Modal Logic and non-well-founded Set Theory: translation, bisimulation, interpolation.

ILLC Dissertation Series 1998-4

institute *for* logic, language *and* computation

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam
phone: +31-20-5256051
fax: +31-20-5255101
e-mail: illc@wins.uva.nl

Modal Logic and non-well-founded Set Theory: translation, bisimulation, interpolation.

Academisch Proefschrift

ter verkrijging van de graad van doctor aan de
Universiteit van Amsterdam,
op gezag van de Rector Magnificus
prof.dr. J.J.M. Franse
ten overstaan van een door het college voor promoties ingestelde
commissie in het openbaar te verdedigen in de
Aula der Universiteit
op vrijdag 4 december 1998 te 9.00 uur

door

Giovanna D'Agostino

geboren te Rome.

Promotor: Prof.dr. J.F.A.K. van Benthem
Faculteit Wiskunde, Informatica, Natuurkunde en Sterrenkunde
Universiteit van Amsterdam
Plantage Muidergracht 24
1018 TV Amsterdam

Co-promotor: Prof.dr. A. Policriti
Facolta' di Scienze MM. FF. NN.
Università di Udine
Via delle Scienze 206,
33100 Udine, Italy

Printed and bound by Print Partners Ipskamp, Amsterdam.

ISBN: 90-5776-014-2

Contents

Acknowledgments

Many people contributed in one way or another to help me in writing this dissertation.
First of all, I would like to thank Johan van Benthem for a number of reasons. He gave me the opportunity of following a PhD program at the ILLC; he was a source of inspiration throughout all these years and his questions often led to new topics of research; he was incredibly patient, especially in the last phase: he made me rewrite and rearrange the material in this thesis a lot of times and I sometimes was not very happy with it... but at the end I must admit that the final version is certainly more readable then the first one (which was unreadable).

I owe the same gratitude to Alberto Policriti, my second supervisor. Johan and Alberto both scolded me, but with different styles. Johan's way is nicely exemplified by his e-mail messages: he always started his correspondence with 'Dear Giovanna, the thesis is quite improved *but*...'; Alberto used to come to me laughing and saying 'you can't possibly have written *that!*'. However, putting their comments together (especially when they were not contradicting each other) improved this dissertation a lot, with regards to both content and presentation.

Then I want to thank my co-authors. Among them we still have Alberto and Johan and Angelo Montanari completes a group with which I worked quite a lot during these last years. Special thanks to Marco Hollenberg who is a wonderful person to work and study with.

Domenico Zambella deserves to be thanked for his help from a scientific point of view (discussing with him is always very profitable) as well as for more practical reasons: he was a sort of *angelo custode* during all my visits to Amsterdam. This role was also shared by Marco de Vries, which helped me finding my way through Holland burocracy. Finally, I want to thank Giacomo Lenzi for reading and commenting the manuscript and last but not least I want to thank my mother for being a very nice and efficient *traveling-baby-sitter*.

Chapter 1

Introduction.

The notion of *bisimulation* is simple, natural, and central for many research fields in Logic and Computer Science. It is the link that connects all the topics discussed in this dissertation.

Bisimulation was proposed independently in many areas in the seventies. Probably the first one was Modal Logic where, with the name of *p-relation* (see [12]) it applies to *Kripke models* and it is an elegant and useful tool to prove a large number of results (interpolation theorems, preservation theorems, etc.).

In Theoretical Computer Science the bisimulation relation applies to *labeled transition systems* (see [54], [52]), which is another word for Kripke models. Labeled transition systems are used to represent *processes*: the basic idea is just to interpret nodes as possible states of the processes, unary relations as properties of states, and binary relations as atomic actions that the processes may undertake. Bisimulation (and its variations) can be considered as equivalence relations on labeled transition systems: two bisimilar transition systems represent the same process.

In this area, logics are used as languages that express behavioral properties of processes. Extended modal logics, being invariant under various forms of bisimulation, turn out to be of particular interest, because we can consider properties expressed by these logics as *process-invariants*, instead of mere transition system properties.

The notion of bisimulation is a central theme also in another area: the Theory of non-well-founded Sets. It was first introduced as an axiom characterizing non-well-founded sets in [31] and since then is a fruitful notion in this field. Working with non-well-founded sets, the usual criterion for equality, the so-called *extensionality* axiom, cannot be applied: the argument becomes circular (!); on the contrary, the notion of bisimulation applies to non-well-founded sets and provides a criterion for equality. Different notions of bisimulation give rise to different theories of non-well-founded sets. In [2], Aczel compares many such theories, obtaining in this way a deep insight in the possible structure of non-well-founded sets.

Altogether bisimulation can be seen as a bridge between Modal Logic, non-well-founded Set Theory, and Process Theory, and one can fruitfully use this bridge to transfer results and techniques from one area to the others. In this dissertation we cross this bridge a number of times.

- We prove results in extended modal logics that can be used in Process Theory.
- We use extended modal formulae to *describe* different non-well-founded universes. This allows us to provide alternative definitions of classes of extended modal logics well-known by the modal logic community.
- We use non-well-founded set theories to study derivability in Modal Logic.

As one naturally expects, the choice of representing a process as an equivalence class of transition systems modulo bisimulation has a strong influence on the choice of the logics used to express properties of such processes. Essentially, we want to restrict our attention to formulae ϕ which are bisimulation invariant in the following sense: if two transition systems are in the same bisimulation class, then they must agree on ϕ. As we shall see, even though not all formulae of First Order or Monadic Second Order Logic have this property, in these environments one can try to isolate those formulae which are bisimulation invariant: the Van Benthem Theorem [12] and the Janin-Walukiewicz Theorem [41] do the job. The first one characterizes the formulae of Basic Modal Logic as those first-order properties of states that are bisimulation invariant; the second one proves a similar results in the second-order setting: the formulae of the Modal μ-Calculus are exactly the monadic second-order properties of states which are bisimulation invariant.

One of the most important aspects of these theorems lies in the fact that Basic Modal Logic and Modal μ-Calculus have desirable properties, such as a complete calculus, decidability, and the finite model property; van Benthem's and Janin-Walukiewicz's theorems guarantee that, by restricting to bisimulation invariant properties, we obtain this nice behavior without losing the expressive power of the logic we started with, First Order or Monadic Second Order. In particular, as far as Process Theory is concerned, the second-order setting is especially interesting, because well-known and much used properties of processes such as *fairness* and *termination* are expressible in it and hence, by Janin-Walukiewicz Theorem, in the Modal μ-Calculus.

Our contribution in this dissertation with respect to this area is an *interpolation* theorem for the Modal μ-Calculus. Interpolation is another desirable property that a well-behaved logic is supposed to have. Since interpolation can be proved whenever a Gentzen-style sequent calculus without cut is available, a failure of interpolation can be seen as a signal that the logic cannot have such an elegant calculus. Moreover, interpolation has a nice consequence known as the Beth property, which says that implicit and explicit definitions in the logic coincide.

Interpolation was also re-discovered in recent years by computer scientists from a more practical point of view, as a useful property in the context of modular databases. A formula ϕ can be seen as a description of a database and interpolation, in its uniform version, says that the database can be split into modules: if we submit a query dealing with a specific aspect of the database, we can restrict ourselves to querying the corresponding module.

Summarizing, the Modal μ-Calculus is a useful logic for application in Computer Science and (uniform) interpolation is a useful property of a logic. In Chapter 3 we show that Modal μ-Calculus enjoys (uniform) interpolation by using automata techniques and the so-called *bisimulation quantifiers*.

Let us now cross another bridge and go into the realm of non-well-founded sets.

As we already said, bisimulation has a well-established role in this theory and hence it is quite surprising that a modal-logic perspective acquired a place in this context only in recent times (see [7]). Using bisimulation, Modal Logic can be used to describe sets: in the most well-known axiomatization of non-well-founded sets, the theory $ZFC^- + AFA$, any set has a precise description in terms of an infinitary modal formula. Moreover, by means of this description sets can be seen as formulae and a model of $ZFC^- + AFA$ consisting of infinitary modal formulae can be built.

In addition to its fairly recent discovery, notice also that the role of Modal Logic in non-well-founded Set Theory has been considered only in connection with the theory $ZFC^- + AFA$. However important, this theory does not certainly exhaust all possible descriptions of non-well-founded sets. Other reasonable theories are known, whose study helped to understand the possible structure of such sets in more depth. Can (extended) Modal Logic be used to describe sets in other non-well-founded theories? In Chapter 4 we show that the Scott anti-foundation theory ([62], [2]) admits such a description in terms of a natural infinitary extension of the so-called *graded modal logics*. We also give some natural variations on this theme, describing other non-well-founded universes having a similar description. A central role in this respect is played by the *expansion operations*, which transform structures into bisimilar trees. The simplest expansion, the *unraveling*, can be used to describe the Scott axiom; more elaborate expansions are used to described different kinds of non-well-founded sets.

Crossing the bridge once more and re-entering the realm of Modal Logic, in Chapter 5 we study interpolation for the class of logics suggested by non-well-founded theories. This is the well-known class of graded modal logics (in its infinitary version) and the problem of interpolation for this class has already been considered by Andréka in [1]. She showed that the behavior of the logics in the class is not uniform: some logics in the class enjoy interpolation and some other do not. In this setting we prove a weak form of interpolation for the class of graded modal logics, the so-called *elementary interpolation* and prove full interpolation whenever possible. The main idea used in this context is that of a *consistency property*, which is often used in the context of infinitary logic. Following an idea proposed by van Benthem ([11]), we tune this notion over bisimulation and introduce *consistency property modulo bisimulation* to prove elementary and Craig interpolation in the class of graded modal logics.

Finally, in Chapter 6 we use sets to describe derivability in Modal Logic.

As we shall see, in non-well-founded Set Theory Kripke frames can be used to represent sets, with the accessibility relation playing the role of the inverse membership relation. From this point of view, we can define a semantics for Modal Logic in which the role of a Kripke frame is taken by the simpler concept of set. Formulae of Modal Logic are then naturally translated into set-terms representing the set of worlds in the Kripke frame which make the formula true. We obtain in this way a natural generalization of the interpretation of propositional connectives as set operations: together with the set interpretation of disjunction as a union, conjunction as intersection, and the other set interpretations of propositional connectives, we prove that we can consider the $\Box$-operator as the powerset operator (to be applied to non-well-founded sets). This leads us to a translation from Basic Modal Logic to a simple theory of non-well-founded sets, the theory Ω.

The above mentioned translation was originally proposed in the context of Automated Theorem Proving for Modal Logic. In this area, translations from modal logics into First Order Logic are often used, since they allow the use of very sophisticated and well performing theorem provers to automatically derive modal logic formulae. From this point of view, the larger the class of translatable logics, the better. However, the most used translations in the field are the standard one and variations of it, which essentially allow one to translate only the class of first-order complete logics. The above mentioned set-translation is, in this sense, a better choice: it works for all complete logics, not necessarily first-order complete. In general, we only claim that the technique for automated deduction in modal logic that opens up with the introduction of the set translation is widely applicable and genuinely new. As far as computationally related issues are concerned, we simply mention the fact that new techniques from Computable Set Theory can be used to produce theorem provers for the set theory Ω (see [16], [58]).

The completeness and soundness of the set translation was first proved in [20] by using tools of non-well-founded Set Theory. In this dissertation, we give a simpler proof based on a comparison between the standard translation and the set translation.

One of the advantage of the theory Ω lies in its simplicity: we just have axioms describing the relationship of $\in$ with the union operator, the set-difference operator, and the powerset operator. However, even though Ω is strong enough to deal with Basic Modal Logic, modern Modal Logic goes toward extensions of the basic formalism and more complex logics arise and are studied. A natural question arises: can we tune the theory Ω to deal with extended modal logics? Is the new theory still a set theory, or do we need some artificial axioms that have nothing to do with sets? What kind of extensions of Basic Modal Logic are we able to cope with?

Extended modal logics may be obtained introducing new operators and, in most cases, they have a first-order definable semantics in the language of Kripke models. Example of such operators are the difference operator, the past operator, the graded operators, and so forth. In this dissertation we consider a (second-order) logic L_2 (see [8], [10]) within which all operators with such a first-order definable semantics can be embedded. We show that we can strengthen the theory Ω in order to capture modal derivability in L_2 and that the new theory is still a genuine set theory: it is obtained by adding to Ω the operators that were introduced by Gödel to describe the universe of constructible sets used to prove the consistency of the continuum hypothesis.

Organization and origin of the chapters.

Chapter 2 gives a general introduction to the topics of this dissertation. Besides notations and basic definitions, it contains five more elaborate sections on the main themes. We have a section on bisimulation and two sections introducing the logics we are mainly interested in: the Modal μ-Calculus and the family of graded modal logics. Then we have a section on the different forms of interpolation that we shall encounter and another section discussing the role played by bisimulation in non-well-founded set theories.

Chapter 3 deals with uniform interpolation for the Modal μ-Calculus. This work is the result of a fruitful collaboration with Marco Hollenberg and, together with other results about the Modal μ-Calculus, it has been accepted for publication in the Journal of Symbolic Logic. It is now available as a preprint ([21]).

Chapter 4 discusses the relationship between non-well-founded sets and Modal Logic. It is still unpublished material.

Chapter 5 considers the problem of interpolation for graded modal logics. As the material in the preceding chapter, it is still unpublished.

Chapter 6 is about the set translation of Modal Logic and extended modal logics. This is again the result of a fruitful collaboration, this time with Johan van Benthem, Angelo Montanari, and Alberto Policriti. The various stages of the work have been published in [20], [14], [15].

Each chapter ends with brief concluding remarks and open questions. The further directions of investigation listed there are by no means exhaustive, but I hope that they do show that our intellectual bridges can support a good deal of traffic.

Finally, I wish to ask the reader a little patience with my roman-english.

Chapter 2

General preliminaries.

In this part we introduce the basic material (definitions, notations, results) that is needed to read this dissertation. In Section 2.1 we describe the logics we shall use, their semantics, and we briefly discuss translations and derivability in Basic Modal Logic. More elaborate sections follow: Section 2.2 is dedicated to a central subject of this thesis, bisimulation, while in Section 2.3 and Section 2.4 we give a brief introduction to the Modal μ-Calculus and to the family of graded modal logics, respectively. Section 2.5 regards the various forms of interpolation that we shall discuss and, finally, Section 2.6 introduces the fundamentals of non-well-founded Set Theory, with special attention to the role of bisimulation.

In these sections, proofs of stated results are given only if they are useful to understand the rest of the thesis.

2.1 Logics.

The languages considered in this thesis essentially contain a set of unary predicates, a set of binary predicates, one constant, and no function symbols. The logics we consider are extensions of Modal Logic in their natural classical environments (first-, or second-order, finitary or infinitary).

2.1.1 Languages and structures.

Fix a countable set PROP of propositional constants and a countable set A of atomic actions. If $\Sigma \subseteq$ PROP and $\Lambda \subseteq A$, we denote by $\mathcal{L} = (\Sigma, \Lambda)$ the *language* built up from Σ and Λ as follows: there is a single constant r (for 'root'), each $p \in \Sigma$ is a unary predicate symbol of $\mathcal{L}$ and for each $a \in \Lambda$ there is a binary relation symbol R_a in $\mathcal{L}$.

If $\mathcal{L} = (\Sigma, \Lambda)$ is a language, then a *structure* $\mathcal{M}$ on $\mathcal{L}$ consists of a domain $D^{\mathcal{M}}$, a root $r^{\mathcal{M}} \in D^{\mathcal{M}}$, a set of unary relations $p^{\mathcal{M}}$ on $D^{\mathcal{M}}$, for $p \in \Sigma$, and a set of binary predicates $R_a^{\mathcal{M}}$ on $D^{\mathcal{M}}$, for $a \in \Lambda$.

If $\mathcal{M}$ is an $\mathcal{L}$-structure, we denote by $(\mathcal{M}, s)$ the structure $\mathcal{M}$ with s as a new root. If $p \in$ PROP and S is a subset of the domain of the structure $\mathcal{M}$, we denote by $\mathcal{M}[p := S]$ the structure which is like $\mathcal{M}$ except that p is now interpreted as S.

If the $\mathcal{L}$-structure $\mathcal{M}$ is clear from the context, we write $s \xrightarrow{a} t$ for $(s,t) \in R_a^{\mathcal{M}}$. For each action a, the set of a-successors of a point s in a structure $\mathcal{M}$ is defined as: $R_a(s) := \{t \mid s \xrightarrow{a} t\}$. The set of successors of s (not relativized to any specific action) is defined as $R(s) := \bigcup_{a\in\Lambda} R_a(s)$.

If $\mathcal{M} = (D^{\mathcal{M}}, R_a^{\mathcal{M}}, \ldots)$ is a structure, we write $\mathcal{M}$ instead of $D^{\mathcal{M}}$ and wR_aw' instead of $wR_a^{\mathcal{M}}w'$. We also use this convention when two models $\mathcal{M}, \mathcal{N}$ are present and we avoid confusion on the binary relation R_a by stipulating that the domains of $\mathcal{M}, \mathcal{N}$ are disjoint: hence uR_au' means $uR_a^{\mathcal{M}}u'$ if $u \in \mathcal{M}$ and $uR_a^{\mathcal{N}}u'$ if $u \in \mathcal{N}$.

2.1.2 Syntax.

We can express various properties of an $\mathcal{L}$-structure by means of different logics on $\mathcal{L}$.

The most natural is *First Order Logic*, which we do not define, being very popular! The set of first-order sentences based on a language $\mathcal{L}$ is still denoted by $\mathcal{L}$.

If k is a cardinal, the *k-Infinitary Logic* based on $\mathcal{L}$ is denoted by $\mathcal{L}_k$ and consists of the extension of $\mathcal{L}$ which allows conjunctions on sets of cardinality strictly smaller than k. The *Infinitary Logic* of $\mathcal{L}$ is denoted by $\mathcal{L}_\infty$ and there we allow conjunctions on arbitrary sets (see [43], [23]).

We will also consider $\mathcal{L}_\infty$ expanded with all *cardinality quantifiers* Q_hx, whose intended meaning is: *there are at least h elements in the domain satisfying* ϕ. This logic is denoted by $\mathcal{L}_\infty(Q_\infty)$. If k is a limit cardinal, we will also use the logic $\mathcal{L}_\infty(Q_k)$ where we restrict the quantifiers Q_h to cardinals $h < k$.

The last classical logic we consider is *Monadic Second Order Logic* (MSO) (see [22], [23]). Here we have an infinite supply of unary predicate variables and we are allowed to quantify over them.

We now turn to Basic Modal Logic and its extensions.

The set of formulae of *Basic Modal Logic* based on $\mathcal{L} = (\Sigma, \Lambda)$ is the smallest set which contains any $p \in \Sigma$ and it is closed under the operators $\phi \wedge \psi, \neg\phi$, and $\langle a\rangle\phi$, for any $a \in \Lambda$. We denote the Basic Modal Logic of $\mathcal{L}$ by $\mathcal{L}^\diamond$. We will also use the defined operator $\phi \vee \psi, \phi \to \psi, [a]\phi := \neg\langle a\rangle\neg\phi$.

Infinitary Modal Logic based on a language $\mathcal{L}$ is the extension of Basic Modal Logic that allows conjunctions on arbitrary sets of formulae. We denote it by

$\mathcal{L}^{\diamond}_{\infty}$. If we restrict conjunctions to sets of cardinality smaller than k we have the logic $\mathcal{L}^{\diamond}_{k}$.

Propositional Dynamic Logic ([59, 30]), PDL for short, based on $\mathcal{L}$ is given by simultaneously defining *propositions* and *programs*: the set of PDL propositions is the least set containing the propositional constants $p \in \Sigma$ and closed under $\neg$, $\wedge$, and the unary operators $\langle \pi \rangle$, where π is a program. The set of PDL programs is the least set containing the atomic propositions $a \in \Lambda$ which is closed under the binary operations on programs ; and $\cup$, the unary Kleene star function on programs *, and the unary function ? from propositions to programs.

Given a formula ϕ in one of the above logics, the *language* $\mathcal{L}(\phi)$ of the formula is the set of propositional constants and atomic actions that appear in ϕ.

2.1.3 Semantics.

Given a structure $\mathcal{M}$ an $\mathcal{L}^{\diamond}$-formula ϕ is interpreted as a subset $[\![\phi]\!]$ of $\mathcal{M}$, defined as follows:

$$\begin{array}{lcl} [\![p]\!] & := & p^{\mathcal{M}} \\ [\![\neg\phi]\!] & := & \mathcal{M} \setminus [\![\phi]\!] \\ [\![\phi \wedge \psi]\!] & := & [\![\phi]\!] \cap [\![\psi]\!] \\ [\![\langle a \rangle \phi]\!] & := & \{s \in \mathcal{M} \mid [\![\phi]\!] \cap R_a(s) \neq \emptyset\} \end{array}$$

The semantics of the defined operators is:

$$\begin{array}{lcl} [\![\phi \vee \psi]\!] & := & [\![\phi]\!] \cup [\![\psi]\!] \\ [\![[a]\phi]\!]_V & := & \{s \in \mathcal{M} \mid R_a(s) \subseteq [\![\phi]\!]_V\} \end{array}$$

The extension of this semantics to formulae of (k-)Infinitary Modal Logic is given by considering

$$[\![\bigwedge \Phi]\!] := \bigcap_{\phi \in \Phi} [\![\phi]\!].$$

As for PDL, we simultaneously interpret programs π as binary relations $[\![\pi]\!]$ of the domain and propositions ϕ as subsets $[\![\phi]\!]$ of the domain. More precisely

$$\begin{array}{lcl} [\![\langle \pi \rangle \phi]\!] & := & \{s \in \mathcal{M} \mid [\![\phi]\!] \cap [\![\pi]\!](s) \neq \emptyset\} \\ [\![a]\!] & := & R_a^{\mathcal{M}} \\ [\![\phi?]\!] & := & \{(s,s) \in \mathcal{M} \mid s \in [\![\phi]\!]\} \\ [\![\pi_1 \cup \pi_2]\!] & := & [\![\pi_1]\!] \cup [\![\pi_2]\!] \\ [\![\pi_1 ; \pi_2]\!] & := & \{(s,t) \mid \exists z \ (s,z) \in [\![\pi_1]\!] \text{ and } \ (z,t) \in [\![\pi_2]\!]\} \\ [\![\pi^\star]\!] & := & \text{reflexive and transitive closure of } [\![\pi]\!]. \end{array}$$

If ϕ is a formula in $\mathcal{L}^{\diamond}$, $\mathcal{L}^{\diamond}_{k}$, $\mathcal{L}^{\diamond}_{\infty}$, PDL, or any of the extensions of $\mathcal{L}^{\diamond}$ defined later in this thesis, $\mathcal{M}, s \models \phi$ stands for $s \in [\![\phi]\!]$: s satisfies ϕ in $\mathcal{M}$. $\mathcal{M} \models \phi$ is used to denote $\mathcal{M}, r^{\mathcal{M}} \models \phi$. Logical consequence is then defined accordingly.

2.1.4 Translations.

Translations allow us to consider our extended modal logics as fragments of classical first-order, monadic second-order, or infinitary logics.

The *standard translation* ST of Basic (or Infinitary) Modal Logic into First Order Logic (Infinitary Logic, respectively) is defined as follows:

$$\begin{array}{lcl} ST(p) & := & p(x); \\ ST(\bigwedge \Phi) & := & \bigwedge_{\phi \in \Phi} ST(\phi); \\ ST(\neg\phi) & := & \neg ST(\phi); \\ ST(\langle a \rangle \phi) & := & \exists y(xR_a y \wedge ST(\phi)(y \mid x)), \text{ if } y \text{ is a variable new for } ST(\phi). \end{array}$$

The extension of ST to PDL into $\mathcal{L}_{\omega_1}$ is given by defining $ST(\langle\pi\rangle\phi)$ inductively on the complexity of π, by means of:

$$\begin{array}{lcl} ST(\langle \psi? \rangle \phi) & := & ST(\psi) \wedge ST(\phi) \\ ST(\langle \pi_1; \pi_2 \rangle \phi & := & ST(\langle \pi_1 \rangle \langle \pi_2 \rangle \phi) \\ ST(\langle \pi_1 \cup \pi_2 \rangle \phi) & := & ST(\langle \pi_1 \rangle \phi) \vee ST(\langle \pi_2 \rangle \phi) \\ ST(\langle \pi^\star \rangle \phi) & := & \bigvee_{n \in \omega} ST(\langle \pi^n \rangle \phi), \end{array}$$

where by $\langle\pi^n\rangle\phi$ we denote the formula $\langle\pi\rangle \ldots \langle\pi\rangle\phi$, in which the operator $\langle\pi\rangle$ is repeated n times.

All these extensions of ST are sound, in the sense that in any structure $\mathcal{M}$ the (modal, infinitary modal, PDL) formula ϕ is equivalent to the sentence $ST(\phi)(r \mid x)$ (which we denote simply by $ST(\phi)(r)$).

2.1.5 Derivability in Basic Modal Logic.

In this section we give a brief account of the axiomatic system K and of its extensions $K + \phi$, for a modal formula ϕ. This material is used in Chapter 6. We restrict here to mono-modal logic, that is, the language only contains an atomic action a; we denote the operator $\langle a \rangle$ by $\Diamond$.

The *Basic Modal Logic* K consists of a set of propositional axioms complete for classical logic, the modal axiom $\Box(P \to Q) \to (\Box(P) \to \Box(Q))$, and the rules of substitution, modus ponens and necessitation (infer $\Box(\phi)$ from ϕ). Deducibility of ψ from ϕ in K (notation: $\vdash_{K+\phi} \psi$) is defined as follows. We add ϕ to the set of axioms K and we close under the rules of K, obtaining the set of formulae $Th(K + \phi)$. Then $\vdash_{K+\phi} \psi$ iff $\psi \in Th(K + \phi)$.

Frame completeness.

A *frame* consists of a domain and of a binary relation on it, that is, it is a structure for a language without propositional constants. A formula $\phi \in \mathcal{L}^\Diamond$ is *valid* in a frame $\mathcal{F}$ if and only if $\mathcal{F}$ satisfies the monadic second-order formula $\forall p_1 \ldots \forall p_n \forall x ST(\phi)$, where $p_1, \ldots p_n$ are all propositional constants in ϕ. We denote this sentence by $\overline{ST}(\phi)$. A formula ψ is a *frame logical consequence* of a formula ϕ ($\phi \models_f \psi$) if and only if, for all frames F, if ϕ is valid in $\mathcal{F}$, then ψ is valid in $\mathcal{F}$, that is, iff $\models \overline{ST}(\phi) \rightarrow \overline{ST}(\psi)$. A formula ϕ is said to be *complete* if and only if, for all ψ, $\vdash_{K+\phi} \psi \Leftrightarrow \phi \models_f \psi$. Examples of incomplete formulae can be found in the literature.

General frame semantics.

A semantic characterization of K is obtained via general frames. A *general frame* is a pair $(\mathcal{F}, \mathcal{W})$, where $\mathcal{F}$ is a frame and $\mathcal{W}$ is a set of subsets of W, which is closed under the boolean operations of union, complementation w.r.t. W, and $\Box(X) = \{w \in W : \forall v\ (wRv \rightarrow v \in X)\}$. Validity in general frames is defined as for frames except for the fact that we only consider second-order quantifiers to range on $\mathcal{W}$. Logical consequence in general frames ($\phi \models_{gf} \psi$) is defined accordingly.

Fact: for all modal formulae ϕ, ψ we have

$$\vdash_{K+\phi} \psi \Leftrightarrow \phi \models_{gf} \psi.$$

An axiomatic system describing general frame semantics is obtained by considering a two-sorted first-order language, with *worlds* and *sets*, with binary predicates R (on worlds) and $\in$ (between worlds and sets), as well as operations $-, \cup$, and $\Box$ on the set sort. The minimal logic describing logical consequence in general frame semantics is obtained by considering the following theory μ:

- first-order principles for both sorts;
- $\forall p \forall w (w \in -p \leftrightarrow \neg w \in p)$;
- $\forall p \forall q \forall w (w \in p \cup q \leftrightarrow w \in p \vee w \in q)$;
- $\forall p \forall w (w \in \Box p \leftrightarrow \forall v (wRv \rightarrow v \in p))$
- $\forall p \forall q (\forall w (w \in p \leftrightarrow w \in q) \rightarrow p = q)$ (extensionality),

where letters w, v denote words and letters p, q denote sets.

It is then clear that the theory μ is sound and complete w.r.t. general frames and by comparing this with the previous result we obtain:

Fact: for all pairs of modal formulae ϕ, ψ, we have:

$$\vdash_{K+\phi} \psi \Leftrightarrow \mu \vdash \overline{ST(\phi)} \rightarrow \overline{ST(\psi)}.$$

2.2 Bisimulation and the like.

The concept of bisimulation is central in this dissertation, being the main link between the different topics of extended modal logics and non-well-founded sets.

We first consider the basic definition and some classical operations on structures preserving bisimulation. We use these operations to characterize bisimilar structures, a task that can also be achieved by using the modal language. We then consider the bounded version of bisimulation and, finally, we state an amalgamation lemma, which is often used in interpolation proofs.

Basic bisimulation.

2.2.1. DEFINITION. Let $\mathcal{M}$ and $\mathcal{N}$ be two $\mathcal{L}$-structures and let $Z \subseteq \mathcal{M} \times \mathcal{N}$. Z is a *bisimulation* between $\mathcal{M}$ and $\mathcal{N}$ (notation $Z : \mathcal{M} \sim \mathcal{N}$) if:

1. $r^{\mathcal{M}} Z r^{\mathcal{N}}$;

2. Z-connected points agree on the proposition constants: if sZt then $\mathcal{M}, s \models p$ iff $\mathcal{N}, t \models p$ for every unary p in $\mathcal{L}$.

3. $s \overset{a}{\to} s'$ and sZt implies that there is a t' such that $s'Zt'$ and $t \overset{a}{\to} t'$;

4. Vice versa: if sZt and $t \overset{a}{\to} t'$ then there is an s' with $s \overset{a}{\to} s'$ and $s'Zt'$.

Two structures $\mathcal{M}$ and $\mathcal{N}$ are *bisimilar* (notation $\mathcal{M} \sim \mathcal{N}$) if there exists a bisimulation between them. If $\mathcal{M}$ and $\mathcal{N}$ are structures of respectively $\mathcal{L}_1$ and $\mathcal{L}_2$, both of which contain $\mathcal{L}$, then an $\mathcal{L}$-*bisimulation* between $\mathcal{M}$ and $\mathcal{N}$ is a relation Z satisfying the above clauses just for the symbols in $\mathcal{L}$. The notion of $\mathcal{L}$-bisimilar structures is defined accordingly (notation $Z : \mathcal{M} \sim_{\mathcal{L}} \mathcal{N}$, $\mathcal{M} \sim_{\mathcal{L}} \mathcal{N}$). □

A sentence ϕ in a language that can speak about structures is *invariant for bisimulation* if whenever $Z : \mathcal{M} \sim \mathcal{N}$ then ϕ holds in $\mathcal{M}$ iff it holds in $\mathcal{N}$.

Basic Modal Logic and any extension of it introduced so far consist of formulae which are invariant for bisimulation.

Characterizations of bisimilar structures.

If two structures are isomorphic, they are obviously bisimilar. The vice versa is easily seen to be false, but it becomes true *modulo* certain operations on structures. Here we consider two different kinds of operations: the collapse (that makes structures *smaller*) and the various forms of expansion (that make structures *larger*).

The first one we consider is the collapse.

2.2.2. DEFINITION. Given a structure $\mathcal{G}$, its *collapse* $Coll(\mathcal{G})$ is defined as follows.

- Its domain consists of the equivalence classes of $\mathcal{G}$ with respect to the equivalence relation $\sim$ defined as follows:

$$w \sim v \text{ iff there exists a bisimulation between } (\mathcal{G}, w) \text{ and } (\mathcal{G}, v).$$

- The relation $R^{Coll(\mathcal{G})}$ is given by

$$[w]R^{Coll(\mathcal{G})}[v] \Leftrightarrow \exists v' \in [v] \; wR^{\mathcal{G}}v'.$$

- The root $r^{Coll(\mathcal{G})}$ is equal to $[r^{\mathcal{G}}]$.
- The unary predicates are defined by

$$p^{Coll(\mathcal{G})}([w]) \Leftrightarrow p^{\mathcal{G}}(w).$$

□

One can easily prove that $Coll(\mathcal{G})$ is bisimilar to $\mathcal{G}$ via the relation $wZ[w]$. Moreover, any bisimulation between collapsed structures must be a bijective function, as it is easy to prove. It follows that if $Coll(\mathcal{G})$ is bisimilar to $Coll(\mathcal{H})$ then they are isomorphic. This gives:

2.2.3. LEMMA. *$\mathcal{G}$ is bisimilar to $\mathcal{H}$ iff $Coll(\mathcal{G})$ is isomorphic to $Coll(\mathcal{H})$.*

The second construction we consider transforms any structure into a bisimilar *tree*.

2.2.4. DEFINITION. For any structure $\mathcal{M}$ in a language (Σ, Λ), the *unraveling* $\mathcal{M}^1$ of $\mathcal{M}$ is defined as follows.

- Its domain $\mathcal{M}^1$ consists of all finite sequences $v_0, a_1, v_1, \ldots, a_k, v_k$ where:

  $a_i \in \Lambda$, for all $i \leq k$;

  $v_0 = r^{\mathcal{M}}$;

  for all $i < k$ we have $v_i R^{\mathcal{M}}_{a_{i+1}} v_{i+1}$.

  We define the function $end : \mathcal{M}^1 \to \mathcal{M}$ as follows: $end(r^{\mathcal{M}}) = r^{\mathcal{M}}$, $end(\sigma a w) = w$.

- The accessibility relations are defined by

$$\sigma R^{\mathcal{M}^1}_a \tau \text{ iff } \tau = \sigma a w.$$

- The root is the sequence $r^{\mathcal{M}}$ having only one element.
- The unary predicates are defined by

$$p^{\mathcal{M}^1}(\sigma) \Leftrightarrow p^{\mathcal{M}}(end(\sigma)).$$

□

Notice that the unraveling $\mathcal{M}^1$ of a structure $\mathcal{M}$ is a tree, which is bisimilar to $\mathcal{M}$ via the relation $Z = \{(\sigma, end(\sigma)) : \sigma \in \mathcal{M}^1\}$. Hence, if $\mathcal{M}^1$ is isomorphic to $\mathcal{N}^1$, then the two structures $\mathcal{M}$ and $\mathcal{N}$ are bisimilar. The converse is not true, however (just consider the full binary tree and a reflexive node), but it becomes true if in the unraveling we are careful enough to *copy* each node a sufficient number of times ($\aleph_0$, in the previous example):

2.2.5. Definition. Let $k > 0$ be a cardinal number. For any structure $\mathcal{M}$ in a language (Σ, Λ), the *k-expansion* $\mathcal{M}^k$ of $\mathcal{M}$ is defined as follows.

- Its domain $\mathcal{M}^k$ consists of all finite sequences $v_0, a_1, \alpha_1, v_1, \ldots, a_h, \alpha_h, v_h$ where:

  $a_i \in \Lambda$, for all i;

  $\alpha_1, \ldots, \alpha_h$ are ordinals strictly less than k;

  $v_0 = r^{\mathcal{M}}$;

  for all $i < h$ we have $v_i R^{\mathcal{M}}_{a_{i+1}} v_{i+1}$.

  We define the function $end : \mathcal{M}^k \to \mathcal{M}$ as follows: $end(r^{\mathcal{M}}) = r^{\mathcal{M}}$, $end(\sigma a \alpha w) = w$.

- The accessibility relations are defined by

$$\sigma R^{\mathcal{M}^n}_a \tau \text{ iff } \tau = \sigma a \alpha w.$$

- The root is the sequence $r^{\mathcal{M}}$ having only one element.
- The unary predicates are defined by

$$p^{\mathcal{M}^k}(\sigma) \Leftrightarrow p^{\mathcal{M}}(end(\sigma)).$$

□

Notice that the k-expansion $\mathcal{M}^k$ of a structure $\mathcal{M}$ is a tree, which is bisimilar to $\mathcal{M}$ via the relation $Z = \{(\sigma, end(\sigma)) : \sigma \in \mathcal{M}^k\}$. Moreover, any a-successor v of a node $w = end(\sigma)$ in $\mathcal{M}$ is bisimilar to any successor of σ of the form $\sigma a \alpha v$, with $\alpha < k$. We say that the node v has been *copied* k times in $\mathcal{M}^k$. Notice that the 1-expansion of a structure is isomorphic to the usual unraveling.

Using expansions we can characterize bisimilar structures.

2.2.6. THEOREM. *Two structure $\mathcal{G}$ and $\mathcal{H}$ are bisimilar iff there exists a cardinal k such that $\mathcal{G}^k$ and $\mathcal{H}^k$ are isomorphic.*

The cardinal k of the above theorem can be chosen to be the maximum between the cardinality of $\mathcal{G}$ and $\mathcal{H}$.

Bisimulation and language invariance.

Another way to characterize bisimilar structures is via *logics*. If two structures are bisimilar then they satisfy the same formulae of all extended modal logics considered above. The converse is false for Basic Modal Logic and PDL if we do not restrict the class of structures, but it is true if we consider Infinitary Modal Logic.

2.2.7. THEOREM. *Two structure $\mathcal{G}$ and $\mathcal{H}$ are bisimilar iff they satisfy the same infinitary modal formulae.*

Moreover, it is possible to prove that any structure $\mathcal{G}$ is characterized modulo bisimulation by a single infinitary formula ([7]):

2.2.8. THEOREM. *For any structure $\mathcal{G}$ there exists an infinitary modal formula $\phi_{\mathcal{G}}$ such that, for any structure $\mathcal{H}$, it holds:*

$$\mathcal{H} \models \phi_{\mathcal{G}} \Leftrightarrow \mathcal{H} \text{ is bisimilar to } \mathcal{G}.$$

Bounded bisimulation.

One disadvantage of the logic $\mathcal{L}_\infty^\diamond$ is that the formulae in this logic do not form a set. It is then sometimes useful to consider fragments of $\mathcal{L}_\infty^\diamond$ that *do* form a set. One way to do this is to bound the modal depth of formulae to a fixed ordinal: the *degree* or *modal depth* of formula is defined as follows:

$$\begin{array}{lcl} d(p) & := & 0 \\ d(\neg\phi) & := & d(\phi) \\ d(\bigvee\Phi) & := & \sup\{d(\phi) \mid \phi \in \Phi\} \\ d(\langle a\rangle\phi) & := & d(\phi)+1 \end{array}$$

Consider the fragment of $\mathcal{L}_\infty^\diamond$ which consists of all formulae whose depth is bounded by a fixed ordinal α. For this logic we have a natural notion of bisimulation:

2.2.9. DEFINITION. Let α be an ordinal. An *α-bisimulation* between two models $\mathcal{M}$ and $\mathcal{N}$ is a set of relations Z_β with $\beta \leq \alpha$ such that:

1. The roots are connected via Z_α;

2. Z_β-connected points agree on the proposition constants;

3. If $s \stackrel{a}{\to} s'$ in $\mathcal{M}$ and $sZ_{\beta+1}t$ then there is a t' in $\mathcal{N}$ with $t \stackrel{a}{\to} t'$ and $s'Z_\beta t'$;

4. Vice versa.

5. if α is a limit then $Z_\alpha = \bigcap_{\beta<\alpha} Z_\beta$. □

All Infinitary Modal Logic formulae ϕ of depth bounded by α are invariant for α-bisimulation. That is, if $(Z_\beta)_{\beta\leq\alpha}$ is an α-bisimulation and $sZ_\alpha t$ then $s \models \phi$ iff $t \models \phi$.

Notice that we have a real bisimulation between $\mathcal{M}$ and $\mathcal{N}$ if and only if we can find a (non-well-founded!) sequence of relations $(Z_i)_{i\in\omega}$ such that

1. The roots are connected via Z_0;

2. If $s \stackrel{a}{\to} s'$ in $\mathcal{M}$ and $sZ_i t$ then there is a t' in $\mathcal{N}$ with $t \stackrel{a}{\to} t'$ and $s'Z_{i+1}t'$;

3. Vice versa.

It is indeed clear that the relation $Z = \bigcup_{i\in\omega} Z_i$ is a bisimulation between $\mathcal{M}$ and $\mathcal{N}$.

Amalgamation.

Amalgamation is a classical tool which is often used to prove interpolation. There is a long tradition in the area of algebraic logic connecting (algebraic) amalgamation properties with (logical) interpolation properties (e.g. [1], [55], [49], and [50]). In this dissertation we shall use a construction (related to the zig-zag products of [50]) that *amalgamates* two structures, which are bisimilar with respect to the common language, in a third one.

2.2.10. LEMMA. *Let $\mathcal{M}, \mathcal{N}$ be structures for the languages $\mathcal{L}$ and $\mathcal{L}'$ respectively. If $\mathcal{M} \sim_{\mathcal{L}\cap\mathcal{L}'} \mathcal{N}$, then there is an $(\mathcal{L}\cup\mathcal{L}')$-structure $\mathcal{K}$ which is $\mathcal{L}$-bisimilar to $\mathcal{M}$ and $\mathcal{L}'$- bisimilar to $\mathcal{N}$.*

Proof.
Fix an $\mathcal{L}\cap\mathcal{L}'$-bisimulation Z between $\mathcal{M}$ and $\mathcal{N}$. $\mathcal{K}$ is defined as follows.

1. The domain is the disjoint union of $\{(m,n) \in \mathcal{M}\times\mathcal{N} : mZn\}$, $\mathcal{M}$, and $\mathcal{N}$; the root is $r^{\mathcal{K}} := (r^{\mathcal{M}}, r^{\mathcal{N}})$.

2. If $k, k' \in \mathcal{M}$ and a is an action in $\mathcal{L}$, then $kR_a^{\mathcal{K}}k'$ iff $kR_a^{\mathcal{M}}k'$.

   If $k, k' \in \mathcal{N}$ and a is an action in $\mathcal{L}'$, then $kR_a^{\mathcal{K}}k'$ iff $kR_a^{\mathcal{N}}k'$.

If a is an action in $\mathcal{L} \cap \mathcal{L}'$, then $(m,n)R_a^{\mathcal{K}}k$ iff $k = (m',n')$, $mR_a^{\mathcal{M}}m'$, and $nR_a^{\mathcal{N}}n'$.

If a is an action in $(\mathcal{L} \setminus \mathcal{L}')$ then $(m,n)R_a^{\mathcal{K}}k$ iff $k \in \mathcal{M}$ and $mR_a^{\mathcal{M}}k$.

If a is an action in $(\mathcal{L}' \setminus \mathcal{L})$ then $(m,n)R_a^{\mathcal{K}}k$ iff $k \in \mathcal{N}$ and $nR_a^{\mathcal{N}}k$.

3. For any $p \in \mathcal{L} \cup \mathcal{L}'$ define: $\mathcal{K},(m,n) \models p$ iff $p \in \mathcal{L}$ and $\mathcal{M},m \models p$ or $p \in \mathcal{L}'$ and $\mathcal{N},n \models p$. Note that this is well-defined, because if $p \in \mathcal{L} \cap \mathcal{L}'$ then m and n must agree on p, since they are $\mathcal{L} \cap \mathcal{L}'$-bisimilar. If $k \in \mathcal{M}$ ($k \in \mathcal{N}$) then $\mathcal{K},k \models p$ iff $\mathcal{M},k \models p$ ($\mathcal{N},k \models p$, respectively).

We claim that $\mathcal{K}$ is $\mathcal{L}$-bisimilar to $\mathcal{M}$ via the relation $S = \{(m,m) : m \in \mathcal{M}\} \cup \{((m,n),m) : (m,n) \in \mathcal{K}\}$ and $\mathcal{L}'$-bisimilar to $\mathcal{N}$ via the relation $T = \{(n,n) : n \in \mathcal{N}\} \cup \{((m,n),n) : (m,n) \in \mathcal{K}\}$. By symmetry we only need to prove that the relation S is an $\mathcal{L}$-bisimulation between $\mathcal{K}$ and $\mathcal{M}$.

1. The roots are connected by definition. If kSm and $k = m \in \mathcal{M}$ then k and m satisfy the same $\mathcal{L}$-propositions by definition of the valuation in $\mathcal{K}$. The same holds if $k = (m,n)$, because then m and n are $\mathcal{L} \cap \mathcal{L}'$-bisimilar and thus satisfy the same propositional constants in $\mathcal{L} \cap \mathcal{L}'$.

2. Suppose that kSm. If $k \in \mathcal{M}$ then $k = m$ and the back and forth conditions for an $\mathcal{L}$-bisimulation are trivially satisfied. If $k = (m,n)$ with mZn, the back and forth conditions for an action a in $\mathcal{L} \setminus \mathcal{L}'$ are satisfied because $(m,n)R_a^{\mathcal{K}}k'$ iff $k' \in \mathcal{M}$ and $mR_a^{\mathcal{M}}k'$. Next, consider $a \in \mathcal{L} \cap \mathcal{L}'$. If $(m,n)R_a^{\mathcal{K}}k'$, then $k' = (m',n')$ with $mR_a^{\mathcal{M}}m'$ and $k'Sm'$. Vice versa, if $mR_a^{\mathcal{M}}m'$ then since Z is a $\mathcal{L} \cap \mathcal{L}'$-bisimulation and mZn, there exists $n' \in \mathcal{N}$ such that $nR_a^{\mathcal{N}}n'$ and $m'Zn'$. Then $(m',n') \in \mathcal{K}$, $(m,n)R_a^{\mathcal{K}}(m',n')$, and $(m',n')Sm'$. □

2.3 A brief introduction to the μ-Calculus.

Basic Modal Logic is computationally attractive but not very expressive; in the previous section to get more power we extended Basic Modal Logic until we reached Infinitary Modal Logic, going through PDL. The latter is still computationally attractive, while Infinitary Modal Logic is highly non-constructive. In this section we introduce another extension of Basic Modal Logic which is stronger than PDL, but it still behaves well from a computational point of view.

The Modal μ-Calculus ([44]) is an extension of Basic Modal Logic with least and greatest fixed point operators. It is a very powerful formalism into which many other well-studied modal logics can be embedded. Examples are Linear Temporal Logic (LTL, [57]), Computation Tree Logic (CTL, [17]), CTL* ([25, 26]), PDL, and of course Basic Modal Logic. Moreover, the Modal μ-Calculus is of great interest to Computer Science since it expresses important properties of processes, such as termination (every run is finite) and fairness (on every infinite run, no action is repeated infinitely often to the exclusion of all others).

The power of the Modal μ-Calculus is also evident from a more theoretical perspective: as we shall see, the Modal μ-Calculus is a fragment of Monadic Second Order Logic containing only formulae that are *invariant for bisimulation*. Janin and Walukiewicz prove the converse: any property which is invariant for bisimulation and expressible in Monadic Second Order Logic is already expressible in the Modal μ-Calculus ([41]). Yet the Modal μ-Calculus enjoys many desirable properties which Monadic Second Order Logic lacks: a complete calculus ([67]), an exponential-time decision procedure, and the finite model property ([63]). Switching from MSO to its bisimulation-invariant fragment gives us these desirable properties.

Given $\mathcal{L} = (\Sigma, \Lambda)$ and an infinite set of variables VAR, the set of formulae of the *Modal μ-Calculus* is defined as the smallest set containing any $p \in \Sigma$, the propositional variables $X \in$ VAR and which is closed under boolean operators, modal operators $\langle a \rangle \phi$, for $a \in \Lambda$, and under the *least fixed point operator* $\mu X.\phi$, if ϕ is positive in X (i.e. all occurrences of X in ϕ fall under an even number of negations).

The dual operator $\nu X.\phi$ is defined as $\nu X.\phi := \neg\mu\neg\phi(\neg X|X)$. A μX- and νX-operator binds occurrences of X in the usual way: the notion of *closed* μ-formula, or *μ-sentence*, is then given as usual.

Given a structure $\mathcal{M}$ and a valuation $V : \text{VAR} \to Pow(\mathcal{M})$, a μ-formula is interpreted as a subset $[\![\phi]\!]_V$ of $\mathcal{M}$, defined as follows:

$$
\begin{array}{lcl}
[\![p]\!]_V & := & p^{\mathcal{M}} \\
[\![\neg p]\!]_V & := & \mathcal{M} \setminus p^{\mathcal{M}} \\
[\![X]\!]_V & := & V(X) \\
[\![\phi \wedge \psi]\!]_V & := & [\![\phi]\!]_V \cap [\![\psi]\!]_V \\
[\![\langle a \rangle \phi]\!]_V & := & \{s \in \mathcal{M} \mid [\![\phi]\!]_V \cap R_a(s) \neq \emptyset\} \\
[\![\mu X.\phi]\!]_V & := & \bigcap\{S \subseteq \mathcal{M} \mid [\![\phi]\!]_{V[X:=S]} \subseteq S\}
\end{array}
$$

where $V[X := S]$ is equal to the valuation function V except that S is assigned to X. Note that $[\![\mu X.\phi]\!]_V$ is the least fixed point of the monotone operator that assigns to S the set $[\![\phi]\!]_{V[X:=S]}$. The semantics of the dual operators is:

$$
\begin{array}{lcl}
[\![\phi \vee \psi]\!]_V & := & [\![\phi]\!]_V \cup [\![\psi]\!]_V \\
[\![[a]\phi]\!]_V & := & \{s \in \mathcal{M} \mid R_a(s) \subseteq [\![\phi]\!]_V\} \\
[\![\nu X.\phi]\!]_V & := & \bigcup\{S \subseteq \mathcal{M} \mid S \subseteq [\![\phi]\!]_{V[X:=S]}\}
\end{array}
$$

Note that if ϕ is a formula containing as free variables only $X_1, \ldots, X_n$ the valuation need only to assign a value to these variables. In such a case we may write $[\![\phi]\!]_{[X_1:=A_1,\ldots,X_n:=A_n]}$ instead of $[\![\phi]\!]_V$. We may leave out the valuation altogether if ϕ

is a sentence. If ϕ is a formula, $X_1, \ldots, X_n$ are all its free variables, and $A_1, \ldots, A_n$ are subsets of the domain of a structure $\mathcal{M}$, we denote by $\phi(A_1, \ldots, A_n)$ the set $[\![\phi]\!]_{[X_1:=A_1,\ldots,X_n:=A_n]}$.

$\mathcal{M}, V, s \models \phi$ stands for $s \in [\![\phi]\!]_V$: s satisfies ϕ under V in $\mathcal{M}$. If ϕ is a sentence we may leave out the valuation. $\mathcal{M} \models \phi$ is used to denote $\mathcal{M}, r^{\mathcal{M}} \models \phi$. Logical consequence is then defined accordingly.

By looking at the semantics of the μ-operator we can easily translate the μ-sentences of a language $\mathcal{L}$ into sentences of the Monadic Second Order Logic of $\mathcal{L}$, by means of an extension of the standard translation ST, defined in Section 2.1.4. We let:

$$ST(\mu X.\phi) = \forall X(\forall x(ST(\phi) \to X(x)) \to X(x)).$$

Then a μ-sentence ϕ is equivalent to the MSO-sentence $ST(\phi)(r)$.

Examples.

- Consider the sentence $\phi = \mu X.(p \vee \langle a \rangle X)$. We claim that the valuation $[\![\phi]\!]$ of a sentence ϕ in a structure $\mathcal{M}$ is equal to the set

  $$A = \{w : \exists w_0 = w, \ldots, w_n(\bigwedge_{i<n} w_i R_a^{\mathcal{M}} w_{i+1} \wedge w_n \models p)\},$$

  that is, ϕ is equivalent to the PDL sentence $\langle a^* \rangle p$.

  Since

  $$[\![\mu X.(p \vee \langle a \rangle X)]\!] = \bigcap \{S \subseteq \mathcal{M} \mid [\![p \vee \langle a \rangle X]\!]_{V[X:=S]} \subseteq S\},$$

  to prove that $A = [\![\mu X.(p \vee \langle a \rangle X)]\!]$ we have to show:

  a) $A \supseteq [\![\mu X.(p \vee \langle a \rangle X)]\!]$, which can be proved by showing that $p^{\mathcal{M}} \cup \{w : \exists v(w R_a^{\mathcal{M}} v \wedge v \in A)\} \subseteq A$;

  b) $[\![\mu X.(p \vee \langle a \rangle X)]\!] \supseteq A$, which can be proved by showing that for any $S \subseteq \mathcal{M}$ with $p^{\mathcal{M}} \cup \{w : \exists v(w R_a^{\mathcal{M}} v \wedge v \in S)\} \subseteq S$ it is true that $A \subseteq S$.

  The first condition is easily verified and we leave it to the reader.

  As for the second one, suppose that $S \subseteq \mathcal{M}$ is such that $p^{\mathcal{M}} \cup \{w : \exists v(w R_a^{\mathcal{M}} v \wedge v \in S) \subseteq S$. If $u \in A$, then there are $w_0 = u, \ldots, w_n$ such that $w_i R_a^{\mathcal{M}} w_{i+1}$ for all $i < n$ and $w_n \models p$. We prove by induction on n that $u \in S$. If $n = 0$, then $u \models p$ and $u \in S$. If $n > 0$, consider the node w_1. By induction, $w_1 \in S$ and from $\{w : \exists v(w R_a^{\mathcal{M}} v \wedge v \in S) \subseteq S$, we obtain $u \in S$.

- Let $\psi(X, Y)$ be the formula $(p \wedge \langle a \rangle X) \vee \langle a \rangle Y$ and consider the sentence $\nu X.\mu Y.\psi(X, Y)$. We prove that $[\![\nu X.\mu Y.\psi(X, Y)]\!]$ is equal to the set A of points w for which there is an infinite chain of a-transitions starting from w in which p holds infinitely often. It is known that expressing this property in the Modal μ-Calculus requires nesting of ν and μ operators (This is proved in [24]. For results related to the nesting of fixed points operators, see [47]).

First, using the preceding example we write the formula $\mu Y.\psi(X, Y)$ as $\langle a^\star\rangle(p \wedge \langle a\rangle X)$. By looking at the semantics of the greatest fixed point operator ν we see that in order to prove that $A = [\![\nu X.\langle a^\star\rangle(p \wedge \langle a\rangle X)]\!]$ we need:

a) $A \subseteq \langle a^\star\rangle(p \wedge \langle a\rangle A)$;

b) if S is a set with $S \subseteq \langle a^\star\rangle(p \wedge \langle a\rangle S)$, then $S \subseteq A$.

In order to prove *a*), suppose that $w \in A$. Then there is an infinite a-sequence $\sigma = (w_i)_{i\in\omega}$ with $w_0 = w$, where p appears infinitely often. Any node in this sequence belongs to A and after a finite number of a-steps from w we reach a node where p holds, hence $w \in \langle a^\star\rangle(p \wedge \langle a\rangle A)$.

We now prove *b*). Suppose $S \subseteq \langle a^\star\rangle(p \wedge \langle a\rangle S)$ and take $s \in S$. Then $s \in \langle a^\star\rangle(p \wedge \langle a\rangle S)$ and there is a finite chain of a-transition starting from s with an end node s_1 which satisfies $p \wedge \langle a\rangle S$, that is, $s_1 \models p$ and have an a-successor $s_2 \in S$. By applying the same argument to s_2 we obtain a finite chain of a-transition starting from s_2 with an end node s_3 which satisfies p and have an a-successor $s_4 \in S$, and so on. This proves that that $s \in A$. □

A basic result applicable to the Modal μ-Calculus is *the Tarski Theorem*, which says that the least fixed point $\mu(F)$ of a monotone operator $F : Pow(A) \to Pow(A)$ is equal to $F_\star = \bigcup_\alpha F_\alpha$, where

$$F_0 = \emptyset,$$

$$F_{\alpha+1} = F(F_\alpha),$$

$$F_\lambda = \bigcup_{\beta<\alpha} F_\alpha, \text{ for a limit ordinal } \lambda.$$

Dually, the greatest fixed point operator of a monotone operator F on a set X is equal to $F^\star = \bigcap_\alpha F^\alpha$, where

$$F^0 = A,$$

$$F^{\alpha+1} = F(F^\alpha),$$

$$F^\lambda = \bigcap_{\beta<\alpha} F^\alpha, \text{ for a limit ordinal } \lambda.$$

The Tarski Theorem can be used as a tool to get a better understanding of the evaluation of a μ-sentence ϕ on a point w of a structure $\mathcal{M}$.

Consider, for example, the formula $\phi = \mu X.[a]X$. Using the Tarski Theorem we can show that the formula $\mu X.[a]X$ is true in w iff there is no infinite chain of a-transitions starting from w. If $w \models \phi$, then w belongs to the fixed point $F_\star$ of the operator $[\![[a]X]\!]$. Since $F_\star$ is a fixed point, we have $F_\star = [a]F_\star$, and $w \in [a]F_\star$. Hence, if w_1 is an a-successor of w we have $w_1 \in F_\star$: we say that the formula ϕ has been *regenerated* in w_1, since in w_1 we have to restart the evaluation of ϕ. However, this regeneration cannot occur infinitely many times in an a-path starting from w. This is true because we can consider the transfinite approximations F_α of $F_\star$, and, by the Tarski Theorem, we know that if $w \in F_\star$ then there exists an ordinal β such that $w \in F_{\beta+1}$. Hence if $w \models \phi$, then $w \in [a]F_\beta$ and all a-successors w_1 of w belongs to $F_\beta \subseteq F_\star$. Hence, the sentence ϕ has been regenerated, but at *an inferior level.* If we perform the same evaluation on an a-successor w_1 of w, we see that all a-successors of w_1 are in F_γ, for $\gamma < \beta$ and this process of regeneration can continue only a finite number of times along a path, because after a finite number of regenerations we must end with $F_0 = \emptyset$. Hence, no infinite chain of a-transitions starts from w.

