

Completeness for Flat Modal Fixpoint Logics*

Luigi Santocanale[†] Yde Venema[‡]

December 12, 2008

Abstract

This paper exhibits a general and uniform method to prove completeness for certain modal fixpoint logics. Given a set Γ of modal formulas of the form $\gamma(x, p_1, \dots, p_n)$, where x occurs only positively in γ , the language $\mathcal{L}_{\#}(\Gamma)$ is obtained by adding to the language of polymodal logic a connective $\#_{\gamma}$ for each $\gamma \in \Gamma$. The term $\#_{\gamma}(\varphi_1, \dots, \varphi_n)$ is meant to be interpreted as the least fixed point of the functional interpretation of the term $\gamma(x, \varphi_1, \dots, \varphi_n)$. We consider the following problem: given Γ , construct an axiom system which is sound and complete with respect to the concrete interpretation of the language $\mathcal{L}_{\#}(\Gamma)$ on Kripke frames. We prove two results that solve this problem.

First, let $\mathbf{K}_{\#}(\Gamma)$ be the logic obtained from the basic polymodal \mathbf{K} by adding a Kozen-Park style fixpoint axiom and a least fixpoint rule, for each fixpoint connective $\#_{\gamma}$. Provided that each indexing formula γ satisfies the syntactic criterion of being untied in x , we prove this axiom system to be complete.

Second, addressing the general case, we prove the soundness and completeness of an extension $\mathbf{K}_{\#}^+(\Gamma)$ of $\mathbf{K}_{\#}(\Gamma)$. This extension is obtained via an effective procedure that, given an indexing formula γ as input, returns a finite set of axioms and derivation rules for $\#_{\gamma}$, of size bounded by the length of γ . Thus the axiom system $\mathbf{K}_{\#}^+(\Gamma)$ is finite whenever Γ is finite.

Keywords. fixpoint logic, modal logic, axiomatization, completeness, least fixpoint, modal algebra, representation theorem

1 Introduction

Suppose that we extend the language of basic (poly-)modal logic with a set $\{\#_{\gamma} \mid \gamma \in \Gamma\}$ of so-called *fixpoint connectives*, which are defined as follows. Each connective $\#_{\gamma}$ is indexed by a modal formula $\gamma(x, p_1, \dots, p_n)$ in which x occurs only positively. The intended meaning of the formula $\#_{\gamma}(\varphi_1, \dots, \varphi_n)$ in a labelled transition system (Kripke model) is the least fixpoint of the formula $\gamma(x, \varphi_1, \dots, \varphi_n)$,

$$\#_{\gamma}(\varphi_1, \dots, \varphi_n) \equiv \mu x. \gamma(x, \varphi_1, \dots, \varphi_n).$$

*Research supported by the Van Gogh research project *Modal Fixpoint Logics*.

[†]Laboratoire d'Informatique Fondamentale de Marseille, Université de Provence, 39 rue F. Joliot Curie, 13453 Marseille Cedex 13, France. *Email:* luigi.santocanale@lif.univ-mrs.fr

[‡]Institute for Logic, Language and Computation, Universiteit van Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, Netherlands. *Email:* Y.Venema@uva.nl

Many logics of interest in computer science are of this kind: Such fixpoint connectives can be found for instance in PDL, propositional dynamic logic [14], in CTL, computation tree logic [11], in LTL, linear temporal logic, and in multi-agent versions of epistemic logic [12]. More concretely, the Kleene iteration diamond $\langle a^* \rangle$ of PDL can be presented (in the case of an atomic program a) as the connective \sharp_δ , where $\delta(x, p)$ is the formula $p \vee \langle a \rangle x$: the formula $\langle a^* \rangle \varphi$ can be interpreted as the parameterized least fixpoint $\mu_x. \delta(x, \varphi)$. As two more examples, let $\theta(x, p, q) := p \vee (q \wedge \diamond x)$, and $\eta(x, p, q) := p \vee (q \wedge \Box x)$, then CTL adds new connectives $\sharp_\theta(p, q), \sharp_\eta(p, q)$ — or $E(pUq), A(pUq)$ in the standard notation — to the basic modal language.

Generalizing these examples we arrive at the notion of a flat modal fixpoint logic. Let $\mathcal{L}_\sharp(\Gamma)$ denote the language we obtain if we extend the syntax of (poly-)modal logic with a connective \sharp_γ for every $\gamma \in \Gamma$. Clearly, every fixpoint connective of this kind can be seen as a macro over the language of the modal μ -calculus. Because the associated formula γ of a fixpoint connective is itself a basic modal formula (which explains our name *flat*), it is easy to see that every flat modal fixpoint language is contained in the alternation-free fragment of the modal μ -calculus [19]. Because of their transparency and simpler semantics, flat modal fixpoint logics such as CTL and LTL are often preferred by end users. In fact, most verification tools implement some flat fixpoint logic rather than the full μ -calculus, regardless of considerations based on the expressive power of these logics.

Despite their wide-spread applications and mathematical interest, up to now general investigations of modal fixpoint logics have been few and far between. In this paper we address the natural problem of *axiomatizing* flat modal fixpoint logics. Here the flat modal fixpoint logic induced by Γ is the set of $\mathcal{L}_\sharp(\Gamma)$ -validities, that is, the collection of formulas in the language $\mathcal{L}_\sharp(\Gamma)$ that are true at every state of every Kripke model.

In general, the problem of axiomatizing fixpoints arising in computer science is recognized to be a nontrivial one. As an example we mention the longstanding problem of axiomatizing regular expressions [9, 7, 22, 20], whereas the monograph [6] is a good general survey on fixpoint theory. More specifically, in the literature on modal logic one may find completeness results for a large number of individual systems. We mention the work of Segerberg [36] and of Kozen & Parikh [21] on PDL, the axiomatization of Emerson & Halpern [10] of CTL, and many results on epistemic logic with the common knowledge operator or similar modalities [12, 29]. In the paper [19] that introduced the modal μ -calculus, Kozen proposed an axiomatization which he proved to be complete for a fragment of the language; the completeness problem of this axiomatization for the full language was solved positively by Walukiewicz [40]. But to our knowledge, no general results or uniform proof methods have been established in the theory of modal fixpoint logics. For instance, the classical filtration methods from modal logic work for relatively simple logics such as PDL [14], but they already fail if this logic is extended with the loop operator [19]. A first step towards a general understanding of flat fixpoint logics is the work [26], where a game-based approach is developed to deal with axiomatization and satisfiability issues for LTL and CTL.

In this paper we contribute to the general theory of flat modal fixpoint logics by providing completeness results that are uniform in the parameter Γ , and modular in the sense that the axiomatizations take care of each fixpoint connective separately. Our research is driven by the

wish to understand the combinatorics of fixpoint logics in their wider mathematical setting. As such it continues earlier investigations by the first author into the algebraic and order-theoretic aspects of fixpoint calculi [34, 35], and work by the second author on coalgebraic (fixpoint) logics [39, 23, 25].

Usually, the difficulty in finding a complete axiomatization problem for a fixpoint logic does not stem from the absence of a natural candidate. In our case, mimicking Kozen's axiomatization of the modal μ -calculus, an intuitive axiomatization for the $\mathcal{L}_{\#}(\Gamma)$ -validities would be to add, to some standard axiomatization \mathbf{K} for (poly-)modal logic, an axiom for each connective $\#_{\gamma}$ stating that $\#_{\gamma}(p_1, \dots, p_n)$ is a prefixpoint of the formula $\gamma(x, p_1, \dots, p_n)$, and a derivation rule which embodies the fact that $\#_{\gamma}(p_1, \dots, p_n)$ is the smallest such.

Definition 1.1. The axiom system $\mathbf{K}_{\#}(\Gamma)$ is obtained by adding to \mathbf{K} the axiom

$$\gamma(\#_{\gamma}(p_1, \dots, p_n), p_1, \dots, p_n) \rightarrow \#_{\gamma}(p_1, \dots, p_n), \quad (\#_{\gamma}\text{-prefix})$$

and the derivation rule¹

$$\frac{\gamma(y, p_1, \dots, p_n) \rightarrow y}{\#_{\gamma}(p_1, \dots, p_n) \rightarrow y} \quad (\#_{\gamma}\text{-least})$$

for each $\gamma \in \Gamma$.

In fact, the first of our two main results, Theorem 5.4, states that for many choices of Γ , $\mathbf{K}_{\#}(\Gamma)$ is indeed a complete axiomatization. More precisely, we identify a class of formulas that we call *untied in x* — these formulas are related to the aconjunctive [19] and disjunctive [40] formulas from the modal μ -calculus. In this paper we shall prove that

if every γ in Γ is untied in x , then $\mathbf{K}_{\#}(\Gamma)$ is a complete axiomatization.

This result takes care of for instance the completeness of CTL.

However, the road to a general completeness result for the system $\mathbf{K}_{\#}(\Gamma)$ is obstructed by a familiar problem, related to the role of conjunctions in the theory of fixpoint logics. Our solution to this problem comprises a *modification* of the intuitive Kozen-style axiomatization, inspired by a construction of Arnold & Niwiński [2]. Roughly speaking, this so-called Subset Construction is a procedure that simulates a suitable system of equations T by a system of equations T_{γ}^{+} that we will call *simple* since it severely restricts occurrences of the conjunction symbol. It is shown in [2, §9.5] that on *complete* lattices, the least solutions of T and T^{+} may be constructed from one another. The key idea of our axiomatization is first to represent γ by an equivalent system of equations T_{γ} , and then to *force* the simulating system T_{γ}^{+} to have a least solution, constructible from $\#_{\gamma}$, on the algebraic models for the logic.

More concretely, we present a simple algorithm that produces, when given as input a modal formula $\gamma(x)$ that is positive in x , a finite set of axioms and rules, of bounded size. Adding these axioms and rules to the basic modal logic \mathbf{K} , we obtain an axiom system $\mathbf{K}_{\#}^{+}(\Gamma)$, which

¹This rule is to be interpreted as stating that if some substitution instance $\gamma(\psi, \varphi_1, \dots, \varphi_n) \rightarrow \psi$ of the premiss is derivable in the system, then so is the corresponding substitution $\#_{\gamma}(\varphi_1, \dots, \varphi_n) \rightarrow \psi$ of the conclusion. Algebraically, it corresponds to the quasi-equation $\gamma(y, p_1, \dots, p_n) \leq y \rightarrow \#_{\gamma}(p_1, \dots, p_n) \leq y$ (or to the Horn formula obtained from this quasi-equation by universally quantifying over the variables y and p_1, \dots, p_n).

is finite if $\mathcal{L}_\#(\Gamma)$ has finitely many fixpoint connectives. Our second main result, Theorem 5.8, states that, for *any* flat fixpoint language,

$\mathbf{K}_\#^+(\Gamma)$ is a complete axiomatization for the validities in $\mathcal{L}_\#(\Gamma)$.

Let us briefly describe the strategy for obtaining the completeness theorem. We work in an algebraic setting for modal logic. Following a well known approach of algebraic logic, we treat formulas as terms over a signature whose function symbols are the logical connectives. Then, axioms correspond to equations and derivation rules to quasi-equations. The algebraic counterpart of the completeness theorem states that the equational theory of the “concrete” algebraic models that arise as complex algebras based on Kripke frames, is the same of the equational theory of the algebraic models of our axiomatization. To obtain such an algebraic completeness theorem, we study the Lindenbaum-Tarski algebras of our logic. Two properties of these structures turn out to be crucial: First, we prove that every Lindenbaum-Tarski algebra is *residuated*, or equivalently, that every diamond of the algebra has a right adjoint. And second, we show that the Lindenbaum-Tarski algebras are *constructive*: every fixpoint operation can be approximated as the join of its finite approximations. Then, we prove an algebraic representation theorem, Theorem 7.1, stating that every countable algebra with these two properties can be represented as a Kripke algebra, that is, as a subalgebra of the complex algebra of a Kripke frame. Putting these observations together, we obtain that the countable Lindenbaum-Tarski algebras have the same equational theory as the Kripke algebras, and this suffices to prove the algebraic version of the completeness theorem.

In order to prove these remarkable properties of the Lindenbaum-Tarski algebras, we switch to a *coalgebraic* reformulation of modal logic, based on the coalgebraic or cover modality ∇ . This connective ∇ takes a finite *set* α of formulas and returns a single formula $\nabla\alpha$, which can be seen as the following abbreviation:

$$\nabla\alpha = \Box(\bigvee \alpha) \wedge \bigwedge \diamond\alpha,$$

where $\diamond\alpha$ denotes the set $\{\diamond a \mid a \in \alpha\}$. The pattern of the definition of ∇ has surfaced in the literature on modal logic, in particular, as Fine’s normal forms [13]. The first explicit occurrences of this modality as a primitive connective, however, appeared not earlier than the 1990s, in the work of Barwise & Moss [3] and of Janin & Walukiewicz [18]. We call this connective “coalgebraic”, because of Moss’ observation [30], that its semantics allows a natural formulation in the framework of Universal Coalgebra, a recently emerging general mathematical theory of state-based evolving systems [32]. Moss’ insight paved the way for the transfer of many concepts, results and methods from modal logic to a far wider setting. As we will see, the main technical advantage of reconstructing modal logic on the basis of the cover modality is that this allows one to, if not completely eliminate conjunctions from the language, then at least *tame* them, so that they become completely harmless. This reduction principle, which lies at the basis of many constructions in the theory of the modal μ -calculus [18], has recently been investigated more deeply [31, 4], and generalized to a coalgebraic level of abstraction [24, 23].

We now briefly discuss how the present work contributes to the existing theory of fixpoint logics. Perhaps the first observation should be that our completeness results does not follow

from Walukiewicz' completeness result for the modal μ -calculus [40]: each *language* $\mathcal{L}_{\sharp}(\Gamma)$ may be a fragment of the full modal μ -calculus, but this does not imply that Kozen's axiomatization of the modal μ -calculus is a *conservative extension* of its restriction to such a language. In this respect, our results should be interpreted by saying that we add to Walukiewicz' theorem the observation that, modulo a better choice of axioms, proofs of validities in any given flat fragments of the modal μ -calculus can be carried out *inside* this fragment.

And second, while our methodology is based on earlier work [35] by the first author, which deals with the alternation-free fragment of the μ -calculus, we extend these results in a number of significant ways. In particular, the idea to use the subset construction of Arnold & Niwiński to *define an axiom system* for flat modal fixpoint logics, is novel. Furthermore, the representation theorem presented in Section 7 strengthens the main result of [35] (which applies to complete algebras only), to a completeness result for *Kripke frames*. With respect to [35], we also emphasize here the role of the coalgebraic cover modality ∇ in the common strategy for obtaining completeness. It is not only that some obscure results of [35] get a specific significance when understood from the coalgebraic perspective, but we also prove some new results on the cover modality ∇ itself, which may be of independent interest. And lastly, we can place an observation similar to the one we made with respect to Walukiewicz' result for the full modal μ -calculus: the results in [35] do not necessarily carry over to arbitrary fragments that are flat fixpoint logics. In fact, we were surprised to observe that it turns out to be possible to find a finitary complete axiomatization of the fixpoint connective \sharp_{γ} without explicitly introducing in the signature the least fixpoint of some other formula δ . This fact contrasts with the method proposed in [33] to equationally axiomatize the prefixpoints.

Finally, our proof method and, consequently, all of our results apply to the framework of *polymodal* logic, and we have formulated our main results accordingly. However, since much of the material presented here requires some rather involved notation, we will frequently choose to work in the setting of *monomodal* logic, in order to keep the text as readable as possible. In those cases where the transition to the polymodal setting is not routine, we always provide explicit details of this transition.

Overview of the paper. In Section 2 we first define flat modal fixpoint logics and then introduce our main tools: the coalgebraic cover modality ∇ , the algebraic approach to modal (fixpoint) logic, the order theoretic notion of a finitary \emptyset -adjoint, and the concept of a system of equations. Section 3 is devoted to the axiomatization $\mathbf{K}_{\sharp}^+(\Gamma)$ which we present as an algorithm producing the axiomatization given as input a set Γ of modal formulas. In Section 4 we give the proof of some algebraic results that relate fixpoints of different functions and that are at the core of the axiomatizations $\mathbf{K}_{\sharp}(\Gamma)$ and $\mathbf{K}_{\sharp}^+(\Gamma)$. With these results at hand, in Section 5 we formulate our two soundness and completeness results, and we sketch an overview of our algebraic proof method, introducing the Lindenbaum-Tarski algebras \mathbf{L} . In Section 6, we show that these Lindenbaum-Tarski algebras \mathbf{L} have a number of properties that make them resemble the power set algebra of a Kripke frame: we prove \mathbf{L} successively to be rigid, residuated, and constructive. Finally, in Section 7, we prove the above-mentioned representation theorem stating that every countable, residuated and constructive algebraic model of our language can be represented as a subalgebra of a powerset algebra of some Kripke frame.

2 Preliminaries

In this section we present some material that we consider background knowledge in the remainder of the paper. We first give a formal definition of the syntax and semantics of flat modal fixpoint logics. We then discuss the reformulation of modal logic in terms of the cover modalities ∇_i . Finally, we introduce modal \sharp -algebras as the key structures of the algebraic setting in which we shall prove our completeness result. For background in the algebraic perspective on modal logic, see [5, 38].

Flat modal fixpoint logic

The flat modal fixpoint logic of language $\mathcal{L}_\sharp(\Gamma)$ will be an extension of polymodal logic. Therefore we shall use I to denote the finite set of atomic actions indexing the modalities of polymodal logic. Next – and throughout this paper – we fix a set Γ of polymodal formulas $\gamma(x, \mathbf{p})$ where the variable x occurs only positively in γ and $\mathbf{p} = (p_1, \dots, p_n)$ is the ordered list of free variables in γ that are distinct from x . As usual x occurs only positively in γ if each occurrence of x appears under an even number of negations. Alternatively, we may decide to present the syntax of polymodal logic so that negation applies to propositional variables only, in which case x occurs positively if it occurs under no negation. The vector \mathbf{p} might be different for each γ , but we decided not to make this explicit in the syntax, in order not to clutter up notation.

First we give a formal definition of the language of flat modal fixpoint logics. Basically we add a new logical connective \sharp_γ to the language, for each $\gamma \in \Gamma$.

Definition 2.1. The set $\mathcal{L}_\sharp(\Gamma)$ of flat modal fixpoint formulas associated with Γ is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \diamond_i\varphi \mid \sharp_\gamma(\varphi),$$

where $p \in P$ is a propositional variable, i and γ range over I and Γ , respectively, and φ is a vector of previously generated formulas indexed by the vector \mathbf{p} .

We move on to the intended semantics of this language. A *labeled transition system* of type I , or equivalently a *Kripke model*, is a structure $\mathbb{S} = \langle S, \{R_i \mid i \in I\} \rangle$, where S is a set of states and, for each $i \in I$, $R_i \subseteq S \times S$ is a transition relation.

Definition 2.2. Given a Kripke model \mathbb{S} and a valuation $\mathbf{v} : P \longrightarrow \mathcal{P}(S)$ of propositional variables as subsets of states, we inductively define the semantics of flat modal fixpoint formulas as follows:

$$\begin{aligned} \|p\|_{\mathbf{v}} &= \mathbf{v}(p), \\ \|\neg\varphi\|_{\mathbf{v}} &= \overline{\|\varphi\|_{\mathbf{v}}}, \\ \|\varphi_1 \wedge \varphi_2\|_{\mathbf{v}} &= \|\varphi_1\|_{\mathbf{v}} \cap \|\varphi_2\|_{\mathbf{v}}, \\ \|\diamond_i\varphi\|_{\mathbf{v}} &= \{x \in S \mid \exists y \in S \text{ s.t. } xR_i y \text{ and } y \in \|\varphi\|_{\mathbf{v}}\}. \end{aligned}$$

In order to define $\|\#_\gamma(\varphi)\|_{\mathbf{v}}$, let x be a variable which is not free in φ and, for $Y \subseteq S$, let $(\mathbf{v}, x \rightarrow Y)$ be the valuation sending x to Y and every other variable y to $\mathbf{v}(y)$. We let

$$\|\#_\gamma(\varphi)\|_{\mathbf{v}} = \bigcap \{ Y \mid \|\gamma(x, \varphi)\|_{(\mathbf{v}, x \rightarrow Y)} \subseteq Y \}. \quad (1)$$

Observe that, by the Knaster-Tarski theorem [37], (1) just says that the interpretation of $\#_\gamma(\varphi)$ is the least fixpoint of the order preserving function sending Y to $\|\gamma(x, \varphi)\|_{(\mathbf{v}, x \rightarrow Y)}$.

