\, p$ and $\mu X.\varphi(X)$ are $\mathcal{M}$-equivalent. By Theorem 3.1.6 we know that $V^*(\langle a^*\rangle\, p)=\bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))$. Thus it remains to prove $V^*(\mu X.(p \vee \langle a\rangle\, X))=\bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))$. We proceed by ordinal unfolding (and use the fact that $m_{R_{\langle a\rangle}}$ is an operator on $\mathcal{P}(W)$):

$$\begin{aligned} \varphi^*(\varnothing) &= V(p) \cup m_{R_{\langle a\rangle}}(\varnothing) \\ &= V(p) \cup \varnothing \\ &= V(p) \\ \varphi^{*2}(\varnothing) &= V(p) \cup m_{R_{\langle a\rangle}}(V(p)) \\ \varphi^{*3}(\varnothing) &= V(p) \cup m_{R_{\langle a\rangle}}(V(p) \cup m_{R_{\langle a\rangle}}(V(p))) \\ &= V(p) \cup m_{R_{\langle a\rangle}}(V(p)) \cup m_{R_{\langle a\rangle}}(m_{R_{\langle a\rangle}}(V(p))) \\ &\vdots \\ \varphi^{*i}(\varnothing) &= \bigcup_{n=0}^{i-1} m^n_{R_{\langle a\rangle}}(V(p)). \end{aligned}$$

Inductively it follows that $\varphi^{*\omega}(\varnothing) = \bigcup\limits_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))$ which gives

$$\begin{aligned} \varphi^{*\omega+1}(\varnothing) &= V(p) \cup m_{R_{\langle a\rangle}}(\bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))) \\ &= \bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p)). \end{aligned}$$

Therefore $\bigcup_{\gamma<|\mathcal{P}(W)|} \varphi^{*\gamma}(\varnothing) = \bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))$ and thus by Lemma 3.2.11 we have $V^*(\mu X.(p \vee \langle a\rangle\, X)) = \bigcup_{n=0}^{\infty} m^n_{R_{\langle a\rangle}}(V(p))$. □

We will now look at a result of Streett [38] which demonstrates how L_μ is strictly larger than PDL; larger in the sense that there is no PDL formula which is $\mathcal{M}$-equivalent to $\mu X.[a]X$ for all τ^*-models $\mathcal{M}$. In the proof we will also see why filtration techniques do not work for L_μ.

3.2.21 Theorem. *Let $a \in \Pi_0$. The L_μ formula $\mu X.[a]X$ is not $\mathcal{M}$-equivalent to a PDL formula for all τ^*-models $\mathcal{M}$.*

Proof. Suppose towards a contradiction that there is a ψ in $T_{\tau^*}(\Phi)$ such that $\mu X.[a]X$ is $\mathcal{M}$-equivalent to ψ for all τ^*-models $\mathcal{M}$. Consider the following τ^*-model $\mathcal{N} = (W, R_{\langle\pi\rangle}, V)_{\pi\in\Pi}$ (where only the $R_{\langle a\rangle}$ elements are shown).

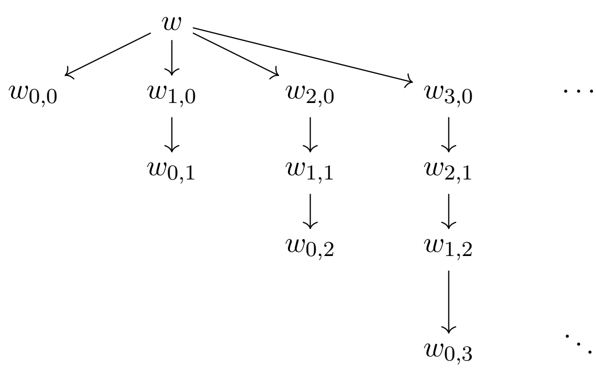

We shall show $\mathcal{N}$ satisfies $\mu X.[a]X$ everywhere: denote $[a]X$ by $\varphi(X)$ (thus we have $\varphi^*(S) = l_{R_{\langle a\rangle}}(S)$) and denote $\bigcup_{i=0}^{\infty}\{w_{n,i}\}$ by S_n[2]. We have:

$$\begin{array}{rcl} l_{R_{\langle a\rangle}}(\varnothing) & = & S_0 \\ l^2_{R_{\langle a\rangle}}(\varnothing) & = & S_0 \cup S_1 \\ & \vdots & \\ l^n_{R_{\langle a\rangle}}(\varnothing) & = & \bigcup_{i=0}^{n-1} S_i \ . \end{array}$$

Inductively $l^{\omega}_{R_{\langle a\rangle}}(\varnothing) = \bigcup_{i=0}^{\infty} S_i$ which gives that $l^{\omega+1}_{R_{\langle a\rangle}}(\varnothing) = \{w\} \cup (\bigcup_{i=0}^{\infty} S_i) = W$ and we thus have $V^*(\mu X.[a]X) = W$ by Lemma 3.2.11. In particular, w satisfies $\mu X.[a]X$ which leads to $w \Vdash \psi$ by supposition. Consider the filtration $\mathcal{N}^{\Sigma}$ of $\mathcal{N}$ through $\Sigma = FL(\{\psi\})$; since $w \Vdash \psi$ we have $\mathcal{N}^{\Sigma}, |w| \Vdash \psi$ and thus by supposition $\mathcal{N}^{\Sigma}, |w| \Vdash \mu X.[a]X$.

Now, the set $FL(\{\psi\})$ is finite so we have $|\mathcal{P}(FL(\{\psi\}))| = m \in \mathbb{N}$ and thus by Lemma 3.2.11 $V^{\Sigma *}(\mu X.[a]X) = \bigcup_{n=0}^{m} l^n_{R^{\Sigma}_{\langle a\rangle}}(\varnothing)$ i.e. $\mathcal{N}^{\Sigma}, |w| \Vdash \mu X.[a]X$ iff $|w|$ is deadlocked or every $R^{\Sigma}_{\langle a\rangle}$-path from $|w|$ ends in a deadlocked state (see Definition 2.1.7). We will show there is an infinite $R^{\Sigma}_{\langle a\rangle}$-path from $|w|$ and hence that $|w| \notin \bigcup_{n=0}^{m} l^n_{R^{\Sigma}_{\langle a\rangle}}(\varnothing)$ i.e. $\mathcal{N}^{\Sigma}, |w| \nVdash \mu X.[a]X$, a contradiction. Consider the $R_{\langle a\rangle}$-path $ww_{m,0}w_{m-1,0}...w_{0,m}$ from w of length $m+2$. For ease of reference we shall denote this path by $(u_i)_{i=0}^{m+1}$. The sequence $(|u_i|)_{i=0}^{m+1}$ is a $R^{\Sigma}_{\langle a\rangle}$-path from $|w|$ since $\mathcal{N}^{\Sigma}$ is a filtration. The filtration $\mathcal{N}^{\Sigma}$ must identify at least three states in $(|u_i|)_{i=0}^{m+1}$ by the pigeonhole principle, since $|\mathcal{N}^{\Sigma}| \leq m$ (see Remark 3.2.22). This will create a loop in $\mathcal{N}^{\Sigma}$: there are $j, k \in \{0, ..., m+1\}$ such that $j < k$-1, $|u_j| = |u_k|$ and $R^{\Sigma}_{\langle a\rangle}|u_i||u_{i+1}|$ for all $i \in \{j, ..., k\text{-}1\}$ i.e.

[2]States which have no neighbours, such as those in S_0, are referred to as *deadlocked states*.

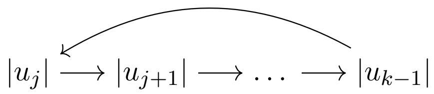

Since $\mathcal{N}^{\Sigma}$ has such a loop, and hence an infinite $R^{\Sigma}_{\langle a \rangle}$-path from $|w|$, we have a contradiction. $\square$

3.2.22 Remark. Note that if we instead consider the $R_{\langle a \rangle}$-path from w of length $m+1$ (call it $(s_i)_{i=0}^{m}$) and hence identify at least two states amongst $(|s_i|)_{i=0}^{m}$ in $\mathcal{N}^{\Sigma}$, we will not necessarily get a loop in the filtration. This is due to the fact that two adjacent states in $(|s_i|)_{i=0}^{m}$ may be identified (for example $|s_0| = |s_1|$) which will not generate a loop. The fact that in the proof three states amongst $(|u_i|)_{i=0}^{m+1}$ are identified in $\mathcal{N}^{\Sigma}$ guarantees that it is possible to identify two non-adjacent states (in the proof the stipulation that $j < k$-1 emphasizes that the two states identified are not adjacent).

4. Algebraic Filtrations

The results in this chapter are taken from the sources [8] and [7]. For each result, the proof is provided when we feel it offers some insight. When a proof is provided, we fill out more details than those found in the source material.

We have seen that a possible way to demonstrate the decidability of a finitely axiomatisable logic Λ is to show that Λ has the finite model property with respect to a class of frames. Analogously in the algebraic setting, if a finitely axiomatisable logic Λ has the finite model property with respect to a class of modal algebras (Definition 4.1.1) then Λ is decidable i.e. Λ is shown to be decidable algebraically.

A natural way to attempt a proof of the finite model property for a logic which is characterised by a class of frames is by using a filtration of a model (see Section 2.4). Similarly, algebraic filtrations can be used to demonstrate the finite model property for a logic which is characterised by a class of algebras i.e. for a given class of modal algebras $\mathfrak{A}$, algebraic filtrations can be used to show that $\mathfrak{A}_{fin} \models \varphi^{\approx}$ implies $\mathfrak{A} \models \varphi^{\approx}$. In this chapter we shall look at the paper of Conradie et al. [7] which defines algebraic filtrations; in particular we will look at algebraic filtrations of τ_0 modal algebras (recall from Definition 3.1.1 that $\tau_0 = \{\langle a \rangle\}_{a \in \Pi_0}$); this will be useful in the next chapter when algebraically showing that **PDL** is decidable (Section 5.1) and that $\Lambda_{\mathfrak{R}}$, the equational theory of $\mathfrak{R}$, has the finite model property with respect to $\mathfrak{R}$ (Section 5.2). We begin with some background needed to discuss algebraic filtrations.

4.1 Preliminaries

4.1.1 Definition. Let Λ be a logic and let $\mathfrak{A}$ be a class of modal algebras. Λ has the *finite model property* (fmp) *with respect to* $\mathfrak{A}$ if the following property holds: $\vdash_\Lambda \varphi$ if and only if $\mathfrak{A}_{fin} \models \varphi^{\approx}$ for all $\varphi \in T(\Phi)$.

4.1.2 Remark. If Λ is sound and complete with respect to $\mathfrak{A}$ and $\mathfrak{A}$ has the property that $\mathfrak{A} \models \varphi^{\approx}$ iff $\mathfrak{A}_{fin} \models \varphi^{\approx}$ then Λ has the finite model property with respect to $\mathfrak{A}$. To see this note that $\vdash_\Lambda \varphi$ iff $\vdash_\Lambda \varphi \leftrightarrow \top$. Since Λ is sound and complete with respect to $\mathfrak{A}$ we have $\vdash_\Lambda \varphi \leftrightarrow \top$ iff $\mathfrak{A} \models \varphi^{\approx}$. But $\mathfrak{A} \models \varphi^{\approx}$ iff $\mathfrak{A}_{fin} \models \varphi^{\approx}$ and we thus have $\vdash_\Lambda \varphi$ iff $\mathfrak{A}_{fin} \models \varphi^{\approx}$ i.e. Λ has the finite model property with respect to $\mathfrak{A}$.

4.1.3 Definition. Let $\mathbf{L} = (L, \vee, \wedge)$ be a bounded lattice. An element $x \in L$ is called an *atom of* $\mathbf{L}$ if $0 < x$ and there does not exist any element $z \in L$ such that $0 < z < x$. The set of atoms of $\mathbf{L}$

is denoted $At\mathbf{L}$. The lattice $\mathbf{L}$ is atomic if for every non-zero $z \in L$ there exists $x \in At\mathbf{L}$ such that $x \leq z$.

4.1.4 Example. Every finite lattice is atomic. Also, for a set W, the power set algebra $\mathcal{P}(W)$ is atomic with $At\mathcal{P}(W) = \{\{w\} \mid w \in W\}$.

4.1.5 Lemma. Let $\mathbf{L}$ be a bounded distributive lattice. If $x \in At\mathbf{L}$ and $x \leq \bigvee_{i=1}^{n} z_i$ then $x \leq z_j$ for some $j \in \{1, ..., n\}$.

Proof. Since $x \leq \bigvee_{i=1}^{n} z_i$ we have $x = x \wedge \bigvee_{i=1}^{n} z_i = \bigvee_{i=1}^{n}(x \wedge z_i)$. Also $x \geq x \wedge z_i$ for all $i \in \{1, ..., n\}$. Suppose towards a contradiction that $x > x \wedge z_i$ for all $i \in \{1, ..., n\}$. Then since $x \in At\mathbf{L}$ we have $x \wedge z_i = 0$ for all $i \in \{1, ..., n\}$ which leads to $x = \bigvee_{i=1}^{n}(x \wedge z_i) = 0$, a contradiction. Thus there must exist a z_j for some $j \in \{1, ..., n\}$ such that $x = x \wedge z_j$ i.e. $x \leq z_j$. □

4.1.6 Lemma. Let $\mathbf{L} = (L, \vee, \wedge, \neg, 0, 1)$ be a finite Boolean algebra. Then for every $z \in L$ we have $z = \bigvee\{x \in At\mathbf{L} \mid x \leq z\}$.

Proof. The element z is obviously an upper bound of $\{x \in At\mathbf{L} \mid x \leq z\}$. Suppose towards a contradiction that z is not the supremum of $\{x \in At\mathbf{L} \mid x \leq z\}$ i.e. there exists a $z' \in L$ such that z' is an upper bound of $\{x \in At\mathbf{L} \mid x \leq z\}$ and $z \not\leq z'$. It can be easily shown that $z \not\leq z'$ implies $z \wedge \neg z' \neq 0$ in $\mathbf{L}$. Since $\mathbf{L}$ is atomic there exists $x \in At\mathbf{L}$ such that $x \leq z \wedge \neg z'$. Since $x \leq z \wedge \neg z'$ we have $x \in \{x \in At\mathbf{L} \mid x \leq z\}$; therefore $x \leq z'$ (since z' is an upper bound of $\{x \in At\mathbf{L} \mid x \leq z\}$) and we have $x \leq z' \wedge \neg z' = 0$, a contradiction. Thus $z = \bigvee\{x \in At\mathbf{L} \mid x \leq z\}$. □

4.2 Constructing a finite falsifying algebra

Recall from Section 2.5 that a τ_0 modal algebra is an algebra $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle a \rangle})_{a \in \Pi_0}$ where $(A, \vee, \wedge, \neg, 0, 1)$ is a Boolean algebra and $f_{\langle a \rangle}$ is an operator on $\mathbf{A}$ for each $a \in \Pi_0$. Let $\mathfrak{A}$ be a class of τ_0 modal algebras for which we want to prove that the validities of $\mathfrak{A}$ and $\mathfrak{A}_{fin}$ are the same. In particular if $\mathbf{A} \in \mathfrak{A}$ with $\mathbf{A} \not\models \sigma^{\approx}$, it needs to be shown that there exists a finite falsifying algebra in $\mathfrak{A}$ i.e. $\mathbf{A}' \in \mathfrak{A}_{fin}$ such that $\mathbf{A}' \not\models \sigma^{\approx}$.

If it can be shown that for an $\mathbf{A}$-assignment h falsifying $\sigma^{\approx}$, the finite set $\{h^*(\psi) \mid \psi \in st(\sigma)\}$ is embedded in an algebra from $\mathfrak{A}_{fin}$ such that the existing fundamental operations on $\{h^*(\psi) \mid \psi \in st(\sigma)\}$ are preserved, then we are done. Motivated by this, we consider the general problem:

For $\mathbf{A} \in \mathfrak{A}$ with $C \subseteq_{fin} A$ show that there exists $\mathbf{A}' \in \mathfrak{A}_{fin}$ and an embedding from C to $\mathbf{A}'$ such that the existing fundamental operations on C are preserved.

It is tempting to simply take the modal algebra generated by C as our $\mathbf{A}'$, but there is no guarantee it will be finite due to the presence of the operators $(f_{\langle a \rangle})_{a \in \Pi_0}$. Instead, towards constructing such an $\mathbf{A}'$, Conradie et al. [7] consider the Boolean algebra generated by C, denoted $\mathbf{A}_C = (A_C, \vee, \wedge, \neg, 0, 1)$. The algebra $\mathbf{A}_C$ is finite since C is finite (see [8]). Also C is embedded in $\mathbf{A}_C$ under the trivial embedding and $\mathbf{A}_C$ certainly preserves all existing Boolean operations on C. It now remains to construct operators $(f'_{\langle a \rangle})_{a \in \Pi_0}$ on $\mathbf{A}_C$ which preserve the existing operators $(f_{\langle a \rangle})_{a \in \Pi_0}$ on C i.e. for $a \in \Pi_0$ and $z \in \underline{C}_a = \{x \in C \mid f_{\langle a \rangle}(x) \in C\}$ we have $f'_{\langle a \rangle}(z) = f_{\langle a \rangle}(z)$. In this case the operators $(f'_{\langle a \rangle})_{a \in \Pi_0}$ are said to extend $(f_{\langle a \rangle}|_{\underline{C}_a})_{a \in \Pi_0}$. Also it must be checked that $(\mathbf{A}_C, f'_{\langle a \rangle})_{a \in \Pi_0}$ is in $\mathfrak{A}$.

The algebra $\mathbf{A}_C$ is atomic and thus by Lemma 4.1.6, operators on $\mathbf{A}_C$ are completely determined by their actions on $At\mathbf{A}_C$: for an operator $f_{\langle a \rangle}$ on $\mathbf{A}_C$ and $z \in A_C$ we have $f_{\langle a \rangle}(z) = f_{\langle a \rangle}(\bigvee\{x \in At\mathbf{A}_C \mid x \leq z\}) = \bigvee\{f_{\langle a \rangle}(x) \mid x \in At\mathbf{A}_C \;\&\; x \leq z\}$. Furthermore, there is a relationship between functions $t : At\mathbf{A}_C \to A_C$ and relations $R \subseteq At\mathbf{A}_C^2$.

4.2.1 Definition. Let $A' \subseteq A_C$ and let $t : A' \to A_C$; define $R^t = \{(x, y) \in A'^2 \mid x \leq t(y)\}$.

4.2.2 Definition. Let $R \subseteq At\mathbf{A}_C^2$. Define the mapping $t^R : At\mathbf{A}_C \to A_C$ by $t^R(y) = \bigvee\{x \in At\mathbf{A}_C \mid Rxy\}$.