Notice that things are different if we consider greatest fixed point operators: here the Tarski Theorem implies that $w \in F^\star$ iff $w \in \bigcap_\alpha F^\alpha$. Hence, on evaluating a greatest fixed point sentence $\phi = \nu X.\psi$, it is possible that we are obliged to regenerate the sentence ϕ infinitely many times along a path.

More complex situations can arise if we have *nested* fixed points operators in a sentence ϕ, as for example in $\phi = \nu X.\mu Y.(p \vee \langle a\rangle X) \vee \langle a\rangle Y$. On evaluating ϕ, the formula $\psi[X := \phi] = \mu Y.(p \vee \langle a\rangle\phi) \vee \langle a\rangle Y$ can be regenerated infinitely many times although it starts with a μ-operator. The reason for this is that the subformula $\psi(X, Y)$ of ϕ contains the variable Y which is in the scope of a greatest fixed point operator νY: if $w \models \phi$ then $w \in [\![\mu Y.(p \vee \langle a\rangle X) \vee \langle a\rangle Y]\!]_{[X:=[\![\phi]\!]]}$ and this implies that after a finite number of a-steps we reach a node w_1 with $w_1 \in [\![p \vee \langle a\rangle X]\!]_{[X:=[\![\phi]\!]]}$ (see the first of the examples given above). Hence, either $w_1 \models p$, or w_1 has an a-successor w_2 such that $w_2 \models \phi$. Then $w_2 \in [\![\mu Y.(p \vee \langle a\rangle X) \vee \langle a\rangle Y]\!]_{[X:=[\![\phi]\!]]}$ and the formula $\psi[X := \phi]$ has been regenerated in w_2: we can have a regeneration of $\psi[X := \phi]$ any time we regenerate ϕ, possibly infinitely many times. □

The range of the ordinals in the definition of $F_\star$ and $F^\star$, for $F : Pow(A) \to Pow(A)$, can be taken to be $\{\alpha : \alpha \leq \gamma\}$, where γ is the least ordinal with $\gamma > |A|$. This allows us to show that on structures of cardinality smaller than a fixed k, any μ-sentence is equivalent to an infinitary modal formula:

2.3.1. Proposition. *Let k be a cardinal. Given a μ-sentence ϕ, there exists an infinitary modal formula ϕ^* in the same language, such that on any structure $\mathcal{M}$ with $|\mathcal{M}| < k$ ϕ^* is equivalent to ϕ.*

Proof.
Let ϕ be of the form $\mu X.\psi(X)$. Consider the sequence defined by $A_0 = \emptyset$,

$A_{\alpha+1} = [\![\psi]\!]_{[X:=A_\alpha]}$, and $A_\alpha = \bigcup_{\beta<\alpha} A_\beta$. By the Tarski Theorem we have that $w \in [\![\phi]\!]$ iff $w \in A_\gamma$, where $\gamma = min\{\alpha \mid |\alpha| > |X|\}$. Hence, if we can use disjunctions on infinite sets, we can inductively define the formula ψ_α by

$$\psi_0 = \bot, \psi_{\alpha+1} = \psi[X := \psi_\alpha], \psi_\lambda = \bigvee_{\alpha<\lambda} \psi_\alpha,$$

and

$$w \in [\![\mu X.\psi]\!] \text{ iff } w \in [\![\psi_\gamma]\!].$$

The thesis is obtained by applying the above reduction to all occurrences of least fixed points operators in ϕ. □

From Proposition 2.3.1 it follows that any sentence of the Modal μ-Calculus is invariant for bisimulation:

2.3.2. Proposition. *If ϕ is a sentence of the Modal μ-Calculus, $\mathcal{M}$ and $\mathcal{N}$ are bisimilar structures, and $\mathcal{M} \models \phi$, then $\mathcal{N} \models \phi$.*

Proof.
Consider a μ-sentence in the language $\mathcal{L}$ and let $\mathcal{M}$, $\mathcal{N}$ be $\mathcal{L}$-structures such that $\mathcal{M} \models \phi$ and $\mathcal{M} \sim \mathcal{N}$. If k is a cardinal such that $k > |\mathcal{M}|, |\mathcal{N}|$, then ϕ is equivalent on $\mathcal{M}$ and $\mathcal{N}$ to an infinitary modal sentence by the preceding proposition. But infinitary modal sentences are invariant for bisimulation. □

Notice however that the Modal μ-Calculus is *not* equivalent to Infinitary Modal Logic on arbitrary structures: non-well-foundedness of a relation R_a is an example of a property which is expressible in the Modal μ-Calculus by $\nu X.\langle a \rangle X$, but not in Infinitary Modal Logic.

2.4 The family of graded modal logics and their semantics.

All extensions of Basic Modal Logic we encountered so far are invariant under the basic notion of bisimulation. We went through PDL and the Modal μ-Calculus until we reached Infinitary Modal Logic, but we did not have to change the notion of bisimulation. However, we could have acted differently, considering from the beginning logics that do need a different notion of bisimulation. None of the logics we encountered so far give us the simple ability of *counting* successors, although we are able to do so in First Order Logic with equality. The basic idea of the graded modalities is to extend Basic Modal Logic with *graded operators* that give this possibility: we define, for any natural number n and action $a \in \Lambda$, the operator $\langle a \rangle_n \phi$, whose interpretation in a structure $\mathcal{M}$ is given by

$$[\![\langle a \rangle_n \phi]\!] = \{s \in \mathcal{M} : |[\![\phi]\!] \cap R_a(s)| \geq n\}.$$

Notice that the new formulae we obtain are not invariant for the basic form of bisimulation.

The *Logic of Finite Graded Modalities* $\mathcal{L}^{\diamond<\aleph_0}$ is defined as Basic Modal Logic extended with the operators $\langle a\rangle_n\phi$ for $n<\omega$ and $a\in\Lambda$. The logic where we only allow the use of a graded operator $\langle a\rangle_m$ if $m<n$, for a fixed natural number n, is denoted by $\mathcal{L}^{\diamond<n}$.

The idea of graded operators, introduced in the seventies (see [34], [42], [28]) and studied in the eighties in [27] where a sound and complete system for the logic $\mathcal{L}^{\diamond<\aleph_0}$ is proposed, is a simple and pervasive notion. It can be naturally connected with many other formalisms coming from the areas of generalized quantifiers theory and knowledge representation. A survey of such connections is presented in [37].

Many technicals results and classical tools from modal model theory can be applied to graded modalities. A detailed study of expressiveness issues using such instruments can be found in [36].

Consider now a sequence $\phi_0,\ldots,\phi_{m-1}$ of formulae and suppose we want to ask for m different successors $w_0,\ldots,w_{m-1}$ such that w_i satisfies ϕ_i, for all $i<m$. This seems a more difficult task than the previous one, but we can prove that we can express this property using the graded operators $\langle a\rangle_h$, for $h\le m$. More formally, consider for any $m\in\omega$ the operator $\{a\}_m$, which applies to a sequence $(\phi_0,\ldots,\phi_{m-1})$ of formulae and is interpreted in a structure $\mathcal{M}$ by:

$$s\models\{a\}_m(\phi_0,\ldots,\phi_{m-1})$$
$$\Updownarrow$$
$$\text{there are } m \text{ different successors } y_0,\ldots,y_{m-1} \text{ of } s \text{ with } y_i\models\phi_i, \text{ for all } i<m.$$

We can prove that the operator $\{a\}_m$ can be expressed using the operators $\langle a\rangle_h$, for $h\le m$. The proof (see [1] and [51], Th 1.1.6.) is an interesting example of an application to Modal Logic of a combinatorial result known as *the Marriage Theorem*, which we state in the following. Let $\{A_0,\ldots,A_{m-1}\}$ be a finite family of sets. A set of *distinct representatives* for this family is a set $\{a_0,\ldots,a_{m-1}\}$ such that $a_i\in A_i$ for all $i<m$ and $a_i\neq a_j$ for $i\neq j$, $i<m, j<m$. The Marriage Theorem gives a necessary and sufficient condition that a family of set must satisfy in order to have a set of distinct representatives.

2.4.1. THEOREM. *A finite family $\{A_0,\ldots,A_{m-1}\}$ of sets has a set of distinct representatives iff for all $k\le m$, $0\le i_1<\ldots<i_k<m$, the set $A_{i_1}\cup\ldots\cup A_{i_k}$ has cardinality greater than or equal to k.*

We use this theorem to show that the operator $\{a\}_m$ can be expressed by means of the operators $\langle a\rangle_i$, for $i<m$.

2.4.2. PROPOSITION. *If $m<n$ and $\phi_0,\ldots,\phi_{m-1}$ are in $\mathcal{L}^{\diamond<n}$, then the formula $\{a\}_m(\phi_0,\ldots,\phi_{m-1})$ is equivalent to a formula in $\mathcal{L}^{\diamond<n}$.*

Proof.
Using the Marriage Theorem, we prove that the formula $\{a\}_m(\phi_0, \ldots, \phi_{m-1})$ is equivalent to the formula

$$\psi = \bigwedge_{i<m} \langle a\rangle \phi_i \wedge \bigwedge_{i_1<i_2<m} \langle a\rangle_2(\phi_{i_1} \vee \phi_{i_2}) \wedge \ldots \wedge$$

$$\bigwedge_{i_1<\ldots<i_{m-1}<m} \langle a\rangle_{m-1}(\phi_{i_1} \vee \ldots \vee \phi_{i_{m-1}}) \wedge \langle a\rangle_m(\phi_0 \vee \ldots \vee \phi_{m-1}).$$

It is clear that ϕ implies ψ. Vice versa, if $\mathcal{M}, w \models \psi$, then apply the Marriage Theorem to the sets $A_i = \{w' : wR_a w' \wedge w' \models \phi_i\}$. □

In this thesis we will consider the natural generalization of the operators $\langle a\rangle_n$ to infinite cardinals: for any cardinal h consider the operator $\langle a\rangle_h\phi$ whose interpretation in a model $\mathcal{M}$ is given by

$$[\![\langle a\rangle_h\phi]\!] = \{s \in \mathcal{M} : |[\![\phi]\!] \cap R_a(s)| \geq h\}.$$

These 'infinite' modalities are studied from an algebraic point of view in [46], where various kind of infinite cardinality quantifiers are added to weakly associative relation algebras and decision results are obtained for the resulting varieties. A study of the finite variants of these operations can be found in [51]. Considering these operators in a modal logic context gives:

2.4.3. DEFINITION. For any cardinal k the logic $\mathcal{L}^{\diamond<k}$ is defined as Basic Modal Logic extended with the operators $\langle a\rangle_h\phi$ for $h < k$ and $a \in \Lambda$.
$\mathcal{L}^{grad}$ is defined as Basic Modal Logic extended with all operators $\langle a\rangle_h\phi$ for all cardinals h and $a \in \Lambda$. □

A generalization of Proposition 2.4.2 to infinite cardinals can be found in Chapter 4.

The extensions of $\mathcal{L}^{\diamond<k}$, $\mathcal{L}^{grad}$ which allow conjunctions over sets of cardinality strictly smaller than k' are denoted respectively by $\mathcal{L}^{\diamond<k}_{k'}$, $\mathcal{L}^{grad}_{k'}$, while in $\mathcal{L}^{\diamond<k}_{\infty}$, $\mathcal{L}^{grad}_{\infty}$ arbitrary set conjunctions are permitted.

As we have seen, in the logics $\mathcal{L}^{\diamond<k}_{\infty}$ we are allowed to count up to cardinal k. In Chapter 4, while considering non-well-founded sets, we shall encounter the dual class of logics, in which we are allowed to count only with cardinals strictly greater than k.

2.4.4. DEFINITION. If k is a cardinal, $\mathcal{L}^{k<\diamond}_{\infty}$ is the fragment of $\mathcal{L}^{grad}_{\infty}$ which allows the use of the operators $\langle a\rangle_h$ only for $k < h$ and $h = 1$ (that is, in these logics we always have the classical diamond operator). □

Given two cardinals $k < k'$ we will also use the logic $\mathcal{L}_\infty^{k<\diamond<k'}$, defined as the fragment of $\mathcal{L}_\infty^{grad}$ which allows the use of the operators $\langle a\rangle_h$ only for $k < h < k'$ and $h = 1$. As usual, the use of a cardinal k'' as a subscript in denoting one of the previous logics means that we are considering the fragment which only allows conjunctions on sets of cardinality strictly less than k''.

Translations of the graded logics.

We will use the following extension of the standard translation.

First, we extend the standard translation ST from Basic Modal Logic $\mathcal{L}^\diamond$ to the Logic of Graded Modalities $\mathcal{L}^{\diamond<\aleph_0}$ as follows:

$$ST(\langle a\rangle_n\phi) := \exists y_0 \dots y_{n-1} \; (\bigwedge_{i<n} xR_a y_i \wedge \bigwedge_{i<j} y_i \neq y_j \wedge \bigwedge_{i<n} ST(\phi)(y_i \mid x)),$$

(where $y_0, \dots, y_{n-1}$ are variables new for $ST(\phi)$). $ST(\phi)(r)$ is then a first-order sentence with equality which is equivalent to ϕ.

The extension of ST to $\mathcal{L}_\infty^{\diamond<\aleph_0}$ is obvious and results in a sentence of $\mathcal{L}_\infty$.

We shall also use the translation ST naturally extended to the full Logic of Graded Modalities $\mathcal{L}_\infty^{grad}$, but in this case the range of the translation is the Infinitary First Order Logic $\mathcal{L}_\infty(Q_\infty)$ (see Section 2.1).

Counting bisimulations

A bisimulation is insensible to the number of successors of a given node: bisimilar nodes can have a different number of successors. It is then clear that, in order to obtain a suitable definition of bisimulation for graded modal logics, we have to restrict the notion of bisimulation. We first consider the case of the logics $\mathcal{L}^{\diamond<k}$, for a cardinal number k and then give the dual definition, for the logics $\mathcal{L}^{k<\diamond}$.

2.4.5. Definition. Let k be a cardinal. A $(< k)-$*counting bisimulation* between two models $\mathcal{M}, \mathcal{N}$ of a language $\mathcal{L}$ is a relation $Z \subseteq \mathcal{M} \times \mathcal{N}$ such that if wZv then:

1. $r^\mathcal{M} Z r^\mathcal{N}$;
2. if wZv then w, v validate the same propositional constants of $\mathcal{L}$;
3. if wZv then for any action $a \in \Lambda$, if X is a subset of $R_a(w)$ of cardinality strictly smaller than k, then there is an injective function f between X and $R_a(v)$ such that $w' Z f(w')$, for all $w' \in X$,
4. and vice versa. □

If Z is a $(< k)$-counting bisimulation between $\mathcal{M}$ and $\mathcal{N}$, we say that $\mathcal{M}$ and $\mathcal{N}$ are $(< k)$-*counting bisimilar.*

If $\mathcal{M}$ is a structure for a language $\mathcal{L}$ and $\mathcal{N}$ is a structure for a language $\mathcal{L}'$ we will also consider $(< k)$-counting bisimulation w.r.t. a language $\mathcal{L}_1$ contained in $\mathcal{L} \cap \mathcal{L}'$: the definition is as above except that in 2) the nodes w, v validate the same propositional constants of $\mathcal{L}_1$ and 3),4) are restricted to atomic actions of this language.

2.4.6. LEMMA. *Let k be cardinal. If $\mathcal{M}$ is an $\mathcal{L}$ structure, $\mathcal{N}$ is an $\mathcal{L}'$ structure, $\mathcal{L}_1 \subseteq \mathcal{L} \cap \mathcal{L}'$, and Z is a $(< k)-$counting bisimulation between $(\mathcal{M}, w)$ and $(\mathcal{N}, v)$ w.r.t. $\mathcal{L}_1$, then w and v satisfy the same $(\mathcal{L}_1)_\infty^{\diamond < k}$ formulae.*

Proof.
By induction on the complexity of formulae. We just deal with one case. If $w \models \langle a \rangle_h \phi$ for $h < k$ we consider the set $\{w' : wR_a w', w' \models \phi\}$, which has a subset X of cardinality h. Let f be the corresponding injective function as in in Definition 2.4.5. By induction all points in $f(X)$ validate ϕ and $v \models \langle a \rangle_h \phi$.

2.4.7. REMARK. Notice that we could give an apparently weaker definition of bisimulation in such a way that Proposition 2.4.6 still holds. We could change points 3),4) of Definition 2.4.5 into the following: 3′) if X is a subset of $R_a(w)$ of cardinality strictly smaller than k then $|\{y \in R_a(v) : \exists x \in X \; xZy\}| \geq |X|$ (without asking the existence of an injective function preserving Z); 4′) and vice versa. However, it can be shown that a bisimulation Z of this type between models $\mathcal{M}, \mathcal{N}$ is a $(< k)$-counting bisimulation as defined in 2.4.5. To see this, we prove that if Z is a bisimulation between $\mathcal{M}$ and $\mathcal{N}$ that satisfies condition 3′), then it satisfies condition 3) of Definition 2.4.5. Suppose that wZv and let $X = \{x_0, \ldots, x_{r-1}\}$ be a subset of $R_a(w)$ of cardinality equal to $r < k$. We prove that the family $\{A_0, \ldots, A_{r-1}\}$, where $A_i = \{y \in R_a(v) : x_i Z y\}$, has a set of distinct representatives by checking the hypotheses of the Marriage Theorem 2.4.1. If $h \leq r$ and $0 \leq i_1 < \ldots < i_h < r$, we have to show that $|A_{i_1} \cup \ldots \cup A_{i_h}| \geq h$. Consider the set $X' = \{x_{i_1}, \ldots, x_{i_h}\}$; since $|X'| < k$ by 3′) we have that $|\{y \in R_a(v) : \exists x' \in X' \; x'Zy\}| \geq |X'|$. But $\{y \in R_a(v) : \exists x' \in X' \; x'Zy\} \subseteq A_{i_1} \cup \ldots \cup A_{i_h}$ and so $|A_{i_1} \cup \ldots \cup A_{i_h}| \geq h$. Since the hypotheses of the Marriage Theorem are satisfied, the family $\{A_0, \ldots, A_{r-1}\}$ has a set of distinct representatives $\{a_0, \ldots, a_{r-1}\}$ and the function $f(x_i) = a_i$ is an injective function between X and $R_a(v)$ with $xZf(x)$ for all $x \in X$. Hence Z satisfies condition 3) of 2.4.5. □

Next, we consider the notion of bisimulation which corresponds to the logic which contains $\langle a \rangle_h$ only if h is bigger than a fixed cardinal k, that is, the logic $\mathcal{L}^{k < \diamond}$:

2.4.8. DEFINITION. A $(k <)$-*counting bisimulation* Z between two structures $\mathcal{G}$ and $\mathcal{H}$ is a relation $Z \subseteq \mathcal{G} \times \mathcal{H}$ such that

1. $r^{\mathcal{M}} Z r^{\mathcal{N}}$;

2. if wZv then w, v validate the same propositional constants;

3. if wZv and if X is a subset of $R_a(w)$ with either $|X| = 1$ or $|X| > k$ then $|\{y \in R_a(v) : \exists x \in X\ xZy\}| \geq |X|$;

4. vice versa. □

Given two structures $\mathcal{G}$ and $\mathcal{H}$, we write $\mathcal{G} \equiv_{\infty}^{k<\diamond} \mathcal{H}$ if the two structures satisfy the same formulae of $\mathcal{L}_{\infty}^{k<\diamond}$.

2.4.9. LEMMA. *If Z is a $(k <)$-counting bisimulation between $\mathcal{G}$ and $\mathcal{H}$ then $\mathcal{G} \equiv_{\infty}^{k<\diamond} \mathcal{H}$.*

We leave the easy proof to the reader.

Finally, the appropriate notion of bisimulation for $\mathcal{L}_{\infty}^{grad}$ is easily guessed.

2.4.10. DEFINITION. A *counting bisimulation* Z between two structures $\mathcal{G}$ and $\mathcal{H}$ is a relation $Z \subseteq \mathcal{G} \times \mathcal{H}$ such that

1. $r^{\mathcal{M}} Z r^{\mathcal{N}}$;

2. if wZv then w, v validate the same propositional constants;

3. if wZv and X is a subset of $R_a(w)$ then $|\{y \in R_a(v) : \exists x \in X\ xZy\}| \geq |X|$;

4. vice versa. □

Given two structures $\mathcal{G}$ and $\mathcal{H}$, we write $\mathcal{G} \equiv_{\infty}^{grad} \mathcal{H}$ if the two structures satisfy the same formulae of $\mathcal{L}_{\infty}^{grad}$.

2.4.11. LEMMA. *If Z is a counting bisimulation between $\mathcal{G}$ and $\mathcal{H}$ then $\mathcal{G} \equiv_{\infty}^{grad} \mathcal{H}$;*

The converse of Lemma 2.4.6, Lemma 2.4.9 and Lemma 2.4.11 can also be proved. In Chapter 5, in the context of non-well-founded set theory, we will prove a strong form of these converses.

2.5 Interpolation.

In this section we give a brief account of the various forms of interpolation that we consider in this thesis, describing the behavior of the logics we introduced so far with respect to it.

2.5.1. DEFINITION. Given a logic, an *interpolant* θ for a pair (ϕ, ψ) of formulae of the logic is a formula of the logic in the language $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$ such that

- $\phi \models \theta$,
- $\theta \models \psi$. □

It is clear that a necessary condition for two formulae to have an interpolant is that $\phi \models \psi$. A logic is said to have *Craig interpolation* if each pair of formulae (ϕ, ψ) with $\phi \models \psi$ have an interpolant.

Craig interpolation (or, simply, interpolation) was first proved for First Order Logic by Craig ([18]). Let us consider the other logics introduced so far with respect to interpolation. Monadic Second Order logic and Infinitary Logic $\mathcal{L}_\infty$ are examples of logics that do not have the interpolation property. If we restrict the use of conjunctions in $\mathcal{L}_\infty$ to countable sets of formulae, however, we obtain the logic $\mathcal{L}_{\omega_1}$ which does have interpolation. This result was first proved by Lopez-Escobar in 1965 ([48]) with a proof-theoretical argument using cut elimination and can also be proved using *consistency properties* for Infinitary Logic. It is the latter proof that one can generalize to get a weaker form of interpolation for $\mathcal{L}_\infty$, as we shall see in Section 2.5.2.

As for the extended modal logics we defined, we have the following situation:

1. Basic Modal Logic enjoys interpolation. There is a proof-theoretic argument using a Gentzen-style calculus of sequents, as well as a model-theoretical argument, using a strong form of amalgamation and bounded bisimulation.

2. Infinitary Modal Logic $\mathcal{L}_\infty^\diamond$ as well as countable Infinitary Modal Logic $\mathcal{L}_{\omega_1}^\diamond$ have the interpolation property. Interpolation for $\mathcal{L}_\infty^\diamond$ is proved in [11] using amalgamation and *consistency properties modulo bisimulation*, that we will encounter in Chapter 5. The countable version is proved in [64], which also gives a completeness theorem for this logic.

3. It is not known whether PDL enjoys interpolation or not.

4. The Modal μ-Calculus enjoys interpolation. We give a proof of this in Chapter 3.

5. In the family of graded modal logics, interpolation is known to be true for $\mathcal{L}^{\Diamond<\aleph_0}$ and false for $\mathcal{L}^{\Diamond<n}$, for a natural number n (see [1]). In Chapter 5 we prove the interpolation property for the logic $\mathcal{L}^{\Diamond<\aleph_0}_{\omega_1}$, as well as for the full Logic of Graded Modalities $\mathcal{L}^{grad}_{\infty}$.

By looking at points 2,4, we see that in these cases the modal fragments behave better than their classical environments: $\mathcal{L}^{\Diamond}_{\infty}$ has interpolation, while $\mathcal{L}_{\infty}$ does not and the same holds for the Modal μ-Calculus with respect to Monadic Second Order Logic.

2.5.1 Uniform interpolation.

We now consider a stronger form of interpolation, the so-called *uniform interpolation*. To introduce this notion, let us first consider classical Propositional Logic. In this logic, given a formula $\phi(p)$ one can easily prove that the formula $\phi(\top/p) \vee \phi(\bot/p)$ serves as an interpolant for all formulae ψ with $\phi \models \psi$ and $\mathcal{L}(\phi) \cap \mathcal{L}(\psi) \subseteq \mathcal{L}(\phi) \setminus \{p\}$, that is, we have $\phi \models \theta$ and for all ψ as above we have $\theta \models \psi$. In other words: the interpolant depends on ϕ and on the common language of ϕ and ψ, but not on ψ. This can be easily generalized from $\mathcal{L}(\phi) \setminus \{p\}$ to arbitrary subsets of $\mathcal{L}(\phi)$: we say that classical Propositional Logic enjoys uniform interpolation, a property that can be stated for general logics:

2.5.2. DEFINITION. Given a logic, a sentence ϕ in it and a language $\mathcal{L}' \subseteq \mathcal{L}(\phi)$, a *uniform interpolant* of ϕ with respect to $\mathcal{L}'$ is a formula θ such that:

1. $\mathcal{L}(\theta) \subseteq \mathcal{L}'$;
2. $\phi \models \theta$;
3. if $\phi \models \psi$ and $\mathcal{L}(\phi) \cap \mathcal{L}(\psi) \subseteq \mathcal{L}'$ then $\theta \models \psi$. □

A logic enjoys *uniform interpolation* if any sentence ϕ of the logic has a uniform interpolant with respect to any $\mathcal{L}' \subseteq \mathcal{L}(\phi)$.

Thus, classical Propositional Logic enjoys uniform interpolation. Intuitionistic Propositional Logic IL is another example of a logic in which uniform interpolation holds, albeit in this context the proof becomes nontrivial. It was first proved by Pitts in [56], by means of proof theoretical methods. Using a *contraction free* variant of the Gentzen sequent calculus for IL, it is possible to define a well-founded relation on sequents in such a way that the hypotheses of a rule of the calculus are always smaller than the conclusion. The proof of the existence of the uniform interpolant is done by recursion on this relation.

If a logic does not have interpolation, then it is clear that it cannot have uniform interpolation. First Order Logic is an example of a logic having interpolation that does not have the uniform version of it (this was first proved in [35]). To see this, consider the first-order sentence $\phi(R)$ expressing the fact that R is irreflexive, transitive and serial (that is: $\forall x \exists y x R y$). Suppose ψ is a uniform interpolant for $\phi(R)$ with respect to $L(\phi) \setminus \{R\}$. ψ may thus only contain $=$. Then one can show that ψ must be true precisely in the infinite models. As infinity cannot be expressed in First Order Logic, we have a contradiction.

Let us see the behavior of the logic we have encounter so far with respect to uniform interpolation (obviously, we only consider the logics that do have interpolation, or for which it is unknown).

1. Basic Modal Logic enjoys uniform interpolation. We will sketch a proof of this fact in Chapter 3.

2. Infinitary Modal Logic $\mathcal{L}^{\Diamond}_{\infty}$, its fragment $\mathcal{L}^{\Diamond}_{\omega_1}$, and PDL do not have uniform interpolation. This will be proved in Chapter 3.

3. The Modal μ-Calculus enjoys uniform interpolation. We give a proof of this in Chapter 3.

4. The logics $\mathcal{L}^{\Diamond<\aleph_0}_{\omega_1}$ and $\mathcal{L}^{grad}_{\infty}$ do not have uniform interpolation (this follows from a result in [21]). A way to prove uniform interpolation for $\mathcal{L}^{\Diamond<\aleph_0}$ could be to use bounded $<\aleph_0$-counting bisimulation and amalgamation, as for $\mathcal{L}^{\Diamond}$. However, amalgamation in this context would probably be a more difficult task than in the basic modal case.

2.5.2 Elementary interpolation.

As we already pointed out, not all logics have interpolation. Can we characterize the pairs of formulae which do have an interpolant in a reasonable way, in logic where interpolation fails? This question can be posed for example for: Infinitary Logic, MSO, First Order Logic on finite models, First Order Logic with the quantifier *there exists uncountably many*, and, on the modal side, Graded Modal Logic with a finite number of operators. In these logics, $\phi \models \psi$ is a necessary but not sufficient condition for (ϕ, ψ) to have an interpolant, hence we have to look for a stronger condition.

The problem of characterizing pairs with interpolant in logic without interpolation is considered in [5] and solved for Infinitary Logic as follows. The idea is to *strengthen* the condition $\phi \models \psi$ in order to obtain all pairs having an interpolant. Since $\phi \models \psi$ means that given a model $\mathcal{M}$ the truth of ϕ implies the truth of ψ, we could strengthen this condition by asking this transfer property not only inside a single model $\mathcal{M}$, but from a model $\mathcal{M}$ to any other model which is related to $\mathcal{M}$ in a *reasonable* way. For example, we could ask that if $\mathcal{M} \models \phi$ then $\mathcal{N} \models \psi$ for any $\mathcal{N}$ which is *isomorphic* to $\mathcal{M}$ w.r.t. the language $\mathcal{L} = \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$. After all, this is certainly a necessary condition for an interpolation pair! However, this transfer property gives nothing new because it is equivalent to $\phi \models \psi$: if $\phi \models \psi$,

$\mathcal{M} \models \phi$ and $\mathcal{M}$ is isomorphic to $\mathcal{N}$, w.r.t. $\mathcal{L}$, consider the model $\mathcal{N}^*$ obtained from $\mathcal{N}$ by changing the interpretation of any predicate in $\mathcal{L}(\phi) \setminus \mathcal{L}(\psi)$ in such a way that $\mathcal{N}^* \models \phi$ (just copy the interpretation of the predicate in $\mathcal{M}$ via the $\mathcal{L}$-isomorphism); by hypothesis, $\mathcal{N}^* \models \psi$, hence $\mathcal{N} \models \psi$ because $\mathcal{N}$ and $\mathcal{N}^*$ agree on $\mathcal{L}(\psi)$.

We are however in the right direction. The next choice is to consider this transfer along $\mathcal{L}$−elementary equivalence, instead of along $\mathcal{L}$-isomorphism. Suppose (ϕ, ψ) *have* an interpolant θ, $\mathcal{M} \models \phi$, and $\mathcal{M}$ and $\mathcal{N}$ satisfy the same formulae of the logic in the language $\mathcal{L}$. Then $\mathcal{M} \models \theta$, $\mathcal{N} \models \theta$ since θ is an $\mathcal{L}$-formula, and from $\theta \models \psi$ we obtain $\mathcal{N} \models \psi$. This kind of transfer is at *least* a necessary condition.

2.5.3. Definition. We say that a logic has *elementary interpolation* if pairs having an interpolant can be characterized as the pairs (ϕ, ψ) such that ϕ entails ψ along $\mathcal{L} = \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$-equivalence, that is:

$$\text{if } \mathcal{M} \models \phi \text{ and } \mathcal{M} \equiv_{\mathcal{L}} \mathcal{N} \text{ then } \mathcal{N} \models \psi.$$

□

The first natural questions to answer are the following. Does a logic without interpolation which has elementary interpolation exist? Does a logic which does not have elementary interpolation exist?

The following well-known argument shows that we must look for a non-compact counterexample.

2.5.4. Proposition. *Any compact logic enjoys elementary interpolation.*

Proof.
Suppose the logic is compact and ϕ entails ψ along $\mathcal{L} = \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$-equivalence w.r.t. this logic. If $Th_{\mathcal{L}}(\phi)$ denotes the set

$$\{\theta : \theta \text{ is in the logic, } \mathcal{L}(\theta) \subseteq \mathcal{L} \text{ and } \ \phi \models \theta\},$$

we prove that

$$Th_{\mathcal{L}}(\phi) \models \psi,$$

and we get an interpolant by compactness. Suppose then that $\mathcal{M} \models Th_{\mathcal{L}}(\phi)$. Using compactness again it follows that $Th_{\mathcal{L}}(\mathcal{M}) \cup \{\phi\}$ has a model: otherwise there would be a $\psi \in Th_{\mathcal{L}}(\mathcal{M})$ with $\phi \models \neg\psi$, that is, $\neg\psi \in Th_{\mathcal{L}}(\phi)$, which contradicts $\mathcal{M} \models Th_{\mathcal{L}}(\phi)$. But a model $\mathcal{N}$ of $Th_{\mathcal{L}}(\mathcal{M}) \cup \{\phi\}$ is such that $\mathcal{N} \models \phi$ and $\mathcal{N}$ satisfies the same $\mathcal{L}$-formulae as $\mathcal{M}$. Hence from the hypothesis we get $\mathcal{M} \models \psi$. □

We find a counterexample to elementary interpolation in Monadic Second Order logic (this counterexample is from [21]):

2.5.5. THEOREM. *Monadic Second Order Logic does not have elementary interpolation.*

Proof.
Let $\phi(<)$ be a sentence expressing that $<$ is a strict linear order. Let $\psi(<)$ be the conjunction of the following:

1. $<$ has a least point (which we'll refer to as `min`) and a greatest point (`max`);
2. Either $x = \mathtt{max}$ or x has a direct successor $S(x)$.
3. 'Induction':

$$\forall P.([P(\mathtt{min}) \wedge \forall x.((x \neq \mathtt{max} \wedge P(x)) \rightarrow P(S(x)))] \rightarrow \forall x.P(x))$$

$\phi(<) \wedge \psi(<)$ can only be true in finite models. Any linear ordering $<$ on a finite model must be of the form $s_0 < s_1 < \ldots < s_n$, so on finite models $\phi(<)$ implies $\psi(<)$. Thus $\phi(R_a) \wedge \psi(R_a) \models \phi(R_b) \rightarrow \psi(R_b)$.

We can also prove that $\phi(R_a) \wedge \psi(R_a)$ entails $\phi(R_b) \rightarrow \psi(R_b)$ along *MSO*-equivalence w.r.t. the language of pure equality. If $\mathcal{M} \models \phi(R_a) \wedge \psi(R_a)$ then $\mathcal{M}$ is finite and if $\mathcal{N} \equiv_{MSO} \mathcal{M}$ then $\mathcal{N}$ is finite too. Hence we have $\mathcal{N} \models \phi(R_b) \rightarrow \psi(R_b)$.

We now show that an interpolant χ for these two sentences would be a sentence expressing finiteness.

For suppose $\mathcal{M}$ is finite, say its universe is $\{s_0, \ldots, s_n\}$. Define R_a by $s_i R_a s_j$ iff $i < j$. Clearly then $\mathcal{M} \models \phi(R_a) \wedge \psi(R_a)$. So $\mathcal{M} \models \chi$.

For the converse, suppose $\mathcal{M} \models \chi$. We may order the domain of $\mathcal{M}$ linearly and denote this ordering by R_b. Because χ is the interpolant, $(\mathcal{M}, R_b) \models \phi(R_b) \rightarrow \psi(R_b)$. As $\phi(R_b)$ holds in $(\mathcal{M}, R_b)$, $\psi(R_b)$ must also hold there. But then $\mathcal{M}$ must be finite.

As finiteness cannot be expressed in Monadic Second Order logic we have a contradiction. □

Hence, not all logics have elementary interpolation. Another counterexample is First Order Logic on finite models, but we omit the proof.

In [5] Barwise and van Benthem proved that Infinitary Logic has elementary interpolation. They actually proved that interpolation pairs can be characterized as the pairs (ϕ, ψ) such that ϕ entails ψ along *potential isomorphism*. Since two structures satisfy the same infinitary sentences iff they are potentially isomorphic, this result gives elementary interpolation for Infinitary Logic. The proof uses essentially the *Boundedness Theorem* for Infinitary Logic, the definability of potential isomorphism as the greatest fixed point of a monotone operator definable inside the logic, and the *Scott property* of this logic. The same proof applies to any logic satisfying these conditions.

As it is shown in [11], elementary interpolation for infinitary logic can also be proved using the notion of *good triples*, which is a generalization of the consistency property of Infinitary Logic. This method can be adapted to prove elementary interpolation for extended modal logics such as the (infinitary) graded modal logics (see Chapter 5).

2.6 Non-well-founded sets and bisimulation.

Finally, we introduce the basic material on non-well-founded sets. Since this subject is probably not very familiar for many readers, in this section we give more proofs than in the preceding ones. Whoever knows Aczel book [2] may just read Theorem 2.6.5 and use this section just for notation when reading Chapter 4.

Since the Russel Paradox, sets that belong to themselves have been looked at suspiciously. The most used axiomatization of set theory, the Zermelo Fraenkel one, contains the axiom of foundation that forces the membership relation between sets to be (inversely) well-founded: under this axiomatization we can never have an infinite chain of sets $a_1, a_2, \ldots$ with $a_{i+1} \in a_i$, for all i. In particular, this axiom does not allow the existence of sets that belong to themselves. However, non-well-founded sets *exist*, in the following sense:

- there are other possible axiomatizations of set theory which imply the negation of the axiom of foundation;
- these axiomatizations are consistent relative to the Zermelo Fraenkel theory without the axiom of foundation (ZFC^-, from now on).

In these theories, sets can belong to themselves, there can be sets a, b with $a \in b$ and $b \in a$, and so on. Non-well-founded axiomatizations, beside being of interest because they give more insight on the structure of sets, have applications to Communicating Systems (see [2]) and Situation Theory (see [6]).

In this section we give a brief introduction to one of the most famous non-well-founded axiomatizations, the theory $ZFC^- + AFA$, considering in particular its relations with the notion of bisimulation and Modal Logic.

In order to describe the anti-foundation axiom AFA, consider a structure $\mathcal{G}$ for the language $\{R, r\}$, where R is a binary relation symbol and r is a constant symbol; suppose also that any element $w \in \mathcal{G}$ can be reached in a finite number of $R^{\mathcal{G}}$ steps from $r^{\mathcal{G}}$. Given a structure $\mathcal{G}$, a *decoration* of it is an assignment of a set to each node of the structure, in such a way that if a is assigned to w then a is the set of all elements that are assigned to $R^{\mathcal{G}}$-successors of w. More precisely, a decoration of $\mathcal{G}$ is a function whose domain contains $\mathcal{G}$, such that

$$d(w) = \{d(v) : wR^{\mathcal{G}}v\}.$$

It is obvious that in ZFC only well-founded structures can have a decoration, where a structure is called *well-founded* if no infinite path $w_0 R^{\mathcal{G}} w_1 R^{\mathcal{G}} \ldots$ exists on it. Working in ZFC^- one can prove the converse, which is known as Mostowski's Collapsing Lemma: any well-founded-structure $\mathcal{G}$ has a *unique* decoration.

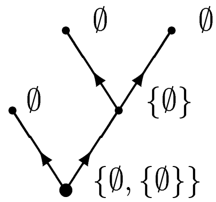

A decorated structure.

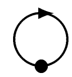

A structure without a ZFC-decoration.

What about other $\{R, r\}$-structures (which we simply call structures, from now on)? If the axiom of foundation *is* present, then there are structures *without* any decoration, for example, a single reflexive node: $\mathcal{G} = \{w\}$, with $wR^{\mathcal{G}}w$. On the other hand, it is possible to show that (a formalization of) the property:

any structure has a unique decoration

is consistent with ZFC^-. This axiom is called AFA in [2]. If we work in $ZFC^- +$ AFA, we will indicate by $d_{\mathcal{G}}$ the (unique) decoration of the structure $\mathcal{G}$.

What is the role of bisimulation in this context? Using only ZFC^-, we see that if two structures $\mathcal{G}$ and $\mathcal{G}'$ *have a common decoration* (that is, a decoration d of $\mathcal{G}$ and $\mathcal{G}'$ such that $d(r^{\mathcal{G}}) = d(r^{\mathcal{G}'})$), then they are bisimilar: the relation $S \subseteq \mathcal{G} \times \mathcal{G}'$ defined by $uSv \Leftrightarrow d(u) = d(v)$ is a bisimulation between $\mathcal{G}$ and $\mathcal{G}'$ w.r.t. the language $\{R, r\}$.

The converse cannot be proved in ZFC^-, but it is true under AFA (see Corollary 2.6.2). To prove this result, we use the notion of collapse, introduced in Definition 2.2.2: under ZFC^- not all structures have a decoration, but we can prove that if $Coll(\mathcal{G})$ has a decoration, then we have a decoration of $\mathcal{G}$ as well.

2.6.1. Proposition. *(ZFC^-) If d is a decoration of $Coll(\mathcal{G})$, then the function d^* defined on $\mathcal{G}$ by $d^*(x) = d([x])$ is a decoration of $\mathcal{G}$.*

Proof.
To prove that d^* is a decoration of $\mathcal{G}$ we have to show that for all $x \in \mathcal{G}$ it holds:

$$d^*(x) = \{d^*(y) : xR^{\mathcal{G}}y\}.$$

If $xR^{\mathcal{G}}y$, then $[x]R^{Coll(\mathcal{G})}[y]$ and $d^*(y) = d([y]) \in d([x]) = d^*(x)$.

Vice versa, if $z \in d^*(x) = d([x])$, then $z = d([y])$, with $[x]R^{Coll(\mathcal{G})}[y]$. By definition of $R^{Coll(\mathcal{G})}$, there exists a y' such that $xR^{\mathcal{G}}y'$ and $[y'] = [y]$. But then $d([y']) = d([y])$ and $z = d([y']) = d^*(y')$ with $xR^{\mathcal{G}}y'$. □

2.6.2. COROLLARY. *(ZFC^- + AFA) Two structures $\mathcal{G}$ and $\mathcal{G}'$ have the same decoration iff they are bisimilar.*

Proof.
We already proved the direction from left to right in ZFC^-. Vice versa, if $\mathcal{G}$ is bisimilar to $\mathcal{G}'$, then $Coll(\mathcal{G})$ is isomorphic to $Coll(\mathcal{G}')$ (see Theorem 2.2.3). By the above proposition and by the uniqueness of decorations we have:

$$d_{\mathcal{G}}(r^{\mathcal{G}}) = d_{Coll(\mathcal{G})}([r^{\mathcal{G}}]) = d_{Coll(\mathcal{G}')}([r^{\mathcal{G}'}]) = d_{\mathcal{G}'}(r^{\mathcal{G}'}),$$

proving that $\mathcal{G}$ and $\mathcal{G}'$ have the same decoration. □

Hence, the concept of bisimulation reappears in a set context, where it is used to decide *when* two structures have the same decoration.

Among all possible structures, we are particularly interested in those that are *pictures* of sets. Given a set a, define its *canonical structure* $\mathcal{G}_a$ as follows. The set of nodes of $\mathcal{G}_a$ is given by the transitive closure of a together with the node a, which represents the root of $\mathcal{G}_a$ and $uR^{\mathcal{G}_a}v$ iff $v \in u$. We say that $\mathcal{G}_a$ is a *picture* of the set a. What structures are pictures of sets, i.e. what structures are (isomorphic to) canonical structures? Notice that a canonical structure has the identity function $d(x) = x$ as decoration, hence, a canonical structure has an injective decoration. Vice versa, if a structure $\mathcal{G}$ has an injective decoration d, then d is an isomorphism between the structure $\mathcal{G}$ and the canonical structure of the set $d(\mathcal{G})$. Hence we can reply to the previous question by saying that a structure is a canonical structure iff it has an injective decoration. However, this characterization is in a sense circular, because it describes structures that are pictures of sets as the ones in which different nodes are decorated by different sets.

We can use the notion of bisimulation to give a characterization of canonical structures that does not involve decorations.

A structure $\mathcal{G}$ is said to be *strongly extensional* if different nodes in it are never bisimilar, that is, if $w, v \in \mathcal{G}$ and w is bisimilar to v, then $w = v$. We have:

2.6.3. PROPOSITION. *(ZFC^-) The axiom AFA is equivalent to:*

AFA^{ext} := A structure is strongly extensional iff it is a canonical structure.

Proof.
Suppose AFA holds. By the discussion above, to prove AFA^{ext} it is enough to show that

A structure is strongly extensional iff it has an injective decoration.

If $\mathcal{G}$ is strongly extensional, then by Corollary 2.6.2 its decoration $d_{\mathcal{G}}$ is injective. Vice versa, if $\mathcal{G}$ has an injective decoration two different nodes of $\mathcal{G}$ cannot be bisimilar by Corollary 2.6.2.

Vice versa, suppose that AFA^{ext} holds and consider a structure $\mathcal{G}$. Then $Coll(\mathcal{G})$ is strongly extensional, hence by AFA^{ext} it is canonical and has an injective decoration $d_{Coll(\mathcal{G})}$. This decoration can be used as in the proof of the preceding proposition to define a decoration $d(x) = d_{Coll(\mathcal{G})}([x])$ of $\mathcal{G}$. In this way we proved that any structure $\mathcal{G}$ has *at least* one decoration.

It remains to prove that a structure cannot have two different decorations. Suppose that d_1, d_2 are decorations of a structure $\mathcal{G}$ and consider the two sets $a = d_1(r^{\mathcal{G}})$ and $b = d_2(r^{\mathcal{G}})$. We must prove that $a = b$.

1. We claim that the canonical structures of a and b are bisimilar.

   Consider the following relation $S \subseteq \mathcal{G}_a \times \mathcal{G}_b$:

   $$uSv \Leftrightarrow \text{ there exists } x \in \mathcal{G} \text{ with } d_1(x) = u \text{ and } d_2(x) = v.$$

   We prove that S is a bisimulation between the canonical structures of a and b. By hypothesis we have aSb. If $c \in a$, then since $a = d_1(r^{\mathcal{G}})$ and d_1 is a decoration of $\mathcal{G}$, there must be a nod $y \in \mathcal{G}$ with $r^{\mathcal{G}} R^{\mathcal{G}} y$ and $d_1(y) = c$. But then, since d_2 is a decoration of $\mathcal{G}$ as well, $d_2(y) \in d_2(r^{\mathcal{G}}) = b$, and $cSd_2(y)$. The symmetric condition of a bisimulation is proved similarly.

2. Using the previous claim, we can now prove that $a = b$. Consider the canonical structure of the set $c = \{a, b\}$. By AFA^{ext} this structure is strongly extensional and since $\mathcal{G}_a$ is bisimilar to $\mathcal{G}_b$ we obtain $a = b$. □

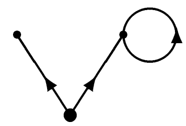

A picture of an AFA-set.

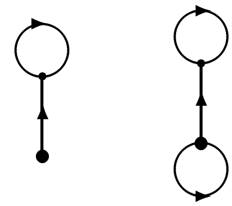

Structures which are not AFA-pictures.

The previous proposition tell us that in universes satisfying AFA, strongly extensional structures can be identified with (canonical structures of) sets. This allows us to prove the consistency of $ZFC^- + AFA$ relatively to ZFC: roughly speaking, one can construct a model of $ZFC^- + AFA$ out of a model of ZFC^- by considering the strongly extensional structures as the new sets, with $\mathcal{G} \in \mathcal{H}$ iff $\mathcal{G}$ is isomorphic to a structure of type $(\mathcal{H}, h)$, with $r^{\mathcal{H}} R^{\mathcal{H}} h$ (a detailed proof can be found in [2]). Another proof of the relative consistency of $ZFC^- + AFA$ with respect to ZFC^- can be found in Chapter 4, where a model for $ZFC^- + AFA$ is constructed by means of infinitary modal formulae.

In the next proposition we show that AFA^{ext} implies a strengthening of the extensionality axiom. Notice first that we can define bisimulation also on sets, via canonical structures: we say that two sets a and b are *bisimilar* if the structures $\mathcal{G}_a$, $\mathcal{G}_b$ are so. The extensionality axiom in ZFC^- says that two sets having the same elements are equal; using AFA^{ext} we see in the next proposition that we can use a weaker hypothesis and reach the same conclusion: under AFA^{ext}, if two sets are *bisimilar*, then they are equal:

2.6.4. PROPOSITION. *(ZFC^-) The following* 1) *and* 2) *are equivalent:*

1. *All canonical structures are strongly extensional.*

2. *If the set a is bisimilar to the set b then $a = b$.*

Proof.
$(1 \Rightarrow 2)$ If a is bisimilar to b, consider the set $c = \{a, b\}$ and its canonical structure $\mathcal{G}_c$. By 1, $\mathcal{G}_c$ must be strongly extensional, hence $a = b$.

$(2 \Rightarrow 1)$ Consider a canonical structure $\mathcal{G}$ If u, v are two bisimilar nodes in $\mathcal{G}$, then u and v are two bisimilar sets and from 2 we obtain $u = v$. This proves that $\mathcal{G}$ is strongly extensional. □

Let us stress the following once again: by the previous proposition it follows that two AFA sets are bisimilar if ad only if they are equal.

What is the role of Modal Logics in all this? Let us consider a very simple language, in which there is a single binary relation R, a constant r, and no propositional constants. An infinitary modal formula in this language can be evaluated in a set a, via the canonical structure of a: we say that the set a satisfy ϕ (in symbols: $a \models \phi$) iff $\mathcal{G}_a \models \phi$. Richer languages which allow a non-empty set of propositional constants can also be interpreted in sets, but in this case we must consider a universe with *urelements*, as it is done in [7]. For the sake of simplicity we restrict here to the language $\{R, r\}$, although all results in Chapter 4 can be generalized (using universes with urelements) to reacher languages.

We know (see Theorem 2.2.8) that for any structure $\mathcal{G}$ there exists a sentence $\phi_{\mathcal{G}} \in \mathcal{L}^{\diamond}_{\infty}$ such that for all structures $\mathcal{H}$ it holds:

$$\mathcal{H} \models \phi_{\mathcal{G}} \Leftrightarrow \mathcal{H} \text{ is bisimilar to } \mathcal{G}.$$

Since in any universe satisfying AFA bisimilar sets are equal, we see that any set admits a complete description in terms of a formula of $\mathcal{L}^{\diamond}_{\infty}$, that is:

2.6.5. Theorem. *For any AFA-set a there exists $\phi_a \in \mathcal{L}^{\diamond}_{\infty}$ such that for all AFA-sets b it holds:*

$$b \models \phi_a \Leftrightarrow b = a.$$

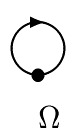

$$\phi_\Omega = \bigwedge_{n\in\omega} \Box^n \Diamond \top.$$

A circular AFA-set and its formula.

Hence, in a universe satisfying AFA sets are *described* by infinitary modal formulae. This suggests the possibility of building a model of $ZFC^{-} + AFA$ using infinitary modal formulae as sets (see Chapter 4).

Further investigations along the lines of Theorem 2.6.5 are considered by Baltag in [3] (see also [7]), where the following questions are addressed: what kind of sets can be described using only formulae or theories of the finitary logic $\mathcal{L}^{\diamond}$? To answer this question, denote by H^0 the set of all well-founded sets whose transitive closure is finite and by H^1 the largest collection of sets C such that every c

in C is a subset of C and has finite size: Baltag proved that for any AFA-set a the following holds:

$$a \in H^0 \Leftrightarrow \ a \text{ is characterizable by a formula of } \mathcal{L}^{\diamond};$$

$$a \in H^1 \Leftrightarrow \ a \text{ is characterizable by a theory of } \mathcal{L}^{\diamond}.$$

The last result was also obtained, in a modal logic setting, by Marco Hollenberg in [38].

Chapter 3

Uniform interpolation for the μ-Calculus.

The main result of this chapter is a proof of uniform interpolation for the Modal μ-Calculus. We prove this theorem by showing that certain non-standard second-order quantifiers, the so-called *bisimulation quantifiers* ([66]), are definable within the Modal μ-Calculus. These quantifiers and their relationships with uniform interpolation are of interest for all logics that are invariant under bisimulation and we will prove some general facts about them in Section 3.1. After that, the next task will be to prove that the Modal μ-Calculus has bisimulation quantifiers. The proof we give uses *μ-automata*, introduced by Janin and Walukiewicz in [40] to demonstrate that μ-sentences can be characterized as the fragment of MSO that is invariant for bisimulation. To prove this result, Janin and Walukiewicz ([40]) define a notion of automaton (a so-called *μ-automaton*) that operates on structures and which corresponds exactly to the Modal μ-Calculus on such structures. We devote Section 3.2 to a brief introduction to the subject. We use automata in Section 3.3, which contains the proof of uniform interpolation for the Modal μ-Calculus via bisimulation quantifiers. Finally, in Section 3.4 we give counterexamples to uniform interpolation in Infinitary Modal Logics and in PDL.

3.1 Bisimulation quantifiers and uniform interpolation

Suppose a logic $\mathcal{L}$ is closed under the propositional existential quantifier $\exists p\phi$, that is, for any formula ϕ of the logic there exists formula θ, in the language $\mathcal{L}(\phi)\setminus\{p\}$, such that for any structure $\mathcal{M}$ it holds:

$$\mathcal{M} \models \theta \Leftrightarrow \text{ there is a valuation of } p \text{ in } \mathcal{M} \text{ such that } \mathcal{M}, p \models \phi.$$

It is then clear that θ would behave as a uniform interpolant of ϕ with respect to $\mathcal{L}(\phi) \setminus \{p\}$. However, it is easy to see that we are asking too much. Consider for example Basic Modal Logic: we will show that this logic enjoys uniform interpolation. However, it is certainly not closed for classical propositional quantifiers: this is easily seen, because for example $\exists p\langle a\rangle p$ is not invariant under $\{a\}$-bisimulation and hence it cannot be equivalent to a modal formula in $\{a\}$.

However, we just have to tune the notion of propositional quantifiers, to obtain a more sensible definition:

3.1.1. Definition. A logic that can speak about structures is said to have *bisimulation quantifiers* if for any formula ϕ of the logic and any subset $\mathcal{L}'$ of $\mathcal{L}(\phi)$ there is a formula θ of the logic such that

1. $\mathcal{L}(\theta) = \mathcal{L}'$;

2. for any structure $\mathcal{M}$ of $\mathcal{L}'$ it holds: $\mathcal{M} \models \theta$ iff there exists a structure $\mathcal{N}$ of $\mathcal{L}(\phi)$ which is bisimilar to $\mathcal{M}$ w.r.t. $\mathcal{L}'$ and $\mathcal{N} \models \phi$. □

The formula θ of the preceding definition is denoted by $\widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi$, which becomes $\widetilde{\exists}p\phi$ or $\widetilde{\exists}a\phi$ if $\mathcal{L}(\phi) \setminus \mathcal{L}'$ is equal to $\{p\}$, for a propositional constant p, or to $\{a\}$, for an action a, respectively.

3.1.2. Proposition. *Any logic that is invariant under bisimulation and it is closed for bisimulation quantifiers enjoys uniform interpolation.*

Proof.
Using the amalgamation property of structures (see Lemma 2.2.10), we show that the uniform interpolant of ϕ w.r.t. $\mathcal{L}' \subseteq \mathcal{L}$ is exactly the formula $\widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi$. From the semantics of $\widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi$, it is clear that $\phi(p) \models \widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi$. We prove that if $\phi \models \psi$ with $\mathcal{L}(\psi) \cap \mathcal{L}(\psi) = \mathcal{L}'$, then $\widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi \models \psi$. If not there is a structure $\mathcal{M}$ such that $\mathcal{M} \models \widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi \wedge \neg\psi$. Then the semantics of $\widetilde{\exists}(\mathcal{L} \setminus \mathcal{L}')\phi$ implies the existence of a structure $\mathcal{N}$ which is $\mathcal{L}'$-bisimilar to $\mathcal{M}$ such that $\mathcal{N} \models \phi$. But then there exists $\mathcal{K}$ which is bisimilar to both $\mathcal{M}$ and $\mathcal{N}$, in the appropriate languages, hence $\mathcal{K} \models \phi \wedge \neg\psi$, contradicting $\phi \models \psi$. □

Proposition 3.1.2 suggests a way to prove uniform interpolation for logics which are invariant under bisimulation: it suffices to prove that the logic is closed under bisimulation quantifiers.

Let us prove, for example, that Basic Modal Logic is closed under bisimulation quantifiers. We just sketch a proof, that can be found in [66]. The proof uses essentially two ingredients:

- for any structure $\mathcal{M}$ and natural number n, there exists a modal formula $\theta_n^{\mathcal{M}}$ of modal depth equal to n that characterizes $\mathcal{M}$ modulo n-bounded bisimulation (see Definition 2.2.9);

- the amalgamation Lemma 2.2.10 can be improved as follows.

  Let $\mathcal{M}$, $\mathcal{N}$ be structures for the languages $\mathcal{L}$ and $\mathcal{L}'$, respectively. If n is a natural number and $\mathcal{M}$ is n-bisimilar to $\mathcal{N}$, with respect to the language $\mathcal{L} \cap \mathcal{L}'$, then there is an $(\mathcal{L} \cup \mathcal{L}')$-structure $\mathcal{K}$ which is (fully!) bisimilar to $\mathcal{M}$ with respect to $\mathcal{L}$ and n- bisimilar to $\mathcal{N}$, with respect to $\mathcal{L}'$.

Suppose that n is the modal depth of ϕ. Using the two observations above, it is not difficult to see that the formula

$$\theta = \bigvee \{\theta_n^{\mathcal{M}} : \mathcal{M} \models \phi\},$$

which is a modal formula because the disjunction is finite, satisfies Definition 3.1.1.

Let us consider now a kind of converse of Proposition 3.1.2. A logic is said to have *bisimulation descriptions* for a class $\mathcal{C}$ of structures, if for all $\mathcal{M} \in \mathcal{C}$ there exists a formula $\tau_{\mathcal{M}}$ in the logic such that $\mathcal{N} \models \tau_{\mathcal{M}}$ iff $\mathcal{N} \sim \mathcal{M}$, for all structure $\mathcal{N}$.

For example, Infinitary Modal Logic can be shown to have bisimulation descriptions for the full class of structures (see [7]), while Propositional Dynamic Logic has bisimulation descriptions for the class of finite models (see [13]).