The cover modality

We will frequently work in a reformulation of the modal language based on the *cover modality* ∇ . This connective, taking a *set* of formulas as their argument, can be defined in terms of the box and diamond operators:

$$\nabla\Phi := \Box \bigvee \Phi \wedge \bigwedge \Diamond\Phi,$$

where $\Diamond\Phi$ denotes the set $\{\Diamond\varphi \mid \varphi \in \Phi\}$. Conversely, the standard diamond and box modalities can be defined in terms of the cover modalities:

$$\Diamond\varphi \equiv \nabla\{\varphi, \top\}, \quad \Box\varphi \equiv \nabla\emptyset \vee \nabla\{\varphi\}. \quad (2)$$

It follows from these observations that we may equivalently base our modal language on ∇ as a primitive symbol.

What makes the cover modality ∇ so useful is that it satisfies two distributive laws:

$$\nabla(\Phi \cup \{\bigvee \Psi\}) = \bigvee_{\emptyset \subset \Psi' \subseteq \Psi} \nabla(\Phi \cup \Psi'), \quad (3)$$

and

$$\nabla\Phi \wedge \nabla\Psi \equiv \bigvee_{Z \in \Phi \bowtie \Psi} \nabla\{\varphi \wedge \psi \mid (\varphi, \psi) \in Z\}, \quad (4)$$

where $\Phi \bowtie \Psi$ denotes the set of relations $R \subseteq \Phi \times \Psi$ that are *full* in the sense that for all $\varphi \in \Phi$ there is a $\psi \in \Psi$ with $(\varphi, \psi) \in R$, and vice versa. The principle (3) clearly shows how the cover modality distributes over disjunctions, but we also call (4) a distributive law since it shows how conjunctions distribute over ∇ .

Remark 2.3. For more information on these distributive laws, the reader is referred to [31, 4], or to [23], where these principles are shown to hold in a very general coalgebraic context. Although to our knowledge it has never been made explicit in the literature on automata theory, equation (4) is in fact the key principle allowing the simulation of alternating automata by non-deterministic ones within the setting of μ -automata [18]. We refer to [17] for an algebraic, or to [24, 25] for a coalgebraic explanation of this.

As a straightforward application of these distributive laws (together with the standard distribution principles of conjunctions and disjunctions), every modal formula can be brought into a normal form, either by pushing conjunctions down to the leaves of the formula construction tree, or by pushing disjunctions up to the root, or by doing both. In order to make this observation more precise, we need some definitions, where we now switch to the polymodal setting in which we have a cover modality ∇_i for each atomic action i .

Definition 2.4. Let X be sets of propositional variables. Then we define the following sets of formulas:

1. $Lit(X)$ is the set $\{x, \neg x \mid x \in X\}$ of *literals* over X ,
2. $\mathcal{L}_\nabla(X)$ is the set of ∇ -formulas over X given by the following grammar:

$$\varphi ::= x \mid \neg x \mid \perp \mid \varphi \vee \varphi \mid \top \mid \varphi \wedge \varphi \mid \nabla_i \Phi$$

where $x \in X$, $i \in I$, and $\Phi \subseteq \mathcal{L}_\nabla(X)$.

3. $\mathcal{D}_\nabla(X)$ is the set of *disjunctive* formulas given by the following grammar:

$$\varphi ::= \perp \mid \varphi \vee \varphi \mid \bigwedge \Lambda \wedge \bigwedge_{j \in J} \nabla_j \Phi_j,$$

where $\Lambda \subseteq Lit(X)$, $J \subseteq I$, and $\Phi_j \subseteq \mathcal{D}_\nabla(X)$ for each $j \in J$. Note the restricted use of the conjunction symbol in disjunctive formulas: a conjunction of the form $\bigwedge \Lambda \wedge \bigwedge_{j \in J} \nabla_j \Phi_j$ will be called a *special conjunction*.

4. $\mathcal{P}_\nabla(X)$ is the set of *pure* ∇ -formulas in X , generated by the following grammar:

$$\varphi ::= \top \mid \bigwedge \Lambda \wedge \nabla \Phi,$$

where Λ is a set of literals, $\Phi = \{\Phi_i \mid i \in I\}$ is a vector such that, for each $i \in I$, Φ_i is a finite subset of $\mathcal{P}_\nabla(X)$, and $\nabla \Phi$ is defined by

$$\nabla \Phi := \bigwedge_{i \in I} \nabla_i \Phi_i. \quad (5)$$

Proposition 2.5. *Let X be a set of proposition letters. There are effective procedures*

1. *associating with each modal formula φ an equivalent ∇ -formula;*
2. *associating with each ∇ -formula $\varphi \in \mathcal{L}_\nabla(X)$ an equivalent disjunctive formula;*
3. *associating with each ∇ -formula $\varphi \in \mathcal{L}_\nabla(X)$ an equivalent disjunction of pure ∇ -formulas.*

Proof. Part 1 of the Proposition is proved by iteratively applying the equivalences of (2), whereas part 2 is obtained by using (4) as well as the distributive law of classical logic to push non special conjunctions to the leaves. For part 3, we first construct a formula $\varphi' \in \mathcal{D}_\nabla(X)$ which is equivalent to φ . Using the fact that \top is equivalent to $\nabla_i \{\top\} \vee \nabla_i \emptyset$, we can suppose that, within φ' , each special conjunction $\bigwedge \Lambda \wedge \bigwedge_{j \in J} \nabla_j \Phi_j$ is such that $J = I$. Then, we iteratively apply the distributive law (3) to φ' to push disjunctions up to the root. \square

Rewriting modal formulas into equivalent disjunctions of pure ∇ -formulas is not strictly necessary for our goals: we could work with disjunctive formulas only. However, we have chosen to consider this further simplification because it drastically improves the exposition of the next section.

Modal algebras and modal \sharp -algebras

We now move on to the algebraic perspective on flat modal fixpoint logic. As usual in algebraic logic, formulas of the logic are considered as terms over a signature whose function symbols are the logical connectives. Thus, from now on, the words “term” and “formula” will be considered as synonyms.

Before we turn to the definition of the key concept, that of a modal \sharp -algebra, we briefly recall the definition of a modal algebra.

Definition 2.6. Let $A = \langle A, \perp, \top, \neg, \wedge, \vee \rangle$ be a Boolean algebra. An operation $f : A \rightarrow A$ is called *additive* if $f(a \vee b) = fa \vee fb$, *normal* if $f\perp = \perp$, and an *operator* if it is both additive and normal. A *modal algebra* (of type I) is a structure $A = \langle A, \perp, \top, \neg, \wedge, \vee, \{\diamond_i^A \mid i \in I\} \rangle$, such that the interpretation \diamond_i^A of each action $i \in I$ is an operator on the Boolean algebra $\langle A, \perp, \top, \neg, \wedge, \vee \rangle$.

Equivalently, a modal algebra is a Boolean algebra expanded with operations that preserve all finite joins.

Let Z be a set of variables containing the free variables of a modal formula φ . If A is a modal algebra, then $\varphi^A : A^Z \rightarrow A$ denotes the term function of φ . Here A^Z is the set of Z -vectors (or Z -records), i.e. functions from the finite set Z to A . Recall that if $\text{card}(Z) = n$, then A^Z is isomorphic to the product of A with itself n times. Next, given $\gamma \in \Gamma$, let us list its free variables as usual, $\gamma = \gamma(x, p_1, \dots, p_n)$. Given a modal algebra A the *term function* of γ is of the form $\gamma^A : A \times A^n \rightarrow A$. Given a vector $\mathbf{b} = (b_1, \dots, b_n) \in A^n$, we let $\gamma_{\mathbf{b}}^A : A \rightarrow A$ denote the map given by

$$\gamma_{\mathbf{b}}^A(a) := \gamma^A(a, \mathbf{b}). \quad (6)$$

Definition 2.7. A *modal \sharp -algebra* is a modal algebra A endowed with an operation \sharp_{γ}^A for each $\gamma \in \Gamma$ such that for each \mathbf{b} , $\sharp_{\gamma}^A(\mathbf{b})$ is the least fixpoint of $\gamma_{\mathbf{b}}^A$ as defined in (6).

Note that modal \sharp -algebras are generally not complete; the definition simply *stipulates* that the least fixpoint exists, but there is no reason to assume that this fixpoint is reached by ordinal approximations.

Recall that $f : A \rightarrow B$ is a modal algebra morphism if the operations $\langle \perp, \top, \neg, \wedge, \{\diamond_i \mid i \in I\} \rangle$ are preserved by f . If A and B are also modal \sharp -algebras then f is a *modal \sharp -algebra morphism* if moreover each \sharp_{γ} , $\gamma \in \Gamma$, is preserved by f . This means that

$$f(\sharp_{\gamma}^A(\mathbf{v})) = \sharp_{\gamma}^B(f \circ \mathbf{v}),$$

for each $\mathbf{v} \in A^n$ and $\gamma \in \Gamma$. A \sharp -algebra morphism is an embedding if it is injective, and we say that A embeds into B if there exists an embedding $f : A \rightarrow B$.

In this paper we will be mainly interested in two kinds of modal \sharp -algebras: the “concrete” or “semantic” ones that encode a Kripke frame, and the “axiomatic” ones that can be seen as algebraic versions of the axiom system \mathbf{K}_{\sharp}^+ to be defined in the next section. We first consider the concrete ones.

Definition 2.8. Let $\mathbb{S} = \langle S, \{R_i \mid i \in I\} \rangle$ be a transition system. Define, for each $i \in I$, the operation $\langle R_i \rangle$ by putting, for each $X \subseteq S$, $\langle R_i \rangle X = \{y \in S \mid \exists x \in X \text{ s.t. } yR_i x\}$. The \sharp -complex algebra is given as the structure

$$\mathbb{S}^\sharp := \langle \mathcal{P}(S), \emptyset, S, \overline{(\cdot)}, \cup, \cap, \{\langle R_i \rangle \mid i \in I\} \rangle.$$

We will also call these structures *Kripke \sharp -algebras*.

Definition 2.9. Let $A = \langle A, \leq \rangle$ be a partial order with least element \perp , and let $f : A \rightarrow A$ be an order-preserving map on A . For $k \in \omega$ and $a \in A$, we inductively define $f^k a$ by putting $f^0 a := a$ and $f^{k+1} a := f(f^k a)$. If f has a least fixpoint $\mu.f$, then we say that this least fixpoint is *constructive* if $\mu.f = \bigvee_{k \in \omega} f^k(\perp)$. A modal \sharp -algebra is called *constructive* if $\sharp_\gamma^A(\mathbf{b})$ is a constructive least fixpoint, for each $\gamma \in \Gamma$ and each \mathbf{b} in A .

Remark 2.10. Our terminology slightly deviates from that in [35], where the least fixpoint of an order-preserving map on a partial order is called constructive if it is equal to the join of all its ordinal approximations, not just of the ω first ones.

\emptyset -adjoints and fixpoints

We now recall the well known concept of adjointness, and briefly discuss its generalization, \emptyset -adjointness.

Definition 2.11. Let $A = (A, \leq)$ and $B = (B, \leq)$ be two partial orders. Suppose that $f : A \rightarrow B$ and $g : B \rightarrow A$ are order-preserving maps such that

$$fa \leq b \text{ iff } a \leq gb, \quad (7)$$

for all $a \in A$ and $b \in B$. Then we call (f, g) an *adjoint pair*, and say that f is the *left adjoint* of, or *residuated* by, g , and that g is the *right adjoint*, or *residual*, of f . We say that f is an \emptyset -adjoint if it satisfies the weaker property that for every $b \in B$ there is a finite set $G_f(b) \subseteq A$ such that

$$fa \leq b \text{ iff } a \leq a' \text{ for some } a' \in G_f(b),$$

for all $a \in A$ and $b \in B$.

Remark 2.12. The terminology ‘ \emptyset -adjoint’ can be explained as follows. Let \mathcal{T} be a functor on the category of partial orders (with order-preserving maps as arrows). Call a morphism $f : (A, \leq) \rightarrow (B, \leq)$ a *left \mathcal{T} -adjoint* if the map $\mathcal{T}f : \mathcal{T}(A, \leq) \rightarrow \mathcal{T}(B, \leq)$ has a right adjoint $G : \mathcal{T}(A, \leq) \rightarrow \mathcal{T}(B, \leq)$ in the sense of (7) above. Let now \mathcal{T} be the functor \mathcal{O}_f defined as follows. On objects, \mathcal{O}_f maps a partial order (A, \leq) to the set $\mathcal{O}_f(A, \leq)$ of finitely generated downsets of (A, \leq) , ordered by inclusion. Alternatively, $\mathcal{O}_f(A, \leq)$ is the free join-semilattice generated by (A, \leq) . To become a functor, \mathcal{O}_f takes an arrow $f : (A, \leq) \rightarrow (B, \leq)$ to the function $\mathcal{O}_f(f)$ that maps a subset $X \in \mathcal{O}_f(A)$ to the set of points that are below some element of the direct image $f(X)$.

We leave it as an exercise for the reader to verify that an order-preserving map f is an \emptyset -adjoint, in the sense of Definition 2.11 iff it is a left \mathcal{O}_f -adjoint in the sense just described. We write \emptyset -adjoint rather than left \mathcal{O}_f -adjoint in order to keep our notation simple.

Finally, observe that to define adjoints, \mathcal{T} -adjoints, and \emptyset -adjoints, we do not need the antisymmetric law of partial order, we can define these notions for quasi orders.

It is well known that left adjoint maps preserve all existing joins of a poset. Similarly, one may prove that \emptyset -adjoints preserve all existing joins of *directed* sets.

\emptyset -adjoints are relevant for the theory of least fixpoints because of the following. If $f : A^n \rightarrow A$ is an \emptyset -adjoint, say that $V \subseteq A$ is f -closed if $y \in V$ and $\mathbf{a} = (a_1, \dots, a_n) \in G_f(y)$ implies $a_i \in V$ for $i = 1, \dots, n$. If \mathcal{F} is a family of \emptyset -adjoints of the form $f : A^n \rightarrow A$, say that V is \mathcal{F} -closed if it is f -closed for each $f \in \mathcal{F}$.

Definition 2.13. A family of \emptyset -adjoints $\mathcal{F} = \{f_i : A^{n_i} \rightarrow A \mid i \in I\}$ is said to be *finitary* if, for each $x \in A$, the least set \mathcal{F} -closed set containing x is finite. The \emptyset -adjoint $f^n : A \rightarrow A$ is finitary if the singleton $\{f\}$ is finitary.

Clearly, if f belongs to a finitary family, then it is finitary.

Proposition 2.14. *If $f : A \rightarrow A$ is a finitary \emptyset -adjoint, then its least prefixpoint, whenever it exists, is constructive.*

See [35, Proposition 6.6] for a proof of the Proposition.

The next Proposition collects the main properties of finitary families of \emptyset -adjoints. Roughly speaking, these properties assert that finitary families may be supposed to be closed under composition, joining, and tupling.

Proposition 2.15. *Let \mathcal{F} be a finitary family of \emptyset -adjoints on a modal algebra A . Suppose also that $f, g \in \mathcal{F}$. Then \mathcal{G} also is a finitary family of \emptyset -adjoints, whenever*

1. $\mathcal{G} \subseteq \mathcal{F}$,
2. $\mathcal{G} = \mathcal{F} \cup \{h\}$, $f : A \times A^Z \rightarrow A$, $g : A^Y \rightarrow A$, and $h = f \circ (g \times A^Z) : A^Y \times A^Z \rightarrow A$,
3. $\mathcal{G} = \mathcal{F} \cup \{h\}$, $f, g : A^Z \rightarrow A$, and $h = f \vee g$,
4. $\mathcal{G} = \{F : A^Z \rightarrow A^Z\}$ and $\{\pi_z \circ F : A^Z \rightarrow A \mid z \in Z\} \subseteq \mathcal{F}$.

Proof. Part 1 of the statement is obvious. For the parts 2 and 4, we invite the reader to consult [35, Lemmas 6.10 to 6.12]. For Part 3, observe that

$$G_{f \vee g}(d) = G_f(d) \wedge G_g(d),$$

where $C \wedge D = \{\mathbf{v} \wedge \mathbf{u} \mid \mathbf{v} \in C \text{ and } \mathbf{u} \in D\}$. Thus, if $v_0 \in A$ and V is a finite \mathcal{F} -closed set with $v_0 \in V$, then V_\wedge , the closure of V under meets, is a finite \mathcal{G} -closed set with $v_0 \in V_\wedge$. QED

Systems of equations

Definition 2.16. A *modal system* or *system of equations* is a pair $T = \langle Z, \{t_z\}_{z \in Z} \rangle$ where Z is a finite set of variables and $t_z \in \mathcal{L}_\nabla(Z \cup P)$ for each $z \in Z$. Such a modal system is *pointed* if it comes with a specified variable $z_0 \in Z$.

Given a modal system T and a modal algebra A , there exists a unique function $T^A : A^Z \times A^P \rightarrow A^Z$ such that, for each projection $\pi_z : A^Z \rightarrow A$, $\pi_z \circ T^A = t_z^A$. We shall say that T^A is the interpretation of T in A . Whenever it exists, we shall denote the least fixpoint of T^A by $\mu_Z.T^A : A^P \rightarrow A^Z$.

In this paper we will be interested in modal systems where every term is in a special syntactic shape.

Definition 2.17. In the monomodal setting, a term $t \in \mathcal{L}_\nabla(Z \cup P)$ is *semi-simple* if it is a disjunction of terms of the form $\Lambda \wedge \nabla\Phi$, where Λ is a set of P -literals, and each $\varphi \in \Phi$ is a finite conjunction of variables in Z (where \top is the empty conjunction). For such a term to be *simple*, we require that each $\varphi \in \Phi$ belongs to the set $Z \cup \{\top\}$. In the polymodal setting, a term t is semi-simple (simple) if it is a disjunction of terms of the form $\Lambda \wedge \bigwedge_{j \in J} \nabla_j \Phi_j$, where $J \subseteq I$ and each of the formulas in $\bigcup_j \Phi_j$ satisfies the respective above-mentioned condition.

A modal system $T = \langle Z, \{t_z\}_{z \in Z} \rangle$ is *semi-simple* (*simple*, respectively), if every term t_z is semi-simple (simple, respectively).

3 The axiomatization $\mathbf{K}_\#^+(\Gamma)$

The axiom system $\mathbf{K}_\#^+(\Gamma)$ that we will define in this section adds, for each $\gamma \in \Gamma$, a number of axioms and derivation rules to the basic (poly-)modal logic \mathbf{K} . We obtain these axioms and rules effectively, via some systems of equations that we will associate with γ . Here is a summary of the procedure.

0. *Preprocess*, rewriting $\gamma(x)$ as a guarded disjunction of special pure ∇ -formulas.
1. *Represent* each such γ by a *semi-simple* system of equations T_γ .
2. *Simulate* T_γ by a *simple* system of equations T_γ^+ .
3. *Read off* the axiomatization for $\#_\gamma$ from T_γ^+ .

The aim of this section is to define and discuss this procedure in full detail — readers who only want to look at the definition of the axiom system can proceed directly via the Definitions 3.10, 3.16 and 3.22. For the sake of readability we work mainly in the monomodal framework.

Before carrying on, let us fix some notation to be used throughout this section. We shall use the capital letters X, Y, Z to denote sets of fixpoint variables. On the other hand, P will denote a set of proposition letters not containing any of these fixpoint variables. If $\tau \in \mathcal{L}_\nabla(X \cup P)$ and $\{\sigma_y \mid y \in Y\} \subseteq \mathcal{L}_\nabla(X)$ is a collection of terms indexed by $Y \subseteq X$, then we shall denote by σ such a collection, and by $\tau[\sigma/\mathbf{y}]$ the result of simultaneously substituting every variable $y \in Y$ with the term σ_y .

Preprocessing γ

Fix a modal formula $\gamma(x)$ in which the variable x occurs only positively. First of all, for our purposes we may assume that each occurrence of x is *guarded* in γ , that is, within the scope

of some modal operator. In the theory of fixpoint logics it is well-known that this assumption is without loss of generality, see for example [40, Proposition 2]. In order to give a quick justification, recall that our goal is to axiomatize the least prefixpoint of $\gamma(x)$. If x is not guarded in γ , then we can find terms γ_1, γ_2 , with x guarded in both γ_1 and γ_2 , and such that the equation

$$\gamma(x, \mathbf{p}) = (x \wedge \gamma_1(x, \mathbf{p})) \vee \gamma_2(x, \mathbf{p}),$$

holds on every modal algebra. It is easily seen that, on every modal algebra, γ and γ_2 have the same set of prefixpoints. Thus, instead of axiomatizing \sharp_γ , we can equivalently axiomatize \sharp_{γ_2} .

Second, given the results mentioned in the previous section, we may assume that γ is a disjunction of pure ∇ -formulas (cf. Proposition 2.5). However, given the special role of the variable x , it will be convenient for us to modify our notation accordingly. We introduce the following abbreviation:

$$\nabla_\Lambda \Phi := \bigwedge \Lambda \wedge \nabla \Phi,$$

in the case that $\Lambda \subseteq Lit(X)$ and x does not occur in Λ .