4.2.3 Theorem. *Consider the set of mappings from $At\mathbf{A}_C$ to A_C, denoted $A_C^{At\mathbf{A}_C}$; the mapping $e : A_C^{At\mathbf{A}_C} \to \mathcal{P}(At\mathbf{A}_C^2)$ defined by $e(t) = R^t$ is a bijection.*

Proof. If it can be shown that the mapping $e' : \mathcal{P}(At\mathbf{A}_C^2) \to A_C^{At\mathbf{A}_C}$ defined by $e'(R) = t^R$, is the inverse of e then we are done. Let $R \in \mathcal{P}(At\mathbf{A}_C^2)$ then $e(e'(R)) = R^{t^R}$. We have $R^{t^R}xy$ iff $x \leq t^R(y)$ iff $x \leq \bigvee\{z \in At\mathbf{A}_C \mid Rzy\}$. By Lemma 4.1.5 $x \leq z$ for some $z \in \{z \in At\mathbf{A}_C \mid Rzy\}$ and thus $x = z$ (otherwise this would contradict z being an atom of $\mathbf{A}_C$). Thus $x \leq \bigvee\{z \in At\mathbf{A}_C \mid Rzy\}$ iff Rxy and $e(e'(R)) = R$. For the other direction $e'(e(t)) = t^{R^t}$; we want to check that $t^{R^t} = t$: let $x \in At\mathbf{A}_C$ then $t^{R^t}(x) = \bigvee\{y \in At\mathbf{A}_C \mid R^t yx\} = \bigvee\{y \in At\mathbf{A}_C \mid y \leq t(x)\} = t(x)$ (where the last equality follows from Lemma 4.1.6). $\square$

Thus there is a one-to-one correspondence between relations $R \subseteq At\mathbf{A}_C^2$ and functions $t : At\mathbf{A}_C \to A_C$. For $R \subseteq At\mathbf{A}_C^2$ the function $t^R : At\mathbf{A}_C \to A_C$ gives rise to an operator on $\mathbf{A}_C$.

4.2.4 Lemma. Let $R \subseteq At\mathbf{A}_C^2$. The mapping t^R has a unique extension to an operator f^R on $\mathbf{A}_C$ i.e. the following diagram commutes (where $e_{At\mathbf{A}_C} : At\mathbf{A}_C \to A_C$ is defined by $e_{At\mathbf{A}_C}(x) = x$).

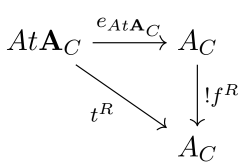

Proof. For $z \in A_C$ define $f^R(z) = \bigvee\{t^R(y) \mid y \in At\mathbf{A}_C \;\&\; y \leq z\}$; therefore

$$\begin{aligned} f^R(z) &= \bigvee\{\bigvee\{x \in At\mathbf{A}_C \mid Rxy\} \mid y \in At\mathbf{A}_C \;\&\; y \leq z\} \\ &= \bigvee\{x \in At\mathbf{A}_C \mid (\exists y \in At\mathbf{A}_C)(Rxy \;\&\; y \leq z)\}. \end{aligned}$$

When $z \in At\mathbf{A}_C$ we have $f^R(z) = \bigvee\{t^R(z)\} = t^R(z)$ and thus f^R extends t^R. The mapping f^R is an operator since $f^R(0) = \bigvee \varnothing = 0$ and for any $y_1, y_2 \in A_C$ we have $f^R(y_1 \vee y_2) = f^R(y_1) \vee f^R(y_2)$ from Lemma 4.1.5. Finally, f^R uniquely extends t^R since operators are uniquely determined by their actions on $At\mathbf{A}_C$. $\square$

Towards constructing operators $(f'_{\langle a\rangle})_{a\in\Pi_0}$ on $\mathbf{A}_C$ such that $(f'_{\langle a\rangle})_{a\in\Pi_0}$ extend $(f_{\langle a\rangle}|_{\underline{C}_a})_{a\in\Pi_0}$, we have seen that we can instead look at relations $R \subseteq At\mathbf{A}_C^2$ since they give rise to functions $t^R : At\mathbf{A}_C \to A_C$ which in turn give rise to operators $f^R : A_C \to A_C$. Thus given a collection $(R_{\langle a\rangle})_{a\in\Pi_0}$ of binary relations on $At\mathbf{A}_C$ we can take $(f^{R_{\langle a\rangle}})_{a\in\Pi_0}$ as our desired $(f'_{\langle a\rangle})_{a\in\Pi_0}$. However in order for $(f^{R_{\langle a\rangle}})_{a\in\Pi_0}$ to extend $(f_{\langle a\rangle}|_{\underline{C}_a})_{a\in\Pi_0}$, the relations $(R_{\langle a\rangle})_{a\in\Pi_0}$ need to each satisfy a precise requirement.

4.2.5 Definition. Let $R \subseteq At\mathbf{A}_C^2$ and let $a \in \Pi_0$. Condition $(R)_a$ is defined as

$$(R)_a:\; (\forall z \in \underline{C}_a)(\forall x \in At\mathbf{A}_C)(x \leq f_{\langle a\rangle}(z) \text{ iff } (\exists y \in At\mathbf{A}_C)(Rxy \;\&\; y \leq z)).$$

4.2.6 Theorem. *Let $R \subseteq At\mathbf{A}_C^2$ and let $a \in \Pi_0$. If R satisfies condition $(R)_a$ then f^R extends $f_{\langle a\rangle}|_{\underline{C}_a}$.*

Proof. Let $R \subseteq At\mathbf{A}_C^2$ such that R satisfies condition $(R)_a$ and let $z \in \underline{C}_a$. We have $f^R(z) = \bigvee\{x \in At\mathbf{A}_C \mid (\exists y \in At\mathbf{A}_C)(Rxy \;\&\; y \leq z)\} = \bigvee\{x \in At\mathbf{A}_C \mid x \leq f_{\langle a\rangle}(z)\}$ where the last equality follows since R satisfies $(R)_a$. But by Lemma 4.1.6 we have $\bigvee\{x \in At\mathbf{A}_C \mid x \leq f_{\langle a\rangle}(z)\} = f_{\langle a\rangle}(z)$. $\square$

4.2.7 Theorem. *Let $a \in \Pi_0$. If f' is an operator on $\mathbf{A}_C$ extending $f_{\langle a\rangle}|_{\underline{C}_a}$ then $R^{f'}$ satisfies condition $(R^{f'})_a$.*

Proof. Suppose f' is an operator on $\mathbf{A}_C$ extending $f_{\langle a\rangle}|_{\underline{C}_a}$, let $z \in \underline{C}_a$ and let $x \in At\mathbf{A}_C$. We need to show that $R^{f'}$ satisfies condition $(R^{f'})_a$. Suppose $x \leq f_{\langle a\rangle}(z)$; since f' extends $f_{\langle a\rangle}|_{\underline{C}_a}$ we have $x \leq f'(z)$. By Lemma 4.1.6 we thus get $x \leq f'(\bigvee\{y \in At\mathbf{A}_C \mid y \leq z\}) = \bigvee\{f'(y) \mid y \in At\mathbf{A}_C \;\&\; y \leq z\}$ where the last equality follows since f' is an operator. Therefore by Lemma 4.1.5 we have $x \leq f'(y')$ for some $y' \in At\mathbf{A}_C$ such that $y' \leq z$. But $x \leq f'(y')$ iff $R^{f'}xy'$ and thus the left to right conditional of $(R^{f'})_a$ holds. For the reverse conditional suppose there exists $y \in At\mathbf{A}_C$ such that $R^{f'}xy$ and $y \leq z$. Since $R^{f'}xy$ we have $x \leq f'(y)$ and since $y \leq z$ and f' is an operator we have $f'(y) \leq f'(z)$. Thus $x \leq f'(z) = f_{\langle a\rangle}(z)$. $\square$

We are now ready to define an algebraic filtration of the τ_0 modal algebra $\mathbf{A}$.

4.2.8 Definition. Let $(R_{\langle a\rangle})_{a\in\Pi_0}$ be a collection of binary relations on $At\mathbf{A}_C$. If $R_{\langle a\rangle}$ satisfies $(R)_a$ for all $a \in \Pi_0$ then the collection $(R_{\langle a\rangle})_{a\in\Pi_0}$ is an *algebraic filtrator of* $\mathbf{A}$ *through* $(C, \underline{C}_a)_{a\in\Pi_0}$ and the τ_0 modal algebra $(\mathbf{A}_C, f^{R_{\langle a\rangle}})_{a\in\Pi_0}$ is called an *algebraic filtration of* $\mathbf{A}$ *through* $(C, \underline{C}_a)_{a\in\Pi_0}$ *with* $(R_{\langle a\rangle})_{a\in\Pi_0}$.

4.2.9 Example. Consider the collection $(R^s_{\langle a\rangle})_{a\in\Pi_0}$ of binary relations on $At\mathbf{A}_C$ defined by $R^s_{\langle a\rangle}xy$ iff $x \wedge f_{\langle a\rangle}(y) \neq 0$ for all $a \in \Pi_0$. The relation $R^s_{\langle a\rangle}$ satisfies condition $(R^s_{\langle a\rangle})_a$ for all $a \in \Pi_0$. To see this let $a \in \Pi_0$, let $z \in \underline{C}_a$ and let $x \in At\mathbf{A}_C$. Suppose $x \leq f_{\langle a\rangle}(z)$; thus $x \leq f_{\langle a\rangle}(\bigvee\{y \in At\mathbf{A}_C \mid y \leq z\}) = \bigvee\{f_{\langle a\rangle}(y) \mid y \in At\mathbf{A}_C \;\&\; y \leq z\}$. By Lemma 4.1.5 $x \leq f_{\langle a\rangle}(y')$ for some $y' \in At\mathbf{A}_C$ such that $y' \leq z$. $R^s_{\langle a\rangle}xy'$ since $x \wedge f_{\langle a\rangle}(y') = x$ and $x \neq 0$ since $x \in At\mathbf{A}_C$. For the reverse direction of $(R^s_{\langle a\rangle})_a$ suppose there exists $y \in At\mathbf{A}_C$ such that $R^s_{\langle a\rangle}xy$ and $y \leq z$. Since $f_{\langle a\rangle}$ is an operator $f_{\langle a\rangle}(y) \leq f_{\langle a\rangle}(z)$ and since $R^s_{\langle a\rangle}xy$ we have $x \wedge f_{\langle a\rangle}(y) \neq 0$ which implies $x \wedge f_{\langle a\rangle}(y) = x$ (i.e. $x \leq f_{\langle a\rangle}(y)$) since $x \in At\mathbf{A}_C$. Therefore $x \leq f_{\langle a\rangle}(z)$.

4.2.10 Example. As mentioned in Chapter 1 the filtration method was first developed by McKinsey [31] in his proof that S_2 and S_4 are decidable logics. Towards constructing a finite falsifying basic modal algebra, McKinsey defined the operator $f'(z) = \bigwedge\{x \in A_C \mid (\exists y \in \underline{C})(x = f(y) \;\&\; z \leq y)\}$ on $\mathbf{A}_C$. It can be shown (see [7]) that the modal algebra $(\mathbf{A}_C, f')$ is the algebraic filtration of $\mathbf{A}$ through $(C, \underline{C})$ with R^l, where R^l is the largest algebraic filtrator of $\mathbf{A}$ through $(C, \underline{C})$ and is defined as $R^l xy$ iff $y \leq z$ implies $x \leq f(z)$ for all $z \in \underline{C}$.

4.2.11 Theorem. *Let* $\mathbf{A}$ *be a* τ_0 *modal algebra with* $\mathbf{A} \not\models \sigma^{\approx}$ *for* $\sigma \in T_{\tau_0}(\Phi)$ *(i.e. there is an* $h : \Phi \to A$ *such that* $h^*(\sigma) \neq 1$*), let* $(\mathbf{A}_C, f^{R_{\langle a\rangle}})_{a\in\Pi_0}$ *be any algebraic filtration of* $\mathbf{A}$ *through* $(C, \underline{C}_a)_{a\in\Pi_0}$ *where* $\{h^*(\psi) \mid \psi \in st(\sigma)\} \subseteq C \subseteq_{fin} A$ *and define* $j : \Phi \to A_C$ *by* $j(p) = h(p)$*. Then* $j^*(\psi) = h^*(\psi)$ *for all* $\psi \in st(\sigma)$ *and thus* $(\mathbf{A}_C, f^{R_{\langle a\rangle}})_{a\in\Pi_0} \not\models \sigma^{\approx}$*.*

Proof. Let $(\mathbf{A}_C, f^{R_{\langle a\rangle}})_{a\in\Pi_0}$ be an algebraic filtration of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ where $\{h^*(\psi) \mid \psi \in st(\sigma)\} \subseteq C \subseteq_{fin} A$. Towards showing $(\mathbf{A}_C, f^{R_{\langle a\rangle}})_{a\in\Pi_0} \not\models \sigma^{\approx}$, define $j : \Phi \to A_C$ by $j(p) = h(p)$[1] and consider the homomorphic extension $j^* : T_{\tau_0}(\Phi) \to A_C$. To demonstrate $j^*(\sigma) \neq 1$, it suffices to show for all $\psi \in st(\sigma)$ that

$$j^*(\psi) = h^*(\psi) \tag{4.2.1}$$

since this would imply $j^*(\sigma) = h^*(\sigma) \neq 1$. Thus we proceed by induction:

Base case: Let $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) = 0$; then $\psi = p \in \Phi$ or $\psi = \bot$. If $\psi = p$ then equation 4.2.1 is true since $j^*(p) = j(p) = h(p) = h^*(p)$ by definition of j. If $\psi = \bot$ then equation 4.2.1 is true since $j^*(\bot) = 0 = h^*(\bot)$.

[1] We only actually require $j(p) = h(p)$ when $p \in \Phi \cap st(\sigma)$.

Inductive hypothesis: Assume that equation 4.2.1 is true for all $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) < n \leq \mathrm{k}(\sigma)$.

Inductive step: Let $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) = n \leq \mathrm{k}(\sigma)$. Then for some $\varphi, \varphi_1, \varphi_2 \in T_{\tau_0}(\Phi)$ and $a \in \Pi_0$:

$$\psi = \begin{cases} \varphi_1 \vee \varphi_2 & (i) \\ \neg\varphi & (ii) \\ \langle a \rangle \varphi & (iii). \end{cases}$$

(i) We have $j^*(\varphi_1 \vee \varphi_2) = j^*(\varphi_1) \vee_{\mathbf{A}_C} j^*(\varphi_2)$. We are permitted to form the element $h^*(\varphi_1) \vee_{\mathbf{A}_C} h^*(\varphi_2)$ since $h^*(\varphi_1), h^*(\varphi_2) \in A_C$ as $\{h^*(\psi) \mid \psi \in st(\sigma)\} \subseteq C$. But by inductive hypothesis $j^*(\varphi_1) = h^*(\varphi_1)$ and $j^*(\varphi_2) = h^*(\varphi_2)$ and therefore $j^*(\varphi_1) \vee_{\mathbf{A}_C} j^*(\varphi_2) = h^*(\varphi_1) \vee_{\mathbf{A}_C} h^*(\varphi_2)$. Since $h^*(\varphi_1), h^*(\varphi_2) \in A_C$ and $\mathbf{A}_C$ is the Boolean subalgebra generated by C and we have $h^*(\varphi_1) \vee_{\mathbf{A}_C} h^*(\varphi_2) = h^*(\varphi_1) \vee_{\mathbf{A}} h^*(\varphi_2) = h^*(\varphi_1 \vee \varphi_2)$.
(ii) Similarly, we have $j^*(\neg\varphi) = \neg_{\mathbf{A_C}} j^*(\varphi) = \neg_{\mathbf{A_C}} h^*(\varphi)$ by the inductive hypothesis; since $\mathbf{A_C}$ is the Boolean algebra generated by C we have $\neg_{\mathbf{A_C}} h^*(\varphi) = \neg_{\mathbf{A}} h^*(\varphi) = h^*(\neg\varphi)$.
(iii) We have $j^*(\langle a \rangle \varphi) = f^{R_{\langle a \rangle}}(j^*(\varphi)) = f^{R_{\langle a \rangle}}(h^*(\varphi))$ by the inductive hypothesis; also $h^*(\langle a \rangle \varphi) = f_{\langle a \rangle}(h^*(\varphi))$. Thus we need to show $f^{R_{\langle a \rangle}}(h^*(\varphi)) = f_{\langle a \rangle}(h^*(\varphi))$. If $h^*(\varphi) \in \underline{C}_a$ we are done since $f^{R_{\langle a \rangle}}$ extends $f_{\langle a \rangle}\big|_{\underline{C}_a}$ as $(R_{\langle a \rangle})_{a \in \Pi_0}$ is an algebraic filtrator of $\mathbf{A}$ through $(C, \underline{C}_a)_{a \in \Pi_0}$. We certainly have $f_{\langle a \rangle}(h^*(\varphi)) \in C$ since $f_{\langle a \rangle}(h^*(\varphi)) = h^*(\langle a \rangle \varphi) \in \{h^*(\psi) \mid \psi \in st(\sigma)\}$. □

4.2.12 Remark. This result demonstrates that in order to show a class of τ_0 modal algebras $\mathfrak{A}$ has the same validities as $\mathfrak{A}_{fin}$, the algebra $(\mathbf{A}_C, f^{R_{\langle a \rangle}})_{a \in \Pi_0}$ will provide a finite falsifying algebra; but it still needs to be checked if $(\mathbf{A}_C, f^{R_{\langle a \rangle}})_{a \in \Pi_0} \in \mathfrak{A}$.

The next lemma will be useful in Chapter 5 when algebraically proving **PDL** is decidable.

4.2.13 Lemma. Let $a \in \Pi_0$, let $R^s_{\langle a \rangle} \subseteq At\mathbf{A}_C^2$ such that $R^s_{\langle a \rangle} xy$ iff $x \wedge f_{\langle a \rangle}(y) \neq 0$ and define $f' : A_C \to A_C$ by $f'(z) = \bigwedge\{x' \in A_C \mid f_{\langle a \rangle}(z) \leq x'\}$. Then $f' = f^{R^s_{\langle a \rangle}}$.