We have:

3.1.3. PROPOSITION. *Suppose that a logic is invariant for bisimulation and has bisimulation descriptions for a class $\mathcal{C}$ of structures. Then, if θ is a uniform interpolant of a formula ϕ w.r.t. $\mathcal{L}' \subseteq \mathcal{L}(\phi)$ and $\mathcal{M} \in \mathcal{C}$, then*
$\mathcal{M} \models \theta$ iff there exists $\mathcal{N}$ which is $\mathcal{L}'-$bisimilar to $\mathcal{M}$ and $\mathcal{N} \models \phi$.

In other words, θ behaves like a bisimulation quantifier w.r.t. $\mathcal{L}'$ on all structures of $\mathcal{C}$. Moreover, if any satisfiable formula of the logic has a model in $\mathcal{C}$, then the above model $\mathcal{N}$ can be taken in the class $\mathcal{C}$.

Proof.
Suppose the above hypotheses hold and consider a structure $\mathcal{M}$ in $\mathcal{C}$.

If there exists a structure $\mathcal{N}$ with $\mathcal{N} \models \phi$ which is $\mathcal{L}'-$bisimilar to $\mathcal{M}$, then $\mathcal{N} \models \theta$ because $\phi \models \theta$ and since $\mathcal{M}$ is $\mathcal{L}'$-bisimilar to $\mathcal{N}$, it follows that $\mathcal{M} \models \theta$.

Vice versa, suppose $\mathcal{M} \models \theta$. and let τ_M be the $\mathcal{L}'$-formula that characterizes the model $\mathcal{M}$ up to bisimulation, that is, for any $\mathcal{L}'$-model $\mathcal{M}'$:

$$\mathcal{M}' \models \tau_{\mathcal{M}} \text{ iff } \mathcal{M}' \text{ is } \mathcal{L}'\text{-bisimilar to } \mathcal{M}.$$

Consider the set of $\mathcal{L}(\phi)$-formulae $\{\phi, \tau_{\mathcal{M}}\}$. This set must have a model, otherwise we would have that $\phi \to \neg\tau_{\mathcal{M}}$ is valid. But since θ is the uniform interpolant of ϕ, then $\theta \to \neg\tau_{\mathcal{M}}$ should be valid as well, which cannot be, as $\mathcal{M} \models \theta \wedge \tau_{\mathcal{M}}$. Hence $\{\phi, \tau_{\mathcal{M}}\}$ is satisfiable and by hypothesis it has a model. Any model $\mathcal{N}$ of $\phi \wedge \tau_{\mathcal{M}}$ is $\mathcal{L}'$-bisimilar to $\mathcal{M}$ and $\mathcal{N} \models \phi$, so we are done.

If any satisfiable formula of the logic has a model in $\mathcal{C}$, then $\phi \wedge \tau_{\mathcal{M}}$ is satisfiable in $\mathcal{C}$ and the above model $\mathcal{N}$ can be taken in this class. □

Proposition 3.1.3 give us a very incomplete answer to the following interesting question: is it true that on logics that are bisimulation invariant the closure under bisimulation quantifiers implies uniform interpolation? Although it is not really a complete answer, Proposition 3.1.3 will be used in Section 3.4 to disprove uniform interpolation for both Infinitary Modal Logic and Propositional Dynamic Logic.

3.2 Automata.

To prove uniform interpolation for the Modal μ-Calculus, we show that this logic is closed under bisimulation quantifiers (see Proposition 3.1.2). By looking at the proof of the closure for Basic Modal Logic that we gave in the preceding section, we see that a basic ingredient in that proof, which is totally missed in the Modal μ-Calculus, is the notion of modal depth. But a new technique is introduced in Janin Walukiewicz paper [40]: they define a class of automata suitable for representing the μ-formulae and one can prove that the Modal μ-Calculus is closed under bisimulation quantifiers by operating on these automata instead that on the formulae.

In the following section we give a brief introduction to these automata, and provide a few examples of their behavior.

In the following we define automata that can *read* arbitrary structures and an automaton will accept or reject a structure according to some acceptance conditions. We first give the definition of an automaton and then we will try to justify the acceptance conditions by keeping in mind that we want automata to correspond exactly to modal μ-sentences: for any μ-sentence ϕ there must be an automaton $\mathcal{A}$ accepting exactly the class of structures in which ϕ holds and vice versa, for any automaton $\mathcal{A}$ such a μ-sentence ϕ must exist. One obvious consequence of this fact is that our acceptance conditions must be *invariant under bisimulation*: if $\mathcal{A}$ accepts $\mathcal{M}$ and $\mathcal{N}$ is bisimilar to $\mathcal{M}$ then $\mathcal{N}$ must be accepted by $\mathcal{A}$ as well.

3.2.1. Definition. A *μ-automaton* $\mathcal{A}$ is a tuple $(Q, \Sigma, \Lambda, q_0, \delta, \Omega)$ such that:

1. Q is a finite set of states;
2. Σ is a finite subset of PROP;
3. Λ is a finite subset of A;
4. $q_0 \in Q$ is the initial state;
5. $\delta : Q \times Pow(\Sigma) \to Pow(Pow(\Lambda \times Q))$ is the transition function;

6. $\Omega : Q \to \mathbb{N}$ is the *parity function.* □

We will also refer to such an automaton as a (Σ, Λ)-automaton, to stress its language.

Roughly speaking, the transition function δ of an automaton $\mathcal{A}$ gives a set of rules for *labeling* with Q-elements the set of successors of a node s of an $\mathcal{L}$-structure $\mathcal{M}$, which has been previously labeled by q. To label such a set, choose a D in $\delta(q, \{p \in \Sigma \mid s \models p\})$. D gives necessary conditions that the labeling of $R_a(s)$ must fulfill: for any action $a \in \Lambda$, the set of states appearing as labels of $a-$successors of s must be exactly the set $\{(a, q) \mid (a, q) \in D\}$.

3.2.2. Definition. Given an automaton $\mathcal{A} = (Q, \Sigma, \Lambda, q_0, \delta, \Omega)$ and a (Σ, Λ)-structure $\mathcal{M}$, a total function $l : \mathcal{M} \to Q$ is called an $\mathcal{A}$-*labeling* of $\mathcal{M}$ if the following conditions are satisfied:

1. $l(r^{\mathcal{M}}) = q_0$;
2. if $l(s) = q$ then:
$$D(s) := \{(a, q') \in \Lambda \times Q \mid \exists t \in R_a(s).\, l(t) = q'\} \in \delta(q, L_{\mathcal{A}}(s)),$$
where $L_{\mathcal{A}}(s) = \{p \in \Sigma \mid s \models p\}$;
3. for any infinite path $r^{\mathcal{M}} = s_0 \stackrel{a_1}{\to} s_1 \stackrel{a_2}{\to} \dots$ in $\mathcal{M}$:
$$min\{\Omega(q) \mid \text{there are infinitely many } s_i \text{ in the path s.t. } l(s_i) = q\}$$
is even. □

Once we have the notion of an $\mathcal{A}$-labeling at hand, we can define the notion of acceptance of a structure $\mathcal{M}$ by an automaton $\mathcal{A}$. Unfortunately, to say that $\mathcal{M}$ is accepted by $\mathcal{A}$ iff $\mathcal{M}$ has an $\mathcal{A}$-labeling does not work, because it is not even invariant under bisimulation. To see this, consider the automaton $\mathcal{A} = (\{q_0, q_1\}, \emptyset, \{a\}, q_0, \delta, \Omega)$ where $\delta(q_0, \emptyset) = \delta(q_1, \emptyset) = \{(a, q_0), (a, q_1)\}$, and the parity function is defined as $\Omega(q_0) = 0$, $\Omega(q_1) = 0$. Let $\mathcal{M}$ be the infinite binary tree and $\mathcal{N}$ a single reflexive node. $\mathcal{M}$ and $\mathcal{N}$ are bisimilar with respect to the language $\mathcal{L} = \{a\}$, but, while the first has an $\mathcal{A}$-labeling, the second has none. In fact, an $\mathcal{A}$-labeling of a structure, via the transition function δ, imposes to a node labeled by q_0 or q_1 to have at least two successors, one labeled by q_0 and the other by q_1: hence $\mathcal{N}$ cannot have an $\mathcal{A}$ labeling.

Thus, having an $\mathcal{A}$-labeling cannot be the right request of acceptance if automata have to be equivalent to (i.e. accept the same structure as) μ-sentences. However, this problem is overcome by considering labelings over expanded structures (see Definition 2.2.5). Suppose that the automaton $\mathcal{A}$ has n states. It can be proved that if $\mathcal{M}$ is bisimilar to $\mathcal{N}$ and $\mathcal{M}^n$ has an $\mathcal{A}$-labeling, then this is true for $\mathcal{N}^n$ as well.

3.2.3. DEFINITION. A (Σ, Λ)-structure $\mathcal{M}$ is *accepted* by $\mathcal{A} = (Q, \Sigma, \Lambda, q_0, \delta, \Omega)$ if the expansion $\mathcal{M}^n$, where $n = |Q|$, has an $\mathcal{A}$-labeling. □

Janin and Walukiewicz ([40]) prove that μ-automata correspond precisely to μ-sentences:

3.2.4. THEOREM. *For every μ-automaton $\mathcal{A}$ there is a μ-sentence $\phi_{\mathcal{A}}$ such that $\mathcal{A}$ accepts $\mathcal{M}$ iff $\mathcal{M} \models \phi_{\mathcal{A}}$. Moreover, if $\mathcal{A}$ is a (Σ, Λ)-automaton then the corresponding sentence is also in this language.*

Vice versa, for every μ-sentence ϕ there is an automaton $\mathcal{A}_\phi$ such that $\mathcal{A}_\phi$ accepts $\mathcal{M}$ iff $\mathcal{M} \models \phi$. If the set of proposition constants in ϕ is Σ and its set of atomic actions is Λ, then the corresponding automaton may be assumed to be a (Σ, Λ)-automaton.

The proof of this theorem is too long to be presented here and the interested reader can consult [40]. We only point out the relation between the Ω-condition on $\mathcal{A}$-labelings and the nesting of fixed point operators. In a very imprecise way, an $\mathcal{A}$-labeling on a structure $\mathcal{M}$ *represents* the evaluation of a sentence in that structure. States q with an odd value of $\Omega(q)$ will represent the regeneration of a μ-subformula, while states with an even value of $\Omega(q)$ will represent the regeneration of a ν-subformula. Moreover, if a fixed point formula ψ is a subformula of a fixed point subformula θ, and q_1, q_2 represent the regeneration of ψ, θ, respectively, then $\Omega(q_1) > \Omega(q_2)$. In this way we see how the Ω-condition of a labeling corresponds to the fact that a μ-subformula can be regenerated infinitely often *only* if it is inside the scope of a ν-one.

The definition of acceptance which uses labelings on expanded structures differs from the one given in [40] where, to decide if a structure $\mathcal{M}$ is accepted by $\mathcal{A}$, games played directly on $\mathcal{M}$ are used. The equivalence of the two notions of acceptance can be derived using Lemma 3.4.3 and Lemma 3.4.4 of [38].

Examples.

- The first example is a μ-automaton that corresponds to the PDL (PDL) formula $\langle a^* \rangle p$ (in Modal μ-Calculus notation: $\mu X.(p \vee \langle a \rangle X)$. Define $\mathcal{A} = (\{q_0, q_1\}, \{p\}, \{a\}, q_0, \delta, \Omega)$ where:

  $$\delta(q, L) := \begin{cases} \{\ \{(a, q_0), (a, q_1)\}\ \} & \text{if } q = q_0 \text{ and } L = \emptyset; \\ \{\ \emptyset\ ,\ \{(a, q_1)\}\ \} & \text{else.} \end{cases}$$

  and the parity function is defined as $\Omega(q_0) = 1$, $\Omega(q_1) = 0$.

  First we prove that any structure $\mathcal{M}$ accepted by $\mathcal{A}$ makes the formula $\langle a^* \rangle p$ true. Suppose there is a labeling l from $\mathcal{M}^2$ to Q that satisfies Definition 3.2.2 and suppose by contradiction that $\mathcal{M}$ (hence $\mathcal{M}^2$) makes $\langle a^* \rangle p$ false.

Then in any node s of $\mathcal{M}^2$ the proposition constant p is false. Using this, we show how to obtain an infinite sequence $r^{\mathcal{M}_n} = s_0 \xrightarrow{a} s_1 \xrightarrow{a} \ldots$ in $\mathcal{M}^2$ such that $l(s_i) = q_0$, in contradiction with the third condition in Definition 3.2.2. Suppose we already have a sequence $r^{\mathcal{M}^2} = s_0 \xrightarrow{a} s_1 \xrightarrow{a} \ldots s_m$ with $l(s_i) = q_0$ for $i = 1, \ldots, m$. We show how we can add a node s_{m+1} with $l(s_{m+1}) = q_0$ to this sequence. By Definition 3.2.3 we have that

$$D(s_m) := \{(a, q') \in \Lambda \times Q \mid \exists t \in R_a(s_m).\, l(t) = q'\} \in \delta(q_0, L_{\mathcal{A}}(s_m)) =$$

$\delta(q_0, \emptyset) = \{\{(a, q_0), (a, q_1)\}\}$. Hence $D(s_m) = \{(a, q_0), (a, q_1)\}$ and there exists a $s_{m+1} \in R_a(s_m)$ with $l(s_{m+1}) = q_0$.

Vice versa, suppose that $\mathcal{M} \models \langle a^* \rangle p$. Then $\mathcal{M}^2 \models \langle a^* \rangle p$ too and there is a sequence $r^{\mathcal{M}^2} = s_0 \xrightarrow{a} s_1 \xrightarrow{a} \ldots s_m$ such that $s_i \models \neg p$ for $i < m$, and $s_m \models p$. Since $\mathcal{M}^2$ is a tree in which every node of $\mathcal{M}$ is *copied* twice, any node in $\mathcal{M}^2$, except the root, has only one immediate predecessor and any node that has at least one successor must have at least two. Define l on $\mathcal{M}^2$ as follows:

$$l(s) := \begin{cases} q_0 & \text{if } s = s_0, \ldots, s_m; \\ q_1 & \text{else.} \end{cases}$$

We prove that l is an $\mathcal{A}$-labeling, that is, it verifies the conditions of Definition 3.2.3. It is clear that any infinite path will eventually be labeled only by q_1, hence the Ω-condition of a labeling is trivially satisfied.

As for the second condition of a labeling, suppose first that s belongs to $\{s_0, \ldots, s_{m-1}\}$. Then $R_a(s) \cap \{s_0, \ldots, s_m\}| = 1$ and $R_a(s)| \geq 2$, because $\mathcal{M}^2$ is two-expanded. Hence $D(s) = \{(a, q_0), (a, q_1)\} \in \delta(q_0, \emptyset) = \delta(l(s), L_{\mathcal{A}}(s))$. If $s = s_m$, then either s has no successors at all and $D(s) = \emptyset \in \delta(q_0, \{p\}) = \delta(l(s), L_{\mathcal{A}}(s))$, or any $s' \in R_a(s)$ is labeled by q_1. In this case $D(s) = \{(a, q_1)\} \in \delta(q_0, \{p\}) = \delta(l(s), L_{\mathcal{A}}(s))$. Finally, if $s \neq s_0, \ldots, s_m$, then $R_a(s) \cap \{s_0, \ldots, s_m\} = \emptyset$ and $D(s)$ is either empty or equal to $\{(a, q_1)\}$. In both cases $D(s)$ belongs to $\delta(q_1, L_{\mathcal{A}}(s)) = \delta(l(s), L_{\mathcal{A}}(s))$. □

Before presenting the second example we just point out the following fact: to prove that $\mathcal{M} \models \phi$ in the previous example, one could have used labeling directly on $\mathcal{M}$ instead of using its expansion $\mathcal{M}^2$. This holds in general: if ϕ is a sentence with corresponding automaton $\mathcal{A}$, then it can be proved that any structure with an $\mathcal{A}$-labeling makes ϕ true. The expansion is needed in the converse direction, because it is not true that $\mathcal{M} \models \phi$ implies the existence of an $\mathcal{A}$-labeling on $\mathcal{M}$. In our example, a single reflexive node w that makes p true cannot have an $\mathcal{A}$-labeling, because the root w should be labeled by q_0, while the second condition of a labeling would force w to be labeled also by q_1.

- Our second example is an automaton corresponding to $\nu X.\langle a\rangle X$, which is true at all nodes from which an infinite a-path emerges. Define $\mathcal{A}$ as $(\{q_0, q_1\}, \emptyset, \{a\}, q_0, \delta, \Omega)$ where:

$$\begin{array}{lcl} \delta(q_0, \emptyset) & := & \{\, \{(a, q_0), (a, q_1)\}\, \} \\ \delta(q_1, \emptyset) & := & \{\, \emptyset\, ,\ \{(a, q_1)\}\, \} \end{array}$$

  and $\Omega(q_0) = \Omega(q_1) = 0$.

  Suppose that a structure $\mathcal{M}$ has an $\mathcal{A}$-labeling. If $l(s) = q_0$ then s has an a-successor t with $l(t) = q_0$, because $D(s)$ must be equal to $\{(a, q_0), (a, q_1)\}$. This enables us to construct an infinite a-path starting from the root $r^{\mathcal{M}}$.

  Vice versa, suppose that $\mathcal{M} \models \nu X.\langle a\rangle X$. Then we want to show that $\mathcal{M}^2$ has an $\mathcal{A}$-labeling. Let $r^{\mathcal{M}^2} = s_0 \stackrel{a}{\to} s_1 \stackrel{a}{\to} \ldots$ be an infinite path starting from the root in the structure $\mathcal{M}^2$. We define the labeling l as follows:

$$l(s) := \begin{cases} q_0 & \text{if } s = s_i, \text{ for } i \in \mathbb{N}\ ; \\ q_1 & \text{else.} \end{cases}$$

  We show that l verifies the conditions of Definition 3.2.3. Since $\Omega(q_0) = \Omega(q_1) = 0$, it is clear that the Ω-condition of a labeling is satisfied.

  As for the second condition of a labeling, suppose first that s belongs to $\{s_0, \ldots, s_i, \ldots\}$. Since $\mathcal{M}^2$ is two-expanded we have that the sets $R_a(s) \cap \{s_0, \ldots, s_i, \ldots\}$, $R_a(s) \cap (\mathcal{M} \setminus \{s_0, \ldots, s_i, \ldots\})$ are both non-empty (and $D(s) = \{(a, q_0), (a, q_1)\} \in \delta(l(s), L_{\mathcal{A}}(s))$.

  If $s \notin \{s_0, \ldots, s_i, \ldots\}$, then $l(s) = q_1$, $R_a(s) \cap \{s_0, \ldots, s_i, \ldots\} = \emptyset$, and $D(s)$ is either empty or equal to $\{(a, q_1)\}$. In both cases it belongs to $\delta(l(s), L_{\mathcal{A}}(s))$.

- Our third example is an automaton corresponding to $\nu X.\mu Y.(p \wedge \langle a\rangle X) \vee \langle a\rangle Y$, which is true in all points w from which an infinite path, where p appears infinitely often, starts. The corresponding automaton is given by $\mathcal{A} = (\{q_0, q_1, q_2\}, \{p\}, \{a\}, q_0, \delta, \Omega)$ where:

$$\delta(q, L) := \begin{cases} \{\, \{(a, q_0), (a, q_1)\}\, \} & \text{if } q = q_0 \text{ or } q = q_2 \text{ and } L = \emptyset; \\ \{(a, q_1), (a, q_2)\}\, \} & \text{if } q = q_0 \text{ or } q = q_2 \text{ and } L = \{p\}; \\ \{\, \emptyset\, ,\ \{(a, q_1)\}\, \} & \text{if } q = q_1; \end{cases}$$

  and the parity function is defined as $\Omega(q_0) = 1$, $\Omega(q_1) = 0$, $\Omega(q_2) = 0$.

  First we suppose that there is a labeling l from $\mathcal{M}$ to Q that satisfies Definition 3.2.2. We claim that from any node w labeled by q_0 or q_2 there is a finite a-path $w = w_0 \stackrel{a}{\to} w_1 \stackrel{a}{\to} \ldots \stackrel{a}{\to} w_n$ with $n \geq 1$, that there exists $i \geq n$ such that $w_i \models p$, and $l(w_n) = q_0$, or $l(w_n) = q_2$. If this claim is true, one can repeat the argument, obtaining an infinite path starting from w, where p appears infinitely often. By applying this to $r^{\mathcal{M}}$ we obtain $r^{\mathcal{M}} \models \nu X.\mu Y.(p \wedge \langle a \rangle X) \vee \langle a \rangle Y$.

The proof of the claim goes as follows. Suppose w is a node in $\mathcal{M}$ with $l(w) = q_0$ or $l(w) = q_2$. We consider two cases.

1. If $w \models p$ then $\delta(l(w), L_{\mathcal{A}}(w)) = \{\{(a, q_1), (a, q_2)\}\}$ and since l is an $\mathcal{A}$-labeling, we must have an a-successor w_1 of w labeled by q_2. Then the path w, w_1 satisfies the claim.
2. If $w \not\models p$, then $(a, q_0) \in D(w)$ and we must have an a-successor of w still labeled by q_0. If this successor does not satisfy p, we can repeat the same argument, but after a finite number of steps a node w_n labeled by q_0 and satisfying p must be reached: otherwise we would have an infinite path labeled by q_0 and since $\Omega(q_0) = 1$ this is in contradiction with the fact that l is an $\mathcal{A}$-labeling. Then the path $w = w_0 \stackrel{a}{\to} w_1 \stackrel{a}{\to} \ldots \stackrel{a}{\to} w_n$ satisfies the claim.

Vice versa, suppose that $\mathcal{M} \models \nu X.\mu Y.(p \wedge \langle a \rangle X) \vee \langle a \rangle Y$. Then $\mathcal{M}^3 \models \nu X.\mu Y.(p \wedge \langle a \rangle X) \vee \langle a \rangle Y$ too and there is an infinite sequence $\sigma = w_0 \stackrel{a}{\to} w_1 \stackrel{a}{\to} \ldots$ with $r^{\mathcal{M}^3} = w_0$ in which p appears infinitely often, say at points $w_{i_1}, w_{i_2}, \ldots$.

Define l on $\mathcal{M}^3$ as follows:

$$l(w) := \begin{cases} q_0 & \text{if } w \in \{w_0, w_1, \ldots\} \setminus \{w_{i_1+1}, w_{i_2+1}, \ldots\}; \\ q_2 & \text{if } w \in \{w_{i_1+1}, w_{i_2+1}, \ldots\}; \\ q_1 & \text{else.} \end{cases}$$

We now check that l is an $\mathcal{A}$-labeling. It is clear that any infinite path different from σ will eventually be labeled only by q_1 and since $\Omega(q_1) = 0$, these paths do not cause any problem. On the other hand,

$$min\{\Omega(q) \mid \text{there are infinitely many } w_i \text{ in } \sigma \text{ s.t. } l(w_i) = q\} = 0,$$

because q_2 appears infinitely often as a label in σ and $\Omega(q_2) = 0$. Hence the Ω-condition of a labeling is satisfied by l.

As for the second condition of a labeling, we consider two cases.

1) Suppose first that $w \notin \{w_0, w_1, \ldots\}$. In this case $l(w) = q_1$, $R_a(w) \cap \{w_0, \ldots\} = \emptyset$ and either w does not have any successor, or all its successors are labeled by q_1. In both cases $D(w) \in \delta(l(w), L(w)) = \{\emptyset, \{(a, q_1)\}\}$.

2)If $w \in \{w_0, w_1, \ldots\}$, then $l(w) = q_0$ or $l(w) = q_2$. If $w \models p$, then there exists j such that $w = w_{i_j}$ and its a-successor w_{i_j+1} is labeled by q_2. Moreover, w has a successor different from w_{i_j+1} (because $\mathcal{M}^3$ is three expanded) and all successors different from w_{i_j+1} do not belong to the path σ and are hence labeled by q_1. This proves that $D(w) \in \delta(l(w), L_{\mathcal{A}}(w)) = \{(a, q_1), (a, q_2)\}$.

If $w \not\models p$, then the a-successor w' of w in the path σ is labeled by q_0 and all successors different from w' do not belong to the path σ and are hence labeled by q_1. This proves that $D(w) \in \delta(l(w), L_{\mathcal{A}}(w)) = \{(a, q_0), (a, q_1)\}$. □

3.3 Uniform interpolation

In this section we prove the Uniform Interpolation Theorem for the Modal μ-Calculus. This is done via an interpretation of second-order quantifiers within the Modal μ-Calculus itself. This approach is inspired by [66] and [56]. The results of this section were proved in collaboration with Marco Hollenberg in [21].

We start by proving that the Modal μ-Calculus allows existential quantification on proposition constants, modulo bisimulation.

3.3.1. THEOREM. *Let ϕ be a μ-sentence and $\mathcal{L}$ its language. Let p be a proposition constant. Then there is a μ-sentence θ in the language $\mathcal{L} \setminus \{p\}$ such that:*

$$\mathcal{M} \models \theta \text{ iff there is a structure } \mathcal{N} \text{ with } \mathcal{M} \sim_{(\mathcal{L}\setminus\{p\})} \mathcal{N} \text{ and } \mathcal{N} \models \phi.$$

Proof.
Clearly, if $p \notin \Sigma$ then ϕ itself does not contain p, so we may simply choose $\theta := \phi$. What remains is the case that $p \in \Sigma$.

Let $\mathcal{A} = (Q, \Sigma, \Lambda, q_0, \delta, \Omega)$ be an automaton for ϕ, with $\mathcal{L} = \Sigma \cup \Lambda$. Starting from $\mathcal{A}$ we construct an automaton $\mathcal{B}$ in the language $(\Sigma \setminus \{p\}, \Lambda)$ and prove that the sentence corresponding to $\mathcal{B}$ behave as required. We first give an informal description of $\mathcal{B}$. This automaton is the same as $\mathcal{A}$ *except* for the transition function δ. When labeling a structure by following the instruction of $\mathcal{B}$ we want to be able, for each node w of the structure, to *choose* a labeling of the set of a-successors of w either according to an *instructions* in $\delta(q, L)$ or to one in $\delta(q, L \cup \{p\})$. Then, if $\mathcal{M}$ is accepted by $\mathcal{B}$, we can define an interpretation of p in $\mathcal{M}^{|Q|}$ as follows. Consider a $\mathcal{B}$-labeling l of $\mathcal{M}^{|Q|}$ and define $s \models p$ iff l labels the set of the a-successors of s with an instruction of $\delta(q, L \cup \{p\})$. In this p-extension of $\mathcal{M}^{|Q|}$ we can easily prove that the labeling l is also an $\mathcal{A}$-labeling. In this way we will see that if $\phi_{\mathcal{B}}$ is true in $\mathcal{M}$, then $\mathcal{M}$ is bisimilar to the structure $(\mathcal{M}^{|Q|}, p)$ in which ϕ holds.

More formally, let $\mathcal{B}$ be the automaton $(Q, \Sigma \setminus \{p\}, \Lambda, q_0, \delta', \Omega)$ with:

$$\delta'(q, L) := \delta(q, L) \cup \delta(q, L \cup \{p\})$$

This is well-defined precisely because p is assumed to be an element of Σ. $\mathcal{B}$ corresponds to a μ-sentence $\phi_{\mathcal{B}}$ in the language $\mathcal{L} \setminus \{p\}$. We will prove that this formula satisfies the requirements, that is,

$$\mathcal{M} \models \phi_{\mathcal{B}} \text{ iff there is a structure } \mathcal{N} \text{ with } \mathcal{M} \sim_{(\mathcal{L}\setminus\{p\})} \mathcal{N} \text{ and } \mathcal{N} \models \phi.$$

.

- We first claim that any structure accepted by $\mathcal{A}$ is also accepted by $\mathcal{B}$. Once we have proved this, it will be clear that if there is a structure $\mathcal{N}$ with $\mathcal{M} \sim_{(\mathcal{L}\setminus\{p\})} \mathcal{N}$ and $\mathcal{N} \models \phi$, then $\mathcal{M} \models \phi_{\mathcal{B}}$: if $\mathcal{N} \models \phi$ then $\mathcal{N}$ is accepted by $\mathcal{A}$, hence by the claim is accepted by $\mathcal{B}$; but then $\mathcal{N} \models \phi_{\mathcal{B}}$ and $\mathcal{M}$ will satisfy it as well, since it is $(\mathcal{L} \setminus \{p\})$-bisimilar to $\mathcal{N}$.

  To prove the claim, notice that if we show that an $\mathcal{A}$-labeling on any structure is also a $\mathcal{B}$-labeling, then we are done: if $\mathcal{M}$ is accepted by $\mathcal{A}$ then $\mathcal{M}^{|Q|}$ has an $\mathcal{A}$-labeling, hence a $\mathcal{B}$ labeling and $\mathcal{B}$ accepts $\mathcal{M}$.

  Suppose that $\mathcal{M}$ is a structure and l an $\mathcal{A}$-labeling on $\mathcal{M}$. We prove that l is a $\mathcal{B}$-labeling. Since the Ω-function of $\mathcal{B}$ is the same as that of $\mathcal{A}$, the Ω-condition of a labeling is trivially satisfied and we just prove the third condition.

  If $l(s) = q$ we have to show that

  $$D(s) := \{(a, q') \in \Lambda \times Q \mid \exists t \in R_a(s).\, l(t) = q'\} \in \delta(q, L_{\mathcal{B}}(s)).$$

  To see this, notice that if $\mathcal{M}, s \models p$ then $\delta(q, L_{\mathcal{A}}(s)) = \delta(q, L_{\mathcal{B}}(s) \cup \{p\})$, otherwise $\delta(q, L_{\mathcal{A}}(s)) = \delta(q, L_{\mathcal{B}}(s))$. Since l is an $\mathcal{A}$-labeling, we know that $D(s) \in \delta(q, L_{\mathcal{A}}(s))$, and in both cases $D(s) \in \delta'(q, L_{\mathcal{B}}(s)) = \delta(q, L_{\mathcal{B}}(s)) \cup \delta(q, L_{\mathcal{B}}(s) \cup \{p\})$.

- Second, we claim that if $\mathcal{B}$ accepts a structure $\mathcal{M}$, then it is $(\mathcal{L} \setminus \{p\})$-bisimilar with some structures accepted by $\mathcal{A}$. Then we easily obtain that if $\mathcal{M} \models \phi_{\mathcal{B}}$, then there is a structure $\mathcal{N}$ with $\mathcal{M} \sim_{(\mathcal{L}\setminus\{p\})} \mathcal{N}$ and $\mathcal{N} \models \phi$.

  By hypothesis the structure $\mathcal{M}^n$, where $n = |Q|$, has a $\mathcal{B}$-labeling l. We define a subset P of the domain of $\mathcal{M}^n$ by: $s \in P$ iff

  $$D(s) := \{(a, q) \in \Lambda \times Q \mid \exists t \in R_a(s).\, l(t) = q\} \in \delta(l(s), L_{\mathcal{B}}(s) \cup \{p\})$$

  If $(\mathcal{M}^n, P)$ is the structure that differs from $\mathcal{M}^n$ only because the interpretation of p is P, then $(\mathcal{M}^n, P)$ is $(\mathcal{L} \setminus \{p\})$-bisimilar to $\mathcal{M}$ and l is an $\mathcal{A}$-labeling for $(\mathcal{M}^n, P)$. To see this, consider a point s in $\mathcal{M}^n$. By definition of labelings, $D(s)$ must be in $\delta'(l(s), L_{\mathcal{B}}(s))$. There are two cases to consider:

1. $D(s) \in \delta(l(s), L_{\mathcal{B}}(s) \cup \{p\})$. Then $s \in P$, $\delta(l(s), L_{\mathcal{A}}(s))$ is equal to $\delta(l(s), L_{\mathcal{B}}(s) \cup \{p\})$ and $D(s) \in \delta(l(s), L_{\mathcal{A}}(s))$.
2. $D(s) \notin \delta(l(s), L_{\mathcal{B}}(s) \cup \{p\})$. Then $D(s) \in \delta(l(s), L_{\mathcal{B}}(s))$, and $s \notin P$ so: $D(s) \in \delta(l(s), L_{\mathcal{A}}(s)) = \delta(l(s), L_{\mathcal{B}}(s))$.

The other conditions of labelings are trivially satisfied, simply because l is a $\mathcal{B}$-labeling and the initial state and the Ω-function are unchanged. $(\mathcal{M}^n, P)$ is thus accepted by $\mathcal{A}$. □

We turn our attention to atomic actions.

3.3.2. THEOREM. *Let ϕ be a μ-sentence with $\mathcal{L}(\phi) = \mathcal{L}$. Let a be an atomic action. Then there is a μ-sentence θ with $\mathcal{L}(\theta) = \mathcal{L} \setminus \{a\}$ such that:*

$$\mathcal{M} \models \theta \text{ iff there is a structure } \mathcal{N} \text{ with } \mathcal{M} \sim_{(\mathcal{L}\setminus\{a\})} \mathcal{N} \text{ and } \mathcal{N} \models \phi.$$

Proof.
Let $\mathcal{A} = (Q, \Sigma, \Lambda, q_0, \delta, \Omega)$ be an automaton for ϕ, with $\mathcal{L} = \Sigma \cup \Lambda$. If ϕ is unsatisfiable we may take $\psi := \bot$; hence we may suppose that ϕ is satisfiable.

Define the automaton $\mathcal{B}$ as $(Q, \Sigma, \Lambda \setminus \{a\}, q_0, \delta', \Omega)$ where:

$$\delta'(q, L) := \{D \cap ((\Lambda \setminus \{a\}) \times Q) \mid D \in \delta(q, L)\}$$

We prove that any structure $\mathcal{M}$ accepted by $\mathcal{A}$ is also accepted by $\mathcal{B}$. As in the previous proof, we just show that an $\mathcal{A}$-labeling on any structure is also a $\mathcal{B}$-labeling. The Ω-condition is trivially satisfied, hence we just have to show that for any $s \in \mathcal{M}$ we have

$$D'(s) := \{(a, q') \in (\Lambda \setminus \{a\}) \times Q \mid \exists t \in R_a(s).\, l(t) = q'\} \in \delta'(l(s), L_{\mathcal{B}}(s)).$$

Since l is a $\mathcal{A}$ labeling, we know that

$$D(s) := \{(a, q') \in (\Lambda) \times Q \mid \exists t \in R_a(s).\, l(t) = q'\} \in \delta(l(s), L_{\mathcal{A}}(s)),$$

and $D'(s) = D(s) \cap ((\Lambda \setminus \{a\}) \times Q)$. Since $L_{\mathcal{B}}(s) = L_{\mathcal{A}}(s)$, we obtain that $D'(s)$ belongs to $\delta'(l(s), L_{\mathcal{B}}(s))$.

Now for the converse. Suppose that $\mathcal{B}$ accepts a structure $\mathcal{M}$: we prove that it is $(\mathcal{L} \setminus \{a\})$-bisimilar to some structure accepted by $\mathcal{A}$.

First notice that since ϕ is satisfiable we may suppose without loss of generality that its automaton $\mathcal{A}$ does not contain any *empty state*: a state q of an automaton $\mathcal{C}$ is said to be empty if the automaton $(\mathcal{C}, q)$, which is equal to $\mathcal{C}$ except that the initial state is q, does not accept any structure. If q is a state of the automaton $\mathcal{A}$ corresponding to ϕ, then consider the following two cases.

1. q is never used in labelings of models of ϕ. Then we can erase it from the states of $\mathcal{A}$ without changing the class of models accepted by the automaton.

2. There is a model $\mathcal{M}$ of ϕ and a labeling l of $\mathcal{M}$ with $l(s) = q$ for some $s \in \mathcal{M}$. Then the automaton $(\mathcal{A}, q)$ accepts $(\mathcal{M}, s)$ and the state q is non empty.

Hence we may suppose that any state of $\mathcal{A}$ is non empty and we may choose, for each $q \in Q$, a structure $\mathcal{M}_q$ with an $\mathcal{A}$-labeling l_q.

Suppose $\mathcal{M}$ is accepted by $\mathcal{B}$. This means that there exists a $\mathcal{B}$-labeling on $\mathcal{M}^n$, where $n = |Q|$. We define a mapping $d : \mathcal{M}^n \to Pow(\Lambda \times Q)$ as follows. By definition of labelings, for every s in the domain of $\mathcal{M}^n$ it holds:

$$D(s) := \{(b, q) \in (\Lambda \setminus \{a\}) \times Q \mid \exists t \in R_b(s).\, l(t) = q\} \in \delta'(l(s), L(s)).$$

So there must be some $D \in \delta(l(s), L_{\mathcal{B}}(s) = L_{\mathcal{A}}(s))$ with $D \cap ((\Lambda \setminus \{a\}) \times Q) = D(s)$. Choose such a D arbitrarily and fix $d(s) := D$.

Let $\mathcal{N}$ be the disjoint union of $\mathcal{M}^n$ and $\mathcal{M}_q$ for all $q \in Q$, with the root of $\mathcal{M}^n$ as its root and with some new a-transitions: we add an a-transition from s to the root of $\mathcal{M}_q$ for every $s \in \mathcal{M}^n$, $q \in Q$ such that $(a, q) \in d(s)$. Clearly the embedding of $\mathcal{M}^n$ into $\mathcal{N}$ is an $(\mathcal{L} \setminus \{a\})$-bisimulation and since $\mathcal{M}$ is in turn $(\mathcal{L} \setminus \{a\})$-bisimilar to $\mathcal{M}^n$ we obtain that $\mathcal{N}$ is $(\mathcal{L} \setminus \{a\})$-bisimilar to $\mathcal{M}$. Furthermore, it is easy to check that $l \cup \bigcup_{q \in Q} l_q$ is an $\mathcal{A}$-labeling of $\mathcal{N}$, which ends the proof of the theorem. □

3.3.3. Corollary. *For any μ-sentence ϕ and every proposition constant p there is a μ-sentence $\widetilde{\exists} p.\phi$ such that:*

1. $\phi \models \widetilde{\exists} p.\phi$;
2. $\mathcal{L}(\widetilde{\exists} p.\phi) = \mathcal{L}(\phi) \setminus \{p\}$;
3. *If* $\phi \models \psi$ *and* $p \notin \mathcal{L}(\psi)$ *then* $\widetilde{\exists} p.\phi \models \psi$.

For any μ-sentence ϕ and every atomic action $a \in A$ there is a formula $\widetilde{\exists} a.\phi$ such that:

1. $\phi \models \widetilde{\exists} a.\phi$;
2. $\mathcal{L}(\widetilde{\exists} a.\phi) = \mathcal{L}(\phi) \setminus \{a\}$;
3. *If* $\phi \models \psi$ *and* $a \notin \mathcal{L}(\psi)$ *then* $\widetilde{\exists} a.\phi \models \psi$.

Proof.
We claim that the sentence θ constructed from ϕ as in Theorem 3.3.1 enjoys the first three properties. It is clear that $\phi \models \theta$ and $\mathcal{L}(\theta) = \mathcal{L}(\phi) \setminus \{p\}$. Let ψ be such that $\phi \models \psi$ and $p \notin \mathcal{L}(\psi)$. We want to prove that for any $(\mathcal{L}(\theta) \cup \mathcal{L}(\psi))$-structure $\mathcal{M}$, if $\mathcal{M} \models \theta$ then $\mathcal{M} \models \psi$.

Let $\mathcal{L}(\theta) \cup \mathcal{L}(\psi) = \mathcal{L}$, $\mathcal{L}(\phi) = \mathcal{L}'$ and let $\mathcal{M}$ be an $\mathcal{L}$-structure such that $\mathcal{M} \models \theta$. If we consider $\mathcal{M}$ as an $\mathcal{L}(\theta)$-structure, then by the property of θ we know that there exists an $\mathcal{L}'$-structure $\mathcal{N}$ which is $\mathcal{L}(\theta)$-bisimilar to $\mathcal{M}$ and satisfies the formula ϕ. Since $\mathcal{L} \cap \mathcal{L}' = \mathcal{L}(\theta)$, we can apply Lemma 2.2.10 to get an $(\mathcal{L} \cup \mathcal{L}')$-structure $\mathcal{K}$ which is $\mathcal{L}$-bisimilar to $\mathcal{M}$ and $\mathcal{L}'$-bisimilar to $\mathcal{N}$. From $\mathcal{K} \sim_{\mathcal{L}'} \mathcal{N}$ it follows $\mathcal{K} \models \phi$, from $\phi \models \psi$ it follows $\mathcal{K} \models \psi$ and from $\mathcal{M} \sim_{\mathcal{L}} \mathcal{K}$ we obtain $\mathcal{M} \models \psi$.

The proof for $\widetilde{\exists} a.\phi$ is completely analogous. □

Notice that the conditions in the above corollary state that $\widetilde{\exists} p.\phi$, $\widetilde{\exists} a.\phi$ are the uniform interpolants of ϕ with respect to ϕ's language without p and without a, respectively. The reason for the existential notation is that the uniform interpolant behaves like a second-order quantification of ϕ, as the conditions demonstrate: the first and the third items correspond to $\exists$-introduction and $\exists$-elimination in natural deduction.

3.3.4. Corollary. *The Modal μ-Calculus enjoys uniform interpolation.*

Proof.
Let ϕ be a μ-sentence. Fix some $\mathcal{L}' \subseteq \mathcal{L}(\phi)$. Suppose

$$\mathcal{L}(\phi) \setminus \mathcal{L}' = \{p_1, \ldots, p_n, a_1, \ldots, a_m\}.$$

Then the uniform interpolant is: $\widetilde{\exists} p_1 \ldots \widetilde{\exists} p_n.\widetilde{\exists} a_1 \ldots \widetilde{\exists} a_m.\phi$. □

3.4 Negative results

In this section we show that Infinitary Modal Logic and PDL do not enjoy uniform interpolation. This was proved in collaboration with Marco Hollenberg in [21] and can be easily derived from Proposition 3.1.3. It is known (see [7]) that Infinitary Modal Logic is bisimulation complete for the class of all structures (in other words, we have, up to bisimulation, a *description* of any structure inside the logic), while PDL is bisimulation complete for the class of finite models (see [13]). Hence, if Infinitary Modal Logic or PDL enjoyed uniform interpolation they would be closed for bisimulation quantifiers, in the full class of models in the first case and in the finite class in the second one.

Consider now the formula $\phi := p \wedge [a^*](p \to \langle a \rangle p)$, which can be expressed in both logics. We show that for any structure $\mathcal{M}$ it holds: $\mathcal{M} \models \widetilde{\exists} p\phi$ iff there exists a sequence $(w_i)_{i \in \omega}$ with $w_0 = r^{\mathcal{M}}$ and $w_i R_a^{\mathcal{M}} w_{i+1}$ (in this case we say that R_a is non-well-founded on $\mathcal{M}$). If $\mathcal{M} \models \widetilde{\exists} p\phi$, consider a model $\mathcal{N} \sim_{\mathcal{L}(\phi)\setminus\{p\}} \mathcal{M}$ with $\mathcal{N} \models \phi$. Then in $\mathcal{N}$ there is an infinite a-chain in which p holds and since $\mathcal{M}$ is bisimilar to $\mathcal{N}$ w.r.t a language containing $\{a\}$, there is an infinite a-chain in $\mathcal{M}$ as well. Hence R_a is non-well-founded on $\mathcal{M}$. Vice versa, if R_a is non-well-founded, then we can find an infinite a-path $(w_i)_{i\in\omega}$ in $\mathcal{M}$. Let p be true in all w_i and false outside $(w_i)_{i\in\omega}$. Then $(\mathcal{M}, p) \models \phi$.

Hence if Infinitary Modal Logic enjoyed uniform interpolation, by Lemma 3.1.3 Infinitary Modal Logic would be closed for bisimulation quantifiers and we could express the non-well-foundedness of R_a. But this is not possible, even in Infinitary Logic:

3.4.1. LEMMA. *Non-well-foundedness is not expressible in Infinitary Logic.*

Proof.
For this we use the notion of bounded bisimulation from Definition 2.2.9. Suppose ψ is an infinitary logic sentence expressing non-well-foundedness. As non-well-foundedness is invariant for bisimulation, by a result of [11], ψ must be equivalent to an infinitary *modal* sentence. Therefore ψ must be invariant for α-bisimulation, for some ordinal α. Consider now two models, $\mathcal{M}$ and $\mathcal{N}$ defined as follows:

1. The domain of both is $\alpha + 1$, that is: $\{\beta \mid \beta \leq \alpha\}$.
2. In $\mathcal{M}$: $m \xrightarrow{a} n$ iff $m > n$.
3. $\mathcal{N}$ has one extra a-transition (besides those that are also present in $\mathcal{M}$), namely from α to α.
4. The root in both is α.

So $\mathcal{M}$ is well-founded (as $>$ is well-founded), while $\mathcal{N}$ is not, because it has a loop at its root. Nevertheless, $\mathcal{M}$ and $\mathcal{N}$ are α-bisimilar, via the sequence of relations $(Z_\beta)_{\beta\leq\alpha}$, with:

$$mZ_\beta n \quad \text{iff} \quad \beta = m \leq \alpha \text{ or } m = n.$$

□

Consider now PDL. As we have seen, if PDL enjoyed interpolation then it would be closed for bisimulation quantifiers in the class of finite models. But:

3.4.2. LEMMA. *PDL is not closed for bisimulation quantifiers, even on finite models.*

Proof.
Consider again the formula $\phi := p \wedge [a^*](p \rightarrow \langle a\rangle p)$. If PDL were closed for bisimulation quantifiers on finite models, then there would be a PDL formula θ such that, for all finite $\mathcal{M}$:
$\mathcal{M} \models \theta$ iff there exists a finite structure $\mathcal{N}$ such that $\mathcal{N} \sim_{\{a\}} \mathcal{M}$ and $\mathcal{N} \models \phi$, iff R_a is non-well-founded on $\mathcal{M}$.

But this formula cannot exists. If θ is such a formula, consider the Fisher-Ladner closure $FL(\theta)$ of θ and an a-chain $\mathcal{M} := w_1, \xrightarrow{a} w_2, \ldots, \xrightarrow{a} w_n$ with $n > |FL(\theta)|$. Then $\mathcal{M} \models \neg\theta$ and by the small model property of PDL we could collapse $\mathcal{M}$ to a model $\mathcal{N}$ by considering as equivalent all nodes which agree on $FL(\theta)$. Since $n > |FL(\theta)|$, two nodes w_i, w_j with $i \neq j$ must agree on $FL(\theta)$ and $\mathcal{N}$ contains an a-loop. Hence R_a is non-well-founded on $\mathcal{N}$ and $\mathcal{N} \models \theta$. But $\mathcal{M}$ and $\mathcal{N}$ agree on all formulae of $FL(\theta)$, hence $\mathcal{N} \models \theta$, a contradiction. □

As a consequence of Lemma 3.4.1 and Lemma 3.4.2 we have:

3.4.3. Corollary. *Uniform interpolation does not hold for Infinitary (Modal) Logic and PDL.*

3.4.4. Remark. Marco Hollenberg showed in [38, 21] that to add bisimulation quantifiers to PDL results in a logic which is equivalent to the Modal μ-Calculus. More precisely, consider the extension of PDL which allows formulae of type $\widetilde{\exists} p\phi$, to be interpreted in the usual way. We denote this logic by BQL. Hollenberg showed that BQL has the same expressive power than the Modal μ-Calculus. □

3.5 Concluding remarks.

In this chapter we proved uniform interpolation for the Modal μ-Calculus, via bisimulation quantifiers.

Some related open questions are the following.

- It is easily seen (Proposition 3.1.2) that a logic which is invariant under bisimulation and has bisimulation quantifiers enjoys uniform interpolation, but what about the inverse implication? Is it true that uniform interpolation gives bisimulation quantifiers in these logics? Proposition 3.1.3 gives only a partial answer and thus this question needs further investigation. In this connection it could be worthwhile to confront our question with a general characterization of uniform interpolation given in categorical terms in [32].

- Existential quantification modulo bisimulation has a natural generalization if we consider other equivalence relations on structures than bisimulation. As an example, one could let potential isomorphism take over the role of bisimulation and move to First Order Logic. Are there significant analogies with the results of this chapter in this new setting?

- As Marco Hollenberg showed in his dissertation, PDL plus bisimulation quantifiers has the same expressive power than the Modal μ-Calculus. However, in applications it is often useful to have results relative to the class of *finite* structures. What about PDL plus bisimulation quantifiers on *finite* models? Lemma 3.4.2 says that this logic is strictly more expressive than

PDL; is it then equally expressive as the Modal μ-Calculus on finite models? A related question is: is the Modal μ-Calculus closed for bisimulation quantifiers in the class of finite structures?

Chapter 4

Modal languages and non-well-founded universes.

In this chapter we use extended modal logics to describe *sets*. As we pointed out in Section 2.6, there is a close connection between infinitary modal formulae and AFA sets: the former can be used to describe the latter completely. It is not too surprising then that we can build a model of $ZFC^- + AFA$ in which sets are infinitary modal formulae.

Although AFA is probably the most famous axiomatization of a non-well-founded universe, it is *not* the only anti-foundation axiom which deserves investigation. Other anti-foundation axioms are known and they can be studied to understand the possible structure of non-well-founded sets in depth. In this chapter we address the following question: are there other theories besides $ZFC^- + AFA$ in which sets have a characterization in terms of (possibly extended) modal formulae?

Aczel in [2] compares AFA with an anti-foundation axiom due to Scott ([62]). In this chapter we consider this axiom and we give a language description for it by means of the Infinitary Logic of Graded Modalities $\mathcal{L}_\infty^{grad}$.

As we shall see, the Scott axiom is defined by using the unraveling operation (see Definition 2.2.4). This operation is the first of a series of *expansions* (see Definition 2.2.5) that in the limit can be used to *define* bisimulation (see Theorem 2.2.6); bisimulation, in turn, can be used to describe the axiom AFA (see Proposition 2.6.3). In this chapter we show that the other expansions can be used in a similar way: any expansion *describes* a non-well-founded universe in which sets have a description in terms of formulae of an extended infinitary modal logic. Hence, results obtained in [7] about AFA-sets can be generalized to deal with other non-well-founded sets (or, on the modal side, results about bisimulation can be generalized to deal with other structure-simulations, such as the Scott one, or the ones obtained by means of expanded structures).

The chapter is organized as follows. In the first section we describe a model for $ZFC^- + AFA$ in which sets are infinitary modal formulae. In Section 4.2 we consider the Scott axiom and prove that $\mathcal{L}_\infty^{grad}$-formulae can be used as descriptions of the Scott-sets. In Section 4.3 we consider an infinite cardinal k and we first characterize isomorphisms between k-expanded structures by means of the fragment $\mathcal{L}_\infty^{k<\diamond}$ of $\mathcal{L}_\infty^{grad}$. Then we use this result to build other non-well-founded universes in which sets are described by graded infinitary formulae.

In this chapter we only consider the simple language $\{r, R\}$, containing no propositional constants and only one binary relation symbol. As we pointed out in Section 2.6, richer languages which allow a non-empty set of propositional constants can also be interpreted in sets, but in this case we must consider a universe with *urelements*. For the sake of simplicity we restrict here to the language $\{R, r\}$, although all results in Chapter 4 can be generalized (using universes with urelements) to richer languages.

4.1 A model of ZFC^-+AFA built with infinitary modal formulae.

As we have seen in Chapter 2 any set in a universe satisfying AFA is characterized by an infinitary modal formula (see Theorem 2.6.5). This suggests a way to build a model for $ZFC^- + AFA$, out of a model of ZFC^-, in which the domain of the universe is made by considering a fragment of the class of infinitary modal formulae. To prove this theorem, we first consider a result of [7] that characterizes the infinitary modal formulae that are descriptions of sets.

4.1.1. DEFINITION. Consider the preorder $\leq$ defined in $\mathcal{L}_\infty^\diamond$ by

$$\psi \leq \phi \Leftrightarrow \models \psi \to \phi,$$

and define the class $Set(\mathcal{L}_\infty^\diamond)$ as the one containing all satisfiable $\mathcal{L}_\infty^\diamond$-formulae which are minimal with respect to $\leq$ on satisfiable formulae, that is:

$$\phi \in Set(\mathcal{L}_\infty^\diamond)$$

$$\Updownarrow$$

$$\phi \text{ is satisfiable and for all satisfiable } \psi \text{ if } \psi \leq \phi \text{ then } \psi \text{ is equivalent to } \phi.$$

□

Given a structure $\mathcal{G}$, we denote by $\phi_\mathcal{G}$ the infinitary modal formula that characterizes $\mathcal{G}$ modulo bisimulation (see Theorem 2.2.8).

4.1.2. PROPOSITION. *([7]) For all $\phi \in \mathcal{L}_\infty^\diamond$,*

$$\phi \in Set(\mathcal{L}_\infty^\diamond) \Leftrightarrow \textit{ there exists a structure } \mathcal{G} \textit{ such that } \phi_\mathcal{G} \textit{ is equivalent to } \phi.$$

Proof.
If $\phi \in Set(\mathcal{L}_\infty^\diamond)$ then by definition of $Set(\mathcal{L}_\infty^\diamond)$ the formula ϕ has a model $\mathcal{G}$. Since ϕ is minimal, we can prove that ϕ is equivalent to $\phi_\mathcal{G}$ by showing that $\models \phi_\mathcal{G} \to \phi$. To this end, let $\mathcal{H}$ be a model of $\phi_\mathcal{G}$: then $\mathcal{H}$ is bisimilar to $\mathcal{G}$ and since $\mathcal{G} \models \phi$, we obtain that $\mathcal{H} \models \phi$.

Vice versa, we want to prove that a formula of type $\phi_\mathcal{G}$ is minimal. $\phi_\mathcal{G}$ is clearly satisfiable. If $\models \psi \to \phi_\mathcal{G}$ and ψ has a model $\mathcal{H}$, we want to prove that $\models \phi_\mathcal{G} \to \psi$. Notice first that two models of $\phi_\mathcal{G}$ are always bisimilar, being both bisimilar to $\mathcal{G}$. But then if $\mathcal{K}$ is an arbitrary model of $\phi_\mathcal{G}$, it is bisimilar to $\mathcal{H}$ and hence it satisfies ψ, proving that $\models \phi_\mathcal{G} \to \psi$. □

We are now ready to build our formula-model:

4.1.3. DEFINITION. Define the universe $\mathcal{U}^\diamond$ as follows:

- The domain is given by the infinitary modal formulae in the class $Set(\mathcal{L}_\infty^\diamond)$, *modulo equivalence.*
- $$\psi \in^{\mathcal{U}^\diamond} \phi \qquad \Leftrightarrow \qquad \models \phi \to \diamond\psi.$$

□

4.1.4. REMARK. When writing *modulo equivalence* in the previous definition, we are cheating a little because the class of formulae which are equivalent to a given one does not form a set. Hence we cannot use it as an element. To be precise, one should reason as follows (see [2] or [7], where a similar discussion is made in connection with the construction of models of $ZFC^- + AFA$). Given an equivalence class A, we consider the *subset* of A of those elements of the class that have the least possible rank in the cumulative hierarchy of well-founded sets (remember that we are building our universe out of a model of ZFC, hence infinitary modal formulae are well-founded sets themselves). This *set* can then be used to represent the class. From now on we will not be so scrupulous on this particular point and use freely classes as elements knowing that we could always consider a subset of any class to represent it. Actually, even the use of the cumulative hierarchy can be eliminated by using some clever tricks, as it is done in [2]. This means that we could actually build our model inside a model of ZFC^-, without using the foundation axiom. We choose to keep things as simple as possible and for this reason we started our construction from a model of ZFC. □

4.1.5. REMARK. If $\phi, \phi' \in \mathcal{U}^\diamond$ have a common model, then they are equivalent. To see this, let $\mathcal{K}$ be the common model. By Proposition 4.1.2 it follows that if $\mathcal{H}$ is another model of ϕ then $\mathcal{K}$ is bisimilar to $\mathcal{H}$ and hence $\mathcal{H}$ is a model for ϕ' as well. Thus, $\models \phi \to \phi'$. The converse is proved similarly. □

To prove that $\mathcal{U}^\diamond$ is a model of ZFC^- we shall use Theorem 4.1.7, which is due to Rieger (see [60]).

4.1.6. Definition. A class structure $\mathcal{U}$ for the language $\{\in\}$ is a *full system* if:

- for all $u \in \mathcal{U}$, the class $\{b \in \mathcal{U} : b \in^{\mathcal{U}} u\}$ forms a set;
- for each set X which is a subset of $\mathcal{U}$, there is a unique $a_X \in \mathcal{U}$ with

$$\{b \in \mathcal{U} : b \in^{\mathcal{U}} a_X\} = X.$$

□

A note of warning is due at this point. In a full system $\mathcal{U}$ there are collections of elements that cannot be sets. For example, one can reproduce Russel's paradox in a full system $\mathcal{U}$ and show that the elements u having the property $u \notin^{\mathcal{U}} u$ do not form a set.

4.1.7. Theorem. *(ZFC^-) Any full system is a model for ZFC^-.*

We omit the proof, that can be found in [2].

4.1.8. Lemma. *(ZFC) $\mathcal{U}^\diamond$ is a full system.*

Proof.
We leave to the reader the verification that the first requirement of a full system is satisfied by $\mathcal{U}^\diamond$. To prove this, notice that, by using Proposition 4.1.2, we can consider any element $u \in \mathcal{U}^\diamond$ as a description of a structure $\mathcal{G}$; then for any $b \in^{\mathcal{U}^\diamond} u$ there is a v with $r^{\mathcal{G}} R^{\mathcal{G}} v$ such that b is equivalent to $\phi_{(\mathcal{G},v)}$.

As for the second requirement of a full system, consider a subset X of the domain of $\mathcal{U}^\diamond$.