Definition 3.1. Given a set P of proposition letters and a variable $x \notin P$, we define the set of pure ∇x -formulas in P by the following grammar:

$$\varphi ::= \top \mid x \mid \nabla_\Lambda \Phi \mid x \wedge \nabla_\Lambda \Phi, \quad (8)$$

where $\Lambda \subseteq Lit(P)$, and Φ is a set of pure ∇x -formulas in P .

Remark 3.2. Recall from equation 5 that, in the polymodal setting, $\nabla \Phi$ denotes the formula $\bigwedge_{i \in I} \nabla_i \Phi_i$, where Φ is the vector $\{\Phi_i \mid i \in I\}$. Now we can define the set of ∇x -formulas in P , in the polymodal setting, by the following grammar:

$$\varphi ::= \top \mid x \mid \nabla_\Lambda \Phi \mid x \wedge \nabla_\Lambda \Phi.$$

Then basically, the algorithm for obtaining the axiomatization in the polymodal case works the same as in the monomodal case, with the polymodal nabla-operator ∇ replacing the monomodal ∇ .

Convention 3.3. In concrete examples we will denote the set Λ in ∇_Λ as a *list* rather than as a set, and write \bar{p} rather than $\neg p$. For instance we will write $\nabla_{\bar{p}\bar{q}} \Phi$ instead of $\nabla_{\{p, \neg q\}} \Phi$. Furthermore, we will write $\nabla \Phi$ instead of $\nabla_\emptyset \Phi$.

Lemma 3.4. *Every modal formula $\gamma \in \mathcal{L}_{\nabla}(P \cup \{x\})$ in which the variable x only occurs positively can be effectively rewritten as an equivalent disjunction γ' of pure ∇x -formulas in P . Furthermore, if x is guarded in γ then x is guarded in γ' as well.*

Proof. In Proposition 2.5 we saw that every modal formula γ can be equivalently rewritten as a disjunction γ' of pure ∇ -formulas. If x occurs only positively in γ , then this formula will have no subformulas of the form $\bigwedge \Lambda \wedge \nabla \Phi$ with $\neg x \in \Lambda$. From this the lemma is immediate. QED

Example 3.5. Consider the formula $(p \wedge \Box x) \vee (\neg p \wedge \Diamond(x \wedge \Diamond x))$. Rewriting this as a disjunction of pure ∇x -formulas, we obtain

$$\gamma(x) = \nabla_p \emptyset \vee \nabla_p \{x\} \vee \nabla_{\bar{p}} \{ \top, x \wedge \nabla \{ \top, x \} \}. \quad (9)$$

Step 1: from formulas to semi-simple systems of equations

In the first step of the procedure, we represent a formula γ as a semi-simple system of equations T_γ . Fix a modal formula $\gamma(x)$ in which the variable x only occurs positively. Without loss of generality we may assume that γ is a disjunction of pure ∇x -formulas, and guarded in x . Roughly speaking, to obtain the modal system T_γ we cut up the formula γ in layers, step by step peeling off its modalities and introducing new variables for (some of) γ 's subformulas of the form $\nabla_\Lambda \Phi$.

Definition 3.6. Let $\gamma(x) \in \mathcal{L}_\nabla(P \cup \{x\})$ be a disjunction of pure ∇x -formulas, and guarded in x . We define SC_γ , the set of *special conjunctions* in γ , as the set of subformulas of γ of the form $\nabla_\Lambda \Phi$. SC'_γ is the set of special conjunctions that occur in the scope of some ∇ -formula. Furthermore, we define $RSF_\gamma := \{\gamma\} \cup SC'_\gamma$ as the set of *relevant subformulas* of γ .

To see the difference between the sets SC'_γ and SC_γ , observe that γ itself is a disjunction of special conjunctions. These disjuncts are elements of SC_γ , but we only put them in SC'_γ if they occur as subformulas of γ deeper in the formula tree as well.

Example 3.7. With γ the formula given by (9), we find that SC_γ consists of the four formulas

$$\begin{aligned} \psi_1 &= \nabla_p \emptyset, \\ \psi_2 &= \nabla_p \{x\}, \\ \psi_3 &= \nabla_{\bar{p}} \{ \top, x \wedge \nabla \{ \top, x \} \}, \\ \psi_4 &= \nabla \{ \top, x \}. \end{aligned}$$

Of these, only ψ_4 makes it into SC'_γ , so $RSF_\gamma = \{\gamma, \psi_4\}$.

The system of equations T_γ will be based on a set of variables that is in one-to-one correspondence with the set of relevant formulas.

Definition 3.8. Let $\gamma(x) \in \mathcal{L}_\nabla(P \cup \{x\})$ be a disjunction of pure ∇x -formulas, and guarded in x . Let

$$Z = \{z_\psi \mid \psi \in RSF_\gamma\}$$

be a set of fresh variables (in one-to-one correspondence with the set RSF_γ), and let $[\psi/z]$ be the natural substitution replacing each variable z_ψ with the formula ψ .

The key observation in the definition of the modal system T_γ is that every disjunction of formulas in SC_γ can be seen as the $[\psi/z]$ -substitution instance of a semi-simple formula $\widehat{\psi}$. For instance, in Example 3.7, writing

$$\widehat{\psi}_3 = \nabla_{\bar{p}} \{ \top, x \wedge z_{\psi_4} \},$$

we have that $\psi_3 = \widehat{\psi}_3[\psi_4/z_{\psi_4}]$.

Lemma 3.9. For every formula $\psi \in RSF_\gamma$ there is a semi-simple formula $\widehat{\psi}$ such that $\psi = \widehat{\psi}[\psi/z]$.

Proof. Given a special conjunction $\nabla_\Lambda \Phi$ in γ , each $\varphi \in \Phi$ has one of the forms \top, x, ψ , or $x \wedge \psi$, where ψ is again a special conjunction. Let $\widehat{\nabla_\Lambda \Phi}$ be the formula we obtain by replacing Φ 's elements of the form ψ and $x \wedge \psi$ with z_ψ and $x \wedge z_\psi$, respectively. It is immediate that $\nabla_\Lambda \Phi = \widehat{\nabla_\Lambda \Phi}[\psi/z]$. This takes care of the formulas $\psi \in SC'_\gamma$, while for γ , which can be written as a disjunction $\bigvee_i \varphi_i$ of special conjunctions, we can simply take the formula $\widehat{\gamma} := \bigvee_i \widehat{\varphi}_i$. It is easy to see that the obtained formulas are semi-simple. QED

Definition 3.10. Let $\gamma(x) \in \mathcal{L}_{\nabla}(P \cup \{x\})$ be a disjunction of pure ∇/x -formulas, and guarded in x . For $z = z_\psi \in Z$, we write $\rho_z := \psi$, and let τ_z denote the term $\rho_z[z_\gamma/x]$. We call the modal system

$$T_\gamma := \langle Z, \{\tau_z \mid z \in Z\} \rangle$$

the *system representation* of γ . T_γ is pointed by the variable z_γ .

The reader will have no difficulties verifying that T_γ is a semi-simple systems of equations.

Example 3.11. For the formula γ of the Examples 3.5/3.7, we obtain (writing z_i rather than z_{ψ_i}) the following system T_γ . As its variables it has the set $\{z_\gamma, z_4\}$, and its equations are the following:

$$\begin{aligned} z_\gamma &= \nabla_p \emptyset \vee \nabla_p \{z_\gamma\} \vee \nabla_p \{\top, z_\gamma \wedge z_4\} \\ z_4 &= \nabla \{\top, z_\gamma\}. \end{aligned}$$

We call the modal system T_γ a *representation* of the formula γ because the least fixpoints of T_γ and γ are mutually expressible — for the precise formulation of this statement we refer to Proposition 4.1 below. Here we just mention the key observation underlying this proposition, which relates the (parametrized) fixpoints of T_γ to those of γ , as follows.

Proposition 3.12. Let γ be a modal formula in which the variable x only occurs positively, let A be a modal algebra, and $\mathbf{v} \in A^P$ a sequence of parameters in A .

1. If $a \in A$ is a fixpoint of $\gamma_{\mathbf{v}}^A$, then the vector $\{\psi^A(a, \mathbf{v}) \mid \psi \in RSF_\gamma\}$ is a fixpoint of $(T_\gamma^A)_{\mathbf{v}}$.
2. If $\{b_\psi \mid \psi \in RSF_\gamma\}$ is a fixpoint of $(T_\gamma^A)_{\mathbf{v}}$, then $b_\gamma \in A$ is a fixpoint of $\gamma_{\mathbf{v}}^A$.

Proof. Immediate by the definitions. QED

Since our main aim is to represent γ by a *simple* set of equations, formulas γ for which T_γ itself is already simple, are clearly of interest. We shall introduce in Section 5 classes of formulas, called *untied* and *harmless*, that have this property. If every formula $\gamma \in \Gamma$ belongs to those classes, then we can prove that $\mathbf{K}_\#(\Gamma)$ is already a complete and sound axiom system.

Step 2: from semi-simple systems of equations to simple ones

The second step of our procedure is based on the *subset construction* of Arnold & Niwiński [2]. The idea behind this construction is that, under some conditions, one may eliminate conjunctions from a system of equations T through *simulating* it by another system, T^+ . Roughly, the idea of the construction is that the variables of the system T^+ correspond to the conjunctions of the non-empty sets of variables of the system T .

Convention 3.13. Given the set of variables Z , we let $Y = \{y_S \mid S \in \mathcal{P}_+(Z)\}$ be a set of new variables in bijection with $\mathcal{P}_+(Z)$, the set of non empty subsets of Z . For $S \in \mathcal{P}_+(Z)$, we denote by z_S the term $\bigwedge S$, and let $[z/\mathbf{y}]$ denote the substitution which replaces each variable $y = y_S \in Y$ with the term z_S .

The following lemma is the heart of the simulation construction.

Proposition 3.14. *Let $\{\tau_i \mid i \in I\}$ be a finite collection of semi-simple terms in Z .*

1. *There is a semi-simple term τ in Z which is equivalent to $\bigwedge_{i \in I} \tau_i$.*
2. *There is a simple term σ in Y , such that the term $\sigma[z/\mathbf{y}]$ is equivalent to $\bigwedge_{i \in I} \tau_i$.*

Proof. We give the proof in the monomodal setting. The first part of the lemma follows easily from successive applications of the distributive law (4) for the cover modality. Obviously it suffices to prove that the conjunction of two semi-simple terms $\bigwedge \Lambda \wedge \nabla \Phi$ and $\bigwedge \Lambda' \wedge \nabla \Phi'$ is semi-simple. But by (4), and the distributive law of classical propositional logic, this conjunction is equivalent to some formula $\bigwedge (\Lambda \cup \Lambda') \wedge \nabla \Psi$, where each formula $\psi \in \Psi$ is of the form $\varphi \wedge \varphi'$, with $\varphi \in \Phi$ and $\varphi' \in \Phi'$, and thus itself a finite conjunction of variables in Z . In other words, the formula $\bigwedge (\Lambda \cup \Lambda') \wedge \nabla \Psi$ is equivalent to a semi-simple formula.

The second part of the proposition is an almost immediate consequence of the first, by the observation that with every semi-simple term τ we may associate a simple term σ such that τ is equivalent to the term $\sigma[z/\mathbf{y}]$. The term σ is obtained from τ simply by replacing, for each disjunct $\Lambda \wedge \nabla \Phi$, each formula $\bigwedge S \in \Phi$ (with $S \neq \emptyset$) by the variable y_S . QED

Remark 3.15. It should be immediate to see how modify the above proof for the setting of polymodal logic. Indeed, recall first from Remark 3.2 the definition of the polymodal ∇ . Trivially, one has

$$\bigwedge \Lambda \wedge \nabla \Phi \wedge \bigwedge \Lambda' \wedge \nabla \Psi = \bigwedge (\Lambda \cup \Lambda') \wedge \bigwedge_{i \in I} \nabla_i \Phi_i \wedge \nabla_i \Psi_i,$$

so that, by applying first the laws (4) for each ∇_i , and then the distributive law of classical propositional logic, a fundamental distributive law for the polymodal ∇ may also be derived.

Definition 3.16. Let $T = \langle Z, \{\tau_z \mid z \in Z\} \rangle$ be a semi-simple modal system. For any $y \in Y$, writing $y = y_S$ with $S \in \mathcal{P}_+(Z)$, let σ_y be the simple term corresponding to the conjunction $\bigwedge_{z \in S} \tau_z$, as provided by Proposition 3.14. The *simulation* of T is defined as the system of equations

$$T^+ := \langle Y, \{\sigma_y \mid y \in Y\} \rangle.$$

Example 3.17. Continuing Example 3.11, we may write

$$\begin{aligned}
z_\gamma \wedge z_4 &= (\nabla_p \emptyset \wedge \nabla \{\top, z_\gamma\}) \vee (\nabla_p \{z_\gamma\} \wedge \nabla \{\top, z_\gamma\}) \vee (\nabla_{\bar{p}} \{\top, z_\gamma \wedge z_4\} \wedge \nabla \{\top, z_\gamma\}) \\
&= \perp \vee \nabla_p \{z_\gamma\} \vee \nabla_{\bar{p}} \{\top, z_\gamma \wedge z_4, z_\gamma\} \\
&= \nabla_p \{z_\gamma\} \vee \nabla_{\bar{p}} \{\top, z_\gamma \wedge z_4, z_\gamma\},
\end{aligned}$$

where we have used some “ ∇ -arithmetic” to simplify the outcome.

Thus we obtain the following as the system T_γ^+ :

$$\begin{aligned}
y_\gamma &= \nabla_p \emptyset \vee \nabla_p \{y_\gamma\} \vee \nabla_{\bar{p}} \{\top, y_{\gamma 4}\} \\
y_4 &= \nabla \{\top, y_\gamma\} \\
y_{\gamma 4} &= \nabla_p \{y_\gamma\} \vee \nabla_{\bar{p}} \{\top, y_{\gamma 4}, y_\gamma\}.
\end{aligned}$$

Here we write y_γ instead of $y_{\{\gamma\}}$, etc.

For a more elaborate example, consider the following.

Example 3.18. Let T be the semi-simple modal system given by

$$\begin{cases}
z_1 &= \nabla_{pq} \{z_1 \wedge z_2, z_1 \wedge z_3\} \vee \nabla_{p\bar{q}} \{z_2\} \\
z_2 &= \nabla_p \{z_1, z_3\} \\
z_3 &= \nabla \{z_2 \wedge z_3\}.
\end{cases}$$

Using the distributive laws for ∇ and some further ∇ -arithmetic, one may derive that

$$\begin{aligned}
z_1 \wedge z_2 &= \nabla_{pq} \{z_1 \wedge z_2, z_1 \wedge z_3\} \vee \nabla_{pq} \{z_1 \wedge z_3, z_1 \wedge z_2 \wedge z_3\} \\
&\quad \vee \nabla_{pq} \{z_1 \wedge z_2, z_1 \wedge z_3, z_1 \wedge z_2 \wedge z_3\} \\
z_1 \wedge z_3 &= \nabla_{pq} \{z_1 \wedge z_2 \wedge z_3\} \vee \nabla_{p\bar{q}} \{z_2 \wedge z_3\} \\
z_2 \wedge z_3 &= \nabla_p \{z_2 \wedge z_3, z_1 \wedge z_2 \wedge z_3\} \\
z_1 \wedge z_2 \wedge z_3 &= \nabla_{pq} \{z_1 \wedge z_2 \wedge z_3\}.
\end{aligned}$$

From this it is easy to see that the simulation T^+ is given by

$$\begin{cases}
y_1 &= \nabla_{pq} \{y_{12}, y_{13}\} \vee \nabla_{p\bar{q}} \{y_2\} \\
y_2 &= \nabla_p \{y_1, y_3\} \\
y_3 &= \nabla \{y_{23}\} \\
y_{12} &= \nabla_{pq} \{y_{12}, y_{13}\} \vee \nabla_{pq} \{y_{13}, y_{123}\} \vee \nabla_{p\bar{q}} \{y_{13}, y_{123}\} \\
y_{13} &= \nabla_{pq} \{y_{123}\} \vee \nabla_{p\bar{q}} \{y_{23}\} \\
y_{23} &= \nabla_p \{y_{23}, y_{123}\} \\
y_{123} &= \nabla_{pq} \{y_{123}\},
\end{cases}$$

where we write y_{12} for $y_{\{1,2\}}$, etc.

The relation between the modal systems T and T^+ is perhaps clarified by a diagram. Let, for some modal algebra A , $\iota^A : A^Z \rightarrow A^Y$ be given by

$$\iota^A(\mathbf{a})(y_S) = \bigwedge_{z \in S} a_z. \tag{10}$$

Then Proposition 3.14(2) maybe understood as stating that, given a semi-simple system T , there exists a simple system T^+ such that, for every modal algebra A and every parameter $\mathbf{v} \in A^P$, the diagram

$$\begin{array}{ccc}
 A^Z & \xrightarrow{T_{\mathbf{v}}^A} & A^Z \\
 \downarrow \iota^A & & \downarrow \iota^A \\
 A^Y & \xrightarrow{(T^+)_{\mathbf{v}}^A} & A^Y
 \end{array} \tag{11}$$

commutes.

On *complete* modal algebras, the modal systems T and T^+ are equivalent in the sense that the respective least fixpoints are mutually definable — this is in fact the point behind the introduction of T^+ in [2]. In general however, the relation between T and T^+ seems to be less tight than that between the formula γ (or rather, the system $\langle \{x\}, \{\gamma\} \rangle$) and the system T_γ . In the next section we discuss this relation in more detail: here we confine ourselves to the following basic observation concerning fixpoints of T and T_γ .

Proposition 3.19. *Let T be a semi-simple modal system, let A be a modal algebra, and $\mathbf{v} \in A^P$ a sequence of parameters in A . If $\{a_z \mid z \in Z\}$ is a fixpoint of $T_{\mathbf{v}}$, then $\{\bigwedge \{a_z \mid z \in S\} \mid S \in \mathcal{P}_+(Z)\}$ is a fixpoint of $T_{\mathbf{v}}^+$.*

Proof. Immediate by (11) and the definitions. QED

Step 3: read off the axiomatization

We are now ready to define the axioms and derivation rules that we associate with a formula $\gamma(x, \mathbf{p})$ in which the variable x occurs only positively. As we will see, these axioms and rules can be easily read off from the simple modal system T_γ^+ that we obtained in the previous step of the procedure. Before going into the syntactic details, let us first take an algebraic perspective.

Let A be a modal \sharp -algebra, and let $\mathbf{v} \in A^P$ be a sequence of parameters in A . Since $\sharp\mathbf{v}$ is the least fixpoint of the map $\gamma_{\mathbf{v}}^A : A \rightarrow A$, it follows from Proposition 4.1 below that the vector

$$\{\psi^A(\sharp\mathbf{v}, \mathbf{v}) \mid \psi \in RSF_\gamma\} \tag{12}$$

is the least fixpoint of $(T_\gamma^A)_{\mathbf{v}}$. In order to arrive at a succinct presentation of our axiom system, it will be convenient to think of the coordinate $\gamma^A(\sharp\mathbf{v}, \mathbf{v})$ of (12) (that is, the case where $\psi = \gamma \in RSF_\gamma$), as the fixpoint $\sharp\mathbf{v}$ itself — this is allowed since A is a modal \sharp -algebra. For this purpose we introduce the following notation, using the one-to-one correspondence between the sets Z and RSF_γ :

$$\chi_z := \begin{cases} x & \text{if } \psi_z = \gamma, \\ \psi_z & \text{otherwise.} \end{cases}$$

We may conclude that on any modal \sharp -algebra A , the set

$$\{\chi_z^A(\sharp \mathbf{v}, \mathbf{v}) \mid z \in Z\} \quad (13)$$

is the least fixpoint of $(T_\gamma^A)_\mathbf{v}$. Then on the basis of Proposition 3.19, the set

$$\left\{ \bigwedge_{z \in S} \chi_z^A(\sharp \mathbf{v}, \mathbf{v}) \mid S \in \mathcal{P}_+(Z) \right\} \quad (14)$$

is *some* fixpoint of $(T_\gamma^+)_\mathbf{v}^A$. In case A is a complete algebra, the results of Arnold & Niwiński [2, §9] imply that (14) is in fact the *least* fixpoint of $(T_\gamma^+)_\mathbf{v}^A$. For a general \sharp -algebra, however, we have no justification for drawing this conclusion. This means that the following is a meaningful definition.

Definition 3.20. A modal \sharp -algebra A is called *regular* if for each $\gamma \in \Gamma$ and each $\mathbf{v} \in A^P$, the set (14) is the least fixpoint of $(T_\gamma^+)_\mathbf{v}^A$.