Proof. Let $z \in A_C$. Denote the following sets $A_2 = \{x' \in A_C \mid f_{\langle a \rangle}(z) \leq x'\}$ and $A_1 = \{x \in At\mathbf{A}_C \mid (\exists y \in At\mathbf{A}_C)(R^s_{\langle a \rangle} xy \;\&\; y \leq z)\}$. We want to prove $\bigwedge A_2 = \bigvee A_1$. To show that $\bigwedge A_2 \geq \bigvee A_1$ it suffices to demonstrate that $\bigwedge A_2$ is an upper bound of A_1. In turn, to show $\bigwedge A_2$ is an upper bound of A_1 it suffices to demonstrate that for any $x \in A_1$, x is a lower bound of A_2. Let $x' \in A_2$; we have by definition of A_2 that $f_{\langle a \rangle}(z) \leq x'$ and so to prove x is a lower bound of A_2, we will show $x \leq f_{\langle a \rangle}(z)$. Since $x \in A_1$ there exists $y \in At\mathbf{A}_C$ such that $R^s_{\langle a \rangle} xy$ and $y \leq z$. $R^s_{\langle a \rangle} xy$ implies $x \leq f_{\langle a \rangle}(y)$ (since x is an atom of $\mathbf{A}_C$) and $y \leq z$ implies $f_{\langle a \rangle}(y) \leq f_{\langle a \rangle}(z)$ (since $f_{\langle a \rangle}$ is an operator) and thus $x \leq f_{\langle a \rangle}(z)$. To show the other inequality $\bigwedge A_2 \leq \bigvee A_1$ we demonstrate that

$\bigvee A_1 \in A_2$. Suppose towards a contradiction that $\bigvee A_1 \notin A_2$ i.e. $f_{\langle a \rangle}(z) \not\leq \bigvee A_1$. This implies $f_{\langle a \rangle}(z) \wedge \neg(\bigvee A_1) \neq 0$. Therefore there exists $x \in At\mathbf{A}_C$ such that $x \leq f_{\langle a \rangle}(z) \wedge \neg(\bigvee A_1)$. Thus $x \leq f_{\langle a \rangle}(z) = \bigvee\{f_{\langle a \rangle}(y) \mid y \in At\mathbf{A}_C \ \& \ y \leq z\}$ giving by Lemma 4.1.5 that $x \leq f_{\langle a \rangle}(y')$ for some $y' \in At\mathbf{A}_C$ such that $y' \leq z$. We have $R^s_{\langle a \rangle}xy'$ since $x \wedge f_{\langle a \rangle}(y') = x \neq 0$ and thus $x \in A_1$. But from $x \leq f_{\langle a \rangle}(z) \wedge \neg(\bigvee A_1)$ we know that $x \leq \neg(\bigvee A_1) = \bigwedge \neg A_1$[2] which implies $x \leq \neg x$ since $x \in A_1$. This gives $0 = x \wedge \neg x = x$, a contradiction since $x \in At\mathbf{A}_C$. □

It is natural to compare algebraic filtrations with filtrations on models and ask whether there is a relationship between the two. Conradie et al. [7] demonstrate that there is a one to one correspondence between algebraic filtrations satisfying an additional "rigidity" condition and model filtrations. In general there are more algebraic filtrations than model filtrations [7].

[2] $\neg A_1$ is defined $\{\neg x \mid x \in A_1\}$.

5. Decidability of Two Least Binder Logics

In this chapter we will first demonstrate that the fragment L^*_μ of the modal μ-calculus is decidable. This will be demonstrated by proving that **PDL** has the finite model property with respect to the class $\mathfrak{D}$ of dynamic algebras (Section 5.1). Then in Section 5.2 we will prove that $\Lambda_\mathfrak{R}$ has the finite model property with respect to $\mathfrak{R}$ where $\mathfrak{R}$ supplies an algebraic semantics for Pratt's least root language L_ρ [34]. These finite model property results are achieved using an algebraic filtration method based on that of Conradie et al. [7].

5.1 L^*_μ is decidable

L^*_μ is the fragment of L_μ where the fixed point binder μ may only be applied to formulae of the form $\psi(X) = \varphi \vee \Diamond X$ (see Definition 3.2.17). As shown in Theorem 3.2.19 *PDL* corresponds to L^*_μ and as such we will show L^*_μ is decidable by proving that the logic **PDL** is decidable. We begin with discussing the $\mathscr{F}$-closure (Definition 5.1.1), we then show that **PDL** is complete with respect to the class of dynamic algebras $\mathfrak{D}$ (Theorem 5.1.11) and finally that the validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same (Theorem 5.1.15).

5.1.1 $\mathscr{F}$-closure.

We now discuss the $\mathscr{F}$-closure, a closure analogous to the Fischer–Ladner closure (Definition 3.1.9). This will be useful when building a finite falsifying dynamic algebra in the proof that the validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same (see Theorem 5.1.15). Roughly speaking, the closure is needed to "educate" the algebraic filtration about the dynamic nature of programs.

5.1.1 Definition. Let $\Gamma \subseteq T_{\tau^*}(\Phi)$. Γ is *$\mathscr{F}$-closed* if for all $\varphi \in T_{\tau^*}(\Phi)$ and all π_1, π_2, π^* in Π:

(i) $\langle \pi_1; \pi_2 \rangle \varphi \in \Gamma$ implies $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi, \langle \pi_2 \rangle \varphi \in \Gamma$;

(ii) $\langle \pi_1 \cup \pi_2 \rangle \varphi \in \Gamma$ implies $\langle \pi_1 \rangle \varphi, \langle \pi_2 \rangle \varphi \in \Gamma$;

(iii) $\langle \pi^* \rangle \varphi \in \Gamma$ implies $\langle \pi \rangle \langle \pi^* \rangle \varphi \in \Gamma$.

The *$\mathscr{F}$-closure of* Γ, denoted $\mathscr{F}(\Gamma)$, is the smallest $\mathscr{F}$-closed set containing Γ (if it exists). For a singleton $\{\psi\} \subseteq T_{\tau^*}(\Phi)$ the set $\mathscr{F}(\{\psi\})$ is denoted $\mathscr{F}(\psi)$.

5.1.2 Example. Let $a, b, c \in \Pi_0$ and $p \in \Phi$; then $\mathscr{F}(\langle a; (b \cup c^*) \rangle p) = \{\langle a; (b \cup c^*) \rangle p, \langle a \rangle \langle b \cup c^* \rangle p,$ $\langle b \cup c^* \rangle p, \langle b \rangle p, \langle c^* \rangle p, \langle c \rangle \langle c^* \rangle p\}$.

5.1.3 Theorem. *For every $\Gamma \subseteq T_{\tau^*}(\Phi)$ the set $\mathscr{F}(\Gamma)$ exists. Furthermore if $\Gamma \subseteq_{fin} T_{\tau^*}(\Phi)$ then $\mathscr{F}(\Gamma) \subseteq_{fin} T_{\tau^*}(\Phi)$.*

Proof. To show the existence of $\mathscr{F}(\Gamma)$ consider the set $C = \{\Sigma \in \mathcal{P}(T_{\tau^*}(\Phi)) \mid \Gamma \subseteq \Sigma \;\&\; \Sigma$ is $\mathscr{F}$-closed$\}$ which is non-empty since $T_{\tau^*}(\Phi)$ is $\mathscr{F}$-closed. The set $\bigcap C \subseteq T_{\tau^*}(\Phi)$ gives the desired result.

Let $\Gamma \subseteq_{fin} T_{\tau^*}(\Phi)$; we will demonstrate $\mathscr{F}(\Gamma) \subseteq_{fin} T_{\tau^*}(\Phi)$. If we can show that for all $\Gamma_1, \Gamma_2 \subseteq T_{\tau^*}(\Phi)$

$$\mathscr{F}(\Gamma_1 \cup \Gamma_2) = \mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2) \tag{5.1.1}$$

then it will follow that $\mathscr{F}(\Gamma) = \bigcup_{\psi \in \Gamma} \mathscr{F}(\psi)$ and thus it will only remain to show $\mathscr{F}(\psi) \subseteq_{fin} T_{\tau^*}(\Phi)$ for any $\psi \in T_{\tau^*}(\Phi)$.

Let $\Gamma_1, \Gamma_2 \subseteq T_{\tau^*}(\Phi)$. Towards demonstrating equation 5.1.1 we prove $\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2)$ is $\mathscr{F}$-closed. Denote the set $\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2)$ as Σ. Suppose we have $\langle \pi \rangle \varphi \in \Sigma$ where $\varphi \in T_{\tau^*}(\Phi)$ and $\pi \in \Pi \setminus \Pi_0$; then for some $\pi_1, \pi_2 \in \Pi$ we have either:

$$\pi = \begin{cases} \pi_1 ; \pi_2 & (i) \\ \pi_1 \cup \pi_2 & (ii) \\ {\pi_1}^* & (iii). \end{cases}$$

(i) $\pi = \pi_1;\pi_2$ and thus $\langle \pi_1;\pi_2 \rangle \varphi \in \Sigma$. Then $\langle \pi_1;\pi_2 \rangle \varphi \in \mathscr{F}(\Gamma_1)$ or $\langle \pi_1;\pi_2 \rangle \varphi \in \mathscr{F}(\Gamma_2)$. If $\langle \pi_1;\pi_2 \rangle \varphi \in \mathscr{F}(\Gamma_1)$ then since $\mathscr{F}(\Gamma_1)$ is $\mathscr{F}$-closed $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi, \langle \pi_2 \rangle \varphi \in \mathscr{F}(\Gamma_1)$ and thus $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi$, $\langle \pi_2 \rangle \varphi \in \Sigma$. Similarly for $\langle \pi_1;\pi_2 \rangle \varphi \in \mathscr{F}(\Gamma_2)$. Thus $\langle \pi_1;\pi_2 \rangle \varphi \in \Sigma$ implies that $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi, \langle \pi_2 \rangle \varphi \in \Sigma$. Similar reasoning follows for cases (ii) and (iii). Thus Σ is $\mathscr{F}$-closed.

We now demonstrate equation 5.1.1. By Definition 5.1.1 of the $\mathscr{F}$-closure $\Gamma_1 \subseteq \mathscr{F}(\Gamma_1)$ and $\Gamma_2 \subseteq \mathscr{F}(\Gamma_2)$ thus $\Gamma_1 \cup \Gamma_2 \subseteq \mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2)$. Since, under set inclusion, $\mathscr{F}$ is an order-preserving function[1] we therefore have $\mathscr{F}(\Gamma_1 \cup \Gamma_2) \subseteq \mathscr{F}(\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2))$. Since we have proven $\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2)$ is $\mathscr{F}$-closed, we have $\mathscr{F}(\Gamma_1 \cup \Gamma_2) \subseteq \mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2)$. For the other direction we have by definition $\Gamma_1, \Gamma_2 \subseteq \Gamma_1 \cup \Gamma_2$ and since $\mathscr{F}$ is an order-preserving function $\mathscr{F}(\Gamma_1), \mathscr{F}(\Gamma_2) \subseteq \mathscr{F}(\Gamma_1 \cup \Gamma_2)$. Thus $\mathscr{F}(\Gamma_1 \cup \Gamma_2)$ is an upper bound of $\mathscr{F}(\Gamma_1)$ and $\mathscr{F}(\Gamma_2)$ and therefore $\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2) \subseteq \mathscr{F}(\Gamma_1 \cup \Gamma_2)$. Thus $\mathscr{F}(\Gamma_1) \cup \mathscr{F}(\Gamma_2) = \mathscr{F}(\Gamma_1 \cup \Gamma_2)$ and equation 5.1.1 is true.

We now finally need to show that for all $\psi \in T_{\tau^*}(\Phi)$ we have $\mathscr{F}(\psi)$ is finite. If ψ is not of the form $\langle \pi \rangle \varphi$, where $\pi \in \Pi$ and $\varphi \in T_{\tau^*}(\Phi)$, then $\mathscr{F}(\psi) = \{\psi\}$ since the set $\{\psi\}$ is $\mathscr{F}$-closed as the three conditionals of the $\mathscr{F}$-closure definition (Definition 5.1.1) are vacuously true. Therefore to complete the proof it needs to be demonstrated that for all $\pi \in \Pi$:

$$\mathscr{F}(\langle \pi \rangle \varphi) \text{ is finite for all } \varphi \in T_{\tau^*}(\Phi). \tag{5.1.2}$$

Base case: Let $\varphi \in T_{\tau^*}(\Phi)$ and let $\pi = a \in \Pi_0$. Statement 5.1.2 is true since vacuously $\mathscr{F}(\langle a \rangle \varphi) = \{\langle a \rangle \varphi\}$.

[1]This can be seen by Definition 5.1.1 of $\mathscr{F}$ as returning the *smallest* $\mathscr{F}$-closed set.

Inductive hypothesis: Assume statement 5.1.2 is true for all $\pi \in \Pi$ with $\mathrm{k}(\pi) < m$.

Inductive step: Let $\varphi \in T_{\tau^*}(\Phi)$ and let $\mathrm{k}(\pi) = m$. Then for some $\pi_1, \pi_2 \in \Pi$ we have either:

$$\pi = \begin{cases} \pi_1; \pi_2 & (I) \\ \pi_1 \cup \pi_2 & (II) \\ {\pi_1}^* & (III). \end{cases}$$

(I) If we can show that

$$\mathscr{F}(\langle \pi_1; \pi_2 \rangle \varphi) = \{\langle \pi_1; \pi_2 \rangle \varphi\} \cup \mathscr{F}(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi) \cup \mathscr{F}(\langle \pi_2 \rangle \varphi)$$

then statement 5.1.2 is true since by the inductive hypothesis $\mathscr{F}(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi)$ and $\mathscr{F}(\langle \pi_2 \rangle \varphi)$ are both finite. Similar to when proving equation 5.1.1 the proof of this equality comes down to showing $\{\langle \pi_1; \pi_2 \rangle \varphi\} \cup \mathscr{F}(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi) \cup \mathscr{F}(\langle \pi_2 \rangle \varphi)$ (denoted Ψ) is $\mathscr{F}$-closed. Thus suppose $\langle \alpha \rangle \chi \in \Psi$ for some $\alpha \in \Pi \setminus \Pi_0$ and $\chi \in T_{\tau^*}(\Phi)$. If $\langle \alpha \rangle \chi \in \mathscr{F}(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi)$ or $\langle \alpha \rangle \chi \in \mathscr{F}(\langle \pi_2 \rangle \varphi)$ then we are done. Otherwise $\langle \alpha \rangle \chi = \langle \pi_1; \pi_2 \rangle \varphi$; but we have $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi \in \mathscr{F}(\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi) \subseteq \Psi$ and $\langle \pi_2 \rangle \varphi \in \mathscr{F}(\langle \pi_2 \rangle \varphi) \subseteq \Psi$. Thus Ψ is $\mathscr{F}$-closed.
(II) Similarly to (I), the result follows by showing

$$\mathscr{F}(\langle \pi_1 \cup \pi_2 \rangle \varphi) = \{\langle \pi_1 \cup \pi_2 \rangle \varphi\} \cup \mathscr{F}(\langle \pi_1 \rangle \varphi) \cup \mathscr{F}(\langle \pi_2 \rangle \varphi).$$

(III) Similarly to (I), the result follows by showing

$$\mathscr{F}(\langle \pi^* \rangle \varphi) = \{\langle \pi^* \rangle \varphi\} \cup \mathscr{F}(\langle \pi \rangle \langle \pi^* \rangle \varphi).$$

Thus statement 5.1.2 is true for all $\pi \in \Pi$ and in all we have that $\Gamma \subseteq_{fin} T_{\tau^*}(\Phi)$ implies $\mathscr{F}(\Gamma) \subseteq_{fin} T_{\tau^*}(\Phi)$. □

5.1.4 Remark. As opposed to the Fischer–Ladner closure, the definition of the $\mathscr{F}$-closure omits the requirement of being subterm closed and thus facilitates a straightforward proof that $\mathscr{F}(\Gamma)$ if finite when Γ is finite. In the proof that the validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same (Theorem 5.1.15), we take the $\mathscr{F}$-closure of a subterm closed set towards building a finite falsifying dynamic algebra.

5.1.2 PDL is complete with respect to $\mathfrak{D}$.

In this subsection we will introduce the class $\mathfrak{D}$ of dynamic algebras; we shall then discuss how dynamic algebras interface with regular frames (see Definition 3.1.3) and finally we will show that the logic **PDL** (see Definition 3.1.8) is complete with respect to $\mathfrak{D}$.

5.1.5 Definition. [33] A *dynamic algebra* is a τ^* modal algebra $(A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi}$ such that :

$$\begin{array}{llcl} (i) & f_{\langle\pi_1;\pi_2\rangle}(x) & = & f_{\langle\pi_1\rangle}(f_{\langle\pi_2\rangle}(x)); \\ (ii) & f_{\langle\pi_1\cup\pi_2\rangle}(x) & = & f_{\langle\pi_1\rangle}(x) \vee f_{\langle\pi_2\rangle}(x); \\ (iii) & f_{\langle\pi^*\rangle}(x) & = & x \vee f_{\langle\pi\rangle}(f_{\langle\pi^*\rangle}(x)); \\ (iv) & f_{\langle\pi^*\rangle}(x) & \leq & x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)) \end{array}$$

for all $\pi, \pi_1, \pi_2 \in \Pi$ and all $x \in A$. The class of dynamic algebras is denoted by $\mathfrak{D}$.

Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi}$ be a dynamic algebra, let $\pi \in \Pi$ and let $x \in A$; property (iii) from the above definition demonstrates that $f_{\langle\pi^*\rangle}(x)$ is a fixed point of the function $t^\pi_{x,\mathbf{A}} : A \to A$ defined by $t^\pi_{x,\mathbf{A}}(y) = x \vee f_{\langle\pi\rangle}(y)$. Furthermore the next result (from the paper [33] by Pratt) shows that $f_{\langle\pi^*\rangle}(x)$ is the least fixed point (and least pre-fixed point) of $t^\pi_{x,\mathbf{A}}$. This is analogous to how PDL formulae involving $\langle\pi^*\rangle$ are interpreted relationally i.e. for a τ^*-model $(W, R_{\langle\pi\rangle}, V)_{\pi\in\Pi}$ and a $T_{\tau^*}(\Phi)$ formula $\langle\pi^*\rangle\,\varphi$, we have $V^*(\langle\pi^*\rangle\,\varphi) = V'(\mu(X)(\varphi \vee \langle\pi\rangle\, X))$ (see Theorem 3.2.19 and Remark 3.2.6). These results tell us that the class of dynamic algebras offers the intended semantics for PDL.

5.1.6 Definition. Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi}$ be a τ^* modal algebra. For a $\pi \in \Pi$ and $x \in A$, the function $t^\pi_{x,\mathbf{A}} : A \to A$ is defined as $t^\pi_{x,\mathbf{A}}(y) = x \vee f_{\langle\pi\rangle}(y)$.

5.1.7 Lemma. [33] Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi} \in \mathfrak{D}$, let $\pi \in \Pi$ and let $x \in A$. Then $f_{\langle\pi^*\rangle}(x)$ is the least fixed point (and least pre-fixed point) of $t^\pi_{x,\mathbf{A}}$.