- We first prove that there is an element $a_X \in \mathcal{U}^\diamond$ with

  $$\{b \in \mathcal{U} : b \in^{\mathcal{U}^\diamond} a_X\} = X.$$

  For each $\phi \in X$, choose a structure $\mathcal{G}_\phi$ such that $\mathcal{G}_\phi \models \phi$ (again, here we use the axiom of choice on classes only apparently, because one could reason as in the above remark). Let $\mathcal{G}^\star$ be the disjoint union of all the $\mathcal{G}_\phi$, for $\phi \in X$, to which we add the new root $r^{\mathcal{G}^\star}$ and the new edges $(r^{\mathcal{G}^\star}, r^{\mathcal{G}_\phi})$, for all $\phi \in X$. Consider the infinitary modal formula $\phi_{\mathcal{G}^\star}$ which characterizes $\mathcal{G}^\star$ modulo bisimulation (see Theorem 2.2.8).

  We prove that $\{\psi : \psi \in^{\mathcal{U}^\diamond} \phi_{\mathcal{G}^\star}\} = X$.

If $\psi \in^{\mathcal{U}^\diamond} \phi_{\mathcal{G}^\star}$, then $\models \phi_{\mathcal{G}^\star} \to \diamond\psi$ and since $\mathcal{G}^\star$ is a model of $\phi_{\mathcal{G}^\star}$, there exists $\phi \in X$ with $r^{\mathcal{G}_\phi} \models \psi$. But then ϕ and ψ have $\mathcal{G}_\phi$ as common model and hence, by Remark 4.1.5, they are equivalent. Therefore, $\psi = \phi \in X$.

Vice versa, if $\psi \in X$, then $r^{\mathcal{G}^\star} R^{\mathcal{G}^\star} r^{\mathcal{G}_\psi}$ and we can prove that $\models \phi_{\mathcal{G}^\star} \to \diamond\psi$. This is shown as follows: if $\mathcal{H}$ is a model of $\phi_{\mathcal{G}^\star}$, then $\mathcal{H}$ is bisimilar to $\mathcal{G}^\star$. Since $r^{\mathcal{G}^\star} R^{\mathcal{G}^\star} r^{\mathcal{G}_\psi}$, there exists a v with $r^{\mathcal{H}} R^{\mathcal{H}} v$ such that $(\mathcal{H}, v)$ is bisimilar to $\mathcal{G}_\psi$. This implies that $(\mathcal{H}, v) \models \psi$ and $\mathcal{H} \models \diamond\psi$, from which $\psi \in^{\mathcal{U}^\diamond} \phi_{\mathcal{G}^\star}$.

- We now prove that for any set X, there is a unique element a_X in $\mathcal{U}^\diamond$ such that

$$\{b \in \mathcal{U}^\diamond : b \in^{\mathcal{U}^\diamond} a_X\} = X.$$

Notice that this is equivalent to prove that the extensionality axiom holds in $\mathcal{U}^\diamond$. If we have two formulae $\phi, \phi' \in \mathcal{U}^\diamond$ with

$$\{\psi : \models \phi \to \diamond\psi\} = \{\psi : \models \phi' \to \diamond\psi\},$$

then we want to prove that ϕ, ϕ' are equivalent. This follows if we show that a model $\mathcal{M}$ of ϕ and a model $\mathcal{N}$ of ϕ' are always bisimilar.

To this end, consider two structures $\mathcal{M}, \mathcal{N}$, with $\mathcal{M} \models \phi$, $\mathcal{N} \models \phi'$, and let w be such that $r^{\mathcal{M}} R^{\mathcal{M}} w$. Then it is easily checked that $\models \phi \to \diamond\phi_{(\mathcal{M},w)}$ and by hypothesis we obtain $\models \phi' \to \diamond\phi_{(\mathcal{M},w)}$. Since $\mathcal{N}$ is a model of ϕ', there exists a successor v of $r^{\mathcal{N}}$ with $(\mathcal{N}, v) \models \phi_{(\mathcal{M},w)}$, that is: $(\mathcal{N}, v)$ is bisimilar to $(\mathcal{M}, w)$. Hence any successor of $r^{\mathcal{M}}$ is bisimilar to a successor of $r^{\mathcal{N}}$; since also the converse holds, $\mathcal{M}$ and $\mathcal{N}$ are bisimilar. □

From Theorem 4.1.7 and Lemma 4.1.8 it follows that $\mathcal{U}^\diamond$ is a model for ZFC^-.
In the following we prove that $\mathcal{U}^\diamond$ is a model for AFA, by using the following version of the axiom (see Proposition 2.6.3):

$AFA^{ext} :=$ A structure is strongly extensional iff it is canonical.

4.1.9. Lemma.

$\mathcal{U}^\diamond \models$ *'if a is a strongly extensional structure then it is canonical'.*

Proof.
Suppose a is a strongly extensional structure w.r.t. $\mathcal{U}^\diamond$. This means that $\mathcal{U}^\diamond$ 'sees' a as a triple, consisting of a domain d_a, a binary relation r_a on d_a, and of an element d (the root) in d_a. To prove that $\mathcal{U}^\diamond$ sees a as a picture of a set, we can make the following steps:

1. We consider the elements a, d_a, r_a, d *from the outside* (i.e. from the point of view of the model of ZFC we are working with). More precisely, we define the structure $\mathcal{O}_a$ having domain equal to $\{b \in \mathcal{U}^\diamond : b \in^{\mathcal{U}^\diamond} d_a\}$, the root equal to d, and

$$R^{\mathcal{O}_a} = \{(b,c) \in \mathcal{O}_a \times \mathcal{O}_a : \langle b,c\rangle^{\mathcal{U}^\diamond} \in^{\mathcal{U}^\diamond} d_a\},$$

   where $\langle b,c\rangle^{\mathcal{U}^\diamond}$ is the ordered pair of b, c inside $\mathcal{U}^\diamond$.

2. We prove that $\mathcal{O}_a$ is strongly extensional.

3. We prove that $\mathcal{U}^\diamond$ sees an isomorphism between a and the formula $\phi_{\mathcal{O}_a} \in \mathcal{U}^\diamond$ which characterizes $\mathcal{O}_a$ modulo bisimulation.

The second point is a delicate one: we are actually claiming that $\mathcal{U}^\diamond$ can see any bisimulation R between the nodes of a that can be seen from the outside. This is true because $\mathcal{U}^\diamond$ is a full system. To prove this, suppose there exists a bisimulation Z between $b, c \in \mathcal{O}_a$. Then Z is a subset of $\mathcal{O}_a \times \mathcal{O}_a$ and we can consider the following *subset* X of $\mathcal{U}^\diamond$:

$$X = \{\langle b,c\rangle^{\mathcal{U}^\diamond} : bRc\}.$$

Since $\mathcal{U}^\diamond$ is full, there exists a $z \in \mathcal{U}^\diamond$ such that $\{e \in \mathcal{U}^\diamond : e \in^{\mathcal{U}^\diamond} z\} = X$. Then it is not difficult to prove that

$$\mathcal{U}^\diamond \models \text{`}z \text{ is a bisimulation between } b \text{ and } c\text{'}.$$

We leave this verification to the reader, as well as the proof of the third point above. □

We are now left with the other direction of the axiom AFA^{ext}, that is, we need to prove that canonical structures are seen by $\mathcal{U}^\diamond$ as strongly extensional structures. This will be proved first in Lemma 4.1.10 from the outside. Then, using the fact that $\mathcal{U}^\diamond$ is full, we will obtain the same result from the point of view of $\mathcal{U}^\diamond$.

How does a canonical structure of $\mathcal{U}^\diamond$ look from the outside? The domain is a set of formulae in $\mathcal{U}^\diamond$ and ϕ is related to ψ iff $\models \phi \to \diamond\psi$. In the following lemma we prove that structures of this kind are strongly extensional.

4.1.10. Lemma. *Any structure $\mathcal{G}$ consisting of formulae in $\mathcal{U}^\diamond$ with $\phi R^{\mathcal{G}} \psi$ iff $\models \phi \to \diamond\psi$, is strongly extensional.*

Proof.
Suppose there is a bisimulation Z that links two points ϕ, ψ of $\mathcal{G}$. We must prove that ϕ is equivalent to ψ. If $\mathcal{H}$ and $\mathcal{K}$ are models of ϕ, ψ, respectively, we claim that they are bisimilar. In this way we prove that ϕ and ψ have a common model

and by Remark 4.1.5 they are equivalent. To prove the claim, define the following relation $S \subseteq \mathcal{H} \times \mathcal{K}$:

$$wSv \text{ iff there exist } \theta, \gamma \text{ with } (\mathcal{H}, w) \models \theta,\ (\mathcal{K}, v) \models \gamma, \text{ and } \theta Z \gamma.$$

We prove that S is a bisimulation between $\mathcal{H}$ and $\mathcal{K}$.

From the hypothesis it follows that $r^{\mathcal{H}} S r^{\mathcal{K}}$.

Suppose wSv, that is, there exist θ, γ with $(\mathcal{H}, w) \models \theta$, $(\mathcal{K}, v) \models \gamma$, and $\theta Z \gamma$. If $wR^{\mathcal{H}}w'$ we show that there exists a v' with $vR^{\mathcal{K}}v'$ and $w'Sv'$. Since $(\mathcal{H}, w) \models \theta$, it is easily proved that $\models \theta \to \Diamond\phi_{(\mathcal{H},w')}$. Since Z is a bisimulation on $\mathcal{G}$ and $\theta Z \gamma$, there exists γ' with $\models \gamma \to \Diamond\gamma'$ and $\phi_{(\mathcal{H},w')} Z \gamma'$. But $(\mathcal{K}, v) \models \gamma$, hence there exists a v' with $vR^{\mathcal{K}}v'$ and $(\mathcal{K}, v') \models \gamma'$. By definition of S we conclude $w'Sv'$. The other condition of a bisimulation is verified similarly. □

4.1.11. LEMMA.

$\mathcal{U}^{\Diamond} \models$ *'if a is a canonical structure then it is strongly extensional'.*

Proof.
By Lemma 4.1.11 a canonical structure of $\mathcal{U}^{\Diamond}$ is strongly extensional from the outside and one can reason as in the second point of Lemma 4.1.9 to prove that it is also strongly extensional from the point of view of $\mathcal{U}^{\Diamond}$. □

4.1.12. THEOREM. *$\mathcal{U}^{\Diamond}$ is a model for $ZFC^- + AFA$*

Proof.
Just put together Lemma 4.1.8, Rieger's Theorem, Lemma 4.1.9, and Lemma 4.1.11. □

4.2 The Scott anti-foundation axiom and the Logic of Graded Modalities.

In this section we show that the Infinitary Logic of Graded Modalities $\mathcal{L}_{\infty}^{grad}$ (see Section 2.1) is the appropriate language description for the theory $ZFC^-{+}SAFA$: for any set a in a universe satisfying $SAFA$ there exists a formula $\phi_a \in \mathcal{L}_{\infty}^{grad}$ such that for all sets b in the same universe it holds:

$$b \models \phi_a \Leftrightarrow b = a.$$

In words: any *set* a in a universe satisfying $SAFA$ admits a description by means of a graded infinitary modal formula. This section is divided into two parts. In the first one we briefly describe the axiom $SAFA$ and in the second one we prove the formula-characterization Theorem (Theorem 4.2.8).

4.2.1 The Scott anti-foundation axiom.

By weakening the definition of a strongly extensional structure, Aczel was able to obtain different anti-foundation axioms and, among them, the Scott axiom (see [2]). In the Aczel-version this axiom is described by defining the following relation $\sim_{Scott}$ between structures.

$$\mathcal{G} \sim_{Scott} \mathcal{H} := \mathcal{G}^1 \simeq \mathcal{H}^1,$$

where $\mathcal{G}^1$ and $\mathcal{H}^1$ are the unravelings of $\mathcal{G}$ and $\mathcal{H}$ respectively (see Definition 2.2.4). It is easy to prove that $\sim_{Scott}$ is a bisimulation and we can define the notion of an extensional structure with respect to it.

4.2.1. DEFINITION. A structure $\mathcal{G}$ is *Scott-extensional* if two different nodes are never $\sim_{Scott}$-bisimilar (i.e. $(\mathcal{G}, w)^1 \simeq (\mathcal{G}, v)^1$ implies $w = v$, for all $w, v \in \mathcal{G}$). □

The Scott anti-foundation axiom $SAFA$ can be stated as follows:

A structure is Scott-extensional iff it is canonical.

It can be proved that $ZFC^- + SAFA$ is consistent and Aczel in [2] proved that $ZFC^- + SAFA$ and $ZFC^- + AFA$ are pairwise incompatible.

Let us compare the two theories with respect to decorations of structures. In $ZFC^- + SAFA$ as in $ZFC^- + AFA$ we can prove that *any* structure $\mathcal{G}$ has a decoration. To see this, consider again the collapse of $\mathcal{G}$. Being strongly extensional, $Coll(\mathcal{G})$ is also Scott-extensional and we can use Proposition 2.6.1 to define a decoration of $\mathcal{G}$ from a decoration of its collapse.

On the other hand, there are structures with more than one decoration: consider for example the structure $\mathcal{G}$, which consists of a single reflexive node $r^{\mathcal{G}}$ and the structure $\mathcal{H}$ which consists of two nodes w, v with $r^{\mathcal{H}} = w$, $wR^{\mathcal{H}}w$, $wR^{\mathcal{H}}v$, and $vR^{\mathcal{H}}v$, as in fig.1.

fig.1: two different AFA-sets.

It is easy to see that both $\mathcal{G}$ and $\mathcal{H}$ are Scott-extensional, hence under the Scott axiom they have injective decorations $d_{\mathcal{G}}$ and $d_{\mathcal{H}}$ respectively. Now, consider the function d with domain $\mathcal{H}$ defined by $d(w) = d(v) = d_{\mathcal{G}}(r^{\mathcal{G}})$. This is easily seen to be a decoration of $\mathcal{H}$ which is not injective, hence it must be different from $d_{\mathcal{H}}$. Thus, $\mathcal{H}$ is a structure with more than one decoration.

Finally, notice that having a common decoration in $ZFC^- + SAFA$ does not imply Scott-bisimilarity. Consider a single reflexive node and the full binary tree: these two structures have the constant function which is always equal to a set $\Omega = \{\Omega\}$ as common decoration, but they are not Scott-bisimilar.

As for $ZFC^- + AFA$ we can prove that the axiom $SAFA$ implies a strengthening of the extensionality axiom. Given two sets a, b, we say that they are Scott-bisimilar if the corresponding canonical structures $\mathcal{G}_a, \mathcal{G}_b$ are Scott-bisimilar. We leave the verification of the following proposition to the reader (compare with Proposition 2.6.4).

4.2.2. PROPOSITION. *(ZFC^-) The following are equivalent:*

1) All structures with an injective decoration are Scott-extensional.

2) If the set a is Scott-bisimilar to the set b then $a = b$.

From this proposition it follows that under the Scott axiom two Scott-bisimilar sets are equal. By comparing this result with the corresponding one in $ZFC^- + AFA$, which says that bisimilar sets are equal, we see that in a certain sense Scott-sets are *more* than AFA-sets. In a Scott-universe, there can be two bisimilar different sets: consider for example the sets $d_{\mathcal{G}}$, $d_{\mathcal{H}}$ defined above.

4.2.2 Scott-sets and the Logic of Graded Modalities.

In this part we prove that the logic $\mathcal{L}_{\infty}^{grad}$ is the appropriate language description for the Scott-sets. First, we show in Lemma 4.2.3 that for each structure $\mathcal{G}$ there exists a formula $\phi_{\mathcal{G}}^{Scott} \in \mathcal{L}_{\infty}^{grad}$ that characterizes $\mathcal{G}$ modulo counting bisimulation (see Definition 2.4.10). Then we prove in Theorem 4.2.7 that two structures are Scott-bisimilar iff they satisfy the same $\mathcal{L}_{\infty}^{grad}$-formulae. Putting these results together we obtain the $\mathcal{L}_{\infty}^{grad}$ description of Scott-sets in Theorem 4.2.8.

4.2.3. LEMMA. *For any structure $\mathcal{G}$ there exists a formula $\phi_{\mathcal{G}} \in \mathcal{L}_{\infty}^{grad}$ such that for any structure $\mathcal{H}$ it holds:,*

$$\mathcal{H} \models \phi_{\mathcal{G}} \Leftrightarrow \ \mathcal{G} \text{ and } \mathcal{H} \text{ are counting bisimilar.}$$

Proof.
Fix a structure $\mathcal{G}$ and let k be a cardinal such that $|\mathcal{G}| < k$. Consider the logic $\mathcal{L}_{k^+}^{\diamond < k^+}$, that is, the logic that allows conjunctions only on sets of cardinality

strictly smaller than k^+ and that uses only the operators $\Diamond_h$ for $h < k^+$. Denote by $\equiv_{k^+}^{\Diamond<k^+}$ the relation *to satisfy the same* $\mathcal{L}_{k^+}^{\Diamond<k^+}$*-formulae.*

Claim. For any model $\mathcal{H}$, if $\mathcal{G}$ and $\mathcal{H}$ satisfy the same $\mathcal{L}_{k^+}^{\Diamond<k^+}$ formulae then the relation $\equiv_{k^+}^{\Diamond<k^+}$ restricted to nodes in $\mathcal{G} \times \mathcal{H}$ is a counting bisimulation between $\mathcal{G}$ and $\mathcal{H}$.

Let us prove the claim.

- The roots are connected by hypothesis.
- To prove the third clause of Definition 2.4.10 (there is no problem with the second clause because the language we consider in this chapter is $\{r, R\}$), we first show that if $(\mathcal{G}, w) \equiv_{k^+}^{\Diamond<k^+} (\mathcal{H}, v)$, then for any $w' \in R(w)$ there exists a formula $\phi_{w'} \in \mathcal{L}_{k^+}^{\Diamond<k^+}$ such that for any $z \in R(w) \cup R(v)$:

  $$(\star) \quad z \models \phi_{w'} \Leftrightarrow z \equiv_{k^+}^{\Diamond<k^+} w'.$$

  This is proved as follows. If $w' \in R(w)$, let $[w'] = \{w^* \in R(w) : w^* \equiv_{k^+}^{\Diamond<k^+} w'\}$. For any $w'' \in R(w) \setminus [w']$, let $\psi_{w''} \in \mathcal{L}_{k^+}^{\Diamond<k^+}$ be such that $w'' \not\models \psi_{w''}$ and $w' \models \psi_{w''}$.

  We define
  $$\phi_{w'} := \bigwedge_{w'' \in R(w) \setminus [w']} \psi_{w''}.$$

  Then $\phi_{w'} \in \mathcal{L}_{k^+}^{\Diamond<k^+}$ and we can prove that $\phi_{w'}$ verifies property $(\star)$ above.

  The direction from right to left is obvious since $w' \models \phi_{w'}$ and $\phi_{w'} \in \mathcal{L}_{k^+}^{\Diamond<k^+}$.

  For the other direction, suppose first that $z \in R(w)$: if $z \not\equiv_{k^+}^{\Diamond<k^+} w'$, then $z \not\models \psi_z$ and hence $z \not\models \phi_{w'}$. If $z \in R(v)$ and $z \models \phi_{w'}$, we prove that for any formula $\theta \in \mathcal{L}_{k^+}^{\Diamond<k^+}$, if $z \models \theta$ then $w' \models \theta$. This is enough to prove that $z \equiv_{k^+}^{\Diamond<k^+} w'$. If $z \models \theta$ then $v \models \Diamond(\theta \wedge \phi_{w'})$, hence $w \models \Diamond(\theta \wedge \phi_{w'})$ and there exists $w'' \in R(w)$ with $w'' \models \theta \wedge \phi_{w'}$. From $w'' \models \phi_{w'}$ it follows $w'' \equiv_{k^+}^{\Diamond<k^+} w'$ and thus $w' \models \theta$.

  Having found a $\phi_{w'}$ with property $(\star)$ for any $w' \in R(w)$, we can now prove the second clause of the definition of a counting bisimulation.

  Let X be a subset of $R(w)$. If $|X| = h$, then $w \models \Diamond_h(\bigvee_{w' \in X} \phi_{w'})$ and since this formula belongs to $\mathcal{L}_{k^+}^{\Diamond<k^+}$, by hypothesis we obtain $v \models \Diamond_h(\bigvee_{w' \in X} \phi_{w'})$. But any successor of v that satisfies $\bigvee_{w' \in X} \phi_{w'}$ must be in relation $\equiv_{k^+}^{\Diamond<k^+}$ with at least an element of X, hence

  $$\{y \in R(v) : \exists x \in X \ x \equiv_{k^+}^{\Diamond<k^+} y\} \supseteq \{y \in R(v) : y \models \bigvee_{w' \in X} \phi_{w'}\}.$$

  This implies that $|\{y \in R(v) : \exists x \in X \ x \equiv_{k^+}^{\Diamond<k^+} y\}| \geq |X|$.

- To prove the last clause of the definition of a counting bisimulation, suppose that X is a subset of $R(v)$. We have to prove that $|Y| \geq |X|$, where Y is the set $Y = \{w' \in R(w) : \exists v' \in X \ w' \equiv_{k^+}^{\Diamond < k^+} v'\}$.

  Let $h = |Y|$. Since $h < k$, the formula

  $$\neg\Diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'})$$

  belongs to $\mathcal{L}_{k^+}^{\Diamond < k^+}$. Since Y is closed under $\equiv_{k^+}^{\Diamond < k^+}$, any $w^* \in R(w)$ which satisfies the formula $\bigvee_{w' \in Y} \phi_{w'}$ belongs to Y and $w \models \neg\Diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'})$. By hypothesis, $v \models \neg\Diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'})$ as well. Notice that for all $v' \in X$ there exists $w' \in R(w)$ with $w' \equiv_{k^+}^{\Diamond < k^+} v'$: this is simply because

  $$(\mathcal{G}, w) \models \Box(\bigvee_{w' \in R(w)} \phi_{w'}),$$

  and $w \equiv_{k^+}^{\Diamond < k^+} v$. By definition, this w' belongs to Y and we obtain that any element of X satisfies $\bigvee_{w' \in Y} \phi_{w'}$. From $v \models \neg\Diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'})$ it then follows $|X| \leq h = |Y|$.

Using the claim we can prove that the formula

$$\phi_{\mathcal{G}}^{Scott} = \bigwedge\{\phi \in \mathcal{L}_{k^+}^{\Diamond < k^+} : \mathcal{G} \models \phi\}$$

characterizes $\mathcal{G}$ modulo counting bisimulation. Notice that even though $\phi_{\mathcal{G}}$ is no longer an $\mathcal{L}_{k^+}^{\Diamond < k^+}$ formula, it is still an $L_{\infty}^{\Diamond < k^+}$ formula (and hence an $\mathcal{L}_{\infty}^{grad}$ formula) because $\mathcal{L}_{k^+}^{\Diamond < k^+}$ forms a set. We want to prove that if $\mathcal{H} \models \phi_{\mathcal{G}}^{Scott}$ then $\mathcal{H}$ is counting bisimilar to $\mathcal{G}$. But $\mathcal{H} \models \phi_{\mathcal{G}}^{Scott}$ implies that $\mathcal{G}$ and $\mathcal{H}$ verify the same $\mathcal{L}_{k^+}^{\Diamond < k^+}$-formulae, hence by the previous claim $\mathcal{H}$ and $\mathcal{G}$ are counting bisimilar. $\Box$

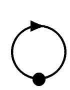

Ω

$$\phi_{\Omega}^{Scott} = \bigwedge_{n \in \omega}(\Box^n \Diamond \top \wedge \Box^n \neg\Diamond^2 \top).$$

A circular *SAFA*-set and its formula.

From the previous lemma it follows:

4.2.4. COROLLARY. *Two structures $\mathcal{G}$ and $\mathcal{H}$ satisfy the same $\mathcal{L}_\infty^{grad}$-formulae iff they are counting bisimilar.*

We are now going to prove (Theorem 4.2.7) that the relation *to satisfy the same $\mathcal{L}_\infty^{grad}$-formulae*, which we denote by $\equiv_\infty^{grad}$, coincides with the Scott-bisimulation defined in the previous part of this section.

We first need some lemmas.

4.2.5. LEMMA. *If $\mathcal{G}$ is a structure and $\sigma \in \mathcal{G}^1$ then $(\mathcal{G}, end(\sigma)) \equiv_\infty^{grad} (\mathcal{G}^1, \sigma)$.*

Proof.
This is true because the relation $Z \subseteq \mathcal{G}^1 \times \mathcal{G}$ defined by $Z = \{(\sigma, end(\sigma)) : \sigma \in \mathcal{G}^1\}$ is a counting bisimulation between $\mathcal{G}^1$ and $\mathcal{G}$. □

4.2.6. LEMMA. *If $\mathcal{G}, \mathcal{H}$ are structures with $(\mathcal{G}, w) \equiv_\infty^{grad} (\mathcal{H}, v)$, then there exists a bijective function $f_v^w : R(w) \to R(v)$ such that $(\mathcal{G}, \overline{w}) \equiv_\infty^{grad} (\mathcal{H}, f_v^w(\overline{w}))$, for any $\overline{w} \in R(w)$.*

Proof.
We first prove that if $\overline{w} \in R(w)$ the two sets $[\overline{w}]_\mathcal{G} = \{w' \in R(w) : (\mathcal{G}, w') \equiv_\infty^{grad} (\mathcal{G}, \overline{w})\}$ and $[\overline{w}]_\mathcal{H} = \{v' \in R(v) : (\mathcal{H}, v') \equiv_\infty^{grad} (\mathcal{G}, \overline{w})\}$ have the same cardinality. Once we have proved this, if $f_{\overline{w}}$ is a bijection between $[\overline{w}]_\mathcal{G}$ and $[\overline{w}]_\mathcal{H}$, for any $\overline{w} \in R(w)$ and $\{w_i\}_{i \in I}$ is a set of representatives for the $\equiv_\infty^{grad}$-equivalence classes of successors of w, we can define $f_v^w = \bigcup_{i \in I} f_{w_i}$.

Since $\equiv_\infty^{grad}$ is a counting bisimulation we have that the set

$$Y = \{v' \in R(v) : \exists w' \in [\overline{w}]_\mathcal{G} \; w' \equiv_\infty^{grad} v'\}$$

has cardinality grater or equal to $[\overline{w}]_\mathcal{G}$. But $[\overline{w}]_\mathcal{H} \supseteq Y$, hence $|[\overline{w}]_\mathcal{H}| \geq |[\overline{w}]_\mathcal{G}|$. The reverse inequality is similarly proved. □

4.2.7. THEOREM. *If $\mathcal{G}, \mathcal{H}$ are structures then*

$$\mathcal{G} \equiv_\infty^{grad} \mathcal{H} \Leftrightarrow \; \mathcal{G}^1 \textit{ is isomorphic to } \mathcal{H}^1.$$

Proof.
If $\mathcal{G} \equiv_\infty^{grad} \mathcal{H}$, then by Lemma 4.2.5 we have $\mathcal{G}^1 \equiv_\infty^{grad} \mathcal{H}^1$. We will define a function $F : \mathcal{G}^1 \mapsto \mathcal{H}^1$ inductively in such a way that $(\mathcal{G}^1, \sigma) \equiv_\infty^{grad} (\mathcal{H}^1, F(\sigma))$.

Let $F(r^{\mathcal{G}^1}) = r^{\mathcal{H}^1}$. Suppose $F(\sigma)$ has already been defined and $(\mathcal{G}^1, \sigma) \equiv_\infty^{grad} (\mathcal{H}^1, F(\sigma))$. By applying Lemma 4.2.6 to the unraveled structures $\mathcal{G}^1, \mathcal{H}^1$, we obtain a bijective function $f_{F(\sigma)}^\sigma$ between $R(\sigma)$ and $R(F(\sigma))$. For any $w' \in R(end(\sigma)$, define $F(\sigma w') = f_{F(\sigma)}^\sigma(\sigma w')$.

We now prove that F is an isomorphism between $\mathcal{G}^1$ and $\mathcal{H}^1$.

1. F is injective.
   We prove that $F(\sigma) = F(\sigma')$ implies $\sigma = \sigma'$ by induction on the length of σ.
   If $\sigma = r^{\mathcal{G}^1}$, then $F(\sigma) = r^{\mathcal{H}^1}$; hence, $F(\sigma') = r^{\mathcal{H}^1}$ and $\sigma' = r^{\mathcal{G}^1}$, because if σ' were a successor of any node, so would be $F(\sigma')$, by definition of F.
   If $\sigma = \sigma_1 w_1$, then $F(\sigma) = f^{\sigma_1}_{F(\sigma_1)}(\sigma_1 w_1)$. If $\sigma' = \sigma_2 w_2$, then $F(\sigma') = f^{\sigma_2}_{F(\sigma_2)}(\sigma_2 w_2)$. Hence, $F(\sigma)$ is a successor of $F(\sigma_1)$, while $F(\sigma')$ is a successors of $F(\sigma_2)$. But $F(\sigma) = F(\sigma')$ and a node in $\mathcal{H}^1$ cannot have two different predecessors, so $F(\sigma_1) = F(\sigma_2)$: by induction $\sigma_1 = \sigma_2$. From $f^{\sigma_1}_{F(\sigma_1))}(\sigma_1 w_1) = f^{\sigma_2}_{F(\sigma_2)}(\sigma_2 w_2)$ it follows $\sigma_1 w_1 = \sigma_2 w_2$, because the function $f^{\sigma_1}_{F(\sigma_1))} = f^{\sigma_2}_{F(\sigma_2)}$ is injective. Since $\sigma_1 = \sigma_2$ we obtain $w_1 = w_2$.

2. $\sigma R \sigma'$ iff $F(\sigma) R F(\sigma')$.
   From left to right by definition of F.
   From right to left. Suppose $F(\sigma) R F(\sigma')$. Since $f^{\sigma}_{F(\sigma)}$ is a bijection between the successors of σ and the successors of $F(\sigma)$, there exists $\tau \in R(\sigma)$ such that $f^{\sigma}_{F(\sigma)}(\tau) = F(\sigma')$. We want to show that $\tau = \sigma'$. By definition of F we have $F(\tau) = F(\sigma')$ and then $\tau = \sigma'$ because F is injective.

3. F is surjective. Left to the reader.

Vice versa, if $\mathcal{G}^1$ is isomorphic to $\mathcal{H}^1$, then by Lemma 4.2.5 $\mathcal{G} \equiv^{grad}_{\infty} \mathcal{H}$. □

Summarizing, in this section we proved that Scott-bisimulation can be characterized via $\mathcal{L}^{grad}_{\infty}$ equivalence. More than this, using counting bisimulations we proved that any structure is characterized modulo Scott-bisimulation by an $\mathcal{L}^{grad}_{\infty}$ formula. Since Scott-sets are equal iff they are Scott-bisimilar, we have proved the following:

4.2.8. THEOREM. *If a is a set in a Scott-universe, then there exists a $\mathcal{L}^{grad}_{\infty}$-formula ϕ_a such that for any set b:*

$$b \models \phi_a \Leftrightarrow b = a.$$

4.2.9. REMARK. Theorem 4.2.8 suggests a possible proof of the consistency of $ZFC^- + SAFA$ w.r.t. ZFC: we could try to build a model for $ZFC^- + SAFA$ starting from formulae of $\mathcal{L}^{grad}_{\infty}$ as we did for $ZFC^- + AFA$ using $\mathcal{L}^{\diamond}_{\infty}$-formulae. However, we cannot use the minimal formulae with respect to implication on satisfiable formulae, with the membership relation defined as in the AFA-case. To see this, notice that these formulae are exactly the ones that characterize structures modulo counting bisimulation: just use the same proof of Proposition 4.1.2. But then we have a problem with the extensionality axiom: consider the structures $\mathcal{G}$ and $\mathcal{H}$ defined as follows: $\mathcal{G} = \{w, v, u\}$ with $r^{\mathcal{G}} = w$, $R^{\mathcal{G}} = \{(w,v),(w,u)\}$, and $\mathcal{H} = \{w', v'\}$ with $r^{\mathcal{H}} = w'$ and $w' R^{\mathcal{H}} v'$ (see figure).

$\mathcal{G}$ $\mathcal{H}$

$$\phi_{\mathcal{G}}^{Scott} = \Diamond_2 \Box \bot \wedge \neg \Diamond_3 \top \qquad \phi_{\mathcal{H}}^{Scott} = \Diamond \Box \bot \wedge \neg \Diamond_2 \top$$

The formula $\phi_{\mathcal{G}}^{Scott}$ (viewed as a set) has the same elements as $\phi_{\mathcal{H}}^{Scott}$, but these two formulae *are not* equivalent. A true formula-model for $ZFC^- + SAFA$ can be obtained by considering just the formulae of type $\phi_{\mathcal{G}}^{Scott}$, for *Scott-extensional* structures $\mathcal{G}$. Notice that in the AFA case the restriction to strongly canonical structures, when considering formulae of type $\phi_{\mathcal{G}}$, would have made no difference, because any structure $\mathcal{G}$ is bisimilar to its collapse $\mathcal{G}$, which is strongly extensional. Unfortunately, if we consider counting bisimulation, then we do not have a corresponding notion of collapse and the restriction to Scott-extensional structures is a true one. Hence, the problem of finding a nice, bisimulation-free description of the sets of this model, as in Definition 4.1.3, is still open. □

4.3 k-expansions and non-well-founded sets.

In the previous section we proved that the unraveling operation can be used to define the Scott anti-foundation axiom, which gives us a different conception of non-well-founded sets than the one obtained using AFA. Moreover, the unraveling construction is just the first of a series of operations, the k-expansions for each cardinal k, that in the limit defines bisimulation: two structures $\mathcal{G}$ and $\mathcal{H}$ are bisimilar if and only if there exists a cardinal k such that the k-expansions $\mathcal{G}^k$ and $\mathcal{H}^k$ are isomorphic (see Theorem 2.2.6). It is then natural to investigate whether we can use the other expansions as well, obtaining different conceptions of non-well-founded sets than the $SAFA$ and AFA ones and if we can obtain for these sets a description in terms of infinitary extended modal formulae. In this section we give a positive answer to this question. The section is divided into two parts. In the first part we show how the technical results under the $SAFA$ characterization of sets by means of graded infinitary formulae, that is Lemma 4.2.3 and Theorem 4.2.7, can be generalized to k-expanded structures, for an infinite cardinal k. In the second part we discuss the possibility of using these results for describing non-well-founded universes in which sets are described by formulae.

4.3.1 Graded modalities and expanded structures

In this section we work on generalizing Theorem 4.2.7 to k-expanded structures (see Definition 2.2.5). We will see that if k is an infinite cardinal then two structures are equivalent with respect to $\mathcal{L}_\infty^{k<\diamond}$ (see Definition 2.4.4) iff their k-expansions are isomorphic. This is easily seen to be false if k is finite (see Lemma 4.3.6). Moreover, we will also prove the analogous of Lemma 4.2.3 for the logics $\mathcal{L}_\infty^{k<\diamond}$: any structure $\mathcal{G}$ can be characterized modulo $(k<)$-counting bisimulation with a formula in $\mathcal{L}_\infty^{k<\diamond}$. The proofs of these results are similar to the ones given in the previous section for the Scott-bisimulation, hence we will merely indicate differences between the corresponding proofs.

4.3.1. LEMMA. *Let k be a cardinal. For any structure $\mathcal{G}$ there exists a formula $\phi_\mathcal{G}^k \in \mathcal{L}_\infty^{k<\diamond}$ such that for any structure $\mathcal{H}$ it holds:,*

$$\mathcal{H} \models \phi_\mathcal{G}^k \Leftrightarrow \mathcal{G} \textit{ and } \mathcal{H} \textit{ are } (k<)\textit{-counting bisimilar.}$$

Proof.
Let us fix a structure $\mathcal{G}$ and let k' be a cardinal greater than k such that $|\mathcal{G}| < k'$. Consider the logic $\mathcal{L}_{k'^+}^{k<\diamond<k'^+}$. As in Theorem 4.2.3, to prove the lemma it is enough to show that, for any other model $\mathcal{H}$, if $\mathcal{G}$ and $\mathcal{H}$ satisfy the same $\mathcal{L}_{k'^+}^{k<\diamond<k'^+}$ formulae (notation: $(\mathcal{G}, w) \equiv_{k'^+}^{k<\diamond<k'^+} (\mathcal{H}, v)$), then they are $(k<)$-counting bisimilar.

To this end, suppose that $(\mathcal{G}, w) \equiv_{k'^+}^{k<\diamond<k'^+} (\mathcal{H}, v)$. As in Theorem 4.2.3, for any $w' \in R(w)$, we can define a formula $\phi_{w'} \in \mathcal{L}_{k'^+}^{k<\diamond<k'^+}$ with the property that for any $z \in R(w) \cup R(v)$ it holds

$$z \models \phi_{w'} \Leftrightarrow z \equiv_{k'^+}^{k<\diamond<k'^+} w'.$$

Then the proof proceeds with minor adjustments as in Theorem 4.2.3. Let us just show that if $X \subseteq R(v)$ and $|X| > k$, then $|Y| \geq |X|$, where

$$Y = \{y \in R(w) : \exists x \in X \; x \equiv_{k'^+}^{k<\diamond<k'^+} y\}.$$

First, let us show that $|Y| > k$. If not, then

$$w \models \neg\diamond_{k^+}(\bigvee_{w' \in Y} \phi_{w'}),$$

and $\neg\diamond_{k^+}(\bigvee_{w' \in Y} \phi_{w'}) \in \mathcal{L}_{k'^+}^{k<\diamond<k'^+}$. By hypothesis v satisfies this formula as well, and from this we obtain $|X| \leq k$, which is not. Then if $|Y| = h > k$, the formula $\neg\diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'})$ belongs to $\mathcal{L}_{k'^+}^{k<\diamond<k'^+}$ and we know that

$$w \models \neg\diamond_{h^+}(\bigvee_{w' \in Y} \phi_{w'}).$$

Hence, v satisfies this formula as well and $|X| \leq h = |Y|$. $\square$

4.3.2. COROLLARY. *Two structures $\mathcal{G}$ and $\mathcal{H}$ satisfy the same $\mathcal{L}_\infty^{k<\diamond}$ formulae iff they are $(k<)$-counting bisimilar.*

If $k \geq \omega$, we are now going to prove (Theorem 4.3.5) that two structures satisfy the same $\mathcal{L}_\infty^{k<\diamond}$-formulae iff they have isomorphic k-expansions. This result generalizes Theorem 4.2.7 (in which k is equal to 1).

We first need two lemmas.

4.3.3. LEMMA. *Let k be an infinite cardinal. If $\mathcal{G}$ is a structure and $\sigma \in \mathcal{G}^k$ then $(\mathcal{G}, end(\sigma)) \equiv_\infty^{k<\diamond} (\mathcal{G}^k, \sigma)$.*

Proof.
Let $Z \subseteq \mathcal{G}^k \times \mathcal{G}$ be the relation defined as $Z = \{(\sigma, end(\sigma)) : \sigma \in \mathcal{G}^k\}$. We show that Z is a $(k<)$-counting bisimulation between $\mathcal{G}^k$ and $\mathcal{G}$. To prove this, suppose that $\sigma \in \mathcal{G}^k$ and let $X \subseteq R(\sigma)$ be such that $|X| > k$. If $Y = \{y \in R(end(\sigma)) : \exists \sigma' \in X\ y = end(\sigma')\}$, we must show that $|Y| \geq |X|$. Since k is an infinite cardinal smaller than $|X|$, if $|Y| < |X|$, the set $k \times Y$ has cardinality smaller than X as well. But there is an injection f between X and $k \times Y$ defined as $f(\sigma\alpha w') = (\alpha, w')$: a contradiction.

Vice versa, if $X \subseteq R(end(\sigma))$ and $|X| > k$, then the set $Y = \{\sigma' \in R(\sigma) : \exists w' \in X\ w' = end(\sigma')\}$ contains, for any $w' \in X$, all elements of the form $\sigma 0 w'$. Hence $|Y| \geq |X|$. □

4.3.4. LEMMA. *If k is an infinite cardinal, $\mathcal{G}$ and $\mathcal{H}$ are structures, $\sigma \in \mathcal{G}^k, \tau \in \mathcal{H}^k$, and $(\mathcal{G}^k, \sigma) \equiv_\infty^{k<\diamond} (\mathcal{H}^k, \tau)$, then there exists a bijective function $f_\tau^\sigma : R(\sigma) \to R(\tau)$ such that $(\mathcal{H}^k, f_\tau^\sigma(\overline{\sigma})) \equiv_\infty^{k<\diamond} (\mathcal{G}^k, \overline{\sigma})$, for any $\overline{\sigma} \in R(\sigma)$.*

Proof.
We first prove that if $\overline{\sigma} \in R(\sigma)$, the two sets $[\overline{\sigma}]_\mathcal{G} = \{\sigma' \in R(\sigma) : (\mathcal{G}^k, \sigma') \equiv_\infty^{k<\diamond} (\mathcal{G}^k, \overline{\sigma})\}$ and $[\overline{\sigma}]_\mathcal{H} = \{\tau' \in R(\tau) : (\mathcal{H}^k, \tau') \equiv_\infty^{k<\diamond} (\mathcal{G}^k, \overline{\sigma})\}$ have the same cardinality. Once we have proved this, if $f_{\overline{\sigma}}$ is a bijection between $[\overline{\sigma}]_\mathcal{G}$ and $[\overline{\sigma}]_\mathcal{H}$ for any $\overline{\sigma} \in R(\sigma)$, and $\{\sigma_i\}_{i \in I}$ is a set of representatives for the $\equiv_\infty^{k<\diamond}$-equivalence classes of successors of σ, we can define $f_\tau^\sigma = \bigcup_{i \in I} f_{\sigma_i}$.

It can be easily proved that

$$x \in [\overline{\sigma}]_\mathcal{G} \Leftrightarrow \exists w' \in R(end(\sigma)) \exists \alpha < k\ (\mathcal{G}, w') \equiv_\infty^{k<\diamond} (\mathcal{G}, end(\overline{\sigma})) \text{ and } x = \sigma\alpha w',$$

and

$$y \in [\overline{\sigma}]_\mathcal{H} \Leftrightarrow \exists v' \in R(end(\tau)) \exists \alpha < k\ (\mathcal{H}, v') \equiv_\infty^{k<\diamond} (\mathcal{G}, end(\overline{\sigma})) \text{ and } y = \tau\alpha v'.$$

Since $\equiv_\infty^{k<\diamond}$ is a $(k<)$-counting bisimulation between $(\mathcal{G}, end(\sigma))$ and $(\mathcal{H}, end(\tau))$ it is easy to see that the two sets

$$\{w' \in R(end(\sigma)) : (\mathcal{G}, w') \equiv_\infty^{k<\diamond} (\mathcal{G}, end(\overline{\sigma}))\}$$

and

$$\{v' \in R(end(\tau)) : (H, v') \equiv_{\infty}^{k<\Diamond} (\mathcal{G}, end(\overline{\sigma}))\}$$

either have cardinality smaller than k or have the same cardinality. From this it follows that $|[\overline{\sigma}]_{\mathcal{G}}| = |[\overline{\sigma}]_{\mathcal{H}}|$. □

4.3.5. Theorem. *If $\mathcal{G}, \mathcal{H}$ are structures then*

$$\mathcal{G} \equiv_{\infty}^{k<\Diamond} \mathcal{H} \Leftrightarrow \mathcal{G}^k \text{ is isomorphic to } \mathcal{H}^k.$$

Proof.
Analogous to the proof of Theorem 4.2.7, using Lemma 4.3.4. □

4.3.2 The case of finite cardinalities.

Let us now see that Lemma 4.3.4 and Theorem 4.3.5 can fail when k is a natural number greater than one.

4.3.6. Lemma. *For any natural number $k > 1$, there is a structure $\mathcal{G}$ such that $\mathcal{G} \not\equiv_{\infty}^{k<\Diamond} \mathcal{G}^k$.*

Proof.
Let the domain of $\mathcal{G}$ be the set $\{w, a, b\}$ and $R^{\mathcal{G}} = \{(w, a), (w, b)\}$. Then $\mathcal{G} \not\models \Diamond_{k+1}(\Box\bot)$, while $\mathcal{G}^k \models \Diamond_{k+1}(\Box\bot)$. □

4.3.7. Lemma. *For any natural number k greater than 1, there are two structures $\mathcal{G}$ and $\mathcal{H}$ such that $\mathcal{G} \equiv_{\infty}^{k<\Diamond} \mathcal{H}$ but $\mathcal{G}^k$ is not isomorphic to $\mathcal{H}^k$.*

Proof.
Let $\mathcal{G} = \{w, w_1, w_2\}$ with $r^{\mathcal{G}} = w$, $wR^{\mathcal{G}}w_1$, and $wR^{\mathcal{G}}w_2$, and let $\mathcal{H} = \{v, v_1\}$ with $r^{\mathcal{H}} = v$, $vR^{\mathcal{H}}v_1$. If k is a natural number greater than 1, then $\mathcal{G}^k$ is not isomorphic to $\mathcal{H}^k$ because the roots do not have the same number of successors. On the other hand, $\mathcal{G} \equiv_{\infty}^{k<\Diamond} \mathcal{H}$, since the relation $Z = \{(w, v), (w_1, v_1), (w_2, v_1)\}$ is a $(k <)$-counting bisimulation between $\mathcal{G}$ and $\mathcal{H}$. □

Why do we have these problems in the case $1 < k < \omega$? Our claim is that the relation *to have isomorphic k-expansion* for a natural number $k > 1$ is the familiar Scott-bisimulation. We present a proof of this claim in the case of image finite structures, that is, for structures in which the set of successors of any node is finite. The proof of this result (and a possible extension of it to the general case) has been suggested by Domenico Zambella.

4.3.8. THEOREM. *For all image finite structures $\mathcal{G}$ and $\mathcal{H}$ and for any natural number k it holds:*

$$\mathcal{G}^k \textit{ is isomorphic to } \mathcal{H}^k \Leftrightarrow \mathcal{G}^1 \textit{ is isomorphic to } \mathcal{H}^1.$$

Proof.
We only sketch the proof, leaving the details to the reader.

By considering $\mathcal{G}^1$ and $\mathcal{H}^1$ instead of $\mathcal{G}$ and $\mathcal{H}$ we may suppose without loss of generality that $\mathcal{G}$ and $\mathcal{H}$ are trees. Suppose F is an isomorphism between $\mathcal{G}^k$ and $\mathcal{H}^k$. By composing F with the projection *end* of $\mathcal{H}^k$ into $\mathcal{H} = \mathcal{H}^1$ we obtain a surjective homomorphism $g = end \circ F$ between $\mathcal{G}^k$ and $\mathcal{H}$.

For any node $\sigma \in \mathcal{G}^k$ if $end(\sigma) = w$ and $R(w) = \{w_0, \ldots, w_n\}$ we divide the set of successors of σ into the sets $X_0, \ldots, X_n$ defined as $X_i = \{\sigma m w_i : m < k\}$.

We claim that it is possible to find a bijective function f_w between the set $\{w_0, \ldots, w_n\}$ and the set $g(X_0 \cup \ldots \cup X_n)$ which *respects* g, in the sense that if $f_w(w_i) = y$ then there exists at least an element $\sigma_i \in X_i$ such that $g(\sigma_i) = y$.

To this end, we apply the Marriage Theorem to the sets $A_i = \{y : g^{-1}(y) \cap X_i \neq \emptyset\}$. In this case the theorem says that $\{A_0, \ldots, A_n\}$ has a set of distinct representatives (i.e. there exists a set $\{a_0, \ldots, a_n\}$ such that, $a_i \in A_i$ and $a_i \neq a_j$ for $i \neq j$) iff for all $h \leq n+1$, $0 \leq i_1 < \ldots < i_h < n+1$, the set $A_{i_1} \cup \ldots \cup A_{i_h}$ has cardinality greater or equal to h (Hall's condition). Once we have proved the existence of the set of distinct representatives $\{a_0, \ldots, a_n\}$, we can safely define $f_w(w_i) = a_i$. Let us check the Hall condition. We have:

- by construction, $|g^{-1}(y)| = k$ and $|X_i| = k$, for all i;
- $g^{-1}(A_{i_1} \cup \ldots \cup A_{i_h}) \supseteq X_{i_1} \cup \ldots \cup X_{i_h}$.

Hence $|g^{-1}(A_{i_1} \cup \ldots \cup A_{i_h})| = k|A_{i_1} \cup \ldots \cup A_{i_h}|$ and $|X_{i_1} \cup \ldots \cup X_{i_h}| = hk$. From $g^{-1}(A_{i_1} \cup \ldots \cup A_{i_h}) \supseteq X_{i_1} \cup \ldots \cup X_{i_h}$ it follows that $|A_{i_1} \cup \ldots \cup A_{i_h}| \geq h$.

It is then easy to check that the function

$$\{(r^{\mathcal{G}}, r^{\mathcal{H}})\} \cup \bigcup_{w \in \mathcal{G}} f_w,$$

is an isomorphism between $\mathcal{G}$ and $\mathcal{H}$. □

4.3.3 Other non-well-founded universes.

In this section we sketch a possible application of the previous results to the theory of non-well-founded sets. Let us fix a cardinal $k \geq \omega$. We say that a structure $\mathcal{G}$ is $(k <)-$*extensional* if two different nodes are never $(k <)$-counting bisimilar. We briefly describe the construction (inside a model of ZFC) of a universe $\mathcal{U}^{k<\Diamond}$ satisfying ZFC^- and the following property:

A structure is $(k <)-$extensional iff it is canonical.

We call this property AFA^k. Notice that AFA^1 coincide with the property expressed by the axiom $SAFA$. Moreover, different properties are incompatible, because if k and k' are different cardinals, we prove in Theorem 4.3.10 that there exists a structure $\mathcal{G}$ which is $(k <)-$extensional but not $(k' <)$-extensional. If both AFA^k and $AFA^{k'}$ were satisfied in a universe of sets $\mathcal{U}$, then $\mathcal{G}$ should have an injective decoration in $\mathcal{U}$, being $(k <)-$extensional, but it cannot have one, because it is not $(k' <)$-extensional.

The model $\mathcal{U}^{k<\diamond}$.

In [2] a general construction of non-well-founded universes defined by means of bisimulations is given. We briefly show here that we can perform this construction starting from our $(k <)$-counting bisimulation. The idea is simply to consider the following interpretation $\mathcal{U}^{k<\diamond}$ of the language $\{\in\}$:

- the domain is given by the class of all $(k <)-$extensional structures, modulo isomorphism;
- $\mathcal{G} \in \mathcal{H} \quad \Leftrightarrow \quad$ there exists v such that $r^{\mathcal{H}} R^{\mathcal{H}} v$ and $\mathcal{G}$ is isomorphic $(\mathcal{H}, v)$.

It is then possible to prove that $\mathcal{U}^{k<\diamond}$ is a full system and hence, by Rieger's Theorem 4.1.7, a model of ZFC^-.

More attention must be given to the verification of AFA^k. We can easily verify the property

a structure is $(k <)-$extensional iff it is canonical,

in the following sense. Consider $a \in \mathcal{U}^{k<\diamond}$ with $\mathcal{U}^{k<\diamond} \models$ 'a is a structure' and suppose that the structure $\mathcal{O}_a$ defined as in Lemma 4.1.9 is $(k <)-$extensional. Then we can prove that $\mathcal{U}^{k<\diamond} \models$ 'a is canonical'.

Vice versa, if $\mathcal{U}^{k<\diamond} \models$ 'a is canonical' then we can prove that $\mathcal{O}_a$ is $(k <)-$extensional. In other words, in general we can verify the property of being $(k <)-$extensional only from the point of view of the model of ZFC. This is so because we cannot hope to have a first-order formula expressing the property of being $(k <)-$extensional for all cardinals k: we only have a denumerable set of possible formulae and for different cardinals we have different properties, as we shall see in Corollary 4.3.10.

Hence, in general we cannot speak of the non-well-founded axiom AFA^k, but only of the property AFA^k.

The AFA^k are pairwise incompatible.

Our next task is to prove that the AFA^k are pairwise incompatible properties. First, we prove the following lemma:

4.3.9. LEMMA. *For any cardinal $\mu > 1$ there exists a structure $\mathcal{G}_\mu$ such that for all cardinals h it holds:*

- *if $h > \mu$ then $\mathcal{G}_\mu$ is not $(h <)$-extensional;*
- *if $1 < h < \mu$ then $\mathcal{G}_\mu$ is $(h <)$-extensional;*
- *if $\mu \geq \aleph_0$ then $\mathcal{G}_\mu$ is not $(\mu <)$-extensional;*
- *if $\mu < \aleph_0$ then $\mathcal{G}_\mu$ is $(\mu <)$-extensional.*

Proof.
If $\mu > 1$, denote by $\mathcal{G}_\mu$ the structure with domain

$$\{a, b\} \cup \{\alpha : \alpha \text{ is an ordinal less than } \mu\},$$

and

$$R^{\mathcal{G}} = \{(a,b),(b,b),(0,b)\} \cup \{(a,\alpha),(\alpha,a) : \alpha < \mu\} \cup \{(\alpha,\beta) : \beta < \alpha < \mu\}.$$

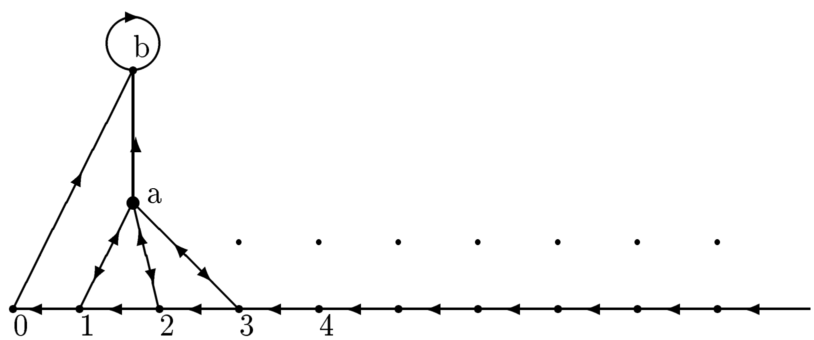

- If $h > \mu$ we prove that $\mathcal{G}_\mu$ is not $(h <)$-extensional. In $\mathcal{G}$ any node has at most h successors and two different nodes are always bisimilar. But a bisimulation between structures in which any node has at most h successors is always an $(h <)$-counting bisimulation, hence $\mathcal{G}_\mu$ is not $(h <)$-extensional.
- If $1 < h < \mu$ we prove that $\mathcal{G}_\mu$ is $(h <)$-extensional by analyzing all possible cases.
  1. We prove that there is no $(h <)$-counting bisimulation between a and b. Suppose there is such a bisimulation Z between a and b. Then for all $\alpha < \mu$ we have $\alpha Z b$, because $a R^{\mathcal{G}} \alpha$ and b is the only successor of b. Consider the set $X = \{w \in R^{\mathcal{G}}(a) : wZb\}$. We have $|X| \geq \mu > h$, while $|\{y \in R^{\mathcal{G}}(b) : \exists x \in X \; xZy\}| = 1$, a contradiction.

2. We prove that there is no $(h <)$-counting bisimulation between $\alpha < \mu$ and b. By contradiction, suppose Z is an $(h <)$-counting bisimulation between $\alpha < \mu$ and b. Since $\alpha R^{\mathcal{G}} a$ we must have aZb and this contradicts point 1.
3. We prove that there is no $(h <)$-counting bisimulation between $\alpha < \mu$ and a (here we use the hypothesis $h > 1$). If $\alpha = 0$ and Z is an $(h <)$-counting bisimulation between 0 and a, then from $aR^{\mathcal{G}}1$ it follows either $1Zb$ or $1Za$ (a, b are the only successors of 0). Since Z is $(h <)$-counting, $1Zb$ is excluded by point 2. Hence $1Za$ and from $aR^{\mathcal{G}}b$ it follows that either $0Zb$ or aZb, but these cases were already excluded.

   If $0 < \alpha < \mu$ and $\alpha Z a$ for an $(h <)$-counting bisimulation Z, then since $aR^{\mathcal{G}}b$ there exists a w such that $\alpha R^{\mathcal{G}} w$ and wZb. But $\alpha R^{\mathcal{G}} w$ only if $w = \beta$ for $\beta < \alpha$, or $w = a$ and we already proved that there is no $(h <)$-counting bisimulation between a and b or β and b.
4. We prove by induction on α that there is no $(h <)$-counting bisimulation between α and γ, for $\mu > \gamma > \alpha$. If $\alpha = 0$ and $0Z\gamma$ for an $(h <)$-counting bisimulation Z then, since $0R^{\mathcal{G}}b$, there exists a w such that $\gamma R^{\mathcal{G}} w$ and bZw. Since $\gamma > 0$, $\gamma R^{\mathcal{G}} w$ iff $w = a$ or $w = \beta$ for $\beta < \gamma$, but bZa, $bZ\beta$ are excluded by the preceding points.

   Suppose that for all $\beta < \alpha$ there is no $(h <)$-counting bisimulation between β and γ, for $\beta < \gamma$. If $\alpha Z \delta$ for $\alpha < \delta$, then from $\delta R^{\mathcal{G}} \alpha$ it follows that there exists w such that $\alpha R^{\mathcal{G}} w$ and $wZ\alpha$. Since $\alpha > 0$, $\alpha R^{\mathcal{G}} w$ iff $w = a$ or $w = \beta$, with $\beta < \alpha$; but $aZ\alpha$ is excluded by point 3, while $\beta Z \alpha$ is excluded by the induction hypothesis.

- We are left to discuss the $(\mu <)$-extensionality of $\mathcal{G}_\mu$. If $\mu \geq \aleph_0$, then we can reason as in the case $h > \mu$: any node has at most μ successors and any bisimulation between the nodes of $\mathcal{G}_\mu$ is a $(\mu <)$-counting bisimulation.