We can now give an intuitive introduction of the axiom system $\mathbf{K}_\sharp^+(\Gamma)$ by saying that it expresses the regularity of modal \sharp -algebras. In other words, our axiomatization requires that the set (14) is the least fixpoint of $(T_\gamma^+)_\mathbf{v}^A$. Thus the above-mentioned result by Arnold & Niwiński will imply the *soundness* of the axiomatization.

Example 3.21. Continuing Example 3.17, we find that $\chi_0 = x$ and $\chi_4 = \nabla\{\top, x\}$. Our axiomatization will express that, for any formula φ (corresponding to the sequence \mathbf{v} of parameters), the vector

$$\begin{pmatrix} \sharp\varphi \\ \nabla\{\top, \sharp\varphi\} \\ \sharp\varphi \wedge \nabla\{\top, \sharp\varphi\} \end{pmatrix} = \begin{pmatrix} \chi_0[\sharp p/x][\varphi/p] \\ \chi_4[\sharp p/x][\varphi/p] \\ (\chi_0 \wedge \chi_4)[\sharp p/x][\varphi/p] \end{pmatrix}$$

is the least fixpoint of the system $T_\gamma^+[\varphi/p]$. It suffices for our axiom system to express this for the proposition letter p : a uniform substitution will then take care of the parameter φ (see footnote 1 on how we formulate and interpret derivation rules). Recall that the following σ_γ , σ_4 and σ_{γ_4} are the terms of the system T_γ^+ :

$$\begin{aligned} \sigma_\gamma &= \nabla_p \emptyset \vee \nabla_p \{y_\gamma\} \vee \nabla_{\bar{p}} \{\top, y_{\gamma_4}\} \\ \sigma_4 &= \nabla \{\top, y_\gamma\} \\ \sigma_{\gamma_4} &= \nabla_p \{y_\gamma\} \vee \nabla_{\bar{p}} \{\top, y_{\gamma_4}, y_\gamma\}. \end{aligned}$$

Thus our axiomatization will contain the axioms

$$\nabla_p \emptyset \vee \nabla_p \{\sharp p\} \vee \nabla_{\bar{p}} \{\top, p \wedge \nabla\{\top, \sharp p\}\} \rightarrow \sharp p \quad (A_\gamma)$$

$$\nabla\{\top, \sharp p\} \rightarrow \nabla\{\top, p\} \quad (A_4)$$

$$\nabla_p \{\sharp p\} \vee \nabla_{\bar{p}} \{\top, p \wedge \nabla\{\top, \sharp p, \sharp p\}\} \rightarrow p \wedge \nabla\{\top, \sharp p\} \quad (A_{\gamma_4})$$

stating that $\begin{pmatrix} \#p \\ \nabla\{\top, \#p\} \\ \#p \wedge \nabla\{\top, \#p\} \end{pmatrix}$ is a prefixpoint of the system T_γ^+ , and the derivation rules

$$\frac{\nabla_p \emptyset \vee \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}\} \rightarrow y_\gamma \quad \nabla\{\top, y_\gamma\} \rightarrow y_4 \quad \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}, y_\gamma\} \rightarrow y_{\gamma 4}}{\#p \rightarrow y_\gamma} \quad (R_\gamma)$$

$$\frac{\nabla_p \emptyset \vee \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}\} \rightarrow y_\gamma \quad \nabla\{\top, y_\gamma\} \rightarrow y_4 \quad \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}, y_\gamma\} \rightarrow y_{\gamma 4}}{\nabla\{\top, \#p\} \rightarrow y_4} \quad (R_4)$$

$$\frac{\nabla_p \emptyset \vee \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}\} \rightarrow y_\gamma \quad \nabla\{\top, y_\gamma\} \rightarrow y_4 \quad \nabla_p \{y_\gamma\} \vee \nabla_{\overline{p}} \{\top, y_{\gamma 4}, y_\gamma\} \rightarrow y_{\gamma 4}}{p \wedge \nabla\{\top, \#p\} \rightarrow y_{\gamma 4}} \quad (R_{\gamma 4})$$

expressing that this same vector is the *least* of the prefixpoints of T_γ^+ .

In order to address the general case, we discuss some notational issues. Given $S \in \mathcal{P}_+(Z)$, let $\chi_S^\#$ denote the following formula

$$\chi_S^\# = \bigwedge_{z \in S} \chi_z[\#_\gamma/x],$$

and, as usual, let $\chi^\#$ be the vector of terms $\{\chi_S^\# \mid S \in \mathcal{P}_+(Z)\}$. Using the one-to-one correspondence between the sets Y and $\mathcal{P}_+(Z)$, we let $[\chi^\#/\mathbf{y}]$ denote the substitution which replaces each variable $y = y_S$ with the formula $\chi_S^\#$. Furthermore, recall that $\{\sigma_S \mid S \in \mathcal{P}_+(Z)\}$ is the vector of terms of the modal system T_γ^+ .

Definition 3.22. The axiom system $\mathbf{K}_\#^+(\gamma)$ is obtained by adding to the axiomatization $\mathbf{K}_\#(\gamma)$ of Definition 1.1, for each $S \in \mathcal{P}_+(Z)$, the following axiom:

$$\sigma_S[\chi^\#/\mathbf{y}] \rightarrow \chi_S^\#, \quad (A_S)$$

as well as the following derivation rule:

$$\frac{\{\sigma_Q \rightarrow y_Q \mid Q \in \mathcal{P}_+(Z)\}}{\chi_S^\# \rightarrow y_S} \quad (R_S)$$

Finally, the axiom system $\mathbf{K}_\#^+(\Gamma)$ is obtained as the union of all the axioms and inference rules of the axiom systems $\mathbf{K}_\#^+(\gamma)$, $\gamma \in \Gamma$.

Remark 3.23. Strictly speaking, we no longer need the axiom ($\#_\gamma$ -prefix) and the rule ($\#_\gamma$ -least) since it can be proved on the basis of Proposition 4.2 and the results in Section 6, that ($\#_\gamma$ -prefix) is derivable and that ($\#_\gamma$ -least) is admissible in the system obtained by deleting ($\#_\gamma$ -prefix) and ($\#_\gamma$ -least) from $\mathbf{K}_\#^+(\Gamma)$.

Remark 3.24. It is not hard to see that the number of rules and axioms that we add to $\mathbf{K}_\#(\Gamma)$ in order to obtain $\mathbf{K}_\#^+(\Gamma)$ is in one-one correspondence with the set of non-finite subsets of RSF_γ , and thus exponential in the size of the formula γ , *provided* that γ has already been pre-processed, that is, γ is a disjunction of pure ∇ -formulas. However, the pre-processing procedure itself, rewriting a modal logic formula into this normal form, involves (at least) an exponential blow-up. We conjecture that the two steps of the procedure could be merged into one single algorithm which would produce an axiomatization of size exponential in the size of the original formula. We did not pursue this matter further since for our purposes it suffices to see that the axiomatization is *finite*, and because we believe that for clarity of exposition our separation of the various steps in the procedure is preferable.

Theorem 5.8 in Section 5 states the soundness and completeness of the axiom system $\mathbf{K}_\#^+(\Gamma)$ with respect to the Kripke semantics of $\mathcal{L}_\#(\Gamma)$, and in the same Section we give an overview of the proof of this result.

4 Comparing least fixpoints of systems of equations

This section is devoted to the proof two rather technical results relating the existence and nature of the least fixpoints of the formulas and systems of equations that we discussed in the previous section. The first proposition substantiates our claim that the semi-simple system of equations T_γ , obtained in step 1 in the procedure, represents the original formula γ , in the sense that in any modal algebra A , the (parametrized) least fixpoints of γ and those of T_γ can be derived from one another.

Proposition 4.1. *Let γ be a modal formula in which the variable x only occurs positively, let A be a modal algebra, and let $\mathbf{v} \in A^P$ be a sequence of parameters in A .*

1. *The least fixpoint $\mu_Z.(T_\gamma^A)_\mathbf{v}$ exists iff the least fixpoint $\mu_x.\gamma_\mathbf{v}^A$ exists.*
2. *If existing, these least fixpoints are related as follows. Writing $\mu_x.\gamma_\mathbf{v}^A = a$ and $\mu_Z.(T_\gamma^A)_\mathbf{v} = \{b_z \mid z \in Z\}$, we have*

$$a = b_{z_\gamma}, \tag{15}$$

$$b_{z_\psi} = \psi_\mathbf{v}^A(a), \quad \text{for all } \psi \in RSF_\gamma. \tag{16}$$

3. *If $\mu_Z.(T_\gamma^A)_\mathbf{v}$ is constructive then so is $\mu_x.\gamma_\mathbf{v}^A$. Conversely, if $\mu_x.\gamma_\mathbf{v}^A$ is constructive, then, provided the operations in γ are continuous, $\mu_Z.(T_\gamma^A)_\mathbf{v}$ is constructive as well.*

Proof. Fix γ , A and \mathbf{v} as in the statement of the proposition. In order to simplify notation, we write γ rather than $\gamma_\mathbf{v}^A$, and T rather than $(T_\gamma^A)_\mathbf{v}$.

First assume that $\mu_x.\gamma_\mathbf{v}^A$ exists, say $a = \mu_x.\gamma_\mathbf{v}^A$. It follows from Proposition 3.12 that the vector $\{\psi^A(a) \mid \psi \in RSF_\gamma\}$ is a solution of T^A . To see that it is in fact the *least* solution, let $\{b_\psi \mid \psi \in RSF_\gamma\}$ be another solution of T^A . Then, again by Proposition 3.12, b_γ is a solution of the equation $x = \gamma^A(x)$, and hence by assumption on a , we find $a \leq b_\gamma$. From this, a formula induction shows that $\psi^A(a) \leq b_\psi$, for each $\psi \in RSF_\gamma$. This proves the direction (\Rightarrow)

of part 1, and the equation (16) of part 2. The other direction of part 1, and the equation (15) of part 2 have a similar proof.

For the proof of part 3, we consider the approximating sequences $\{(\gamma^A)^n(\perp) \mid n \in \omega\}$ and $\{(T^A)^n(\perp) \mid n \in \omega\}$. Abbreviate $c_n := (\gamma^A)^n(\perp)$ and $t_n := \pi_{z_\gamma}((T^A)^n(\perp))$. The main claim in the proof is the following.

Claim 1. The sequences $(c_n)_{n \in \omega}$ and $(t_n)_{n \in \omega}$ are mutually cofinal:

1. For all $n \in \omega$ there is an $m \in \omega$ such that $t_n \leq c_m$.
2. For all $n \in \omega$ there is an $m \in \omega$ such that $c_n \leq t_m$.

PROOF OF CLAIM. For the first statement of the claim, by induction on n we prove that

$$\text{for all } n \in \omega : T^n(\perp) \leq \{\psi(\gamma^n(\perp)) \mid \psi \in RSF_\gamma\}. \quad (17)$$

The base case is immediate by the fact that $T^0(\perp) = \perp$. Inductively, for $\chi \in RSF_\gamma$ we have

$$\begin{aligned} \pi_\chi(T^{n+1}(\perp)) &= \widehat{\chi}[T^n(\perp)/z] \\ &\leq \widehat{\chi}[\psi[\gamma^n(\perp)/x]/z] \\ &= \chi(\gamma^n(\perp)). \end{aligned}$$

This proves (17), and so in particular we obtain

$$\text{for all } n \in \omega : t_{n+1} = \pi_\gamma(T^{n+1}(\perp)) \leq \gamma(\gamma^n(\perp)) = \gamma^{n+1}(\perp).$$

From this the first part of the claim is immediate.

Part 2 of the claim is a little harder to prove. Given a modal formula φ , let $d(\varphi)$ denote the modal *depth* of φ , and put $k := d(\gamma)$. Then by induction on n we prove that

$$\text{for all } n \in \omega : c_n \leq t_{kn}. \quad (18)$$

Whereas the base case of (18) is immediate by the fact that $c_0 = \perp$, for the inductive case we need a subinduction to prove the following.

$$\text{for all } \chi \in \{x\} \cup RSF_\gamma : \chi^A(c_n) \leq \pi_\chi(T^{kn+d(\chi)}(\perp)), \quad (19)$$

where we let π_x denote π_{z_γ} .

The proof of (19) proceeds by induction on the depth of χ . For the *base step* we must have $\chi = x$. So in this case we see that $\chi^A(c_n) = c_n$, while $\pi_\chi(T^{kn+d(\chi)}(\perp)) = \pi_{z_\gamma}(T^{kn}(\perp)) = t_{kn}$, where the latter equality is nothing but the definition of t_{kn} . So in this case, (19) follows from the main inductive hypothesis.

For the *inductive step*, fix a formula $\chi \in RSF_\gamma$. We may write $\chi = \widehat{\chi}(\psi_1, \dots, \psi_n)$, where each $\psi_i \in RSF_\gamma$ has depth properly smaller than $d(\chi)$, and $\widehat{\chi} = t_{z_\chi}$ is a depth 1 formula such that

$$\text{for all } \mathbf{a} \in A^Z, \widehat{\chi}^A(a_{\psi_1}, \dots, a_{\psi_n}) = \pi_\chi(T(\mathbf{a})). \quad (20)$$

Then we obtain

$$\begin{aligned}
\chi^A(c_n) &= \widehat{\chi}^A((\psi_1^A(c_n), \dots, \psi_n^A(c_n))) && \text{by definition of } \widehat{\chi} \\
&\leq \widehat{\chi}^A(\pi_{\psi_1}(T^{kn+d(\psi_1)}(\perp)), \dots, \pi_{\psi_n}(T^{kn+d(\psi_n)}(\perp))) && \text{by the IH} \\
&\leq \widehat{\chi}^A(\pi_{\psi_1}(T^{kn+d(\chi)-1}(\perp)), \dots, \pi_{\psi_n}(T^{kn+d(\chi)-1}(\perp))) && \text{by monotonicity} \\
&= \pi_\chi(T(T^{kn+d(\chi)-1}(\perp))) && \text{by (20)} \\
&= \pi_\chi(T^{kn+d(\chi)}(\perp)),
\end{aligned}$$

which proves (19).

To obtain the inductive case of (18) from this, take $\chi := \gamma$ in (19). This gives

$$c_{n+1} = \gamma(c_n) \leq \pi_\gamma(T^{kn+d(\gamma)}(\perp)) \leq \pi_\gamma(T^{(k+1)n}(\perp)) = t_{(k+1)n},$$

as required. ◀

It easily follows from Claim 1 that

$$\bigvee_{n \in \omega} c_n \text{ exists iff } \bigvee_{n \in \omega} t_n \text{ exists; and if existing, } \bigvee_{n \in \omega} c_n = \bigvee_{n \in \omega} t_n. \quad (21)$$

Now suppose that T has a constructive fixpoint $\mu_Z.T = \bigvee_{n \in \omega} T^n(\perp)$. It follows from part 1 that $\mu_x.\gamma$ exists and that $\mu_x.\gamma = \pi_\gamma(\mu_Z.T)$. But by the continuity of the projection operation π_γ we obtain that $\pi_\gamma(\mu_Z.T) = \bigvee_{n \in \omega} \pi_\gamma(T^n(\perp)) = \bigvee_{n \in \omega} t_n$, and so by (21) we may derive that $\mu_x.\gamma = \bigvee_{n \in \omega} \gamma^n(\perp)$. That is, γ has a constructive fixpoint indeed.

Conversely, suppose that γ has a constructive fixpoint: $\mu_x.\gamma = \bigvee_{n \in \omega} \gamma^n(\perp)$; write $c_\omega := \mu_x.\gamma$. Then by (16), $\mu_Z.T = \{\psi^A(c_\omega) \mid \psi \in RSF_\gamma\}$. But if all the operations in γ are continuous, then each $\psi \in RSF_\gamma$ is continuous, implying that

$$\psi^A(c_\omega) = \bigvee_{n \in \omega} \psi(c_n).$$

Then it follows from (19) and the continuity of the projections that

$$\bigvee_{n \in \omega} \psi(c_n) \leq \bigvee_{m \in \omega} \pi_\psi(T^m(\perp)) \leq \pi_\psi\left(\bigvee_{m \in \omega} T^m(\perp)\right).$$

Since this applies to all formulas $\psi \in RSF_\gamma$ we obtain that

$$T\left(\bigvee_{n \in \omega} T^n(\perp)\right) = \bigvee_{n \in \omega} T^n(\perp).$$

In other words, T has a constructive fixpoint as well. QED

The second proposition in this section relates the least fixpoint of a semi-simple system of equations to that of its simple simulation. It justifies the third step in the procedure of defining the axiomatization which we defined in the previous section.

Proposition 4.2. *Let T be a semi-simple modal system, let A be a modal algebra, and $\mathbf{v} \in A^P$ a sequence of parameters in A .*

1. *If A is complete, then $\mu_Z.T_{\mathbf{v}}^A$ and $\mu_Y.(T^+)_{\mathbf{v}}^A$ both exist, and they are related as follows. Writing $\mu_Z.T_{\mathbf{v}}^A = \{a_z \mid z \in Z\}$ and $\mu_Y.(T^+)_{\mathbf{v}}^A = \{b_y \mid y \in Y\}$, we have:*

$$\begin{aligned} a_z &= b_{\{z\}} && \text{for } z \in Z \\ b_{y_S} &= \bigwedge_{z \in S} a_z && \text{for } S \in \mathcal{P}_+(Z). \end{aligned}$$

2. *If $\mu_Y.(T^+)_{\mathbf{v}}^A$ exists and is constructive, then $\mu_Z.T_{\mathbf{v}}^A$ exists, and is constructive as well. Writing, again, $\mu_Z.T_{\mathbf{v}}^A = \{a_z \mid z \in Z\}$ and $\mu_Y.(T^+)_{\mathbf{v}}^A = \{b_y \mid y \in Y\}$, we have:*

$$a_z = b_{\{z\}} \quad \text{for } z \in Z.$$

Proof. Part 1 of the proposition is the main statement of Arnold & Niwiński in [2, §9].

Part 2 is a special case of Lemma 4.3 below. To see why we may apply this lemma, take $P := A^Z$, $Q := A^Y$, and let f and g be the maps $T_{\mathbf{v}}^A$ and $(T^+)_{\mathbf{v}}^A$, respectively. Let $\iota : A^Z \rightarrow A^Y$ be as in (10), and let $\pi : A^Y \rightarrow A^Z$ be given by

$$\pi(\mathbf{b})(z) := b_{\{z\}}.$$

Then it is obvious that all maps involved are order preserving, that $\iota(\perp) = \perp$, and that $\pi(\iota(\mathbf{a})) = \mathbf{a}$, for all $\mathbf{a} \in A^Z$. It is straightforward to prove that ι is continuous, and, finally, we already discussed the commutativity of the diagram (11). QED

We have isolated the following lemma from the proof of the previous Proposition since it may have some independent interest.

Lemma 4.3. *Let P, Q be posets with a least element \perp and consider a commuting diagram of the form*

$$\begin{array}{ccc} P & \xrightarrow{f} & P \\ \downarrow \iota & & \downarrow \iota \\ Q & \xrightarrow{g} & Q \end{array}$$

where f and g are order preserving, and ι is continuous and preserves \perp . Moreover, suppose that there exists an order preserving $\pi : Q \rightarrow P$ such that $\pi \circ \iota$ is the identity on P . If g has a constructive least prefixpoint $\mu.g$, then f has also has a constructive least prefixpoint $\mu.f$ given by the formula

$$\mu.f = \pi(\mu.g).$$

Proof. We shall prove that, for each ordinal α , the following holds:

$$\text{if } g^\alpha(\perp) \text{ exists, then } f^\alpha(\perp) \text{ exists, and } \iota(f^\alpha(\perp)) = g^\alpha(\perp). \quad (22)$$

Let us first see how to derive the Lemma from this. To start with, we may infer from (22) that for all α such that $g^\alpha(\perp)$ exists, we have $f^\alpha(\perp) = \pi(\iota(f^\alpha(\perp))) = \pi(g^\alpha(\perp))$. So if $g^{\omega+1}(\perp) =$

$g^\omega(\perp)$, then we immediately obtain that $f^{\omega+1}(\perp) = \pi(g^{\omega+1}(\perp)) = \pi(g^\omega(\perp)) = f^\omega(\perp)$. In other words, if $\mu.g$ is constructive then so is $\mu.f$.

We prove (22) by ordinal induction on α . If $\alpha = 0$, then $\iota(f^0(\perp)) = g^0(\perp)$ amounts to saying that ι preserves the least element. If α is a successor ordinal $\beta + 1$, then the existence of $f^\alpha(\perp)$ is not an issue. The second part of (22) follows from

$$\iota(f^\alpha(\perp)) = \iota(f(f^\beta(\perp))) = g(\iota(f^\beta(\perp))) = g(g^\beta(\perp)) = g^\alpha(\perp).$$

Here the second identity follows by the commutativity of the diagram, and the third identity, by the inductive hypothesis.