Proof. (This proof utilises the same proof strategy as found in [33] but we fill out more details). The fact that $f_{\langle\pi^*\rangle}(x)$ is a fixed point (and thus a pre-fixed) of $t^\pi_{x,\mathbf{A}}$ is immediate from property (iii) of dynamic algebras (Definition 5.1.5). It suffices to show $f_{\langle\pi^*\rangle}(x)$ is the least pre-fixed point of $t^\pi_{x,\mathbf{A}}$ (this will imply $f_{\langle\pi^*\rangle}(x)$ is the least fixed point of $t^\pi_{x,\mathbf{A}}$ since any fixed point of $t^\pi_{x,\mathbf{A}}$ is a pre-fixed point of $t^\pi_{x,\mathbf{A}}$). Thus consider an arbitrary pre-fixed point z of $t^\pi_{x,\mathbf{A}}$:

$$x \vee f_{\langle\pi\rangle}(z) \leq z.$$

Therefore $x \leq z$ which gives $f_{\langle\pi^*\rangle}(x) \leq f_{\langle\pi^*\rangle}(z)$ since $f_{\langle\pi^*\rangle}$ is an operator on $\mathbf{A}$. Thus it remains to show $f_{\langle\pi^*\rangle}(z) \leq z$. We have from property (iv) of dynamic algebras that

$$f_{\langle\pi^*\rangle}(z) \leq z \vee f_{\langle\pi^*\rangle}(\neg z \wedge f_{\langle\pi\rangle}(z)).$$

If we can show $\neg z \wedge f_{\langle\pi\rangle}(z) = 0$ then we are done since this implies $z \vee f_{\langle\pi^*\rangle}(\neg z \wedge f_{\langle\pi\rangle}(z)) = z \vee f_{\langle\pi^*\rangle}(0) = z \vee 0 = z$ and hence $f_{\langle\pi^*\rangle}(z) \leq z$ (where $f_{\langle\pi^*\rangle}(0) = 0$ follows since $f_{\langle\pi^*\rangle}$ is an operator on **A**). Since z is pre-fixed point of $t^\pi_{x,\mathbf{A}}$ we have $x \vee f_{\langle\pi\rangle}(z) \leq z$ and from the fact that meets preserve order we get

$$\neg z \wedge (x \vee f_\pi(z)) \leq \neg z \wedge z = 0.$$

Since **A** is a distributive lattice we have

$$(\neg z \wedge x) \vee (\neg z \wedge f_\pi(z)) \leq 0$$

and thus $\neg z \wedge f_{\langle\pi\rangle}(z) = 0$ since $\neg z \wedge f_{\langle\pi\rangle}(z) \leq (\neg z \wedge x) \vee (\neg z \wedge f_\pi(z)) \leq 0$. □

The following result regarding τ^* modal algebras is also from Pratt's paper [33].

5.1.8 Lemma. [33] Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi}$ be a τ^* modal algebra. Suppose $f_{\langle\pi^*\rangle}(x)$ is the least pre-fixed point of $t^\pi_{x,\mathbf{A}}$ for all $\pi \in \Pi$ and all $x \in A$; then $f_{\langle\pi^*\rangle}(x) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$ for all $\pi \in \Pi$ and all $x \in A$.

Proof. (This proof utilises the same proof strategy as found in [33] but we fill out more details). Let $\pi \in \Pi$ and $x \in A$. Since $f_{\langle\pi^*\rangle}(x)$ is the least pre-fixed point of $t^\pi_{x,\mathbf{A}}$, to show $f_{\langle\pi^*\rangle}(x) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$ it suffices to prove that $x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$ is a pre-fixed point of $t^\pi_{x,\mathbf{A}}$, in particular it suffices to prove:

$$t^\pi_{x,\mathbf{A}}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x));$$

which is

$$x \vee f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)).$$

We will now prove this last inequality transitively by deriving a sequence of inequalities starting from the left hand side. Since **A** is distributive we have for all $y \in A$:

$$\begin{aligned} x \vee y &= 1 \wedge (x \vee y) \\ &= (x \vee \neg x) \wedge (x \vee y) \\ &= x \vee (\neg x \wedge y) \end{aligned}$$

i.e. for all $y \in A$ we have the identity $x \vee y = x \vee (\neg x \wedge y)$. Using this identity with $y = f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)))$ we get:

$$x \vee f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))) = x \vee (\neg x \wedge f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)))).$$

Since $f_{\langle\pi\rangle}$ is an operator on **A** we have

$$x \vee (\neg x \wedge f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)))) = x \vee (\neg x \wedge (f_{\langle\pi\rangle}(x) \vee f_{\langle\pi\rangle}f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))))$$

and since **A** is distributive

$$x \vee (\neg x \wedge (f_{\langle\pi\rangle}(x) \vee f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)))) = x \vee (\neg x \wedge f_{\langle\pi\rangle}(x)) \vee (\neg x \wedge f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))).$$

By the definition of a meet we know $\neg x \wedge f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)) \leq f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$ and thus, since joins preserve order, we have

$$x \vee (\neg x \wedge f_{\langle\pi\rangle}(x)) \vee (\neg x \wedge f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))) \leq x \vee (\neg x \wedge f_{\langle\pi\rangle}(x)) \vee (f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))).$$

By our initial assumption, we have for all $y \in A$ that $f_{\langle\pi^*\rangle}(y)$ is a pre-fixed point of $t^\pi_{y,\mathbf{A}}$ and thus in particular for $y = \neg x \wedge f_{\langle\pi\rangle}(x)$ we have $(\neg x \wedge f_{\langle\pi\rangle}(x)) \vee f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)) \leq f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$. This implies, since joins preserve order, that

$$x \vee (\neg x \wedge f_{\langle\pi\rangle}(x)) \vee f_{\langle\pi\rangle} f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x)).$$

In all we have proven that $x \vee f_{\langle\pi\rangle}(x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))) \leq x \vee f_{\langle\pi^*\rangle}(\neg x \wedge f_{\langle\pi\rangle}(x))$ and thus the result holds. □

We can use this last result to show that for any regular frame $\mathcal{F}$, the τ^* modal algebra $\mathcal{F}^+$ is a dynamic algebra.

5.1.9 Theorem. *Let* $\mathcal{F} = (W, R_{\langle\pi\rangle})_{\pi\in\Pi} \in RFrm$. *Then* $\mathcal{F}^+ = (\mathcal{P}(W), \cup, \cap, -, \varnothing, W, m_{R_{\langle\pi\rangle}})_{\pi\in\Pi}$ *is a dynamic algebra.*

Proof. $\mathcal{F}^+$ is a τ^* modal algebra. We have that $\mathcal{F}^+$ obeys properties (i) and (ii) of dynamic algebras from Theorem 3.1.6. It can be shown that Theorem 3.2.19 implies $m_{R_{\langle\pi^*\rangle}}(S)$ is the least fixed point (and least pre-fixed point) of $t^\pi_{S,\mathcal{F}^+}$ for all $\pi \in \Pi$ and all $S \in \mathcal{P}(W)$. Therefore $\mathcal{F}^+$ has property (iii) of dynamic algebras and by Lemma 5.1.8 property (iv) also holds. □

5.1.10 Remark. Although the complex algebra of a regular frame is a dynamic algebra, it is not the case that every dynamic algebra is the complex algebra of a regular frame. To see this let $\mathcal{F} = (W, R_{\langle\pi\rangle})_{\pi\in\Pi}$ be a countably infinite regular frame and consider the τ^* modal algebra $\mathbf{A} = (FC(W), \cup, \cap, -, \varnothing, W, m_{R_{\langle\pi\rangle}})_{\pi\in\Pi}$ where $FC(W) = \{S \in \mathcal{P}(W) \mid S \text{ is finite or cofinite}\}$. The operators on **A** obey the dynamic algebra properties since $\mathcal{F}$ is regular. Thus **A** is a dynamic algebra. However **A** is not the complex algebra of any frame since it can be checked that $FC(W)$ is countably infinite and by Cantor's theorem, no power set is countably infinite. In this way dynamic algebras are more general than regular frames.

5.1.11 Theorem. PDL *is complete with respect to* $\mathfrak{D}$.

Proof. From Theorem 2.6.9 it suffices to show that the Lindenbaum–Tarski algebra over **PDL** is a dynamic algebra i.e. $\mathcal{L}_{\mathbf{PDL}}(\Phi) \in \mathfrak{D}$. The algebra $\mathcal{L}_{\mathbf{PDL}}(\Phi)$ is a τ^* modal algebra (see Lemma 2.6.8). The operators on $\mathcal{L}_{\mathbf{PDL}}(\Phi)$ will be denoted by $f^{\mathcal{L}}_{\langle\pi\rangle}$ where $\pi \in \Pi$. Let $\varphi \in T_{\tau^*}(\Phi)$:

$$\begin{aligned} f^{\mathcal{L}}_{\langle \pi_1;\pi_2\rangle}[\varphi] &= [\langle \pi_1;\pi_2\rangle \varphi] \\ &= [\langle \pi_1\rangle \langle \pi_2\rangle \varphi] \\ &= f^{\mathcal{L}}_{\langle \pi_1\rangle} f^{\mathcal{L}}_{\langle \pi_2\rangle}[\varphi] \end{aligned}$$

where the second equality follows since $\vdash_{\mathbf{PDL}} \langle \pi_1;\pi_2\rangle \varphi \leftrightarrow \langle \pi_1\rangle \langle \pi_2\rangle \varphi$. Similarly since $\vdash_{\mathbf{PDL}}$ $\langle \pi_1 \cup \pi_2\rangle \varphi \leftrightarrow \langle \pi_1\rangle \varphi \vee \langle \pi_2\rangle \varphi$ it follows that

$$f^{\mathcal{L}}_{\langle \pi_1\cup\pi_2\rangle}[\varphi] = f^{\mathcal{L}}_{\langle \pi_1\rangle}[\varphi] \vee f^{\mathcal{L}}_{\langle \pi_2\rangle}[\varphi]$$

and since $\vdash_{\mathbf{PDL}} \langle \pi^*\rangle \varphi \leftrightarrow (\varphi \vee \langle \pi\rangle \langle \pi^*\rangle \varphi)$ it follows

$$f^{\mathcal{L}}_{\langle \pi^*\rangle}[\varphi] = [\varphi] \vee f^{\mathcal{L}}_{\langle \pi\rangle} f^{\mathcal{L}}_{\langle \pi^*\rangle}[\varphi].$$

It remains to demonstrate that $\mathcal{L}_{\mathbf{PDL}}(\Phi)$ satisfies property (iv) of dynamic algebras i.e. $f^{\mathcal{L}}_{\langle \pi^*\rangle}[\varphi] \leq [\varphi] \vee f^{\mathcal{L}}_{\langle \pi^*\rangle}(\neg[\varphi] \wedge f^{\mathcal{L}}_{\langle \pi\rangle}[\varphi])$. Note that

$$\begin{aligned} f^{\mathcal{L}}_{\langle \pi^*\rangle}[\varphi] &\leq [\varphi] \vee f^{\mathcal{L}}_{\langle \pi^*\rangle}(\neg[\varphi] \wedge f^{\mathcal{L}}_{\langle \pi\rangle}[\varphi]) \\ \text{iff } [\langle \pi^*\rangle \varphi] &\leq [\varphi] \vee f^{\mathcal{L}}_{\langle \pi^*\rangle}[\neg\varphi \wedge \langle \pi\rangle \varphi] \\ \text{iff } [\langle \pi^*\rangle \varphi] &\leq [\varphi \vee \langle \pi^*\rangle (\neg\varphi \wedge \langle \pi\rangle \varphi)]. \end{aligned}$$

For any $\varphi_1, \varphi_2 \in T_{\tau^*}(\Phi)$ we know $[\varphi_1] \leq [\varphi_2]$ iff $[\varphi_1] \vee [\varphi_2] = [\varphi_2]$ iff $[\varphi_1 \vee \varphi_2] = [\varphi_2]$ iff $\vdash_{\mathbf{PDL}}$ $(\varphi_1 \vee \varphi_2) \leftrightarrow \varphi_2$ iff $\vdash_{\mathbf{PDL}} \varphi_1 \to \varphi_2$ and thus it remains to show

$$\vdash_{\mathbf{PDL}} \langle \pi^*\rangle \varphi \to (\varphi \vee \langle \pi^*\rangle (\neg\varphi \wedge \langle \pi\rangle \varphi)).$$

We shall now give an informal proof of why this is so:

(1)	$\vdash_{\mathbf{PDL}} [\pi^*](\neg\varphi \to [\pi]\neg\varphi) \to (\neg\varphi \to [\pi^*]\neg\varphi)$	u.s. on axiom (vi) of **PDL** (Definition 3.1.8)
(2)	$\vdash_{\mathbf{PDL}} \neg \langle \pi^*\rangle (\neg\varphi \wedge \langle \pi\rangle \varphi) \to (\varphi \vee \neg \langle \pi^*\rangle \varphi)$	$[\pi^*], \to$ definition, u.s. on tautology
		$\neg\neg p \leftrightarrow p$ and u.s. on tautology
		$\neg(p \vee q) \leftrightarrow (\neg p \wedge \neg q)$
(3)	$\vdash_{\mathbf{PDL}} \langle \pi^*\rangle (\neg\varphi \wedge \langle \pi\rangle \varphi) \vee \varphi \vee \neg \langle \pi^*\rangle \varphi$	$\to$ definition
(4)	$\vdash_{\mathbf{PDL}} \langle \pi^*\rangle \varphi \to (\varphi \vee \langle \pi^*\rangle (\neg\varphi \wedge \langle \pi\rangle \varphi))$	$\to$ definition.

□

5.1.3 The validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same.

We need the next definition for the discussion that follows.

5.1.12 Definition. Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle \pi\rangle})_{\pi\in\Pi}$ be a τ^* modal algebra, let $C \subseteq A$, let $\Gamma \subseteq T_{\tau^*}(\Phi)$ and let $\pi \in \Pi$. Define $\underline{C}_\pi = \{x \in C \mid f_{\langle \pi\rangle}(x) \in C\}$ and $\overline{\Gamma}_\pi = \{\varphi \in \Gamma \mid \langle \pi\rangle \varphi \in \Gamma\}$.

5.1.13 Remark. For a homomorphism $g : T_{\tau^*}(\Phi) \to A$ we have $g(\overline{\Gamma}_\pi) \subseteq \underline{g(\Gamma)}_\pi$ (with equality when g is injective). To see this let $g(\psi) \in g(\overline{\Gamma}_\pi)$ i.e. $\psi, \langle\pi\rangle\,\psi \in \Gamma$. Since $\psi \in \Gamma$ we know $g(\psi) \in g(\Gamma)$. As g is a homomorphism $f_{\langle\pi\rangle}(g(\psi)) = g(\langle\pi\rangle\,\psi)$. But $g(\langle\pi\rangle\,\psi) \in g(\Gamma)$ as $\langle\pi\rangle\,\psi \in \Gamma$ and therefore $g(\psi) \in \underline{g(\Gamma)}_\pi$.

To complete the algebraic proof that **PDL** is decidable it needs to be demonstrated that the validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same i.e. $\mathfrak{D} \not\models \varphi^\approx$ implies $\mathfrak{D}_{fin} \not\models \varphi^\approx$. In particular we must demonstrate that if a dynamic algebra falsifies a PDL equation then there is a finite dynamic algebra which also falsifies the equation. The strategy for the proof is as follows:

- Suppose $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi} \in \mathfrak{D}$ with $\mathbf{A} \not\models \sigma^\approx$ for $\sigma \in T_{\tau^*}(\Phi)$, i.e. there is an $h : \Phi \to A$ such that $h^*(\sigma) \neq 1$. Define $\Gamma = \mathscr{F}(st(\sigma)) \subseteq_{fin} T_{\tau^*}(\Phi)$ and $C = h^*(\Gamma)$.
- Towards building a finite falsifying dynamic algebra, start with $(\mathbf{A}_C, f^{R^s_{\langle a\rangle}})_{a\in\Pi_0}$: the algebraic filtration of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ with $(R^s_{\langle a\rangle})_{a\in\Pi_0}$ (strictly speaking we are taking the algebraic filtration of the τ_0 reduct of $\mathbf{A}$).
- Inductively define a set of operators $(\tilde{f}_{\langle\pi\rangle})_{\pi\in\Pi}$ on $\mathbf{A}_C$ starting with $\tilde{f}_{\langle a\rangle} = f^{R^s_{\langle a\rangle}}$ for all $a \in \Pi_0$ and thus show $(\mathbf{A}_C, \tilde{f}_{\langle\pi\rangle})_{\pi\in\Pi} \in \mathfrak{D}_{fin}$.
- Define $j : \Phi \to A_C$ by $j(p) = h(p)$. Show by induction that $j^*(\psi) = h^*(\psi)$ for all $\psi \in st(\sigma)$. Thus it will hold that $j^*(\sigma) \neq 1$.

Regarding the induction in the last step: to prove $j^*(\langle\pi\rangle\,\varphi) = h^*(\langle\pi\rangle\,\varphi)$ where $\langle\pi\rangle\,\varphi \in st(\sigma)$, it suffices to show that $\tilde{f}_{\langle\pi\rangle}$ extends $f_{\langle\pi\rangle}\big|_{h^*(\overline{\Gamma}_\pi)}$ i.e. the conditional

$$x \in h^*(\overline{\Gamma}_\pi) \text{ implies } \tilde{f}_{\langle\pi\rangle}(x) = f_{\langle\pi\rangle}(x).$$

This gives $j^*(\langle\pi\rangle\,\varphi) = h^*(\langle\pi\rangle\,\varphi)$ since $j^*(\langle\pi\rangle\,\varphi) = \tilde{f}_{\langle\pi\rangle}(j^*(\varphi))$ and $h^*(\langle\pi\rangle\,\varphi) = f_{\langle\pi\rangle}(h^*(\varphi))$ and by inductive hypothesis $j^*(\varphi) = h^*(\varphi)$; but $h^*(\varphi) \in h^*(\overline{\Gamma}_\pi)$ (since $\varphi, \langle\pi\rangle\,\varphi \in \Gamma$) and thus by the conditional $\tilde{f}_{\langle\pi\rangle}(h^*(\varphi)) = f_{\langle\pi\rangle}(h^*(\varphi))$. Therefore as part of the last step in the proof strategy we shall prove $x \in h^*(\overline{\Gamma}_\pi)$ implies $\tilde{f}_{\langle\pi\rangle}(x) = f_{\langle\pi\rangle}(x)$ for all $\pi \in \Pi$ using induction on the length of programs.