  If $\mu < \aleph_0$ and Z is a $(\mu <)$-counting bisimulation between a and b then we have $|\{w \in R^{\mathcal{G}}(a) : wZb\}| = \mu + 1 > \mu$, while $|\{y \in R^{\mathcal{G}}(b) : \exists x \in X\ xZy\}| = 1$.

  We can exclude all other possible cases as we did in the case of $h < \mu$. □

4.3.10. Corollary. *For $k = 1 or k \geq \omega$, the properties AFA^k and AFA are pairwise incompatible.*

Proof.
First we show that AFA^k is incompatible with AFA^h, for $h > k$, by producing a $(k <)$-extensional structure which is not $(h <)$-extensional.

$k = 1$. Consider the structure $\mathcal{G} = \{a, b\}$ with $R^{\mathcal{G}} = \{(a, a), (a, b), (b, b)\}$ (this example was used in [2] to prove that $SAFA = AFA^1$ is incompatible with

AFA). This is a Scott-extensional structure because $(\mathcal{G}^1, a)$ is not isomorphic to $(\mathcal{G}^1, b)$. On the other hand, the relation $Z = \{(a,a),(a,b),(b,b)\}$ is an $(h <)$-bisimulation if $h \geq 2$, hence $(\mathcal{G}, a)$ is not $(h <)$-extensional for $h \geq 2$.

$k \geq \omega$. Consider the structure $(\mathcal{G}_h, a)$ of Lemma 4.3.9. From the hypothesis it follows that $(\mathcal{G}_h, a)$ is $(k <)$-extensional but not $(h <)$-extensional.

The cases above show also that AFA^k is incompatible with AFA, for any $k \geq \omega$, because the structures defined there are not strongly-extensional. □

4.4 Concluding remarks.

In this chapter we considered formulae of extended modal logics as (descriptions of) non-well-founded sets. We presented a model for the theory $ZFC^- + AFA$ built of infinitary modal logic formulae and we investigated the role of extended modal logics w.r.t. the Scott anti-foundation axiom. In this case it turned out that the logic corresponding to this axiom is Infinitary Graded Logic. The following points need further investigation.

- Can we find a bisimulation-free characterization of the infinitary graded formulae which are description of Scott-sets, as Proposition 4.1.2 does for AFA-sets?

- As we have seen, the AFA-axiom implies a strengthening of the extensionality axiom by saying that *bisimilar sets are equal.* In its modal logic version this axiom becomes: *two sets that verify the same infinitary modal formulae are equal.* Aczel proved that we can restrict the notion of bisimulation, ending with new non-well-founded axioms that are consistent with ZFC^-. In this respect, by investigating the Scott axiom we showed how we could ask that *two sets that verify the same infinitary graded modal formulae are equal.* Within this theory, we *identify* less objects than in $ZFC^- + AFA$, because two sets to be equal have to verify the same formulae of a logic which strictly contains Infinitary Modal Logic. What about considering logics that go in the opposite direction (that is, that will identify more objects)? As an example, consider the Modal μ-Calculus. If two sets are bisimilar they verify the same formulae of this logic, but the converse is not true. Can we prove that to identify sets having the same Modal μ-Calculus behavior is consistent with ZFC^-?

- As we have seen, the class of minimal infinitary modal formulae is a model of $ZFC^- + AFA$. Can we use this result to transfer results from modal logic to non-well-founded sets? This seems not to be the case of interpolation, even if we consider sets with urelements and languages with propositional constants (see the discussion about urelements at the end of Section 2.6). Being minimal formulae, sets can never be entailed by other sets. On the other hand, we can give a set-meaning to generic infinitary modal formulae (non necessarily minimal) as a collection of sets or as a partial description of a set. What is the meaning of interpolation in this context?

- In Lemma 4.2.3 and Lemma 4.3.1 we proved that structures can be described by extended modal formulae modulo counting bisimulations. Is there a general pattern here? For what kind of structure-simulations and languages is this description possible? One possibility here is to consider the bisimulations described by means of *pebble games* as in [5].

Chapter 5

Interpolation for graded modal logics.

In this chapter we consider the problem of interpolation in the family of graded modal logics. As we already pointed out in the preliminaries, the behavior of these logics w.r.t. interpolation is not uniform: the finitary logic $\mathcal{L}^{\Diamond<n}$, which contains only a finite number of graded operators, lacks interpolation, while the logic $\mathcal{L}^{\Diamond<\aleph_0}$ has Craig interpolation. Here we consider the extensions of these logics that allow countable conjunctions, $\mathcal{L}_{\omega_1}^{\Diamond<n}$ and $\mathcal{L}_{\omega_1}^{\Diamond<\aleph_0}$ and prove elementary interpolation for the former and Craig interpolation for the latter. We also give preservation theorems for these logics w.r.t. Countable Infinitary First Order Logic $\mathcal{L}_{\omega_1}$. We prove these results for the countable case and not for the full infinitary logics $\mathcal{L}_{\infty}^{\Diamond<n}$ and $\mathcal{L}_{\infty}^{\Diamond<\aleph_0}$ for the sake of simplicity in the proofs; in the final section, however, we give a proof of interpolation for the logic $\mathcal{L}_{\infty}^{grad}$ which we already encountered in Chapter 4 in connection with the Scott anti-foundation axiom, showing that the restriction to the countable case is avoidable.

The main technique used in this chapter is the use of *consistency property modulo bisimulation*, that was first introduced in [11] in the context of Infinitary Logic. It consists of a generalization of the notion of consistency property for First Order and Countable Infinitary Logic. Let us recall briefly the use of this technique in proving interpolation. Suppose a pair (ϕ, ψ) does not have an interpolant. We sketch a construction of a model for $\phi \wedge \neg\psi$, showing that $\not\models \phi \to \psi$, as follows. Starting from the two sets $s_1 = \{\phi\}$, $s_2 = \{\neg\psi\}$ we can try to *simplify* the formulae in these two sets, towards atomic sentences, being careful to preserve the non existence of an interpolant between the formulae $\wedge s_1$ and $\neg \wedge s_2$. For example, if $\exists x\theta$ belongs to s_1, we can add a constant c to the language, taken from a countable set of C of new constants and consider the set $s_1' = s_1 \cup \{\theta(c)\}$; if $\wedge s_1$ and $\neg \wedge s_2$ do not have an interpolant, then $\wedge s_1'$ and $\neg \wedge s_2$ do not have one as well. We can then consider the two sets s_1', s_2 and reduce another formula in their union and so forth: by making appropriate reductions we can build an increasing sequence of sets of formulae and in ω-steps we reach a set s^* whose sentences are either atomic or have been reduced in previous steps. Then the atomic sentences of s^* can be used to build a model from the set of constants C which verifies all sentences in s^*. In particular, this model verifies $\phi \wedge \neg\psi$.

Let us see with an example what happens to this construction if the logic does not have interpolation. Consider the two $\mathcal{L}^{\diamond<3}$ sentences

$$\phi(p) \equiv (\langle a\rangle_2 p \wedge \langle a\rangle \neg p),$$

$$\psi(q) \equiv (\langle a\rangle_2 q \vee \langle a\rangle_2 \neg q).$$

We have $\phi(p) \models \psi(q)$: if $\mathcal{M}$ satisfies $\phi(p)$, then the root $r^{\mathcal{M}}$ has at least three a-successors and hence it satisfies $\psi(q)$. We can also show that ϕ and ψ do not have an interpolant. Consider the following two structures $\mathcal{M}$ and $\mathcal{N}$:

$$\mathcal{M} = \{r^{\mathcal{M}}, w_1, w_2, w_3\}, r^{\mathcal{M}} R_a w_i \text{ for } i = 1, 2, 3,\ w_i \models p \text{ for } i = 1, 2\ ;$$

$$\mathcal{N} = \{r^{\mathcal{N}}, v_1, v_2\}, r^{\mathcal{N}} R_a v_i \text{ for } i = 1, 2 \text{ and } v_1 \models q.$$

It is easy to see that $\mathcal{M}$ and $\mathcal{N}$ are $(< 3)-$counting bisimilar w.r.t. the language $\{a\}$ and that $\mathcal{M} \models \phi$ and $\mathcal{N} \models \neg\psi$. Suppose by contradiction that ϕ and ψ have an interpolant $\theta \in \mathcal{L}^{\diamond<3}$. Then $\mathcal{M} \models \theta$ and $\mathcal{N} \models \theta$ as well, being $(< 3)-$counting bisimilar to $\mathcal{M}$. But since θ is an interpolant of ϕ, ψ, it follows that $\mathcal{N} \models \psi$, a contradiction.

Hence we have a pair ϕ, ψ without a garded interpolant, but the previous construction of a model for $\phi \wedge \neg\psi$ will fail, because $\phi(p) \models \psi(q)$. By simplifying $\{\phi\}$, after a few steps three new constants are introduced, and the set $\{\phi\}$ is extended with the sentences saying that these constants are distinct successors of the root. But then when we simplify the set $\{\neg\psi\}$ we have to decide which constants have to go in q and which not and this generates a contradiction at the atomic level. Hence, we cannot hope to simplify the sets $\{\phi\}, \{\neg\psi\}$ and at the same time preserve the absence of a graded interpolant. But if we reduce our expectations and we give up the idea of building a model for $\phi \wedge \neg\psi$, then we can still use the method of consistency property to characterize pairs without interpolants in $\mathcal{L}^{\diamond<3}$. What we shall do is to build *two different* $(< 3)-$counting bisimilar structures: the first, $\mathcal{M}$, which verifies ϕ and the second, $\mathcal{N}$, which verifies $\neg\psi$, thereby showing that ϕ does not entail ψ along $\mathcal{L}^{\diamond<3}$-equivalence. We start again from $s_1 = \{\phi\}$ and $s_2 = \{\neg\psi\}$, but this time we use a set C of countable constants to simplify s_1 and a different set D to simplify s_2; at the same time we create *links* between C and D, in such a way to obtain after ω-steps the desired models $\mathcal{M}$ and $\mathcal{N}$ out of C and D, respectively, together with a $(< 3)-$counting bisimulation described by the links we created.

This is just the general idea of a proof of interpolation in the family of graded logics using consistency property modulo bisimulation. For the sake of simplicity, we divide the proof in two parts. The first deals with a general notion of consistency property modulo bisimulation and in the second we apply this notion to give a proof of interpolation. The whole chapter is organized as follows. In

Section 5.1 we give the notion of a consistency property modulo bisimulation and prove a Model Existence Lemma with respect to it. In Section 5.2 we use this notion to prove elementary interpolation for the logics $\mathcal{L}_{\omega_1}^{\diamond<k}$, for $k \geq \aleph_0$ and we provide preservation theorems for these logics with respect to First Order Logic. In Section 5.3 we consider the Infinitary Logic of Finite Graded Modalities $\mathcal{L}_{\omega_1}^{\diamond<\aleph_0}$: in this case we improve the result of the previous section by showing that this logic enjoys Craig interpolation (this result was already proved in the case of $\mathcal{L}^{\diamond<\aleph_0}$ by Andréka in [1]). Finally, Section 5.4 proves Craig interpolation for the full logic of graded modalities $\mathcal{L}_{\infty}^{grad}$.

5.1 Consistency properties for graded modalities

In the classical case, a *consistency property* is a set S of countable sets of sentences in a language $\mathcal{L}_C$ containing constants from a countable set C, which satisfy a set $\mathcal{P}$ of natural closure properties, such as if $s \in S$ and $\exists x\phi(x) \in s$ then there exists a $c \in C$ with $s \cup \{\phi(c)\} \in S$. These properties are extrapolated form the properties of the set

$$S_0 = \{s \subseteq (\mathcal{L}_C)_{\omega_1} : s \text{ is countable}, \exists \mathcal{M}(\mathcal{M} \models s \wedge \mathcal{M} = \{c^{\mathcal{M}} : c \in C\})\}.$$

Then a *Model Existence Lemma* shows that the satisfiability of the properties $\mathcal{P}$ in S is a sufficient condition for the existence of a model for any element s_0 in S. To build such a model, we first close the set s_0 with respect to $\mathcal{P}$, that is, we build a sequence $s_0, s_1, \ldots$ in such a way that, for example, if $\exists x\phi(x) \in s_i$ for some i then there exists a $j > i$ and a $c \in C$ such that $\{\phi(c)\} \in s_j$; then we take all atomic sentences $R(c_1, \ldots, c_n)$ in $\bigcup_{i\in\omega} s_i$ and use them to describe an interpretation $\mathcal{M}$ with domain equal to C (or equivalence classes in C, if the language contains equality).

In this section, we generalize this idea to fit the logics $\mathcal{L}^{\diamond<k}$, for $k = \aleph_0$ or a natural number and for their countable infinitary extensions. We introduce the method of consistency property modulo bisimulation, which will be used to prove, besides interpolation, a preservation theorem of $\mathcal{L}_{\omega_1}^{\diamond<k}$ ($\mathcal{L}^{\diamond<k}$) w.r.t. countable Infinitary Logic $\mathcal{L}_{\omega_1}$ (First Order Logic, respectively). This method is introduced for the modal case in [11] and it is used to prove an interpolation theorem for $\mathcal{L}_{\omega_1}^{\diamond}$ in [64].

Let k be a natural number or $\aleph_0$. We consider two languages $\mathcal{L}, \mathcal{L}'$ and their extensions $\mathcal{L}_C$ and $\mathcal{L}'_D$ relative to two countable sets of constants C, D with $\{r\} = C \cap D$. We think now about a model $\mathcal{M}$ for $\mathcal{L}_C$ with domain $\{c^{\mathcal{M}} : c \in \mathcal{M}\}$ and a model $\mathcal{N}$ for $\mathcal{L}'_D$ with domain $\{d^{\mathcal{N}} : d \in D\}$, such that $\mathcal{M}$ and $\mathcal{N}$ are $(< k)$-counting bisimilar w.r.t. the language $\mathcal{L} \cap \mathcal{L}'$. Our *motivating* example consists here in the set S of triples (Z, Σ, Δ) for which there exists $\mathcal{M}$, $\mathcal{N}$ as above and such that:

1. $\Sigma \subseteq (\mathcal{L}_C)_{\omega_1}$, $\Delta \subseteq (\mathcal{L}_D)_{\omega_1}$;
2. $\mathcal{M} \models \Sigma$;
3. $\mathcal{N} \models \Delta$;
4. Z is a graded bisimulation w.r.t. $\mathcal{L} \cap \mathcal{L}'$ between $\mathcal{M}$ and $\mathcal{N}$.

Again, some natural properties are extrapolated from this example: for example, if $p(c) \in \Sigma$, for $p \in \mathcal{L} \cap \mathcal{L}'$ and $(c, d) \in Z$, then the triple $(Z, \Sigma, \Delta \cup \{p(d)\})$ belongs to S. This example leads to the notion of consistency properties modulo bisimulation and a Model Existence Lemma proves that any element (Z, Σ, Δ), which belongs to such a set, is *satisfiable* in the following sense: there are $\mathcal{M}$, $\mathcal{N}$ whose domains are given by (equivalence classes in) C, D, respectively, such that

1. $\mathcal{M} \models \Sigma$;
2. $\mathcal{N} \models \Delta$;
3. Z describes a graded bisimulation w.r.t. $\mathcal{L} \cap \mathcal{L}'$ between $\mathcal{M}$ and $\mathcal{N}$.

We are now ready for details. Let C, D be countable sets of constants symbols which are disjoint, except for the root constant $r \in C \cap D$. Without loss of generality we suppose that the formulae of $(\mathcal{L}_C)_{\omega_1}$ and $(\mathcal{L}'_D)_{\omega_1}$ are constructed from atoms and negated atoms using $\forall, \exists, \bigwedge_{n\in\omega}, \bigvee_{n\in\omega}$. If ϕ is a formula, we denote by $\neg\phi$ the formula obtained by 'moving the negation inside' (e.g. $\neg(p(x) \wedge q(x)) \equiv \neg p(x) \vee \neg q(x)$).

5.1.1. Definition. Let k be a natural number or $\aleph_0$. A *consistency property modulo* $(< k)-$*counting bisimulation* is a set S of triples (A, Σ, Δ) where

- $(r, r) \in A \subseteq C \times D$,
- Σ is a countable set of $(\mathcal{L}_C)_{\omega_1}$ sentences, and Δ is a countable set of $(\mathcal{L}'_D)_{\omega_1}$ sentences, which satisfy the following *closure properties.*

Left consistency properties:

(c_1) If ϕ is an atomic sentence of $(\mathcal{L}_C)_{\omega_1}$, then either $(A, \Sigma \cup \{\phi\}, \Delta) \in S$ or $(A, \Sigma \cup \{\neg\phi\}, \Delta) \in S$;

(c_2) If ϕ is an atomic sentence of $(\mathcal{L}_C)_{\omega_1}$, either $\phi \notin \Sigma$ or $\neg\phi \notin \Sigma$;

(c_3) if $\bigwedge \Phi \in \Sigma$ then for all $\phi \in \Phi$ $(A, \Sigma \cup \{\phi\}, \Delta) \in S$;

(c_4) if $\forall x \phi(x) \in \Sigma$, then for all $c \in C$ $(A, \Sigma \cup \{\phi(c/x)\}, \Delta) \in S$;

(c_5) if $\bigvee \Phi \in \Sigma$, then there is a $\phi \in \Phi$ such that $(A, \Sigma \cup \{\phi\}, \Delta) \in S$;

(c_6) if $\exists x\phi \in \Sigma$, then there is a $c \in C$ such that $(A, \Sigma \cup \{\phi(c/x)\}, \Delta) \in S$.

Right consistency properties:

The same as above, with Δ, $\mathcal{L}'$, and D instead of Σ, $\mathcal{L}$, and C.

Left Bisimulation properties:

(b_1) if $(c, d) \in A$, $p \in \mathcal{L} \cap \mathcal{L}'$, and $p(c) \in \Sigma$ then $(A, \Sigma, \Delta \cup \{p(d)\}) \in S$;

(b_2) for all $h < k$, if $(c, d) \in A$ and $c_0, \ldots, c_{h-1}$ are constants in C such that

$$\{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\} \subseteq \Sigma,$$

then there are constants $d_0, \ldots, d_{h-1}$ in D such that

$$(A \cup \{(c_i, d_i) : i < h\}, \Sigma, \Delta \cup \{R_a(d, d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\}) \in S.$$

Right Bisimulation properties

The same as above exchanging the role of Σ, C with Δ, D.

Left Equality properties:

($=_1$) If $c = c' \in \Sigma$ then $(A, \Sigma \cup \{c' = c\}, \Delta) \in S$;

($=_2$) if $c = c' \in \Sigma$ and $\phi(c') \in \Sigma$ then $(A, \Sigma \cup \{\phi(c)\}, \Delta) \in S$;

($=_3$) for all $c \in C$, $(A, \Sigma \cup \{c = c\}, \Delta) \in S$;

($=_4$) If $c, c' \in C$, $d \in D$, $(c, d) \in A$, and $(c = c') \in \Sigma$, then the triple $(A \cup \{(c', d)\}, \Sigma, \Delta)$ belongs to S.

Right Equality properties:

The same as above exchanging the role of Σ and Δ. □

Notice how these properties are extrapolated from our first example, in which S is the set of triples (Z, Σ, Δ) such that there exist two ($< k$)-counting bisimilar models $\mathcal{M}$, $\mathcal{N}$ with $\mathcal{M} \models \Sigma$, $\mathcal{N} \models \Delta$, and $D^{\mathcal{M}} = \{c^{\mathcal{M}} : c \in C\}, D^{\mathcal{N}} = \{d^{N} : d \in D\}$. On the other hand, we can prove that this example of a consistency property is *paradigmatic*, in the sense given by the following theorem. We give the proof of this theorem in full details and this will keep us busy until the end of the paragraph. In the next sections we will not be so pedantic and refer to this proof to fill the gaps in the exposition.

5.1.2. LEMMA. *If S is a consistency property modulo $(< k)$-counting bisimulation and $(A, \Sigma, \Delta) \in S$ then there are:*

- *a countable model $\mathcal{M}$ of $\mathcal{L}_C$ such that $\mathcal{M} \models \Sigma$,*
- *a countable model $\mathcal{N}$ of $\mathcal{L}'_D$ such that $\mathcal{N} \models \Delta$,*
- *a $(< k)$-counting bisimulation Z w.r.t. $\mathcal{L} \cap \mathcal{L}'$ between $\mathcal{M}$ and $\mathcal{N}$ such that for all $(c, d) \in A$, the pair $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$.*

Proof.
We say that a consistency property S is *closed* if whenever $(A, \Sigma, \Delta) \in S$ and $A' \subseteq A$, $\Sigma' \subseteq \Sigma$ and $\Delta' \subseteq \Delta$ then $(A', \Sigma', \Delta') \in S$. Notice that if S is a consistency property modulo $(< k)$-counting bisimulation then

$$S' = \{(A', \Sigma', \Delta') : \exists (A, \Sigma, \Delta) \in S \text{ with } (r, r) \in A' \subseteq A, \Sigma' \subseteq \Sigma, \Delta' \subseteq \Delta\}$$

is also a consistency property modulo $(< k)$-counting bisimulation. Just to give an example, we check that property (c_1) holds for $(A', \Sigma', \Delta') \in S'$. By definition of S', there exists $(A, \Sigma, \Delta) \in S$ such that $A' \subseteq A$, $\Sigma' \subseteq \Sigma$, $\Delta' \subseteq \Delta$. Since (A, Σ, Δ) satisfies c_1), if ϕ is an atomic sentence of $(\mathcal{L}_C)_{\omega_1}$, then either $(A, \Sigma \cup \{\phi\}, \Delta) \in S$ or $(A, \Sigma \cup \{\neg\phi\}, \Delta) \in S$. In the first case, we have $(A', \Sigma' \cup \{\phi\}, \Delta') \in S'$, while in the second $(A', \Sigma' \cup \{\neg\phi\}, \Delta') \in S'$.

Hence, we may suppose without loss of generality that S is closed.

Let Y be the least set of sentences of $(\mathcal{L}_C)_{\omega_1}$ and $(\mathcal{L}'_D)_{\omega_1}$ such that:

$\Sigma \subseteq Y$, $\Delta \subseteq Y$;

Y is closed under C-substitution and D-substitution instances of subformulae;

if $c, c' \in C$ and $d, d' \in D$ then $c = c'$, $c \neq c'$, $d = d'$, and $d \neq d'$ belong to Y;

For all $c, c' \in C, d, d' \in D$ if $\phi(c) \in Y$ then $\phi(c') \in Y$, if $\psi(d) \in Y$ then $\psi(d') \in Y$.

Let T be a set of tuples containing:

L_1) all tuples (c, c', d), with $c, c' \in C$, $d \in D$;

R_1) all tuples (d, d', c), with $d, d' \in D$, $c \in C$;

L_2) all tuples $(a, c, d, c_0, \ldots, c_{h-1})$, where $h < k$, a is an action in $\mathcal{L} \cap \mathcal{L}'$, $c, c_0, \ldots c_{h-1}$ in C, and d in D;

R_2) all tuples $(a, c, d, d_0, \ldots, d_{h-1})$, where $h < k$, a is an action in $\mathcal{L} \cap \mathcal{L}'$, c in C, and $d, d_0, \ldots, d_{h-1}$ in D;

L_3) all tuples (c, d, p), where $c \in C$, $d \in D$, and p is unary predicate symbol of $\mathcal{L} \cap \mathcal{L}'$;

R_3) all tuples (d, c, p), where $c \in C$, $d \in D$, and p is unary predicate symbol of $\mathcal{L} \cap \mathcal{L}'$.

These tuples correspond to the various requirements that a $(< k)-$counting bisimulation Z between models $\mathcal{M}, \mathcal{N}$ must satisfy, on the left (L) and on the right (R). We will construct our models and our bisimulation step by step, following an enumeration that contains these tuples. Whenever a tuple of type L_1 (R_1) is met, we continue the construction in such a way that if $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$ and $c^{\mathcal{M}} = c'^{\mathcal{M}}$ then $(c'^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$ (if $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$ and $d^{\mathcal{N}} = d'^{\mathcal{N}}$ then $(c^{\mathcal{M}}, d'^{\mathcal{N}}) \in Z$, respectively). The second pair corresponds to the zig-zag requirement of a $(< k)-$counting bisimulation and the third pair to preservation of propositions: the tuple (c, d, p) is used to remind us that if $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$ and $\mathcal{M} \models p(c)$ then $\mathcal{N} \models p(d)$ must hold, while the tuple (d, c, p) means that if $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$ and $\mathcal{M} \models p(d)$ then $\mathcal{N} \models p(c)$ must be true.

Consider an enumeration $G_0, G_1, \ldots, G_n, \ldots$ of all sentences in Y and of all tuples in T.

We shall construct a sequence $t_0 = (A_0, \Sigma_0, \Delta_0), t_1 = (A_1, \Sigma_1, \Delta_1), \ldots,$ of triples in S as follows.

$(A_0, \Sigma_0, \Delta_0)$ is the given (A, Σ, Δ). If $(A_n, \Sigma_n, \Delta_n)$ is already defined, then $(A_{n+1}, \Sigma_{n+1}, \Delta_{n+1})$ is constructed using the definition of a consistency property, according to the element G_n.

More precisely, if G_n is a formula in $(\mathcal{L}_C)_{\omega_1}$ then:

- if $(A_n, \Sigma_n \cup \{G_n\}, \Delta) \notin S$, then $t_{n+1} = t_n$;
- if G_n is an atomic formula, or $G_n = \bigwedge \Phi$, or $G_n = \forall x \phi$ and $(A_n, \Sigma_n \cup \{G_n\}, \Delta) \in S$, then $t_{n+1} = (A_n, \Sigma_n \cup \{G_n\}, \Delta)$;
- if $G_n = \bigvee \Phi$ and $(A_n, \Sigma_n \cup \{G_n\}, \Delta) \in S$ then by (c_5) there is a $\phi \in \Phi$ such that $(A_n, \Sigma_n \cup \{G_n\} \cup \{\phi\}, \Delta)$ belongs to S. We take this triple as t_{n+1};
- if $G_n = \exists x \phi$ and $(A_n, \Sigma_n \cup \{G_n\}, \Delta) \in S$ then by (c_6) there is a $c \in C$ such that $(A_n, \Sigma_n \cup \{G_n\} \cup \{\phi(c)\}, \Delta)$ belongs to S. We take this triple as t_{n+1};

The same construction is performed if G_n is a formula in $(\mathcal{L}'_D)_{\omega_1}$, using the right consistency properties of S.

On the other hand, if G_n is a tuple then:

if $G_n = (c, c', d)$ with $c, c' \in C$, $d \in D$, and the triple

$$(A_n \cup \{(c, d)\}, \Sigma_n \cup \{c = c'\}, \Delta_n)$$

belongs to S, then by $(=_4)$ the triple

$$(A_n \cup \{(c,d),(c',d)\}, \Sigma_n \cup \{c = c'\}, \Delta_n)$$

belongs to S as well and we choose t_{n+1} equal to this triple.

If $(A_n \cup \{(c,d)\}, \Sigma_n \cup \{c = c'\}, \Delta_n) \notin S$, choose $t_{n+1} = t_n$.

If $G_n = (c,d,p)$ with $c \in C, d \in D$, and p is a unary predicate of $\mathcal{L} \cap \mathcal{L}'$, then if $(A_n \cup \{(c,d)\}, \Sigma_n \cup \{p(c)\}, \Delta_n)$ belongs to S, by the left bisimulation property (b_1) we can take

$$t_{n+1} = (A_n \cup \{(c,d)\}, \Sigma_n \cup \{p(c)\}, \Delta_n \cup \{p(d)\}).$$

If $(A_n \cup \{(c,d)\}, \Sigma_n \cup \{p(c)\}, \Delta_n)$ does not belongs to S, take $t_{n+1} = t_n$.

if $G_n = (a,c,d,c_0,\ldots,c_{h-1})$ where $h < k$, a is an action in $\mathcal{L} \cap \mathcal{L}'$, $c, c_0, \ldots, c_{h-1}$ belong to C, $d \in D$, and the triple $(A_n^*, \Sigma_n^*, \Delta_n)$ belongs to S, where $A_n^* = A_n \cup \{(c,d)\}, \Sigma_n^* = \Sigma_n \cup \{R_a(c,c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\}$, then by the left bisimulation property (b_2) there are $d_0, \ldots, d_{h-1}$ such that the triple

$$(A_n^*, \Sigma_n^*, \Delta_n \cup \{R_a(d,d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\})$$

belongs to S. We take this triple as t_{n+1}. If $(A_n^*, \Sigma_n^*, \Delta_n)$ does not belong to S, take $t_{n+1} = t_n$.

The same construction is performed on the right, when G_n is a tuple of type (d,d',c) or (d,c,p) or $(c,d,d_0,\ldots,d_{h-1})$.

Our next task is to define the models $\mathcal{M}$, $\mathcal{N}$ and the bisimulation Z between them.

Define

$$A_\omega = \bigcup_{n<\omega} A_n, \qquad \Sigma_\omega = \bigcup_{n<\omega} \Sigma_n, \qquad \Delta_\omega = \bigcup_{n<\omega} \Delta_n.$$

5.1.3. Remark. a) If G_n is a sentence in Σ_ω, then $(A_n, \Sigma_n \cup \{G_n\}, \Delta_n)$ belongs to S and $G_n \in \Sigma_{n+1}$.

b) If G_n is a sentence and $r \geq n$ is such that $(A_r, \Sigma_r \cup \{G_n\}, \Delta_r)$ belongs to S, then $G_n \in \Sigma_{n+1}$.

c) If G_n is a sequence (c,d,p) with $(c,d) \in A_\omega$, $p \in \mathcal{L} \cap \mathcal{L}'$, and $p(c) \in \Sigma_\omega$ then $p(d) \in \Delta_{n+1}$;

d) If G_n is a sequence $(a, c, d, c_0 \ldots, c_{h-1})$ where $h < k$, a is an action in $\mathcal{L} \cap \mathcal{L}'$, $(c, d) \in A_\omega$ and $\{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\} \subseteq \Sigma_\omega$, then there are $d_0, \ldots, d_{h-1}$ in D such that $\{R_a(d, d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\} \subseteq \Delta_{n+1}$ and $\{(c_i, d_i) : i < h\} \subseteq A_{n+1}$.

The same properties hold 'on the right'. □

We just prove *a*) and *d*).

a) If $G_n \in \Sigma_\omega$, let $r \geq n$ be such that $G_n \in \Sigma_r$. Then $\Sigma_n \cup \{G_n\} \subseteq \Sigma_r$ and since S is closed the triple $(A_n, \Sigma_n \cup \{G_n\}, \Delta_n)$ belongs to S. By construction, $G_n \in \Sigma_{n+1}$.

d) Let $r \geq n$ be such that $(c, d) \in A_r$ and $\{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\} \subseteq \Sigma_r$. Then $A_n \cup \{(c, d)\} \subseteq A_r$, $\Sigma_n \cup \{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\} \subseteq \Sigma_r$, and, since S is closed, the triple

$$(A_n \cup \{(c, d)\}, \Sigma_n \cup \{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\}, \Delta_n)$$

belongs to S. By construction of t_{n+1} there are $d_0, \ldots, d_{h-1}$ in D such that $\{R_a(d, d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\} \subseteq \Delta_{n+1}$ and $\{(c_i, d_i) : i < h\} \subseteq A_{n+1}$.

Construction of the models $\mathcal{M}$, $\mathcal{N}$.

To define the model $\mathcal{M}$, let E be the following relation on C:

$$cEc' \qquad \text{iff} \;\; c = c' \in \Sigma_\omega.$$

From the equality properties of S it follows that E is an equivalence relation. For example, suppose cEc' and $c'Ec''$. Let n be such that G_n is $c = c''$, and let $r \geq n$ be such that $c = c'$ and $c' = c''$ belong to Σ_r. By applying $(=)_2$ with $\phi(c')$ equal to $c' = c''$ we obtain that the triple

$$(A_r, \Sigma_r \cup \{c = c''), \Delta_r)$$

belongs to S. By Remark 5.1.3.b we obtain $c = c'' \in \Sigma_{n+1}$. Hence, cEc''. Reflexivity and symmetry of E are proved in the same way using properties $(=)_1, (=)_3$.

For $c \in C$, let $[c]$ be its equivalence class w.r.t. E. The domain of $\mathcal{M}$ is the set $\{[c] : c \in C\}$ and each constant $c \in C$ is interpreted in $\mathcal{M}$ as $[c]$. By the equality property $(=)_2$, if $R_a(c_1, c_2) \in \Sigma_\omega$ and c_1Ec_1', c_2Ec_2', then $R_a(c_1', c_2') \in \Sigma_\omega$ and the definition

$$R_a^{\mathcal{M}}([c_1], [c_2]) \qquad \text{iff} \qquad R_a(c_1, c_2) \in \Sigma_\omega$$

does not depend on the representatives of the classes $[c_1], [c_2]$.

The same is true if we define

$$p^{\mathcal{M}}([c]) \qquad \text{iff} \qquad p(c) \in \Sigma_\omega,$$

for all unary predicate $p \in \mathcal{L}$.

Finally, let $[r]$ be the root $r^{\mathcal{M}}$ of $\mathcal{M}$.

We now show that all formulae of Σ_ω are true in $\mathcal{M}$. If ϕ is an atomic formula, then by definition we have

$$\phi \in \Sigma_\omega \qquad \text{iff} \qquad \mathcal{M} \models \phi.$$

Using the consistency properties $(c_1), (c_2)$ we can easily extend this fact to literals. For example, if ϕ is $c_i \neq c_j$ and $\phi \in \Sigma_\omega$, then $\mathcal{M} \models \phi$, because otherwise $c_i = c_j$ would also belong to Σ_ω, contradicting (c_2). Vice versa, if $\mathcal{M} \models \phi$, G_n is $c_i = c_j$, and G_m is $c_i \neq c_j$, consider $r \geq m, n$. The triple $(A_r, \Sigma_r \cup \{G_n\}, \Delta_r)$ does not belong to S otherwise by Remark 5.1.3.b we would have $G_n \in \Sigma_{n+1}$ and $\mathcal{M} \models G_n$, which is not. By (c_1) the triple $(A_r, \Sigma_r \cup \{G_m\}, \Delta_r)$ belongs to S and by Remark 5.1.3.b it follows that $G_m \in \Sigma_{m+1} \subseteq \Sigma_\omega$.

In case of more complex sentences, we can prove by induction on the complexity of a sentence ϕ *belonging to* Y that

$$\phi \in \Sigma_\omega \qquad \text{implies} \qquad \mathcal{M} \models \phi.$$

- If ϕ is a literal, we have already proved a stronger result.
- If $\phi = \forall x \psi \in \Sigma_\omega$ and $c \in C$, let $G_n = \psi(c)$ ($\psi(c)$ is in Y because Y is closed under C-substitution instances of subformulae). If $r \geq n$ is such that $\forall x \psi \in \Sigma_r$, by (c_4) we have that the triple $(A_r, \Sigma_r \cup \{\psi(c)\}, \Delta_r)$ belongs to S. From Remark 5.1.3.b it follows that $\psi(c) \in \Sigma_{n+1} \subseteq \Sigma_\omega$ and by induction $\mathcal{M} \models \psi(c)$. Since this holds for all $c \in C$ and any element of $\mathcal{M}$ is denoted by an element in C, it follows that $\mathcal{M} \models \forall x \psi$.
- If $\phi = G_n = \exists x \psi$ belongs to Σ_ω, then by Remark 5.1.3.a we have that the triple $(A_n, \Sigma_n \cup \{\exists x \psi\}, \Delta_n)$ belongs to S. By construction of t_{n+1}, there is a $c \in C$ such that $\psi(c) \in \Sigma_{n+1} \subseteq \Sigma_\omega$ and by induction $\mathcal{M} \models \psi(c)$. Hence, $\mathcal{M} \models \exists x \psi$.
- The remaining cases are similar and left to the reader.

In this way we proved that $\mathcal{M}$ is a model for all sentences in Σ_ω.

In a similar way we can prove that there is a model $\mathcal{N}$, whose definition is based on the equivalence relation

$$d \sim d' \qquad \text{iff} \qquad d = d' \in \Delta_\omega,$$

which is a model for all sentences in Δ_ω.

Construction of the $(<k)-$counting bisimulation Z between $\mathcal{M}$ and $\mathcal{N}$.

Let

$$[c]Z[d] \qquad \text{iff} \;\; (c,d) \in A_\omega.$$

We prove that the definition of Z does not depend on the representatives of $[c],[d]$. By the equality property $(=_4)$ we have that if $(c,d) \in A_\omega$ and $c = c' \in \Sigma_\omega$, then $(c',d) \in A_\omega$: if $G_n = (c,c',d)$ and $r \geq n$ is such that $(c,d) \in A_r$, $c = c' \in \Sigma_r$, then, since S is closed, the triple $(A_n \cup \{(c,d)\}, \Sigma_n \cup \{c = c'\}, \Delta_n)$ belongs to S. By definition of t_{n+1} we have $(c',d) \in A_{n+1} \subseteq A_\omega$.

In the same way we prove that if $(c,d) \in A_\omega$ and $d = d' \in \Delta_\omega$, then $(c,d') \in A_\omega$.

Finally, we show that Z is a graded bisimulation w.r.t. $\mathcal{L} \cap \mathcal{L}'$ between $\mathcal{M}$ and $\mathcal{N}$.

- From $(r,r) \in A_0 \subseteq A_\omega$ we have $r^{\mathcal{M}} Z r^{\mathcal{N}}$.
- If $[c]Z[d]$ and $\mathcal{M} \models p(c)$ for $p \in \mathcal{L} \cap \mathcal{L}'$, then $(c,d) \in A_\omega$, $p(c) \in \Sigma_\omega$ and by Remark 5.1.3.c it follows that $p(d) \in \Delta_{n+1}$. Hence by definition we also have $\mathcal{N} \models p(d)$.
- If $[c]Z[d]$, $[c_0], \ldots, [c_{h-1}]$ with $h < k$ are such that $\mathcal{M} \models R_a(c, c_i)$, for all $i < h$ and a in $\mathcal{L} \cap \mathcal{L}'$, and $\mathcal{M} \models c_l \neq c_j$ for $l < j < h$, then $R_a(c,c_i) \in \Sigma_\omega$ and $(c_l \neq c_j) \in \Sigma_\omega$ (remember that for literals we have $\phi \in \Sigma_\omega \Leftrightarrow \mathcal{M} \models \phi$). By Remark 5.1.3 d there are $d_0, \ldots, d_{h-1}$ in D such that $\{R_a(d,d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\} \subseteq \Delta_\omega$ and $\{(c_i,d_i) : i < h\} \subseteq A_\omega$. It follows that $\mathcal{N} \models R_a(d,d_i)$, $\mathcal{N} \models d_l \neq d_j$, and $[c_i]Z[d_i]$ for $i < h$, $l < j < h$, as required to a $(<k)-$counting bisimulation.

Hence, if $(A, \Sigma, \Delta) \in S$ we showed how to construct two models $\mathcal{M}, \mathcal{N}$ with $\mathcal{M} \models \Sigma$, $\mathcal{N} \models \Delta$, and how to *expand* A to a $(<k)$-counting bisimulation Z between $\mathcal{M}$ and $\mathcal{N}$. □

5.2 Preservation and elementary interpolation

In this section we characterize the graded (infinitary) logic $\mathcal{L}^{\diamond<k}_{\omega_1}$, for $k \leq \aleph_0$, as the fragment of Infinitary Logic $\mathcal{L}_{\omega_1}$ which is invariant for $(<k)-$counting bisimulation and prove elementary interpolation for it. All proofs can be specialized to the finitary fragment $\mathcal{L}^{\diamond<k}$.

First of all, let us recall how consistency properties are used in $\mathcal{L}_{\omega_1}$ to prove interpolation. Suppose $\phi \models \psi$ and suppose by contradiction that the pair (ϕ, ψ) does not have an interpolant.

Consider the set S of all finite sets of sentences s which can be decomposed as $s = s_1 \cup s_2$, with:

1. $\mathcal{L}(s_1) \subseteq \mathcal{L}(\phi)$, $\mathcal{L}(s_2) \subseteq \mathcal{L}(\psi)$;

2. $\wedge s_1$, $\neg \wedge s_2$ do not have an interpolant.

It can be proved that the set S defined above is a consistency property and since by hypothesis the pair $(\{\phi\}, \{\neg\psi\})$ belongs to S, the Model Existence Lemma tells us that there exists a model for $\phi \wedge \neg\psi$, a contradiction.

To fit our definition of consistency property modulo bisimulation (see Definition 5.1.1), we generalize this idea in the proof of Theorem 5.2.3. In this context the interpolant is a graded formula, but for technical reasons that will become clear during the proof, we need an extension inside $\mathcal{L}_{\omega_1}$ of the graded fragment, the so-called extended graded formulae over a set of variables X (these formulae where first introduced for the modal case by van Benthem). Intuitively, extended graded formulae consist on boolean combination of ordinary graded modal formulae *evaluated* on different nodes.

5.2.1. Definition. A formula $\theta(X) \in \mathcal{L}_{\omega_1}$ is an *extended k-graded formula* over the set of variables X if there are graded infinitary formulae $\phi_{j,l} \in \mathcal{L}_{\omega_1}^{\diamond<k}$ and variables $x_{j,l} \in X$ for $j \in J, l \in L$ such that

$$\theta = \bigvee_{j\in J} \bigwedge_{l \in L} ST(\phi_{j,l})(x_{j,l}/x).$$

□

Notice that if $|X| = 1$ the extended graded formulae over X are equivalent to standard translation formulae in $\mathcal{L}_{\omega_1}^{\diamond<k}$.

In order to state the interpolation theorem, we need the following definition.

5.2.2. Definition. Let ϕ, ψ be sentences in $\mathcal{L}_{\omega_1}$. We say that ϕ *entails ψ along $(< k)$−counting bisimulation* if whenever $\mathcal{M}$ is an $\mathcal{L}(\phi)$ model such that $\mathcal{M} \models \phi$ and $\mathcal{N}$ is an $\mathcal{L}(\psi)$ models which is $(< k)$−counting bisimilar to $\mathcal{M}$ with respect to $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$, then $\mathcal{N} \models \psi$. □

Suppose that the two sentences ϕ, ψ have a graded interpolant, that is, there exists a formula $\theta \in (\mathcal{L}(\phi) \cap \mathcal{L}(\psi))_{\omega_1}^{\diamond<k}$ such that

- $\phi \models \theta$ (i.e. if $\mathcal{M} \models \phi$ then $\mathcal{M} \models ST(\theta)(r)$),

- $\theta \models \psi$ (i.e. if $\mathcal{N} \models ST(\theta)(r)$ then $\mathcal{N} \models \psi$).

If ϕ, ψ have a graded interpolant then it is easy to show that ϕ entails ψ along $(< k)-$counting bisimulation: if $\mathcal{M} \models \phi$ and $\mathcal{N}$ is $(< k)-$counting bisimilar to $\mathcal{M}$ w.r.t. $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$, then $\mathcal{M} \models \theta$, $\mathcal{N} \models \theta$ (because k-graded formulae are preserved under $(< k)$-counting bisimulation) and $\mathcal{N} \models \psi$.

The converse is given by the following theorem.

5.2.3. THEOREM. *Let ϕ, ψ be sentences in $\mathcal{L}_{\omega_1}$. If ϕ entails ψ along $(< k)$-counting bisimulation, then there is a graded infinitary formula θ in the logic $(\mathcal{L}(\phi) \cap \mathcal{L}(\psi))_{\omega_1}^{\diamond < k}$ such that $\phi \models \theta$ and $\theta \models \psi$.*

Proof.
By contraposition. If ϕ and ψ are infinitary sentences having no graded interpolant, then using the notion of consistency property modulo $(< k)-$counting bisimulation we prove that there are two models $\mathcal{M}, \mathcal{N}$ which are $(< k)-$counting bisimilar w.r.t. $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$, such that $\mathcal{M} \models \phi$, $\mathcal{N} \models \neg\psi$.

Let C, D be countable sets of constant symbols with $\{r\} = C \cap D$. Consider the set S consisting of all triples (A, Σ, Δ) such that

- $(r, r) \in A$ and A is a finite subset of $C \times D$;
- Σ is a finite set of C-substitution instances of subformulae of ϕ, or of type $c = c'$, $c \neq c'$, for $c, c' \in C$;
- Δ is a finite set of D-substitution instances of subformulae of $\neg\psi$, or of type $d = d'$, $d \neq d'$, for $d, d' \in D$.
- there is no extended graded formula $\theta(X)$ over variables X in the language $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$ and no substitution functions $\sigma_1 : X \to C$, $\sigma_2 : X \to D$ such that:

  for each $x \in X$ the pair $(\sigma_1 x, \sigma_2 x)$ is in A;

  $\Sigma \models \theta(\sigma_1)$ and $\Delta \models \neg\theta(\sigma_2)$.

Notice that only finitely many constants appear in Σ, Δ, and A (this is because if χ is a sentence of $\mathcal{L}_{\omega_1}$ then its subformulae contain only a finite number of free variables).

5.2.4. PROPOSITION. *S is a consistency property modulo $(< k)$-counting bisimulation.*

Proof.
We prove some of the left properties, leaving the others and the right properties to the reader.

(c_1) If ϕ is an atomic sentence in $(\mathcal{L}_C)_{\omega_1}$, suppose by contradiction that $(A, \Sigma \cup \{\phi\}, \Delta) \notin S$, and $(A, \Sigma \cup \{\neg\phi\}, \Delta) \notin S$. Then there are extended graded formulae $\alpha(X), \beta(Y)$ over set of variables X, Y (which we may suppose disjoint) and substitutions $\delta_1, \delta_2, \tau_1, \tau_2$ such that

$$\Sigma \cup \{\phi\} \models \alpha(\delta_1), \qquad \Delta \models \neg\alpha(\delta_2),$$
$$\Sigma \cup \{\neg\phi\} \models \beta(\tau_1), \qquad \Delta \models \neg\beta(\tau_2).$$

If $\theta(X, Y) = \alpha(X) \vee \beta(Y)$ and $\sigma_1 = \delta_1 \cup \tau_1$, $\sigma_2 = \delta_2 \cup \tau_2$, then

$$\Sigma \models \theta(\sigma_1), \qquad \Delta \models \neg\theta(\sigma_2),$$

in contradiction with $(A, \Sigma, \Delta) \in S$.

(c_5) If $\bigvee \Phi \in \Sigma$ and for no $\phi \in \Phi$ the triple $(A, \Sigma \cup \{\phi\}, \Delta)$ belongs to S, then for each $\phi \in \Phi$ there is an extended graded formula θ_ϕ over a set of variables X_ϕ and two substitutions $\sigma_1^\phi, \sigma_2^\phi$ such that

$$\Sigma \cup \{\phi\} \models \theta_\phi(\sigma_1^\phi),$$
$$\Delta \models \neg\theta_\phi(\sigma_2^\phi),$$

and $(\sigma_1^\phi(x), \sigma_2^\phi(x)) \in A$, for all $x \in X_\phi$. We may suppose without loss of generality that the sets of variables are pairwise disjoint.

Consider the formula $\theta = \bigvee_{\phi\in\Phi} \theta_\phi$ over the set of variables $X = \bigcup_{\phi\in\Phi} X_\phi$ and the two substitutions $\sigma_1 = \bigcup_{\phi\in\Phi} \sigma_1^\phi$, $\sigma_2 = \bigcup_{\phi\in\Phi} \sigma_2^\phi$. Then

$$\Sigma \models \theta(\sigma_1),$$
$$\Delta \models \neg\theta(\sigma_2),$$

and $(\sigma_1(x), \sigma_2(x)) \in A$, for all $x \in X$, proving that $(A, \Sigma, \Delta) \notin S$.

(c_6) If $\exists x\phi \in \Sigma$, consider a $c \in C$ which is new for A, Σ (this is possible because A and Σ only contain a finite number of constants). We prove that $(A, \Sigma \cup \{\phi(c/x)\}, \Delta) \in S$. If $(A, \Sigma \cup \{\phi(c/x)\}, \Delta) \notin S$ then there are an extended graded formula $\theta(X)$ and substitutions σ_1, σ_2 such that $\Sigma \cup \{\phi(c/x)\} \models \theta(\sigma_1)$, $\Delta \models \neg\theta(\sigma_2)$, and $(\sigma_1 x, \sigma_2 x) \in A$, for all $x \in X$. Since c is new for A, Σ we also have $\Sigma \models \theta(\sigma_1)$, $\Delta \models \neg\theta(\sigma_2)$, a contradiction.

We prove the left bisimulation properties.

(b_1) If $(c, d) \in A$, $p \in \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$, and $p(c) \in \Sigma$ suppose by contradiction that there are an extended graded formula $\theta(X)$ and substitutions σ_1, σ_2 such that $\Sigma \models \theta(\sigma_1)$, $\Delta \cup \{p(d)\} \models \neg\theta(\sigma_2)$, and $(\sigma_1 x, \sigma_2 x) \in A$, for all $x \in X$.

Then $\Sigma \models p(c) \wedge \theta(\sigma_1)$, $\Delta \models \neg(p(d) \wedge \theta(\sigma_2))$, and $(A, \Sigma, \Delta) \notin S$, witnessed by the extended formula $\theta'(y, X) = p(y) \wedge \theta(X)$ (which can be rewritten in the form required to be an extended graded formula by moving $p(y)$ inside over disjunctions) and by the substitutions $\sigma_1 \cup \{(y, c)\}$, $\sigma_2 \cup \{(y, d)\}$.

(b_2) for all $h < k$, if $(c, d) \in A$, and a is an action in $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$ with

$$\{R_a(c, c_i) : i < h\} \cup \{c_i \neq c_j : i < j < h\} \subseteq \Sigma,$$

then consider some constants $d_0, \ldots, d_{h-1}$ new for A, Δ.

Suppose there is an extended graded formula $\theta(X)$ and substitutions σ_1, σ_2 such that

$$\Sigma \models \theta(\sigma_1), \qquad \Delta \cup \{R_a(d, d_i) : i < h\} \cup \{d_i \neq d_j : i < j < h\} \models \neg\theta(\sigma_2),$$

and $(\sigma_1 x, \sigma_2 x) \in A \cup \{(c_i, d_i) : i < h\}$, for all $x \in X$. For each $i < h$, let $Y_i = \{x \in X : \sigma_1 x = c_i, \sigma_2 x = d_i\}$ and consider the set of variables $X' = X \setminus \bigcup_{i<h} Y_i$. If $\theta'(X' \cup \{x\})$ is defined as the formula

$$\exists y_0 \ldots y_{h-1} \ (\bigwedge_{i<h} x R_a y_i \wedge \bigwedge_{i<j} y_i \neq y_j \wedge \theta(y_0/Y_0, \ldots, y_{h-1}/Y_{h-1})),$$

(where each y_i substitutes all variables in Y_i) and σ'_1, σ'_2 are the restrictions of σ_1, σ_2 to X', then

$$\Sigma \models \theta'(\sigma'_1 \cup \{(x, c)\}),$$
$$\Delta \models \neg\theta'(\sigma'_2 \cup \{(x, d)\}).$$

However, θ' is not in the form of an extended graded formula. The problem is that in θ' we require the existence of different successors $y_0 \ldots y_{h-1}$ satisfying the formula θ, where all y_i may occur simultaneously. Since θ *is* an extended graded formula, we can divide out θ so that each y_i has to satisfy a formula without the occurrence of other variables and then we can use Proposition 2.4.2 to show that the obtained formula is equivalent to an extended graded formula over X'.

Formally, if $\theta = \bigvee_{l \in L} \bigwedge_{j \in J} ST\phi_{l,j}(x_{l,j})$, then θ' is equivalent to the formula

$$\bigvee_{l \in L} \exists y_0 \ldots y_{h-1} \ (\bigwedge_{i<h} x R_a y_i \wedge \bigwedge_{i<j} y_i \neq y_j \wedge \bigwedge_{j \in J} ST\phi_{l,j}(x_{l,j})[\ldots, y_i/Y_i, \ldots]).$$

A variable $x_{l,j}$ is substituted by a y-variable only if $x_{l,j} \in \bigcup_{i<h} Y_i$, proving the equivalence with the formula

$$\bigvee_{l \in L} \exists y_0 \ldots y_{h-1} \ (\bigwedge_{i<h} x R_a y_i \wedge \bigwedge_{i<j} y_i \neq y_j \wedge$$

$$\bigwedge_{x_{l,j} \notin \bigcup_{i<h} Y_i} ST\phi_{l,i}(x_{l,j}) \wedge \bigwedge_{i<h} \bigwedge_{x_{l,j} \in Y_i} ST\phi_{l,j}(y_i)),$$

which is in turn equivalent to

$$\bigvee_{l\in L}\ \bigwedge_{x_{l,j}\notin\bigcup_{i<h}Y_i} ST\phi_{l,i}(x_{l,j})\wedge\gamma_l(x),$$

where

$$\gamma_l(x)=\exists y_0\ldots y_{h-1}\ (\bigwedge_{i<h} xR_a y_i\wedge\bigwedge_{i<j} y_i\neq y_j\wedge\bigwedge_{i<h}\ \bigwedge_{x_{l,j}\in Y_i} ST\phi_{l,j}(y_i)).$$

If $\{a\}_h$ is the operator defined in Section 2.4, then $\gamma_l(x)$ is equivalent to the formula

$$\{a\}_h(\bigwedge_{x_{i_j}\in Y_0}\phi_{i_j}\wedge\ldots\wedge\bigwedge_{x_{i_j}\in Y_{n-1}}\phi_{i_j}).$$

Since the formulae ϕ_{i_j} belongs to $\mathcal{L}^{\Diamond<k}_{\omega_1}$, by Proposition 2.4.2 it follows that $\gamma_l(x)$ is equivalent to the (standard translation of) a formula ϕ_l in $\mathcal{L}^{\Diamond<k}_{\omega_1}$. Hence, θ' is equivalent to the k-extended graded formula

$$\bigvee_{l\in L}\ \bigwedge_{x_{l,j}\notin\bigcup_{i<h}Y_i} ST\phi_{l,i}(x_{l,j})\wedge ST\phi_l(x).$$

Then, $(A,\Sigma,\Delta)\notin S$.

Next, we prove the equality property $(=)_4$, leaving the other properties to the patient reader.

$(=_4)$ If $(c,d)\in A$ and $(c=c')\in\Sigma$, suppose that there exists an extended graded formula $\theta(X)$ substitutions σ_1,σ_2 such that

$$\Sigma\models\theta(\sigma_1),\qquad\Delta\models\neg\theta(\sigma_2),$$

and $(\sigma_1 x,\sigma_2 x)\in A\cup\{(c',d)\}$, for all $x\in X$. Consider the set $Y=\{x\in X:\sigma_1 x=c',\sigma_2 x=d\}$, and the substitution

$$\tau_1 x\ =\ \begin{cases}\sigma_1 x & \text{if } x\in X\setminus Y\\ c & \text{otherwise;}\end{cases}$$

Then $\Sigma\models\theta(\tau_1)$, $\Delta\models\neg\theta(\sigma_2)$, and $(\tau_1 x,\sigma_2 x)\in A$, for all $x\in X$, contradicting $(A,\Sigma,\Delta)\in S$.

Having proved that S is a consistency property modulo $(< k)-$counting bisimulation, we can conclude the proof of our theorem. If ϕ, ψ are sentences in $(\mathcal{L}_C)_{\omega_1}$ that have no graded interpolant, the triple $\{(r,r), \phi, \neg\psi\}$ belongs to S. Otherwise there are an extended graded formula $\theta(X)$ and two substitutions σ_1, σ_2 such that

$$\phi \models \theta(\sigma_1), \qquad \neg\psi \models \neg\theta(\sigma_2),$$

and $(\sigma_1 x, \sigma_2 x) = (r,r)$, for all $x \in X$. Hence $\sigma_1(x) = r, \sigma_2(x) = r$, for all $x \in X$, and we may suppose that $|X| = 1$. Therefore $\theta(\sigma_1) = \theta(\sigma_2)$ is equivalent to a graded formula θ which is an interpolant for ϕ, ψ, a contradiction. Hence $\{(r,r), \phi, \neg\psi\}$ belongs to S and by Lemma 5.1.2 it follows that there are two $(< k)$-counting bisimilar models $\mathcal{M}, \mathcal{N}$ with $\mathcal{M} \models \phi$, $\mathcal{N} \models \neg\psi$. This proves that ϕ does not entails ψ modulo $(< k)-$counting bisimulation. □

The finitary version of the invariance part of the following corollary is already known; this and related results can be found in [61].

5.2.5. COROLLARY. *Let k be a natural number or $\aleph_0$.*

(Invariance) A sentence $\phi \in \mathcal{L}_{\omega_1}$ is invariant for $(< k)-$counting bisimulation iff it is equivalent to a formula in $\mathcal{L}_{\omega_1}^{\diamond<k}$. If ϕ is a first-order sentence, then the equivalent formula can be found in $\mathcal{L}^{\diamond<k}$.

(Interpolation) If ϕ, ψ are sentences in $\mathcal{L}_{\omega_1}^{\diamond<k}$ and ϕ entails ψ along $(< k)$-counting bisimulation w.r.t. $\mathcal{L}(\phi)\cap\mathcal{L}(\psi)$, there is a graded infinitary formula $\theta \in (\mathcal{L}(\phi) \cap \mathcal{L}(\psi))_{\omega_1}^{\diamond<k}$ such that $\phi \models \theta$ and $\theta \models \psi$ (and the finitary version as above).

Proof.
For the first part, notice that a sentence $\phi \in \mathcal{L}_{\omega_1}$ is invariant for $(< k)-$counting bisimulation iff ϕ entails ϕ along $(< k)-$counting bisimulation.