If α is a limit ordinal then we will prove first that the approximant $f^\alpha(\perp)$ exists. We will actually show that $\pi(g^\alpha(\perp)) = \bigvee_{\beta < \alpha} f^\beta(\perp)$, so that $\pi(g^\alpha(\perp)) = f^\alpha(\perp)$. Observe that, for $\beta < \alpha$, $g^\beta(\perp) \leq g^\alpha(\perp)$ implies $f^\beta(\perp) = \pi(g^\beta(\perp)) \leq \pi(g^\alpha(\perp))$. Also, if $f^\beta(\perp) \leq x$ for all $\beta < \alpha$, then $g^\beta(\perp) = \iota(f^\beta(\perp)) \leq \iota(x)$, hence $g^\alpha(\perp) \leq \iota(x)$ and $\pi(g^\alpha(\perp)) \leq \pi(\iota(x)) = x$. We are now ready to argue that $\iota(f^\alpha(\perp)) = g^\alpha(\perp)$:

$$\iota(f^\alpha(\perp)) = \iota\left(\bigvee_{\beta < \alpha} f^\beta(\perp)\right) = \bigvee_{\beta < \alpha} \iota(f^\beta(\perp)) = \bigvee_{\beta < \alpha} g^\beta(\perp) = g^\alpha(\perp),$$

where we need ι to be continuous in the second identity. QED

5 Soundness and Completeness

In this section we state the two main soundness and completeness results of the paper, and we outline the proofs.

As mentioned already, our completeness proofs are algebraic in nature, crucially involving the *Lindenbaum-Tarski algebra* $\mathbf{L}^{\mathbf{S}}(X)$ associated with a system \mathbf{S} of axioms and deductive rules, and with a set X of variables. In the next two subsections \mathbf{S} will denote $\mathbf{K}_\#(\Gamma)$ and $\mathbf{K}_\#^+(\Gamma)$, respectively, and, if \mathbf{S} and X are understood, we shall write simply \mathbf{L} in place of $\mathbf{L}^{\mathbf{S}}(X)$. The definition of \mathbf{L} is based on the standard construction of an algebra from the syntax of a logic [5]. The *elements* of this algebra are equivalence classes of the formulas/terms that are generated from the set X of variables. Here two terms t, s are declared to be equivalent if $\vdash t \leftrightarrow s$ is derivable in the system \mathbf{S} . The *operations* of our Lindenbaum algebra also have a standard definition. For example we shall have

$$[t] \wedge^{\mathbf{L}} [s] = [t \wedge s]$$

or, for the fixpoint connective $\#_\gamma$,

$$\#_\gamma^{\mathbf{L}}([t_1], \dots, [t_n]) = [\#_\gamma(t_1, \dots, t_n)].$$

Clearly, for the correctness of the latter definition we use the fact that the congruence rule

$$\frac{\{s_i \leftrightarrow t_i\}_{1 \leq i \leq n}}{\#_\gamma(s_1, \dots, s_n) \leftrightarrow \#_\gamma(t_1, \dots, t_n)}$$

is derivable in $\mathbf{K}_\#(\Gamma)$ — and a fortiori in $\mathbf{K}_\#^+(\Gamma)$ — as a straightforward derivation reveals.

Lindenbaum-Tarski algebras are of fundamental importance, both logically and algebraically. In logic, they are the algebraic incarnation of the associated derivation system \mathbf{S} , in the sense that two formulas s and t are equivalent with respect to \mathbf{S} iff the equation $s = t$ holds in the algebra $\mathbf{L}^{\mathbf{S}}(X)$ (provided that X contains all the variables occurring in s and t). Algebraically, they are the free algebras in the class of algebraic models for the logic.

More specifically, in our setting, both $\mathbf{L}^{\mathbf{K}_\#(\Gamma)}(X)$ and $\mathbf{L}^{\mathbf{K}_\#^+(\Gamma)}(X)$ are modal $\#$ -algebras, and, moreover, the latter algebra is regular. Also, in both cases, there is a canonical interpretation of the variables in X as elements in \mathbf{L} , sending the variable x to the equivalence class $[x]$ of the term x . Now first let \mathbf{L} be $\mathbf{L}^{\mathbf{K}_\#(\Gamma)}(X)$ and observe that whenever A is a modal $\#$ -algebra and $\mathbf{v} : X \rightarrow A$ is a valuation of the variables in x as elements of A , then there exists a unique modal $\#$ -algebra morphism $f : \mathbf{L} \rightarrow A$ such that $f[x] = \mathbf{v}(x)$ for all $x \in X$. In universal algebraic, or categorical terms, \mathbf{L} is the *free $\#$ -algebra over X* , and this property, freeness, determines \mathbf{L} up to isomorphism of modal $\#$ -algebras. Next, if we let \mathbf{L} be $\mathbf{L}^{\mathbf{K}_\#^+(\Gamma)}(X)$, then an analogous property holds: \mathbf{L} is the *free regular $\#$ -algebra over X* .

Returning to the proof sketch, we will underpin our completeness results algebraically by a representation theorem stating that

Theorem 5.1. *If X is countable, then $\mathbf{L}(X)$ embeds in a Kripke $\#$ -algebra.*

We shall see that such a theorem holds if $\mathbf{S} = \mathbf{K}_\#^+(\Gamma)$, so that $\mathbf{L}(X)$ is the free regular $\#$ -algebra over X , or if $\mathbf{S} = \mathbf{K}_\#(\Gamma)$ is the standard Kozen-Park axiomatization and all the formulas in Γ are subject to some syntactic constraints.

In both cases, such a result implies completeness as follows. Let X be the set of variables of a term/formula t . If the formula t is valid in every Kripke frame, then the equation $t = \top$ holds in every Kripke $\#$ -algebra, and thus certainly in the one that $\mathbf{L}(X)$ embeds into. Consequently, the equation $t = \top$ holds in the Lindenbaum algebra $\mathbf{L}(X)$, and by our earlier observation that \mathbf{L} incarnates the associated logic, this means that the formula $\top \leftrightarrow t$ is a derivable theorem of the associated logic. As usual, this implies that $\vdash t$ is derivable as well, which establishes the completeness of the logic.

In turn, the proof of Theorem 5.1 is subdivided in many steps, which we here collect into some main results, to be proved successively in the next two sections.

1. First we show that the modal operators $\diamond_i^{\mathbf{L}}$ of \mathbf{L} are residuated (Corollary 6.12).
2. Then we prove that \mathbf{L} is constructive (Theorem 6.18).
3. Finally, Theorem 7.1 states that every countable modal $\#$ -algebra that has residuated modalities and constructive fixpoint connectives, can be embedded in a Kripke $\#$ -algebra.

Since $\mathbf{L}(X)$ is countable whenever X is countable, Theorem 5.1 follows immediately from this.

The proof of Theorem 5.1 will be carried out almost in parallel for the two systems $\mathbf{K}_\#(\Gamma)$ and $\mathbf{K}_\#^+(\Gamma)$. For the sake of readability, we shall give the details of the proof in the monomodal setting but discuss also in extent the steps that have to be taken to generalize the proof to the polymodal setting.

5.1 Completeness of the Kozen-Park axiomatization

As we mentioned in the Introduction, in many cases the relatively simple Kozen axiomatization is already sound and complete with respect to the Kripke semantics. This applies to flat modal fixpoint languages in which each connective \sharp_γ can be defined as the least fixpoint of a formula γ' which is *untied* with respect to x . This notion is closely related to those of the aconjunctive formulas of Kozen [19] and the disjunctive formulas of Walukiewicz [40], but it is fine-tuned to the fact that we are focussing on the special role of the variable x .

Definition 5.2. A modal ∇ -formula $\gamma(x) \in \mathcal{L}_\nabla(X)$ is *untied in x* if it can be obtained from the following grammar:

$$\varphi ::= x \mid \top \mid \varphi \vee \varphi \mid \psi \wedge \bigwedge_{j \in J} \nabla_j \Phi_j$$

Here ψ is a formula in which x does not occur, $J \subseteq I$, and each Φ_j is a set of x -untied formulas.

Example 5.3. The key point of untied formulas in x is that we restrict the use of conjunctions to formulas of the form $\psi \wedge \bigwedge_{j \in J} \nabla_j \Phi_j$ where x may not occur in ψ , and no two ∇ operators in $\bigwedge_{j \in J} \nabla_j \Phi_j$ may be indexed by the same atomic action. Thus, for instance, the formulas $\nabla_1 \{ \nabla_2 \{ p \} \} \wedge \nabla_1 \{ x \}$ and $\nabla_1 \{ \nabla_2 \{ x \} \} \wedge \nabla_2 \{ x \}$ are untied in x , but the formula $\nabla_1 \{ \nabla_2 \{ x \} \} \wedge \nabla_1 \{ x \}$ is not. For a slightly more elaborate example, the formula

$$\varphi := (\nabla_1 \{ \top, x, \nabla_1 \{ \top, x \} \} \wedge \nabla_2 \emptyset) \vee (\nabla_1 \{ \top, x, \nabla_1 \{ \top, x \} \} \wedge \nabla_2 \{ \nabla_1 \{ x, \top \} \}) \quad (23)$$

can be parsed by the above grammar and therefore is untied in x .

We can now formulate the first result, returning to its proof at the end of this subsection.

Theorem 5.4. *Suppose that each $\gamma(x) \in \Gamma$ is untied with respect to x . Then the axiom system $\mathbf{K}_\sharp(\Gamma)$ is sound and complete with respect to the Kripke semantics of $\mathcal{L}_\sharp(\Gamma)$.*

For readers that are not familiar with the cover modalities, we give a corollary of Theorem 5.4 that is phrased in terms of the classical presentation of modal logic using diamonds and boxes. We leave it for the reader to verify that this corollary covers fixpoint connectives \sharp_γ indexed by a formula γ in which x has exactly one, positive, occurrence. This takes care of for instance the computation tree logic, CTL.

Definition 5.5. A modal formula $\gamma(x)$ is *harmless with respect to x* if it can be obtained from the following grammar:

$$\varphi ::= x \mid \top \mid \varphi \vee \varphi \mid \psi \wedge \varphi \mid \diamond_i \varphi \mid \square_i \varphi \mid \bigwedge_{j \in J} \varphi_j.$$

Here ψ is a formula in which x does not occur, $J \subseteq I$, and $\bigwedge_{j \in J} \varphi_j$ is a *harmless conjunction*. This means that for each $j \in J$, the conjunct φ_j is either of the form $\square_j \chi$, or itself a conjunction of the form $\bigwedge_{\ell \in L} \diamond_j \chi_\ell$ (with χ and each χ_ℓ being harmless in x).

Example 5.6. The formula $\diamond_1(x \wedge \diamond_2 x)$ is not harmless, and neither is $\diamond_1 x \wedge \square_1 \diamond_2 x$. The formula $\diamond_1 x \wedge \diamond_1 \diamond_1 x \wedge \square_2 \diamond_1 x$ is, on the other hand, harmless, and this also applies to $\diamond_1 x \wedge \diamond_1 \diamond_1 x \wedge \square_1 \diamond_1 p$.

Corollary 5.7. *Let Γ be a set of modal formulas each of which is harmless with respect to x . Then the axiom system $\mathbf{K}_{\sharp}(\Gamma)$ is sound and complete with respect to the Kripke semantics of $\mathcal{L}_{\sharp}(\Gamma)$.*

Proof. A straightforward induction shows that every $\gamma(x)$ which is harmless with respect to x , is equivalent to a ∇ -formula that is untied in x . (For instance, the harmless formula $\diamond_1 x \wedge \square_1 \diamond_2 x$ of Example 5.6 is equivalent to the untied formula (23).) Then the Corollary is immediate by Theorem 5.4. QED

Proof of Theorem 5.4. The axiomatization $\mathbf{K}_{\sharp}(\Gamma)$ certainly is sound. To argue about completeness, we need Theorem 5.1 for \mathbb{L} the Lindenbaum algebra associated with $\mathbf{K}_{\sharp}(\Gamma)$. We proceed along the path sketched above and, to this goal, the key observation is that if γ is untied in x , then for any vector of parameters \mathbf{v} , the term function $\gamma_{\mathbf{v}}^{\mathbb{L}}$ on the Lindenbaum-Tarski algebra $\mathbb{L} = \mathbb{L}^{\mathbf{K}_{\sharp}(\Gamma)}(X)$ is a finitary \emptyset -adjoint. This implies that the least fixpoint $\gamma_{\mathbf{v}}^{\mathbb{L}}$ is constructive, see Theorem 6.18 for more details of this argument. QED

5.2 The general case

We leave it as an open problem whether, in the general case, the system $\mathbf{K}_{\sharp}(\Gamma)$ is complete. However, for its extension $\mathbf{K}_{\sharp}^+(\Gamma)$ we have the following uniform soundness and completeness result.

Theorem 5.8. *The axiom system $\mathbf{K}_{\sharp}^+(\Gamma)$ is sound and complete with respect to the Kripke semantics of $\mathcal{L}_{\sharp}(\Gamma)$.*

In the sequel we shall use the phrase “free regular \sharp -algebra” as a synonym of the Lindenbaum algebra, and $\mathbb{L}, \mathbb{L}(X)$ shall be short notation for $\mathbb{L}^{\mathbf{K}_{\sharp}^+(\Gamma)}(X)$.

Proof of Theorem 5.8. As we mentioned already in the previous section, the *soundness* of $\mathbf{K}_{\sharp}^+(\Gamma)$ follows from the main result of Arnold & Niwiński in [2, §9], here mentioned as Proposition 4.2. For, it is an immediate consequence of this result that all Kripke \sharp -algebras are regular. But from this and the fact that the Lindenbaum-Tarski algebra is the free regular \sharp -algebra, the soundness of $\mathbf{K}_{\sharp}^+(\Gamma)$ follows by a standard algebraic logic argument.

To argue for completeness, we need Theorem 5.1 for \mathbb{L} the Lindenbaum algebra associated with $\mathbf{K}_{\sharp}^+(\Gamma)$. We proceed again along the path sketched above but this time the path is less direct.

To argue that the least fixpoint of $\gamma_{\mathbf{v}}^{\mathbb{L}}$ is constructive, we first observe that the least fixpoint of $(T_{\gamma}^+)_{\mathbf{v}}^{\mathbb{L}}$ – which by regularity exists – is constructive, since (T_{γ}^+) is a simple system and its interpretation in \mathbb{L} is a finitary \emptyset -adjoint. Then, the property of constructiveness of the respective fixpoints can be transferred from $(T_{\gamma}^+)_{\mathbf{v}}^{\mathbb{L}}$ to $(T_{\gamma})_{\mathbf{v}}^{\mathbb{L}}$ and from $(T_{\gamma})_{\mathbf{v}}^{\mathbb{L}}$ to $\gamma_{\mathbf{v}}^{\mathbb{L}}$, using the results of Section 4. A detailed account of this process will be given in the proof of Theorem 6.18. QED

6 Properties of the Lindenbaum Algebras

The goal of this section is to prove that the Lindenbaum algebra $L^{\mathbf{S}}$, where \mathbf{S} is one of the axiom systems $\mathbf{K}_{\sharp}(\Gamma)$ and $\mathbf{K}_{\sharp}(\Gamma)$, is constructive, cf. Definition 2.9. We shall obtain this result by subsequently analyzing properties of this algebra.

6.1 Rigidness

We start with showing that L is *rigid* with respect to X .

Definition 6.1. Let A be a modal algebra generated by a set X . A is *rigid* with respect to X if

$$\bigwedge G \wedge \nabla Y = \perp \text{ implies } \bigwedge G = \perp \text{ or } \exists y \in Y \text{ s.t. } y = \perp, \quad (24)$$

where G and Y are finite, possibly empty, sets of elements of A , with $G \subseteq \{x, \neg x \mid x \in X\}$.

Remark 6.2. In a polymodal setting we say that A is *rigid* with respect to X if

$$\bigwedge G \wedge \bigwedge_{i \in I} \nabla_i Y_i = \perp \text{ implies } \bigwedge G = \perp \text{ or } \exists i \in I \text{ and } y \in Y_i \text{ s.t. } y = \perp.$$

Remark 6.3. To gather some intuitions about this property, we first prove rigidness of the free *modal* algebra $\mathcal{M}(X)$ generated by a set X of variables. Reformulating the property in terms of formulas, and reasoning by contraposition, it suffices to show that whenever Λ is a consistent set of X -literals, and $\Phi = \{\varphi_1, \dots, \varphi_n\}$ is a set of consistent modal formulas, then the formula $\bigwedge \Lambda \wedge \nabla \Phi$ is consistent as well.

So let Λ and Φ be as indicated. Then by completeness there is a pointed Kripke model (\mathbb{M}_i, r_i) for each formula φ_i . Now create a new model \mathbb{M} as follows. Take the disjoint union of the models $\mathbb{M}_1, \dots, \mathbb{M}_n$, and add one single new point r . Let $\{r_1, \dots, r_n\}$ be the successor set of r , and define a valuation for r so that the propositional formula $\bigwedge \Lambda$ is true at r . Clearly then $\mathbb{M}, r \Vdash \bigwedge \Lambda \wedge \nabla \Phi$, witnessing the consistency of the formula $\bigwedge \Lambda \wedge \nabla \Phi$.

Second, for readers that are familiar with the *duality theory* of modal algebras [38], the notion of rigidness has a very natural formulation in terms of the dual relational space A_* of A . Let A be a modal algebra generated by some set X . Then A is rigid with respect to X iff for every finite set $G \subseteq \{x, \neg x \mid x \in X\}$ such that $\bigwedge G > \perp$, and every finite set $U = \{u_1, \dots, u_n\}$ of ultrafilters in A_* , there is an ultrafilter $u \supseteq G$ which has U as its collection of successors.

Theorem 6.4. Let L denote either the free modal \sharp -algebra or the free regular modal \sharp -algebra. Then L is rigid with respect to X .

The proof of the Theorem depends on the following construction.

Definition 6.5. Let A be some modal algebra, and let $\Pi = \{\pi_\ell : A \longrightarrow 2 \mid 0 < \ell \leq n\}$ be a finite (possibly empty) set of Boolean algebra homomorphisms. We define the operation $\diamond^\Pi : A \longrightarrow 2$ by putting

$$\diamond^\Pi(a) := \bigvee \{\pi(a) \mid \pi \in \Pi\}. \quad (25)$$

We define the operation $\diamond^{A^\Pi} : A \times 2 \longrightarrow A \times 2$ as follows:

$$\diamond^{A^\Pi}(a, d) := (\diamond^A(a), \diamond^\Pi(a)), \quad (26)$$

and let A^Π be the algebra obtained by expanding the Boolean algebra $A \times 2$ with this operation.

For future reference we define the cover operation associated with \diamond^Π as the map $\nabla^\Pi : \mathcal{P}_\omega A \longrightarrow 2$ given by

$$\nabla^\Pi \alpha := \square^\Pi \bigvee \alpha \wedge \bigwedge \diamond^\Pi \alpha, \quad (27)$$

where of course $\square^\Pi x = \neg \diamond^\Pi \neg x$.

Remark 6.6. In a polymodal setting the construction has to be parameterized by a collection of the form $\{\Pi_i \mid i \in I\}$. Then \diamond^{Π_i} and $\diamond_i^{A^\Pi}$ are defined from Π_i as in the equations (25) and (26), respectively.

Remark 6.7. Again, a dual perspective on this construction may be illuminating. Recall that Boolean homomorphisms may be identified with ultrafilters. In a nutshell (and again, presupposing familiarity with the duality theory of modal algebras), we obtain the dual structure of A^Π by *adding* an ultrafilter u to the dual structure A_* of A , making the set Π of Boolean homomorphisms/ultrafilters its successor set.

It is not difficult to verify that the operation \diamond^{A^Π} is additive, so that A^Π is a modal algebra. But in fact, as we will see in the Proposition below, the construction preserves many other properties as well.

Proposition 6.8. *Let A be a modal algebra, and let $\Pi = \{\pi_i : A \longrightarrow 2 \mid 0 < i \leq n\}$ be a finite set of Boolean algebra homomorphisms.*

1. *If A is a modal \sharp -algebra for some fixpoint connective \sharp_γ , then so is A^Π .*
2. *Let T be a semi-simple modal system, and let $\mathbf{v} \in A^P$ be some parameter for T . If $T_{\mathbf{v}}^A$ has a least solution on A , then so does $T_{(\mathbf{v}, \mathbf{w})}^{A^\Pi}$, for each parameter $\mathbf{w} \in 2^P$.*
3. *If A is regular with respect to some semi-simple modal system T , then so is A^Π .*

Proof. Since part 1 of the proposition is a direct consequence of part 2 and Proposition 4.1, we start with proving part 2. Let $T = \langle Z, \{t_z \mid z \in Z\} \rangle$ be a semi-simple system of equations. Since the carrier of A^Π is the set $A \times 2$, we may see $T^{A^\Pi} : (A^\Pi)^Z \times (A^\Pi)^P \longrightarrow (A^\Pi)^Z$ as a map

$$T^{A^\Pi} : (A^Z \times A^P) \times (2^Z \times 2^P) \longrightarrow (A^Z \times 2^Z).$$

Let π_A and π_2 denote the projections of A^Π onto A and 2 , respectively.