It seems natural in light of Chapter 4 to rather prove (by induction on the length of programs) the stronger statement[2] that $\tilde{f}_{\langle\pi\rangle}$ extends $f_{\langle\pi\rangle}\big|_{\underline{C}_\pi}$ for all $\pi \in \Pi$ (and indeed this would imply $\tilde{f}_{\langle\pi\rangle}(h^*(\varphi)) = f_{\langle\pi\rangle}(h^*(\varphi))$). This will work in the base case (i.e. with basic programs) but will not work for complex programs as $\underline{C}_\pi$ is not necessarily sensitive to program constructs. For

[2]Stronger since $h^*(\overline{\Gamma}_\pi) \subseteq \underline{C}_\pi$ by Remark 5.1.13.

example suppose $x \in \underline{C}_{a;b}$ and we want to prove $\tilde{f}_{\langle a;b\rangle}(x) = f_{\langle a;b\rangle}(x)$. In particular we want to show $\tilde{f}_{\langle a\rangle}\tilde{f}_{\langle b\rangle}(x) = f_{\langle a\rangle}f_{\langle b\rangle}(x)$ (since $\mathbf{A}$ and $\mathbf{A}_C$ are dynamic algebras (will be proven)). However $x \in \underline{C}_{a;b}$ does not necessarily imply $x \in \underline{C}_b$ or $f_{\langle b\rangle}(x) \in \underline{C}_a$ (since h^* is not injective) and thus we will not be able to conclude $\tilde{f}_{\langle b\rangle}(x) = f_{\langle b\rangle}(x)$ or $\tilde{f}_{\langle a\rangle}\tilde{f}_{\langle b\rangle}(x) = f_{\langle a\rangle}f_{\langle b\rangle}(x)$ respectively. The following lemma demonstrates that $h^*(\overline{\Gamma}_\pi)$ is sensitive to program constructs in this way.

5.1.14 Lemma. Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi}$ be a τ^* modal algebra, let $g : T_{\tau^*}(\Phi) \to A$ be a homomorphism, let Γ be a $\mathscr{F}$-closed subset of $T_{\tau^*}(\Phi)$ and let $\pi \in \Pi \setminus \Pi_0$. Then

$$x \in g(\overline{\Gamma}_\pi) \text{ implies } \begin{cases} x \in g(\overline{\Gamma}_{\pi_2}) \text{ and } f_{\langle\pi_2\rangle}(x) \in g(\overline{\Gamma}_{\pi_1}) & \text{if } \pi = \pi_1;\pi_2 \\ x \in g(\overline{\Gamma}_{\pi_1}) \text{ and } x \in g(\overline{\Gamma}_{\pi_2}) & \text{if } \pi = \pi_1 \cup \pi_2 \\ f_{\langle\pi_1^*\rangle}(x) \in g(\overline{\Gamma}_{\pi_1}) & \text{if } \pi = \pi_1^*. \end{cases}$$

Proof. Consider the first case. Let $x \in g(\overline{\Gamma}_{\pi_1;\pi_2})$, in particular, $x = g(\psi)$ for some $\psi \in \Gamma$ with $\langle\pi_1;\pi_2\rangle\,\psi \in \Gamma$. As Γ is $\mathscr{F}$-closed $\langle\pi_1\rangle\langle\pi_2\rangle\,\psi \in \Gamma$ and $\langle\pi_2\rangle\,\psi \in \Gamma$; the latter giving that $\psi \in \overline{\Gamma}_{\pi_2}$ and thus $x \in g(\overline{\Gamma}_{\pi_2})$. As g is a homomorphism we have $f_{\langle\pi_2\rangle}(x) = g(\langle\pi_2\rangle\,\psi)$; furthermore since $\langle\pi_2\rangle\,\psi \in \Gamma$ and $\langle\pi_1\rangle\langle\pi_2\rangle\,\psi \in \Gamma$ we get $\langle\pi_2\rangle\,\psi \in \overline{\Gamma}_{\pi_1}$ and therefore $g(\langle\pi_2\rangle\,\psi) \in g(\overline{\Gamma}_{\pi_1})$.
Consider the second case. Let $x \in g(\overline{\Gamma}_{\pi_1\cup\pi_2})$, thus $x = g(\psi)$ for some $\psi \in \Gamma$ with $\langle\pi_1 \cup \pi_2\rangle\,\psi \in \Gamma$. As Γ is $\mathscr{F}$-closed $\langle\pi_1\rangle\,\psi \in \Gamma$ and $\langle\pi_2\rangle\,\psi \in \Gamma$. Thus $\psi \in \overline{\Gamma}_{\pi_1}$ and $\psi \in \overline{\Gamma}_{\pi_2}$ which gives that $x \in g(\overline{\Gamma}_{\pi_1})$ and $x \in g(\overline{\Gamma}_{\pi_2})$.
For the final case let $x \in g(\overline{\Gamma}_{\pi_1^*})$. Therefore $x = g(\psi)$ for some $\psi \in \Gamma$ such that $\langle\pi_1^*\rangle\,\psi \in \Gamma$; therefore $\langle\pi_1\rangle\langle\pi_1^*\rangle\,\psi \in \Gamma$ since Γ is $\mathscr{F}$-closed. As g is a homomorphism $f_{\langle\pi_1^*\rangle}(x) = g(\langle\pi_1^*\rangle\,\psi)$; since $\langle\pi_1\rangle\langle\pi_1^*\rangle\,\psi \in \Gamma$ we know that $\langle\pi_1^*\rangle\,\psi \in \overline{\Gamma}_{\pi_1}$ and therefore $g(\langle\pi_1^*\rangle\,\psi) \in g(\overline{\Gamma}_{\pi_1})$. □

5.1.15 Theorem. *Let $\sigma \in T_{\tau^*}(\Phi)$. Then $\mathfrak{D} \not\models \sigma^\approx$ implies $\mathfrak{D}_{fin} \not\models \sigma^\approx$.*

Proof. Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle\pi\rangle})_{\pi\in\Pi} \in \mathfrak{D}$ with $\mathbf{A} \not\models \sigma^\approx$ for $\sigma \in T_{\tau^*}(\Phi)$ i.e. there is an $h : \Phi \to A$ such that $h^*(\sigma) \neq 1$. Define $\Gamma = \mathscr{F}(st(\sigma))$ and $C = h^*(\Gamma)$. The set C is finite as Γ is finite by Theorem 5.1.3. Consider $(\mathbf{A}_C, f^{R^s_{\langle a\rangle}})_{a\in\Pi_0}$: the algebraic filtration of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ with $(R^s_{\langle a\rangle})_{a\in\Pi_0}$ (recall Example 4.2.9 for the definition of $R^s_{\langle a\rangle}$). Towards forming a dynamic algebra, recursively define the functions $(\tilde{f}_{\langle\pi\rangle})_{\pi\in\Pi}$ on $\mathbf{A}_C$ as:

$$\tilde{f}_{\langle\pi\rangle}(x) = \begin{cases} f^{R^s_{\langle a\rangle}}(x) & \pi = a \in \Pi_0 \\ \tilde{f}_{\langle\pi_1\rangle}(\tilde{f}_{\langle\pi_2\rangle}(x)) & \pi = \pi_1;\pi_2 \\ \tilde{f}_{\langle\pi_1\rangle}(x) \vee \tilde{f}_{\langle\pi_2\rangle}(x) & \pi = \pi_1 \cup \pi_2 \\ \displaystyle\bigvee_{n=0}^{|A_C|} \tilde{f}^n_{\langle\pi_1\rangle}(x) & \pi = \pi_1^* \end{cases}$$

where $x \in A_C$ and $\tilde{f}^0_{\langle \pi_1 \rangle}(x) = x$. For $a \in \Pi_0$ the function $\tilde{f}_{\langle a \rangle} = f^{R^s_{\langle a \rangle}}$ is an operator on $\mathbf{A}_C$. This forms the base case of a straightforward inductive proof that the functions $(\tilde{f}_\pi)_{\pi \in \Pi}$ are operators on $\mathbf{A}_C$. It still needs to be checked that $(\mathbf{A}_C, \tilde{f}_{\langle \pi \rangle})_{\pi \in \Pi} \in \mathfrak{D}$. Let $x \in A_C$ and let $\pi, \pi_1, \pi_2 \in \Pi$; from their definitions it is immediate that $\tilde{f}_{\langle \pi_1 ; \pi_2 \rangle}(x)$ and $\tilde{f}_{\langle \pi_1 \cup \pi_2 \rangle}(x)$ obey properties (i) and (ii) of dynamic algebras. We will now check that $\tilde{f}_{\langle \pi^* \rangle}(x)$ has properties (iii) and (iv) of dynamic algebras. To see that $\tilde{f}_{\langle \pi^* \rangle}(x)$ obeys property (iii), it needs to be demonstrated that

$$\tilde{f}_{\langle \pi^* \rangle}(x) = x \vee \tilde{f}_{\langle \pi \rangle}(\tilde{f}_{\langle \pi^* \rangle}(x)).$$

Since $\tilde{f}_{\langle \pi \rangle}$ preserves joins we have

$$\begin{aligned} x \vee \tilde{f}_{\langle \pi \rangle}(\tilde{f}_{\langle \pi^* \rangle}(x)) &= x \vee \tilde{f}_{\langle \pi \rangle}(\bigvee_{n=0}^{|A_C|} \tilde{f}^n_{\langle \pi \rangle}(x)) \\ &= x \vee \bigvee_{n=1}^{|A_C|+1} \tilde{f}^n_{\langle \pi \rangle}(x) \\ &= \bigvee_{n=0}^{|A_C|+1} \tilde{f}^n_{\langle \pi \rangle}(x). \end{aligned}$$

We have $\tilde{f}^{|A_C|+1}_{\langle \pi \rangle}(x) \in \{\tilde{f}^n_{\langle \pi \rangle}(x)\}^{|A_C|}_{n=0}$. This is true because if $\tilde{f}^{|A_C|+1}_{\langle \pi \rangle}(x) \notin \{\tilde{f}^n_{\langle \pi \rangle}(x)\}^{|A_C|}_{n=0}$ then there would be no repetitions amongst $\{\tilde{f}^n_{\langle \pi \rangle}(x)\}^{|A_C|}_{n=0}$ and the algebra $\mathbf{A}_C$ would have $|A_C| + 1$ elements, a contradiction. Therefore

$$\bigvee_{n=0}^{|A_C|+1} \tilde{f}^n_{\langle \pi \rangle}(x) = \bigvee_{n=0}^{|A_C|} \tilde{f}^n_{\langle \pi \rangle}(x) = \tilde{f}_{\langle \pi^* \rangle}(x).$$

To check that $\tilde{f}_{\langle \pi^* \rangle}(x)$ obeys property (iv) of dynamic algebras it suffices to check that $\tilde{f}_{\langle \pi^* \rangle}(x)$ is the least pre-fixed point of $t^\pi_{x,\mathbf{A}_C}$ (by Lemma 5.1.8). We know $\tilde{f}_{\langle \pi^* \rangle}(x)$ is a pre-fixed point of $t^\pi_{x,\mathbf{A}_C}$ since it has just been shown that $\tilde{f}_{\langle \pi^* \rangle}(x)$ is a fixed point of $t^\pi_{x,\mathbf{A}_C}$. Suppose $z \in A_C$ is a pre-fixed point of $t^\pi_{x,\mathbf{A}_C}$ i.e. $x \vee \tilde{f}_{\langle \pi \rangle}(z) \leq z$. Since $\tilde{f}_{\langle \pi \rangle}$ is order-preserving, we get

$$\tilde{f}_{\langle \pi \rangle}(x) \vee \tilde{f}_{\langle \pi \rangle}(\tilde{f}_{\langle \pi \rangle}(z)) \leq \tilde{f}_{\langle \pi \rangle}(z)$$

and since joins preserve order

$$x \vee \tilde{f}_{\langle \pi \rangle}(x) \vee \tilde{f}_{\langle \pi \rangle}(\tilde{f}_{\langle \pi \rangle}(z)) \leq x \vee \tilde{f}_{\langle \pi \rangle}(z) \leq z$$

where the last inequality follows since z is a pre-fixed point of $t^\pi_{x,\mathbf{A}_C}$. Inductively it follows

$$x \vee \tilde{f}_{\langle\pi\rangle}(x) \vee ... \vee \tilde{f}_{\langle\pi\rangle}^{|A_C|}(x) \vee \tilde{f}_{\langle\pi\rangle}^{|A_C|+1}(z) = \tilde{f}_{\langle\pi^*\rangle}(x) \vee \tilde{f}_{\langle\pi\rangle}^{|A_C|+1}(z) \leq z$$

which gives $\tilde{f}_{\langle\pi^*\rangle}(x) \leq z$ and therefore $\tilde{f}_{\langle\pi^*\rangle}(x)$ is the least pre-fixed point of $t^{\pi}_{x,\mathbf{A}_C}$ i.e. $\tilde{f}_{\langle\pi^*\rangle}(x)$ obeys property (iv) of dynamic algebras by Lemma 5.1.8. Thus we have that $(\mathbf{A}_C, \tilde{f}_{\langle\pi\rangle})_{\pi\in\Pi} \in \mathfrak{D}_{fin}$.

Finally it needs to be checked that $(\mathbf{A}_C, \tilde{f}_{\langle\pi\rangle})_{\pi\in\Pi}$ falsifies $\sigma^{\approx}$. To this end define $j : \Phi \to A_C$ by $j(p) = h(p)$; to prove that $j^*(\sigma) \neq 1$ it suffices to show $j^*(\psi) = h^*(\psi)$ for all $\psi \in st(\sigma)$. However, since $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0}$ is an algebraic filtration, we already have $j^*(\psi) = h^*(\psi)$ for all $\psi \in st(\sigma) \cap T_{\tau_0}(\Phi)$ (by Theorem 4.2.11). We therefore only prove that for all $\langle\pi\rangle\,\varphi \in st(\sigma)$:

$$j^*(\langle\pi\rangle\,\varphi) = h^*(\langle\pi\rangle\,\varphi). \tag{5.1.3}$$

As discussed in the introduction to this subsection, if it can be shown that for all $\pi \in \Pi$ that

$$x \in h^*(\overline{\Gamma}_\pi) \text{ implies } \tilde{f}_{\langle\pi\rangle}(x) = f_{\langle\pi\rangle}(x) \tag{5.1.4}$$

then equation 5.1.3 holds for all $\langle\pi\rangle\,\varphi \in st(\sigma)$. Thus we will now prove conditional 5.1.4 for all $\pi \in \Pi$.

Base case: Let $\mathrm{k}(\pi) = 0$ i.e. $\pi = a \in \Pi_0$ and let $x \in h^*(\overline{\Gamma}_a)$. We know by Remark 5.1.13 that $h^*(\overline{\Gamma}_a) \subseteq \underline{C}_a$ and hence $x \in \underline{C}_a$. Since $f^{R^s_{\langle a\rangle}}$ extends $f_{\langle a\rangle}\big|_{\underline{C}_a}$ we have $f^{R^s_{\langle a\rangle}}(x) = f_{\langle a\rangle}(x)$ i.e. $\tilde{f}_{\langle a\rangle}(x) = f_{\langle a\rangle}(x)$.

Inductive hypothesis: Assume conditional 5.1.4 is true for all $\pi \in \Pi$ such that $\mathrm{k}(\pi) < m$.

Inductive step: Let $\pi \in \Pi$ such that $\mathrm{k}(\pi) = m$ and let $x \in h^*(\overline{\Gamma}_\pi)$. Therefore for some $\pi_1, \pi_2 \in \Pi$ we have either:

$$\pi = \begin{cases} \pi_1;\pi_2 & (I) \\ \pi_1 \cup \pi_2 & (II) \\ \pi_1^* & (III). \end{cases}$$

(I) Since $x \in h^*(\overline{\Gamma}_{\pi_1;\pi_2})$ we have by Lemma 5.1.14 that $x \in h^*(\overline{\Gamma}_{\pi_2})$ and thus by the inductive hypothesis

$$\tilde{f}_{\langle\pi_2\rangle}(x) = f_{\langle\pi_2\rangle}(x).$$

Therefore $\tilde{f}_{\langle\pi_1\rangle}(\tilde{f}_{\langle\pi_2\rangle}(x)) = \tilde{f}_{\langle\pi_1\rangle}(f_{\langle\pi_2\rangle}(x))$. By Lemma 5.1.14 we have $f_{\langle\pi_2\rangle}(x) \in h^*(\overline{\Gamma}_{\pi_1})$ which gives by the inductive hypothesis

$$\tilde{f}_{\langle\pi_1\rangle}(f_{\langle\pi_2\rangle}(x)) = f_{\langle\pi_1\rangle}(f_{\langle\pi_2\rangle}(x)).$$

(II) Since $x \in h^*(\overline{\Gamma}_{\pi_1\cup\pi_2})$ we have by Lemma 5.1.14 that $x \in h^*(\overline{\Gamma}_{\pi_1}) \cap h^*(\overline{\Gamma}_{\pi_2})$ and thus by the inductive hypothesis

$$\tilde{f}_{\langle\pi_1\rangle}(x) = f_{\langle\pi_1\rangle}(x)$$

and

$$\tilde{f}_{\langle \pi_2 \rangle}(x) = f_{\langle \pi_2 \rangle}(x)$$

which gives

$$\tilde{f}_{\langle \pi_1 \rangle}(x) \vee \tilde{f}_{\langle \pi_2 \rangle}(x) = f_{\langle \pi_1 \rangle}(x) \vee f_{\langle \pi_2 \rangle}(x).$$

(III) Let $x \in h^*(\overline{\Gamma}_{\pi_1^*})$. Firstly we will show

$$\tilde{f}_{\langle \pi_1^* \rangle}(x) \leq f_{\langle \pi_1^* \rangle}(x).$$

Since $\tilde{f}_{\langle \pi_1^* \rangle}(x)$ is the least pre-fixed point of the function $t^{\pi}_{x,\mathbf{A}_C}$ (by Lemma 5.1.7) it will suffice to show that $f_{\langle \pi_1^* \rangle}(x)$ is a pre-fixed point of $t^{\pi}_{x,\mathbf{A}_C}$ i.e.

$$x \vee \tilde{f}_{\langle \pi \rangle}(f_{\langle \pi_1^* \rangle}(x)) \leq f_{\langle \pi_1^* \rangle}(x).$$

Since $x \in h^*(\overline{\Gamma}_{\pi_1^*})$ we have by Lemma 5.1.14 that $f_{\langle \pi_1^* \rangle}(x) \in h^*(\overline{\Gamma}_{\pi_1})$ and thus by the inductive hypothesis

$$\tilde{f}_{\langle \pi_1 \rangle}(f_{\langle \pi_1^* \rangle}(x)) = f_{\langle \pi_1 \rangle}(f_{\langle \pi_1^* \rangle}(x)).$$

Therefore

$$\begin{aligned} x \vee \tilde{f}_{\langle \pi_1 \rangle}(f_{\langle \pi_1^* \rangle}(x)) &= x \vee f_{\langle \pi_1 \rangle}(f_{\langle \pi_1^* \rangle}(x)) \\ &= f_{\langle \pi_1^* \rangle}(x) \end{aligned}$$

where the second equality follows since $\mathbf{A}$ is a dynamic algebra. Thus $f_{\langle \pi_1^* \rangle}(x)$ is a pre-fixed point of $t^{\pi}_{x,\mathbf{A}_C}$ and $\tilde{f}_{\langle \pi_1^* \rangle}(x) \leq f_{\langle \pi_1^* \rangle}(x)$. To complete the proof that $\tilde{f}_{\langle \pi_1^* \rangle}(x) = f_{\langle \pi_1^* \rangle}(x)$ it remains to show that $f_{\langle \pi_1^* \rangle}(x) \leq \tilde{f}_{\langle \pi_1^* \rangle}(x)$. We thus finish by showing for all $\alpha \in \Pi$ that

$$f_{\langle \alpha \rangle}(x) \leq \tilde{f}_{\langle \alpha \rangle}(x) \text{ for all } x \in A_C. \tag{5.1.5}$$

Base case: Let $\mathrm{k}(\alpha) = 0$ i.e. $\alpha = a \in \Pi_0$ and let $x \in A_C$. By Lemma 4.2.13 we have that

$$f^{R^s_{\langle a \rangle}}(x) = \bigwedge\{x' \in A_C \mid f_{\langle a \rangle}(x) \leq x'\}$$

and thus $f_{\langle a \rangle}(x) \leq f^{R^s_{\langle a \rangle}}(x) = \tilde{f}_{\langle a \rangle}(x)$.