The second part is just Theorem 5.2.3 applied to the sentences $\phi, \psi \in \mathcal{L}_{\omega_1}^{\diamond<k}$, which can be considered $\mathcal{L}_{\omega_1}$ sentences via the standard translation. □

Notice that this corollary gives elementary interpolation for the logic $\mathcal{L}_{\omega_1}^{\diamond<k}$: since two $(< k)$-counting bisimilar structures satisfy the same $\mathcal{L}_{\omega_1}^{\diamond<k}$-sentences, if ϕ entails ψ along $(\mathcal{L}(\phi) \cap \mathcal{L}(\psi))_{\omega_1}^{\diamond<k}$-equivalence, then ϕ entails ψ along $(< k)$-counting bisimulation w.r.t. $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$ and we can apply the corollary to get the interpolant.

Finally, we remark that we cannot improve Corollary 5.2.5 to full interpolation in the logics $\mathcal{L}_{\omega}^{\diamond<n}$ for $n \geq 3$. At the beginning of this chapter we proved that $\mathcal{L}_{\omega}^{\diamond<3}$ does not have interpolation and it is easy to modify the counterexample we gave for this logic in order to fit the logics $\mathcal{L}_{\omega}^{\diamond<n}$: just consider the two formulae

$$\phi(p) \equiv (\langle a \rangle_{n-1} p \wedge \langle a \rangle \neg p),$$

$$\psi(q) \equiv (\langle a\rangle_2 q \vee \langle a\rangle_{n-1}\neg q).$$

We have $\phi(p) \models \psi(q)$, but $\phi(q)$ does not entail $\psi(q)$ along $(< n)-$counting bisimulation. Hence ϕ, ψ cannot have an interpolant. The situation is different if we allow the use of the operators $\langle a\rangle_n$ for all natural n: indeed, an interpolant for the formulae ϕ, ψ above is the formula $\langle a\rangle_n\top$. In the next section we prove Craig interpolation for the Infinitary Logic $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$.

5.3 Craig interpolation for $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$

In the previous section we proved a weak form of interpolation for the graded logics $\mathcal{L}^{\diamond<k}_{\omega_1}$, $\mathcal{L}^{\diamond<k}$, for $k \leq \aleph_0$ and we provided a counterexample to interpolation in the logics $\mathcal{L}^{\diamond<n}$, $\mathcal{L}^{\diamond<n}_{\omega_1}$ for a natural number n. The situation is different for the logics $\mathcal{L}^{\diamond<\aleph_0}$, $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$: these logics enjoy the classical Craig interpolation property. This result is already known for $\mathcal{L}^{\diamond<\aleph_0}$ ([1]) and it will be proved here for the logic $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ by using Corollary 5.2.5 and the notion of unraveling (see Definition 2.2.5).

For the sake of simplicity, in the following two sections we just deal with languages containing as binary relation symbols only one binary relation R_a, which we denote by R. Similarly, the graded operators $\langle a\rangle_h$ are denoted by $\diamond_h$. However, this restriction is not necessary and the proof could be carried out with minor adaptations also in the case of a countable number of such relations.

By Corollary 5.2.5 we know that a sufficient condition for a pair (ϕ, ψ) in $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ to have an interpolant is that ϕ entails ψ along $(< \aleph_0)$-counting bisimulation w.r.t. the language $\mathcal{L}_1 = \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$. To prove interpolation for $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ it is then enough to prove the following lemma:

5.3.1. LEMMA. *Let ϕ, ψ be in $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$. If ϕ entails ψ then ϕ entails ψ along $(< \aleph_0)$-counting bisimulation w.r.t. the language $\mathcal{L}_1 = \mathcal{L}(\phi) \cap \mathcal{L}(\psi)$.*

Proof.
Suppose ϕ does not entails ψ along $(< \aleph_0)$-counting bisimulation w.r.t. $\mathcal{L}_1$, that is, suppose there are two models $\mathcal{M}, \mathcal{N}$ which are $(< \aleph_0)$-counting bisimilar w.r.t. $\mathcal{L}_1$, such that $\mathcal{M} \models \phi$ and $\mathcal{N} \models \neg\psi$. We prove that ϕ does not entails ψ, either.

Claim: we can suppose without loss of generality that $\mathcal{M}$ and $\mathcal{N}$ are countable. To show this claim, consider the language $L = \{M, N, r_M, r_N, R, Z\} \cup \{p : p \in \mathcal{L}_1\}$, where M, N are unary predicates, R, Z are binary, r_M, r_N are constants, and $\{p : p \in \mathcal{L}_1\}$ is the set of unary predicates of $\mathcal{L}_1$. We can then write a $\mathcal{L}_{\omega_1}$- sentence Θ which expresses the existence of models $\mathcal{M}, \mathcal{N}$ as above. For example, to say that Z is an $(< \aleph_0)$-counting bisimulation w..r.t. $\mathcal{L}_1$ between the models $\mathcal{M}$ and $\mathcal{N}$ we use the conjunction of the sentences $M(r) \wedge N(r)$, $r_M Z r_N$,

$\forall x, y(xZy \to \bigwedge_{p\in\mathcal{L}_1}(p(x) \leftrightarrow p(y))$, and

$$\forall x, y(xZy \to \bigwedge_{n<\aleph_0}\Big(\exists x_1 \ldots \exists x_n(\bigwedge_i xRx_i \wedge \bigwedge_{i<j} x_i \neq x_j) \to$$

$$\exists y_1 \ldots y_n(\bigwedge_i (yRy_i \wedge x_iZy_i) \wedge \bigwedge_{i<j} y_i \neq y_j)\Big) \wedge \text{ and vice versa,}$$

while to say that the formula ϕ is true in $\mathcal{M}$ we just use the standard translation $ST(\phi)$ of ϕ with r_M in the place of x.

Since Θ is a consistent $\mathcal{L}_{\omega_1}$-sentence, by the Lowenheim-Skolem Theorem for $\mathcal{L}_{\omega_1}$ we obtain that Θ has a countable model.

Hence, there are two *countable* models $\mathcal{M}, \mathcal{N}$ which are $(<\aleph_0)$-counting bisimilar w.r.t. $\mathcal{L}_1$, such that $\mathcal{M} \models \phi$ and $\mathcal{N} \models \neg\psi$. But an $(<\aleph_0)$-counting bisimulation between countable models is easily seen to be a full counting bisimulation (see Definition 2.4.10) and two counting bisimilar models verify the same formulae of $\mathcal{L}^{grad}_{\infty}$ (see Lemma 2.4.11). By Theorem 4.2.7 (or, better, by its straightforward generalization to languages with propositional constants) it follows that the unravelings $\mathcal{M}^1$ and $\mathcal{N}^1$ are isomorphic. Hence, we can copy the evaluation of all propositional constants of $\mathcal{L}(\psi)$ from $\mathcal{N}^1$ to $\mathcal{M}^1$, ending in a model $\mathcal{K}$ for the language $\mathcal{L}(\phi) \cup \mathcal{L}(\psi)$ which is isomorphic to both the unravelings of $\mathcal{M}$ and $\mathcal{N}$ w.r.t. the language $\mathcal{L}(\phi) \cup \mathcal{L}(\psi)$. Then, again by Theorem 4.2.7, $\mathcal{K}$ satisfies the same formulae of $(\mathcal{L}(\phi) \cup \mathcal{L}(\psi))^{grad}_{\infty}$ as $\mathcal{M}$ and $\mathcal{N}$, hence $\mathcal{K} \models \phi$ and $\mathcal{K} \models \neg\psi$. This proves that ϕ does not entails ψ. □

From Lemma 5.3.1 and Corollary 5.2.5 we obtain:

5.3.2. COROLLARY. *The logic $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ enjoys Craig-interpolation.*

5.4 Interpolation for $\mathcal{L}^{grad}_{\infty}$.

In this section we prove Craig interpolation for the full logic of graded modalities $\mathcal{L}^{grad}_{\infty}$, that is, Infinitary Modal Logic extended with the operators $\diamond_h\phi$, for all cardinals h.

In the previous section we proved interpolation for $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ using the method of consistency property modulo bisimulation. By looking at the proof of this result we see that a central step in it is given by the following property: given n formulae $\phi_0, \ldots, \phi_{n-1} \in \mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$, there exists a formula of $\mathcal{L}^{\diamond<\aleph_0}_{\omega_1}$ that in any structure $\mathcal{M}$ defines the set

$$\{w : \exists w_0, \ldots, \exists w_{n-1} \in R(w) \bigwedge_{i\neq j} w_i \neq w_j \wedge \bigwedge_{i<n} w_i \models \phi_i\}.$$

This was proved in Proposition 2.4.2 using the Marriage Theorem. Unfortunately, the same proof fails when we consider an infinite set of formulae, because a straightforward generalization of the Marriage Theorem does not hold for infinite families (consider for example $\{\{0\},\{0,1\},\ldots,\{0,\ldots,n\},\ldots,\omega\}$). However, as we shall see in Corollary 5.4.7, a generalization of Proposition 2.4.2 still holds and can be used to prove the interpolation property for $\mathcal{L}_\infty^{grad}$.

The proof of interpolation will consist in the following steps.

1. First, in Definition 5.4.1 we generalize the notion of a bounded bisimulation (see Definition 2.2.9) to fit the logics $\mathcal{L}_\infty^{\Diamond<k}$.

2. Second, we prove that we can use $\mathcal{L}_\infty^{\Diamond<k}$ formulae to characterize equivalence classes defined via *bounded* bisimulation and that for a fixed depth α we just need a set of formulae to describe all equivalence classes.

3. Third, using the preceding points we give the generalization to infinite cardinals of Proposition 2.4.2.

4. Finally, we use consistency property and the preceding points to give a proof of interpolation for $\mathcal{L}_\infty^{grad}$.

First Step: Bounded Bisimulation for $\mathcal{L}_\infty^{\Diamond<k}$.

5.4.1. DEFINITION. Fix a limit cardinal k and an ordinal α. An (α,k)-*counting bisimulation* between two models $\mathcal{M}$, $\mathcal{N}$ of a language $\mathcal{L}$ is a sequence $(Z_\beta)_{\beta\leq\alpha}$ of relations $Z_\beta \subseteq \mathcal{M}\times\mathcal{N}$ such that the following holds:

1. $r^{\mathcal{M}} Z_\beta r^{\mathcal{N}}$, for all $\beta\leq\alpha$.

2. For all $\beta\leq\alpha$, if $uZ_\beta v$ then u,v agree on all propositional constants of $\mathcal{L}$.

3. If $\beta<\gamma\leq\alpha$ then $Z_\gamma\subseteq Z_\beta$.

4. If $uZ_{\beta+1}v$, uRu', and X is a subset of $R(w)$ with $|X|<k$, then there exists an injection f between X and $R(v)$ such that $w'Z_\beta f(w')$, for all $w'\in X$.

5. Vice versa.

6. If β is a limit, then $Z_\beta=\bigcap_{\gamma<\beta} Z_\gamma$. □

If $(Z_\beta)_{\beta\leq\alpha}$ is such a sequence we write $\mathcal{M}\sim_{\alpha,k}\mathcal{N}$ and if $wZ_\alpha v$ we write $(\mathcal{M},w)\sim_{\alpha,k}(\mathcal{N},v)$ or simply $w\sim_{\alpha,k} v$. It is not difficult to check that the relation $\sim_{\alpha,k}$ is an equivalence relation on the class of all structures.

We define an (α, k)-counting bisimulation *w.r.t. a language* $\mathcal{L}' \subseteq \mathcal{L}$ as in Definition 5.4.1, except that point 2) is now restricted to propositional constants in $\mathcal{L}'$.

To define the class of graded formulae that are invariant for (α, k)-counting bisimulation, we need the notion of modal depth and modal width of a graded formula:

5.4.2. DEFINITION. The *(graded) modal depth* $d(\phi)$ of a formula $\phi \in \mathcal{L}_\infty^{grad}$ is the ordinal defined by $d(p) = 0$, $d(\phi \vee \psi) = d(\phi \wedge \psi) = d(\phi) + d(\psi)$, $d(\neg\phi) = d(\phi)$, $d(\bigwedge \Theta) = d(\bigvee \Theta) = sup\{d(\phi : \phi \in \Theta\}$, and $d(\Diamond_h \phi) = d(\phi) + 1$, for all cardinal h.

The *modal width* $w(\phi)$ of ϕ is the least cardinal k such that $\phi \in \mathcal{L}_\infty^{\Diamond < k}$ (that is, if ϕ only uses the operators $\Diamond_h$ with $h < k$). □

5.4.3. LEMMA. *Suppose $\phi \in \mathcal{L}_\infty^{grad}$, $d(\phi) \leq \alpha$, and $w(\phi) \leq k$. If $(Z_\beta)_{\beta \leq \alpha}$ is an (α, k)-counting bisimulation between $\mathcal{M}$ and $\mathcal{N}$ such that $uZ_\alpha v$, then*

$$(\mathcal{M}, u) \models \phi \qquad \Leftrightarrow \qquad (\mathcal{N}, v) \models \phi.$$

Proof.
By induction on α. The case $\alpha = 0$ is obvious.

We suppose the proposition is true for all $\beta < \alpha$ and prove it for α by structural induction on ϕ.

The only interesting case is $\phi = \Diamond_h \psi$, with $h < k$. Since $d(\phi) = d(\psi) + 1$, if α is a limit we have $d(\phi) < \alpha$ and we obtain the desired result by induction on α because $(Z_\beta)_{\beta \leq d(\phi)}$ is a $(d(\phi), k)$-counting bisimulation between u and v.

If α is the successor ordinal $\gamma + 1$, then $d(\psi) \leq \gamma$. Suppose $(\mathcal{M}, u) \models \Diamond_h \psi$, and consider a set X of successors of u satisfying ψ such that $|X| = h$. By definition of (α, k)-counting bisimulation, there exists an injection f between X and the set of successors of v such that $w' Z_\gamma f(w')$, for all $w' \in X$. By induction, $(\mathcal{N}, v) \models \Diamond_h \psi$.

Second Step: Bounded bisimulation descriptions.

5.4.4. LEMMA. *Let k be a limit cardinal. For any ordinal α and model $\mathcal{M}$ there is a formula $\theta_{\alpha,k}^{\mathcal{M}} \in \mathcal{L}_\infty^{grad}$ of modal depth α and modal width k such that for all model $\mathcal{N}$ it holds*

$$\mathcal{N} \models \theta_{\alpha,k}^{\mathcal{M}} \quad \Leftrightarrow \quad \mathcal{N} \sim_{\alpha,k} \mathcal{M}.$$

For fixed α and k, the collection of all these formulae, when $\mathcal{M}$ ranges over all models, forms a set modulo semantical equivalence.

Proof.
We prove the proposition by induction on α.

- If $\alpha = 0$, define
$$\theta^{\mathcal{M}}_{0,k} = \bigwedge_{\mathcal{M}\models R} R,$$
where R is a literal of the language.

- Let $\alpha = \beta + 1$. Consider the set $R(w)$ of successors of w and a set X of representatives for the equivalence classes of $\sim_{\beta,k}$ over w. In other words, X is such that if $w^* \in R(w)$ then there is a $w' \in X$ with $w' \sim_{\beta,k} w^*$ and if $x \neq y$ with $x, y \in X$ then $x \not\sim_{\beta,k} y$.

  To proceed, we need some notation: if $w' \in R(w)$, let $c(w') = |\{w^* \in R(w) : w^* \sim_{\beta,k} w'\}|$ and $Y = \{w' \in X : c(w') < k\}$. Let $\theta^{\mathcal{M}}_{\beta+1,k}$ be the conjunction of the following four formulae:
$$\theta^{\mathcal{M}}_{\beta,k}$$
$$\bigwedge_{w'\in Y} \Diamond_{c(w')}\theta^{(\mathcal{M},w')}_{\beta,k} \wedge \neg\Diamond_{c(w')+}\theta^{(\mathcal{M},w')}_{\beta,k},$$
$$\bigwedge_{w'\in X\setminus Y}\ \bigwedge_{h<k} \Diamond_h\theta^{(\mathcal{M},w')}_{\beta,k},$$
$$\Box(\bigvee_{w'\in X} \theta^{(\mathcal{M},w')}_{\beta,k}).$$
  The induction step then consists of showing that $\theta^{\mathcal{M}}_{\beta+1,k}$ describes the $\sim_{\alpha,k}$-class of $\mathcal{M}$. The proof is not difficult and we leave it to the reader. Just notice that since the relation $\sim_{\beta,k}$ is symmetric one should have that if $\mathcal{N} \models \theta^{\mathcal{M}}_{\beta,k}$ then $\mathcal{M} \models \theta^{\mathcal{N}}_{\beta,k}$ and this can be used to prove point 4 of Definition 5.4.1. Notice where the hypothesis that k is a limit is used: if $w' \in X \setminus Y$ we want that the cardinality of the set of successors satisfying $\theta^{(\mathcal{M},w')}_{\beta,k}$ is at least k and this can be done with a formula of $\mathcal{L}^{\Diamond<k}_{\infty}$ (as in the third conjunct) only if k is a limit.

- If α is a limit, just put
$$\theta^{\mathcal{M}}_{\alpha,k} = \bigwedge_{\beta<\alpha} \theta^{\mathcal{M}}_{\beta,k}.$$

Let us prove that for fixed α and k the collection of all these formulae forms a set modulo semantical equivalence, when $\mathcal{M}$ ranges over all structures. The reader can reason by induction on α and show that the collection
$$\{\theta^{(\mathcal{M},w)}_{\alpha,k} : \mathcal{M} \text{ is an } \mathcal{L}\text{-structure and } w \in \mathcal{M}\},$$
is contained in the set
$$\mathcal{L}^{\Diamond<k}_{max\{|\alpha|,k^+,|\{\theta^{(\mathcal{M},w)}_{\beta,k}:\beta<\alpha,\mathcal{M} \text{ is an } \mathcal{L}\text{-structure and } w\in\mathcal{M}\}|\}}.$$

Third Step: On definability of $\{\ \}_h$ inside $\mathcal{L}_\infty^{grad}$.

We now generalize to infinite cardinals Proposition 2.4.2. For each cardinal h, consider the operator $\{\ \}_h$ which applies to a transfinite sequence of formulae $\{\phi_i : i \in h\}$ and whose intended meaning is: *there are at least h different successors $\{y_i : i \in h\}$ with $y_i \models \phi_i$, for all $i \in h$.* We prove in Corollary 5.4.7 that the operator $\{\ \}_h$ is expressible in the logic $\mathcal{L}_\infty^{grad}$, for any cardinal h. The proof is based on the following:

5.4.5. DEFINITION. A class $\mathcal{C}$ of models is said to be *closed for (α, k) equivalence* if whenever $\mathcal{M} \in \mathcal{C}$ and $\mathcal{N} \sim_{\alpha,k} \mathcal{M}$ then $\mathcal{N} \in \mathcal{C}$. □

5.4.6. THEOREM. *A class $\mathcal{C}$ is closed for (α, k) equivalence iff there exists a formula θ of $\mathcal{L}_\infty^{grad}$ with $d(\phi) \leq \alpha, w(\phi) \leq k$ such that $\mathcal{C} = \{\mathcal{M} : \mathcal{M} \models \theta\}$.*

Proof.
Suppose $\mathcal{C}$ is closed for (α, k) equivalence and consider the formula

$$\theta = \bigvee_{\mathcal{M}\in\mathcal{C}} \theta_{\alpha,k}^{\mathcal{M}}.$$

To prove that $\mathcal{N} \in \mathcal{C}$ iff $\mathcal{N} \models \theta$, just apply Lemma 5.4.4. □

5.4.7. COROLLARY. *For any cardinal h the operator $\{\ \}_h$ is expressible in the logic $\mathcal{L}_\infty^{grad}$.*

Proof.
By Theorem 5.4.6 it is enough to show that if $\{\phi_i : i \in h\}$ is a set of $\mathcal{L}_\infty^{grad}$ formulae, $\alpha = sup\{d(\phi) : i \in h\}$, and $k = max\{h^+, sup\{w(\phi) : i \in h\}\}$, then the class $\{\mathcal{M} : \mathcal{M} \models \{\ \}_h(\{\phi_i : i \in h\})\}$ is closed for $(\alpha+1, k)$ equivalence. Consider $\mathcal{M}$ with $\mathcal{M} \models \{\ \}_h(\{\phi_i : i \in h\})$ and a model $\mathcal{N}$ with $\mathcal{N} \sim_{\alpha+1,k} \mathcal{M}$: let $(Z_\beta)_{\beta\leq\alpha+1}$ be an $(\alpha+1, k)$-counting bisimulation between $\mathcal{N}$ and $\mathcal{M}$. If $X = \{m_i : i \in h\}$ is a set of different successors of $r^{\mathcal{M}}$ such that $(\mathcal{M}, m_i) \models \phi_i$, then there exists an injective function between X and $R(r^{\mathcal{N}})$ such that $m_i Z_\alpha f(m_i)$. By Lemma 5.4.3 we have $f(m_i) \models \phi_i$ and $\mathcal{N} \models \{\ \}_h(\{\phi_i : i \in h\})$. □

We have now completed the first three steps towards a proof of interpolation for $\mathcal{L}_\infty^{grad}$. The final step consists of generalizing the notion of consistency property modulo bisimulation (see Definition 5.1.1) and on proceeding in the proof of interpolation as in sections 5.1, 5.2. We will not give all the details of the proofs and definitions involved, but merely indicate how the corresponding proofs and definitions in sections 5.1, 5.2 must be changed in order to fit the logic $\mathcal{L}_\infty^{grad}$. We consider the logic $\mathcal{L}_\infty^{grad}$ as a fragment (via the standard translation) of Infinitary Logic expanded with all quantifiers $Q_h x\ \phi$, whose intended meaning is: *there are* *at least h elements in the domain satisfying ϕ.* This logic is denoted by $\mathcal{L}_\infty(Q_\infty)$. If k is a limit cardinal, we will also use the logic $\mathcal{L}_\infty(Q_k)$ where we restrict the quantifiers Q_h to cardinals $h < k$.

Forth Step: Consistency property for $\mathcal{L}_\infty^{grad}$.

Let k be a limit cardinal and let C, D be sets of constants of cardinality equal to k with $\{r\} = C \cap D$.

5.4.8. DEFINITION. A *consistency property modulo* $(< k)-$*counting bisimulation* is a set of triples (A, Σ, Δ) where

- $(r, r) \in A \subseteq C \times D$,
- Σ is a set of $(\mathcal{L}_C)_\infty(Q_k)$ sentences and Δ is a set of $(\mathcal{L}'_D)_\infty(Q_k)$ sentences (hence only the quantifiers Q_h with $h < k$ can appear in Σ, Δ).
- The cardinality of the set of the subformulae of Σ (Δ) is at most k.

Moreover, this set of triples satisfies the following *closure properties* (we only indicate the differences and only on the left, with Definition 5.1.1):

(c_6) if $Q_h x\ \phi \in \Sigma$, then there is a subset C' of C such that $|C'| = h$ and

$$(A, \Sigma \cup \{\phi(c/x) : c \in C'\} \cup \{c \neq c' : \text{ for different } \ c, c' \in C'\}, \Delta) \in S.$$

(b_2) Suppose that $(c, d) \in A$, $C' \subseteq C$ is such that $|C'| < k$, and

$$\{R(c, c') : c' \in C'\} \cup \{c_i \neq c_j : \text{ for different } \ c_i, c_j \in C'\} \subseteq \Sigma.$$

Then there are a subset D' of D and a bijection f between C' and D' such that $(A^*, \Sigma, \Delta^*) \in S$, where

$$A^* = (A \cup \{(c', f(c')) : c' \in C'\},$$

$$\Delta^* = \Delta \cup \{R(d, d') : d' \in D'\} \cup \{d_i \neq d_j : \text{ for different } \ d_i, d_j \in D'\}.$$

□

5.4.9. THEOREM. *If S is a consistency property modulo $(< k)$-counting bisimulation and $(A, \Sigma, \Delta) \in S$ then there are:*

- *a model $\mathcal{M}$ of $\mathcal{L}_C$ such that $\mathcal{M} \models \Sigma$ and $|\mathcal{M}| \leq k$;*
- *a model $\mathcal{N}$ of $\mathcal{L}'_D$ such that $\mathcal{N} \models \Delta$ and $|\mathcal{N}| \leq k$;*
- *a $(< k)$-counting bisimulation Z w.r.t. $\mathcal{L} \cap \mathcal{L}'$ between $\mathcal{M}$ and $\mathcal{N}$ such that if $(c, d) \in A$ then $(c^{\mathcal{M}}, d^{\mathcal{N}}) \in Z$.*

Proof.
We follow the lines of the proof of Lemma 5.1.2. As in that proof, we may suppose without loss of generality that S is closed.

Let Y be the least set of sentences of $(\mathcal{L}_C)_\infty(Q_k)$ and $(\mathcal{L}'_D)_\infty(Q_k)$ such that:

$\Sigma \subseteq Y$, $\Delta \subseteq Y$;

Y is closed under C-substitution and D-substitution instances of subformulae;

if $c, c' \in C$ and $d, d' \in D$ then $c = c'$, $c \neq c'$, $d = d'$, and $d \neq d'$ belong to Y;

For all $c, c' \in C, d, d' \in D$, if $\phi(c) \in Y$ then $\phi(c') \in Y$, if $\psi(d) \in Y$ then $\psi(d') \in Y$.

Let T be a set of tuples containing:

all tuples (c, c', d), with $c, c' \in C$, $d \in D$;

all tuples (d, d', c), with $d, d' \in D$, $c \in C$;

all tuples (c, d, C'), with $C' \subseteq C$, $|C'| < k$, $c \in C$, and d in D;

all tuples (c, d, D'), with c in C, $d \in D$, and $D' \subseteq D$, $|D'| < k$;

all tuples (c, d, p), with $c \in C$, $d \in D$, and p a unary predicate symbol of $\mathcal{L} \cap \mathcal{L}'$;

all tuples (d, c, p), with $c \in C$, $d \in D$, and p a unary predicate symbol of $\mathcal{L} \cap \mathcal{L}'$.

The set $T \cup Y$ has cardinality at most k and we can enumerate it in a transfinite sequence $G_0, \ldots, G_\alpha, \ldots$ with $\alpha < k$.

We shall construct a transfinite sequence

$$t_0 = (A_0, \Sigma_0, \Delta_0), \ldots, t_\alpha = (A_\alpha, \Sigma_\alpha, \Delta_\alpha), \ldots$$

(with $\alpha < k$) of triples in S as follows.

t_0 is the given (A, Σ, Δ). If α is a limit, then

$$t_\alpha = (\bigcup_{\beta<\alpha} A_\beta, \bigcup_{\beta<\alpha} \Sigma_\beta, \bigcup_{\beta<\alpha} \Delta_\beta).$$

If $(A_\alpha, \Sigma_\alpha, \Delta_\alpha)$ is already defined, then $(A_{\alpha+1}, \Sigma_{\alpha+1}, \Delta_{\alpha+1})$ is constructed using the definition of a consistency property, according to the element G_α. The construction follows the corresponding one in the proof of Lemma 5.1.2, except for the following differences (which we indicate only on the left):

- if $G_\alpha = Q_h x\ \phi$ and $(A_\alpha, \Sigma_\alpha \cup \{G_\alpha\}, \Delta) \in S$ then by (c_6) there is a set $C' \subseteq C$ such that $|C'| = h$ and

$$(A_\alpha, \Sigma_\alpha \cup \{G_\alpha\} \cup \{\phi(c) : c \in C'\} \cup \{c \neq c' : \text{ for different } \ c, c' \in C'\}, \Delta)$$

belongs to S. We take this triple as $t_{\alpha+1}$;

- if $G_\alpha = (c, d, C')$, with $c \in C$, $d \in D$, $C' \subseteq C$, $|C'| < k$, and $(A^*_\alpha, \Sigma^*_\alpha, \Delta_\alpha) \in S$, where

$$A^*_\alpha = A_\alpha \cup \{(c, d)\},$$
$$\Sigma^*_\alpha = \Sigma_\alpha \cup \{R(c, c') : c' \in C'\} \cup \{c_i \neq c_j : \text{ for different } \ c_i, c_j \in C'\},$$

then by the left bisimulation property (b_2) there is a set $D' \subseteq D$ and a bijective function f between D' and C' such that

$$(A^*_\alpha \cup \{(c', f(c')) : c' \in C'\}, \Sigma^*_\alpha, \Delta^*_\alpha) \in S,$$

where

$$\Delta^*_\alpha = \Delta_\alpha \cup \{R(d, d') : d' \in D'\} \cup \{d_i \neq d_j : \text{ for different } \ d_i, d_j \in D'\}).$$

We take this triple as $t_{\alpha+1}$.

Define

$$A_k = \bigcup_{\alpha<k} A_\alpha, \quad \Sigma_k = \bigcup_{\alpha<k} \Sigma_\alpha, \quad \Delta_k = \bigcup_{\alpha<k} \Delta_\alpha.$$

Finally, the construction of the two models $\mathcal{M}$ and $\mathcal{N}$ from C, D, the verification that $\mathcal{M} \models \Sigma_k$, $\mathcal{N} \models \Delta_k$, as well as the fact that there is a $(< k)$-counting bisimulation between $\mathcal{M}$ and $\mathcal{N}$ containing A_k, are proved with minor adjustments as in Lemma 5.1.2 and is left to the reader.

We now apply the method of consistency property to prove Craig interpolation for $\mathcal{L}^{grad}_\infty$. We first prove that this logic has elementary interpolation.

5.4.10. LEMMA. *Let $\phi, \psi \in \mathcal{L}^{grad}_\infty$. If ϕ, ψ do not have an interpolant, then there are a cardinal k and two models $\mathcal{M}$, $\mathcal{N}$ with $|\mathcal{M}|, |\mathcal{N}| \leq k$ such that $\mathcal{M} \models \phi$, $\mathcal{N} \models \neg\psi$ and $\mathcal{M}$ is $(< k)$-counting bisimilar to $\mathcal{N}$. In particular, if ϕ entails ψ along $(< k)$-counting bisimulation, then ϕ, ψ have an interpolant.*

Proof.
Let $\phi, \psi \in \mathcal{L}^{grad}_\infty$ and let k be the first limit cardinal which is bigger than $\aleph_0, w(\phi), w(\psi), |subformulae(\phi)|, |subformulae(\psi)|$. Via the standard translation (see Section 2.4 in the preliminaries) we can consider ϕ, ψ as $\mathcal{L}_\infty(Q_k)$-sentences.

Let C, D be disjoint sets of constant symbols of cardinality equal to k. Consider the set S consisting of all triples (A, Σ, Δ) such that

- $(r, r) \in A \subseteq C \times D$ and $|A| < k$;
- Σ is a set of C-substitution instances of subformulae of ϕ with $|\Sigma| < k$;
- Δ is a set of D-substitution instances of subformulae of $\neg\psi$ with $|\Delta| < k$;
- there is no extended graded formula $\theta(X)$ over variables X in the language $\mathcal{L}(\phi) \cap \mathcal{L}(\psi)$ and no substitution functions $\sigma_1 : X \to C$, $\sigma_2 : X \to D$ such that:

  for each $x \in X$ the pair $(\sigma_1 x, \sigma_2 x)$ is in A;

  $\Sigma \models \theta(\sigma_1)$ and $\Delta \models \neg\theta(\sigma_2)$.

As in the proof of Proposition 5.2.4, one can prove that S is a consistency property modulo $(< k)$-counting bisimulation. Here we only point out the major differences.

(c_6) becomes: if $Q_h x\ \phi \in \Sigma$, consider a set C' with $|C'| = h$ which is new for A, Σ. This set exists because the set of constants appearing in Σ and A has cardinality strictly smaller than k. Since C' is new we have

$$(A, \Sigma \cup \{\phi(c/x) : c \in C'\} \cup \{c \neq c' \text{ for different } \ c, c' \in C'\}, \Delta) \in S.$$

(b_2) is generalized similarly, using Corollary 5.4.7.

Suppose now that the two sentences ϕ, ψ do not have an interpolant: it follows that the triple $(\{(r, r)\}, \{\phi\}, \{\neg\psi\})$ belongs to S and from Theorem 5.4.9 we obtain two models $\mathcal{M}, \mathcal{N}$ both having cardinality at most k with $\mathcal{M} \models \phi$, $\mathcal{N} \models \neg\psi$. Moreover, $\mathcal{M}$ is $(< k)$-counting bisimilar to $\mathcal{N}$ and we obtain that ϕ does not entail ψ along $(< k)-$counting bisimulation. □

We can now give the proof of interpolation:

5.4.11. Theorem. *The logic $\mathcal{L}_\infty^{grad}$ enjoys interpolation.*

Proof.
Let $\phi, \psi \in \mathcal{L}_\infty^{grad}$. Suppose by contradiction that ϕ entails ψ but the pair (ϕ, ψ) does not have an interpolant. Then from Lemma 5.4.10 it follows that there are a cardinal k and two models $\mathcal{M}, \mathcal{N}$ both having cardinality at most k with $\mathcal{M} \models \phi$, $\mathcal{N} \models \neg\psi$. Moreover, $\mathcal{M}$ is $(< k)$-counting bisimilar to $\mathcal{N}$. Since $|\mathcal{M}| \leq k, |\mathcal{N}| \leq k$, and k is a limit cardinal, it is easy to see that a $(< k)$-counting bisimulation between $\mathcal{M}$ and $\mathcal{N}$ is a counting bisimulation (see Definition 2.4.10). Since two counting bisimilar models verify the same formulae of $\mathcal{L}_\infty^{grad}$ (see Lemma 2.4.11), by Theorem 4.2.7 (or, better, by its straightforward generalization to languages with propositional constants) it follows that the unravelings $\mathcal{M}^1$ and

$\mathcal{N}^1$ are isomorphic. As in the proof of Lemma 5.3.1, we can then find a third model $\mathcal{K}$ which is counting bisimilar to $\mathcal{M}$ w.r.t. $\mathcal{L}(\mathcal{M})$ and counting bisimilar to $\mathcal{N}$ w.r.t. $\mathcal{L}(\mathcal{N})$. From this it follows that $\mathcal{K} \models \phi$ and $\mathcal{K} \not\models \psi$, a contradiction. □

5.5 Concluding remarks.

In this chapter we investigated interpolation in the family of (infinitary) graded modal logics, proving elementary interpolation and Craig interpolation whenever possible. One question we left unsolved here is whether the logic $\mathcal{L}^{\Diamond<\aleph_0}$ enjoys uniform interpolation or not. We also have some general questions regarding the method of consistency property modulo bisimulation. We list them below.

- Can we fruitfully use this technique to prove results in other logics as well? In any logic which is invariant under bisimulation, the amalgamation Lemma 2.2.10 allows to prove interpolation by showing that if ϕ entails ψ *modulo bisimulation* then ϕ, ψ have an interpolant. In this chapter we showed how this weak form of interpolation can be proved via consistency property. What about this technique in PDL?

- In this chapter we characterized pairs of formulae having an interpolant in logics which do not have Craig interpolation. What about doing the same with uniform interpolation? For example, we know that PDL does not enjoy uniform interpolation. Can we characterize formulae of PDL having a uniform interpolant? Between them we certainly have all basic modal logic formulae, but these are not all. For example, $\langle a^* \rangle p$ is not a basic modal logic formula but it has $\top$ as uniform interpolant w.r.t. the language $\{a\}$ (more generally, all positive $\phi(p)$ have $\phi(\top)$ as a uniform interpolant).

- In classical logics consistency property are often used to prove completeness theorems. Can we do the same with consistency property modulo bisimulation? In this respect, there are two directions that we think are of some interest. The first one is the more straightforward: just use consistency property modulo bisimulation to prove completeness theorems for infinitary modal logics. The second one is perhaps more interesting. Consider entailment along bisimulation in First Order Logic, that is, consider the relation $\models_\sim$ defined as

  $$\phi \models_\sim \psi \Leftrightarrow \text{ for all bisimilar structures } \mathcal{M}\ \mathcal{N}, \text{ if } \mathcal{M} \models \phi \text{ then } \mathcal{N} \models \psi.$$

  One can show that this relation is r.e. and not decidable. Can we turn the closure conditions of a consistency property into proof rules for this kind of entailment?

Chapter 6

A Set translation method for Modal Logic.

In this chapter we consider the relation between modal logic and set theory in a specular way with respect to our discussion in Chapter 4, where modal logics where used to *describe* non-well-founded sets. Our aim here is to use non-well-founded sets to describe (derivability in extended) modal logic.

The starting point of our discussion is the correspondence between propositional connectives and set operations. Propositional connectives have a natural interpretation as set operations: disjunction corresponds to set-union, conjuction to set-intersection, and negation to set-complementation. If we limit ourselves to these operations, we only use the boolean perspective on sets: sets are seen as flat objects and the membership relation is not even necessary. However, sets can be viewed (as Cantor and von Neumann taught us) as a much more complex collection of objects, among which the membership relation plays such a significant role that the hole mathematics can be built in a language containing only $\in$. One of the first operators on sets that is introduced after $\cap, \cup$, and $\setminus$ is the powerset, that given a set produces the collection of its subsets. In this chapter we show that the modal $\Box$ can be seen as a powerset operating among the elements of a (in general non-well-founded) transitive set, thereby obtaining a correspondence between $\Box$ and *Pow* as natural as the ones between $\wedge$ and $\cap$, $\vee$ and $\cup$, and $\neg$ and $\setminus$.

More precisely, in Chapter 4 we have seen that there are non-well founded set theories in which Kripke frames can be seen as sets, with the membership relation acting as the accessibility relation. A Kripke model is then a set x, together with a collection of sets $x_1, \ldots, x_n$ that are supposed to represent the valuation of the propositional constants $p_1, \ldots, p_n$. In this chapter we show that we can translate a modal formula $\phi(p_1, \ldots, p_n)$ by a set term $\phi^*(x, x_1, \ldots, x_n)$: this term represents the set of those worlds (in the frame x) in which the formula ϕ holds. In other words, we want the following to hold:

$$w \models \phi(p_1, \ldots, p_n) \Leftrightarrow w \in \phi^\star(x, x_1, \ldots, x_n),$$

where on the left we consider w as a world in a Kripke model, and on the right we look at it as a set.

We begin with the correspondence given by translating a propositional letter into a variable:

$$w \models p_i \Leftrightarrow w \in x_i.$$

The following steps are easy: as in propositional logic we translate $\phi \vee \psi$ with $\phi^* \cup \psi^*$. The formula $\neg\phi$ can be translated with $x \setminus \phi^*$, but how should we translate $\Box\phi$ in such a way that the corresponding term $(\Box\phi)^*$ still satisfies the correspondence $w \models \Box\phi \Leftrightarrow w \in (\Box\phi)^*$? Kripke semantics together with our decision to consider the accessibility relation as the membership relation between sets gives us the answer:

$$w \models \Box\phi \Leftrightarrow \forall v(wRv \rightarrow v \models \phi) \Leftrightarrow$$

$$\forall v(v \in w \models v \in \phi^*) \Leftrightarrow w \subseteq \phi^* \Leftrightarrow w \in Pow(\phi^*).$$

Hence, it seems we can translate $(\Box\phi)$ as $Pow(\phi^*)$, achieving our goal to find a natural set operator for the $\Box$.

However, if ϕ is always true in a model, so is $\Box\phi$, while if x is a subset of ϕ^* nothing forces x to be a subset of $Pow(\phi^*)$. Since this property corresponds to the necessitation-rule of modal logic, we *have* to force it for our sets and this is the reason for restricting our correspondence to *transitive* sets, which are subsets of their own powerset: if x is transitive and $x \subseteq \phi^*$ then $x \subseteq Pow(x) \subseteq Pow(\phi^*)$ and the necessitation-rule will be built-in in our system.

We define the universe of sets in which the correspondence is carried out axiomatically. Our set theory Ω is a theory in the language $\{\in, \cup, \setminus, \subset, Pow\}$ and consists simply in a description of the elements of a union, of a difference, and of a powerset. As we will see, neither the axiom of foundation, nor the axiom of extensionality are present in the theory. Modal formulae will be translated as terms of this language with the powerset operator corresponding to the $\Box$-operator; transitive sets will correspond to Kripke frames and all these ingredients will be used to build a sentence $SetTr(\phi, \psi)$ for each pair of formulae ϕ, ψ, such that the derivability of ψ from $K + \phi$ corresponds for *frame complete* ϕ to the derivability of $SetTr(\phi, \psi)$ in Ω.

We obtain this result by comparing the set translation with the standard translation. Standard modal deduction via the standard translation is supported by a simple two-sorted first-order theory μ of general frames (see Chapter 2 Section 2.1.5). Comparing μ with Ω, we establish an effective equivalence between the latter and a suitable extension of μ, that we call μ^+. Such an extension is obtained by adding to μ an axiom reflecting the "uniformity" of our set-theoretic approach which uses just $\in$ for both set-membership and the accessibility relation. We then prove that μ^+ is non-conservative over μ. On the ground of the equivalence between Ω and μ^+, this allows us to conclude that the set translation is not adequate for logics which are not frame complete. However, by varying our translation we could obtain a full equivalence between general modal deduction from K (that is: deduction from $K + \phi$, for all ϕ) and Ω-deduction.

One possible line of generalization of the previous result is toward extended modal logic and there are at least two directions one can explore. One is to leave the standard modal language (or, equivalently, its standard translation in Monadic Second Order logic) unchanged and strengthen the deductive power of the logic over its intended frames; the other is to enhance the expressive power of the logic by extending the basic language (for instance, by adding new propositional operators). In this dissertation we consider the first solution which, in the limit, leads us to standard monadic second-order theories of frames, of which a prime example is the logic L_2 [10, 9]. In particular, such a logic suffices for carrying out the usual Sahlqvist-style substitution arguments that are so characteristic for much modern work in modal logic. The semantics of L_2 is obtained by considering general frames in which the existence of all valuations defined by means of suitable first-order formulae is required. For the case of "pure" general-frame semantics considered before–that is, for the case in which no condition is imposed on the set of possible valuations–the underlying set theory is still Ω. This cannot be the case with L_2, whose expressive power forces a vigorous departure from Ω; this theory allows us to establish a deeper link between second-order and set-theoretic derivability. Our main result in this context is the proof that modal derivability in L_2 can be captured via (a suitable adaptation of) the set translation, extending the underlying set theory by adding closure axioms for the so-called Gödel constructible operations [4, 33].

This chapter is organized as follows. In Section 6.1 we prove the soundness and completeness of the set translation, for frame complete ϕ. In the same section we prove that the theory μ^+, which is used to prove this result, is non-conservative over μ. In Section 6.2 we show how to modify the set translation, in order to discard the hypothesis of frame completeness of ϕ.

In Section 6.3 we consider the problem of extending the set theory Ω in such a way to capture modal derivability in the monadic second-order theory L_2. We begin presenting the L_2 theory and its semantics, then we define the theory Ω_c obtained from Ω adding one axiom for each of the operations introduced by Gödel to describe the constructible universe and then we rewrite the set translation within the language of Ω_c. Next, we prove the completeness of the translation with respect to L_2-derivability, showing that whenever there exists a model of Ω_c containing a set-theoretic frame satisfying a formula, then a true frame satisfying the same formula can be built. Subsequently, we prove the correctness of the translation by showing that a constructible universe (a model of Ω_c) can always be generated starting from a general frame closed under first-order definitions.

6.1 The basic translation.

Let us begin by defining the theory of sets Ω and the translation $SetTr$, taking a pair of modal formulae to a first-order sentence in the language of Ω, in such a way to be able to prove that, for any frame complete formula ϕ:

$$\vdash_{K+\phi} \psi \Leftrightarrow \Omega \vdash SetTr(\phi, \psi).$$

This result will be proved by comparing modal derivability in Ω (via the set translation) with modal derivability in a suitable extension μ^+ of the theory μ (via the standard translation) (see Chapter 2 Section 2.1.5). From this we derive our main result, that is, the equivalence, for frame complete ϕ, between derivability of ψ from $K + \phi$ and derivability of $SetTr(\phi, \psi)$ in Ω. Finally, we prove that the theory μ^+ is stronger than μ, with respect to modal derivability.

6.1.1 The set theory Ω and the modal theory μ^+.

First of all, let us introduce the first component of our correspondence, that is, the set theory Ω. Its axioms, in the language with relational symbols $\in$ and $\subseteq$, and functional symbols $\cup, \setminus$, and Pow, are the following:

$x \in y \cup z \Leftrightarrow x \in y \vee x \in z;$

$x \in y \setminus z \Leftrightarrow x \in y \wedge x \notin z;$

$x \subseteq y \Leftrightarrow \forall z(z \in x \rightarrow z \in y);$

$x \in Pow(y) \Leftrightarrow x \subseteq y.$

Notice that neither the extensionality axiom nor the axiom of foundation are in Ω.

Next, we define the translation $SetTr$. Given a modal formula $\phi(p_1, \ldots, p_n)$, we define its translation as the set-theoretic *term* $\phi^*(x, x_1, \ldots, x_n)$, with variables $x, x_1, \ldots, x_n$, built using $\cup, \setminus$, and Pow. Intuitively the term $\phi^*(x, x_1, \ldots, x_n)$ represents the set of those worlds (in the frame x) in which the formula ϕ holds. The inductive definition of $\phi^*(x, x_1, \ldots, x_n)$ is the following:

- $p_i^* = x_i;$
- $(\phi \vee \psi)^* = \phi^* \cup \psi^*;$
- $(\neg\phi)^* = x \setminus \phi^*;$
- $(\Box\phi)^* = Pow(\phi^*).$

To say that a formula ϕ is true in a Kripke frame x, which we see in our correspondence as a transitive set, we write

$$\forall x_1, \ldots, x_n \; x \subseteq \phi^*(x, x_1, \ldots, x_n).$$

Finally, we define $SetTr(\phi, \psi)$ as the first-order sentence

$$\forall x(Trans(x) \wedge \forall \vec{y}(x \subseteq \phi^*(x, \vec{y})) \to \forall \vec{z}(x \subseteq \psi^*(x, \vec{z}))),$$

where $\vec{y}, \vec{z}$ denote finite sequences of variables and $\forall \vec{y}, \forall \vec{z}$ denote the universal closure w.r.t. these variables.

Let us now consider the second component of our correspondence, that is, modal logic. We know that derivability in modal logic can be seen via the standard translation and the first-order theory μ defined in Section 2.1.5. For all pair of modal formulae ϕ, ψ, the following holds:

$$\vdash_{K+\phi} \psi \Leftrightarrow \mu \vdash \overline{ST(\phi)} \to \overline{ST(\psi)}.$$

Can we compare μ and the standard translation with Ω and the set translation? In other words, can we prove that the derivability of $\overline{ST(\phi)} \to \overline{ST(\psi)}$ from μ is equivalent to the derivability of $SetTr(\phi, \psi)$ from Ω?

We will see in Lemma 6.1.6 that the answer is negative, if we don't put any restriction on ϕ. We will also define an extension μ^+ of the theory μ which corresponds exactly to the theory Ω with respect to derivability of modal formulae. The key idea is the following: in the theory μ we have two sorts of elements, worlds and sets, and two binary relations, the membership relation $\in$ and the accessibility operation R, while in the theory Ω we have just one sort of individuals and only the membership relation $\in$. To compare the two theories we extend μ by requiring that for each element w of the world-sort there exists an element p of the set-sort which has the set of R-successors of w as elements. Formally, μ^+ is the theory obtained by adding to μ the axiom

$$\forall w \exists p \forall v \; (v \in p \leftrightarrow wRv),$$

that links R and $\in$.

In the following we will show that the standard translation method w.r.t. μ^+ and the set-theoretic one w.r.t. Ω are equivalent. More precisely, we will prove that, for all modal formulae ϕ, ψ,

$$\mu^+ \vdash \overline{ST}(\phi) \to \overline{ST}(\psi) \Leftrightarrow \Omega \vdash SetTr(\phi, \psi).$$

6.1.2 Completeness of the translation.

We begin by proving that derivability in μ^+ implies derivability in Ω. Given an Ω-model $\mathcal{U}$ and a transitive set x in that model, we show how to construct a model for μ^+.

6.1.1. DEFINITION. Let $\mathcal{U}$ be a model for Ω and let x be a transitive element of the domain of $\mathcal{U}$ (that is: $x \subseteq^{\mathcal{U}} Pow^{\mathcal{U}}(x)$). The *general frame* $(W, R, \mathcal{W})$ *corresponding to* $\mathcal{U}$ is defined as follows:

- the *world* sort W is given by all the elements (w.r.t. the model $\mathcal{U}$) of x, that is, by the set $\{y \in \mathcal{U} : y \in^{\mathcal{U}} x\}$;
- the *set* sort $\mathcal{W}$ is given by the equivalence classes in $\mathcal{U}$ with respect to the equivalence relation of "having the same elements of the world sort". More precisely, we state that $[y] = [z]$ if and only if $w \in^{\mathcal{U}} y \Leftrightarrow w \in^{\mathcal{U}} z$, for every $w \in^{\mathcal{U}} x$. An element of the set sort is then an equivalence class $[y]$, for $y \in \mathcal{U}$.
- the operations over $\mathcal{W}$ are defined as follows: $[y] \cup [z] = [y \cup^{\mathcal{U}} z]$, $-[y] = [x \setminus^{\mathcal{U}} y]$, and $\Box[y] = [Pow^{\mathcal{U}}(y)]$.
- for worlds y, z we define yRz iff $z \in^{\mathcal{U}} y$, while if y is a world and $[z]$ is a set, we define $y \in [z]$ iff $y \in^{\mathcal{U}} z$. □

Notice that we cannot restrict ourselves to equivalence classes of elements in x, because x may be not closed under the set operations. For instance, the transitive set $\{\emptyset, \{\emptyset\}\}$ is not closed under the Pow operation.

It is easy to show that the resulting model verifies all principles of μ^+.

6.1.2. THEOREM. *For each pair of formulae* ϕ, ψ,

$$\mu^+ \vdash \overline{ST}(\phi) \to \overline{ST}(\psi) \Rightarrow \Omega \vdash SetTr(\phi, \psi).$$

Proof.
Let $\mathcal{U}$ be a model of Ω and let x be a transitive element in it. Consider the corresponding general frame $(W, R, \mathcal{W})$, as in Definition 6.1.1. We first prove that, for any modal formula $\theta(p_1, \ldots, p_n)$, $y \in^{\mathcal{U}} x$, and $y_1, \ldots, y_n \in \mathcal{U}$,

$$\mathcal{U} \models y \in \theta^*(x, y_1, \ldots, y_n) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, [y_1], \ldots, [y_n]).$$

The proof is by induction on the structure of θ.

- if $\theta = p_i$, then $\mathcal{U} \models y \in \theta^*(x, y_i) \Leftrightarrow \mathcal{U} \models y \in y_i \Leftrightarrow (W, R, \mathcal{W}) \models y \in [y_i] \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, [y_i])$;
- if $\theta = \neg\alpha$, then $\mathcal{U} \models y \in \theta^*(x, y_1, \ldots, y_n) \Leftrightarrow \mathcal{U} \models y \in x \setminus \alpha^*(x, y_1, \ldots, y_n) \Leftrightarrow \mathcal{U} \models y \in x \land y \notin \alpha^*(x, y_1, \ldots, y_n) \Leftrightarrow (W, R, \mathcal{W}) \not\models ST(\alpha)(y, [y_1], \ldots, [y_n]) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, [y_1], \ldots, [y_n])$;
- the case of $\cup$ is left to the reader;

- if $\theta = \Box\alpha$ then $\mathcal{U} \models y \in \theta^*(x, y_1, \ldots, y_n) \Leftrightarrow \mathcal{U} \models y \in Pow(\alpha^*(x, y_1, \ldots, y_n))$ $\Leftrightarrow \mathcal{U} \models \forall w\ (w \in y \to w \in \alpha^*(x, y_1, \ldots, y_n)) \Leftrightarrow (W, R, \mathcal{W}) \models \forall w\ (yRw \to ST(\alpha)(w, [y_1], \ldots, [y_n])) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, [y_1], \ldots, [y_n])$ holds.

This implies that, for any modal formula $\theta(p_1, \ldots, p_n)$,

$$\mathcal{U} \models \forall \vec{y}\ (x \subseteq \theta^*(x, \vec{y})) \text{ iff } (W, R, \mathcal{W}) \models \overline{ST}(\phi).$$

Finally, suppose that $\mu^+ \vdash \overline{ST}(\phi) \to \overline{ST}(\psi)$. Let $\mathcal{U}$ be a model of Ω and let x be a transitive element in it. If $\mathcal{U} \models \forall \vec{y}\ (x \subseteq \phi^*(x, \vec{y}))$, then the corresponding general frame $(W, R, \mathcal{W})$ satisfy $(W, R, \mathcal{W}) \models \overline{ST}(\phi)$. By hypothesis, we obtain $(W, R, \mathcal{W}) \models \overline{ST}(\psi)$ and $\mathcal{U} \models \forall \vec{y}(x \subseteq \psi^*(x, \vec{y}))$. □

6.1.3 Soundness of the translation.

In this section we prove that derivability in Ω implies derivability in μ^+. Let $(W, R, \mathcal{W})$ be a general frame satisfying μ^+. This means that $\mathcal{W}$ includes all subsets of W defined by the additional axiom characterizing μ^+, plus the subsets that can be obtained from them by union, (relative) complementation, and the $\Box$ operator. For reason that will become clear during the proof of the next theorem and without loss of generality we suppose that the elements of W are standard sets of the same infinite rank α (the definition of the rank of a set can be found in any book of set theory, see for example [45]). In particular, we have $W \cap \mathcal{W} = \emptyset$. An Ω-model can be obtained from $(W, R, \mathcal{W})$ in the following way:

6.1.3. Definition. The structure $\mathcal{U}$ *generated* by $(W, R, \mathcal{W})$ is defined as follows. Let

$$\begin{aligned} U_0 &= W, \\ &\vdots \\ U_n &= Pow(U_{n-1}) \cup U_{n-1}, \end{aligned}$$

and let $U = \bigcup_{n=0}^{\omega} U_n$. Define a function $F : U \to U$ as:

$$F(x) = \begin{cases} \{v \in W : xRv\} & \text{if } x \in W \\ x \setminus W & \text{if } x \notin W \wedge x \cap W \notin \mathcal{W} \\ x & \text{if } x \notin W \wedge x \cap W \in \mathcal{W}. \end{cases}$$

The structure $\mathcal{U}$ for the language $\{\in\}$ is obtained by taking U as domain and by defining

$$x \in^{\mathcal{U}} y \Leftrightarrow x \in F(y).$$

□

Since $(W, R, \mathcal{W})$ is a general frame satisfying μ^+, we have $\{v \in W : xRv\} \in \mathcal{W}$ for any $x \in W$ and the elements of $\mathcal{W}$ are contained in U_1. The following properties are easily verified:

- if $y \in W$, then $x \in^{\mathcal{U}} y$ iff yRx;
- if $y \in \mathcal{W}$, then $x \in^{\mathcal{U}} y$ iff $x \in y$.

6.1.4. LEMMA. *The interpretation $\mathcal{U}$ generated by a general frame $(W, R, \mathcal{W})$ satisfying μ^+ is an Ω-model.*

Proof.
To show that $\mathcal{U}$ is a model of Ω, it is enough to prove that

$$\begin{array}{lcl} \mathcal{U} & \models & \forall x \forall y \exists z \forall u \; (u \in z \leftrightarrow u \in x \vee u \in y); \\ \mathcal{U} & \models & \forall x \forall y \exists z \forall u \; (u \in z \leftrightarrow u \in x \wedge u \notin y); \\ \mathcal{U} & \models & \forall x \exists z \forall y \; (y \in z \leftrightarrow \forall s \; (s \in y \rightarrow s \in x)). \end{array}$$

Notice that for any i, the set $U_i \backslash W$ only contains a finite number of elements of finite rank, or elements whose rank is strictly greater than α. From this, it follows that the intersection of the range of F and W is empty: if $x \in W$, $F(x) \in \mathcal{W}$ and $\mathcal{W} \cap W = \emptyset$; if $x \notin W$, then either $F(x) = x$ and $F(x) \notin W$ or $F(x) = x \setminus W$. In this case, since $x \subseteq U_i$ for some i, $F(x) \notin W$ as well, for rank reasons. More precisely we have:

$$range(F) = \{y \in \mathcal{U} \setminus W : y \cap W \in \mathcal{W}\}.$$

For any $x, y \in \mathcal{U}$, the set $F(x) \cup F(y)$ is in the range of F. Indeed $(F(x) \cup F(y)) \cap W = (F(x) \cap W) \cup (F(y) \cap W)$ and both $(F(x) \cap W)$ and $(F(y) \cap W)$ belong to $\mathcal{W}$, which is closed under union. Therefore, take a z such that $F(z) = F(x) \cup F(y)$ and observe that

$$t \in^{\mathcal{U}} z \leftrightarrow t \in F(z) \leftrightarrow t \in F(x) \vee t \in F(y) \leftrightarrow t \in^{\mathcal{U}} x \vee t \in^{\mathcal{U}} y.$$

Similarly, we can show that the set $F(x) \setminus F(y)$ is in the range of F. In fact, $(F(x) \setminus F(y)) \cap W = (F(x) \cap W) \setminus (F(y) \cap W)$ and both $(F(x) \cap W)$ and $(F(y) \cap W)$ belong to $\mathcal{W}$, which is closed under difference. Any z such that $F(z) = F(x) \setminus F(y)$ will verify the axiom for $\setminus$.