Given the definition of the modal operator of A^Π , the first coordinate $\pi_A \circ T^{A^\Pi}$ of the map T^{A^Π} is identical to the map $T^A \circ \pi_A$. Furthermore, since T is semi-simple, in each term t_z the unguarded variables are all from P , while the guarded variables are all from Z , and

each occurrence of these is in the scope of exactly one modality. As a consequence, the second coordinate of T^{A^Π} is the compose of

$$(A^Z \times 2^Z) \times (A^P \times 2^P) \xrightarrow{\pi} A^Z \times 2^P \xrightarrow{\widetilde{T}_2} 2^Z.$$

Here \widetilde{T}_2 is best understood by observing that its terms are obtained from those of T by replacing every occurrence of the symbol ∇ with the formal symbol ∇^Π .

Summarizing, we may write $T^{A^\Pi} = \langle T^A \circ \pi_A, \widetilde{T}_2 \circ \pi \rangle$. It follows by Bekić' property that, for each $\mathbf{v} \in A^P$ and $\mathbf{w} \in 2^P$, the least fixpoint of $T^{A^\Pi}_{(\mathbf{v}, \mathbf{w})}$ exists, and can be written as

$$\mu_Z.T^{A^\Pi}_{(\mathbf{v}, \mathbf{w})} = \langle \mu_Z.T_{\mathbf{v}}^A, \widetilde{T}_2(\mu_Z.T_{\mathbf{v}}^A, \mathbf{w}) \rangle. \quad (28)$$

Part 3 also follows from part 2, but it needs more work. We first prove that the following diagram commutes, for every $\mathbf{w} \in 2^P$:

$$\begin{array}{ccc} A^Z & \xrightarrow{(\widetilde{T}_2)_{\mathbf{w}}} & 2^Z \\ \downarrow \iota^A & & \downarrow \iota^2 \\ A^Y & \xrightarrow{(\widetilde{T}_2^+)_{\mathbf{w}}} & 2^Y \end{array} \quad (29)$$

Recall that in Section 3 we showed the diagram (11) to commute because of Proposition 3.14(2). A careful analysis of that proposition reveals that the only property needed for its proof is that the diamond \diamond underlying the operation ∇ (in the sense that $\nabla\alpha = \square \vee \alpha \wedge \diamond\alpha$ with $\square a = \neg \diamond \neg a$) preserves finite joins. Now the operation \diamond^Π underlying the operation ∇^Π of \widetilde{T}_2 also preserves finite joins, and so we prove that the diagram (29) commutes in exactly the same manner.

Now we establish the regularity of A^Π as follows. First, it follows from part 2 of this proposition that for each $\mathbf{v} \in A^P$ and $\mathbf{w} \in 2^P$, the least fixpoint $\mu_Y.(T^+)_{(\mathbf{v}, \mathbf{w})}^{A^\Pi}$ exists. Moreover, we may calculate

$$\begin{aligned} \mu_Y.(T^+)_{(\mathbf{v}, \mathbf{w})}^{A^\Pi} &= \langle \mu_Y.(T^+)_{\mathbf{v}}^A, \widetilde{T}_2^+(\mu_Y.(T^+)_{\mathbf{v}}^A, \mathbf{w}) \rangle && \text{by (28),} \\ &= \langle \iota^A(\mu_Z.T_{\mathbf{v}}^A), \widetilde{T}_2^+(\iota^A(\mu_Z.T_{\mathbf{v}}^A), \mathbf{w}) \rangle && \text{since } A \text{ is regular,} \\ &= \langle \iota^A(\mu_Z.T_{\mathbf{v}}^A), \iota^2(\widetilde{T}_2(\mu_Z.T_{\mathbf{v}}^A, \mathbf{w})) \rangle && \text{since diagram (29) commutes,} \\ &= \iota^{A^\Pi}(\langle \mu_Z.T_{\mathbf{v}}^A, \widetilde{T}_2(\mu_Z.T_{\mathbf{v}}^A, \mathbf{w}) \rangle) && \text{since } \iota^{A^\Pi} = \iota^A \times \iota^2, \\ &= \iota^{A^\Pi}(\mu_Z.T^A(\mathbf{v}, \mathbf{w})) && \text{again, by (28).} \end{aligned}$$

This finishes the proof of the third and final part of the proposition. QED

We can now prove the rigidness of \mathbb{L} , on the basis of Proposition 6.8 and the fact that \mathbb{L} is the *free* \sharp -algebra over X . Moreover, part 3 of Proposition 6.8 ensures that the same proof works if \mathbb{L} is the *free regular* \sharp -algebra over X .

Proof of Theorem 6.4. Suppose for contradiction that \mathbf{L} is not rigid with respect to X . Then there is a finite set Λ of X -literals, and a finite subset $\alpha \subseteq_{\omega} A$ such that $\bigwedge \Lambda > \perp$ and $b > \perp$ for all $b \in \alpha$, while $\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha = \perp$.

By the prime filter theorem, we may find a set $\Pi = \{\pi_b : \mathbf{L} \rightarrow 2 \mid b \in \alpha\}$ of Boolean homomorphisms such that $\pi_b(b) = \top$ for all $b \in \alpha$. Now consider the algebra \mathbf{L}^{Π} , and let $f : X \rightarrow \mathbf{L}^{\Pi}$ be some map satisfying

$$f(x) = \begin{cases} (x, \top) & \text{if } x \in \Lambda \\ (x, \perp) & \text{if } \neg x \in \Lambda. \end{cases} \quad (30)$$

Clearly, such a map exists by the consistency of Λ , and since \mathbf{L} is the free (regular) \sharp -algebra generated by X , f can be extended to a modal \sharp -homomorphism \tilde{f} from \mathbf{L} to \mathbf{L}^{Π} . Then it follows from our assumption that $\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha) = \tilde{f}(\perp^{\mathbf{L}}) = \perp^{\mathbf{L}^{\Pi}}$.

On the other hand, we claim that

$$\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha) = (\perp, \top), \quad (31)$$

which provides us with the desired contradiction. For the proof of (31), using the fact that \tilde{f} is a homomorphism, we find

$$\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha) = \tilde{f}(\bigwedge \Lambda) \wedge \tilde{f}(\nabla^{\mathbf{L}} \alpha).$$

From the assumption (30) on f , and the fact that \tilde{f} is an extension of f , it follows that $\tilde{f}(a) = (a, \top)$ for all a in Λ , so that $\tilde{f}(\bigwedge \Lambda) = \bigwedge \{\tilde{f}(a) \mid a \in \Lambda\} = (\bigwedge \Lambda, \top)$, while $\tilde{f}(\nabla^{\mathbf{L}} \alpha) = (\nabla^{\mathbf{L}} \alpha, \nabla^{\Pi} \alpha)$, where ∇^{Π} is the cover modality associated with \diamond^{Π} , see (27). The point of the construction of \mathbf{L}^{Π} is that

$$\nabla^{\Pi} \alpha = \top, \quad (32)$$

as we shall prove now. The relation (32) trivially holds if α is empty, since then $\diamond^{\Pi} x = \perp$ for all $x \in A$ and so $\nabla^{\Pi} \alpha = \square^{\Pi} \perp = \neg \diamond^{\Pi} \top = \top$. So let us now assume that α is not empty. Then we compute

$$\begin{aligned} \square^{\Pi} \bigvee \alpha &= \neg \diamond^{\Pi} (\bigwedge \{\neg b \mid b \in \alpha\}) \\ &= \neg \bigvee_{a \in \alpha} \pi_a (\bigwedge \{\neg b \mid b \in \alpha\}) && \text{by (25)} \\ &\geq \neg \bigvee_{a \in \alpha} \pi_a (\neg a) && (\pi_a \text{ is monotone}) \\ &\geq \neg \bigvee_{a \in \alpha} \neg \pi_a (a) && (\pi_a \text{ is a homomorphism}) \\ &= \neg \bigvee_{a \in \alpha} \neg \top && (\text{by assumption on } \pi_a) \\ &= \top \end{aligned}$$

and

$$\bigvee \diamond^{\Pi} \alpha = \bigvee_{b \in \alpha} \diamond^{\Pi} b = \bigvee_{b \in \alpha} \bigvee_{a \in \alpha} \pi_a(b) \geq \bigvee_{b \in \alpha} \pi_b(b) = \bigvee_{b \in \alpha} \top = \top$$

so that we find

$$\nabla^{\Pi} \alpha = \square^{\Pi} \bigvee \alpha \wedge \bigvee \diamond^{\Pi} \alpha = \top,$$

which proves (32). Continuing our computation of $\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha)$, we now have that

$$\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha) = (\bigwedge \Lambda, \top) \wedge (\nabla^{\mathbf{L}} \alpha, \top) = (\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha, \top \wedge \top) = (\perp, \top).$$

This finishes the proof of (31), and thus, of the Theorem. QED

Remark 6.9. In a polymodal setting, by the same sort of computations, we shall have

$$\tilde{f}(\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha) = (\bigwedge \Lambda, \top) \wedge \bigwedge_{i \in I} (\nabla_i^{\mathbf{L}} \alpha_i, \top) = (\bigwedge \Lambda \wedge \nabla^{\mathbf{L}} \alpha, \top \wedge \bigwedge_{i \in I} \top) = (\perp, \top).$$

Thus, in presence of many modalities, a contradiction with the regularity of \mathbf{L} is obtained in a similar way.

6.2 Finitary \emptyset -adjoints

We now turn to the notion of a finitary \emptyset -adjoint and to its generalization, that of a finitary family of \emptyset -adjoints, see Definition 2.13. The use of these notions lies in an earlier result by the first author [35], which roughly states that fixpoints of finitary \emptyset -adjoints, if existing, are constructive. In order to apply this result we aim to show that simple systems of equations on the Lindenbaum algebra give rise to finitary \emptyset -adjoints. To reach this goal we only need \mathbf{L} to be rigid with respect to X and to be generated by X . Therefore the next results apply both to the Lindenbaum algebra $\mathbf{L}^{\mathbf{K}_i^+(\Gamma)}$ and to the Lindenbaum algebra $\mathbf{L}^{\mathbf{K}_\#^+(\Gamma)}$.

Our first observation is that the cover modality $\nabla^{\mathbf{L}}$ on the Lindenbaum algebra is itself a finitary \emptyset -adjoint. In order to turn this into a meaningful mathematical statement, we need to endow the domain $\mathcal{P}_\omega(\mathbf{L})$ of the operation $\nabla^{\mathbf{L}}$ with a quasi-order, see Remark 2.12. Thus, let us define the relation $\overline{\leq}$ on $\mathcal{P}_\omega(\mathbf{L})$ by saying that $\alpha \overline{\leq} \beta$ iff for all $a \in \alpha$ there is a $b \in \beta$ such that $a \leq b$, and for all $b \in \beta$ there is an $a \in \alpha$ such that $a \leq b$. It is not hard to see that $\overline{\leq}$ is a quasi-order on $\mathcal{P}_\omega(\mathbf{L})$.

Theorem 6.10. *Let L denote either the free modal $\#$ -algebra or the free regular modal $\#$ -algebra. Then each cover modality $\nabla_i^{\mathbf{L}} : \mathcal{P}_\omega(L) \longrightarrow L$ is an \emptyset -adjoint.*

Proof. Given an element $d \in L$, we need to define a finite set $G_\nabla(d) \in \mathcal{P}_\omega \mathcal{P}_\omega(L)$ such that for all $\alpha \in \mathcal{P}_\omega(L)$, we have

$$\nabla \alpha \leq d \text{ iff } \alpha \overline{\leq} \beta \text{ for some } \beta \in G_\nabla(d). \quad (33)$$

First we confine our attention to the so-called *weakly irreducible elements* of \mathbf{L} , that is, the ones of the form

$$\bigvee \Pi \vee \diamond b \vee \bigvee_{c \in C} \square c, \quad (34)$$

where Π is some set of X -literals, b is an element of \mathbf{L} , and C is a finite set of elements of \mathbf{L} .

For a weakly irreducible element $d = \bigvee \Pi \vee \diamond b \vee \bigvee_{c \in C} \square c$ we let

$$G_{\nabla}(d) := G^{\Pi}(d) \cup G^{\diamond}(d) \cup G^{\square}(d), \quad (35)$$

where

$$G^{\Pi}(d) := \begin{cases} \{\{\top\}, \emptyset\} & \text{if } \bigvee \Pi = \top, \\ \emptyset & \text{otherwise,} \end{cases} \quad (36)$$

$$G^{\diamond}(d) := \{b, \top\},$$

$$G^{\square}(d) := \bigcup_{c \in C} \{b \vee c, \emptyset\}.$$

The correctness of this definition follows from the following Claim.

Claim 1. Let $d = \bigvee \Pi \vee \diamond b \vee \bigvee_{c \in C} \square c$ be weakly irreducible. Then the following are equivalent, for any $\alpha \in \mathcal{P}_{\omega}(\mathbf{L})$:

1. $\nabla \alpha \leq d$;
2. (a) $\bigvee \Pi = \top$, or
(b) $a \leq b$ for some $a \in \alpha$, or
(c) $\bigvee \alpha \leq b \vee c$ for some $c \in C$;
3. $\alpha \bar{\leq} \beta$, for some $\beta \in G_{\nabla}(d)$.

PROOF OF CLAIM. (1 \Rightarrow 2) Reasoning by contraposition, we assume that (2) does not hold. Then (a') the set $\Lambda := \{\neg \pi \mid \pi \in \Pi\}$ of literals is consistent, (b') $\neg b \wedge a > \perp$ for every $a \in \alpha$, and (c') $\neg b \wedge \neg c \wedge \bigvee \alpha > \perp$ for every $c \in C$. Now consider the element

$$e := \bigwedge \Lambda \wedge \nabla(\{\neg b \wedge a \mid a \in \alpha\} \cup \{\neg b \wedge \neg c \wedge \bigvee \alpha \mid c \in C\}).$$

It is immediate that $e \leq \bigwedge \Lambda$, and easy to verify that $e \leq \bigwedge_{c \in C} \diamond \neg c$. In addition, considering that

$$\begin{aligned} e &\leq \square \bigvee (\{\neg b \wedge a \mid a \in \alpha\} \cup \{\neg b \wedge \neg c \wedge \bigvee \alpha \mid c \in C\}) \\ &= \square(\neg b \wedge \bigvee(\{a \mid a \in \alpha\} \cup \{\neg c \wedge \bigvee \alpha \mid c \in C\})) \end{aligned}$$

we have $e \leq \square \neg b$. Combining these observations, we find that $e \leq \neg d$. But it is also easily seen that $e \leq \nabla \alpha$. On the other hand, we may apply the rigidness of \mathbf{L} to derive from (a')–(c') that $e > \perp$. From this it follows that $\nabla \alpha \not\leq d$; that is, (1) fails, as required.

(2 \Rightarrow 1) In each of the cases (2a)–(2c) it is obvious that $\nabla\alpha \leq d$.

(2 \Rightarrow 3) Suppose that (2) holds, and distinguish cases. (a) If $\bigvee \Pi = \top$ then both $\{\top\}$ and \emptyset belong to $G_{\nabla}(d)$. Then $\alpha \overline{\leq} \{\top\}$ if $\alpha \neq \emptyset$, and $\alpha \overline{\leq} \emptyset$ if $\alpha = \emptyset$, so there is always some $\beta \in G_{\nabla}(d)$ with $\alpha \overline{\leq} \beta$. (b) If $a \leq b$ for some $a \in \alpha$, then it is easy to see that $\alpha \overline{\leq} \{b, \top\}$, and this suffices to prove (3) since in this case $\{b, \top\}$ belongs to $G_{\nabla}(d)$. (c) If $\bigvee \alpha \leq b \vee c$, with $c \in C$, then $\alpha \overline{\leq} \{b \vee c\}$ if $\alpha \neq \emptyset$ and $\alpha \overline{\leq} \emptyset$ if $\alpha = \emptyset$. In both cases we have proved (3), since both \emptyset and $\{b \vee c\}$ belong to $G_{\nabla}(d)$.

(3 \Rightarrow 2) Assume that $\alpha \overline{\leq} \beta$, with $\beta \in G_{\nabla}(d)$, and again distinguish cases. If $\beta \in G^{\Pi}(d)$, then in particular $G^{\Pi}(d)$ is nonempty; this can only be the case if $\bigvee \Pi = \top$, so (2a) holds. If $\beta \in G^{\diamond}(d)$, then $\beta = \{b, \top\}$, so from $\alpha \overline{\leq} \beta$ it follows that there is some $a \in \alpha$ such that $a \leq b$, so (2b) holds. Finally, if $\beta \in G^{\square}(d)$, then C is not empty. If $\beta = \emptyset$, then $\alpha = \emptyset$. Let $c \in C$ be arbitrary, then $\bigvee \alpha = \perp \leq b \vee c$. If $\beta = \{b \vee c\}$ for some $c \in C$, then from $\alpha \overline{\leq} \beta$ we may deduce that $a \leq b \vee c$ for all $a \in \alpha$. This implies $\bigvee \alpha \leq b \vee c$. In both cases (2c) holds. \blacktriangleleft

Finally, let d be an arbitrary element of \mathbf{L} . It is not hard to show that d can be written as a finite meet $d = \bigwedge_{\ell=1, \dots, n} d_{\ell}$ of weakly irreducible elements. Thus, in order to define $G_{\nabla}(d)$ for such a meet, it is enough to define $G_{\nabla}(\top)$ and $G_{\nabla}(d_1 \wedge d_2)$ assuming that we have already defined $G_{\nabla}(d_1)$ and $G_{\nabla}(d_2)$. We let

$$\begin{aligned} G_{\nabla}(\top) &= \{ \{\top\}, \emptyset \}, \\ G_{\nabla}(d_1 \wedge d_2) &= \{ \{b_1 \wedge b_2 \mid (b_1, b_2) \in Z\} \mid \exists \beta_i \in G_{\nabla}(d_i), i = 1, 2, \text{ and } Z \in \beta_1 \bowtie \beta_2 \}. \end{aligned} \quad (37)$$

We leave it to the reader to verify that, with the above definition, $G_{\nabla}(\top)$ satisfies (33). For $G_{\nabla}(d_1 \wedge d_2)$ we argue as follows. If $\nabla\alpha \leq d_1 \wedge d_2$ then, for $i = 1, 2$, $\nabla\alpha \leq d_i$ and $\alpha \overline{\leq} \beta_i$ for some $\beta_i \in G_{\nabla}(d_i)$. Define Z by putting $(b_1, b_2) \in Z$ iff there exists $a \in \alpha$ such that $a \leq b_1$ and $a \leq b_2$. Then $Z \in \beta_1 \bowtie \beta_2$ and $\alpha \overline{\leq} \{b_1 \wedge b_2 \mid (b_1, b_2) \in Z\}$. Conversely, if for $i = 1, 2$, some $\beta_i \in G_{\nabla}(d_i)$ and some $Z \in \beta_1 \bowtie \beta_2$, the relation $\alpha \overline{\leq} \{b_1 \wedge b_2 \mid (b_1, b_2) \in Z\}$ holds, then $\alpha \overline{\leq} \beta_i$, so that $\nabla\alpha \leq d_i$, $i = 1, 2$, and $\nabla\alpha \leq d_1 \wedge d_2$. QED

Remark 6.11. In a polymodal setting the vectorial nabla $\nabla = \bigwedge_{i \in I} \nabla_i$ is an \emptyset -adjoint on the Lindenbaum algebra \mathbf{L} . Recalling that $\nabla^{\mathbf{L}} : \mathcal{P}_{\omega}(\mathbf{L})^I \rightarrow \mathbf{L}$, then we need to define $G_{\nabla^{\mathbf{L}}}(d)$ as a finite set of vectors (of finite subsets of \mathbf{L}), that is, $G_{\nabla^{\mathbf{L}}}(d) \subseteq_{\omega} \mathcal{P}_{\omega}(\mathbf{L})^I$. To this aim, we proceed as before: we first define it on weakly irreducible elements and then we extend its definition to meets of weakly irreducible elements. Now, in a polymodal setting, d is weakly irreducible if it can be written as

$$d = \bigvee \Pi \vee \bigvee_{i \in I} (\diamond_i b_i \vee \bigvee \square_i C_i).$$

For d weakly irreducible, we let

$$G_{\nabla^{\mathbf{L}}}(d) = \mathbf{G}^{\Pi}(d) \cup \bigcup_{i \in I} \mathbf{G}_i(d)$$

where

$$\begin{aligned} \beta \in \mathbf{G}^{\Pi}(d) &\text{ iff } \beta_i \in G^{\Pi}(d) \text{ for all } i \in I, \\ \beta \in \mathbf{G}_i(d) &\text{ iff } \beta_i \in \{ \{b_i, \top\} \} \cup \{ \{b_i \vee c\} \mid c \in C_i \} \text{ and } \beta_k \in \{ \{\top\}, \emptyset \} \text{ for } k \neq i, \end{aligned}$$

where $G^\Pi(d)$ is defined as in equation (36).