Inductive hypothesis: Assume statement 5.1.5 is true for all $\alpha \in \Pi$ such that $\mathrm{k}(\alpha) < m$.

Inductive step: Let $\alpha \in \Pi$ such that $\mathrm{k}(\alpha) = m$ and let $x \in A_C$. Therefore for some $\alpha_1, \alpha_2 \in \Pi$:

$$\alpha = \begin{cases} \alpha_1; \alpha_2 & (i) \\ \alpha_1 \cup \alpha_2 & (ii) \\ \alpha_1^* & (iii). \end{cases}$$

(i) By inductive hypothesis we have

$$f_{\langle\alpha_1\rangle}(f_{\langle\alpha_2\rangle}(x)) \leq \tilde{f}_{\langle\alpha_1\rangle}(f_{\langle\alpha_2\rangle}(x))$$

and

$$f_{\langle\alpha_2\rangle}(x) \leq \tilde{f}_{\langle\alpha_2\rangle}(x).$$

Since $\tilde{f}_{\langle\alpha_1\rangle}$ is order-preserving, applying it on the last inequality gives

$$\tilde{f}_{\langle\alpha_1\rangle}(f_{\langle\alpha_2\rangle}(x)) \leq \tilde{f}_{\langle\alpha_1\rangle}(\tilde{f}_{\langle\alpha_2\rangle}(x)).$$

Thus overall

$$f_{\langle\alpha_1;\alpha_2\rangle}(x) = f_{\langle\alpha_1\rangle}(f_{\langle\alpha_2\rangle}(x)) \leq \tilde{f}_{\langle\alpha_1\rangle}(\tilde{f}_{\langle\alpha_2\rangle}(x)) = \tilde{f}_{\langle\alpha_1;\alpha_2\rangle}(x).$$

(ii) By inductive hypothesis we have

$$f_{\langle\alpha_1\rangle}(x) \leq \tilde{f}_{\langle\alpha_1\rangle}(x)$$

and

$$f_{\langle\alpha_2\rangle}(x) \leq \tilde{f}_{\langle\alpha_2\rangle}(x)$$

thus giving

$$f_{\langle\alpha_1\cup\alpha_2\rangle}(x) = f_{\langle\alpha_1\rangle}(x) \vee f_{\langle\alpha_2\rangle}(x) \leq \tilde{f}_{\langle\alpha_1\rangle}(x) \vee \tilde{f}_{\langle\alpha_2\rangle}(x) = \tilde{f}_{\langle\alpha_1\cup\alpha_2\rangle}(x).$$

(iii) To prove $f_{\langle\alpha^*\rangle}(x) \leq \tilde{f}_{\langle\alpha^*\rangle}(x)$ it suffices to show that $\tilde{f}_{\langle\alpha^*\rangle}(x)$ is a pre-fixed point of $t^{\alpha}_{x,\mathbf{A}}$ (by Lemma 5.1.7) i.e.

$$x \vee f_{\langle\alpha\rangle}(\tilde{f}_{\langle\alpha^*\rangle}(x)) \leq \tilde{f}_{\langle\alpha^*\rangle}(x).$$

By hypothesis

$$f_{\langle\alpha\rangle}(\tilde{f}_{\langle\alpha^*\rangle}(x)) \leq \tilde{f}_{\langle\alpha\rangle}(\tilde{f}_{\langle\alpha^*\rangle}(x))$$

but since joins preserve order and $\tilde{f}_{\langle\alpha^*\rangle}(x)$ is a fixed point of $t^{\alpha}_{x,\mathbf{A}_C}$ we have

$$x \vee f_{\langle\alpha\rangle}(\tilde{f}_{\langle\alpha^*\rangle}(x)) \leq x \vee \tilde{f}_{\langle\alpha\rangle}(\tilde{f}_{\langle\alpha^*\rangle}(x)) = \tilde{f}_{\langle\alpha^*\rangle}(x).$$

□

5.2 $\Lambda_{\mathfrak{R}}$ has the fmp with respect to $\mathfrak{R}$

We lastly turn our attention to L_ρ: a modal language with least binders introduced by Pratt [34]. The minimization binder, ρ, of L_ρ is interpreted not as a least fixed point but as a least root of a function. This allows L_ρ to express all the PDL formulae as well as formulae involving the converse modality $\langle\pi^-\rangle$, where $\langle\pi^-\rangle$ is interpreted relationally by means of $R_{\langle\pi^-\rangle}w_1w_2$ iff $R_{\langle\pi\rangle}w_2w_1$.

In [34] Pratt posed the open problem of whether it is possible to express every L_ρ formula as an L_μ formula. We now know by the work of Janin and Walukiewicz [20] that this is not the case. In particular Janin and Walukiewicz demonstrated that L_μ is bisimulation invariant; since L_ρ can express the converse modality, which can easily be shown *not* to be bisimulation invariant, we know that L_ρ cannot be a fragment of L_μ in this way.

In this final section we will first look at the syntax and algebraic semantics of L_ρ (Definitions 5.2.1 and 5.2.10 respectively). We will then finally discuss Pratt's result (Theorem 3 of [34]) that $\Lambda_{\mathfrak{R}}$ has the finite model property with respect to $\mathfrak{R}$ (Theorem 5.2.22). In particular, we will offer a proof based on that of Pratt's but differing in the construction of the finite falsifying algebra.

5.2.1 Definition. [34] The set Var denotes a set of *variables* disjoint from Φ. The formulae of L_ρ are defined recursively as $\varphi = \bot \mid p \mid X \mid \neg\varphi_1 \mid \varphi_1 \vee \varphi_2 \mid \langle a\rangle\, \varphi_1 \mid \rho X.\varphi(X)$ where $p \in \Phi$, $X \in Var$, $a \in \Pi_0$ and $\varphi(X)$ obeys the following two syntactic restrictions: X does not occur conjunctively in $\varphi(X)$ (Definition 5.2.2) and $\varphi(\top)$ evaluates syntactically to $\bot$ (see Definition 5.2.4). The *dual of* $\rho X.\varphi(X)$ is the formula $\neg\rho X.\neg\varphi(X)$ and is denoted by $\iota X.\varphi(X)$. The collection of L_ρ formulae is denoted by $L_\rho(\Phi, Var)$.

5.2.2 Definition. [34] Let $X, \varphi(X) \in L_\rho(\Phi, Var)$. The formula X is said to *occur conjunctively in* $\varphi(X)$ if any of the following hold:

- (i) likesigned occurrences of X in $\varphi(X)$ occur in different arguments of a conjunction;
- (ii) $[a]\psi(X)$ occurs in $\varphi(X)$;
- (iii) $\iota X.\psi(X)$ occurs in $\varphi(X)$

where $a \in \Pi_0$ and $\psi(X) \in L_\rho(\Phi, Var)$ such that X occurs in $\psi(X)$.

5.2.3 Example. The formula $\rho X.(\neg X \wedge \langle a\rangle\, X)$ is a L_ρ formula whilst $\rho X.(\neg X \wedge [a]X)$ is not as X occurs conjunctively in $\neg X \wedge [a]X$ by condition (ii).

The following definition is based on a definition found in [34].

5.2.4 Definition. Define the binary relation $\rightsquigarrow$ on $L_\rho(\Phi, Var)$ by

$$\begin{array}{rcl}
\neg\top, \psi \wedge \bot, \bot \wedge \psi, \neg\psi \wedge \psi, \bot \vee \bot, \langle a\rangle\, \bot & \rightsquigarrow & \bot \\
\psi, \psi \vee \bot, \bot \vee \psi, \psi \wedge \top, \top \wedge \psi & \rightsquigarrow & \psi \\
\neg\bot, \psi \vee \top, \top \vee \psi, \neg\psi \vee \psi, \top \vee \top & \rightsquigarrow & \top \\
\varphi_1 \rightsquigarrow \psi_1 & \text{implies} & \neg\varphi_1 \rightsquigarrow \neg\psi_1 \\
\varphi_1 \rightsquigarrow \varphi_2 \;\&\; \psi_1 \rightsquigarrow \psi_2 & \text{implies} & \varphi_1 \vee \varphi_2 \rightsquigarrow \psi_1 \vee \psi_2 \\
\varphi_1 \rightsquigarrow \psi_1 & \text{implies} & \langle a\rangle\, \varphi_1 \rightsquigarrow \langle a\rangle\, \psi_1 \\
\psi(\bot) \rightsquigarrow \bot & \text{implies} & \rho X.\psi(X) \;\rightsquigarrow\; \bot
\end{array}$$

where $\psi, \psi_1, \psi_2, \varphi_1, \varphi_2 \in L_\rho(\Phi, Var)$ and $a \in \Pi_0$. The formula φ *evaluates syntactically to* $\bot$ if $\varphi \rightsquigarrow^* \bot$ (where $\rightsquigarrow^*$ is the transitive closure of $\rightsquigarrow$).

5.2.5 Remark. It can be shown that this relation is defined in such a way that $\varphi \rightsquigarrow^* \psi$ implies $\mathbf{A} \models \varphi \approx \psi$ for any τ_0 modal algebra $\mathbf{A}$.

Given a L_ρ formula $\rho X.\varphi(X)$, the restriction that $\varphi(\top) \rightsquigarrow^* \bot$ guarantees that 1 will be a root of $\varphi(X)$ whenever it is interpreted on a τ_0 modal algebra i.e. the set of roots of $\varphi(X)$ is not empty. The fact that X cannot occur conjunctively in $\varphi(X)$ makes L_ρ amenable to a filtration argument, in part by ensuring that the set of roots of $\varphi(X)$ is a lattice (see Lemma 5.2.13).

5.2.6 Definition. The *bounded variables of a L_ρ formula* (denoted b.v.()) are defined inductively as b.v.($\bot$)=b.v.(p)=b.v.(X)=$\varnothing$, b.v.($\neg\varphi$)=b.v.($\langle a\rangle \varphi$)=b.v.(φ), b.v.($\varphi_1 \vee \varphi_2$)=b.v.(φ_1) $\cup$ b.v.(φ_2) and b.v.($\rho X.\varphi(X)$)=b.v.($\varphi(X)$) $\cup \{X\}$ where $p \in \Phi$, $X \in Var$, $a \in \Pi_0$ and $\varphi, \varphi_1, \varphi_2, \rho X.\varphi(X) \in L_\rho(\Phi, Var)$. The *free variables of a L_ρ formula* (f.v.()) are defined analogously except f.v.(X)=$\{X\}$ and f.v.($\rho X.\varphi(X)$)=f.v.($\varphi(X)$) $\setminus \{X\}$.

5.2.7 Example. For $\sigma = p \vee \rho Y.(\langle a\rangle \neg Y \wedge X)$, we have b.v.($\sigma$) $= \{Y\}$ and f.v.(σ) $= \{X\}$.

5.2.8 Definition. Let $\sigma \in L_\rho(\Phi, Var)$. The formula σ is in *normal form* if all of the negations in σ appear in front of variables i.e. σ is in normal form if $\neg\psi \in st(\sigma)$ implies $\psi \in \Phi \cup Var$ for all $\psi \in st(\sigma)$.

5.2.9 Remark. Any L_ρ formula is equivalent to a formula in normal form. For example the L_ρ formula $\rho X.(\neg X \wedge \neg \langle a\rangle X)$ is equivalent to the L_ρ formula $\rho X.(\neg X \wedge [a]\neg X)$. We thus assume all L_ρ formulae to be in normal form.

We now look at the algebraic semantics for L_ρ (recall Definition 3.2.7 of an X-variant of a map) as well as formally define the notion of a root.

5.2.10 Definition. [34] Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle a\rangle})_{a\in\Pi_0}$ be a τ_0 modal algebra and let $h : \Phi \cup Var \to A$. The map $h^* : L_\rho(\Phi, Var) \to A$ is the *extension of h to $L_\rho(\Phi, Var)$* and is defined as

$$h^*(\varphi) = \begin{cases} 0 & \varphi = \bot \\ h(p) & \varphi = p \in \Phi \cup Var \\ \neg h^*(\varphi_1) & \varphi = \neg\varphi_1 \\ h^*(\varphi_1) \vee h^*(\varphi_2) & \varphi = \varphi_1 \vee \varphi_2 \\ f_{\langle a\rangle}(h^*(\varphi_1)) & \varphi = \langle a\rangle \varphi_1 \\ \text{smallest } x \in A \text{ s.t. } h^*_{X\to x}(\varphi(X)) = 0 & \varphi = \rho X.\varphi(X) \end{cases}$$

where $h^*_{X\to x}$ is the extension of $h_{X\to x}$ to $L_\rho(\Phi, Var)$. An element $x \in A$ is a *root of* $\varphi(X)$ *in* A *under* h if $h^*_{X\to x}(\varphi(X)) = 0$; as such $h^*(\rho X.\varphi(X))$ is the smallest root of $\varphi(X)$ in A under h.

5.2.11 Remark. By the restriction on $\varphi(X)$ that $\varphi(\top) \rightsquigarrow^* \bot$ we know that $\mathbf{A} \models \varphi(\top) \approx \bot$ i.e. $h^*_{X\to 1}(\varphi(X)) = 0$. However h^* may still be a partial map since $h^*(\rho X.\varphi(X))$ may fail to exist i.e. the map $h^*_{X\to}(\varphi(X))$ may not have a *smallest* root.

5.2.12 Definition. [34] Let $\mathbf{A}$ be a τ_0 modal algebra. $\mathbf{A}$ is a ρ-*algebra* when $h^* : L_\rho(\Phi, Var) \to A$ is a total mapping for all $h : \Phi \cup Var \to A$ i.e. $h^*(\rho X.\varphi(X))$ is well defined for all $h : \Phi \cup Var \to A$ and $\rho X.\varphi(X) \in L_\rho(\Phi, Var)$. The class of ρ-algebras will be denoted by $\mathfrak{R}$.

In the next lemma we take for granted the non-trivial result of Pratt that for a τ_0 modal algebra $\mathbf{A}$, a mapping $h : \Phi\cup Var \to A$ and a formula $\rho X.\varphi(X) \in L_\rho(\Phi, Var)$ the set $\{x \in A \mid h^*_{X\to x}(\varphi(X)) = 0\}$ is a lattice [34]. The result relies heavily on the fact that X does not occur conjunctively in $\varphi(X)$ and hence that likesigned occurrences of X in $\varphi(X)$ do not occur within distinct arguments of a conjunction.

5.2.13 Lemma. [34] Let $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle a\rangle})_{a\in\Pi_0}$ be a finite τ_0 modal algebra. Then $\mathbf{A} \in \mathfrak{R}$.

Proof. Let $\rho X.\varphi(X) \in L_\rho(\Phi, Var)$. For any τ_0 modal algebra $\mathbf{A}$ and $h : \Phi \cup Var \to A$, the set $C = \{x \in A \mid h^*_{X\to x}(\varphi(X)) = 0\}$ is a lattice (see appendix of [34]). When $\mathbf{A}$ is finite the element $\bigwedge C$ will exist in $\mathbf{A}$ and will be the least root of $h^*_{X\to}(\varphi(X))$ i.e. $h^*(\rho X.\varphi(X))$ will be well defined making $\mathbf{A}$ a ρ-algebra. □

5.2.1 The validities of $\mathfrak{R}$ and $\mathfrak{R}_{fin}$ are the same.

In this final subsection we will show that $\Lambda_{\mathfrak{R}}$, the equational theory of $\mathfrak{R}$, has the finite model property with respect to $\mathfrak{R}$. By definition $\Lambda_{\mathfrak{R}}$ is characterised by $\mathfrak{R}$. Thus to prove that $\Lambda_{\mathfrak{R}}$ has the finite model property with respect to $\mathfrak{R}$, we will demonstrate that if a L_ρ formula σ is a validity on $\mathfrak{R}_{fin}$ then σ is a validity on $\mathfrak{R}$. This will be done by proving the contrapositive: if there a ρ-algebra $\mathbf{A}$ such that $\mathbf{A} \not\models \sigma^{\approx}$ then there is a finite ρ-algebra $\mathbf{A}'$ such that $\mathbf{A}' \not\models \sigma^{\approx}$. Pratt gives a proof of this in Theorem 3 of [34]: suppose $\mathbf{A}{=}(A, \vee, \wedge, \neg, 0, 1, f_{\langle a\rangle})_{a\in\Pi_0} \in \mathfrak{R}$ such that $\mathbf{A} \not\models \sigma^{\approx}$; towards building a finite falsifying ρ-algebra, Pratt takes the Boolean algebra $\mathbf{A}_C$ generated by $C = \{h^*(\psi) \mid \psi \in st(\sigma)\} \subseteq A$ where $h : \Phi \cup Var \to A$ with $h^*(\sigma) \neq 1$. For each $a \in \Pi_0$ he then defines operators $\tilde{f}_{\langle a\rangle} : A_C \to A_C$ by $\tilde{f}_{\langle a\rangle}(z) = \bigwedge\{x' \in A_C \mid f_{\langle a\rangle}(z) \leq x'\}$ giving rise to the τ_0 modal algebra $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0}$. In particular he takes the algebraic filtration of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ with $(R^s_{\langle a\rangle})_{a\in\Pi_0}$ (by Lemma 4.2.13). This algebra is in $\mathfrak{R}$ by Theorem 5.2.13. Pratt then proceeds to demonstrate that $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0}$ also falsifies $\sigma^{\approx}$ by considering the extension of j to $L_\rho(\Phi, Var)$

where $j : \Phi \cup Var \to A_C$ is defined by $j(p) = h(p)$. In particular Pratt shows $j^*(\psi) = h^*(\psi)$ for all $\psi \in st(\sigma)$ and when $\psi = \rho X.\varphi(X)$, he argues that $j^*(\psi) = h^*(\psi)$ by showing $j^*(\psi)$ is a root of $h^*_{X\to}\varphi(X)$ (hence $h^*(\psi) \leq j^*(\psi)$) and $h^*(\psi)$ is a root of $j^*_{X\to}\varphi(X)$ (hence $h^*(\psi) \geq j^*(\psi)$).