Consider now the set $\{y \in \mathcal{U} : F(y) \subseteq F(x)\}$. This set is in the range of F because

$$(\{y \in \mathcal{U} : F(y) \subseteq F(x)\}) \cap W = \{y \in W : F(y) \subseteq F(x)\} =$$

$$\{y \in W : \forall z (yRz \rightarrow z \in F(x))\} = \Box(F(x) \cap W).$$

Since $F(x) \cap W$ is in $\mathcal{W}$ and $\mathcal{W}$ is closed under $\Box$, the set $\Box(F(x) \cap W)$ is in $\mathcal{W}$. Any z such that $F(z) = \{y \in \mathcal{U} : F(y) \subseteq F(x)\}$ will verify the axiom for $\Box$, because

$$F(y) \subseteq F(x) \Leftrightarrow \forall s(s \in^{\mathcal{U}} y \to s \in^{\mathcal{U}} x).$$

□

Using the previous lemma we obtain:

6.1.5. THEOREM. *For each pair of formulae* ϕ, ψ,

$$\Omega \vdash SetTr(\phi, \psi) \Rightarrow \mu^+ \vdash \overline{ST}(\phi) \to \overline{ST}(\psi).$$

Proof.
Let $(W, R, \mathcal{W})$ be a general frame satisfying μ^+ and let $\mathcal{U}$ be the Ω-model generated by $(W, R, \mathcal{W})$ as in Definition 6.1.3. For any modal formula $\theta(p_1, \ldots, p_n)$, $y \in W$, and $y_1, \ldots, y_n$ in $\mathcal{U}$, we prove:

$$\mathcal{U} \models y \in \theta^*(W, y_1, \ldots, y_n) \text{ iff } (W, R, \mathcal{W}) \models ST(\theta)(y, F(y_1) \cap W, \ldots, F(y_n) \cap W)$$

(recall that $F(y_i) \cap W \in \mathcal{W}$).

The proof is by induction on the structure of θ. In the following, we denote the sequence $y_1, \ldots, y_n$ by $\vec{y}$.

- if $\theta = p_i$, then $\mathcal{U} \models y \in \theta^*(W, y_i) \Leftrightarrow y \in^{\mathcal{U}} y_i \Leftrightarrow y \in F(y_i) \Leftrightarrow y \in F(y_i) \cap W \Leftrightarrow (W, R, \mathcal{W}) \models ST(\phi)(y, F(y_i) \cap W)$;
- if $\theta = \neg\psi$, then $\mathcal{U} \models y \in \theta^*(W, \vec{y}) \Leftrightarrow \mathcal{U} \models y \in W \setminus \psi^*(W, \vec{y}) \Leftrightarrow \mathcal{U} \models y \in W \wedge y \notin \psi^*(W, \vec{y}) \Leftrightarrow (W, R, \mathcal{W}) \not\models ST(\psi)(y, F(y_1) \cap W, \ldots, F(y_n) \cap W) \Leftrightarrow (W, R, \mathcal{W}) \models -ST(\psi)(y, F(y_1) \cap W, \ldots, F(y_n) \cap W) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, F(y_1) \cap W, \ldots, F(y_n) \cap W)$;
- the case of $\cup$ is left to the reader;
- if $\theta = \Box\psi$, then $\mathcal{U} \models y \in \theta^*(W, \vec{y}) \Leftrightarrow \mathcal{U} \models y \in Pow(\psi^*(W, \vec{y})) \Leftrightarrow \mathcal{U} \models \forall v(v \in y \to v \in \psi^*(W, \vec{y})) \Leftrightarrow (W, R, \mathcal{W}) \models \forall v\ (yRv \to ST(\psi)(v, F(y_1) \cap W, \ldots, F(y_n) \cap W) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(y, F(y_1) \cap W, \ldots, F(y_n) \cap W)$.

This implies that, for any modal formula $\theta(p_1, \ldots, p_n)$,

$$\mathcal{U} \models \forall y_1 \ldots \forall y_n\ (W \subseteq \theta^*(W, y_1, \ldots, y_n)) \Leftrightarrow (W, R, \mathcal{W}) \models \overline{ST}(\theta).$$

Let us prove that W is transitive, that is,

$$\mathcal{U} \models \forall x(x \in W \to \forall y(y \in x \to y \in W)).$$

From $\mathcal{U} \models x \in W$, it follows that x is an element of the world sort in $(W, R, \mathcal{W})$ and from $\mathcal{U} \models y \in x$ it follows that xRy in $(W, R, \mathcal{W})$; therefore, $\mathcal{U} \models y \in W$.

Thus, if $\overline{ST}(\phi)$ holds in $(W, R, \mathcal{W})$, then $\mathcal{U} \models \forall \vec{y}\ (W \subseteq \phi^*(W, \vec{y}))$ and W is transitive in $\mathcal{U}$. By hypothesis, we obtain $\mathcal{U} \models \forall \vec{y}(W \subseteq \psi^*(W, \vec{y}))$ and then $(W, R, \mathcal{W}) \models \overline{ST}(\psi)$. □

6.1.4 Comparing μ^+ and μ.

Theorems 6.1.2, 6.1.5 give a complete correspondence (via the respective translations) between derivability of modal formulae in the theories Ω and μ^+. However, we did not prove yet that these theories are really stronger, with regards to modal formulae, than μ, which we know to correspond to modal derivability in K. This is precisely the goal of the next proposition.

6.1.6. PROPOSITION. *The theory $\mu^+ = \mu + \forall x \exists p \forall y\ (y \in p \leftrightarrow xRy)$ is non-conservative w.r.t. modal (the translation of) formulae over μ.*

Proof.
The proof is a variant of the proof given in [8], where it is proved that the theory μ^s equal to μ plus the axiom $\forall x \exists p \forall y\ (y \in p \leftrightarrow x = y)$ is non-conservative over μ. The modal incomplete logic and the formula showing its incompleteness are indeed the same as in [8].

We prove that, if $\psi = \Box\Diamond\top \rightarrow \Box(\Box(\Box p \rightarrow p) \rightarrow p)$ and $\phi = \Box\Diamond\top \rightarrow \Box\perp$, then

$$\mu \not\vdash \overline{ST}(\psi) \rightarrow \overline{ST}(\phi) \quad \text{(or, equivalently, } \psi \not\models_{gf} \phi\text{)},$$

but

$$\mu^+ \vdash \overline{ST}(\psi) \rightarrow \overline{ST}(\phi).$$

The proof that $\psi \not\models_g \phi$ is given in [8]. Let us prove now that $\mu^+ \vdash \overline{ST}(\psi) \rightarrow \overline{ST}(\phi)$: we suppose that a general frame $(W, R, \mathcal{W})$ satisfying μ^+ is given in which ϕ is not valid and then prove that such a general frame invalidates ψ.

If ϕ is not valid, then there exists a world x in the general frame such that $\Box\Diamond\top$ is true at x and $\exists y\ xRy$. There are two possible cases:

- all y such that xRy are reflexive;
- there exists y such that xRy and y is not reflexive.

In the first case, we know that there exists $X = \{y : xRy\}$, together with its relative complement $-X$. Let us take a valuation $\models$ such that $\{w \in W : w \models \neg p)\} = X$. Consider a world y such that xRy. For all z such that yRz, either $z \in -X$, or xRz and z is reflexive; in both cases $z \models \Box p \rightarrow p$. Thus, $y \models \Box(\Box p \rightarrow p)$, but $y \not\models p$.

Consider now the second case. Let $\models$ be such that $\{w \in W : w \models p\} = \{z : yRz\}$. Since y is not reflexive, again $y \not\models p$; moreover, for all z such that yRz, $z \models p$, and then $y \models \Box(\Box p \rightarrow p)$.

In both cases, there exists at least one y such that xRy and $y \not\models \Box(\Box p \rightarrow p) \rightarrow p$. Thus,

$$x \not\models \Box\Diamond T \rightarrow \Box(\Box(\Box p \rightarrow p) \rightarrow p).$$

□

We can actually prove something stronger, namely, that μ^+ plus $\overline{ST}(\psi)$ is complete with respect to the class of frames for ψ.

6.1.7. PROPOSITION. *For any formula θ,*

$$\psi \models_f \theta \Rightarrow \mu^+ \vdash \overline{ST}(\psi) \to \overline{ST}(\theta)$$

Proof.
Consider the formula ϕ of Proposition (6.1.6). We already proved that

$$\mu^+ \vdash \overline{ST}(\psi) \to \overline{ST}(\phi),$$

and it is straightforward to show that $\phi \models_f \psi$. Therefore, in order to prove that $\psi \models_f \theta$ implies $\mu^+ \vdash \overline{ST}(\psi) \to \overline{ST}(\theta)$, it suffices to show that $\mu \vdash \overline{ST}(\phi) \to \overline{ST}(\theta)$. Since μ^+ is an extension of μ, the thesis follows from $\mu^+ \vdash \overline{ST}(\psi) \to \overline{ST}(\phi)$.

Let $(W, R, \mathcal{W})$ be a general frame satisfying ϕ. Since ϕ does not involve any propositional variable, the underlying frame (W, R) validates ϕ. Hence, it validates ψ and, from $\psi \models_f \theta$, it follows that it validates θ. This implies that also the general frame $(W, R, \mathcal{W})$ validates θ. □

6.2 Capturing full K-derivability.

In the previous section we proved that modal derivability in μ^+ and modal derivability in Ω are comparable by means of the set translation and the standard translation, but we also showed that μ^+ is a non-conservative extension of μ. Hence, a direct comparison between derivability from $K + \phi$ and Ω needs the frame completeness of ϕ as an hypothesis. This gives rise to a natural question: can we modify the translation to make Ω and μ comparable? In such a case, we would be able to capture derivability in $K + \phi$ not only for complete formulae, but in general. The answer is positive. In this section we define a generalization of the proposed set-theoretic translation which works with any (possibly incomplete) modal formula ϕ. The idea is to represent general frames in Ω as a pair satisfying some *closure conditions* (to be denoted by $Cl(x, y)$).

Define $Cl(y, x)$ as the conjunction of the following formulae:

- $\forall z\ (z \in y \to Pow(z) \cap x \in y)$;
- $\forall z\ (z \in y \to x \setminus z \in y)$;
- $\forall z \forall u\ (z \in y \wedge u \in y \to z \cup u \in y)$.

We also slightly modify the translation of the $\Box$:

$$(\Box\phi)^* = Pow(\phi^*) \cap x.$$

We prove that, for all modal formulae ϕ, ψ, the following holds:

$$\vdash_{K+\phi} \psi$$

if and only if

$$\Omega \vdash \forall x \forall y\ (Trans(x) \wedge Cl(y,x) \wedge \forall \vec{z} \in y\ (x \subseteq \phi^*(x,\vec{z})) \rightarrow \forall \vec{z} \in y\ (x \subseteq \psi^*(x,\vec{z}))),$$

where if $\alpha(p_1, \ldots p_n)$ is a modal formula, $\forall \vec{z} \in y\ (x \subseteq \alpha^*(x,\vec{z}))$ stands for $\forall z_1 \ldots \forall z_n\ (z_1 \in y \wedge \ldots \wedge z_n \in y \rightarrow x \subseteq \alpha^*(x, z_1, \ldots, z_n))$.

6.2.1. THEOREM. **(Completeness)** *For each pair of formulae ϕ, ψ,*

$$\vdash_{K+\phi} \psi$$

$$\Downarrow$$

$$\Omega \vdash \forall x \forall y\ (Trans(x) \wedge Cl(y,x) \wedge \forall \vec{z} \in y\ (x \subseteq \phi^*(x,\vec{z})) \rightarrow \forall \vec{z} \in y\ (x \subseteq \psi^*(x,\vec{z}))).$$

Proof.
The proof is by induction on the length of the derivation of ψ from $K + \phi$. We skip the verification of the cases relative to tautologies, K-axioms, and the rule of necessitation and modus ponens. We only discuss the case of the substitution rule.

Let $\psi(p_1, ..., p_n)$ such that

$$\Omega \vdash \forall x \forall y\ (Trans(x) \wedge Cl(y,x) \wedge \forall \vec{z} \in y\ (x \subseteq \phi^*(x,\vec{z})) \rightarrow \forall \vec{z} \in y\ (x \subseteq \psi^*(x,\vec{z})))$$

and let $\tau_1(p_1, .., p_m), ..., \tau_n(p_1, .., p_m)$ be modal formulae. We want to prove that

$$\Omega \vdash \forall x \forall y\ (Trans(x) \wedge Cl(y,x) \wedge \forall \vec{z} \in y\ (x \subseteq \phi^*(x,\vec{z})) \rightarrow$$
$$\forall \vec{z} \in y\ (x \subseteq \psi(\tau_1, \ldots, \tau_n)^*(x,\vec{z}))).$$

The following two facts can be easily proved by induction:

(i) for any modal formula $\sigma(p_1, ..., p_m)$, $\Omega \vdash Cl(y,x) \rightarrow \forall \vec{x} \in y(\sigma^*(\vec{x}) \in y)$;

(ii) $\psi(\tau_1, ..., \tau_n)^*$ is the same (syntactically) as $\psi^*(\tau_1^*|x_1, ..., \tau_n^*|x_n)$.

By simultaneous substitution, it holds that

$$\Omega \vdash \forall \vec{x} \in y\ (x \subseteq \psi^*(x,\vec{x}) \rightarrow \forall \vec{x}\ (\tau_1^* \in y \wedge ... \wedge \tau_n^* \in y \rightarrow x \subseteq \psi^*(x, \tau_1^*, .., \tau_n^*)).$$

By applying *(i)* to the terms $\tau_1^*, .., \tau_n^*$, it follows that

$$\Omega \vdash Cl(y,x) \wedge \forall \vec{x} \in y\ (x \subseteq \psi^*(x,\vec{x})) \rightarrow \forall \vec{x} \in y\ (x \subseteq \psi^*(x, \tau_1^*, .., \tau_n^*)),$$

and from *(ii)* the result follows. $\Box$

Let us now prove the soundness of the modified translation method. First of all, we show how an Ω model can be obtained from $(W, R, \mathcal{W})$. Without loss of generality, we suppose that the elements of W are sets of the same rank $\alpha > 0$.

6.2.2. DEFINITION. The Ω-model $\mathcal{U}$ *generated* by $(W, R, \mathcal{W})$ is defined as follows.
Let

$$\begin{aligned} U_0 &= W, \\ &\vdots \\ U_n &= Pow(U_{n-1}) \cup U_{n-1}, \end{aligned}$$

and let $U = \bigcup_{n=0}^{\omega} U_n$.

Define a function $F : U \to U$ as

$$F(x) = \begin{cases} \{v \in W : xRv\} & \text{if } x \in W \\ x & \text{otherwise} \end{cases}$$

The structure $\mathcal{U}$ for the language $\{\in\}$ is obtained by taking U as domain and by defining

$$x \in^{\mathcal{U}} y \Leftrightarrow x \in F(y).$$

□

To prove that $\mathcal{U}$ is a model for Ω, we first notice that, if $y \in W$ we have:

$$x \in^{\mathcal{U}} y \Leftrightarrow x \in W \wedge yRx.$$

Moreover, we can easily see that the range of F is the set $U \setminus W$. This set is closed under union and difference and F is the identity on $U \setminus W$. Therefore, we can define

$$x \cup^{\mathcal{U}} y = F(x) \cup F(y) \qquad x \setminus^{\mathcal{U}} y = F(x) \setminus F(y)$$

and the axioms for $\cup$ and $\setminus$ are satisfied. Finally, the set $\{y \in U : F(y) \subseteq F(x)\}$ is in $U \setminus W$ too and we can safely define

$$Pow^{\mathcal{U}}(x) = \{y \in U : F(y) \subseteq F(x)\}.$$

6.2.3. THEOREM. **(Soundness)** *For each pair of formulae ϕ, ψ,*

$$\Omega \vdash \forall x \forall y\ (Trans(x) \wedge Cl(y, x) \wedge \forall \vec{z} \in y\ (x \subseteq \phi^*(x, \vec{z})) \to \forall \vec{z} \in y\ (x \subseteq \psi^*(x, \vec{z})))$$

$$\Downarrow$$

$$\vdash_{K+\phi} \psi.$$

Proof.
We prove that $\vdash_{K+\phi} \psi$ by showing that, for any general frame $(W, R, \mathcal{W})$, if ϕ is valid in $(W, R, \mathcal{W})$, then ψ is valid as well. Given a general frame $(W, R, \mathcal{W})$, let $\mathcal{U}$ be the model of Ω generated by $(W, R, \mathcal{W})$ (see Definition 6.2.2).

As before, $\mathcal{U} \models Trans(W)$; moreover, it is straightforward to show that $Closure(\mathcal{W}, W)$ holds in $\mathcal{U}$: closure under union and complementation w.r.t. W follows from the corresponding closure properties of $\mathcal{W}$, while, for all $x \in \mathcal{W}$,

$$Pow^{\mathcal{U}}(x) \cap W = \{y \in W : F(y) \subseteq F(x)\} = \{y \in W : \{v \in W : yRv\} \subseteq x\} =$$

$$= \Box(x) \in \mathcal{W}.$$

It is now easy to show that, for all $\vec{y}$ in $\mathcal{W}$ and for any modal formula $\theta(p_1, \ldots, p_n)$,

$$\mathcal{U} \models W \subseteq \theta^*(W, \vec{y})) \Leftrightarrow (W, R, p_1 := y_1, \ldots, p_n := y_n) \models \theta(p_1, .., p_n).$$

By using this correspondence, if ϕ is valid in $(W, R, \mathcal{W})$ one easily shows that ψ is valid as well. □

6.3 Beyond K: L_2-derivability

In this section we consider the monadic second-order theory L_2 and we find an extension Ω_c of Ω such that modal derivability (via the standard translation) in L_2 corresponds exactly to modal derivability (via the set translation) in this extension.

We begin by giving a brief introduction to the theory L_2 and its semantics. Then we present the theory Ω_c obtained from Ω adding one axiom for each of the operations introduced by Gödel to describe the constructible universe and we 'rewrite' the set translation within the language of Ω_c. Next, we prove the completeness of the translation with respect to L_2-derivability, showing that whenever there exists a model of Ω_c containing a set-theoretic frame satisfying a formula, then a true frame satisfying the same formula can be built. Subsequently, we prove the correctness of the translation showing that a constructible universe (a model of Ω_c) can always be generated starting from a general frame closed for first-order definitions.

6.3.1 The theory L_2.

The system L_2 is defined by means of the rule of *first-order substitution*. This rule can be expressed by the following schema:

$$\forall p \alpha \rightarrow \alpha(\gamma | p),$$

where

- α is a monadic second-order formula;
- $\gamma(x)$ is a first-order formula;

- $\alpha(\gamma|p)$ is the formula obtained from α by replacing the occurrences of $p(u)$ in which the p is free in α by $\gamma(u|x)$ (after renaming bound variables as necessary).

The system L_2 contains a set of axioms which is complete for first-order predicate logic, plus the rule of first-order substitution. This system is sound and complete with respect to a general frame semantics in which *all* first-order definable sets can be used as possible evaluations. More precisely, we say that a general frame $(W, R, \mathcal{W})$ is *closed under first-order definitions* if, for all first-order formulae $\gamma(p_1, \ldots, p_m)$ in the language $\{R, =, p_1, \ldots, \}$, with free variables $x, x_1, \ldots, x_n$, for all $w_1, \ldots, w_n \in W$, and $A_1, \ldots, A_m \in \mathcal{W}$, the set

$$\{w \in W : (W, R, p_1 := A_1, \ldots, p_n := A_n) \models \gamma(w, w_1, \ldots, w_n)\}$$

belongs to $\mathcal{W}$. The sets $A_1, \ldots, A_n$ are called the *parameters* of the definition.

Given two sentences α and β, we have:

$\alpha \vdash_{L_2} \beta \;\Leftrightarrow\;$ for all general frames $(F, \mathcal{W})$ closed under first-order definitions and all assignments σ of worlds in W to individual variables and of set of worlds in $\mathcal{W}$ to unary predicate variables, if $(F, \mathcal{W}), \sigma \models \alpha$, then $(F, \mathcal{W}), \sigma \models \beta$.

6.3.2 The theory Ω_c

We will work with a set theory whose axioms are closely related to the so-called Gödel operations for defining the constructible universe. Among all possible approaches, we choose a weaker version of the one proposed by Barwise in [4] with the aim of stating our results for a theory as weak as possible. We introduce functions defining the singleton operator, suitable cartesian products $(\times, \times_=, \times_\in)$, together with the associated projections (Dom, Rng), plus some operations allowing us to manipulate argument positions in ordered sequences (C_1, C_2).

We call Ω_c the resulting theory. The language of Ω_c consists of $=, \in$, and $\subseteq$ as predicate symbols, $\{\}$, Dom, and Rng as unary functional symbols, and $\cup, \setminus, \times, \times_\in, \times_=, C_1$, and C_2 as binary functional symbols. The axioms for Ω_c consist of a set of axioms complete for predicate logic with equality plus the axioms, already in Ω, describing $\subseteq, \cup$, and $\setminus$ in terms of $\in$, plus the axioms defining the operators $\{\}, \times, \times_\in, \times_=, Dom, Rng, C_1$, and C_2:

$$\begin{aligned}
t \in x \cup y &\;\Leftrightarrow\; t \in x \vee t \in y;\\
t \in x \setminus y &\;\Leftrightarrow\; t \in x \wedge t \notin y;\\
x \subseteq y &\;\Leftrightarrow\; \forall t(t \in x \rightarrow t \in y);\\
t \in \{x\} &\;\leftrightarrow\; t = x;\\
t \in x \times y &\;\leftrightarrow\; \exists a \in x \exists b \in y\ (t = \langle a, b\rangle);\\
t \in x \times_\in y &\;\leftrightarrow\; \exists a \in x \exists b \in y\ (t = \langle a, b\rangle \wedge a \in b);\\
t \in x \times_= y &\;\leftrightarrow\; \exists a \in x \exists b \in y\ (t = \langle a, b\rangle \wedge a = b));
\end{aligned}$$

$$\begin{aligned}
t \in Dom(x) &\leftrightarrow \exists s\ (\langle t, s\rangle \in x);\\
t \in Rng(x) &\leftrightarrow \exists s\ (\langle s, t\rangle \in x);\\
t \in C_1(x, y) &\leftrightarrow \exists a \exists b \exists c\ (\langle a, b\rangle \in x \wedge c \in y \wedge t = \langle a, \langle b, c\rangle\rangle);\\
t \in C_2(x, y) &\leftrightarrow \exists a \exists b \exists c\ (\langle a, c\rangle \in x \wedge b \in y \wedge t = \langle a, \langle b, c\rangle\rangle),
\end{aligned}$$

where $\langle x, y\rangle$ is the usual encoding for ordered pairs, namely $\{\{x\}\} \cup \{\{x\} \cup \{y\}\}$. Inductively, we denote by $\langle x_1, \ldots, x_n\rangle$ the pair $\langle x_1, \langle x_2, \ldots, x_n\rangle\rangle$.

The theory Ω_c differs from the classical approach to constructible sets mainly in the following points:

1. we do not introduce the axioms of choice, infinity, replacement, and powerset, since they are not needed;
2. we do not introduce any form of extensionality, since our aim is to use the membership relation to mimic the accessibility relation in frames.

The translation function $(\cdot)^*$ is the same as that given in Section 6.1, except for the $\Box$ operator. This is obviously the case, because we do not have Pow among the symbols in the language of Ω_c. Moreover, $Trans(x)$ disappears from the antecedent of the translated sentence.

To understand these changes, it can be easily checked that, in the case of Ω, it would have made no difference to work with $(\Box\phi)^*$ defined either as $Pow(\phi^*)$ or as $Pow(\phi^*) \cap x$ (see Section 6.1). Actually, we chose the first alternative only to maintain the translated terms simpler. It is also easy to see that $Pow(\phi^*) \cap x = \{y \in x : y \cap x \subseteq \phi^*\}$, whenever x is transitive; as a matter of fact, we will see that the set $\{y \in x : y \cap x \subseteq \phi^*\}$ can always be used to translate $\Box\phi$ (even in the case in which x is not transitive). Finally, since the following holds:

$$\{y \in x : y \cap x \subseteq \phi^*\} = x \setminus Rng((x \setminus \phi^*) \times_{\in} x),$$

we put:

$$(\Box\phi)^* = x \setminus Rng((x \setminus \phi^*) \times_{\in} x).$$

We will prove the soundness and completeness of the translation by showing that for any pair of modal formulae ϕ, ψ,

$$L_2 \vdash \overline{ST}(\phi) \to \overline{ST}(\psi) \Leftrightarrow \Omega_c \vdash \forall x(\forall \vec{z}\ (x \subseteq \phi^*(x, \vec{z})) \to \forall \vec{z}\ (x \subseteq \psi^*(x, \vec{z}))).$$

6.3.3 Completeness.

We start proving the completeness of the translation, that is,

$$L_2 \vdash \overline{ST}(\phi) \to \overline{ST}(\psi) \Rightarrow \Omega_c \vdash \forall x(\forall \vec{z}\ (x \subseteq \phi^*(x, \vec{z})) \to \forall \vec{z}\ (x \subseteq \psi^*(x, \vec{z}))).$$

Given a model of Ω_c and an element in the model, we preliminary construct a general frame and show that it is closed under first-order definitions. To this end it is useful to introduce the notion of *generated general frame.*

6.3.1. DEFINITION. Let $\mathcal{U}$ be a model for Ω_c and let x_0 be an element of $\mathcal{U}$. The *general frame generated* by x_0 in $\mathcal{U}$ is the triplet $\mathcal{F}_{x_0} = (W, R, \mathcal{W})$ where:

- $W = \{w \in \mathcal{U} : w \in^{\mathcal{U}} x_0\}$,
- for all $w, w' \in W$, $wRw' \Leftrightarrow w' \in^{\mathcal{U}} w$,
- $\mathcal{W} = \{\{w \in \mathcal{U} : w \in^{\mathcal{U}} y\} : y \in \mathcal{U} \wedge y \subseteq^{\mathcal{U}} x_0\}$. □

In order to show that a generated general frame is a general frame closed under first-order definitions, we first prove that for any first-order formula $\alpha(x_1, \ldots, x_n)$, the set of n-tuples of worlds defined by it in $(W, R, \mathcal{W})$ can be obtained in $\mathcal{U}$ by means of the set-theoretic operations of Ω_c.

In the following, given $y \in \mathcal{U}$ such that $y \subseteq^{\mathcal{U}} x_0$, we will identify it with the element $\{w \in \mathcal{U} : w \in^{\mathcal{U}} y\}$ of $\mathcal{W}$ in the generated general frame; in particular, we will identify x_0 with W. Hence, it makes sense to evaluate set-terms on elements of $\mathcal{W}$. We also denote by $\gamma(x_1, \ldots, x_n, p_1, \ldots, p_k)$ a first-order formula in the language $\{R, =, p_1, \ldots, p_n\}$ whose free variables are among $x_1, \ldots, x_n$.

6.3.2. LEMMA. *For every first-order formula $\gamma(x_1, \ldots, x_n, p_1, \ldots, p_k)$, there exists a term $F_\gamma(X_1, \ldots, X_n, Y_1, \ldots, Y_k)$ in the language of Ω_c such that, for all $V_1, \ldots, V_n, W_1, \ldots, W_k$ in $\mathcal{W}$ and $u \in \mathcal{U}$,*

$$\mathcal{U} \models u \in F_\gamma(V_1, \ldots, V_n, W_1, \ldots, W_k)$$

if and only if
there exist $w_1 \in V_1, \ldots, w_n \in V_n$ such that

a) $\mathcal{U} \models u = \langle w_n, w_{n-1} \ldots, w_1 \rangle$,
b) $(W, R, p_1 := W_1, \ldots, p_k := W_k) \models \gamma(w_1, \ldots, w_n)$.

The proof of Lemma 6.3.2 is given in the appendix.

On the ground of Lemma 6.3.2, we can now prove that any generated general frame is indeed a general frame closed under first-order definitions.

6.3.3. LEMMA. *Let $\mathcal{U}$ be a model for Ω_c and let x_0 be an element of $\mathcal{U}$. The general frame $\mathcal{F}_{x_0}$ generated by x_0 in $\mathcal{U}$ is a general frame closed under first-order definitions.*

Proof.
Let $\gamma(x_1, \ldots, x_n, p_1, \ldots, p_k)$ be a first-order formula, $a_2, \ldots, a_n \in W$, and W_1, $\ldots, W_k \in \mathcal{W}$. We have to show that the set $\{a_1 \in W : (W, R, p_1 := W_1, \ldots, p_k := W_k)) \models \gamma(a_1, a_2 \ldots, a_n)\}$ belongs to $\mathcal{W}$. Since the sets $\{a_2\}, \ldots, \{a_n\}$ belong to $\mathcal{W}$, by applying Lemma 6.3.2 with $V_1 = W, V_2 = \{a_2\}, \ldots, V_n = \{a_n\}$, we have that, for every $u \in \mathcal{U}$,

$$\mathcal{U} \models u \in F_\gamma(W, \{a_2\}, \ldots, \{a_n\}, W_1, \ldots, W_k)$$

if and only if
there exists $a_1 \in W$ such that

a) $\mathcal{U} \models u = \langle a_n, a_{n-1} \ldots, a_1 \rangle$,
b) $(W, R, p_1 := W_1, \ldots, p_k := W_k) \models \gamma(a_1, \ldots, a_n)$.

It follows that, for every $a_1 \in \mathcal{U}$,

$$\mathcal{U} \models a_1 \in Rng^{n-1}(F_\gamma(W, \{a_2\}, \ldots, \{a_n\}, W_1, \ldots, W_k))$$

if and only if

$$a_1 \in W \text{ and } (W, R, p_1 := W_1, \ldots, p_k := W_k) \models \gamma(a_1, a_2 \ldots, a_n).$$

Hence, $Rng^{n-1}(F_\gamma(W, \{a_2\}, \ldots, \{a_n\}, W_1, \ldots, W_k)) \subseteq^{\mathcal{U}} W$ and the set

$$\{a_1 \in W : (W, R, p_1 := W_1, \ldots, p_k := W_k) \models \gamma(a_1, a_2 \ldots, a_n)\}$$

is equal to the set $\{a_1 \in \mathcal{U} : a_1 \in^{\mathcal{U}} Rng^{n-1}(F_\gamma(W, \{a_2\}, \ldots, \{a_n\}, W_1, \ldots, W_k))\}$. By definition of generated general frame, it follows that the set

$$\{a_1 \in W : (W, R, p_1 := W_1, \ldots, p_k := W_k) \models \gamma(a_1, a_2 \ldots, a_n)\}$$

belongs to $\mathcal{W}$. Hence $\mathcal{F}_{x_0} = (W, R, \mathcal{W})$ is closed under first-order definability (we do not need to verify the closure under union, relative complementation, and $\Box$ because, if $W_1, W_2 \in \mathcal{W}$, then $W_1 \cup W_2$, $W \setminus W_1$, and $\Box(W_1)$ are characterized by first-order definitions with parameters in $\{W_1, W_2\}$). □

On the basis of the previous results, we can now prove the completeness of the translation.

6.3.4. Theorem. **(Completeness)** *For any pair of formulae ϕ, ψ,*

$$L_2 \vdash \overline{ST}(\phi) \to \overline{ST}(\psi) \Rightarrow \Omega_c \vdash \forall x(\forall \vec{z}\ (x \subseteq \phi^*(x, \vec{z})) \to \forall \vec{z}\ (x \subseteq \psi^*(x, \vec{z}))).$$

Proof.
Let $\mathcal{U}$ be a model for Ω_c and let x_0 be an element of $\mathcal{U}$. We already proved that the general frame $\mathcal{F}_{x_0} = (W, R, \mathcal{W})$ generated by x_0 is closed under first-order

definitions. To prove the thesis it suffices to show that a modal formula θ is valid in $\mathcal{F}_{x_0}$ (equivalently, $\mathcal{F}_{x_0} \models \overline{ST}(\theta)$) if and only if $\mathcal{U} \models \forall \vec{z}\ (x_0 \subseteq \theta^*(x_0, \vec{z}))$ (equivalently, $\mathcal{U} \models \forall \vec{z}\ (W \subseteq \theta^*(W, \vec{z}))$, by the identification of x_0 with W).

It is not difficult to prove that for all $w \in W$ and $z_1, \ldots, z_n \in \mathcal{U}$,

$$\mathcal{U} \models w \in \theta^*(W, z_1, \ldots, z_n) \Leftrightarrow \mathcal{F}_{x_0} \models ST(\theta)(w, z_1 \cap W, \ldots, z_n \cap W),$$

where, for $i = 1, \ldots, n$, $z_i \cap W$ stands for the element $\{v \in W : v \in^{\mathcal{U}} z_i\}$ of $\mathcal{W}$. From this, the thesis easily follows. We just check the property for the $\Box$ operator, leaving the remaining cases to the reader:

$$\mathcal{U} \models w \in (\Box\theta)^*(W, z_1, \ldots, z_n) \Leftrightarrow \mathcal{U} \models w \in W \setminus Rng((W \setminus \theta^*) \times_{\in} W) \Leftrightarrow$$

$$\mathcal{U} \models w \notin Rng((W \setminus \theta^*) \times_{\in} W) \Leftrightarrow \mathcal{U} \models \forall s\ (\langle s, w\rangle \notin (W \setminus \theta^*) \times_{\in} W) \Leftrightarrow$$

$$\mathcal{U} \models \forall s\ (s \in W \wedge s \in w \rightarrow s \in \theta^*) \Leftrightarrow$$

$$\forall s \in W\ (wRs \rightarrow \mathcal{F}_{x_0} \models ST(\theta)(s, z_1 \cap W, \ldots, z_n \cap W)) \Leftrightarrow$$

$$\mathcal{F}_{x_0} \models ST(\Box\theta)(w, z_1 \cap W, \ldots, z_n \cap W).$$

□

6.3.4 Soundness.

Let us prove now the soundness of the translation. We will first define a suitable notion of model *generated* by a general frame $\mathcal{F} = (W, R, \mathcal{W})$ closed under first-order definability; then, we will prove the following two claims.

1. The generated model is a model of Ω_c, which contains W as an element;

2. A modal formula θ is valid in the general frame $\mathcal{F}$ if and only if $\forall \vec{z}(W \subseteq \theta^*(W, \vec{z}))$ holds in the generated Ω_c-model.

In the following, we fix a general frame $\mathcal{F} = (W, R, \mathcal{W})$ closed under first-order definitions.

Definition of Ω_c. The general idea of the construction is to build the model of Ω_c as an Herbrand model starting from the element $W, \mathcal{W}$ of the general frame $\mathcal{F}$. However, we have to be careful in this construction: to achieve point 2 above, in this model any element w of W must have exactly the R-successors of w in $\mathcal{F}$ as elements; moreover, in the domain of the model we do not have to add any element whose intersection with W in the model produces a subset not already in $\mathcal{W}$. For these reasons, we will proceed as follows: without loss of generality, we suppose that

a) all the elements of W are sets of the same rank k,

b) there exists a set X_0 of rank k which is not in W,

c) no element in $W \cup \{X_0\}$ is the empty set, a singleton, a pair, or contains such elements.

The above three assumptions allow us to build the domain of our model in such a way that all its elements have rank greater than or equal to k. More precisely, at rank k we will only have the elements of W and the set X_0 playing the role of the empty-set. The remaining elements of the model will be obtained closing the set $(\mathcal{W} \setminus \{\emptyset\}) \cup \{X_0\}$ under the singleton and the following operations:

$\cup'$ and $\setminus'$, defined as follows:

$z \cup' r$	$r = \emptyset \vee r = X_0$	$r \neq \emptyset \wedge r \neq X_0$
$z = \emptyset \vee z = X_0$	X_0	r
$z \neq \emptyset \wedge z \neq X_0$	z	$z \cup r$

$z \setminus' r$	$r = \emptyset \vee r = X_0$	$r \neq \emptyset \wedge r \neq X_0$
$z = \emptyset \vee z = X_0$	X_0	X_0
$z \neq \emptyset \wedge z \neq X_0$	z	u

$$\text{where} \quad u = \begin{cases} z \setminus r & \text{if } z \setminus r \neq \emptyset \\ X_0 & \text{otherwise} \end{cases}$$

$$z \circ' r = \begin{cases} z \circ r & \text{if } \ z \circ r \neq \emptyset \text{ and } \ z \neq X_0, r \neq X_0 \\ X_0 & \text{otherwise} \end{cases}, \text{ for } \circ \in \{\times, \times_=, C_1, C_2\};$$

$$z \times'_\in r = \begin{cases} z \times_{\in,R} r & \text{if } \ z \times_{\in,R} r \neq \emptyset \text{ and } \ z \neq X_0, r \neq X_0 \\ X_0 & \text{otherwise} \end{cases},$$

where $z \times_{\in,R} r = \{\langle a, b\rangle : a \in z, b \in r, a \in b \vee bRa\}$;

$$Dom'(z) = \begin{cases} Dom(z) & \text{if } \ Dom(z) \neq \emptyset \text{ and } \ z \neq X_0 \\ X_0 & \text{otherwise} \end{cases};$$

$$Rng'(z) = \begin{cases} Rng(z) & \text{if } \ Rng(z) \neq \emptyset \text{ and } \ z \neq X_0 \\ X_0 & \text{otherwise} \end{cases}.$$

Finally, we define the domain U_c of our model as follows:

$\mathcal{W}_0 = (\mathcal{W} \setminus \{\emptyset\}) \bigcup \{X_0\}$;

$\mathcal{W}_{n+1} = \{\{t_1\}, Dom'(t_1), Rng'(t_1) : t_1 \in \bigcup_{i=0}^{n} \mathcal{W}_i\} \cup$

$\{t_1 \circ' t_2 : \circ \in \{\cup, \setminus, \times, \times_\in, \times_=, C_1, C_2\}, t_1, t_2 \in \bigcup_{i=0}^{n} \mathcal{W}_i\}$;

$\overline{\mathcal{W}} = \bigcup_{i=0}^{\omega} \mathcal{W}_i$;

$U_c = W \cup \overline{\mathcal{W}}$.

6.3.5. DEFINITION. The model $\mathcal{U}_c$ *generated* by $(W, R, \mathcal{W})$ is defined as follows.

The domain of interpretation is U_c.

The interpretation $\in^{\mathcal{U}_c}$ of the membership relation is:

$$x \in^{\mathcal{U}_c} y \text{ iff } x \in F(y),$$

where

$$F(y) = \begin{cases} y & \text{if } y \in \overline{\mathcal{W}}, \\ \{w \in W : yRw\} = R(y) & \text{if } y \notin \overline{\mathcal{W}}, y \in W, R(y) \neq \emptyset, \\ X_0 & \text{if } y \notin \overline{\mathcal{W}}, y \in W, R(y) = \emptyset. \end{cases}$$

The interpretation of function symbols is:

$$\begin{array}{rcl} \{z\}^{\mathcal{U}_c} & = & \{z\}, \\ z \circ^{\mathcal{U}_c} r & = & F(z) \circ' F(r), \text{ for } \circ \in \{\cup, \setminus, \times, \times_=, \times_\in, C_1, C_2\}, \\ Dom^{\mathcal{U}_c}(z) & = & Dom'(F(z)), \\ Rng^{\mathcal{U}_c}(z) & = & Rng'(F(z)). \end{array}$$

□

6.3.6. REMARK. Since $F(u) \in \overline{\mathcal{W}}$ for all $u \in U_c$, it can be easily seen that the range of the defined interpretation of the function symbols is contained in $\overline{\mathcal{W}}$. Moreover, $\emptyset \notin U_c$, and, from the inductive definition of $\overline{\mathcal{W}}$, it follows that $u \in \overline{\mathcal{W}}$ if and only if there exist a term $\tau(X_1, \ldots, X_n)$ and $A_1, \ldots, A_n \in \mathcal{W}_0$, such that $U_c \models u = \tau(A_1, \ldots, A_n)$. In the following, we denote by $\vec{A}$ the sequence $A_1, \ldots, A_n$, and $\vec{A} \in \mathcal{W}_0$ stands for $A_1, \ldots, A_n \in \mathcal{W}_0$. □

Our next task consists of proving (Lemma 6.3.11) that $\mathcal{U}_c$ is a model for Ω_c. To show this, we first prove that any element of $\overline{\mathcal{W}} \setminus \{X_0\}$ has rank strictly greater than k (Lemma 6.3.9) and that the element X_0 behaves like the empty-set in $\mathcal{U}_c$, that is, $\mathcal{U}_c \models \forall x(x \notin X_0)$ (Corollary 6.3.10). To show these two results we need the notion of *height* of a term.

6.3.7. DEFINITION. The *height* h of a term is defined as follows:

$h(x) = 1$ if x is a variable,

$h(t \cup s) = h(t \setminus s) = max(h(t), h(s)) + 1,$

$h(\{t\}) = h(Dom(t)) = h(Rng(t)) = h(t) + 1,$

$h(t \times s) = h(t \times_{\in} s) = h(t \times_{=} s) = max(h(t), h(s)) + 5,$

$h(C_1(t, s)) = h(C_2(t, s)) = max(h(t), h(s)) + 9.$

The *height* of an element $u \in \overline{\mathcal{W}}$ is defined as follows:

$$h(u) = min\{h(\tau) : \exists \vec{A} \in \mathcal{W}_0, U_c \models u = \tau(\vec{A})\}.$$

□

By using the height we can prove that U_c is transitive outside $W \cup \{X_0\}$ and that the height decreases according to the standard membership relation.

6.3.8. LEMMA. *For all $u \in \overline{\mathcal{W}} \setminus \{X_0\}$, if $a \in u$ then $a \in U_c$ and either $a \in W$ or $a \in \overline{\mathcal{W}}$ and $h(a) < h(u)$.*

The proof of Lemma 6.3.8 is given in the appendix.

From the above lemma, using the hypothesis on the rank of the elements in W, it follows that any element in $\overline{\mathcal{W}} \setminus \{X_0\}$ has rank strictly greater than k:

6.3.9. LEMMA. *For all u in $\overline{\mathcal{W}} \setminus \{X_0\}$, $rk(u) > k$.*

Proof.
By induction on the height of u. First of all we recall that U_c does not contain the empty set.

If $h(u) = 1$, then $u \in \mathcal{W}$ and $a \in u$ implies $a \in W$. The thesis follows from the fact that all elements of W have rank equal to k.

If $h(u) > 1$ and $a \in u$, then from Lemma 6.3.8 either $a \in W$ or $a \in \overline{\mathcal{W}}$ and $h(a) < h(u)$. If $a \in W$ or $a = X_0$, then $rk(a) = k$. Otherwise, $a \in \overline{\mathcal{W}} \setminus \{X_0\}$ and the thesis follows from the inductive hypothesis. □

6.3.10. COROLLARY. *$\mathcal{U}_c \models \forall x(x \notin X_0)$ and $W \cap \overline{\mathcal{W}} = \emptyset$.*

Proof.
The proof is by contradiction. Since $F(X_0) = X_0$ and $rk(X_0) = k$, it follows that if $u \in U_c$ and $u \in^{\mathcal{U}_c} X_0$, then $rk(u) < k$. Moreover, since any element of W has rank equal to k, u does not belong to W. Therefore, $u \in \overline{\mathcal{W}} \setminus \{X_0\}$, which, together with $rk(u) < k$, contradicts Lemma 6.3.9.

As for the second part of the corollary, since any element of W is different from X_0 and has rank equal to k, we have that $W \cap \overline{\mathcal{W}} = \emptyset$. □

On the ground of the above results, we can prove the first of our initial claims.

6.3.11. LEMMA. *The model generated by $(W, R, \mathcal{W})$ is a model of Ω_c. Moreover, this model is extensional on elements of $\overline{\mathcal{W}}$, i.e. if $u, u' \in \overline{\mathcal{W}}$ are such that $a \in^{\mathcal{U}_c} u \leftrightarrow a \in^{\mathcal{U}_c} u'$ for all $a \in U_c$, then $u = u'$.*

Proof.
In the following equivalences, we will exploit the fact that the range of an interpreted function is a subset of $\overline{\mathcal{W}}$ (see Remark 6.3.6) and hence the function F is the identity on this range. Let a, z, and r be elements of U_c. By definition of F we have that $F(a), F(z), F(r)$ are all different from $\emptyset$.

- $a \in^{\mathcal{U}_c} \{z\}^{\mathcal{U}_c}$ iff $a \in F(\{z\}^{\mathcal{U}_c})$ iff $a \in \{z\}^{\mathcal{U}_c}$ iff $a \in \{z\}$ iff $a = z$. Hence the singleton axiom holds in $\mathcal{U}_c$.

- $a \in^{\mathcal{U}_c} z \cup^{\mathcal{U}_c} r$ iff $a \in F(z \cup^{\mathcal{U}_c} r)$ iff $a \in z \cup^{\mathcal{U}_c} r$ iff $a \in F(z) \cup' F(r)$. We consider the various cases, depending on the definition of $F(z) \cup' F(r)$.

  If $F(z), F(r) \neq X_0$, then $a \in^{\mathcal{U}_c} z \cup^{\mathcal{U}_c} r$ iff $a \in F(z) \cup F(r)$ iff $a \in F(z)$ or $a \in F(r)$ iff $a \in^{\mathcal{U}_c} z$ or $a \in^{\mathcal{U}_c} r$.

  If $F(z) = X_0$ then $F(z) \cup' F(r) = F(r)$. Hence $a \in^{\mathcal{U}_c} z \cup^{\mathcal{U}_c} r$ iff $a \in F(r)$ iff $a \in^{\mathcal{U}_c} r$. On the other hand, by Corollary 6.3.10 $a \notin^{\mathcal{U}_c} z$ and the thesis follows.

  The remaining cases are left to the reader, who can thus verify the axiom of union.

- The verification of the set-difference axiom is similar to the above one.

- $a \in^{\mathcal{U}_c} z \times_{\in}^{\mathcal{U}_c} r$ iff $a \in F(z \times_{\in}^{\mathcal{U}_c} r)$ iff $a \in z \times_{\in}^{\mathcal{U}_c} r$ iff $a \in F(z) \times'_{\in} F(r)$.

  If $F(z) \times_{\in,R} F(r) = \emptyset$, or $F(z) = X_0$, or $F(r) = X_0$, then there is no pair $\langle b, c\rangle$ of elements of U_c such that $b \in^{\mathcal{U}_c} z, c \in^{\mathcal{U}_c} r$ and $b \in^{\mathcal{U}_c} c$. On the other hand, no element of U_c belongs to $X_0 = z \times_{\in}^{\mathcal{U}_c} r$.

  If $F(z) \times_{\in,R} F(r) \neq \emptyset$, $F(z) \neq X_0$, and $F(r) \neq X_0$, then $a \in^{\mathcal{U}_c} z \times_{\in}^{\mathcal{U}_c} r$ iff $a \in F(z) \times_{\in,R} F(r)$ iff $\exists b, c\ (a = \langle b, c\rangle \wedge b \in F(z) \wedge c \in F(r) \wedge (b \in c \vee cRb))$ iff (use Lemma 6.3.8) $\exists b, c \in U_c\ (a = \langle b, c\rangle^{\mathcal{U}_c} \wedge b \in^{\mathcal{U}_c} z \wedge c \in^{\mathcal{U}_c} r \wedge b \in^{\mathcal{U}_c} c)$.

  Hence, the axiom for $\times_{\in}$ is verified.

- The cases of the axioms for $\times, \times_{=}, C_1, C_2$ are similar.

- $a \in^{\mathcal{U}_c} Dom^{\mathcal{U}_c}(z)$ iff $a \in F(Dom^{\mathcal{U}_c}(z))$ iff $a \in Dom^{\mathcal{U}_c}(z)$ iff $a \in Dom'(F(z))$.

  If $Dom(F(z)) = \emptyset$ or $F(z) = X_0$, then there is no pair $\langle a, b\rangle$ in $F(z)$ as well as $a \notin X_0 = Dom'(F(z))$.

  If $Dom(F(z)) \neq \emptyset$ and $F(z) \neq X_0$, we have $a \in^{\mathcal{U}_c} Dom^{\mathcal{U}_c}(z)$ iff $a \in Dom(F(z))$ iff $\exists b\ (\langle a, b\rangle \in F(z))$ iff (use Lemma 6.3.8) $\exists b \in U_c\ (\langle a, b\rangle^{\mathcal{U}_c} \in F(z))$ iff $\exists b \in U_c\ (\langle a, b\rangle^{\mathcal{U}_c} \in^{\mathcal{U}_c} z)$.

  Hence, the axiom for the *Dom* function is verified.

- In a similar way it is possible to verify the axiom for Rng.

Finally, to see that extensionality holds on elements of $\overline{\mathcal{W}}$, suppose we have $u, u' \in \overline{\mathcal{W}}$ such that for all $a \in U_c$ it holds $a \in^{\mathcal{U}_c} u$ iff $a \in^{\mathcal{U}_c} u'$. Then, either $u = u' = X_0$, or, for example, $u \neq X_0$, u is not empty and, by Lemma 6.3.8, if $d \in u$, then $d \in U_c$. Hence $u \neq X_0$ implies $u' \neq X_0$ and $u = u'$, by the extensionality property of the standard universe. □

Let us now consider our second claim, that is, that a modal formula θ is valid in the general frame $\mathcal{F}$ if and only if $\forall \vec{z}(W \subseteq \theta^*(W, \vec{z}))$ holds in the generated Ω_c-model $\mathcal{U}_c$. To prove this it is essential to show that $\mathcal{U}_c$ does not see any new subset of W which is not already in $\mathcal{W}$. To show this we will prove (Lemma 6.3.13) that $F(u) \cap W \in \mathcal{W}$, for all $u \in U_c$. To this end, we extend the definition of height to the elements of W by putting $h(w) = 0$ for all $w \in W$. In the following lemma we show that we can isolate the elements of any element $\tau(\vec{A})$ of $\overline{\mathcal{W}}$ by using smaller terms and Δ_0-formulae. Using the closure of $\overline{\mathcal{W}}$ under first-order definitions, we will then prove in Lemma 6.3.13 that with elements in U_c we do not get more subsets of W that the ones already in $\overline{\mathcal{W}}$.

6.3.12. LEMMA. *For any set term $\tau(\vec{y})$ there are:*

1. *terms $\tau_1(\vec{x}, \vec{y}), \ldots, \tau_{k_\tau}(\vec{x}, \vec{y},)$, such that $h(\tau_i) \leq h(\tau)$ and if $h(\tau) > 1$, then $h(\tau_i) < h(\tau)$,*
2. *Δ_0-formulae in the language $\{=, \in\}$, $\alpha_1^\tau(\vec{x}, \vec{y}), \ldots, \alpha_{k_\tau}^\tau(\vec{x}, \vec{y})$, with possibly W as parameter,*

such that, for all $\vec{A} \in \mathcal{W}_0 \cup W$ and for all $u \in U_c$,

$$\mathcal{U}_c \models u \in \tau(\vec{A}) \leftrightarrow \bigvee_{i=1}^{k_\tau} \exists \vec{w} \in W \ (u = \tau_i(\vec{w}, \vec{A}) \wedge \alpha_i^\tau(\vec{w}, \vec{A})).$$

The proof of Lemma 6.3.12 is given in the appendix.

6.3.13. LEMMA. *For all $u \in U_c$, $F(u) \cap W \in \mathcal{W}$.*

Proof.
Consider first the case in which $u \in W \cup \{X_0\}$; in this case $F(u) \cap W = R(u)$ or $F(u) \cap W = X_0 \cap W = \emptyset$. In both cases, $F(u) \cap W \in \mathcal{W}$, since $\mathcal{W}$ is closed under first-order definitions. If $u \in \overline{\mathcal{W}} \setminus X_0$, then the element $u' = u \setminus^{\mathcal{U}_c} (u \setminus^{\mathcal{U}_c} W) = u \setminus' (u \setminus' W)$ is in $\overline{\mathcal{W}}$. Moreover, for all $a \in U_c$, by the axiom of Ω_c it holds that $a \in^{\mathcal{U}_c} u'$ if and only if $a \in F(u) \cap W$ and we can consider u' instead of $F(u) \cap W$ (recall that $x \cap y = x \setminus (x \setminus y)$). Since $u' \in \overline{\mathcal{W}}$, there exists $\vec{A}$ in $\mathcal{W}_0$ such that $\mathcal{U}_c \models u' = \tau(\vec{A})$, where τ is a term in the language of Ω_c. From Lemma 6.3.12, it follows that there exist

- terms $\tau_1(\vec{x}, \vec{Y}), \ldots, \tau_k(\vec{x}, \vec{Y})$,
- Δ_0-formulae $\alpha_1^\tau(\vec{x}, \vec{Y}), \ldots, \alpha_k^\tau(\vec{x}, \vec{Y})$

such that for all $a \in U_c$,

$$\mathcal{U}_c \models a \in \tau(\vec{A}) \leftrightarrow \bigvee_{i=1}^{k_\tau} \exists \vec{w} \in W \left(a = \tau_i(\vec{w}, \vec{A}) \wedge \alpha_i^\tau(\vec{w}, \vec{A})\right).$$

Notice that if $h(\tau_i) > 1$, then, for any $\vec{w}$ in W, to build the element $\tau_i(\vec{w}, \vec{A})$ we apply a function of the model at least once and $\tau_i(\vec{w}, \vec{A}) \notin W$. Since all the elements of $\tau(\vec{A})$ are in W and $\overline{\mathcal{W}} \cap W = \emptyset$, by Corollary 6.3.10 we must have $h(\tau_i) = 1$ or $\mathcal{U}_c \models \neg\alpha_i^\tau(\vec{w}, \vec{A})$, for all $\vec{w}$. Moreover, if $\tau_i = y_j$, then $\tau_i(A_1, \ldots A_n) = A_j \notin W$. Thus, if the variables $\vec{x}$ used in the terms $\tau_1, \ldots, \tau_k$ are $x_1, \ldots, x_m$, then

$$\mathcal{U}_c \models a \in \tau(\vec{A}) \leftrightarrow \bigvee_{\substack{1 \le i \le k_\tau \\ \tau_i \in \{x_1, \ldots, x_m\}}} \exists \vec{w} \in W \left(a = \tau_i(\vec{w}, \vec{A}) \wedge \alpha_i^\tau(\vec{w}, \vec{A})\right).$$

Notice now that the validity of the right-hand side formula above is equivalent to the validity in (W, R) of a corresponding first-order formula (just read the atomic subformulae of type $x_i \in x_j$ in α_i^τ as $x_j R x_i$ and the parameter X_0 as the empty set). Since $\mathcal{W}$ is closed under first-order definitions, this implies that the set $\{a \in U_c : a \in^{\mathcal{U}_c} u'\} = F(u) \cap W$ belongs to $\mathcal{W}$. □

On the basis of the previous results we can now prove the soundness of the translation.

6.3.14. THEOREM. **Soundness**

$$\Omega_c \vdash \forall x(\forall \vec{z}\ (x \subseteq \phi^*(x, \vec{z})) \rightarrow \forall \vec{z}\ (x \subseteq \psi^*(x, \vec{z}))) \Rightarrow L_2 \vdash \overline{ST}(\phi) \rightarrow \overline{ST}(\psi).$$

Proof.
Consider a general frame $(W, R, \mathcal{W})$ closed under first-order definitions and the generated general model of Definition 6.3.5. We only need to show that a modal formula θ is valid in $(W, R, \mathcal{W})$ if and only if $\mathcal{U}_c \models \forall \vec{z}\ (W \subseteq \theta^*(W, \vec{z}))$. This follows from Lemma 6.3.13. Indeed, for all $w \in W$, $u_1, \ldots, u_n \in U_c$, we can prove by induction on θ that

$$\mathcal{U}_c \models w \in \theta^*(W, u_1, \ldots, u_n) \Leftrightarrow (W, R, \mathcal{W}) \models ST(\theta)(w, F(u_1) \cap W, \ldots, F(u_n) \cap W).$$

We prove the equivalence above only for the cases $\theta = p_i$ and $\theta = \Box\psi$, leaving the remaining cases to the reader.

- If $\theta = p_i$ then $\mathcal{U}_c \models w \in \theta^*(W, u_1, \dots, u_n)$ iff $\mathcal{U}_c \models w \in u_i$ iff $w \in^{\mathcal{U}_c} u_i$ iff $w \in F(u_i)$ iff $w \in F(u_i) \cap W$ iff $(W, R, \mathcal{W}) \models p_i(x)(w, F(u_i) \cap W)$ iff $(W, R, \mathcal{W}) \models ST(p_i)(w, F(u_i) \cap W)$ iff $(W, R, \mathcal{W}) \models ST(\theta)(w, F(u_i) \cap W)$.

- If $\theta = \Box\psi$, then $\mathcal{U}_c \models w \in \theta^*(W, u_1, \dots, u_n)$ iff $\mathcal{U}_c \models w \in W \setminus Rng((W \setminus \psi^*) \times_\in W)$ iff $\mathcal{U}_c \models \forall u \langle u, w \rangle \notin (W \setminus \psi^*) \times_\in W$ iff $\mathcal{U}_c \models \forall u (u \in W \wedge u \in w \rightarrow u \in \psi^*)$ iff $(W, R, \mathcal{W}) \models \forall u (wRu \rightarrow ST(\psi)(u, F(u_1) \cap W, \dots, F(u_n) \cap W)$ iff $(W, R, \mathcal{W}) \models ST(\Box\psi)(u, F(u_1) \cap W, \dots, F(u_n) \cap W)$ iff $(W, R, \mathcal{W}) \models ST(\theta)(u, F(u_1) \cap W, \dots, F(u_n) \cap W)$.