To see that this is a correct definition, it suffices to observe that $\nabla\beta \leq d$ if $\beta \in G_{\nabla\mathbf{L}}(d)$, and that, conversely, $\nabla^{\mathbf{L}}\alpha \leq d$ implies the existence of some $\beta \in G_{\nabla\mathbf{L}}(d)$ such that $\alpha_i \leq \beta_i$ for all $i \in I$. The first of these two observations is straightforward; the second follows from an analog to Claim 1 in the proof of Theorem 6.10 stating that by the rigidity of \mathbf{L} , $\nabla^{\mathbf{L}}\alpha \leq d$ implies one of the following three cases: (1) either $\bigvee \Pi = \top$, or (2) there exists $i \in I$ and $a \in \alpha_i$ such that $a \leq b_i$, or (3) there exists $i \in I$ and $c \in C_i$ such that $\bigvee \alpha_i \leq b_i \vee c$.

To extend the definition of $G_{\nabla\mathbf{L}}$ to all elements of \mathbf{L} , we let

$$\begin{aligned} \beta \in G_{\nabla\mathbf{L}}(\top) &\text{ iff } \beta_i \in \{\{\top\}, \emptyset\} \text{ for all } i \in I, \\ \beta \in G_{\nabla\mathbf{L}}(d_1 \wedge d_2) &\text{ iff } \exists \beta^j \in G_{\nabla\mathbf{L}}(d_j), j = 1, 2, \\ &\text{ and } Z_i \in \beta_i^1 \boxtimes \beta_i^2 \text{ s.t. } \beta_i = \{b_1 \wedge b_2 \mid (b_1, b_2) \in Z_i\}. \end{aligned}$$

We leave it for the reader to verify the correctness of this definition along the ideas given for formulas (37).

As an immediate corollary of Theorem 6.10, we obtain the following.

Corollary 6.12. *The Lindenbaum algebra L is residuated, that is, each operation $\diamond_i^{\mathbf{L}} : L \rightarrow L$ is a left adjoint.*

Proof. Recall that $\diamond x = \nabla\{x, \top\}$ and observe that the correspondence $\{\cdot, \top\} : \mathbf{L} \rightarrow \mathcal{P}_\omega(\mathbf{L})$, sending $x \in \mathbf{L}$ to $\{x, \top\} \in \mathcal{P}_\omega(\mathbf{L})$ is an \emptyset -adjoint: We can define

$$G_{\{\cdot, \top\}} = \begin{cases} \{\wedge \alpha\}, & \top \in \alpha, \\ \emptyset, & \text{otherwise,} \end{cases}$$

leaving it for the reader that this definition is indeed correct. As \emptyset -adjoints compose, it follows from Theorem 6.10 that $\diamond^{\mathbf{L}}$ is an \emptyset -adjoint. But then it must be a left adjoint since it preserves finite joins, see [35, Proposition 6.3]. QED

Remark 6.13. In passing we note that the same results apply to the free *modal* algebra, which can be identified with the Lindenbaum-Tarski algebra of the basic (poly-)modal logic \mathbf{K} . In particular, simplified versions of the proofs given here will show that the coalgebraic modality of the free modal algebra is an \emptyset -adjoint.

In order to prove the main result of this section, viz., Proposition 6.17 dealing with constructiveness of simple systems of equations, we need to adapt the definition of the cover modality so that it has as its domain a product set of the form A^Z . Formally, for a finite set of variables Z , we introduce the operation $\nabla_Z^A : A^Z \rightarrow A$, defined by the formula

$$\nabla_Z(\mathbf{v}) = \bigwedge_{z \in Z} \diamond \mathbf{v}_z \wedge \square \bigvee_{z \in Z} \mathbf{v}_z.$$

If $Y \subseteq Z$, then we shall write $\nabla_Y^A : A^Z \rightarrow A$ for the compose $\nabla_Y^A \circ \pi_Y$, where $\pi_Y : A^Z \rightarrow A^Y$ denotes the obvious projection.

It is not difficult to see that $\nabla_Z^A = \nabla^A \circ S_Z^A$, where $S_Z^A : A^Z \longrightarrow \mathcal{P}_\omega(A)$ transforms a vector into a finite subset, $S_Z^A(\mathbf{v}) = \{\mathbf{v}_z \mid z \in Z\}$. Now, S_Z^A is an \mathcal{O} -adjoint for every modal algebra A , since we can define

$$G_{S_Z^A}(\beta) = \{\mathbf{v}^R \mid R \in Z \bowtie \beta\}, \text{ with } \mathbf{v}_z^R = \bigwedge \{b \in \beta \mid zRb\}. \quad (38)$$

The first part of the next Lemma is an immediate consequence of our previous observations. The second part of the Lemma will be needed when arguing about constructiveness of a simple system of equations.

Lemma 6.14. *For every pair (Z, Y) with Z a finite set of variables Z and $Y \subseteq Z$, the following holds:*

1. *The vectorial cover modality $\nabla_Y^L : L^Z \longrightarrow L$ is an \mathcal{O} -adjoint on the Lindenbaum algebra L .*
2. *Let $d = \bigwedge_{\ell=1, \dots, n} d_\ell$, where each d_ℓ is a weakly irreducible element of the form $\bigvee \Lambda_\ell \vee \diamond b_\ell \vee \bigvee \square C_\ell$. If $\mathbf{v} \in G_{\nabla_Y^L}(d)$ and $z \in Z$, then \mathbf{v}_z is a conjunction of elements from the set $\bigcup_{\ell=1, \dots, n} \{b_\ell\} \cup \{b_\ell \vee c \mid c \in C_\ell\}$.*

Proof. The first part of the Lemma is an immediate consequence of the facts that π_Y^L, S_Y^L , and ∇^L are all \mathcal{O} -adjoints, that \mathcal{O} -adjoints compose, and that $\nabla_Y^L = \nabla^L \circ S_Y^L \circ \pi_Y^L$:

$$L^Z \xrightarrow{\pi_Y^L} L^Y \xrightarrow{S_Y^L} \mathcal{P}_\omega(L) \xrightarrow{\nabla^L} L.$$

For the second part of the Lemma we argue as follows. Let D be the set $\bigcup_{\ell=1, \dots, n} \{b_\ell\} \cup \{b_\ell \vee c \mid c \in C_\ell\}$. From the equations (35) and (37) we prove, by induction on n , that if $a \in \alpha \in G_{\nabla^L}(\bigwedge d_\ell)$, then a is a (possibly empty) conjunction of elements from D . Then we use the formula that witnesses that \mathcal{O} -adjoints compose, $G_{g \circ f}(d) = \bigcup_{c \in G_f(d)} G_g(c)$ and the expressions for $G_{S_Y^L}$ and $G_{\pi_Y^L}$. From equation (38) it is immediately seen that if $\mathbf{v} \in G_{S_Y^L \circ \nabla^L}(d)$ and $y \in Y$, then \mathbf{v}_y is a conjunction of elements from D . We leave it for the reader to determine an expression for $G_{\pi_Y^L}$ and to conclude that \mathbf{v}_z is a conjunction of elements from D if $\mathbf{v} \in G_{\nabla_Y^L}(d)$ and $z \in Z$. QED

On the basis of the results obtained until now, we can use Proposition 6.3 of [35] to prove that, if $T = \langle Z, \{t_z \mid z \in Z\} \rangle$ is a *simple* system of equations, then $T_{\mathbf{v}}^L : L^Z \longrightarrow L^Z$ is an \mathcal{O} -adjoint, for each parameter \mathbf{v} . However, our real goal is to argue that $T_{\mathbf{v}}^L$ is a *finitary* \mathcal{O} -adjoint and hence, by Proposition 2.14, that the least fixpoint $\mu_Z.T_{\mathbf{v}}^L$ is constructive. To this goal, we shift the focus of our discussion from \mathcal{O} -adjoints to families of \mathcal{O} -adjoints.

Definition 6.15. A modal algebra A is said to be ∇ -*finitary* if any family \mathcal{F} of the form

$$\mathcal{F} = \{k_\ell \wedge \nabla_{Y_\ell}^A : A^Z \longrightarrow A \mid \ell = 1, \dots, n\} \quad (39)$$

is a finitary family of \mathcal{O} -adjoints – where Z is a finite set of variables and for each $\ell = 1, \dots, n$ $Y_\ell \subseteq Z$ and $k_\ell \in A$.

Proposition 6.16. *The Lindenbaum algebra L is ∇ -finitary.*

Proof. Let us define the Fischer-Ladner closure $FL(\varphi)$ of a formula φ as the least set of formulas satisfying the following equations:

$$\begin{aligned} FL(p) &= \{p\} \\ FL(\neg\varphi) &= \{\neg\varphi\} \cup FL(\varphi) \\ FL(\varphi_1 \wedge \varphi_2) &= \{\varphi_1 \wedge \varphi_2\} \cup FL(\varphi_1) \cup FL(\varphi_2) \\ FL(\diamond\varphi) &= \{\diamond\varphi\} \cup FL(\varphi) \\ FL(\#_\gamma(\varphi)) &= \{\#_\gamma(\varphi)\} \cup FL(\gamma(\#_\gamma(\varphi), \varphi)). \end{aligned}$$

It is a standard argument to prove that $FL(\varphi)$ is a finite set.

Next, consider a family \mathcal{F} as in equation (39). We shall first argue that the family of \emptyset -adjoints

$$\mathcal{F}' = \{ \nabla_{Y_\ell}^L : L^Z \longrightarrow L \mid \ell = 1, \dots, n \} \cup \{ k_\ell \wedge \cdot : L \longrightarrow L \mid \ell = 1, \dots, n \}$$

is finitary. To this goal, we fix an arbitrary formula φ_0 and need to construct a finite set V such that $[\varphi_0] \in V$ and V is \mathcal{F}' -closed. We begin by fixing formulas φ_ℓ , $\ell = 1, \dots, n$, such that $[\varphi_\ell] = k_\ell$. Next we let $V \subseteq L$ be the Boolean algebra generated by the set $\bigcup_{\ell=0, \dots, n} \{ [\psi] \mid \psi \in FL(\varphi_\ell) \}$. Clearly V is finite and contains $[\varphi_0]$. In order to show that V is \mathcal{F}' -closed, we observe first that V is generated by the modal equivalence classes, i.e. equivalence classes $[\psi]$, where $\psi \in \bigcup_{\ell=0, \dots, n} FL(\varphi_\ell)$ is such that $\psi = p$ is a propositional variable or $\psi = \diamond\psi'$ for some $\psi' \in \bigcup_{\ell=0, \dots, n} FL(\varphi_\ell)$. Hence, if $d \in V$, then d is a conjunction of disjunctions of modal equivalence classes and their negations. Therefore d is a conjunction of weakly irreducible elements of the form $\bigvee \Lambda \vee \diamond b \vee \bigvee_{c \in C} \square c$ with $\{b\} \cup C \subseteq V$.

We can now argue that V is \mathcal{F}' -closed. If $d \in V$, then write d as a conjunction of weakly irreducible elements d_j of the form $\bigvee \Lambda_j \vee \diamond b_j \vee \bigvee_{c \in C_j} \square c$ with $\{b_j\} \cup C_j \subseteq V$. Then, by Lemma 6.14, if $z \in Z$ and $\mathbf{v} \in G_{\nabla_{Y_\ell}^L}(d)$, then $\mathbf{v}_z \in V$, since \mathbf{v}_z is a conjunction of elements that belong to $\bigcup_j \{b_j\} \cup \{b_j \vee v \mid c \in C_j\}$, so that $\mathbf{v}_z \in V$. This shows that V is $\nabla_{Y_\ell}^L$ -closed. Similarly, since the map $(k_\ell \wedge \cdot)$ is left adjoint to the map $(\neg k_\ell \vee \cdot)$, $G_{k_\ell \wedge \cdot}(d) = \{\neg k_\ell \vee d\} = \{\neg[\varphi_\ell] \vee d\} \subseteq V$ provided $d \in V$. This shows that V is also $(k_\ell \wedge \cdot)$ -closed, and therefore we have established that \mathcal{F}' is a finitary family.

Finally, since finitary families are closed under composition and a sub-family of a finitary family is a finitary family, see Proposition 2.15, we may deduce that \mathcal{F} is itself a finitary family of \emptyset -adjoints. QED

Proposition 6.17. *Let $T = \langle Z, \{t_z \mid z \in Z\} \rangle$ be a simple system of equations, let A be a ∇ -finitary modal algebra, and let \mathbf{v} be a set of parameters for T . Then $\mu_Z.T_{\mathbf{v}}^A$, if existing, is constructive.*

Proof. Let T , A , and \mathbf{v} be as stated, and recall that each $(t_z)_{\mathbf{v}}^A$ is of the form $\bigvee_{\ell \in L} k_\ell \wedge \nabla_{Y_\ell}^A$. Since families of finitary \emptyset -adjoints can be closed under joins, it follows from the assumptions that the family

$$\{(t_z)_{\mathbf{v}}^A : A^Z \longrightarrow A \mid z \in Z\}$$

is a family of finitary \emptyset -adjoints. Hence, by Proposition 2.15, $T_{\mathbf{v}}^A$ is itself a finitary \emptyset -adjoint, and hence its least fixpoint, if existing, is constructive by Proposition 2.14. QED

As a specific example of Proposition 6.17, we see that on a regular \sharp -algebra, the modal system T_{γ}^+ is constructive. Together with the results in Section 4, this is the key to prove constructiveness of the least fixpoint \sharp_{γ} itself.

6.3 Constructiveness of \mathbf{L}

We have now gathered sufficient material to prove the main result of this section.

Theorem 6.18. *The Lindenbaum algebra L of the system $\mathbf{K}_{\sharp}^+(\Gamma)$ is constructive. If every $\gamma \in \Gamma$ is equivalent to an untied formula, then the Lindenbaum algebra L of the simpler system $\mathbf{K}_{\sharp}(\Gamma)$ is constructive.*

Proof. For the first part of the statement we argue as follows. We have seen in Section 5 that L is the free regular modal \sharp -algebra. In particular L is regular and $(T_{\gamma}^+)_{\mathbf{v}}$ has a least fixpoint $\mu_Z.(T_{\gamma}^+)_{\mathbf{v}}$ for each parameter $\mathbf{v} \in L^P$. Since T_{γ}^+ is a *simple* system of equations, it follows from Proposition 6.17 that each of these least fixpoints $\mu_Z.(T_{\gamma}^+)_{\mathbf{v}}$ is constructive. But then it follows by successive applications of the Propositions 4.2 and 4.1 that all parametrized least fixpoints on L of T_{γ} and γ , respectively, are constructive as well.

The second part is even simpler: L is, in this case, the free modal \sharp -algebra. Being rigid, the operations that can be constructed using substitution starting from ∇^L , constants, conjunctions with constants, and disjunctions, are finitary \emptyset -adjoints on L . If $\gamma \in \Gamma$ – so that γ is untied – then $\gamma^L(x, \mathbf{p})$ is among these operations. Thus, $\gamma_{\mathbf{p}}^L(x)$ is a finitary \emptyset -adjoint and its least fixpoint is constructive. QED

7 A representation theorem

The aim of this section is to prove that every countable modal \sharp -algebra A in which each diamond modality is residuated, and each fixpoint connective is constructive, can be embedded in a Kripke \sharp -algebra (Theorem 7.1 below). Our proof method consists of building a representation for A via a step-by-step approximation process and can be seen as a version of more general game-based methods for building structures in model theory (see [16, 15] for an overview). It has a long history in modal and algebraic logic, see [27, 28, 8] for some early references.

Theorem 7.1. *Let A be a countable modal \sharp -algebra. Assume that each \sharp_{γ} is constructive on A , and that each \diamond_i^A is residuated. Then A can be embedded in a Kripke \sharp -algebra.*

Fix an algebra A as in Theorem 7.1. For simplicity we restrict attention to a language with a single diamond \diamond , and a single fixpoint connective \sharp . We let $\gamma(x, \mathbf{p})$ denote the associated formula of \sharp , where $\mathbf{p} = (p_1, \dots, p_n)$. We will say that $a \in A$ is *nonzero* if $a \neq \perp$. The main lemma in the proof of Theorem 7.1 is the following.

Lemma 7.2. *For each nonzero $a \in A$ there is a Kripke frame S_a and a modal \sharp -homomorphism $\rho_a : A \rightarrow S_a^\sharp$ such that $\rho_a(a) > \perp$.*

The key notion involved in the step-by-step approximation process leading up to Lemma 7.2 is that of a *network*. Let ω^* denote the set of finite sequences of natural numbers. We denote concatenation of such sequences by juxtaposition, and write ϵ for the empty sequence. If $t = sk$ for some $k \in \omega$ we say that s is the parent of t and write either $s = t^-$ or $s \triangleleft t$. A *tree* is a subset T of ω^* which is both downward and leftward closed; that is, if $t \neq \epsilon$ belongs to T , then so does t^- , and if $sm \in T$ then $sk \in T$ for all $k < m$. Obviously, a tree T , together with the relation \triangleleft , forms a Kripke frame; this frame will also be denoted as T , and its complex \sharp -algebra, as T^\sharp .

An *A-network* is a pair $N = \langle T, L \rangle$ such that T is a tree and $L : T \rightarrow \mathcal{P}(A)$ is some labelling. Such a network N induces a map $r_N : A \rightarrow \mathcal{P}(T)$, given by

$$r_N(a) := \{ t \in T \mid a \in L(t) \}. \quad (40)$$

The aim of the proof will be to construct, for an arbitrary nonzero $a \in A$, a network $N = \langle T, L \rangle$, with $a \in L(\epsilon)$, and such that r_N is a modal \sharp -homomorphism from A to T^\sharp . We need some definitions.

A network $N = \langle T, L \rangle$ is called *locally coherent* if $\bigwedge X > \perp$, whenever X is a finite subset of $L(t)$ for some $t \in T$; *modally coherent* if $\bigwedge X \wedge \diamond \bigwedge Y > \perp$, for all $s, t \in T$ such that $s \triangleleft t$ and all finite subsets X and Y of respectively $L(s)$ and $L(t)$; and *coherent* if it satisfies both coherence conditions. N is *prophetic* if for every $s \in T$, and for every $\diamond a \in L(s)$, there is a *witness* $t \in T$ such that $s \triangleleft t$ and $a \in L(t)$; *decisive* if either $a \in L(t)$ or $\neg a \in L(t)$, for every $t \in T$ and $a \in A$; and *\sharp -constructive* if, for every $t \in T$, and every sequence \mathbf{a} in A such that $\sharp \mathbf{a} \in L(t)$, there is a natural number n such that $(\gamma_{\mathbf{a}}^A)^n(\perp) \in L(t)$. A network is *perfect* if it has all of the above properties.

Lemma 7.3. *If N is a perfect A-network, then $r_N : A \rightarrow T^\sharp$ is a modal \sharp -homomorphism from the modal \sharp -algebra A to the complex algebra T^\sharp of the Kripke model $\langle T, \triangleleft \rangle$.*

Clearly, we shall have that $r_N(a) \neq \emptyset$ for all $a \in A$ for which there is a $t \in T$ with $a \in L(t)$.

Proof. Let $N = \langle T, L \rangle$ be a perfect network. It is fairly easy to derive from local coherence and decisiveness that each $L(t)$ is an *ultrafilter* of (the Boolean reduct of) A . From this it is immediate that r_N is a Boolean homomorphism.

In order to prove that r_N is a modal homomorphism, we need to show that

$$r_N(\diamond a) = \{ t \in T \mid t \triangleleft s \text{ for some } s \in r_N(a) \}, \quad (41)$$

for all $a \in A$. The inclusion \subseteq holds because N is prophetic. For the opposite inclusion, assume that $t \triangleleft s$ and $a \in L(s)$. Suppose for contradiction that $t \notin r_N(\diamond a)$, so that $\diamond a \notin L(t)$. Then by decisiveness, $\neg \diamond a \in L(t)$. This gives the desired contradiction with the assumed modal coherence of N , so that indeed we may conclude that (41) holds.