It seems however that there is an alteration required in the construction of the set C in Pratt's proof. To elaborate on this point, we present the following two scenarios:

• Consider $\sigma = \rho X.\varphi(X)$ where $\varphi(X) = p \wedge \langle a \rangle \neg X$. The formula σ is a legal L_ρ formula since X does not occur conjunctively in $\varphi(X)$ and $p \wedge \langle a \rangle \neg \top \leadsto^* p \wedge \langle a \rangle \bot \leadsto^* p \wedge \bot \leadsto^* \bot$. Let $h^*(\sigma) = x_1$; therefore $C = \{x_1, h(p) \wedge f_{\langle a \rangle}(\neg h(X)), h(p), f_{\langle a \rangle}(\neg h(X)), \neg h(X), h(X)\}$. The algebra $\mathbf{A}_C$ will falsify $\sigma^{\approx}$ if $j^*(\sigma) = h^*(\sigma)$.
$\underline{j^*(\sigma) \geq h^*(\sigma)}$ since $j^*(\sigma)$ is a root of $h^*_{X\to}(\varphi(X))$:

$$\begin{aligned} f_{\langle a \rangle}(\neg j^*(\sigma)) \wedge h(p) &= f_{\langle a \rangle}(\neg j^*(\sigma)) \wedge j(p) && \text{(defn. of } j) \\ &\leq \tilde{f}_{\langle a \rangle}(\neg j^*(\sigma)) \wedge j(p) && \text{(defn. of } \tilde{f}) \\ &= 0 && \text{(defn. of } j^*(\sigma)). \end{aligned}$$

$\underline{j^*(\sigma) \leq h^*(\sigma)}$ if x_1 is a root of $j^*_{X\to}(\varphi(X))$:

$$\begin{aligned} \tilde{f}_{\langle a \rangle}(\neg x_1) \wedge j(p) &= \tilde{f}_{\langle a \rangle}(\neg x_1) \wedge h(p) && \text{(defn. of } j) \\ &\stackrel{?}{=} f_{\langle a \rangle}(\neg x_1) \wedge h(p) \\ &= 0 && \text{(defn. of } x_1). \end{aligned}$$

We do not necessarily have $f_{\langle a \rangle}(\neg x_1) \in C$ so it is not guaranteed that $\tilde{f}_{\langle a \rangle}(\neg x_1) = f_{\langle a \rangle}(\neg x_1)$. This is fixed by instead building C as $\{h^*_{X\to x_1}(\psi) \mid \psi \in st(\sigma)\}$.

• The situation becomes more complicated with nested least root formulae. Consider the formula σ where $\sigma{=}\rho X_1.(\langle a \rangle \neg X_1 \wedge \rho X_2.(\langle a \rangle \neg X_2 \wedge X_1))$ and let $\varphi_2(X_2){=}\langle a \rangle \neg X_2 \wedge X_1$ and $\varphi_1(X_1) = \langle a \rangle \neg X_1 \wedge \rho X_2.\varphi_2(X_2)$. It can be checked that σ is indeed a L_ρ formula. Let $x_1 = h^*(\sigma)$ and let $z_1 = j^*(\sigma)$. For now the set C is constructed as $\{h^*_{X_1\to x_1}(\psi) \mid \psi \in st(\sigma)\}$. The algebra $\mathbf{A}_C$ will falsify $\sigma^{\approx}$ if $z_1 = x_1$.
$\underline{z_1 \geq x_1}$ since z_1 is a root of $h^*_{X_1\to}(\varphi_1(X_1))$:

$$\begin{aligned} f_{\langle a \rangle}(\neg z_1) \wedge h^*_{X_1\to z_1}(\rho X_2.\varphi_2(X_2)) &\leq \tilde{f}_{\langle a \rangle}(\neg z_1) \wedge h^*_{X_1\to z_1}(\rho X_2.\varphi_2(X_2)) && \text{(defn. of } \tilde{f}) \\ &\leq \tilde{f}_{\langle a \rangle}(\neg z_1) \wedge j^*_{X_1\to z_1}(\rho X_2.\varphi_2(X_2)) && \text{(inequality 5.2.1)} \\ &= 0 && \text{(defn. of } z_1). \end{aligned}$$

We will show

$$h^*_{X_1\to z}(\rho X_2.\varphi_2(X_2)) \leq j^*_{X_1\to z}(\rho X_2.\varphi_2(X_2)) \text{ for all } z \in A_C. \tag{5.2.1}$$

Let $z \in A_C$ and let $z' = j^*_{X_1 \to z}(\rho X_2.\varphi_2(X_2))$. By definition, inequality 5.2.1 is true if z' is a root of $h^*_{X_1 \to z, X_2 \to}(\varphi_2(X_2))$:

$$\begin{aligned} f_{\langle a \rangle}(\neg z') \wedge z &\leq \tilde{f}_{\langle a \rangle}(\neg z') \wedge z & \text{(defn. of } \tilde{f}) \\ &= 0 & \text{(defn. of } z'). \end{aligned}$$

Therefore inequality 5.2.1 is true which gives $z_1 \geq x_1$.
$\underline{z_1 \leq x_1}$ if x_1 is a root of $j^*_{X_1 \to}(\varphi_1(X_1))$:

$$\begin{aligned} \tilde{f}_{\langle a \rangle}(\neg x_1) \wedge j^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2)) &= f_{\langle a \rangle}(\neg x_1) \wedge j^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2)) & (f_{\langle a \rangle}(\neg x_1) \in C) \\ &\stackrel{?}{=} f_{\langle a \rangle}(\neg x_1) \wedge h^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2)) \\ &= 0 & \text{(defn. of } x_1). \end{aligned}$$

Thus $z_1 \leq x_1$ if it can be demonstrated that $j^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2)) = h^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2))$. Let $x_2 = h^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2))$. We know from inequality 5.2.1 that $x_2 \leq j^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2))$. It remains to show that $j^*_{X_1 \to x_1}(\rho X_2.\varphi_2(X_2)) \leq x_2$. This is true if x_2 is a root of $j^*_{X_1 \to x_1, X_2 \to}(\varphi_2(X_2))$:

$$\begin{aligned} \tilde{f}_{\langle a \rangle}(\neg x_2) \wedge x_1 &\stackrel{?}{=} f_{\langle a \rangle}(\neg x_2) \wedge x_1 \\ &= 0 & \text{(defn. of } x_2). \end{aligned}$$

We are not guaranteed that $\tilde{f}_{\langle a \rangle}(\neg x_2) = f_{\langle a \rangle}(\neg x_2)$ since $f_{\langle a \rangle}(\neg x_2)$ is not necessarily in C. Therefore the proof will follow if we adjust our C to $\{h^*_{X_1 \to x_1, X_2 \to x_2}(\psi) \mid \psi \in st(\sigma)\}$.

The alterations of C are necessary since we are evaluating roots from the algebra $\mathbf{A}$ in the algebra $\mathbf{A}_C$ in the proof that j^* and h^* agree on least root formulae. Notice that in proving $j^*(\sigma) \geq h^*(\sigma)$ the alterations to C are not needed. We will see that in general (Lemma 5.2.19) this is due both to the construction of the $\mathbf{A}_C$ operators $(\tilde{f}_{\langle a \rangle})_{a \in \Pi_0}$ and the prevention of conjunctive occurrence in least root formulae.

Motivated by the previous discussion we now propose a general proof strategy for proving that $\mathfrak{R} \not\models \sigma^{\approx}$ implies $\mathfrak{R}_{fin} \not\models \sigma^{\approx}$ for $\sigma \in L_\rho(\Phi, Var)$. However we will first need to make the following two definitions.

5.2.14 Definition. Let $\sigma \in L_\rho(\Phi, Var)$ and let $X_1, X_2 \in \text{b.v.}(\sigma)$. Define $X_1 \prec X_2$ iff X_2 is in the scope of the binder binding X_1.

5.2.15 Example. Consider $\sigma = \rho X.(\neg X \wedge \rho Y.(X \wedge \langle b \rangle \neg Y))$. We have that $X \prec Y$.

5.2.16 Remark. For a $\sigma \in L_\rho(\Phi, Var)$ we do not always have that $\prec$ is a linear order on b.v.(σ). For example this is true when $\sigma = \rho X.(\neg X \wedge p) \vee \rho Y.(\neg Y \wedge q)$. Therefore we consider a linear extension of $\prec$ (denoted $\prec'$) on b.v.(σ).

5.2.17 Definition. A formula $\sigma \in L_\rho(\Phi, Var)$ is *clean* if

(i) no variable occurs both bound and free in σ and
(ii) no two binders bind the same variable in σ.

5.2.18 Remark. Consider the unclean L_ρ formula $\rho X.(\neg X \wedge \rho X.(\neg X \wedge p))$ and recall the second scenario discussed above (which involved a nested least root formula). We would not be able to construct C as desired since X needs to be mapped to two different elements. Any unclean formula can be "baptised" into a clean formula by a suitable relabelling of the bound variables. In particular σ can be rewritten as $\rho X.(\neg X \wedge \rho Y.(\neg Y \wedge p))$. Without loss of generality, we assume all L_ρ formulae to be clean.

We now present our general strategy for demonstrating that $\mathfrak{R} \not\models \sigma^{\approx}$ implies $\mathfrak{R}_{fin} \not\models \sigma^{\approx}$ for all $\sigma \in L_\rho(\Phi, Var)$. This strategy is based on Pratt's proof strategy in [34] but differs in the construction of C and hence also $j : \Phi \cup Var \to A_C$.

- Let $\sigma \in L_\rho(\Phi, Var)$. Suppose $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle a \rangle})_{a \in \Pi_0} \in \mathfrak{R}$ such that $\mathbf{A} \not\models \sigma^{\approx}$ i.e. there is an $h : \Phi \cup Var \to A$ with $h^*(\sigma) \neq 1$.
- Let b.v.$(\sigma) = \{X_1, ..., X_m\}$ ordered as $X_1 \prec' ... \prec' X_m$. Inductively define the following variants of h up to m:

$$\begin{array}{rclcrcl} h_1 & = & h_{X_1 \to x_1} & \text{where} & x_1 & = & h^*(\rho X_1.\varphi_1(X_1)) \\ h_2 & = & h_{1,\, X_2 \to x_2} & \text{where} & x_2 & = & h_1^*(\rho X_2.\varphi_2(X_2)) \\ \vdots & & & \vdots & & & \\ h_i & = & h_{i-1,\, X_i \to x_i} & \text{where} & x_i & = & h_{i-1}^*(\rho X_i.\varphi_i(X_i)). \end{array}$$

- Let $C = \{h_m^*(\psi) \mid \psi \in st(\sigma)\}$ and consider the algebraic filtration $(\mathbf{A}_C, f^{R^s_{\langle a \rangle}})_{a \in \Pi_0}$ of $\mathbf{A}$ through $(C, \underline{C}_a)_{a \in \Pi_0}$ with $(R^s_{\langle a \rangle})_{a \in \Pi_0}$. For each $a \in \Pi_0$ denote the $\mathbf{A}_C$ operator $f^{R^s_{\langle a \rangle}}$ by $\tilde{f}_{\langle a \rangle}$; it needs to be checked that $(\mathbf{A}_C, \tilde{f}_{\langle a \rangle})_{a \in \Pi_0} \in \mathfrak{R}$ and that $\mathbf{A}_C \not\models \sigma^{\approx}$. To this end define $j : \Phi \cup Var \to A_C$ by $j(p) = h_m(p)$ and prove that $j^*(\sigma) = h^*(\sigma)$.

Note that this construction has already partly been seen when using the $\mathscr{F}$-closure in the proof that **PDL** is decidable (Section 5.1). For example if we consider the PDL formula $\langle a^* \rangle\, p$, which is $\mu X.(p \vee \langle a \rangle\, X)$ by Theorem 3.2.19, the set C is constructed as $\{h^*(\psi) \mid \psi \in \mathscr{F}(st(\langle a^* \rangle\, p))\} =$

$\{f_{\langle a^*\rangle}(h(p)), f_{\langle a\rangle}f_{\langle a^*\rangle}(h(p)),\ h(p)\}$ i.e. $C = \{h^*_{X\to f_{\langle a^*\rangle}(h(p))}(\psi) \mid \psi \in st(\mu X.(p \vee \langle a\rangle X))\}$. The following two lemmas (and corollary) will be used in the last step of the proof strategy i.e. when proving that j^* and h^* agree on least root formulae. For the first lemma recall the definition of an $\overline{X}$-variant (or vector variant) of a mapping from Definition 3.2.7.

5.2.19 Lemma. Suppose $\eta = \rho X.\varphi(X) \in L_\rho(\Phi, Var)$. Then $h^*_m(\eta) \leq j^*(\eta)$.

Proof. The result will be proved by induction on the length of L_ρ formulas. For the induction to work we actually prove a stronger statement i.e. we will show that for every $\psi \in st(\eta)$:

$$h^*_{m,\ \overline{Y}\to\overline{y}}(\psi) \leq j^*_{\overline{Y}\to\overline{y}}(\psi) \text{ for all } \overline{Y} = (Y_1, ..., Y_n) \in (\text{b.v.}(\eta))^n, \overline{y} = (y_1, ..., y_n) \in A^n_C \text{ and } n \in \mathbb{N}. \tag{5.2.2}$$

Base case: Let $n \in \mathbb{N}$, let $\overline{Y} = (Y_1, ..., Y_n) \in (\text{b.v.}(\eta))^n$, let $\overline{y} = (y_1, ..., y_n) \in A^n_C$ and consider $\psi \in st(\eta)$ such that $\text{k}(\psi) = 0$ i.e. $\psi = \bot$ or $\psi = p \in \Phi \cup Var$. For either case we have by definition that $h^*_{m,\ \overline{Y}\to\overline{y}}(\psi) = j^*_{\overline{Y}\to\overline{y}}(\psi)$. As all L_ρ formulae are assumed to be in normal form (see Definition 5.2.8) we deal with negation in the base case and in particular we have $h^*_{m,\ \overline{Y}\to\overline{y}}(\neg\psi) = \neg(h^*_{m,\ \overline{Y}\to\overline{y}}(\psi)) = \neg(j^*_{\overline{Y}\to\overline{y}}(\psi)) = j^*_{\overline{Y}\to\overline{y}}(\neg\psi)$.

Inductive hypothesis: Assume statement 5.2.2 is true for all $\psi \in st(\eta)$ such that $\text{k}(\psi) < l \leq \text{k}(\eta)$.

Inductive step: Let $n \in \mathbb{N}$, let $\overline{Y} = (Y_1, ..., Y_n) \in (\text{b.v.}(\eta))^n$, let $\overline{y} = (y_1, ..., y_n) \in A^n_C$ and let $\psi \in st(\eta)$ such that $\text{k}(\psi) = l$. Thus for some $\varphi_1, \varphi_2, \theta(Z) \in L_\rho(\Phi, Var)$ we have:

$$\psi = \begin{cases} \varphi_1 \vee \varphi_2 & (i) \\ \varphi_1 \wedge \varphi_2 & (ii) \\ \langle a\rangle\,\varphi & (iii) \\ \rho Z.\theta(Z) & (iv). \end{cases}$$

For the cases of $[a]\varphi$ and $\iota Z.\theta(Z)$, we actually have equality in statement 5.2.2. This is due to the fact that Y_i cannot occur in these subterms for all $1 \leq i \leq n$, otherwise this would serve as a conjunctive occurance of bound variables in η.

(i) By the inductive hypothesis $h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi_1) \leq j^*_{\overline{Y}\to\overline{y}}(\varphi_1)$ and $h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi_2) \leq j^*_{\overline{Y}\to\overline{y}}(\varphi_2)$. Thus

$$\begin{aligned} h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi_1 \vee \varphi_2) &= h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi_1) \vee h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi_2) \\ &\leq j^*_{\overline{Y}\to\overline{y}}(\varphi_1) \vee j^*_{\overline{Y}\to\overline{y}}(\varphi_2) \\ &= j^*_{\overline{Y}\to\overline{y}}(\varphi_1 \vee \varphi_2). \end{aligned}$$

(ii) Similar to case (i).

(iii) By the inductive hypothesis $h^*_{m,\overline{Y}\to\overline{y}}(\varphi) \leq j^*_{\overline{Y}\to\overline{y}}(\varphi)$. Since $f_{\langle a\rangle}$ is an operator on $\mathbf{A}$ (and hence order-preserving) $f_{\langle a\rangle}(h^*_{m,\ \overline{Y}\to\overline{y}}(\varphi)) \leq f_{\langle a\rangle}(j^*_{\overline{Y}\to\overline{y}}(\varphi))$. But by Lemma 4.2.13 we have that

$f_{\langle a\rangle}(j^*_{\overline{Y}\to\overline{y}}(\varphi)) \leq \tilde{f}_{\langle a\rangle}(j^*_{\overline{Y}\to\overline{y}}(\varphi))$ and hence the result holds.
(iv) Let $z = j^*_{\overline{Y}\to\overline{y}}(\rho Z.\theta(Z))$. We know that $h^*_{m,\,\overline{Y}\to\overline{y}}(\rho Z.\theta(Z)) \leq z$ if z is a root of the mapping $h^*_{m,\,\overline{Y}\to\overline{y},\,Z\to}(\theta(Z))$. This is true since

$$\begin{aligned} h^*_{m,\,\overline{Y}\to\overline{y},\,Z\to z}(\theta(Z)) &\leq j^*_{\overline{Y}\to\overline{y},\,Z\to z}(\theta(Z)) && \text{(by inductive hypothesis)} \\ &= 0 && \text{(by defn. of } z\text{).} \end{aligned}$$

□

5.2.20 Lemma. Let $\mathbf{A} \in \mathfrak{R}$ and let $\varphi \in L_\rho(\Phi, Var)$. If $g_1, g_2 : \Phi \cup Var \to A$ are two maps that agree on $\text{f.v.}(\varphi) \cup \Phi \cup \{\perp\}$ (i.e. $g_1(X) = g_2(X)$ for all $X \in \text{f.v.}(\varphi) \cup \Phi \cup \{\perp\}$) then $g_1^*(\varphi) = g_2^*(\varphi)$.