Since by Lemma 6.3.13 we have $F(u) \cap W \in \mathcal{W}$, for all $u \in U_c$, it easily follows that

$$\mathcal{U}_c \models \forall \vec{z}\ (W \subseteq \theta^*(W, \vec{z})) \Leftrightarrow (W, R, \mathcal{W}) \models \overline{ST}(\theta).$$

□

6.4 Concluding remarks.

In this chapter we used non-well-founded sets to give a semantics for Modal Logic in which the □-operator is represented by the powerset operator in transitive non-well-founded sets. We presented this result via a translation from Basic Modal Logic and its complete extensions to a set theory Ω. Then we showed how we can tune this theory to deal with different notions of modal derivability, ending with the second-order theory of general frame closed under first-order definitions. Here we showed that the corresponding set theory is nothing but Ω reinforced with axioms for the Gödel constructible operators. We only considered mono-modal languages and this could appear as an intrinsic limitation of the set translation since we insist on translating the accessibility relation with the membership relation. However it can be proved that the set translation can be generalized to deal with poly-modal logics as well, using a technique similar to the one introduced in [65]. This generalization can be found in [20]. Another generalization goes towards graded modal logics: in [53] it is showed how the set theoretic translation has to be modified to deal with these logics.

Further directions of investigation in this area could be:

- Being a theory of sets, one can think of various other extensions of Ω, besides Gödel constructible operators. Just to make an example, one could add to Ω the axiom of choice. Do we have other natural extensions of Ω which have a natural modal counterpart, as Gödel constructible operators did?

- Can we establish some more *uniform information* on the modal side using the set translation? As an example, various decidability results are known in Computable Set Theory which apply to classes of formulae described by their syntactical complexity. Can we use this results and the set translation to obtain (uniform) decidability results in Modal Logic? A toy example here is the proof of the decidability of K. $SetTr(\top, \psi)$ is a universal sentence in the Ω-language and one can prove (see [19]) that the universal theory of (a variation of) Ω is decidable. Less trivial examples could be obtained by considering the class of subframe logics defined by Kit Fine in [29].

6.5 Appendix

LEMMA 6.3.2 *For every first-order formula $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$, there exists a term $F_\alpha(X_1, \ldots, X_n, Y_1, \ldots, Y_k)$ in the language of Ω_c such that, for all $V_1, \ldots, V_n, W_1, \ldots, W_k$ in $\mathcal{W}$, and $u \in U$,*

$$U \models u \in F_\alpha(V_1, \ldots, V_n, W_1, \ldots, W_k)$$

if and only if
there exist $w_1 \in V_1, \ldots, w_n \in V_n$ such that

a) $U \models u = \langle w_n, w_{n-1} \ldots, w_1 \rangle$,
b) $(W, R, p_1 := W_1, \ldots, p_k := W_k) \models \alpha(w_1, \ldots, w_n)$.

Proof.
The proof is by induction on the structural complexity of α, and it is very similar to the proof of Lemma 6.1, p. 64, in [4].

By possibly renaming bound variables, we can assume that the formula α verifies the following property: if $\forall x_i$ or $\exists x_i$ occur in α, then i is the largest index among the indices of the free variables in the scope of the quantifier. Moreover, notice that a formula in the language $\{R, =, p_1, \ldots, p_j\}$ with free variables among $x_1, \ldots, x_i$ can be considered either as a formula $\alpha(x_1, \ldots, x_i, p_1, \ldots, p_j)$ or as a formula $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$, for any $n \geq i$, $k \geq j$. Statement (1) below shows that we can consider every possible denotation for the same formula without loosing the property expressed by the lemma.

(1) For any m, r, n, k, if $\beta(x_1 \ldots x_m, p_1, \ldots, p_r) = \alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$ and the thesis is true for the formula $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$, then it is true for $\beta(x_1, \ldots, x_m, p_1, \ldots, p_r)$.

Statement (1) follows from (1a)–(1c) below.

(1a) For any $r \geq 1$, if $\beta(x_1, \ldots, x_n, p_1, \ldots, p_k) = \alpha(x_1, \ldots, x_n, p_1, \ldots, p_r)$ and the thesis is true for the formula $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_r)$, then it is true for $\beta(x_1, \ldots, x_n, p_1, \ldots, p_k)$, with

$$F_\beta(X_1, \ldots, X_n, Y_1, \ldots, Y_k) = F_\alpha(X_1, \ldots, X_n, Y_1, \ldots, Y_r)$$

(the verification is left to the reader).

(1b) If $\beta(x_1,\dots,x_n,x_{n+1},p_1,\dots,p_k)=\alpha(x_1,\dots,x_n,p_1,\dots,p_k)$ and the thesis is true for the formula $\alpha(x_1,\dots,x_n,p_1,\dots,p_k)$, then it is true for the formula $\beta(x_1,\dots,x_n,x_{n+1},p_1,\dots,p_k)$, with

$$F_\beta(X_1,\dots,X_{n+1},Y_1,\dots,Y_k)=X_{n+1}\times F_\alpha(X_1,\dots,X_n,Y_1,\dots,Y_k).$$

To prove this, denote by $\overline{V}$ a sequence $V_1,\dots,V_n$ and by $\overline{W}$ a sequence $W_1,\dots,W_k$, where $V_1,\dots,V_n,\dots,W_1,\dots,W_n$ are elements in $\mathcal{W}$. Then, if $V_{n+1}\in\mathcal{W}$ and $u\in U$ we have:

$$U\models u\in F_\beta(\overline{V},V_{n+1},\overline{W})\Leftrightarrow U\models u\in V_{n+1}\times F_\alpha(\overline{V},\overline{W})\Leftrightarrow$$

$$U\models\exists w_{n+1}\exists b\ (w_{n+1}\in V_{n+1}\wedge b\in F_\alpha(\overline{V},\overline{W})\wedge u=\langle w_{n+1},b\rangle)\Leftrightarrow$$

there exist $w_1\in V_1,\dots,w_{n+1}\in V_{n+1}$ such that

a) $U\models u=\langle w_{n+1},w_n,\dots,w_1\rangle$,
b) $(W,R,p_1:=W_1,\dots,p_k:=W_k)\models\beta(w_1,\dots,w_{n+1})$.

(1c) If $\beta(x_1,\dots,x_{n-1},p_1,\dots,p_k)=\alpha(x_1,\dots,x_n,p_1,\dots,p_k)$ and the thesis is true for the right-hand side formula $\alpha(x_1,\dots,x_n,p_1,\dots,p_k)$, then it is true for $\beta(x_1,\dots,x_{n-1},p_1,\dots,p_k)$, with

$$F_\beta(X_1,\dots,X_{n-1},p_1,\dots,p_k)=Rng(F_\alpha(X_1,\dots,X_{n-1},W,p_1,\dots,p_k))$$

(we leave the verification to the reader).

Two further preliminary steps are needed before we can start the inductive proof. In each step, we just determine the term F_β, leaving the verification of the thesis to the reader. The notation $\alpha(\dots,x_i|x_j,\dots,\dots)$ is used to denote uniform substitution of x_i for x_j in the formula α.

(2) If $\beta(x_1,\dots,x_{n+1},p_1,\dots,p_k)=\alpha(x_1,\dots,x_{n+1}|x_n,p_1,\dots,p_k)$ and the thesis is true for α, then it is true for β. If $n=1$, it is enough to take $F_\beta(X_1,X_2,Y_1,\dots,Y_k)=F_\alpha(X_2,Y_1,\dots,Y_k)\times X_1$. If $n>1$, let

$$F_\beta(X_1,\dots,X_{n+1},Y_1,\dots,Y_k)=C_2(F_\alpha(X_1,\dots,X_{n-1},X_{n+1},Y_1,\dots,Y_k),X_n).$$

(3) If $\beta(x_1,\dots,x_n,p_1,\dots,p_k)=\alpha(x_{n-1}|x_1,x_n|x_2,p_1,\dots,p_k)$, with $n\geq 2$, and the thesis is true for α, then it is true for β, with

$$F_\beta(X_1,\dots,X_n,Y_1,\dots,Y_k)=C_1(F_\alpha(X_{n-1},X_n,Y_1,\dots,Y_k),X_{n-2}\times\cdots\times X_1).$$

Let us prove now the thesis by induction on the structural complexity of $\alpha(x_1,\dots,x_n,p_1,\dots,p_k)$. We start with atomic formulae.

(4a) If $\alpha(x_1, x_2)$ is the formula $x_1 = x_2$, then let $F_\alpha = X_2 \times_= X_1$. If we need to consider $x_1 = x_2$ as a formula in $x_1, \ldots, x_n, p_1, \ldots, p_k$, then the thesis follows from (1).

(4b) If $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$ is the formula $x_{n-1} = x_n$, then the thesis follows from (4a) and (3).

(4c) If $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$ is the formula $x_m = x_r$, with $m \leq r \leq n$, there are two possibilities: if $m = r$, then $F_\alpha = X_n \times \ldots \times X_1$; if $m < r$, then the thesis follows from (4a) and (2) by induction on r.

(5a) If $\alpha(x_1, x_2)$ is the formula $x_1 R x_2$, then $F_\alpha = X_2 \times_\in X_1$. If we need to consider $x_1 R x_2$ as a formula in the variables $x_1, \ldots, x_n, p_1, \ldots, p_k$, then the thesis follows from (1).

(5b) If $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$ is the formula $x_{n-1} R x_n$, the thesis follows from (5a) and (3).

(5c) If $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k)$ is the formula $x_i R x_j$, consider the formulae

$$\begin{aligned} \gamma_1(x_1, \ldots, x_{n+2}, p_1, \ldots, p_k) &= x_i = x_{n+1}; \\ \gamma_2(x_1, \ldots, x_{n+2}, p_1, \ldots, p_k) &= x_j = x_{n+2}; \\ \gamma_3(x_1, \ldots, x_{n+2}, p_1, \ldots, p_k) &= x_{n+1} R x_{n+2}. \end{aligned}$$

From (4c) and (5b), it follows that the thesis is true for γ_i, $i = 1, 2, 3$. Then, it is easy to verify that the thesis is true for α, with F_α equal to

$$Rng(Rng(F_{\gamma_1}(X_1, \ldots, X_n, X_i, X_j, Y_1, \ldots, Y_k) \cap$$

$$F_{\gamma_2}(X_1, \ldots, X_n, X_i, X_j, Y_1, \ldots, Y_k) \cap F_{\gamma_3}(X_1, \ldots, X_n, X_i, X_j, Y_1, \ldots, Y_k))).$$

6) If $\alpha(x_1, p_1)$ is the formula $p_1(x_1)$, then $F_\alpha(X_1, Y_1) = X_1 \cap Y_1$. It can be easily generalized to the case of $\alpha(x_1, \ldots, x_n, p_1, \ldots, p_k) = p_j(x_i)$ as in (5c).

In the inductive step, the case of boolean combinations of formulae is dealt with $\cup$ and $\setminus$ (whit respect to $X_n \times \ldots \times X_1$, where n is the number of world variables displayed in the formula), and left to the reader.

Consider the case in which $\beta(x_1, \ldots, x_n, p_1, \ldots, p_k)$ is the formula

$$\exists x_{n+1} \alpha(x_1, \ldots, x_{n+1}, p_1, \ldots, p_k).$$

Let $F_\beta(X_1, \ldots, X_n, Y_1, \ldots, Y_k) = Rng(F_\alpha(X_1, \ldots, X_n, W, Y_1, \ldots, Y_k))$. Then, for all $V_1, \ldots, V_n, W_1, \ldots, W_k$ in $\mathcal{W}$, and $u \in U$,

$$U \models u \in F_\beta(V_1, \ldots, V_n, W_1, \ldots, W_k)$$

if and only if

there exist $w_1 \in V_1, \ldots, w_n \in V_n, s \in W$ such that

a) $U \models u = \langle w_n, w_{n-1}, \dots, w_1 \rangle$,
b) $(W, R, p_1 := W_1, \dots, p_k := W_k) \models \alpha(w_1, \dots, w_n, s)$

if and only if
there exist $w_1 \in V_1, \dots, w_n \in V_n$ such that

a) $U \models u = \langle w_n, w_{n-1}, \dots, w_1 \rangle$,
b) $(W, R, p_1 := W_1, \dots, p_k := W_k) \models \exists x_{n+1} \alpha(w_1, \dots, w_n, x_{n+1})$.

□

LEMMA 6.3.8 *For all $u \in \overline{\mathcal{W}} \setminus \{X_0\}$, if $a \in u$ then $a \in U_c$, and either $a \in W$ or $a \in \overline{\mathcal{W}}$ and $h(a) < h(u)$.*

Proof.
The proof is by induction on $h(u)$, starting from $h(u) = 1$, which implies $u \in \mathcal{W}_0 = (\mathcal{W} \setminus \{\emptyset\}) \cup \{X_0\}$. Since $u \neq X_0$, we have that $u \in \mathcal{W} \setminus \{\emptyset\}$ and $a \in u$ implies $a \in W$.

Suppose that the thesis is true for all $u' \in \overline{\mathcal{W}} \setminus \{X_0\}$ with $h(u') < h(u)$, and let $\tau(\vec{X})$ be a term such that $\exists \vec{A} \in \mathcal{W}_0$ with $\mathcal{U}_c \models u = \tau(\vec{A})$ and $h(u) = h(\tau)$.

- If $\tau = \mu \cup \nu$, then $u = u_1 \cup^{\mathcal{U}_c} u_2$, with $\mathcal{U}_c \models u_1 = \mu(\vec{A})$, $\mathcal{U}_c \models u_2 = \nu(\vec{A})$, and $u_1, u_2 \in \overline{\mathcal{W}}$. Moreover, both u_1, u_2 are different from X_0: if u_1 were equal to X_0, then $u = X_0 \cup^{\mathcal{U}_c} u_2 = F(X_0) \cup' F(u_2) = X_0 \cup' u_2 = u_2$, and $h(u) = h(u_2) \leq h(\nu) < h(\tau))$ (the same argument shows that $u_2 \neq X_0$). Since $\emptyset \notin U_c$, u_1, u_2 are also different from $\emptyset$, and $u = F(u_1) \cup' F(u_2) = u_1 \cup' u_2 = u_1 \cup u_2$. Moreover, $h(\mu), h(\nu) < h(\tau) = h(u)$ and since $h(u_1) \leq h(\mu)$ and $h(u_2) \leq h(\nu)$, it follows that $h(u_1), h(u_2) < h(u)$. If $a \in u$, then $a \in u_1$ or $a \in u_2$ and the thesis follows by the inductive hypothesis.

- The case $\tau = \mu \setminus \nu$ is similar to the previous case, and it is left to the reader.

- If $\tau = \{\mu\}$, then $u = \{u'\}^{\mathcal{U}_c} = \{u'\}$, with $\mathcal{U}_c \models u' = \mu(\vec{A})$, $u' \in \overline{\mathcal{W}}$, and $h(u') < h(u)$. Then, $a \in u$ implies $a = u'$, and the thesis follows.

- If $\tau = Dom(\mu)$, then $u = Dom^{\mathcal{U}_c}(u')$, with $\mathcal{U}_c \models u' = \mu(\vec{A})$, $u' \in \overline{\mathcal{W}}$, and $h(u') < h(u)$. Since $u \neq X_0$ and $u = Dom^{\mathcal{U}_c}(u') = Dom'(F(u'))$, it follows that $u' \neq X_0$ and $u = Dom(u')$. If $a \in u$, then there exists b such that $\langle a, b \rangle \in u'$. By induction, $\langle a, b \rangle \in U_c$ and since no element of $W \cup \{X_0\}$ is a pair, it follows that $\langle a, b \rangle \in \overline{\mathcal{W}} \setminus \{X_0\}$, and $h(\langle a, b \rangle) < h(u') < h(u)$. Then, we can apply the inductive hypothesis to $\langle a, b \rangle$. Since $\{a\} \in \langle a, b \rangle$, we have that $\{a\} \in U_c$, and since no element of $W \cup \{X_0\}$ is a singleton, we obtain $\{a\} \in \overline{\mathcal{W}} \setminus \{X_0\}$, and $h(\{a\}) < h(\langle a, b \rangle) < h(u)$. Another application of the inductive hypothesis to $\{a\}$ guarantees that $a \in U_c$ and either $a \in W$ or $h(a) < h(\{a\}) < h(u)$.

- The case $\tau = Rng(\mu)$ is similar to the previous case, and it is left to the reader.

- If $\tau = \mu \times \nu$, then $u = u_1 \times^{\mathcal{U}_c} u_2$, with $u_1, u_2 \in \overline{\mathcal{W}}$, and $h(u_1), h(u_2) < h(u)$. Since $u \neq X_0$, we have that $u_1, u_2 \neq X_0$, and $u = u_1 \times^{\mathcal{U}_c} u_2 = u_1 \times u_2$. If $a \in u$, there exist $b \in u_1$ and $c \in u_2$ such that $a = \langle b, c\rangle$. Moreover, $a = \langle b, c\rangle = \langle b, c\rangle^{\mathcal{U}_c}$, because no element of $W \cup \{X_0\}$ is a singleton or a pair. Hence, it follows that $a \in \overline{\mathcal{W}}$, because, by induction, we have $b \in U_c$ and $c \in U_c$.

  To see that $h(a) < h(u)$, suppose that neither b nor c are in W. Then, $h(b) < h(u_1)$, $h(c) < h(u_2)$, and $h(a) \leq max(h(b), h(c)) + 4$, which is in turn strictly less than $max(h(u_1), h(u_2)) + 4 < h(u)$. Suppose that b is in W, but c is not. Hence $\{b\}$ is in $\mathcal{W}_0$ and $h(\{b\}) = 1$. Therefore, $h(a) \leq h(c) + 4 < max(h(u_1), h(u_2)) + 4 < h(u)$. The remaining cases are similar.

- If $\tau = \mu \times_= \nu$ or $\tau = \mu \times_\in \nu$, we can proceed as in the previous case.

- If $\tau = C_1(\mu, \nu)$, then $u = C_1^{\mathcal{U}_c}(u_1, u_2)$, with $u_1, u_2 \in \overline{\mathcal{W}}$, and $h(u_1), h(u_2) < h(u)$. Since $u \neq X_0$, we have that $u_1, u_2 \neq X_0$ and $u = C_1^{\mathcal{U}_c}(u_1, u_2) = C_1(u_1, u_2)$. If $a \in u$, then there are b, c, and d such that $a = \langle b, \langle c, d\rangle\rangle$ and $\langle b, c\rangle \in u_1$, $d \in u_2$. From the inductive hypothesis and the fact that no element of $W \cup \{X_0\}$ is a pair, it follows that $\langle b, c\rangle \in \overline{\mathcal{W}} \setminus \{X_0\}$, $h(\langle b, c\rangle) < h(u_1)$, $d \in U_c$, and either $d \in W$ or $d \in \overline{\mathcal{W}}$ and $h(d) < h(u_2)$. Using the inductive hypothesis twice on $b \in \{b\} \in \langle b, c\rangle$ as in the case of *Dom*, we can prove that either b is in W or $b \in \overline{\mathcal{W}}$ and $h(b) < h(u_1)$, and the same is true for c. Suppose that b, c, and d are not in W. Thus, $h(b), h(c) < h(u_1)$ and $h(a) \leq max(h(b), h(c), h(d)) + 8 < max(h(u_1), h(u_2)) + 8 < h(u)$. The case in which one of b, c, and d is in W is proved as in the preceding case.

- The case of $\tau = C_2(\mu, \nu)$ is similar to the previous case, and thus it is left to the reader.

□

LEMMA 6.3.12 *For any set term $\tau(\vec{y})$ there are:*

1. *terms $\tau_1(\vec{x}, \vec{y}), \ldots, \tau_{k_\tau}(\vec{x}, \vec{y},)$, such that $h(\tau_i) \leq h(\tau)$ and if $h(\tau) > 1$, then $h(\tau_i) < h(\tau)$,*

2. *Δ_0-formulae in the language $\{=, \in\}$, $\alpha_1^\tau(\vec{x}, \vec{y}), \ldots, \alpha_{k_\tau}^\tau(\vec{x}, \vec{y})$, with possibly W as parameter,*

such that, for all $\vec{A} \in \mathcal{W}_0 \cup W$, and for all $u \in U_c$,

$$U_c \models u \in \tau(\vec{A}) \leftrightarrow \bigvee_{i=1}^{k_\tau} \exists \vec{w} \in W \ (u = \tau_i(\vec{w}, \vec{A}) \wedge \alpha_i^\tau(\vec{w}, \vec{A})).$$

Proof.
The proof is by induction on the height of the term τ.

If $h(\tau) = 1$ and $\tau = y$, then put $k_\tau = 1$, $\tau_1 = x$, and $\alpha_1^\tau = x \in y$.

Suppose $h(\tau) > 1$, and let μ, ν be terms such that $h(\mu), h(\nu) < h(\tau)$. By induction, the lemma holds for μ, ν; let $\mu_1, \dots, \mu_{k_\mu}$ and $\nu_1, \dots, \nu_{k_\nu}$ be the terms, described in (1), for μ and ν, respectively.

We begin by showing the following four facts:

(a) there exists a Δ_0-formula $\alpha_=^{\mu,\nu}$ such that for all $\vec{A} \in \mathcal{W}_0 \cup W$,

$$\mathcal{U}_c \models \Big(\mu(\vec{A}) = \nu(\vec{A})\Big) \leftrightarrow \alpha_=^{\mu,\nu}(\vec{A});$$

(b) there exists a Δ_0-formula $\alpha_\in^{\mu,\nu}$ such that for all $\vec{A} \in \mathcal{W}_0 \cup W$,

$$\mathcal{U}_c \models \Big(\mu(\vec{A}) \in \nu(\vec{A})\Big) \leftrightarrow \alpha_\in^{\mu,\nu}(\vec{A});$$

(c) there exist Δ_0-formulae $\beta_{1,1}^\mu(\vec{x}, \vec{x}', \vec{y}), \beta_{1,2}^\mu(\vec{x}, \vec{x}', \vec{y}) \dots, \beta_{k_\mu,k_\mu}^\mu(\vec{x}, \vec{x}', \vec{y},)$, such that for all $\vec{A} \in \mathcal{W}_0 \cup W$ and for all $a, b \in U_c$,

$$\mathcal{U}_c \models \{a, b\} = \mu(\vec{A}) \leftrightarrow$$

$$\bigvee_{i,j=1}^{k_\mu} \exists \vec{w}', \vec{w}'' \in W \ (a = \mu_i(\vec{w}', \vec{A}) \wedge b = \mu_j(\vec{w}'', \vec{A}) \wedge \beta_{i,j}^\mu(\vec{w}', \vec{w}'', \vec{A}));$$

(d) there exist Δ_0-formulae $\sigma_1^\mu(\vec{x}, \vec{y}), \dots, \sigma_{k_\mu}^\mu(\vec{x}, \vec{y})$ such that for all $\vec{A} \in \mathcal{W}_0 \cup W$ and $a \in U_c$,

$$\mathcal{U}_c \models \{a\} = \mu(\vec{A}) \leftrightarrow \bigvee_{i=1}^{k_\mu} \exists \vec{w} \in W \ (a = \mu_i(\vec{w}, \vec{A}) \wedge \sigma_i^\mu(\vec{w}, \vec{A})).$$

Each of the above facts is proved by induction using the main inductive hypothesis.

(a) if $h(\mu) = h(\nu) = 1$, the formula $\alpha_=^{\mu,\nu}$ is $\mu = \nu$.

If $h(\mu) = 1$ and $h(\nu) > 1$ (or vice versa), then μ is a variable, say y_i. For all $\vec{A} \in \mathcal{W}_0 \cup W$, $\mu(\vec{A}) = A_i$, $\nu(\vec{A}) \notin W$ (see Corollary 6.3.10). Since $\mathcal{U}_c$ is extensional on $\overline{\mathcal{W}}$ we obtain:

$$\mathcal{U}_c \models A_i = \nu(\vec{A}) \leftrightarrow A_i \notin W \wedge \forall t\ (t \in A_i \leftrightarrow t \in \nu(\vec{A})).$$

Since $1 < h(\nu) < h(\tau)$, the main inductive hypothesis holds for ν, with terms ν_j such that $h(\nu_j) < h(\nu)$, and

$$\mathcal{U}_c \models t \in \nu(\vec{A}) \leftrightarrow \bigvee_{j=1}^{k_\nu} \exists \vec{w} \in W \left(t = \nu_j(\vec{w}, \vec{A}) \wedge \alpha_j^\nu(\vec{w}, \vec{A})\right).$$

Thus $\forall t\ (t \in A_i \leftrightarrow t \in \nu(\vec{A}))$ is equivalent in $\mathcal{U}_c$ to

$$\forall t \in A_i \left(\bigvee_{j=1}^{k_\nu} \exists \vec{w} \in W \left(t = \nu_j(\vec{w}, \overline{A}) \wedge \alpha_j^\nu(\vec{w}, \vec{A})\right)\right) \wedge$$

$$\bigwedge_{j=1}^{k_\nu} \forall \vec{w} \in W \left(\alpha_j^\nu(\vec{w}, \vec{A}) \rightarrow \nu_j(\vec{w}, \vec{A}) \in A_i\right).$$

Notice now that if $h(\nu_j) > 1$, then the element $\nu_j(\vec{w}, \vec{A})$ does not belong to A_i in the model $\mathcal{U}_c$ (since it does not belong to W), and thus $\alpha_j^\nu(\vec{w}, \vec{A})$ is false in $\mathcal{U}_c$. This implies that $A_i = \nu(\vec{A})$ is equivalent in $\mathcal{U}_c$ to

$$A_i \notin W \wedge \forall t \in A_i \left(\bigvee_{j=1}^{k_\nu} \exists \vec{w} \in W \left(t = \nu_j(\vec{w}, \vec{A}) \wedge \alpha_j^\nu(\vec{w}, \vec{A})\right)\right) \wedge$$

$$\wedge \bigwedge_{\substack{1 \le j \le k_\nu \\ h(\nu_j)=1}} \forall \vec{w} \in W \left(\alpha_j^\nu(\vec{w}, \vec{A}) \rightarrow \nu_j(\vec{w}, \vec{A}) \in A_i\right),$$

which is a Δ_0-formula except for the subformula $t = \nu_j(\vec{w}, \vec{A})$. Since $h(\nu_j) < h(\nu)$, we can proceed inductively on $t = \nu_j(\vec{w}, \vec{A})$ to obtain the required Δ_0-formula.

If $h(\mu), h(\nu) > 1$, then for all $\vec{A} \in \mathcal{W}_0 \cup W$, $\mu(\vec{A}), \nu(\vec{A}) \notin W$ and by Lemma 6.3.11 (extensionality),

$$\mathcal{U}_c \models \mu(\vec{A}) = \nu(\vec{A}) \leftrightarrow \mu(\vec{A}) \subseteq \nu(\vec{A}) \wedge \nu(\vec{A}) \subseteq \mu(\vec{A}).$$

By the inductive hypothesis, the lemma holds for μ, ν with terms μ_i, ν_j such that $h(\mu_i) < h(\mu)$ and $h(\nu_i) < h(\nu)$. Thus, the first conjunct on the right-hand side of the above formula is equivalent to

$$\bigwedge_{i=1}^{k_\mu} \forall \vec{w} \in W \left[\alpha_i^\mu(\vec{w}, \vec{A}) \rightarrow \bigvee_{j=1}^{k_\nu} \exists \vec{w}' \in W\ (\mu_i(\vec{w}, \vec{A}) = \nu_j(\vec{w}', \vec{A}) \wedge \alpha_j^\nu(\vec{w}', \vec{A}))\right].$$

Consider the subformula $\mu_i(\vec{w}, \vec{A}) = \nu_j(\vec{w}', \vec{A})$. Since $h(\mu_i) < h(\mu)$ and $h(\nu_j) < h(\nu)$, we can proceed inductively to obtain the required Δ_0-formula. The same argument applies to the second conjunct.

(b) It follows directly from (a) and the inductive hypothesis of the lemma.

(c) If $h(\mu) = 1$, then $\mu = y_i$ and we can take $\beta_{1,1}(x, x', y)$ equal to the formula

$$y \notin W \wedge x \in y \wedge x' \in y \wedge \forall v \in W\ (v \in y \rightarrow v = x \vee v = x').$$

If $h(\mu) > 1$, then $\mu(\vec{A}) \notin W$ and, by Lemma 6.3.11, we have

$$\mathcal{U}_c \models \{a, b\} = \mu(\vec{A}) \leftrightarrow a \in \mu(\vec{A}) \wedge b \in \mu(\vec{A}) \wedge \forall t\ (t \in \mu(\vec{A}) \rightarrow t = a \vee t = b).$$

Since $h(\mu) < h(\tau)$, by the main inductive hypothesis this is easily seen to be equivalent (in $\mathcal{U}_c$) to the formula

$$\bigvee_{i,j=1}^{k_\mu} \exists \vec{w}' \vec{w}'' \in W$$

$$\Big[a = \mu_i(\vec{w}', \vec{A}) \wedge b = \mu_j(\vec{w}'', \vec{A}) \wedge \alpha_i^\mu(\vec{w}', \vec{A}) \wedge \alpha_j^\mu(\vec{w}'', \vec{A}) \wedge$$

$$\bigwedge_{l=1}^{k_\mu} \forall \vec{s} \in W \left(\alpha_l^\mu(\vec{s}, \vec{A}) \rightarrow \mu_l(\vec{s}, \vec{A}) = \mu_i(\vec{w}', \vec{A}) \vee \mu_l(\vec{s}, \vec{A}) = \mu_j(\vec{w}'', \vec{A}) \right) \Big].$$

Since $h(\mu_i) < h(\mu) < h(\tau)$ for any $i \in \{1, \ldots, k_\mu\}$, from (a) it follows that the equalities between the terms above are equivalent to Δ_0-formulae.

(d) It easily follows from (c).

To complete the proof of the lemma, we consider the various cases corresponding to the structure of the term τ. To simplify the notation, we sometime avoid to explicitly indicate the elements $\vec{A}$. We suppose that the lemma holds for terms μ, ν with $h(\mu), h(\nu) < h(\tau)$, and we denote by μ_i, ν_j, $\alpha_i^\mu, \alpha_j^\nu$ the corresponding terms and Δ_0-formulae.

- The case $\tau = \mu \cup \nu$ directly follows from the inductive hypothesis.
- Let $\tau = \mu \setminus \nu$. In $\mathcal{U}_c$ we have$a \in \tau(\vec{A}) \leftrightarrow a \in \mu(\vec{A}) \wedge a \notin \nu(\vec{A})$, which is easily seen to be equivalent to

$$\bigvee_{i=1}^{k_\mu} \exists \vec{w} \in W$$

$$\left[a = \mu_i(\vec{w}) \wedge \alpha_i^{\mu}(\vec{w}) \wedge \bigwedge_{j=1}^{k_\nu} \forall \vec{w}' \in W \left(\alpha_j^{\nu}(\vec{w}') \rightarrow \mu_i(\vec{w}) \neq \nu_j(\vec{w}')\right)\right].$$

The last formula can be reduced to a Δ_0-formula applying (a) to each subformula $\mu_i(\vec{w}) \neq \nu_j(\vec{w}')$. Notice that $h(\mu_i) \leq h(\mu) < h(\tau)$, $h(\nu_j) \leq h(\nu) < h(\tau)$.

- Let $\tau = \{\mu\}$. In this case, $\mathcal{U}_c \models a \in \tau(\vec{A}) \leftrightarrow a = \mu(\vec{A})$. Since $h(\mu) < h(\tau)$, it is enough to take $\tau_1(\vec{x}, \vec{y}) = \mu(\vec{y})$ and $y_1 = y_1$ as Δ_0-formula.

- Let $\tau = \mu \times \nu$. By the inductive hypothesis, $a \in \tau(\vec{A})$ is equivalent in $\mathcal{U}_c$ to

$$\bigvee_{i=1}^{k_\mu} \bigvee_{j=1}^{k_\nu} \exists \vec{w}\vec{w}' \in W \ \left[a = \langle \mu_i(\vec{w}), \nu_j(\vec{w}')\rangle \wedge \alpha_i^{\mu}(\vec{w}) \wedge \alpha_j^{\nu}(\vec{w}')\right].$$

  Since $h(\langle \mu_i, \nu_j \rangle) < h(\tau)$, the thesis follows.

- The case of $\tau = \mu \times_= \nu$ is proved as the case of $\tau = \mu \times \nu$ using (a).

- The case $\tau = \mu \times_\in \nu$ is proved as the case of $\tau = \mu \times \nu$ using (b).

- Let $\tau = Dom(\mu)$. Suppose that $h(\mu) = 1$. In this case, $\mu = y_i$ and $\mathcal{U}_c \models \forall x(x \notin Dom(A_i))$, because no element of an element of $\mathcal{W}_0 \cup W$ is a pair. It is then enough to take $\tau_1 = x$, and $x \neq x$ as α_1^{τ}.

  Consider now the case in which $h(\mu) > 1$. In $\mathcal{U}_c$, $a \in \tau(\vec{A})$ is equivalent to

$$\exists b \ \left(\langle a, b\rangle = \{\{a\}, \{a, b\}\} \in \mu(\vec{A})\right),$$

  which, in turn, is equivalent to

$$\exists b, c, d \ \left(c = \{a\} \wedge d = \{a, b\} \wedge \{c, d\} \in \mu(\vec{A})\right).$$

  By the inductive hypothesis and (c), this can be rewritten as

$$\exists b, c, d(c = \{a\} \wedge d = \{a, b\} \wedge \bigvee_{l,m=1}^{k_{\mu_i}} \wedge \bigvee_{i=1}^{k_\mu} \exists \vec{w} \in W(\alpha_i^{\mu}(\vec{w}) \wedge$$

$$\bigvee_{l,m=1}^{k_{\mu_i}} \exists \vec{w}', \vec{w}'' \in W(c = \mu_{i,l}(\vec{w}', \vec{w}) \wedge d = \mu_{i,m}(\vec{w}'', \vec{w}) \wedge \beta_{l,m}^{\mu_i}(\vec{w}', \vec{w}'', \vec{w})))],$$

  in which c and d can be eliminated obtaining

$$\exists b \bigvee_{i=1}^{k_\mu} \bigvee_{l,m=1}^{k_{\mu_i}} \exists \vec{w}, \vec{w}', \vec{w}'' \in W[\mu_{i,l}(\vec{w}', \vec{w}) = \{a\} \wedge \mu_{i,m}(\vec{w}'', \vec{w}) = \{a, b\} \wedge$$

$$\alpha_i^{\mu}(\vec{w}) \wedge \beta_{l,m}^{\mu_i}(\vec{w}', \vec{w}'', \vec{w})].$$

Using (c) again, we can eliminate b, obtaining the following formula:

$$\bigvee_{i=1}^{k_\mu} \bigvee_{l,m=1}^{k_{\mu_i}} \bigvee_{r,s=1}^{k_{\mu_{i,m}}} \exists \vec{w}, \vec{w}', \vec{w}'', \vec{v}', \vec{v}'' \in W$$

$$[a = \mu_{i,m,r}(\vec{v}', \vec{w}'', \vec{w}) \wedge \beta_{r,s}^{\mu_{i,m}}(\vec{v}', \vec{v}'', \vec{w}'', \vec{w}) \wedge$$

$$\beta_{l,m}^{\mu_i}(\vec{w}', \vec{w}'', \vec{w}) \wedge \mu_{i,l}(\vec{w}', \vec{w}) = \{\mu_{i,m,r}(\vec{v}', \vec{w}'', \vec{w})\} \wedge \alpha_i^{\mu}(\vec{w}) \wedge \beta_{l,m}^{\mu_i}(\vec{w}', \vec{w}'', \vec{w})$$

Since $h(\mu) > 1$, we have that $h(\mu_i) < h(\mu)$, for all $i \in \{1, \ldots, k_\mu\}$. It follows that $h(\mu_{i,l}), h(\{\mu_{i,m,r}\}) < h(\tau)$, and a further application of (a) to the subformula $\mu_{i,l}(\vec{w}', \vec{w}) = \{\mu_{i,m,r}(\vec{v}', \vec{w}'', \vec{w})\}$ gives the desired result.

- The case of $\tau = Rng(\mu)$ is proved as the case of Dom.
- Let $\tau = C_1(\mu, \nu)$. In $\mathcal{U}_c$, we have

$$a \in \tau(\vec{A}) \leftrightarrow \exists b, c, d\ (a = \langle b, \langle c, d\rangle\rangle \wedge \langle b, c\rangle \in \mu(\vec{A}) \wedge d \in \nu(\vec{A})).$$

Let $\mu' = Dom(\mu)$ and $\mu'' = Rng(\mu)$. Notice that $\langle b, c\rangle \in \mu(\vec{A})$ implies $b \in \mu'(\vec{A})$ and $c \in \mu''(\vec{A})$. Hence the above formula is equivalent to

$$\exists b, c, d\ (a = \langle b, \langle c, d\rangle\rangle \wedge b \in \mu'(\vec{A}) \wedge c \in \mu''(\vec{A}) \wedge \langle b, c\rangle \in \mu(\vec{A}) \wedge d \in \nu(\vec{A})).$$

Since $h(\mu'), h(\mu'') < h(\tau)$, we can conclude as in the case of $\times$.

- The case of $\tau = C_2(\mu, \nu)$ is proved as the case of C_1.

□

Bibliography

[1] H. Andréka, I. Nemeti, and I. Sain. On interpolation, amalgamation, universal algebra and Boolean algebras with operators. Technical report, Mathematical Institute, Budapest, 1995.

[2] P. Aczel. *Non-well-founded sets.* CSLI, Lecture Notes 14, 1988.

[3] A. Baltag. Modal characterization for sets and Kripke models. To appear.

[4] J. Barwise. *Admissible sets and structure.* Springer-Verlag, 1975.

[5] J. Barwise, J.F.A.K. van Benthem. Interpolation, Preservation, and Pebble Games. Thecnical Report of the ILLC, University of Amsterdam, ML-96-12 (to appear in the Journal of Symbolic Logic).

[6] J. Barwise and J. Etchemendy. *The Liar: An Essay on Truth and Circular Propositions.* Oxford University Press, Oxford, 1987.

[7] J. Barwise and L. Moss. *Vicious Circle.*, volume 60 of *CSLI Lecture Notes*, Stanford California, 1996.

[8] J.F.A.K. van Benthem. Syntactic aspects of modal incompleteness theorems. *Theoria*, 45, 1979, pp. 67-81.

[9] J.F.A.K. van Benthem. *Modal Logic and Classical Logic.* Bibliopolis, Napoli, 1985.

[10] J.F.A.K. van Benthem and K. Doets. Higher-Order Logic. In *Handbook of Philosophical Logic*, Vol. I, D. Gabbay and F. Guenthner (eds.); D. Reidel Pub. Comp., Dordrecht-Holland, 1983, pp. 275-329.

[11] J.F.A.K. van Benthem. Modalities, Bisimulation and Interpolation in Infinitary Logic. Thecnical Report of the ILLC, University of Amsterdam, ML-97-06 (to appear in the Annals of Pure and Applied Logic).

[12] J.F.A.K. van Benthem. *Modal Correspondence Theory.* PhD thesis, Mathematisch Instituut voor Grondslagenonderzoek, University of Amsterdam, 1976.

[13] J.F.A.K. van Benthem. Dynamic Odds and Ends. Thecnical Report of the ILLC, University of Amsterdam, ML-1998-08.

[14] J.F.A.K. van Benthem, G. D'Agostino, A. Montanari, and A. Policriti. Modal Deduction in Second-Order Logic and Set Theory-I. *Journal of Logic and Computation* Vol 7 Nr.2 pp.251-265, 1997.

[15] J.F.A.K. van Benthem, G. D'Agostino, A. Montanari, and A. Policriti. Modal Decuction in Second-Order Logic and Set Theory-II. *Studia Logica* Vol. 60 n.3 pp. 387-420.

[16] D. Cantone, A. Ferro, E.G. Omodeo. *Computable Set Theory. Vol 1.* Oxford University Press, 1989. Int. Series of Monographs on Computer Science.

[17] E.M. Clarke, E.A. Emerson, and A.P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1986.

[18] W. Craig. Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *Journal of Symbolic Logic*, 22:269–285, 1957.

[19] G. D'Agostino, A. Montanari, and A. Policriti. Set-theoretic decidability result for modal logic. Thecnical Report of the University of Udine, UDMI/21/95/RR.

[20] G. D'Agostino, A. Montanari, and A. Policriti. A Set-Theoretic Translation Method for Polymodal Logics. *Journal of Automated Reasoning* 15: 317-337, 1995.

[21] G. D'Agostino, M. Hollenberg. Logical questions concerning the μ-calculus: Interpolation, Lyndon and Łoś-Tarski. (to appear in the *Journal of Symbolic Logic).*

[22] D. van Dalen. *Logic and Structures.* Springer Verlag, 1991.

[23] H. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer, Berlin, 1995.

[24] E.A. Emerson and E. M. Clarke. Characterizing correctness properties of parallel programs using fixpoints. 7th *ICALP, LNCS* n. 85, 1980, pp. 169-181.

[25] E.A. Emerson and J.Y. Halpern. Sometimes and not never revisited: on branching versus linear time. *Journal of the ACM*, 33(1):151–178, 1986.

[26] E.A. Emerson and Chin-Laung Lei. Modalities for model checking: branching time logic strikes back. *Science of Computer Programming*, 8:275–306, 1987.

[27] M. Fattorosi-Barnaba and F. de Caro. Graded Modalities I. *Studia Logica* 44:197-221, 1985.

[28] K. Fine. In so many possible words. *Notre Dame Journal of Formal Logic*, 13:516-520, 1972.

[29] K. Fine. Logics containing K4. Part II. *Journal of Symbolic Logic*, Vol 50 N 3 :619-651, 1985.

[30] M.J. Fisher and R.E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[31] M. Forti and F. Honsell. Set Theory with Free Construction Principles. In *Annali Scuola Normale Superiore-Pisa Classe di Scienza*, 10:493-522, Serie IV, 1983.

[32] S. Ghilardi and Zawadowski. Model completions and r-Heyting categories. *Annals of Pure and Applied Logic*, Vol. 88 n.1 pp.27-46, 1997.

[33] K. Gödel. The Consistency of the Axiom of Choice and of the Generalized Continuum Hypothesis with the Axioms of Set Theory. *Annals of Mathematical Studies*, 3, Princeton University Press, Princeton, N.J., 1940.

[34] L.F. Goble. Grades of modalities. *Logioque et Analyse*, 13:323-324, 1970.

[35] L. Henkin. An extension of the Craig-Lyndon interpolation theorem. *Journal of Symbolic Logic*, 28:201–216, 1963.

[36] W. van der Hoek On the semantics of graded modalities. *Journal of Applied Non-Classical Logics,* 2, 81-123, 1992.

[37] W. van der Hoek, M. de Rijke. Counting Objects. *Journal of Logic and Computation* Vol.5 No. 3, 325-345, 1995.

[38] M. Hollenberg. *Logic and Bisimulation*. PhD thesis, Quaestiones Infinitae, Zeno Institute of Philosophy, Vol. XXIV, 1998.

[39] D. Janin. *Propriétés Logiques du NonDéterminisme et μ-Calcul Modal.* PhD thesis, LaBRI – Université de Bordeaux I, 1996

[40] D. Janin and I. Walukiewicz. Automata for the Modal μ-Calculus and related results. In *Proceedings MFCS'95*, pages 552–562. Springer, 1995.

[41] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus w.r.t. Monadic Second-Order logic (to appear in the *Proceedings of CONCUR'96*, 1996).

[42] D. Kaplan. S5 with multiple possibility. *Journal of Symbolic Logic*, 35:355-356, 1970.

[43] H.J. Keisler. *Model Theory for Infinitary Logic.* North Holland, 1971.

[44] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27:234–241, 1983.

[45] J.L. Krivine. *Introduction to axiomatic Set Theory.* D. Reidel Pub. Comp., Dordrecht-Holland, 1971.

[46] A. Kurucz. Decision Problems in Algebraic Logic PhD thesis, Mathematical Institute, Budapest, 1997.

[47] G. Lenzi. The μ-calculus and the hierarchy problem. PhD thesis, Scuola Normale Superiore, Pisa, 1997.

[48] E. Lopez-Escobar. An interpolation theorem for denumerably long sentences. *Fund. Math.*, LXII, 253-272, 1965.

[49] L. Maksimova. Amalgamation and interpolation in normal modal logics. *Studia Logica*, L(3/4):457-471, 1991.

[50] M. Marx. *Algebraic Relativization and Arrow Logic.* PhD thesis, ILLC-dissertation series 1995-3, 1995.

[51] S. Mikulas. *Taming Logics.* PhD thesis, ILLC Dissertation series 1995-12.

[52] R. Milner. *Communication and concurrency.* Prentice Hall, Englewood Cliffs, New Jersey, 1989.

[53] A. Montanari, A. Policriti. A Set Theoretic Approach to Automated Deduction in Graded Modal Logics. In *Prooceedings of the International Joint Conference on Artificial Intelligence, IJCAI 97, Nagoya, Japan, 1997.*

[54] D. Park. Concurrency and automata on infinite sequences. In *Proceedings 5th GI Conference*, pages 167–183, Springer 1981.

[55] D. Pigozzi. Amalgamation, congruence extension and interpolation properties in algebras. *Algebra Universalis* 1(3):269-349, 1972.

[56] A. Pitts. On an interpretation of second-order quantification in first-order intuitionistic propositional logic. *Journal of Symbolic Logic*, 57:33–52, 1992.

[57] A. Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science*, pages 46–57, 1977.

[58] A. Policriti, J.T. Schwartz. T-Theorem Proving I. *Journal of Symbolic Computation* 20:315-342, 1995.

[59] V. Pratt. Semantical considerations on Floyd-Hoare logic. In *Proceedings 17th IEEE Symposium on Foundations of Computer Science*, pages 109–121, 1976.

[60] L. Rieger. A contribution to Göedel's Axiomatic Set Theory, I. *Czechoslovak Math.* J. 7: 323-357.

[61] M. de Rijke. A note on Graded Modal Logics. Technical report, Department of Computer Science, University of Warwick, Report RR318, 1996.

[62] D. Scott. A different kind of model for set theory. Unpublished paper, given at the *1960 Stanford Congress of Logic, Methodology and Philosophy of Science.*

[63] R.S. Streett and E.A. Emerson. An automata theoretic procedure for the propositional mu-calculus. *Information and Computation*, 81:249–264, 1989.

[64] H. Sturm. *Modale Fragmente von* $\mathcal{L}_{\omega\omega}$ *und* $\mathcal{L}_{\omega_1\omega}$. PhD thesis, University of Munich, CIS, Munich 1997

[65] S.K. Thomason. Reduction of tense logic to modal logic I. *Journal of Symbolic Logic*, 39(3) 549-551, 1974.

[66] A. Visser. Bisimulations, model descriptions and propositional quantifiers. Logic Group Preprint Series 161, Dept. of Philosophy, Utrecht University, 1996.

[67] I. Walukiewicz. Completeness of Kozen's axiomatization of the propositional μ-calculus. In *Proceedings LICS'95*, pages 14–24, 1995.

Index

Samenvatting

Bisimulatie is een eenvoudige en natuurlijke notie van equivalentie tussen verschillende weergaven van processen en verwante structuren, die sinds 1980 een belangrijk begrip is geworden in logica en informatica. Met name verbindt het de twee voornaamste gebieden die in dit proefschrift worden samengebracht: "uitgebreide modale logica's" en "niet-gefundeerde verzamelingen". Modale logica's beschrijven z.g. *proces-invarianten*, d.w.z. eigenschappen die processen blijven bewaren, ook onder een andere manier van beschrijven van hun toestanden en toestandsovergangen, mits die laatste een geschikte bisimulatie toelaat met de eerdere weergave. Dit verband blijkt overigens heel flexibel: verschillende noties van simulatie blijken equivalent met verschillende visies op een proces, en deze weer met verschillende modale talen om invarianten te beschrijven.

Verzamelingen hebben althans formeel eenzelfde structuur als processen: hun leden zijn op te vatten als toestanden, en overgangen zijn mogelijk hiertussen langs de relatie "element van". In de moderne "niet-gefundeerde verzamelingenleer" blijkt deze parallel essentieel. Het gaat hier om een uitbreiding van de klassieke verzamelingenleer met de mogelijkheid van circulariteit en oneindige regressie in lidmaatschap –iets wat op vele gebieden in de moderne semantiek en informatica nuttig blijkt. Twee weergaven van een niet-gefundeerde verzameling gelden nu als "hetzelfde" wanneer er een bisimulatie tussen bestaat. Ook dit verband blijkt weer flexibel: verschillende noties van simulatie genereren verschillende genres "verzamelingen", en daarmee verschillende grondlagentheorieen.

De aldus geslagen bruggen tussen modale logica, verzamelingstheorie, en procestheorie kunnen gebruikt worden om resultaten en technieken de een oever over te hevelen naar de andere. Hierbij ontstaan vaak verrassende nieuwe inzichten. In dit proefschrift steken we een aantal malen over. Onze voornaamste resultaten zijn daarbij als volgt –na een eerste inleiding in het totale werkterrein.

In hoofdstuk 3 bewegen we van modale logica naar procestheorie. We bestuderen de z.g. "μ-calculus", een modale procestaal uit de vroege 80er jaren waarin veel rekenprocessen en handelingen in het algemeen op een natuurlijke wijze kunnen worden weergegeven, zodat hun gedrag exact kan worden bestudeerd. Ons hoofdresultaat is dat de μ-calculus de logische "uniforme interpolatie-eigenschap" heeft. Dit zegt, heel ruwweg, dat gegevens opgeslagen in deze taal een modulaire structuur hebben, zodat afleiden van verdere informatie zich steeds kan beperken tot de relevante module. Ons bewijs van deze wenselijke eigenschap gebruikt algemene modale logica, maar in het byzonder ook technieken uit de automatentheorie en z.g. "bisimulatie-kwantoren", een recente uitbreiding van standaard modale talen bedacht door Marco Hollenberg.

In hoofdstuk 4 introduceren we modale talen als techniek in de verzamelingenleer. Bekend is uit recent werk van Jon Barwise en Larry Moss [7] dat elke niet-gefundeerde verzameling een unieke beschrijving heeft door middel van een modale formule met oneindige conjuncties en disjuncties. Wij generalizeren deze opmerking, en tonen aan dat zulke modale formules een alternatief model vormen voor het gangbare verzamelingstheoretisch universum, beschreven door Peter Aczel [2]. Maar er zijn ook alternatieve visies, waaronder een van Dana Scott, die verzamelingen definieert via een andere simulatie. We identificeren ook hiervoor de juiste modale taal, die de gewone modale logica uitbreidt met z.g. tel-modaliteiten, die aantallen toegankelijke alternatieven tellen vanuit een geven toestand. Dit generalizeert het hoofdresultaat van Barwise & Moss, en maakt ook een spectrum aan verschillende niet-gefundeerde verzamelingsnoties zichtbaar.

Tel-modaliteiten (E: "graded modalities") vormen ook op zich een natuurlijk formalisme. Ze komen bijvoorbeeld natuurlijk op bij de veel voorkomende techniek van "uitrollen" van procesmodellen tot bisimulerende boomstructuren. Hajnal Andréka bewees dat deze uitgebreide modale taal in het algemeen geen interpolatie kent. Wij laten echter zien in hoofdstuk 5 dat er wel een gemodificeerde vorm van deze eigenschap opgaat (voorgesteld voor oneindige talen door Jon Barwise en Johan van Benthem [5]), die bovendien voor sommige talen met tel-modaliteiten te versterken is tot de gebruikelijke versie. Wij bepalen alle gevallen waarin deze "opwaardering" mogelijk is.

Tenslotte steken we nogmaals over van verzamelingenleer naar modale logica. In hoofdstuk 6 laten we zien hoe modale talen kunnen worden vertaald in een verzamelingenleer, met als voornaamste idee de behandeling van een "universele modaliteit" als een machtsverzamelingsoperator. Als gevolg kunnne we geldigheid in allerlei modale logica's, na vertaling, opvatten als afleidbaarheid in een eenvoudige axiomatische theorie van niet-gefundeerde verzamelingen: Ω. Deze theorie werd in 1995 voorgesteld door D'Agostino, Montanari en Policriti in [20] als een computationeel aantrekkelijke "kern" van de standaard verzamelingstheorie. Om met deze techniek ook allerlei uitgebreide modale logica's te kunnen behandelen, breiden we de basis modale logica uit tot een natuurlijke tweede-orde logica L_2. Tegelijkertijd versterken we Omega met axioma's die ons alle verzamelingstheoretische operaties geven in het "construeerbare universum", gebruikt door Kurt Gödel voor zijn bewijs van consistentie van de continuum hypothese. Ons voornaamste resultaat is een getrouwe vertaling van afleidbaarheid in L_2 naar afleidbaarheid in deze sterkere verzamelingenleer. Verzamelingenleer valt dus te bestuderen met modale logica's, maar het omgekeerde geldt evenzeer.

Onze voornaamste conclusie is dat modale logica, procestheorieen, en verzamelingsleer vele bruikbare analogieen bezitten. Onze aanbeveling is om deze gebieden dan ook meer in samenhang te bestuderen.

Titles in the ILLC Dissertation Series:

ILLC DS-1993-01: **Paul Dekker**
Transsentential Meditations; Ups and downs in dynamic semantics

ILLC DS-1993-02: **Harry Buhrman**
Resource Bounded Reductions

ILLC DS-1993-03: **Rineke Verbrugge**
Efficient Metamathematics

ILLC DS-1993-04: **Maarten de Rijke**
Extending Modal Logic

ILLC DS-1993-05: **Herman Hendriks**
Studied Flexibility

ILLC DS-1993-06: **John Tromp**
Aspects of Algorithms and Complexity

ILLC DS-1994-01: **Harold Schellinx**
The Noble Art of Linear Decorating

ILLC DS-1994-02: **Jan Willem Cornelis Koorn**
Generating Uniform User-Interfaces for Interactive Programming Environments

ILLC DS-1994-03: **Nicoline Johanna Drost**
Process Theory and Equation Solving

ILLC DS-1994-04: **Jan Jaspars**
Calculi for Constructive Communication, a Study of the Dynamics of Partial States

ILLC DS-1994-05: **Arie van Deursen**
Executable Language Definitions, Case Studies and Origin Tracking Techniques

ILLC DS-1994-06: **Domenico Zambella**
Chapters on Bounded Arithmetic & on Provability Logic

ILLC DS-1994-07: **V. Yu. Shavrukov**
Adventures in Diagonalizable Algebras

ILLC DS-1994-08: **Makoto Kanazawa**
Learnable Classes of Categorial Grammars

ILLC DS-1994-09: **Wan Fokkink**
Clocks, Trees and Stars in Process Theory

ILLC DS-1994-10: **Zhisheng Huang**
Logics for Agents with Bounded Rationality

ILLC DS-1995-01: **Jacob Brunekreef**
On Modular Algebraic Protocol Specification

ILLC DS-1995-02: **Andreja Prijatelj**
Investigating Bounded Contraction

ILLC DS-1995-03: **Maarten Marx**
Algebraic Relativization and Arrow Logic

ILLC DS-1995-04: **Dejuan Wang**
Study on the Formal Semantics of Pictures

ILLC DS-1995-05: **Frank Tip**
Generation of Program Analysis Tools

ILLC DS-1995-06: **Jos van Wamel**
Verification Techniques for Elementary Data Types and Retransmission Protocols

ILLC DS-1995-07: **Sandro Etalle**
Transformation and Analysis of (Constraint) Logic Programs

ILLC DS-1995-08: **Natasha Kurtonina**
Frames and Labels. A Modal Analysis of Categorial Inference

ILLC DS-1995-09: **G.J. Veltink**
Tools for PSF

ILLC DS-1995-10: **Giovanna Cepparello**
Studies in Dynamic Logic

ILLC DS-1995-11: **W.P.M. Meyer Viol**
Instantial Logic. An Investigation into Reasoning with Instances

ILLC DS-1995-12: **Szabolcs Mikulás**
Taming Logics

ILLC DS-1995-13: **Marianne Kalsbeek**
Meta-Logics for Logic Programming

ILLC DS-1995-14: **Rens Bod**
Enriching Linguistics with Statistics: Performance Models of Natural Language

ILLC DS-1995-15: **Marten Trautwein**
Computational Pitfalls in Tractable Grammatical Formalisms

ILLC DS-1995-16: **Sophie Fischer**
The Solution Sets of Local Search Problems

ILLC DS-1995-17: **Michiel Leezenberg**
Contexts of Metaphor

ILLC DS-1995-18: **Willem Groeneveld**
Logical Investigations into Dynamic Semantics

ILLC DS-1995-19: **Erik Aarts**
Investigations in Logic, Language and Computation

ILLC DS-1995-20: **Natasha Alechina**
Modal Quantifiers

ILLC DS-1996-01: **Lex Hendriks**
Computations in Propositional Logic

ILLC DS-1996-02: **Angelo Montanari**
Metric and Layered Temporal Logic for Time Granularity

ILLC DS-1996-03: **Martin H. van den Berg**
Some Aspects of the Internal Structure of Discourse: the Dynamics of Nominal Anaphora

ILLC DS-1996-04: **Jeroen Bruggeman**
Formalizing Organizational Ecology

ILLC DS-1997-01: **Ronald Cramer**
Modular Design of Secure yet Practical Cryptographic Protocols

ILLC DS-1997-02: **Nataša Rakić**
Common Sense Time and Special Relativity

ILLC DS-1997-03: **Arthur Nieuwendijk**
On Logic. Inquiries into the Justification of Deduction

ILLC DS-1997-04: **Atocha Aliseda-LLera**
Seeking Explanations: Abduction in Logic, Philosophy of Science and Artificial Intelligence

ILLC DS-1997-05: **Harry Stein**
The Fiber and the Fabric: An Inquiry into Wittgenstein's Views on Rule-Following and Linguistic Normativity

ILLC DS-1997-06: **Leonie Bosveld - de Smet**
On Mass and Plural Quantification. The Case of French 'des'/'du'-NP's.

ILLC DS-1998-01: **Sebastiaan A. Terwijn**
Computability and Measure

ILLC DS-1998-02: **Sjoerd D. Zwart**
Approach to the Truth: Verisimilitude and Truthlikeness

ILLC DS-1998-03: **Peter Grunwald**
not yet available