From this it follows that, for all sequences $\mathbf{a} \in A^n$, and all modal formula φ :

$$\varphi^{T^\sharp}(r_N(\mathbf{a})) = r_N(\varphi^A(\mathbf{a})), \quad (42)$$

where for a vector $\mathbf{a} = (a_1, \dots, a_n)$ $r_N(\mathbf{a})$ denotes – here and in the sequel – the vector $(r_N(a_1), \dots, r_N(a_n))$.

In particular, for $\varphi = \gamma$, (42) implies that for all \mathbf{b} :

$$r_N(\sharp^A \mathbf{b}) = r_N(\gamma^A(\sharp^A \mathbf{b}, \mathbf{b})) = \gamma^{T^\sharp}(r_N(\sharp^A \mathbf{b}), r_N(\mathbf{b})).$$

In other words, $r_N(\sharp^A \mathbf{b})$ is a *fixpoint* of the map $\gamma_{r_N(\mathbf{b})}^{T^\sharp}$. But we can also prove that $r_N(\sharp^A \mathbf{b})$ is the ω -approximation of $\sharp^{T^\sharp}(r_N(\mathbf{b}))$. To see why this is so, we start from the definition of $r_N(\sharp^A \mathbf{b})$:

$$r_N(\sharp^A \mathbf{b}) = \{t \in T \mid \sharp \mathbf{b} \in L(t)\}. \quad (43)$$

Since $L(t)$ is an ultrafilter and the network T is \sharp -constructive, $\sharp \mathbf{b} \in L(t)$ if and only if, for some n , $(\gamma_{\mathbf{b}}^A)^n(\perp) \in L(t)$, and hence

$$r_N(\sharp^A \mathbf{b}) = \bigcup_{n < \omega} \{t \in T \mid (\gamma_{\mathbf{b}}^A)^n(\perp) \in L(t)\}. \quad (44)$$

Recall that, by definition of r_N , $(\gamma_{\mathbf{b}}^A)^n(\perp) \in L(t)$ if and only if $t \in r_N((\gamma_{\mathbf{b}}^A)^n(\perp))$. Moreover, a straightforward inductive proof, on the basis of (42), will show that

$$r_N((\gamma_{\mathbf{b}}^A)^n(\perp)) = \left(\gamma_{r_N(\mathbf{b})}^{T^\sharp}\right)^n(\perp).$$

Hence equation (44) becomes

$$r_N(\sharp^A \mathbf{b}) = \bigcup_{n < \omega} \left(\gamma_{r_N(\mathbf{b})}^{T^\sharp}\right)^n(\perp^{T^\sharp}).$$

But if $r_N(\sharp^A \mathbf{b})$ is both a fixpoint of the map $\gamma_{r_N(\mathbf{b})}^{T^\sharp}$ and an ordinal approximation of $\sharp^A(r_N(\mathbf{b}))$, then it must be the *least* fixpoint of the map $\gamma_{r_N(\mathbf{b})}^{T^\sharp}$, or, equivalently,

$$r_N(\sharp^A \mathbf{b}) = \sharp^{T^\sharp}(r_N(\mathbf{b})).$$

Having shown that r_N is also a homomorphism with respect to \sharp , we have completed the proof of the Lemma. QED

From the previous Lemma it follows that, in order to prove Lemma 7.2, it suffices to construct a perfect network with $a \in L(\epsilon)$ for an arbitrary nonzero $a \in A$. Our construction will be carried out in a step-by-step process, where at each stage we are dealing with a finite *approximation* of the final network. Since these approximations are not perfect themselves, they will suffer from certain *defects*. We will only be interested in those defects that can be *repaired* in the sense that the network can be extended to a bigger version that is lacking the defect.

Formally we define a *defect* of a network $N = \langle T, L \rangle$ to be an object d of one of the following three kinds:

1. $d = (t, a, \neg)$, with $t \in T$ and $a \in A$ such that neither a nor $\neg a$ belongs to $L(t)$,

2. $d = (t, a, \diamond)$, with $t \in T$ and $a \in A$ such that $\diamond a \in L(t)$, but there is no witness s such that $t \triangleleft s$ and $a \in L(s)$,
3. $d = (t, \mathbf{a}, \sharp)$, with $t \in T$ and $\mathbf{a} \in A^n$ such that $\sharp \mathbf{a} \in L(t)$, but there is no $n \in \omega$ such that $(\gamma_{\mathbf{a}}^A)^n(\perp) \in L(t)$.

These three types of defects witness a network's failure to be decisive, prophetic, and \sharp -constructive, respectively.

In our proof we will construct a perfect network as a limit of coherent networks, one by one repairing the defects of the approximants. In order to guarantee the coherence of these approximants in the long run, we need them to satisfy a stronger, *global* version of coherency. To define this notion we extend the *local* labelling function L of the network to a global one, \tilde{L} . This global labelling gathers all relevant information concerning the network at one single node. Since N is finite, it is straightforward to define such a global labelling map for the root ϵ of the tree: if we let

$$\Delta_{\downarrow}(t) := \bigwedge L(t) \wedge \bigwedge_{t \triangleleft s} \diamond \Delta_{\downarrow}(s),$$

then the set $\Delta_{\downarrow}(\epsilon)$ on its own collects all relevant information from the full network. The *residuateness* of the modality \diamond allows us to access the global information on the network at *each* of its nodes, not just at the root. The resulting labelling $\tilde{L} : T \rightarrow A$ will considerably simplify the process of repairing defects.

Turning to the technical details, for the definition of \tilde{L} we use the *conjugate* of \diamond , which can be defined as the unique map $\blacklozenge : A \rightarrow A$ satisfying

$$a \wedge \diamond b > \perp \text{ iff } \blacklozenge a \wedge b > \perp, \quad (45)$$

for all $a, b \in A$. This map exists by the fact that \diamond is residuated; in fact, it is the Boolean dual of the residual (or right adjoint) of \diamond . Using this operation \blacklozenge , we can define the global labelling \tilde{L} as follows:

$$\begin{aligned} \tilde{L}(t) &:= \Delta_{\downarrow}(t) \wedge \Delta_{\uparrow}(t), \\ \Delta_{\downarrow}(t) &:= \bigwedge L(t) \wedge \bigwedge_{t \triangleleft s} \diamond \Delta_{\downarrow}(s), \\ \Delta_{\uparrow}(t) &:= \begin{cases} \top & \text{if } t = \epsilon, \\ \blacklozenge(\Delta_{\uparrow}(t^-) \wedge \Delta_{\downarrow,-t}(t^-)) & \text{otherwise,} \end{cases} \\ \Delta_{\downarrow,-u}(t) &:= \bigwedge L(t) \wedge \bigwedge_{t \triangleright s, s \neq u} \diamond \Delta_{\downarrow}(s). \end{aligned}$$

The idea behind this definition is straightforward: for $\tilde{L}(t)$, we start by collecting the local information $\bigwedge L(t)$ and then move on to t 's neighbors, both its predecessor (with $\Delta_{\uparrow}(t)$) and its successors (with $\Delta_{\downarrow}(s)$). The role of $\Delta_{\downarrow,-u}$ is to ensure termination of the procedure, avoiding a loop between $\Delta_{\uparrow}(t)$ and $\Delta_{\downarrow}(u)$ when $t \triangleleft u$.

Alternatively, we can understand the formula for $\tilde{L}(t)$ as follows. Given $t \in T$, we consider the unoriented tree T' which is obtained by forgetting the orientation of the edges of the form

$u \triangleleft v$. Using a basic result in graph theory, we obtain a unique new orientation \rightarrow on T' by taking t as a new root. Observe that $u \rightarrow v$ implies that either $u \triangleleft v$ or $v \triangleleft u$. Then the formula for $\tilde{L}(u)$ can be defined inductively on the basis of the new orientation, analogous to the definition of $\Delta_{\downarrow}(u)$, with the proviso that the conjunct contributed by a \rightarrow -successor v is modalized by \diamond if $u \triangleleft v$, and by \blacklozenge if $v \triangleleft u$. More precisely:

$$\tilde{L}(u) = \bigwedge L(u) \wedge \bigwedge \{ \diamond \tilde{L}(v) \mid u \rightarrow v \ \& \ u \triangleleft v \} \wedge \bigwedge \{ \blacklozenge \tilde{L}(v) \mid u \rightarrow v \ \& \ v \triangleleft u \}.$$

One of the key observations in the proof is the following claim.

Lemma 7.4. *Let N be a finite network. Then $\tilde{L}(s) > \perp$ iff $\tilde{L}(t) > \perp$, for any $s, t \in N$.*

Proof. It clearly suffices to prove the following special case:

$$\tilde{L}(t) > \perp \text{ iff } \tilde{L}(t^-) > \perp, \quad (46)$$

for an arbitrary $t \neq \epsilon$. But it is straightforward to derive from the definitions that

$$\tilde{L}(t^-) = \Delta_{\uparrow}(t^-) \wedge \Delta_{\downarrow, -t}(t^-) \wedge \diamond \Delta_{\downarrow}(t),$$

and

$$\tilde{L}(t) = \blacklozenge(\Delta_{\uparrow}(t^-) \wedge \Delta_{\downarrow, -t}(t^-)) \wedge \Delta_{\downarrow}(t).$$

Hence, (46) follows from the conjugacy of \diamond and \blacklozenge : simply take $a = \Delta_{\uparrow}(t^-) \wedge \Delta_{\downarrow, -t}(t^-)$ and $b = \Delta_{\downarrow}(t)$ in (45). QED

Call a finite network $N = \langle T, L \rangle$ *globally coherent* if $\tilde{L}(t) > \perp$ for all $t \in T$. We can now prove our *repair lemma*. We say that N' *extends* N , notation: $N \leq N'$, if $T \subseteq T'$ and $L(t) \subseteq L'(t)$ for every $t \in T$.

Lemma 7.5 (Repair Lemma). *Let $N = \langle T, L \rangle$ be a globally coherent A -network. Then for any defect d of N there is a globally coherent extension N^d of N which lacks the defect d .*

Proof. We will take action depending on the type of the defect d . In each case we will make heavily use of the global extension \tilde{L} of L .

1. If $d = (t, a, \neg)$ is a defect of the first kind, then we define $N^d := \langle T, L^d \rangle$, where $L^d(s) := L(s)$ for $s \neq t$, while we put

$$L^d(t) := \begin{cases} L(t) \cup \{a\} & \text{if } \tilde{L}^N(t) \wedge a > \perp, \\ L(t) \cup \{\neg a\} & \text{if } \tilde{L}^N(t) \wedge \neg a > \perp. \end{cases}$$

Then clearly the triple (t, a, \neg) is no longer a defect, and so all that is left to show is the global coherence of N^d . But since $\tilde{L}^N(t) > \perp$ by assumption, we will have either $\tilde{L}^N(t) \wedge a > \perp$ or $\tilde{L}^N(t) \wedge \neg a > \perp$. It is easy to check that in either case, we have $\tilde{L}^{N^d}(t) = \tilde{L}^N(t) \wedge x$ with $x \in \{a, \neg a\}$, and from this coherence follows easily.

2. Now suppose that $d = (t, a, \diamond)$ is a type 2 defect. Let k be the least number such that $tk \notin T$, and define $N^d := \langle T^d, L^d \rangle$, where $T^d = T \cup \{tk\}$, and L^d is given by putting $L^d(s) := L(s)$ for $s \neq t$, while $L^d(tk) := \{a\}$. In this case it is easy to prove that $\tilde{L}^{N^d}(t) = \tilde{L}^N(t)$, so N^d is certainly globally coherent. It is likewise simple to see that (t, a, \diamond) is no longer a defect of N .
3. Finally, suppose that $d = (t, \mathbf{a}, \sharp)$ is a defect of the third kind. By global coherency we have that $\tilde{L}^N(t) > \perp$. Suppose for contradiction that $\tilde{L}^N(t) \wedge (\gamma_{\mathbf{a}}^A)^n(\perp) = \perp$ for all numbers n . Then for all n we have $(\gamma_{\mathbf{a}}^A)^n(\perp) \leq \neg \tilde{L}^N(t)$, and so by constructiveness of \sharp on A it follows that $\sharp^A \mathbf{a} \leq \neg \tilde{L}^N(t)$. But this contradicts the fact that N is coherent. It follows that $\tilde{L}^N(t) \wedge (\gamma_{\mathbf{a}}^A)^n(\perp) > \perp$ for some natural number n . Now proceed as in the first case, defining $L^d(t) := L(t) \cup \{(\gamma_{\mathbf{a}}^A)^n(\perp)\}$.

QED

Lemma 7.6. *Every globally coherent A -network can be extended to a perfect network.*

Proof. We will define a sequence of networks $N = N_0 \leq N_1 \leq N_2 \leq \dots$ such that for each $i \in \omega$ and each defect d of N_i there is a $j > i$ such that d is not a defect of N_j .

For the details of this construction, define

$$D := \omega^* \times A \times \{\neg, \diamond\} \cup \omega^* \times A^n \times \{\sharp\}.$$

Informally we shall say that D is the set of *potential defects*. Clearly, since D is countable, we may assume the existence of an enumeration $(d_n)_{n < \omega}$ such that every element of D occurs infinitely often.

Now we set

$$\begin{aligned} N_0 &:= N \\ N_{i+1} &:= \begin{cases} N_i^{d_i} & \text{if } d_i \text{ is actually a defect of } N_i, \\ N_i & \text{otherwise.} \end{cases} \end{aligned}$$

Finally, define $N' := \langle T', L' \rangle$, with $T' := \bigcup_{i < \omega} T_i$ and, for each $t \in T'$, $L'(t) := \bigcup_{i < \omega} L_i(t)$. It is then straightforward to verify that N' is a perfect extension of N . For instance, suppose for contradiction that N' would have some defect d . It readily follows from the definitions that there must be some approximation N_k in the sequence for which d is also a defect. But then the next time i such that $d = d_i$, this defect will be repaired. As a consequence, d is not a defect of N_{d_i+1} , and so it cannot be a defect of N' either. This provides the desired contradiction. QED

Proof of Lemma 7.2. Consider an arbitrary nonzero element $a \in A$, and let N_a be the network $\langle \{\epsilon\}, L_a \rangle$, L_a given by $L_a(\epsilon) := \{a\}$. It is obvious that N_a is globally coherent, so Lemma 7.2 follows by a direct application of the Lemmas 7.6 and 7.3. QED

Proof of Theorem 7.1. Let S be the disjoint union of the family $\{S_a \mid \perp \neq a \in A\}$, where for each nonzero $a \in A$, S_a is given by Lemma 7.2. It is straightforward to verify that A can be embedded into the product $\prod_{a \neq \perp} S_a^\sharp$, and that this latter product is isomorphic to S^\sharp , the complex \sharp -algebra of S . QED

References

- [1] Proceedings of the Twentieth IEEE Symposium on Logic in Computer Science (LICS 2005), IEEE Computer Society Press, 2005 (2005).
- [2] A. Arnold, D. Niwiński, Rudiments of μ -calculus, No. 146 in Studies in Logic and the Foundations of Mathematics, North-Holland Publishing Co., Amsterdam, 2001.
- [3] J. Barwise, L. Moss, Vicious Circles, vol. 60 of CSLI Lecture Notes, CSLI Publications, 1996.
- [4] M. Bílková, A. Palmigiano, Y. Venema, Proof systems for the coalgebraic cover modality, in: C. Areces, R. Goldblatt (eds.), Advances in Modal Logic, Volume 7, College Publications, 2008.
- [5] P. Blackburn, M. de Rijke, Y. Venema, Modal Logic, No. 53 in Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 2001.
- [6] S. L. Bloom, Z. Ésik, Iteration theories, Springer-Verlag, Berlin, 1993.
- [7] M. Boffa, Une condition impliquant toutes les identités rationnelles, RAIRO Informatique Théorique et Applications 29 (6) (1995) 515–518.
- [8] J. Burgess, Axioms for tense logic I: ‘since’ and ‘until’, Notre Dame Journal of Formal Logic 23 (1982) 375–383.
- [9] J. Conway, Regular Algebra and Finite Machines, Chapman and Hall, 1971.
- [10] E. Emerson, J. Halpern, Decision procedures and expressiveness in the temporal logic of branching time, Journal of Computer and System Sciences 30 (1985) 1–24.
- [11] E. A. Emerson, Temporal and modal logic, in: Handbook of theoretical computer science, Vol. B, Elsevier, Amsterdam, 1990, pp. 995–1072.
- [12] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning about knowledge, MIT Press, Cambridge, MA, 1995.
- [13] K. Fine, Normal forms in modal logic, Notre Dame Journal of Formal Logic 16 (1975) 229–234.
- [14] D. Harel, D. Kozen, J. Tiuryn, Dynamic Logic, MIT Press, Cambridge, MA, 2000.
- [15] R. Hirsch, I. Hodkinson, Relation Algebras by Games, No. 147 in Studies in Logic and the Foundations of Mathematics, North-Holland Publishing Co., Amsterdam, 2002.
- [16] W. Hodges, Model Theory, Cambridge University Press, 1993.
- [17] D. Janin, Automata, tableaux and a reduction theorem for fixpoint calculi in arbitrary complete lattices, in: Proceedings of the Twelfth Annual IEEE Symposium on Logic in Computer Science (LICS 1997), IEEE Computer Society Press, 1997.

- [18] D. Janin, I. Walukiewicz, Automata for the modal μ -calculus and related results, in: Proc. MFCS'95, Springer, Berlin, 1995, INCS 969.
- [19] D. Kozen, Results on the propositional μ -calculus, Theoretical Computer Science 27 (3) (1983) 333–354.
- [20] D. Kozen, A completeness theorem for Kleene algebras and the algebra of regular events, Inform. and Comput. 110 (2) (1994) 366–390.
- [21] D. Kozen, R. Parikh, An elementary proof of the completeness of PDL, Theoretical Computer Science 14 (1981) 113–118.
- [22] D. Krob, Complete systems of B-rational identities, Theoret. Comput. Sci. 89 (2) (1991) 207–343.
- [23] C. Kupke, A. Kurz, Y. Venema, A complete coalgebraic logic, in: C. Areces, R. Goldblatt (eds.), Advances in Modal Logic, Volume 7, College Publications, 2008.
- [24] C. Kupke, Y. Venema, Closure properties of coalgebra automata, in: LICS 2005 [1], pp. 199–208.
- [25] C. Kupke, Y. Venema, Coalgebraic automata theory: basic results, Logical Methods in Computer Science.
- [26] M. Lange, C. Stirling, Focus games for satisfiability and completeness of temporal logic, in: Proceedings of the Sixteenth Annual IEEE Symposium on Logic in Computer Science (LICS 1997), IEEE Computer Society Press, 2001.
- [27] R. Lyndon, The representation of relation algebras, Annals of Mathematics 51 (1950) 707–729.
- [28] R. Maddux, Some varieties containing relation algebras, Transaction of the American Mathematical Society 272 (1982) 501–526.
- [29] J.-J. Meyer, W. van der Hoek, Epistemic Logic for AI and Computer Science, No. 41 in Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 1995.
- [30] L. Moss, Coalgebraic logic, Annals of Pure and Applied Logic 96 (1999) 277–317, (Erratum published *Ann.P.Appl.Log.* 99:241–259, 1999).
- [31] A. Palmigiano, Y. Venema, Nabla algebras and Chu spaces, in: Algebra and Coalgebra in Computer Science (CALCO 2007), Springer-Verlag, Berlin, 2007, INCS 4624.
- [32] J. Rutten, Universal coalgebra: A theory of systems, Theoretical Computer Science 249 (2000) 3–80.
- [33] L. Santocanale, On the equational definition of the least prefixed point, Theoretical Computer Science 295 (1-3) (2003) 341–370.
- [34] L. Santocanale, Completions of μ -algebras, in: LICS 2005 [1], pp. 219–228.

- [35] L. Santocanale, Completions of μ -algebras, *Annals of Pure and Applied Logic* 154 (2008) 27–50.
- [36] K. Segerberg, A completeness theorem in the modal logic of programs, in: T. Traczyk (ed.), *Universal Algebra and Applications*, vol. 9 of Banach Centre Publications, PWN–Polish Scientific Publishers, 1982, pp. 31–46.
- [37] A. Tarski, A lattice-theoretical fixpoint theorem and its applications, *Pacific Journal of Mathematics* 5 (1955) 285–309.
- [38] Y. Venema, Algebras and coalgebras, in: P. Blackburn, J. van Benthem, F. Wolter (eds.), *Handbook of Modal Logic*, Elsevier, 2006, pp. 331–426.
- [39] Y. Venema, Automata and fixed point logic: a coalgebraic perspective, *Information and Computation* 204 (2006) 637–678.
- [40] I. Walukiewicz, Completeness of Kozen’s axiomatisation of the propositional μ -calculus, *Inform. and Comput.* 157 (1-2) (2000) 142–182, (An earlier version of the proof appeared in the Proceedings of LICS 1995.).