Proof. We will show for all $\psi \in L_\rho(\Phi, Var)$ that:

$$\text{If } g_1, g_2 : \Phi \cup Var \to A \text{ are any two maps agreeing on f.v.}(\psi) \cup \Phi \cup \{\perp\} \text{ then } g_1^*(\psi) = g_2^*(\psi). \tag{5.2.3}$$

Base case: Let $\psi \in L_\rho(\Phi, Var)$ such that $\text{k}(\psi) = 0$ and let $g_1, g_2 : \Phi \cup Var \to A$ be two maps agreeing on $\text{f.v.}(\psi) \cup \Phi \cup \{\perp\}$. We have that either $\psi = \perp$, $\psi = p \in \Phi$ or $\psi = X \in Var$. When $\psi = \perp$ or $\psi = p$, conditional 5.2.3 is true as g_1 and g_2 agree on $\Phi \cup \{\perp\}$. If $\psi = X$, we have that $g_1(\psi) = g_2(\psi)$ since $\text{f.v.}(\psi) = \{\psi\}$, and thus $g_1^*(\psi) = g_2^*(\psi)$.

Inductive hypothesis: Assume conditional 5.2.3 is true for all $\psi \in L_\rho(\Phi, Var)$ with $\text{k}(\psi) < n$.

Inductive step: Let $\psi \in L_\rho(\Phi, Var)$ such that $\text{k}(\psi) = n$. Thus for some $\varphi_1, \varphi_2, \varphi(X) \in L_\rho(\Phi, Var)$ and $a \in \Pi_0$ we have:

$$\psi = \begin{cases} \varphi_1 \vee \varphi_2 & (i) \\ \neg\varphi & (ii) \\ \langle a\rangle\,\varphi & (iii) \\ \rho X.\varphi(X) & (iv). \end{cases}$$

Let $g_1, g_2 : \Phi \cup Var \to A$ be two maps agreeing on $\text{f.v.}(\psi) \cup \Phi \cup \{\perp\}$.
(i) By definition it is true that

$$\begin{aligned} \text{f.v.}(\varphi_1) &\subseteq \text{f.v.}(\varphi_1 \vee \varphi_2) \\ \text{f.v.}(\varphi_2) &\subseteq \text{f.v.}(\varphi_1 \vee \varphi_2). \end{aligned}$$

Thus by the inductive hypothesis $g_1^*(\varphi_1) = g_2^*(\varphi_1)$ and $g_1^*(\varphi_2) = g_2^*(\varphi_2)$ and therefore

$$\begin{aligned} g_1^*(\varphi_1 \vee \varphi_2) &= g_1^*(\varphi_1) \vee g_1^*(\varphi_2) \\ &= g_2^*(\varphi_1) \vee g_2^*(\varphi_2) \\ &= g_2^*(\varphi_1 \vee \varphi_2). \end{aligned}$$

Cases (ii) and (iii) follow similarly.
(iv) Let $z_1 = g_1^*(\rho X.\varphi(X))$ and let $z_2 = g_2^*(\rho X.\varphi(X))$. We need to show that $z_1 = z_2$.
$\underline{z_1 \leq z_2}$ if z_2 is a root of $g_1^*{}_{X\to}(\varphi(X))$ i.e. if $g_1^*{}_{X\to z_2}(\varphi(X)) = 0$. By assumption the two maps g_1 and g_1 agree on f.v.$(\rho X.\varphi(X))$. Thus g_1 and g_2 agree on f.v.$(\varphi(X)) \setminus \{X\}$ as f.v.$(\rho X.\varphi(X)) =$ f.v.$(\varphi(X)) \setminus \{X\}$. Therefore the two maps $g_1{}_{X\to z_2}$ and $g_2{}_{X\to z_2}$ agree on f.v.$(\varphi(X))$ and thus by the inductive hypothesis $g_1^*{}_{X\to z_2}(\varphi(X)) = g_2^*{}_{X\to z_2}(\varphi(X))$. But by definition of z_2 we know $g_2^*{}_{X\to z_2}(\varphi(X)) = 0$. The proof of $z_1 \geq z_2$ follows similarly. □

Recall from the proof strategy that b.v.$(\sigma) = \{X_1, ..., X_m\}$ where $X_1 \prec' X_2 \prec' ... \prec' X_m$.

5.2.21 Corollary. Let $i \in \{1, ..., m\}$; then $h_m^*(\rho X_i.\varphi_i(X_i)) = h_{i-1}^*(\rho X_i.\varphi_i(X_i))$.

Proof. The maps h_m and h_{i-1} can only disagree on elements within Γ where $\Gamma = \{X_i, ..., X_m\}$. But $\Gamma \cap$ f.v.$(\rho X_i.\varphi_i(X_i)) = \varnothing$ and the result follows by Lemma 5.2.20. □

5.2.22 Theorem. *Let $\sigma \in L_\rho(\Phi, Var)$. Then $\mathfrak{R} \not\models \sigma^{\approx}$ implies $\mathfrak{R}_{fin} \not\models \sigma^{\approx}$.*

Proof. Recall from the proof strategy that $\sigma \in L_\rho(\Phi, Var)$ and $\mathbf{A} = (A, \vee, \wedge, \neg, 0, 1, f_{\langle a\rangle})_{a\in\Pi_0} \in \mathfrak{R}$ such that $\mathbf{A} \not\models \sigma^{\approx}$ i.e. there is an $h : \Phi \cup Var \to A$ with $h^*(\sigma) \neq 1$. Recall $h_m : \Phi \cup Var \to A$ and define $C = \{h_m^*(\psi) \mid \psi \in st(\sigma)\} \subseteq_{fin} A$. Consider the algebraic filtration $(\mathbf{A}_C, f^{R^s_{\langle a\rangle}})_{a\in\Pi_0}$ of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ with $(R^s_{\langle a\rangle})_{a\in\Pi_0}$. For each $a \in \Pi_0$ denote the $\mathbf{A}_C$ operator $f^{R^s_{\langle a\rangle}}$ by $\tilde{f}_{\langle a\rangle}$; we will show $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0}$ is a finite falsifying ρ-algebra. Since $\mathbf{A}_C$ is a finite τ_0 modal algebra we know by Lemma 5.2.13 that $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0} \in \mathfrak{R}$. To prove that $(\mathbf{A}_C, \tilde{f}_{\langle a\rangle})_{a\in\Pi_0} \not\models \sigma^{\approx}$ we define $j : \Phi \cup Var \to A_C$ by $j(p) = h_m(p)$ and will demonstrate that $j^*(\sigma) = h^*(\sigma)$. Since h and h_m agree on f.v.$(\sigma) \cup \Phi \cup \{\bot\}$ we have $h^*(\sigma) = h_m^*(\sigma)$ by Lemma 5.2.20. Thus we will show for all $\psi \in st(\sigma)$ that

$$j^*(\psi) = h_m^*(\psi). \tag{5.2.4}$$

Base case: Let $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) = 0$ i.e. $\psi = \bot$ or $\psi = p \in \Phi \cup Var$. If $\psi = \bot$ then equation 5.2.4 holds as $j^*(\bot) = 0 = h_m^*(\bot)$. The same is true for $\psi = p \in \Phi \cup Var$ since $j^*(p) = j(p) = h_m(p) = h_m^*(p)$.

Inductive hypothesis: Assume equation 5.2.4 is true for all $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) < n \leq \mathrm{k}(\sigma)$.

Inductive step: Let $\psi \in st(\sigma)$ such that $\mathrm{k}(\psi) = n$. Thus for some $\varphi_1, \varphi_2, \varphi(X) \in L_\rho(\Phi, Var)$ and $a \in \Pi_0$ we have:

$$\psi = \begin{cases} \varphi_1 \vee \varphi_2 & (i) \\ \neg\varphi & (ii) \\ \langle a \rangle \varphi & (iii) \\ \rho X.\varphi(X) & (iv). \end{cases}$$

Cases $(i), (ii)$ and (iii) follow since $\mathbf{A}_C$ is an algebraic filtration of $\mathbf{A}$ through $(C, \underline{C}_a)_{a\in\Pi_0}$ (see Lemma 4.2.11).

(iv) Recall from the proof strategy that b.v.$(\sigma) = \{X_1, ..., X_m\}$; we thus have that $\rho X.\varphi(X) = \rho X_i.\varphi_i(X_i)$ where $1 \leq i \leq m$. We need to show that $j^*(\rho X_i.\varphi_i(X_i)) = h^*_m(\rho X_i.\varphi_i(X_i))$. From Lemma 5.2.19 we immediately know that $j^*(\rho X_i.\varphi_i(X_i)) \geq h^*_m(\rho X_i.\varphi_i(X_i))$.

$\underline{j^*(\rho X_i.\varphi_i(X_i)) \leq h^*_m(\rho X_i.\varphi_i(X_i))}$:

Note that by Corollary 5.2.21 we have $h^*_m(\rho X_i.\varphi_i(X_i)) = h^*_{i-1}(\rho X_i.\varphi_i(X_i))$. Furthermore by definition we have $h^*_{i-1}(\rho X_i.\varphi_i(X_i)) = x_i$. Thus we need to show that $j^*(\rho X_i.\varphi_i(X_i)) \leq x_i$. This is true if x_i is a root of $j^*_{X_i\to}(\varphi_i(X_i))$ i.e. if $j^*_{X_i\to x_i}(\varphi_i(X_i)) = 0$. By definition of j we have that $j = j_{X_i\to x_i}$ and thus we are looking to prove $j^*(\varphi_i(X_i)) = 0$. By the inductive hypothesis we know $j^*(\varphi_i(X_i)) = h^*_m(\varphi_i(X_i))$. Since h_m and h_i agree on f.v.$(\varphi_i(X_i)) \cup \Phi \cup \{\perp\}$ we know that $h^*_m(\varphi_i(X_i)) = h^*_i(\varphi_i(X_i))$ by Lemma 5.2.20. But $h^*_i(\varphi_i(X_i)) = 0$ since $h^*_i(\varphi_i(X_i)) = h^*_{i-1,\ X_i\to x_i}(\varphi_i(X_i)) = 0$ by the definition of x_i. Thus $j^*(\varphi_i(X_i)) = 0$ and therefore $j^*(\rho X_i.\varphi_i(X_i)) \leq h^*_m(\rho X_i.\varphi_i(X_i))$. Hence, in all we have that $j^*(\rho X_i.\varphi_i(X_i)) = h^*_m(\rho X_i.\varphi_i(X_i))$. □

6. Conclusion

We began this work by looking at the tools needed to understand decidability of a logic. We then analysed the question of decidability for two logics which contain least binders. The first such logic we looked at was **PDL** and we reviewed how it can be shown that **PDL** is be decidable via a relational argument [11]. We then gave a new proof that **PDL** is decidable by showing that **PDL** is complete with respect to $\mathfrak{D}$ (Theorem 5.1.11) and that the validities of $\mathfrak{D}$ and $\mathfrak{D}_{fin}$ are the same (Theorem 5.1.15). The proof that $\mathfrak{D}_{fin} \models \sigma^{\approx}$ implies $\mathfrak{D} \models \sigma^{\approx}$ built on the algebraic filtration work of Conradie et al. [7] reviewed (in part) in Chapter 4.

Notice that in the proof of $\mathfrak{D}_{fin} \models \sigma^{\approx}$ implies $\mathfrak{D} \models \sigma^{\approx}$ (Theorem 5.1.15), no reference is made to the underlying algebraic filtrators and we could have simply taken $(\mathbf{A}_C, f'_{\langle a \rangle})_{a \in \Pi_0}$ as a starting point for our falsifying algebra, where $f'_{\langle a \rangle}(x) = \bigwedge\{y \in A_C \mid f_{\langle a \rangle}(x) \leq y\}$ for $x \in A_C$. However relating the construction to the algebraic filtrations of Conradie et al. opens the door towards a comparison of the algebraic and relational proofs that **PDL** is decidable and indeed this could be a possible future research direction. Also notice that, analogously to the finite model property proof in Section 5.2, we could have directly proved L^*_μ is decidable. We instead proved **PDL** is decidable since this allows for a more direct comparison with the relational proof in [11].

We then looked at the issue of decidability for $\Lambda_{\mathfrak{R}}$. We demonstrated that $\Lambda_{\mathfrak{R}}$ has the finite model property with respect to $\mathfrak{R}$ (Theorem 5.2.22). Our proof was based on that of Pratt but differed in the construction of the generating set used to build the finite falsifying ρ-algebra [34]. If it can be proven that $\Lambda_{\mathfrak{R}}$ is finitely axiomatisable then this serves to demonstrate that $\Lambda_{\mathfrak{R}}$ is decidable.

Comparing the two proofs that **PDL** and $\Lambda_{\mathfrak{R}}$ have the finite model property, it is easy to see the advantage of expressing all of the PDL modalities using only basic modalities and least root formulae. For example, if instead we were to prove directly that L^*_μ is decidable, we would have avoided the use of Theorem 5.1.3 pertaining to the $\mathscr{F}$-closure.

Possible future research directions include finding richer fragments of L_μ for which decidability can be shown with algebraic filtrations.

References

[1] P. Blackburn, M. de Rijke, and Y. Venema. Modal Logic. Cambridge University Press, 2001.

[2] W. Blok. The Lattice of Modal Logics: an Algebraic Investigation. *The Journal of Symbolic Logic*, 45:221–236, 1980.

[3] G. Boole. An Investigation of the Laws of Thought. Walton and Maberley, 1854.

[4] J. Bradfield and C. Stirling. *Modal μ-Calculi.* Handbook of Modal Logic. Editors Blackburn, Wolter and van Benthem. Elsevier, 2006.

[5] S. Burris and H. Sankappanavar. A Course in Universal Algebra. Springer, 1981.

[6] E. M. Clarke and E. A. Emerson. Design and Synthesis of Synchronization Skeletons using Branching Time Temporal Logic. *LNCS*, 131:52–71, 1981.

[7] W. Conradie, W. Morton, and C. Van Alten. An Algebraic Look at Filtrations in Modal Logic. *Logic Journal of the IGPL*, 21:788–811, 2013.

[8] B. Davey and H. Priestley. Introduction to Lattices and Order. Cambridge University Press, 2002.

[9] F. Donini, M. Lenzerini, D. Nardi, and A. Schaerf. *Reasoning in Description Logics.* Principles of Knowledge Representation, Studies in Logic, Language and Information. Editor G. Brewka. CSLI Publications, 1996.

[10] H. Enderton. A Mathematical Introduction to Logic. Academic Press, 2001.

[11] M. Fischer and R. Ladner. Propositional Dynamic Logic of Regular Programs. *Journal of Computer and System Sciences*, 18:194–211, 1979.

[12] D. Gabbay. On Decidable, Finitely Axiomatizable, Modal and Tense Logics without the Finite Model Property, part I. *Israel Journal of Mathematics*, 11:478–495, 1972.

[13] L. Giacomo. The Modal μ-Calculus: A Survey. *Task Quarterly*, 9:293–316, 2004.

[14] K. Gödel. Eine Interpretation des Intuitionistischen Aussagenkalküls. *Ergebnisse eines mathematischen Kolloquiums*, 4:39–40, 1933.

[15] R. Goldblatt. *Mathematical Modal Logic: A View of its Evolution.* Handbook of the History of Logic. Editors D.M. Gabbay and J. Woods. Elsevier, Amsterdam, 2006.

[16] R. Harrop. On the Existence of Finite Models and Decision Procedures for Propositional Calculi. *Proceedings of the Cambridge Philosophical Society*, 54:1–13, 1958.

[17] L. Henkin. The Completeness of First-Order Functional Calculus. *The Journal of Symbolic Logic*, 14:159–166, 1949.

[18] M. Hennessy and R. Milner. *On Observing Nondeterminism and Concurrency.* Automata, Languages and Programming, volume 85 of Lecture Notes in Computer Science. Editors J.W. de Bakker and J. van Leeuwen. Springer–Verlag, 1980.

[19] C. Hoare. An Axiomatic Basis for Computer Programming. *Communications of the Association for Computer Machinery*, 12:576–580, 1969.

[20] D. Janin and I. Walukiewicz. On the Expressive Completeness of the Modal Mu-Calculus with respect to Monadic Second Order Logic. *Conf. on Concurrency Theory (CONCUR '96)*, LNCS 1119:263–277, 1996.

[21] B. Jónsson and A. Tarski. Boolean Algebras with Operators, Part I. *American Journal of Mathematics*, 73:891–939, 1952.

[22] D. Kozen. Results on the Propositional μ-Calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[23] D. Kozen and R. Parikh. *A Decision Procedure for the Propositional μ-Calculus.* Logics of Programs. Editors E. Clarke and D. Kozen. Springer–Verlag, 1984.

[24] S. Kripke. A Completeness Theorem in Modal Logic. *The Journal of Symbolic Logic*, 24:1–14, 1959.

[25] S. Kripke. Semantical Analysis of Modal Logic I, Normal Propositional Calculi. *Zeitschrift für Mathemathische Logic und Grundlagen der Mathematik*, 9:67–96, 1963.

[26] E. Lemmon and D. Scott. Intensional Logic. Stanford University, 1966.

[27] C. Lewis. Implication and the Algebra of Logic. *Mind*, 21:522–531, 1912.

[28] C. Lewis and C. Langford. Symbolic Logic. The Century Co., 1932.

[29] H. MacColl. Symbolic Reasoning. *Mind*, 5:54, 1880.

[30] Z. Manna and A. Pnueli. Formalization of Properties of Recursively Defined Functions. *Proceedings ACM STOC*, pages 201–210, 1969.

[31] J. McKinsey. A Solution to the Decision Problem for the Lewis Systems S2 and S4 with an Application to Topology. *The Journal of Symbolic Logic*, 6:117–134, 1941.

[32] V. Pratt. Semantical Considerations on Floyd–Hoare Logic. *Proceedings of the 17th Annual IEEE Symposium on Foundations of Computer Science*, pages 109–121, 1976.

[33] V. Pratt. Dynamic Algebras and the Nature of Induction. *Proc. 12th Symp. Theory of Computing*, pages 22–28, 1980.

[34] V. Pratt. A Decidable Mu-Calculus: Preliminary Report. *Foundations of Computer Science 22nd Annual Symposium*, pages 421–427, 1981.

[35] M. Rabin. Decidability of Second-Order Theories and Automata on Infinite Trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.

[36] D. Scott and J. De Bakker. A Theory of Programs. *Unpublished manuscript, IBM, Vienna*, 1969.

[37] K. Segerberg. A Completeness Theorem in the Modal Logic of Programs. *Notices of the American Mathematical Society*, 24:A–552, 1977.

[38] R. Streett. Propositional Dynamic Logic of Looping and Converse. *Proc. 13th ACM Symp. on Theory of Computing*, pages 375–383, 1981.

[39] S. Thomason. Semantic Analysis of Tense Logics. *Journal of Symbolic Logic*, 37:150–158, 1972.

[40] I. Walukiewicz. Completeness of Kozen's Axiomatisation of the Propositional μ-Calculus. *Proceedings of the Tenth Annual IEEE Symposium on Logic in Computer Science. Editor D. Kozen*, IEEE Computer Science Press, 1995.